

Power Systems

SAS RAID controllers for Linux

IBM

Power Systems

SAS RAID controllers for Linux

IBM

Note

Before using this information and the product it supports, read the information in “Safety notices” on page vii, “Notices” on page 153, the *IBM Systems Safety Notices* manual, G229-9054, and the *IBM Environmental Notices and User Guide*, Z125-5823.

This edition applies to IBM Power Systems™ servers that contain the POWER8® processor and to all associated models.

© Copyright IBM Corporation 2014, 2017.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Safety notices	vii
SAS RAID controllers for Linux	1
General information	1
Feature comparison of SAS RAID cards	2
PCI-X SAS RAID card comparison	2
PCIe SAS RAID card comparison	7
PCIe2 SAS RAID card comparison	14
PCIe3 SAS RAID card comparison	17
SAS overview	23
SAS architecture summary	24
Disk arrays	25
Supported RAID levels	27
RAID 0	27
RAID 5	28
RAID 6	29
RAID 10	30
RAID 5T2	31
RAID 6T2	32
RAID 10T2	33
Estimating disk array capacities	34
RAID level summary	34
Stripe-unit size	35
Disk array overview	35
Disk array states	38
Physical disk states	39
I/O adapter states	39
Auxiliary write cache adapter	39
Related information	41
RAID controller software	41
Verifying installation of the controller software	42
Linux ipr device driver updates	44
Updating the iprutils package	44
Common IBM SAS RAID controller tasks	45
Starting the iprconfig utility	45
Status of devices, arrays, and paths	46
Viewing device status	46
Viewing array status	49
Viewing path status	50
RAID and JBOD formats	52
Formatting to advanced function	52
Formatting to JBOD	52
Creating and deleting disk arrays	53
Creating an IBM SAS RAID disk array	53
Deleting an IBM SAS RAID disk array	54
Adding disks to an existing array	55
Migrating an existing disk array to a new RAID level	55
Managing the volatile adapter write cache	58
Hot spare disks	59
Creating hot spare disks	59
Deleting hot spare disks	60
Drive queue depth	60
Linux command-line interface	60
Considerations for solid-state drives (SSD)	61
Multi-initiator and high availability	62

Possible HA configurations	64
Controller functions	65
Controller function attributes	67
Viewing HA controller attributes	68
HA cabling considerations	69
HA performance considerations	69
Configuration and serviceability considerations for HA RAID configurations	69
HA asymmetric access optimization	70
Enabling asymmetric access	73
Asymmetric access status of disk arrays	73
Installing high availability	74
Installing an HA single-system RAID configuration	74
Installing an HA two-system RAID configuration	76
Functions requiring special attention in an HA two-system RAID configuration	78
Installing an HA two-system JBOD configuration	79
IBM SAS RAID controller maintenance	81
Usage tips	81
Updating the controller microcode	82
Rechargeable battery maintenance	82
Displaying rechargeable battery information	82
Forcing a rechargeable battery error	83
Replacing the rechargeable battery pack	84
Replacing a nonconcurrently maintainable battery pack	85
Replacing a concurrently maintainable battery pack	87
Replacing the cache directory card	89
Replacing an SSD module on the PCIe RAID and SSD SAS adapter	91
Replacing a disk in a SAS RAID adapter	94
Physical disks	95
Removing a failed disk	95
Installing a new disk	96
Disk failure recovery	98
RAID 0 failure	98
RAID 5 disk recovery	98
Recovering a RAID 5 single-disk failure	98
RAID 5 multiple-disk failure	98
RAID 6 disk recovery	99
Recovering a RAID 6 single-disk or dual-disk failure	99
RAID 6 failure of three or more disks	99
RAID 10 disk recovery	99
Recovering a RAID 10 single-disk failure	99
RAID 10 multiple-disk failure	100
Viewing SAS fabric path information	100
Viewing SCSI location and physical resource path	101
Reclaiming IOA cache storage	102
Problem determination and recovery	103
Analyzing error logs	104
Basic vi commands	105
Searching logs	105
Samples: Error logs	105
Sample: Generic IOA or device error log	105
Sample: Device configuration error log	106
Sample: Array error log	107
Sample: Cache error log	108
Disk array problem identification	108
Unit reference code tables	109
Maintenance analysis procedures	116
MAP 3300	116
MAP 3310	117
MAP 3311	118
MAP 3312	118
MAP 3313	119

MAP 3320	120
MAP 3321	120
MAP 3330	121
MAP 3331	121
MAP 3332	121
MAP 3333	122
MAP 3334	122
MAP 3335	124
MAP 3337	124
MAP 3340	125
MAP 3341	126
MAP 3342	126
MAP 3343	127
MAP 3344	127
MAP 3345	127
MAP 3346	128
MAP 3347	128
MAP 3348	128
MAP 3349	128
MAP 3350	129
MAP 3351	131
MAP 3352	132
MAP 3353	133
MAP 3390	134
MAP 3410	134
MAP 3411	135
MAP 3412	136
MAP 3413	137
MAP 3420	137
MAP 3421	138
MAP 3430	138
MAP 3431	139
MAP 3432	139
MAP 3433	139
MAP 3434	140
MAP 3435	141
MAP 3440	142
MAP 3441	143
MAP 3442	143
MAP 3443	143
MAP 3444	143
MAP 3445	144
MAP 3446	144
MAP 3447	144
MAP 3448	145
MAP 3449	145
MAP 3450	145
MAP 3452	148
MAP 3453	148
MAP 3454	149
MAP 3460	149
MAP 3461	149
MAP 3490	150
MAP 3495	150

Notices	153
Accessibility features for IBM Power Systems servers	154
Privacy policy considerations	155
Trademarks	156
Electronic emission notices	156
Class A Notices	156

Class B Notices.	160
Terms and conditions.	163

Safety notices

Safety notices may be printed throughout this guide:

- **DANGER** notices call attention to a situation that is potentially lethal or extremely hazardous to people.
- **CAUTION** notices call attention to a situation that is potentially hazardous to people because of some existing condition.
- **Attention** notices call attention to the possibility of damage to a program, device, system, or data.

World Trade safety information

Several countries require the safety information contained in product publications to be presented in their national languages. If this requirement applies to your country, safety information documentation is included in the publications package (such as in printed documentation, on DVD, or as part of the product) shipped with the product. The documentation contains the safety information in your national language with references to the U.S. English source. Before using a U.S. English publication to install, operate, or service this product, you must first become familiar with the related safety information documentation. You should also refer to the safety information documentation any time you do not clearly understand any safety information in the U.S. English publications.

Replacement or additional copies of safety information documentation can be obtained by calling the IBM Hotline at 1-800-300-8751.

German safety information

Das Produkt ist nicht für den Einsatz an Bildschirmarbeitsplätzen im Sinne § 2 der Bildschirmarbeitsverordnung geeignet.

Laser safety information

IBM® servers can use I/O cards or features that are fiber-optic based and that utilize lasers or LEDs.

Laser compliance

IBM servers may be installed inside or outside of an IT equipment rack.

DANGER: When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- If IBM supplied the power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
 - For AC power, disconnect all power cords from their AC power source.
 - For racks with a DC power distribution panel (PDP), disconnect the customer's DC power source to the PDP.
- When connecting power to the product ensure all power cables are properly connected.

- For racks with AC power, connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.
- For racks with a DC power distribution panel (PDP), connect the customer's DC power source to the PDP. Ensure that the proper polarity is used when attaching the DC power and DC power return wiring.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Do not attempt to switch on power to the machine until all possible unsafe conditions are corrected.
- Assume that an electrical safety hazard is present. Perform all continuity, grounding, and power checks specified during the subsystem installation procedures to ensure that the machine meets safety requirements.
- Do not continue with the inspection if any unsafe conditions are present.
- Before you open the device covers, unless instructed otherwise in the installation and configuration procedures: Disconnect the attached AC power cords, turn off the applicable circuit breakers located in the rack power distribution panel (PDP), and disconnect any telecommunications systems, networks, and modems.

DANGER:

- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

To Disconnect:

1. Turn off everything (unless instructed otherwise).
2. For AC power, remove the power cords from the outlets.
3. For racks with a DC power distribution panel (PDP), turn off the circuit breakers located in the PDP and remove the power from the Customer's DC power source.
4. Remove the signal cables from the connectors.
5. Remove all cables from the devices.

To Connect:

1. Turn off everything (unless instructed otherwise).
2. Attach all cables to the devices.
3. Attach the signal cables to the connectors.
4. For AC power, attach the power cords to the outlets.
5. For racks with a DC power distribution panel (PDP), restore the power from the Customer's DC power source and turn on the circuit breakers located in the PDP.
6. Turn on the devices.

Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

(R001 part 1 of 2):

DANGER: Observe the following precautions when working on or around your IT rack system:

- Heavy equipment—personal injury or equipment damage might result if mishandled.
- Always lower the leveling pads on the rack cabinet.
- Always install stabilizer brackets on the rack cabinet.
- To avoid hazardous conditions due to uneven mechanical loading, always install the heaviest devices in the bottom of the rack cabinet. Always install servers and optional devices starting from the bottom of the rack cabinet.
- Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices. In addition, do not lean on rack mounted devices and do not use them to stabilize your body position (for example, when working from a ladder).



- Each rack cabinet might have more than one power cord.
 - For AC powered racks, be sure to disconnect all power cords in the rack cabinet when directed to disconnect power during servicing.
 - For racks with a DC power distribution panel (PDP), turn off the circuit breaker that controls the power to the system unit(s), or disconnect the customer's DC power source, when directed to disconnect power during servicing.
- Connect all devices installed in a rack cabinet to power devices installed in the same rack cabinet. Do not plug a power cord from a device installed in one rack cabinet into a power device installed in a different rack cabinet.
- An electrical outlet that is not correctly wired could place hazardous voltage on the 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.

(R001 part 2 of 2):

CAUTION:

- Do not install a unit in a rack where the internal rack ambient temperatures will exceed the manufacturer's recommended ambient temperature for all your rack-mounted devices.
- Do not install a unit in a rack where the air flow is compromised. Ensure that air flow is not blocked or reduced on any side, front, or back of a unit used for air flow through the unit.
- Consideration should be given to the connection of the equipment to the supply circuit so that overloading of the circuits does not compromise the supply wiring or overcurrent protection. To provide the correct power connection to a rack, refer to the rating labels located on the equipment in the rack to determine the total power requirement of the supply circuit.
- *(For sliding drawers.)* Do not pull out or install any drawer or feature if the rack stabilizer brackets are not attached to the rack. Do not pull out more than one drawer at a time. The rack might become unstable if you pull out more than one drawer at a time.



- *(For fixed drawers.)* This drawer is a fixed drawer and must not be moved for servicing unless specified by the manufacturer. Attempting to move the drawer partially or completely out of the rack might cause the rack to become unstable or cause the drawer to fall out of the rack.

CAUTION:

Removing components from the upper positions in the rack cabinet improves rack stability during relocation. Follow these general guidelines whenever you relocate a populated rack cabinet within a room or building.

- Reduce the weight of the rack cabinet by removing equipment starting at the top of the rack cabinet. When possible, restore the rack cabinet to the configuration of the rack cabinet as you received it. If this configuration is not known, you must observe the following precautions:
 - Remove all devices in the 32U position (compliance ID RACK-001 or 22U (compliance ID RR001) and above.
 - Ensure that the heaviest devices are installed in the bottom of the rack cabinet.
 - Ensure that there are little-to-no empty U-levels between devices installed in the rack cabinet below the 32U (compliance ID RACK-001 or 22U (compliance ID RR001) level, unless the received configuration specifically allowed it.
- If the rack cabinet you are relocating is part of a suite of rack cabinets, detach the rack cabinet from the suite.
- If the rack cabinet you are relocating was supplied with removable outriggers they must be reinstalled before the cabinet is relocated.
- Inspect the route that you plan to take to eliminate potential hazards.
- Verify that the route that you choose can support the weight of the loaded rack cabinet. Refer to the documentation that comes with your rack cabinet for the weight of a loaded rack cabinet.
- Verify that all door openings are at least 760 x 230 mm (30 x 80 in.).
- Ensure that all devices, shelves, drawers, doors, and cables are secure.
- Ensure that the four leveling pads are raised to their highest position.
- Ensure that there is no stabilizer bracket installed on the rack cabinet during movement.
- Do not use a ramp inclined at more than 10 degrees.
- When the rack cabinet is in the new location, complete the following steps:
 - Lower the four leveling pads.
 - Install stabilizer brackets on the rack cabinet.
 - If you removed any devices from the rack cabinet, repopulate the rack cabinet from the lowest position to the highest position.
- If a long-distance relocation is required, restore the rack cabinet to the configuration of the rack cabinet as you received it. Pack the rack cabinet in the original packaging material, or equivalent. Also lower the leveling pads to raise the casters off of the pallet and bolt the rack cabinet to the pallet.

(R002)

(L001)



DANGER: Hazardous voltage, current, or energy levels are present inside any component that has this label attached. Do not open any cover or barrier that contains this label. (L001)

(L002)



DANGER: Rack-mounted devices are not to be used as shelves or work spaces. (L002)

(L003)



or




or



or



or



DANGER: Multiple power cords. The product might be equipped with multiple AC power cords or multiple DC power cables. To remove all hazardous voltages, disconnect all power cords and power cables. (L003)

(L007)



CAUTION: A hot surface nearby. (L007)

(L008)



CAUTION: Hazardous moving parts nearby. (L008)

All lasers are certified in the U.S. to conform to the requirements of DHHS 21 CFR Subchapter J for class 1 laser products. Outside the U.S., they are certified to be in compliance with IEC 60825 as a class 1 laser product. Consult the label on each part for laser certification numbers and approval information.

CAUTION:

This product might contain one or more of the following devices: CD-ROM drive, DVD-ROM drive, DVD-RAM drive, or laser module, which are Class 1 laser products. Note the following information:

- Do not remove the covers. Removing the covers of the laser product could result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.
- Use of the controls or adjustments or performance of procedures other than those specified herein might result in hazardous radiation exposure.

(C026)

CAUTION:

Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. Although shining light into one end and looking into the other end of a disconnected optical fiber to verify the continuity of optic fibers may not injure the eye, this procedure is potentially dangerous. Therefore, verifying the continuity of optical fibers by shining light into one end and looking at the other end is not recommended. To verify continuity of a fiber optic cable, use an optical light source and power meter. (C027)

CAUTION:

This product contains a Class 1M laser. Do not view directly with optical instruments. (C028)

CAUTION:

Some laser products contain an embedded Class 3A or Class 3B laser diode. Note the following information: laser radiation when open. Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam. (C030)

CAUTION:

The battery contains lithium. To avoid possible explosion, do not burn or charge the battery.

Do Not:

- ___ Throw or immerse into water
- ___ Heat to more than 100°C (212°F)
- ___ Repair or disassemble

Exchange only 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. (C003)

CAUTION:

Regarding IBM provided VENDOR LIFT TOOL:

- Operation of LIFT TOOL by authorized personnel only.
- LIFT TOOL intended for use to assist, lift, install, remove units (load) up into rack elevations. It is not to be used loaded transporting over major ramps nor as a replacement for such designated tools like pallet jacks, walkies, fork trucks and such related relocation practices. When this is not practicable, specially trained persons or services must be used (for instance, riggers or movers).
- Read and completely understand the contents of LIFT TOOL operator's manual before using. Failure to read, understand, obey safety rules, and follow instructions may result in property damage and/or personal injury. If there are questions, contact the vendor's service and support. Local paper manual must remain with machine in provided storage sleeve area. Latest revision manual available on vendor's web site.
- Test verify stabilizer brake function before each use. Do not over-force moving or rolling the LIFT TOOL with stabilizer brake engaged.
- Do not move LIFT TOOL while platform is raised, except for minor positioning.
- Do not exceed rated load capacity. See LOAD CAPACITY CHART regarding maximum loads at center versus edge of extended platform.
- Only raise load if properly centered on platform. Do not place more than 200 lb (91 kg) on edge of sliding platform shelf also considering the load's center of mass/gravity (CoG).
- Do not corner load the platform tilt riser accessory option. Secure platform riser tilt option to main shelf in all four (4x) locations with provided hardware only, prior to use. Load objects are designed to slide on/off smooth platforms without appreciable force, so take care not to push or lean. Keep riser tilt option flat at all times except for final minor adjustment when needed.
- Do not stand under overhanging load.
- Do not use on uneven surface, incline or decline (major ramps).
- Do not stack loads.
- Do not operate while under the influence of drugs or alcohol.
- Do not support ladder against LIFT TOOL.
- Tipping hazard. Do not push or lean against load with raised platform.
- Do not use as a personnel lifting platform or step. No riders.
- Do not stand on any part of lift. Not a step.
- Do not climb on mast.
- Do not operate a damaged or malfunctioning LIFT TOOL machine.
- Crush and pinch point hazard below platform. Only lower load in areas clear of personnel and obstructions. Keep hands and feet clear during operation.
- No Forks. Never lift or move bare LIFT TOOL MACHINE with pallet truck, jack or fork lift.
- Mast extends higher than platform. Be aware of ceiling height, cable trays, sprinklers, lights, and other overhead objects.
- Do not leave LIFT TOOL machine unattended with an elevated load.
- Watch and keep hands, fingers, and clothing clear when equipment is in motion.
- Turn Winch with hand power only. If winch handle cannot be cranked easily with one hand, it is probably over-loaded. Do not continue to turn winch past top or bottom of platform travel. Excessive unwinding will detach handle and damage cable. Always hold handle when lowering, unwinding. Always assure self that winch is holding load before releasing winch handle.
- A winch accident could cause serious injury. Not for moving humans. Make certain clicking sound is heard as the equipment is being raised. Be sure winch is locked in position before releasing handle. Read instruction page before operating this winch. Never allow winch to unwind freely. Freewheeling will cause uneven cable wrapping around winch drum, damage cable, and may cause serious injury. (C048)

Power and cabling information for NEBS (Network Equipment-Building System) GR-1089-CORE

The following comments apply to the IBM servers that have been designated as conforming to NEBS (Network Equipment-Building System) GR-1089-CORE:

The equipment is suitable for installation in the following:

- Network telecommunications facilities
- Locations where the NEC (National Electrical Code) applies

The intrabuilding ports of this equipment are suitable for connection to intrabuilding or unexposed wiring or cabling only. The intrabuilding ports of this equipment *must not* be metalically connected to the interfaces that connect to the OSP (outside plant) or its wiring. These interfaces are designed for use as intrabuilding interfaces only (Type 2 or Type 4 ports as described in GR-1089-CORE) and require isolation from the exposed OSP cabling. The addition of primary protectors is not sufficient protection to connect these interfaces metalically to OSP wiring.

Note: All Ethernet cables must be shielded and grounded at both ends.

The ac-powered system does not require the use of an external surge protection device (SPD).

The dc-powered system employs an isolated DC return (DC-I) design. The DC battery return terminal *shall not* be connected to the chassis or frame ground.

The dc-powered system is intended to be installed in a common bonding network (CBN) as described in GR-1089-CORE.

SAS RAID controllers for Linux

The Peripheral Component Interconnect-X (PCI-X) and PCI Express (PCIe) serial-attached SCSI (SAS) RAID controller is available for various versions of the Linux kernel. Here you will find out how to use and maintain the controller.

General information

This section provides general information about the IBM SAS RAID controllers for Linux.

The controllers have the following features:

- PCI-X 266 system interface or PCI Express (PCIe) system interface.
 - Physical link speed of 3 Gb per second serial-attached SCSI (SAS) supporting transfer rates of 300 MB per second on PCI-X and PCIe controllers.
 - Physical link speed of 6 Gb per second SAS supporting transfer rates of 600 MB per second on PCI Express 2.0 (PCIe2) controllers.
 - Supports SAS devices and non-disk Serial Advanced Technology Attachment (SATA) devices.
 - Optimized for SAS disk configurations which utilize dual paths thru dual expanders for redundancy and reliability.
 - Controller managed path redundancy and path switching for multiported SAS devices.
 - Embedded PowerPC[®] RISC processor, hardware XOR DMA engine, and hardware Finite Field Multiplier (FFM) DMA Engine (for Redundant Array of Independent Disks (RAID) 6).
 - Supports nonvolatile write cache for RAID disk arrays on some adapters.
 - Support for RAID 0, 5, 6, 10, 5T2, 6T2, and 10T2 disk arrays.
 - Supports attachment of other devices such as non-RAID disks, tape, and optical devices.
 - RAID disk arrays and non-RAID devices supported as a bootable device.
 - Advanced RAID features:
 - Hot spares for RAID 5, 6, 10, 5T2, 6T2, and 10T2 disk arrays
 - Ability to increase the capacity of an existing RAID 5 or 6 disk array by adding disks
 - Background parity checking
 - Background data scrubbing
 - Disks formatted to 528 or 4224 bytes per sector, providing cyclical redundancy checking (CRC) and logically bad block checking on PCI-X and PCIe controller
 - Disks formatted to 528 or 4224 bytes per sector, providing SCSI T10 standardized data integrity fields along with logically bad block checking on PCIe2 controllers
 - Optimized hardware for RAID 5 and 6 sequential write workloads
 - Optimized skip read/write disk support for transaction workloads
 - Support for a maximum of 240 advanced function disks with a total device support maximum of 1023 on PCIe2 controllers
- Note:** The number of all physical SAS and SATA devices plus the number of logical RAID disk arrays must be less than 1023 per controller.
- Supports a maximum of 64 advanced function disks with a total device support maximum of 255 (the number of all physical SAS and SATA devices plus the number of logical RAID disk arrays must be less than 255 per controller).

Note: This information refers to various hardware and software features and functions. The realization of these features and functions depends on the limitations of your hardware and software. The Linux

operating system supports all functions mentioned. If you are using another operating system, consult the appropriate documentation for that operating system regarding support for the mentioned features and functions.

Related reference:

“Related information” on page 23

Many other sources of information about the Linux operating system, RAID, and other associated topics are available.

“References to Linux” on page 22

Three different versions of the Linux operating system are referred to in this topic collection.

Feature comparison of SAS RAID cards

Compare the main features of PCI-X, PCI Express (PCIe), PCIe2, and PCIe3 SAS RAID cards.

These tables provides a breakdown of the main features of the SAS RAID PCI-X and PCIe controller cards.

PCI-X SAS RAID card comparison

This table compares the main features of PCI-X SAS RAID cards.

Table 1. PCI-X SAS RAID controller cards.

CCIN (custom card identification number)	2BD9	2BE0	2BE1	572A	572B	572C	572F / 575C	57B8
Description	PCI-X266 Planar 3 Gb SAS RAID Adapter (RAID/Cache Storage Controller)	PCI-X266 Planar 3 Gb SAS Adapter (RAID-10 Storage Controller)	PCI-X266 Planar 3 Gb SAS RAID Adapter (RAID/Cache Enablement)	PCI-X 266 Ext Dual-x4 3 Gb SAS Adapter	PCI-X 266 Ext Dual-x4 3 Gb SAS RAID Adapter	PCI-X 266 Planar 3 Gb SAS Adapter	PCI-X 266 Ext Tri-x4 3 Gb SAS RAID Adapter	PCI-X 266 Planar 3 Gb SAS RAID Adapter
Form factor	Planar unique 64-bit PCI-X	Planar unique 64-bit PCI-X	Planar RAID enablement	Low profile 64-bit PCI-X	Long 64-bit PCI-X	Planar integrated	Long 64-bit PCI-X, double-wide card set	Planar RAID enablement
Adapter failing function code LED value	2D18	2D16	2D17	2515	2517	2502	2519 / 251D	2505
Physical links	6 (two 2x wide ports to shared SAS Drives and one 2x wide port to 2BE1 adapter)	3 (Direct attach SAS Drives)	8 (two 2x wide ports to shared SAS Drives, one 2x wide port to 2BD9 adapter, one phy to DVD, and optionally one physical link to Tape Drive)	8 (two mini SAS 4x connectors)	8 (two mini SAS 4x connectors)	8 ¹	12 (bottom 3 mini SAS 4x connectors) and 2 (top mini SAS 4x connector for HA only)	8 ¹
RAID levels supported	RAID 0, 5, 6, 10	RAID 0, 5 ⁴ , 10	RAID 0, 5, 6, 10	RAID 0, 5 ⁴ , 6 ⁴ , 10	RAID 0, 5, 6, 10	RAID 0	RAID 0, 5, 6, 10	RAID 0, 5, 6, 10
Adding disks to an existing disk array supported RAID levels	RAID 5, 6	RAID 5	RAID 5, 6	RAID 5, 6	RAID 5, 6		RAID 5, 6	RAID 5, 6
Write cache size	175 MB		175 MB		175 MB		Up to 1.5 Gb (compressed)	175 MB
Read cache size							Up to 1.6 Gb (compressed)	
Cache battery pack technology	Lilon		Lilon		Lilon		Lilon	Not applicable ²
Cache battery pack FFC	2D1B		2D1B		2D03		2D06 ⁵	Not applicable ²

Table 1. PCI-X SAS RAID controller cards (continued).

CCIN (custom card identification number)	2BD9	2BE0	2BE1	572A	572B	572C	572F / 575C	57B8
Cache battery concurrent maintenance	No	No	No	No	No	No	Yes	Not applicable ²
Cache data present LED	Yes	No	Yes	No	No	No	No	No
Removable cache card	No	No	No	No	Yes	No	No	No
Auxiliary write cache (AWC) support	No	No	No	No	No	No	Yes	Yes
High availability (HA) two system RAID	Yes	No	Yes	Yes ³	Yes	No	No	No
HA two system JBOD	No	No	No	Yes ³	No	No	No	No
HA single system RAID	Yes	No	Yes	Yes ³	Yes	No	No	No
Requires HA RAID configuration	Yes	No	Yes	No	Yes	No	No	No
JBOD support	No	Yes	No	Yes	No	Yes	No	Yes

¹ Some systems provide an external mini SAS 4x connector from the integrated backplane controller.

² The controller contains battery-backed cache, but the battery power is supplied by the 57B8 controller through the backplane connections.

³ Multi-initiator and high availability is supported on the CCIN 572A adapter except for part numbers of either 44V4266 or 44V4404 (feature code 5900).

⁴ The write performance of RAID 5 and RAID 6 might be poor on adapters that do not provide write cache. Consider using an adapter that provides write cache when using RAID 5 or RAID 6, or use solid-state drives (SSDs) where supported to improve write performance.

⁵ The cache battery pack for both adapters is contained in a single battery field-replaceable unit (FRU), which is physically located on the 575C auxiliary cache card.

⁶ Nonvolatile write cache is supported for RAID disk arrays only.




Figure 1. CCIN 572A PCI-X266 External Dual-x4 3 Gb SAS adapter




Figure 2. CCIN 572B PCI-X266 Ext Dual-x4 3 Gb SAS RAID adapter




Figure 3. CCIN 572F PCI-X266 Ext Tri-x4 3 Gb SAS RAID adapter and CCIN 575C PCI-X266 auxiliary cache adapter




Figure 4. CCIN 57B8 planar RAID enablement card

RAID/Cache Enablement




Figure 5. CCIN 2BE1 PCI-X266 Planar 3 Gb SAS Adapter

Planar RAID 10 Storage Controller





Figure 6. CCIN 2BE0 PCI-X266 Planar 3 Gb SAS Adapter

Planar RAID/Cache Storage Controller



P7EJ531-1

Figure 7. CCIN 2BD9 PCI-X266 Planar 3 Gb SAS RAID Adapter

PCIe SAS RAID card comparison

Use the tables provided here to compare the main features of PCI Express (PCIe) SAS RAID cards.

Table 2. PCIe SAS RAID controller cards.

CCIN (custom card identification number)	2B4C	2B4F	574E	57B3	57B7	57B9
Description	PCIe x4 Internal 3 Gb SAS RAID Adapter	PCIe x4 Internal 3 Gb SAS Adapter	PCIe x8 Ext Dual-x4 3 Gb SAS RAID Adapter	PCIe x8 Ext Dual-x4 3 Gb SAS Adapter	PCIe x1 Auxiliary Cache Adapter	PCIe x8 Ext Dual-x4 3 Gb SAS Adapter and Cable Card
Form factor	Planar-unique PCIe	Planar-unique PCIe	PCIe x8	PCIe x8	Planar Auxiliary Cache	Combination PCIe x8 and Cable Card
Adapter failing function code LED value	2D28	2D27	2518	2516	2504	2D0B
Physical links	6 (two 2x wide ports to shared SAS drives and one 2x wide port to 57CB)	3 (Directly attached SAS drives)	8 (two mini SAS 4x connectors)	8 (two mini SAS 4x connectors)	2	4 (bottom mini-SAS 4x connector required to connect by an external AI cable to the top mini-SAS 4x Cable Card connector)
RAID levels supported	RAID 0, 5, 6, 10	RAID 0, 5 ¹ , 10	RAID 0, 5, 6, 10	RAID 0, 5 ¹ , 6 ¹ , 10		RAID 0, 5 ¹ , 6 ¹ , 10
Adding disks to an existing disk array supported RAID levels	RAID 5, 6	RAID 5	RAID 5, 6	RAID 5, 6		RAID 5, 6
Write cache size ⁴	175 MB		380 MB		175 MB	
Read cache size						
Cache battery pack technology	LiIon		LiIon		LiIon	
Cache battery pack FFC	2D1B		2D0E		2D05	
Cache battery concurrent maintenance	No	No	Yes	No	Yes	No

Table 2. PCIe SAS RAID controller cards (continued).

CCIN (custom card identification number)	2B4C	2B4F	574E	57B3	57B7	57B9
Cache data present LED	Yes	No	Yes	No	Yes	No
Removable cache card	No	No	Yes	No	No	No
Auxiliary write cache (AWC) support	No	No	No	No	Yes	No
High availability (HA) two system RAID	Yes	No	Yes	Yes	No	No
HA two system JBOD	No	No	No	Yes	No	No
HA single system RAID	Yes	No	Yes	Yes	No	No
Requires HA RAID configuration	Yes	No	Yes	No	No	No
JBOD support	No	Yes	No	Yes	No	Yes

¹ The write performance of RAID 5 and RAID 6 might be poor on adapters that do not provide write cache. Consider using an adapter that provides write cache when using RAID 5 or RAID 6 or use solid-state drives (SSDs) where supported to improve write performance.

² RAID 0 requires the use of AIX® logical volume manager (LVM) mirroring.

³ The integrated SSDs are allowed to run in JBOD mode.

⁴ Nonvolatile write cache is supported only for RAID disk arrays.

Table 3. PCIe SAS RAID controller cards.

CCIN (custom card identification number)	57BA	57C7	57CB	57CC	57CD	57CF
Description	PCIe x8 Ext Dual-x4 3 Gb SAS Adapter and Cable Card	PCI Express x8 Planar 3 Gb SAS Adapter (Disk/Media Backplane)	PCIe x4 Planar 3 Gb SAS RAID Adapter	PCIe x8 Internal 3 Gb SAS Adapter	PCIe SAS RAID and SSD Adapter 3 Gb x8	PCI Express x8 Planar 3 Gb SAS RAID Adapter (with 175 MB Cache RAID - Dual IOA Enablement Card)
Form factor	Combination PCIe x8 and Cable Card	Planar	Planar cache enablement	PCIe unique to IBM PureSystems®	Double Wide PCIe x8 with 1 - 4 integrated SSDs	Planar and Enablement Card
Adapter failing function code LED value	2D0B	2D14	2D26	2D29	2D40	2D15
Physical links	8 (two mini SAS 4x connectors, one required to connect by an external AI cable to the top mini-SAS 4x Cable Card connector)	8	8 (two 2x wide ports to shared SAS drives and one 2x wide port to 2B4F, one physical link to DVD, and optionally one physical link to tape drive)	1	4 (one direct SAS physical link to each integrated SSD)	8
RAID levels supported	RAID 0, 5 ¹ , 6 ¹ , 10	RAID 0, 5 ¹ , 6 ¹ , 10	RAID 0, 5, 6, 10	RAID 0	RAID 0 ² , 5, 6	RAID 0, 5, 6, 10
Adding disks to an existing disk array supported RAID levels	RAID 5, 6	RAID 5, 6	RAID 5, 6		RAID 5, 6	RAID 5, 6
Write cache size ⁴			175 MB			175 MB
Read cache size						
Cache battery pack technology			LiIon			LiIon
Cache battery pack FFC			2D1B			2D19
Cache battery concurrent maintenance	No	No	No	No	No	Yes
Cache data present LED	No	No	Yes	No	No	Yes
Removable cache card	No	No	No	No	No	No
Auxiliary write cache (AWC) support	No	No	No	No	No	No

Table 3. PCIe SAS RAID controller cards (continued).

CCIN (custom card identification number)	57BA	57C7	57CB	57CC	57CD	57CF
High availability (HA) two system RAID	No	No	Yes	No	No	Yes
HA two system JBOD	No	No	No	No	No	No
HA single system RAID	No	No	Yes	No	No	Yes
Requires HA RAID configuration	No	No	Yes	No	No	Yes
JBOD support	Yes	Yes	No	Yes	Yes ³	No

¹ The write performance of RAID 5 and RAID 6 might be poor on adapters that do not provide write cache. Consider using an adapter that provides write cache when using RAID 5 or RAID 6 or use solid-state drives (SSDs) where supported to improve write performance.

² RAID 0 requires the use of AIX logical volume manager (LVM) mirroring.

³ The integrated SSDs are allowed to run in JBOD mode.

⁴ Nonvolatile write cache is supported only for RAID disk arrays.




Figure 8. CCIN 574E PCIe x8 Ext Dual-x4 3 Gb SAS RAID adapter (disk/media backplane)




Figure 9. CCIN 57B3 PCIe x8 Ext Dual-x4 3 Gb SAS adapter




Figure 10. CCIN 57B7 Planar auxiliary cache




Figure 11. CCIN 57B9 PCIe x8 Ext Dual-x4 3 Gb SAS adapter and cable card




Figure 12. CCIN 57BA PCIe x8 Ext Dual-x4 3 Gb SAS adapter and cable card




Figure 13. CCIN 57C7 PCI Express x8 Planar 3 Gb SAS adapter (with 175 MB Cache RAID - Dual IOA Enablement card)




Figure 14. CCIN 57CD PCIe SAS RAID and 3 Gb x8 SSD adapter




Figure 15. CCIN 57CF PCI Express x8 Planar 3 Gb SAS RAID adapter (with 175 MB Cache RAID - Dual IOA Enablement card)

Planar RAID/Cache Storage Controller




Figure 16. CCIN 2B4C PCIe x4 internal 3 Gb SAS adapter

Planar RAID 10 Storage Controller




Figure 17. CCIN 2B4F PCIe x4 internal 3Gb SAS RAID adapter

PCIe2 SAS RAID card comparison

This table compares the main features of PCI Express 2.0 (PCIe2) SAS RAID cards.

Table 4. PCIe2 SAS RAID controller cards.

CCIN (custom card identification number)	57B5 ¹	57BB	57C4 ²	57C3
Description	PCIe2 1.8 GB Cache RAID SAS adapter Tri-port 6 Gb	PCIe2 1.8 GB Cache RAID SAS adapter Tri-port 6 Gb	PCIe2 RAID SAS adapter Dual-port 6 Gb	PCIe2 3.1GB Cache RAID SAS enclosure 6Gb x8
Form factor	PCIe2 x8	PCIe2 x8	PCIe2 x8	PCIe2 x8 Storage Enclosure
Adapter failing function code LED value	2D20	2D1F	2D1D	2D24
Physical links	11 (three mini SAS HD 4x connectors with top connector containing three physical links)	11 (three mini SAS HD 4x connectors with top connector containing three physical links)	8 (two mini SAS HD 4x connectors)	11 internally integrated with two external mini SAS HD 4x connectors, each containing three physical links
RAID levels supported	RAID 0, 5, 6, 10	RAID 0, 5, 6, 10	RAID 0, 5, 6, 10	RAID 0, 5, 6, 10
Adding disks to an existing disk array supported RAID levels	None ³	None ³	None ³	None ³
Write cache size	1.8 GB	1.8 GB		3.1 GB
Read cache size				
Cache battery pack technology	None (uses supercapacitor technology)	None (uses supercapacitor technology)		None (uses supercapacitor technology)
Cache battery pack FFC				
Cache battery concurrent maintenance				
Cache data present LED				
Removable cache card				
Auxiliary write cache (AWC) support	No	No	No	No
High availability (HA) two-system RAID	Yes	Yes	Yes	Yes

Table 4. PCIe2 SAS RAID controller cards (continued).

CCIN (custom card identification number)	57B5 ¹	57BB	57C4 ²	57C3
HA two-system JBOD	No	No	No	No
HA single-system RAID	Yes	Yes	Yes	Yes
Requires HA RAID configuration	Yes	Yes	No	Yes
JBOD SAS disk support	No	No	No	No
SAS tape support	No	No	No	No
SAS DVD support	No	No	No	No
Native 4K block device support	No	No	No	No

Notes:

1. Feature 5913 (CCIN 57B5) adapters installed in POWER6 servers must be placed in I/O expansion units. Feature 5913 (CCIN 57B5) adapters are not supported in POWER6 system units. Feature 5913 (CCIN 57B5) is supported for the full set of SAS adapter functions on POWER6 servers with the exception of controlling boot drives or load source drives.
2. Feature ESA1 or ESA2 (CCIN 57C4) only supports connecting to SSD devices.
3. You must stop and restart the array to add additional disk units.






Figure 18. CCIN 57B5 PCIe2 1.8 GB cache RAID SAS adapter Tri-port 6 Gb



P7EJ523-0

Figure 19. CCIN 57C3 PCIe2 3.1 GB Cache RAID SAS enclosure 6 Gb x8



P7EJ522-0

Figure 20. CCIN 57C4 PCIe2 RAID SAS adapter Dual-port 6 Gb




Figure 21. CCIN 57BB PCIe2 1.8 GB cache RAID SAS adapter Tri-port 6 Gb

PCIe3 SAS RAID card comparison

This table compares the main features of PCI Express 3.0 (PCIe3) SAS RAID cards.

Table 5. PCIe3 SAS RAID controller cards.

CCIN (custom card identification number)	57B4	57CE	57D7	57D8	2CCA	2CCD	2CD2	57B1
Description	PCIe3 RAID SAS adapter Quad-port 6 Gb x8	PCIe3 12 GB Cache RAID SAS adapter Quad-port 6 Gb x8	PCIe3 x8 SAS RAID internal adapter 6 Gb	PCIe3 x8 Cache SAS RAID internal adapter 6 Gb	PCIe3 x8 Cache SAS RAID internal adapter 6 Gb	PCIe3 x8 SAS RAID internal adapter 6 Gb	PCIe3 x8 SAS RAID internal adapter 6 Gb	PCIe3 12 Gb Cache RAID+ SAS adapter Quad-port 6 Gb
Form factor	PCIe3 x8	PCIe3 x8	Planar-unique PCIe3 x8	Planar-unique PCIe3 x8	Planar integrated	Planar integrated	Planar integrated	PCIe3 x8
Adapter failing function code LED value	2D11	2D21	2D35	2D36	2509	250D	250B	2D22
Physical links	16 (four mini SAS HD 4x connectors)	16 (four mini SAS HD 4x connectors)	16 (internal connection to directly attached SAS drives)	16 (internal connections to directly attached SAS drives and remote adapter link) and 4 (one mini SAS HD 4x connector for external SAS attach)	20 internally integrated	13 internally integrated	20 internally integrated	16 (four mini SAS HD 4x connectors)
Supported RAID levels	RAID 0, 5, 6, 10	RAID 0, 5, 6, 10, 5T2, 6T2, and 10T2	RAID 0, 5, 6, 10	RAID 0, 5, 6, 10, 5T2, 6T2, and 10T2	RAID 0, 5, 6, 10, 5T2, 6T2, and 10T2	RAID 0, 5, 6, 10, and 10T2	RAID 0, 5, 6, 10, 5T2, 6T2, and 10T2	RAID 0, 5, 6, 10, 5T2, 6T2, and 10T2

Table 5. PCIe3 SAS RAID controller cards (continued).

CCIN (custom card identification number)	57B4	57CE	57D7	57D8	2CCA	2CCD	2CD2	57B1
Adding disks to an existing disk array	None ⁵	None ⁵	None ⁵	None ⁵	None ⁵	None ⁵	None ⁵	None ⁵
Write cache size	Up to 1 GB (compressed) (disabled for HA RAID configurations)	Up to 12 GB (compressed)	Up to 1 GB (compressed)	Up to 7.2 GB (compressed)	Up to 7.2 GB (compressed)			Up to 7.2 GB (compressed). See
Read cache size								
Cache battery pack technology	None (uses synchronize cache)	None (uses supercapacitor technology)	None (uses synchronize cache)	None (uses supercapacitor technology)	None (uses supercapacitor technology)			None (uses supercapacitor technology)
High-availability (HA) two-system RAID	Yes	Yes	No	Yes	Yes	Yes	No	No
HA two-system JBOD	No	No	No	No	No	No	No	No
Auxiliary write cache (AWC) support	No	No	No	No	No	No	No	No
HA single-system RAID	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Requires HA RAID configuration	No	Yes	No	Yes	Yes	No	Yes	Yes
JBOD SAS disk support	Yes ²	No	Yes ²	No	No	Yes	No	No
SAS tape support	Yes ¹	No	No	No	No	No	No	No
SATA DVD support	Yes ^{1, 3}	No	Yes	Yes	Yes	Yes	Yes	No
Native 4K block device support	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Easy Tier [®] function	No	Yes	No	Yes	Yes	Yes	Yes	Yes

Note:

1. SAS tape and SATA DVD is only supported in a single adapter configuration and cannot be mixed with SAS disk on the same adapter.
2. JBOD is not supported on SSDs.
3. SATA DVD is supported on all CCIN 57B4 adapters, except for those with initial part numbers of either 00FX843, 00MH900, 00FX846, or 00MH903.
4. The 8247-21L, 8247-22L, and 8247-42L systems support volatile adapter write cache feature with the 57D7 integrated SAS controllers or the 57B4 PCIe3 RAID SAS adapter. For more information, see “Managing the volatile adapter write cache” on page 58.
5. You must stop and restart the disk array to add additional disk units.




Figure 22. PCIe3 x8 Cache SAS RAID internal adapter 6 Gb (CCIN can be 2CCA, 2CCD, and 2CD2)




Figure 23. CCIN 57B1 PCIe3 12 GB cache RAID+ SAS adapter Quad-port 6 Gb x8




Figure 24. CCIN 57B4 PCIe3 RAID SAS adapter Quad-port 6 Gb x8, four units




Figure 25. CCIN 57B4 PCIe3 RAID SAS adapter Quad-port 6 Gb x8, two units




Figure 26. CCIN 57CE PCIe3 12 GB cache RAID SAS adapter Quad-port 6 Gb x8




Figure 27. CCIN 57D7 PCIe3 x8 SAS RAID internal adapter 6 Gb




Figure 28. CCIN 57D8 PCIe3 x8 Cache SAS RAID internal adapter 6 Gb for 8286-41A or 8286-42A systems




Figure 29. CCIN 57D8 PCIe3 x8 Cache SAS RAID internal adapter 6 Gb for 5148-21L, 5148-22L, 8247-21L, 8247-22L, or 8284-22A systems

References to Linux

Three different versions of the Linux operating system are referred to in this topic collection.

References to the Linux operating system in this topic collection include the Linux kernel versions 2.6, and later, SUSE Linux Enterprise Server 10, SUSE Linux Enterprise Server 11, Red Hat Enterprise Linux 4, Red Hat Enterprise Linux 5, and Red Hat Enterprise Linux 6. Make sure you are consulting the appropriate section of this topic collection for the operating system you are using.

This document might describe hardware features and functions. While the hardware supports them, the realization of these features and functions depends upon support from the operating system. The Linux operating system 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 information

Many other sources of information about the Linux operating system, RAID, and other associated topics are available.

The following publications contain related information:

- System unit documentation for information specific to your hardware configuration
- The IPR Linux device driver website
- *RS/6000® eServer™ pSeries Adapters, Devices, and Cable Information for Multiple Bus Systems*, order number SA38-0516 (<http://www-01.ibm.com/support/docview.wss?uid=pub1sa38051616>)
- The Linux Documentation Project website
- Linux for IBM eServer pSeries website
- *RS/6000 eServer pSeries Diagnostic Information for Multiple Bus Systems*, order number SA38-0509 (<http://www-01.ibm.com/support/docview.wss?uid=pub1sa38050923>)
- *The RAIDbook: A Handbook of Storage Systems Technology*, Edition 6, Editor: Paul Massiglia
- OpenPOWER website and IBM developerWorks® website for information about Linux on PowerPC architecture

SAS overview

The term *Serial-attached SCSI (SAS)* refers to a set of serial device interconnect and transport protocols.

This set of protocols defines the rules for information exchange between devices. SAS is an evolution of the parallel SCSI device interface into a serial point-to-point interface.

SAS physical links are a set of four wires used as two differential signal pairs. One differential signal transmits in one direction while the other differential signal transmits in the opposite direction. Data can be transmitted in both directions simultaneously.

Physical links are contained in *ports*. A port contains one or more physical links. A port is a *wide* port if there are multiple physical links in the port. A port is a *narrow* port if there is only one physical link in the port. A port is identified by a unique SAS worldwide name (also called a SAS address). A SAS controller contains one or more SAS ports.

A *path* is a logical point-to-point link between a SAS initiator port in the controller and a SAS target port in the I/O device (for example, a disk).

A *connection* is a temporary association between a controller and an I/O device through a path. A connection enables communication to a device. The controller can communicate to the I/O device over this connection using either the SCSI command set or the ATA/ATAPI command set, depending on the device type.

An *expander* facilitates connections between a controller port and multiple I/O device ports. An expander routes connections between the expander ports. There is only a single connection through an expander at any given time. Using expanders creates more nodes in the path from the controller to the I/O device.

If an I/O device supports multiple ports, then it is possible to have more than one path to the device when there are expander devices on the path.

A *SAS fabric* refers to the summation of all paths between all controller ports and all I/O device ports in the SAS subsystem.

SAS architecture summary

Elements that interact to enable the structure of the SAS architecture include controllers, ports, and expanders.

The following points are applicable to this description of general SAS architecture:

- A SAS fabric describes all possible paths between all SAS controllers and I/O devices including cables, enclosures and expanders.
- A SAS controller, expander and I/O device contains one or more SAS ports.
- A SAS port contains one or more physical links.
- A SAS path is a logical connection between a SAS controller port and I/O device ports.
- SAS devices use the SCSI command set and SATA devices use the ATA/ATAPI command set.




Figure 30. Example of the SAS subsystem

The example of a SAS subsystem in the preceding figure illustrates some general concepts.

This controller has eight SAS physical link connections. Four of those physical links are connected into two different wide ports. (One connector contains four physical links grouped into two ports; the connectors signify a physical wire connection.) The four-physical link connector can contain between one and four ports depending on the type of cabling used.

The uppermost port in the figure shows a controller wide port number 6 consisting of physical link numbers 6 and 7. Port 6 connects to an expander which attaches to one of the I/O devices dual ports.

The dashed red line indicates a path between the controller and an I/O device. There is another path from the controller's port number 4 to the other port of the I/O device. These two paths provide two

different possible connections for increased reliability by using redundant controller ports, expanders and I/O device ports. The SCSI Enclosure Services (SES) is a component of each expander.

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, the technique of storing data across a group of disks 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 file systems on it.

A tier is a grouping of physical disks within an Easy Tier disk array all of which have the same performance characteristics. For example, an Easy Tier disk array may contain a tier of SSDs and a tier of HDDs. A data band is the block of data in an Easy Tier disk array that is being analyzed for I/O activity. This data band is the block of data that may move between tiers to better match the I/O activity within the band with the performance characteristics of the tier. The size of the data band can be 1 MB to 8 MB in size depending on the configuration of the Easy Tier disk array.

The SAS controller and I/O devices are managed by the `iprconfig` utility. The `iprconfig` utility is the interface to the RAID configuration, monitoring, and recovery features of the controller and I/O devices.

If a disk array is to be used as the boot device, it might be necessary 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.




Figure 31. Disk array configuration

Easy Tier function

Easy Tier function works with specific RAID levels (i.e., 5T2, 6T2, and 10T2) that support grouping disks with different performance characteristics but similar RAID block formats into tiers within a single array. The Easy Tier function automatically optimizes storage performance across the tiers by moving the physical data placement between the tiers while keeping the external disk view of the Disk array logical block locations unchanged. The Easy Tier function logically divides the Disk array into data bands and continually analyzes the I/O activity in each band. Based on the current I/O activity in each band, the Easy Tier function optimizes performance and resource utilization by automatically and non-disruptively swapping data bands between physical disk tiers containing the most appropriate performance characteristics for the current I/O activity in the band (for example, moving the hottest data to fastest tier.) The tiers are automatically organized such that the best performing tier aligns with disk array LBA 0 (the beginning of the array) when a new array is created before any of the data bands become swapped. It is important to note that a hot spare disk only replaces a disk in the tier that has the similar performance characteristics as the hot spare. Therefore, you need different hot spare disks to fully cover all tiers in a tiered RAID level. For example, an SSD hot spare and a HDD hot spare.

Easy Tier function supports tiers with different performance characteristics by using the following disk drive technologies:

- SSDs that have a high write endurance
- Read Intensive (RI) SSDs that are intended to be used for read-intensive workloads
- HDDs or Enterprise Nearline (ENL) HDDs

A tiered RAID array might be created with the following combinations of disk drive technologies:

- SSDs and HDDs
- RI SSDs and HDDs
- SSDs and ENL HDDs
- RI SSDs and ENL HDDs

When SSDs are used with HDDs in a tiered RAID array, hot data is the frequently accessed read data and write data, and will be moved to the SSDs. However, when RI SSDs are used with HDDs in a tiered RAID array, hot data is only the frequently accessed read data and will be moved to the RI SSDs, while frequently accessed write data will be moved to the HDDs. This policy enables RI SSDs to maintain the reliability for a long time even when write-intensive workloads exist. When using RAID adapters that have a write cache, write performance is likely to be very good, irrespective of whether the write data is placed on SSDs, RI SSDs, or HDDs.

Notes:

- All tiers in the Easy Tier array must contain devices with the same block size. All the SSDs and HDDs in the array must either be 528 bytes per sector or 4224 bytes per sector.
- Each tier in an Easy Tier array must contain at least 10% of the total disk capacity. For more information, see “Estimating disk array capacities” on page 34.

Supported RAID levels

The level of a disk array refers to the manner in which data is stored on the disk and the level of protection that is provided.

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 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 SAS RAID controller supports RAID 0, 5, 6, and 10. Not all controllers support all RAID levels. See the “Feature comparison of SAS RAID cards” on page 2 table for more information. Each RAID level supported by the SAS RAID controller has its own attributes and uses a different method of writing data. The following information details each supported RAID level.

Related concepts:

“RAID 0”

RAID 0 stripes data across the disks in the array for optimal performance.

“RAID 5” on page 28

RAID 5 stripes data across all disks in the array.

“RAID 6” on page 29

RAID 6 stripes data across all disks in the array.

“RAID 10” on page 30

RAID 10 uses mirrored pairs to redundantly store data.

RAID 0

RAID 0 stripes data across the disks in the array for optimal performance.

For a RAID 0 array of three disks, data would be written in the following pattern.




Figure 32. RAID 0

RAID 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 0 array as degraded as the result of a disk failure. If a physical disk fails in a RAID 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.

Related concepts:

“Supported RAID levels” on page 27

The level of a disk array refers to the manner in which data is stored on the disk and the level of protection that is provided.

RAID 5

RAID 5 stripes data across all disks in the array.

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




Figure 33. RAID 5

If a disk fails in a RAID 5 array, you can continue to use the array normally. A RAID 5 array operating with a single failed disk is said to be operating in degraded mode. Whenever data is read from a degraded disk array, the array controller recalculates the data on the failed disk by using data and parity blocks on the operational disks. If a second disk fails, the array will be placed in the failed state and will not be accessible.

Related concepts:

“Supported RAID levels” on page 27

The level of a disk array refers to the manner in which data is stored on the disk and the level of protection that is provided.

RAID 6

RAID 6 stripes data across all disks in the array.

In addition to data, RAID 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 6 array of four disks, array data and parity information are written in the following pattern:




Figure 34. RAID 6

If one or two disks fail in a RAID 6 array, you can continue to use the array normally. A RAID 6 array operating with 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 disks by using data and parity blocks on the operational disks. A RAID 6 array with a single failed disk has similar protection to that of a RAID 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.

Related concepts:

“Supported RAID levels” on page 27

The level of a disk array refers to the manner in which data is stored on the disk and the level of protection that is provided.

RAID 10

RAID 10 uses mirrored pairs to redundantly store data.

The array must contain an even number of disks. Two is the minimum number of disks that you need to create a RAID 10 array. A two-disk RAID 10 array is equal to a RAID 1 array. The data is striped across the mirrored pairs. For example, a RAID 10 array of four disks would have data written to it in the following pattern:




Figure 35. RAID 10

RAID 10 can tolerate multiple disk failures. If one disk in each mirrored pair fails, the array is still 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.

Related concepts:

“Supported RAID levels” on page 27

The level of a disk array refers to the manner in which data is stored on the disk and the level of protection that is provided.

RAID 5T2

Learn how data is written to a RAID 5T2 array when using the Easy Tier function.

RAID 5T2 is a RAID level that provides RAID 5 protection when using the Easy Tier function utilizing two different tiers of physical disk that have unique performance characteristics. Each tier functions as a single redundancy group and stripes data across all disks in the tier. Each tier is RAID 5 protected and writes array parity data across all the disks in the tier. For a RAID 5T2 array that has one tier of three SSD disks and another tier of four HDD disks, array data and parity information is written in the following pattern:

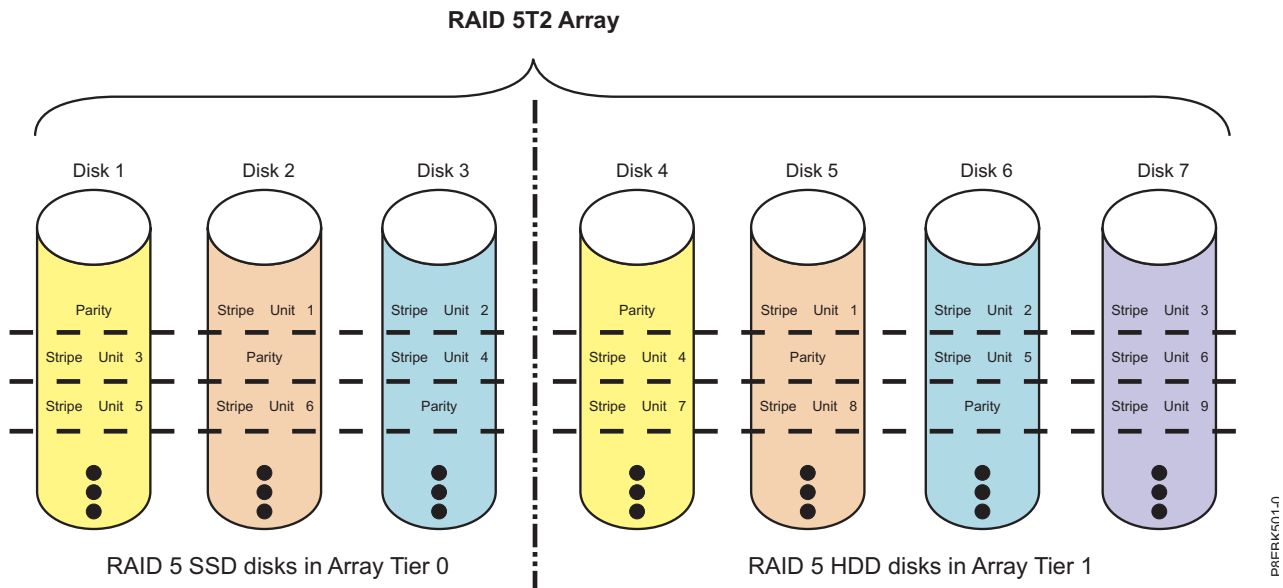


Figure 36. RAID 5T2

If a disk fails in either RAID 5 tier, you can continue to use the entire array. Each tier can contain a failed disk and the array continues to function. A RAID 5T2 array that is operating with a single failed disk in either or both tiers is said to be operating in degraded mode. Whenever data is read from a degraded disk array, the array controller recalculates the data on the failed disk by using data and parity blocks on the operational disks. If a second disk fails in either tier, the entire array is placed in the failed state and is not accessible.

RAID 6T2

Learn how data is written to a RAID 6T2 array when using the Easy Tier function.

RAID 6T2 is a RAID level that provides RAID 6 protection when using the Easy Tier function utilizing two different tiers of physical disk that have unique performance characteristics. Each tier functions as a single redundancy group and stripes data across all disks in the tier. Each tier is RAID 6 protected and writes P and Q parity data across all the disks in the tier. For a RAID 6T2 array that has one tier of four SSD disks and another tier of 5 HDD disks, array data and parity information is written in the following pattern:

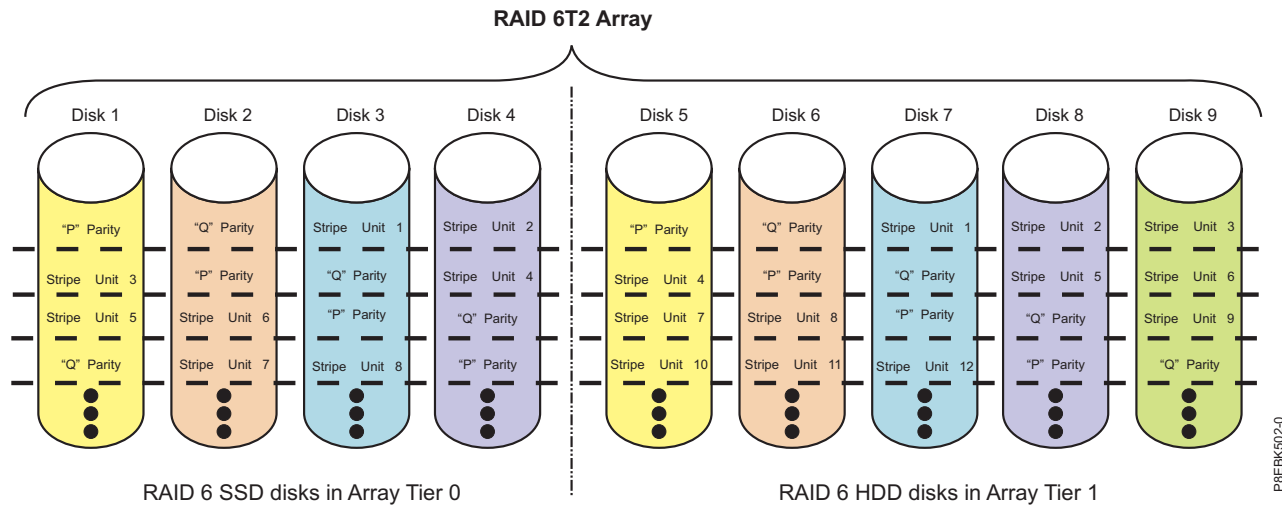


Figure 37. RAID 6T2

If one or two disks fail in either RAID 6 tier, you can continue to use the entire array normally. A RAID 6T2 array that is operating with one or two failed disks in either or both tiers 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 disks by using data and parity blocks on the operational disks. A tier in a RAID 6T2 array with a single failed disk has similar protection as that of a RAID 5 array with no disk failures. If a third disk fails in either tier, the entire array is placed in the failed state and is not be accessible.

RAID 10T2

Learn how data is written to a RAID 10T2 array when using the Easy Tier function.

RAID 10T2 is a RAID level that provides RAID 10 mirrored pair redundancy when using the Easy Tier function utilizing two different tiers of physical disk that have unique performance characteristics. Each tier must contain an even number of disks. A minimum of two disks are needed to create a RAID 10T2 tier. The data is striped across the mirrored pairs in each tier. For example, a RAID 10T2 array that has one tier of four SSD disks and another tier of 6 HDD disks would have data written on it in the following pattern:

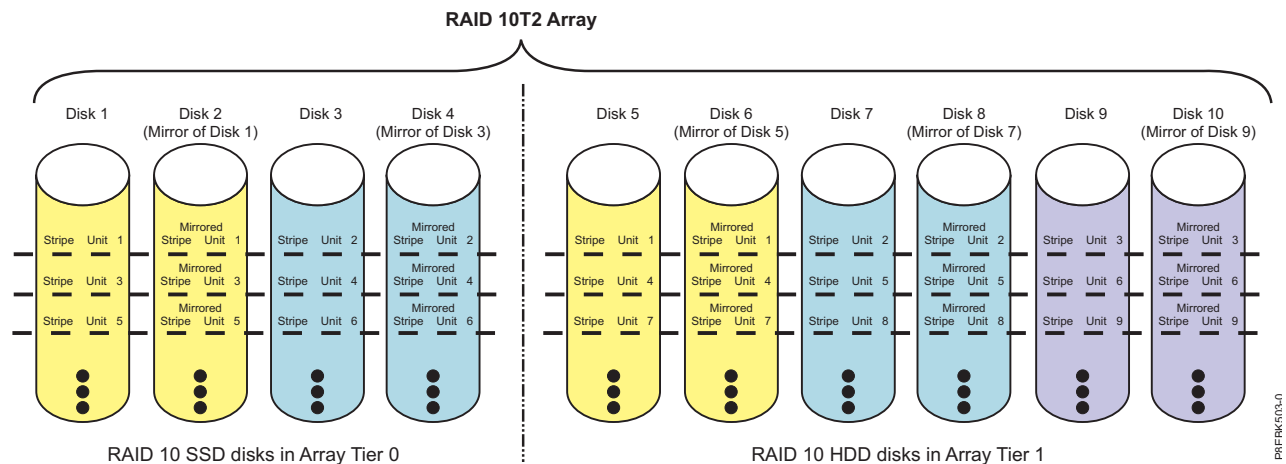


Figure 38. RAID 10T2

RAID 10T2 tolerates multiple disk failures. If one disk in each mirrored pair fails, the array continues to function, operating in degraded mode. You can continue to use the array 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 10T2 disk array is created, the controller automatically attempts to select the disks for each mirrored pair from a different controller connector (a different cable to a different device enclosure). For example, if four disks selected for a disk array are located on one of the controller connectors and another four disks selected are located on another of the controller connectors, the controller automatically attempts to create each mirrored pair from one disk on each controller connector. In the case of a controller port, cable, or enclosure failure, each mirrored pair continues to operate in a degraded mode. Such redundancy requires careful planning when you are determining where to place devices.

Estimating disk array capacities

The capacity of a disk array depends on the capacity of the advanced function disks used and the RAID level of the array.

In order to estimate the capacity of a disk array, you must know the capacity of the advanced function disks and the RAID level of the array.

1. For RAID 0, multiply the number of disks by the disk capacity.
2. For RAID 5, multiply one fewer than the number of disks by the disk capacity.
3. For RAID 6, multiply two fewer than the number of disks by the disk capacity.
4. For RAID 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.
 - SAS RAID controllers support up to 18 member disks in each RAID array.
5. For RAID 5T2, 6T2, and 10T2, each tier in the array follows the capacity rules for the base RAID level of the tier. Note that each tier must contain at least 10% of the total disk capacity. The disk capacity per tier is calculated by taking the smallest drive in each tier multiplied by the total number of physical disks in that tier. Divide the disk capacity of each tier by the total disk capacity. The result must be greater than 10%.

Related concepts:

“RAID level summary”

Compare RAID levels according to their capabilities.

RAID level summary

Compare RAID levels according to their capabilities.

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

Table 6. RAID level summary

RAID level	Data redundancy	Usable disk capacity	Read performance	Write performance	Min/Max devices per array on PCI-X and PCIe controllers	Min/Max devices per array on PCIe2 and PCIe3 controllers
RAID 0	None	100%	Very good	Excellent	1/18	1/32
RAID 5	Very good	67% to 94%	Very good	Good	3/18	3/32
RAID 6	Excellent	50% to 89%	Very good	Fair to good	4/18	4/32
RAID 10	Excellent	50%	Excellent	Very good	2/18 (even numbers only)	2/32 (even numbers only)

RAID 0

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

RAID 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 6

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

RAID 10

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

Note: A two-drive RAID level 10 array is equivalent to RAID level 1.

RAID 5T2, 6T2, and 10T2

Each tier in the array follows the capabilities of the base RAID level of the tier, except that the total maximum number of devices in both tiers combined cannot exceed the max number of devices for that base RAID level.

Related tasks:

"Estimating disk array capacities" on page 34

The capacity of a disk array depends on the capacity of the advanced function disks used and the RAID level of the array.

Stripe-unit size

With RAID technology, data is striped across an array of physical disks.

Striping data across an array of physical disks 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 disk array to 16 KB, 64 KB, 256 KB or 512 KB. PCIe2 and PCIe3 controllers only support a stripe size of 256 KB. You might be able to maximize the performance of your 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 or 512 KB. The recommended stripe size for most applications is 256 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.

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 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 topic and the `iprconfig` utility use common terminology for disk formats:

- **JBOD**

A Just a Bunch Of Disks (JBOD) disk is formatted to 512 or 4096 bytes per sector. A JBOD disk is assigned a `/dev/sdX` name and can be used by the Linux operating system.

- **Advanced function**

An *advanced function* disk is formatted to 528 or 4224 bytes per sector. This format allows disks to be used in disk arrays. An advanced function disk cannot be used by the Linux operating system directly. The Linux operating system can use an advanced function disk only if it is configured into a disk array.

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 disk arrays, see “Formatting to JBOD” on page 52.

An advanced function disk can be configured as:

Array Member

A 528 bytes per sector HDD disk that is configured as a member of an array.

Hot Spare

A 528 bytes per sector HDD disk that can be used by the controller to automatically replace a Failed disk in a Degraded RAID 5, 6, 10, 5T2, 6T2, or 10T2 disk array. A Hot Spare disk is useful only if its capacity is greater than or equal to the capacity of the smallest disk in an array that becomes Degraded. For more information about hot spare disks, see “Hot spare disks” on page 59.

Array Candidate

A 528 bytes per sector HDD disk that is a candidate for becoming an Array Member or a Hot Spare.

SSD Array Member

A 528 bytes per sector Solid-State disk that is configured as a member of an array.

SSD Hot Spare

A 528 bytes per sector Solid-State disk that can be used by the controller to automatically replace a Failed disk in a Degraded RAID 5, 6, 10, 5T2, 6T2, or 10T2 disk array. A Hot Spare disk is useful only if its capacity is greater than or equal to the capacity of the smallest disk in an array that becomes Degraded. For more information about hot spare disks, see “Hot spare disks” on page 59.

SSD Array Candidate

A 528 bytes per sector solid-state disk that is a candidate for becoming an Array Member or a Hot Spare.

4K Array Member

A 4224 bytes per sector HDD disk that is configured as a member of an array.

4K Hot Spare

A 4224 bytes per sector HDD disk that can be used by the controller to automatically replace a Failed disk in a Degraded RAID 5, 6, 10, 5T2, 6T2, or 10T2 disk array. A Hot Spare disk is useful only if its capacity is greater than or equal to the capacity of the smallest disk in an array that becomes Degraded. For more information about hot spare disks, see “Hot spare disks” on page 59.

4K Array Candidate

A 4224 bytes per sector HDD disk that is a candidate for becoming an Array Member or a Hot Spare.

4K SSD Array Member

A 4224 bytes per sector Solid-State disk that is configured as a member of an array.

4K SSD Hot Spare

A 4224 bytes per sector Solid-State disk that can be used by the controller to automatically replace a failed disk in a Degraded RAID 5, 6, 10, 5T2, 6T2, or 10T2 disk array. A Hot Spare disk is useful only if its capacity is greater than or equal to the capacity of the smallest disk in an array that becomes Degraded. For more information about hot spare disks, see “Hot spare disks” on page 59.

4K SSD Array Candidate

A 4224 bytes per sector Solid-State disk that is a candidate for becoming an Array Member or a Hot Spare.

The **Display Hardware Status** option in the `iprconfig` utility can be used to display these disks and their associated resource names. For details regarding how to view the disk information, see “Viewing device status” on page 46. 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              |
|-----|-----|-----|-----|-----|
|          0000:00:01.0/0:         PCI-X SAS RAID Adapter  Operational          |
|  sda     0000:00:01.0/0:4:2:0    Physical Disk           Active                |
|  sdb     0000:00:01.0/0:4:5:0    Physical Disk           Active                |
|          0000:00:01.0/0:4:10:0   Enclosure               Active                |
|          0000:00:01.0/0:6:10:0   Enclosure               Active                |
|          0000:00:01.0/0:8:0:0    Enclosure               Active                |
|          0002:00:01.0/1:         PCI-X SAS RAID Adapter  Operational          |
|  sdc     0002:00:01.0/1:0:1:0    Physical Disk           Active                |
|  sdd     0002:00:01.0/1:0:2:0    Physical Disk           Active                |
|          0002:00:01.0/1:0:4:0    Advanced Function Disk  Active                |
|          0002:00:01.0/1:0:5:0    Advanced Function Disk  Active                |
|          0002:00:01.0/1:0:6:0    Advanced Function Disk  Active                |
|          0002:00:01.0/1:0:7:0    Hot Spare               Active                |
|  sde     0002:00:01.0/1:255:0:0  RAID 0 Disk Array       Active                |
|          0002:00:01.0/1:0:0:0    RAID 0 Array Member     Active                |
|  sdf     0002:00:01.0/1:255:1:0  RAID 6 Disk Array       Active                |
|          0002:00:01.0/1:0:10:0   RAID 6 Array Member     Active                |
|-----|-----|-----|-----|-----|
```

```

|          0002:00:01.0/1:0:11:0        RAID 6 Array Member    Active    |
|          0002:00:01.0/1:0:8:0         RAID 6 Array Member    Active    |
|          0002:00:01.0/1:0:9:0         RAID 6 Array Member    Active    |
|          0002:00:01.0/1:0:24:0        Enclosure               Active    |
|          0002:00:01.0/1:2:24:0        Enclosure               Active    |
|
|
|
|
|
|
|e=Exit  q=Cancel  r=Refresh  t=Toggle  |
+-----+

```

Disk array, physical disk, and I/O adapter (IOA) states are displayed in the fifth column of the **Display Hardware Status** screen.

Disk array states

A disk array can be in one of seven states.

The seven valid states for disk arrays are described in the following table:

Table 7. Disk array states

State	Description
Active	The disk array is functional and fully protected (RAID level 5, 6, 10, 5T2, 6T2, and 10T2) 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 be fully protected against disk failures. The Degraded state indicates the protection against disk failures is less than optimal. 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. The Degraded state indicates the performance is less than optimal.
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.

Related concepts:

“I/O adapter states” on page 39

There are three possible states for the I/O adapter.

“Physical disk states”

There are six possible states for the physical disks.

Physical disk states

There are six possible states for the physical disks.

The six possible states for physical disks are described in the following table:

Table 8. Physical disk states

State	Description
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.

Related concepts:

“I/O adapter states”

There are three possible states for the I/O adapter.

“Disk array states” on page 38

A disk array can be in one of seven states.

I/O adapter states

There are three possible states for the I/O adapter.

The three possible states for I/O adapters are described in the following table:

Table 9. I/O adapter states

State	Description
Operational	The IOA is functional.
Not Operational	The device driver cannot successfully communicate with this IOA.
Not Ready	The IOA requires a microcode download.

Related concepts:

“Disk array states” on page 38

A disk array can be in one of seven states.

“Physical disk states”

There are six possible states for the physical disks.

Auxiliary write cache adapter

The Auxiliary Write Cache (AWC) adapter provides a duplicate, nonvolatile copy of write cache data of the RAID controller to which it is connected.

Protection of data is enhanced by having two battery-backed (nonvolatile) copies of write cache, each stored on separate adapters. If a failure occurs to the write cache portion of the RAID controller, or if the RAID controller itself fails in such a way that the write cache data is not recoverable, the AWC adapter

provides a backup copy of the write cache data to prevent data loss during the recovery of the failed RAID controller. The cache data is recovered to the new replacement RAID controller and then written out to disk before resuming normal operations.

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

It is important to understand the difference between multi-initiator connections and AWC connections. Connecting controllers in a multi-initiator environment refers to multiple RAID controllers connected to a common set of disk enclosures and disks. The AWC controller is not connected to the disks, and it does not perform device media accesses.

Should a failure of either the RAID controller or the Auxiliary Cache occur, it is extremely important that the maintenance analysis procedures (MAPs) for the errors in the Linux error log be followed precisely. Required service information can be found at “Problem determination and recovery” on page 103.

The RAID controller and the AWC adapter each require a PCI-X or PCIe bus connection. They are required to be in the same partition. The two adapters are connected by an internal SAS connection. For the Planar RAID Enablement and Planar Auxiliary Cache features, the dedicated SAS connection is integrated into the system planar.

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

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

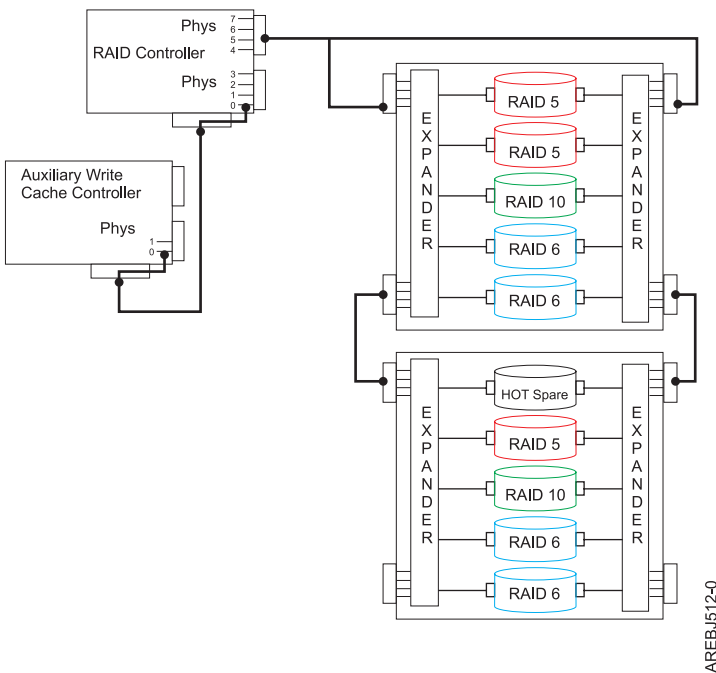


Figure 39. Example RAID and AWC controller configuration

Related information

Many other sources of information about the Linux operating system, RAID, and other associated topics are available.

The following publications contain related information:

- System unit documentation for information specific to your hardware configuration
- The IPR Linux device driver website
- *RS/6000 eServer pSeries Adapters, Devices, and Cable Information for Multiple Bus Systems*, order number SA38-0516 (<http://www-01.ibm.com/support/docview.wss?uid=pub1sa38051616>)
- The Linux Documentation Project website
- Linux for IBM eServer pSeries website
- *RS/6000 eServer pSeries Diagnostic Information for Multiple Bus Systems*, order number SA38-0509 (<http://www-01.ibm.com/support/docview.wss?uid=pub1sa38050923>)
- *The RAIDbook: A Handbook of Storage Systems Technology*, Edition 6, Editor: Paul Massiglia
- OpenPOWER website and IBM developerWorks website for information about Linux on PowerPC architecture

RAID controller software

A device driver and a set of utilities must be installed so that the controller can be identified and configured by Linux.

Note: References to the Linux operating system in this topic collection include the versions Linux 2.6, SUSE Linux Enterprise Server 10 and SUSE Linux Enterprise Server 11, Red Hat Enterprise Linux 4, Red Hat Enterprise Linux 5, and Red Hat Enterprise Linux 6. Make sure you are consulting the appropriate section of this set of information for the operating system you are using.

This document might 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.

For the controller to be identified and configured by Linux, the requisite device support software must be installed. Software for the controller consists of a device driver and a set of utilities.

The device driver is usually compiled as a kernel module named `ipr.ko`. The user utilities are usually packaged in a Red Hat Package Manager (RPM) called `iprutils`. The requisite software for the controller is often preinstalled as part of the normal Linux installation.

If the software package is not installed, software verification will fail. The missing packages 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 SourceForge.net.

The controller executes onboard microcode. The `iprconfig` utility in `iprutils` RPM can be used to update the microcode being used by the controller. For more information regarding the `iprconfig` utility, see “Updating the controller microcode” on page 82.

Verifying installation of the controller software

Verify that the `ipr` device driver for the controller is installed.

Consult the following table for the minimum `ipr` device driver version required for each supported adapter:

Table 10. Minimum ipr device driver support

Supported adapter	Minimum supported main line Linux version		Minimum supported Red Hat Enterprise Linux version		Minimum supported SUSE Enterprise Linux version	
	Device driver ipr version	Kernel version	Device driver ipr version	RHEL Version	Device driver ipr version	SLES Version
Custom card identification number (CCIN)						
572A No dual adapter	2.1.2	2.6.16	2.0.11.5 2.2.0.1	RHEL4 U6 RHEL5 U1	2.2.0.1	SLES10 SP1
572A With dual adapter	2.4.1	2.6.22	2.0.11.6 2.2.0.2	RHEL4 U7 RHEL5 U2	2.2.0.2	SLES10 SP2
572B	2.4.1	2.6.22	2.0.11.6 2.2.0.2	RHEL4 U7 RHEL5 U2	2.2.0.2	SLES10 SP2
572C	2.1.2	2.6.16	2.0.11.4 2.2.0	RHEL4 U5 RHEL5	2.2.0.1	SLES10 SP1
572F/575C	2.4.1	2.6.22	2.0.11.6 2.2.0.2	RHEL4 U7 RHEL5	2.2.0.2	SLES10 SP2
574E	2.4.1	2.6.22	2.0.11.6 2.2.0.2	RHEL4 U7 RHEL5 U2	2.2.0.2	SLES10 SP2
57B3	2.4.1	2.6.22	2.0.11.6 2.2.0.2	RHEL4 U7 RHEL5 U2	2.2.0.2	SLES10 SP2

Table 10. Minimum ipr device driver support (continued)

Supported adapter	Minimum supported main line Linux version		Minimum supported Red Hat Enterprise Linux version		Minimum supported SUSE Enterprise Linux version	
57B4	2.6.0	3.0	2.5.4	RHEL6 U4	2.6.0	SLES11 SP3
57B5	2.3.2	2.6.34	2.3.0.1	RHEL5 U8	2.5.2	SLES11 SP2
			2.5.0	RHEL6 U1		
57B7	2.3.0	2.6.20	2.0.11.5	RHEL4 U6	2.2.0.1	SLES10 SP1
			2.2.0.1	RHEL5 U1		
57B8	2.3.0	2.6.20	2.0.11.5	RHEL4 U6	2.2.0.1	SLES10 SP1
			2.2.0.1	RHEL5 U1		
57B9	2.1.2	2.6.16	2.0.11.5	RHEL4 U6	2.2.0.1	SLES10 SP1
			2.2.0.1	RHEL5 U1		
57BA	2.1.2	2.6.16	2.0.11.5	RHEL4 U6	2.2.0.1	SLES10 SP1
			2.2.0.1	RHEL5 U1		
57C3	2.5.2	3.2.0	2.3.0.1	RHEL5 U8	2.5.2	SLES11 SP2
			2.2.0.2	RHEL6 U3		
57C4	2.5.1	2.6.37	2.3.0.1	RHEL5 U8	2.5.2	SLES11 SP2
			2.5.0	RHEL65 U1		
57C7	2.4.1	2.6.22	2.0.11.6	RHEL4 U7	2.2.0.2	SLES10 SP2
			2.2.0.2	RHEL5 U2		
57CD	2.4.1	2.6.22	2.2.0.5	RHEL5 U6	2.2.0.3 /2.4.3	SLES10 SP3
			2.4.3	RHEL6		SLES11 SP1
57CE	2.6.0	3.0	2.5.4	RHEL6 U4	2.6.0	SLES11 SP3
57CF	2.4.1	2.6.22	2.0.11.6	RHEL4 U7	2.2.0.2	SLES10 SP2
			2.2.0.2	RHEL5 U2		
57D7	2.6.0	3.0	2.6.0	RHEL6 U5	2.6.0	SLES11 SP3
57D8	2.6.0	3.0	2.6.0	RHEL6 U5	2.6.0	SLES11 SP3

Consult the following table for the minimum version of the iprconfig utility required for each supported function:

Table 11. Versions of the iprconfig utility needed for SAS functions

SAS function	Version of the iprconfig utility
Base SAS support	2.0.15.8
Displaying SAS drive path status	2.2.5
High availability dual adapter support	2.2.8
Migrating RAID array between RAID level 0 and 10	2.2.10
Change asymmetric access mode	2.2.11
The latest version	2.4.1

You can verify several parameters of the device driver software.

1. To verify that the ipr device driver for the controller is installed, type `modinfo ipr | grep version`. Data is displayed in this form:

```
version:          2.6.0
```

```
srcversion:       D720FC6935DA2B2638B177A
```

2. To verify the ipr device driver version, type `# modinfo -F version ipr`. You will see data in this form:
2.6.0
3. To verify that the iprconfig utility is installed, type `iprconfig --version`. Output from this command will indicate if the utility is installed, and if so, version information will be displayed. You will see data in this form:
iprconfig: 2.4.0 (Feb 07, 2014)

Linux ipr device driver updates

The controller device driver for Linux, ipr, is part of the Linux kernel. Linux distributors periodically release updated device drivers for their Linux kernels.

Updated device drivers are usually packaged in an RPM. The updated device drivers provided by the Linux distributors are well-tested with other parts of the Linux distribution. Contact your Linux distributor for detailed information about the procedure for obtaining and updating the ipr device driver.

The Linux device drivers provided by your Linux distribution might be a few versions behind due to their test and release cycles. The most up-to-date ipr device driver can be found in the main line Linux kernel tree. Main line Linux kernel trees can be downloaded from Linux Kernel Archives.

The ipr device driver is located in the `drivers/scsi/` directory. Users can obtain the ipr device driver from a Linux kernel source tree and follow the instructions in `Documentation/kbuild/modules.txt` in the Linux kernel source tree to recompile the ipr device driver.

Attention:

- The ipr device driver depends on other parts of Linux kernel to work properly. Updating ipr device driver without updating other parts of Linux kernel might cause compiling errors or runtime errors and result in system crash and data loss. Only experienced users should update the ipr device driver from source code.
- A device driver is considered part of the Linux kernel. Updating a device driver may void a Linux service contract provided by your Linux distributor or service provider. Check your Linux distributor or service provider for their device driver update policy.

Updating the iprutils package

You can update the iprutils package.

You can update the iprutils package using the following methods:

- Downloaded the latest prebuilt iprutils rpms from <http://www14.software.ibm.com/support/customer/sas/f/lopdiags/home.html>.

By clicking on the **Learn more and download tools**, you can access the **IBM POWER® Linux Tools Repository** website. The tools repository supports the RHEL6, RHLE7, SLES11, and SLES12. Follow the instructions on the website to download the latest iprutils rpm packages.

- Download the user space package iprutils from Open Build Service. Click **Download Package** and select the required GNU or Linux distribution. Follow the instructions on the page to set up the OBS repository and download the package.
- Download the user space package source code from Open Build Service. You can build iprutils from the source using **tar.gz** provided by the OBS. Download the source code from OBS by clicking the

Download link for the latest version of iprutils in the **Actions** column. Unpack the source code and follow the build and install instructions provided in the README file.

You must install certain development headers and libraries before compiling iprutils:

Table 12. Required development headers and libraries

RHEL4, RHEL5, and RHEL6	SLES10 and SLES11
<ul style="list-style-type: none"> • glibc-headers • kernel-headers • ncurses-devel • pciutils-devel • libsysfs-devel 	<ul style="list-style-type: none"> • glibc-devel • ncurses-devel • pci-utils-devel • sysfsutils

When building iprutils in an RPM package, use the correct <RPM_DIR> for your system. If your system is running RHEL6 or RHEL7, use \$HOME/rpmbuild. If your system is running SLES11 or SLES12, use /usr/src/packages.

1. Log on to the GNU or Linux system.
2. Enter rpmbuild -rebuild -target=ppc iprutils-x.y.z-1.src.rpm where x.y.z is the iprutils version to build the iprutils package.
3. Enter cd <RPM_DIR>/RPMS/ppc.
4. Enter rpm -U iprutils-x.y.z-1.*ppc*.rpm to update iprutils.

Common IBM SAS RAID controller tasks

The instructions in this set of topics pertain to the various tasks that can be performed to manage disk arrays. The interface for working with the IBM SAS RAID controller is iprconfig.

When managing the IBM SAS RAID controller, you work with the iprconfig interface.

Starting the iprconfig utility

Use the iprconfig utility to work with the RAID controller.

When you start the iprconfig utility, you can select from various tasks.

To start the iprconfig utility, type the command iprconfig. A menu offers options for working with your 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. Work with adapter configuration |
| 8. Download microcode |
| 9. Analyze log |
| |
| |
|Selection: |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
|e=Exit |
+-----+

```

Status of devices, arrays, and paths

You can check on the status of your devices, arrays, and paths.

Use the methods described here to check the status of your devices, arrays, and paths.

Viewing device status

View the status of all devices controlled by the ipr device driver.

With this procedure you can view information about the disks and disk arrays on your system.

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
	0000:00:01.0/0:	PCI-X SAS RAID Adapter	Operational
sda	0000:00:01.0/0:4:2:0	Physical Disk	Active
sdb	0000:00:01.0/0:4:5:0	Physical Disk	Active
	0000:00:01.0/0:4:10:0	Enclosure	Active
	0000:00:01.0/0:6:10:0	Enclosure	Active
	0000:00:01.0/0:8:0:0	Enclosure	Active
	0002:00:01.0/1:	PCI-X SAS RAID Adapter	Operational
sdc	0002:00:01.0/1:0:1:0	Physical Disk	Active
sdd	0002:00:01.0/1:0:2:0	Physical Disk	Active
	0002:00:01.0/1:0:4:0	Advanced Function Disk	Active
	0002:00:01.0/1:0:5:0	Advanced Function Disk	Active
	0002:00:01.0/1:0:6:0	Advanced Function Disk	Active
	0002:00:01.0/1:0:7:0	Hot Spare	Active
sde	0002:00:01.0/1:255:0:0	RAID 0 Disk Array	Active
	0002:00:01.0/1:0:0:0	RAID 0 Array Member	Active
sdf	0002:00:01.0/1:255:1:0	RAID 6 Disk Array	Active
	0002:00:01.0/1:0:10:0	RAID 6 Array Member	Active
	0002:00:01.0/1:0:11:0	RAID 6 Array Member	Active
	0002:00:01.0/1:0:8:0	RAID 6 Array Member	Active
	0002:00:01.0/1:0:9:0	RAID 6 Array Member	Active
	0002:00:01.0/1:0:24:0	Enclosure	Active
	0002:00:01.0/1:2:24:0	Enclosure	Active

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

+-----+

The five columns of the Display Hardware Status screen have the following meanings:

- 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. Type t to 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, disk array and I/O adapter hardware statuses, see "Disk array overview" on page 35.
3. To view information on a specific device, select the desired device with a 1 and press Enter. If multiple pages of information are available, you can press f to page down (forward) or b to page up (backward).

When you view information on a specific device, the result will depend on which device you select. Information is presented similar to the following:

+-----+

```
|          Disk Unit Hardware Resource Information Details          |
|                                                                    |
|                                                                    |
|Manufacturer . . . . . : IBM                                       |
|Product ID . . . . . : ST373455SS                                  |
|Firmware Version . . . . . : 45303035 (E005)                       |
|Serial Number. . . . . : 3LQ0P24C                                  |
|Capacity . . . . . : 73.41 GB                                       |
|Resource Name. . . . . : /dev/sda                                    |
|                                                                    |
|Physical location                                                  |
|PCI Address. . . . . : 0000:00:01.0                                  |
|SCSI Host Number . . . . . : 0                                       |
|SCSI Channel . . . . . : 4                                           |
|SCSI Id. . . . . : 2                                                 |
|SCSI Lun . . . . . : 0                                               |
|                                                                    |
|Extended Details                                                  |
|FRU Number . . . . . : 10N7199                                       |
|EC Level . . . . . : D76038                                           |
|Part Number. . . . . : 10N7200                                       |
```

```

|Device Specific (Z0) . . . . . : 000005229F001002 |
|Device Specific (Z1) . . . . . : 0402E005 |
|Device Specific (Z2) . . . . . : 0021 |
|Device Specific (Z3) . . . . . : 07082 |
|Device Specific (Z4) . . . . . : |
|Device Specific (Z5) . . . . . : 22 |
| | More... |
|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 might be available. Press f to page down (forward) or b to page up (backward).

Viewing array status

You can view the status of the disk array by using the procedure presented here.

1. Run the `iprconfig` utility by typing `iprconfig`.
2. Select **Work with disk arrays**.
3. Select **Display disk array status**. You will see a screen similar to the following:

```

+-----+
|                               |
|           Display Disk Array |
|                               |
|Type option, press Enter.     |
| 1=Display hardware resource |
|                               |
|OPT Name  PCI/SCSI Location   |
|-----|-----|-----|-----|
|          0002:00:01.0/1:0:7:0 | Hot Spare | Active |
| sde      0002:00:01.0/1:255:0:0 | RAID 0 Disk Array | Active |
|          0002:00:01.0/1:0:0:0   | RAID 0 Array Member | Active |
| sdf      0002:00:01.0/1:255:1:0 | RAID 6 Disk Array | Active |
|          0002:00:01.0/1:0:10:0  | RAID 6 Array Member | Active |
|          0002:00:01.0/1:0:11:0  | RAID 6 Array Member | Active |
|          0002:00:01.0/1:0:8:0   | RAID 6 Array Member | Active |
|          0002:00:01.0/1:0:9:0   | RAID 6 Array Member | Active |
|                               |
+-----+

```

```

|
|
|
|
|
|
|
|
|
|
|
|
|
|
|
|
|
|e=Exit  q=Cancel  r=Refresh  t=Toggle
+-----+

```

Viewing path status

Use the `iprconfig` utility to view the SAS path status.

1. Run the `iprconfig` utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Display SAS path status**. The **Display SAS Path Status** screen that is displayed is similar to the following:

```

+-----+
|
|                               Display SAS Path Status
|
| Type option, press Enter.
|  1=Display SAS Path routing details
|
| OPT Name   PCI/SCSI Location      Description          Status
+-----+
|   sda     0000:c0:01.0/0:0:2:0    Physical Disk      Redundant Paths
|   sdb     0000:c0:01.0/0:0:3:0    Physical Disk      Redundant Paths
|   sdc     0000:c0:01.0/0:0:4:0    Physical Disk      Redundant Paths
|
|
|

```

```

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

```

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

```

+-----+
|
|          Display SAS Path Details
|
|Device . . . . . : /dev/sda
|Location . . . . . : 0000:c0:01.0/0:0:2:0
|
|
|SAS Port/SAS Address  Description      Active Status      Info
|-----|-----|-----|-----|
| 0/                   Physical Path      No   Healthy
| 0/5005076C03023F01  IOA port          Functional        3.0Gbps
| 0/5005076C0400C703  Expander port     Functional        3.0Gbps
| 2/5005076C0400C703  Expander port     Functional        3.0Gbps
| 0/5000CCA00376F4AC  Device port       Functional        3.0Gbps
| 2/5000CCA00336F4AC  Device LUN        Functional        Enabled
| 2/                   Physical Path      Yes  Healthy
| 2/5005076C03023F03  IOA port          Functional        3.0Gbps
| 0/5005076C0400C783  Expander port     Functional        3.0Gbps
| 2/5005076C0400C783  Expander port     Functional        3.0Gbps
| 1/5000CCA003B6F4AC  Device port       Functional        3.0Gbps
| 2/5000CCA00336F4AC  Device LUN        Functional        Enabled
|
|
|
|
|e=Exit  q=Cancel
|
+-----+

```

RAID and JBOD formats

Disks must be formatted before they can be used and recognized in a disk array.

Before a disk can be used in an IBM SAS 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. Procedures for performing these actions are contained in this set of information.

Formatting to advanced function

Before a disk can be used in a disk array or as a hot spare disk, it must be formatted for advanced function. Advanced function disks are physical disks that are formatted to a block size that is compatible with SAS RAID. The RAID block size is larger than a JBOD block size due to the SCSI T10 standardized data integrity fields along with logically bad block checking stored on each block with the data. The SAS RAID adapters support disk blocks based on 512 Bytes of data or 4K Bytes of data. The RAID block size for the 512 disks is 528 Bytes per sector and the RAID block size for the 4K disks is 4224 bytes per sector.

The steps in this procedure result in a disk formatted for advanced function.

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. To proceed with the format, type `c` to confirm. To return to the previous menu without formatting the disks, type `q`.

Notes:

- After the formatting is complete, the disks are ready for use in disk arrays. Resource names (such as `/dev/sdb`) might change when the system is rebooted. This could affect kernel command line entries and `fstab` entries. You might want to reboot the server now.
- This task requires rewriting data to the whole disk. The required time varies depending on the disk capacities. For large disks, it might take hours to complete.

Formatting to JBOD

Before a disk is recognized as a standalone disk, it must be formatted to JBOD.

The steps in this procedure result in a disk formatted for JBOD.

1. Run the `iprconfig` utility by typing `iprconfig`.
2. Select **Work with disk arrays**.
3. Select **Format device for JBOD function**.
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. To proceed with the format, type `c` to confirm. To return to the previous menu without formatting the disks, type `q`.

Notes:

- After the formatting is complete, the disks are ready for use as stand-alone disks. Resource names (such as `/dev/sdb`) might change when the system is rebooted. This could affect kernel command line entries and `fstab` entries. You might want to reboot the server now.

- This task requires rewriting data to the whole disk. The required time varies depending on the disk capacities. For large disks, it might take hours to complete.

Note: Resource names (such as /dev/sdb) might change as a result of reformatting. This could affect kernel command line entries and fstab entries. You might want to reboot the server now.

Creating and deleting disk arrays

Create and delete IBM SAS RAID disk arrays by following the procedures in this information set.

Follow the instructions provided in this section to create and delete IBM SAS RAID disk arrays.

Creating an IBM SAS 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 5, 6, 10, 5T2, 6T2, and 10T2) are 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.

Complete the following steps to create an IBM SAS RAID disk array:

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, and then press Enter. If the disks do not have a status of Zeroed, first format them using the instructions provided at Formatting to advanced function. Formatting the disks reduces the overall RAID creation time and ensures the RAID array is fully optimized following RAID creation. The information that you see will be similar to the following screen:

```
+-----+
|
|           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 . . . . . : 256 k                              |
|Queue Depth (default = 12). . . . . : 12                  |
|
|
|
```

```

|
|
|
|
|Press Enter to Continue
|
|e=Exit  q=Cancel
+-----+

```

6. Select the protection level that you want for the array. For more information about selecting an appropriate RAID level, see “Supported RAID levels” on page 27.
7. Select the stripe size that you want, in kilobytes, for the array. For more information about selecting an appropriate stripe size, see “Stripe-unit size” on page 35.
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 is displayed until the operation is completed. If you want to create additional disk arrays, you can exit the status screen by typing e.

When a disk array has been built, it is available to Linux like any SCSI disk. To find the resource name that has been assigned to the disk array, see the **Display Disk Array Status** screen.

Note: Resource names (such as /dev/sdb) might change as a result of creating a new disk array. This could affect kernel command line entries and fstab entries. You might want to reboot the server now.

Deleting an IBM SAS RAID disk array

Delete a disk array carefully to avoid losing data.

To preserve the data on the disk array, you must first back up all data that you want to save.

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

A disk array cannot be deleted if it is currently being rebuilt or synched. However, it is possible to delete a disk array that is currently in use. Make sure the disk array is not being used before deleting it.

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 want to delete, and 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.

Remember: Resource names (such as /dev/sdb) might change when the system is rebooted. This could affect kernel command line entries and fstab entries.

Adding disks to an existing array

Some controllers support adding disks to existing RAID level 5 or 6 disk arrays, which allows you to dynamically increase the capacity of a disk array while preserving its existing data.

Adding disks to an existing RAID 5 or 6 disk array can dynamically increase the capacity of a disk array while preserving existing data in the disk array. After you add disks to an existing disk array, the disks are protected and become part of the disk array, but will not contain parity and the data will not be restriped. Extensive use of this feature results in performance penalties because of the following reasons.

1. Not all the disks in the disk array contain parity and therefore the disks with parity are accessed more often for parity updates.
2. Data is not restriped and this reduces the ability to use the hardware assisted stripe write features.

Note: Some controllers might not support adding of disks to an existing array. See the feature comparison tables of PCIe3, PCIe2, and PCIe, and PCI-X cards to identify controllers that support this feature.

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 disk array to which you want to add the disks.
5. Select the eligible disks to be included into the disk array, and then press Enter.

Attention: All data on the selected drives will be lost when the disks are added to the existing array.

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.

6. To include the disks, press Enter. To exclude the disks from the disk array, type `q` to cancel.

Migrating an existing disk array to a new RAID level

The SAS RAID controller supports migrating an existing RAID 0 or 10 disk array to RAID 10 or 0, respectively. This allows you to dynamically change the level of protection of a disk array while preserving its existing data.

When migrating RAID 0 to RAID 10, additional disks must be included into the RAID 10 disk array in order to provide the additional level of protection. The number of additional disks will be equal to the number of disks in the original RAID 0 disk array. The capacity of the disk array will remain unchanged and the disk array remains accessible during the migration. The disk array is not protected by RAID 10 until the migration completes.

When migrating RAID 10 to RAID 0, no additional disks are included into the RAID 0 disk array. The number of disks in the resulting RAID 0 disk array will be reduced to half the number of disks in the original RAID 10 disk array. The capacity of the disk array will remain unchanged and the disk array remains accessible during the migration.

To migrate an existing array to a new level, do the following:

1. Run the `iprconfig` utility by typing `iprconfig`.
2. Select **Work with disk arrays**.
3. Select **Migrate disk array protection**.
4. Select the disk array to which you want to migrate the disk array protection level.

5. Press c and use up arrow and down arrow keys to change the protection level, and press Enter. A screen will display similar to the following:

```
+-----+
|                                     |
|                               Select Protection Level                               |
|                                     |
| Current RAID protection level is shown. To change                               |
| setting hit "c" for options menu. Highlight desired                             |
| option then hit Enter                                                         |
|                                     |
| c=Change Setting                                                             |
|                                     |
| /dev/sdf - Protection Level . . . . . : RAID 0                               |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
| Press Enter to Continue                                                       |
|                                     |
| e=Exit  q=Cancel                                                             |
+-----+
```

- Optional: If more member disks are required, for example when migrating from RAID 0 to 10, move the cursor by using the up arrow and down arrow keys at the list of candidate disks, and press 1 to select the desired disks. When you have completed your disk selection, press Enter. A screen will display similar to the following:

```
+-----+
|                                     |
|                   Select Disk Units for Migration                   |
|                                     |
| A minimum of 2 disks must be selected.                             |
| A maximum of 2 disks must be selected.                             |
| The number of disks selected must be a multiple of 2.             |
|                                     |
| OPT Name  PCI/SCSI Location      Vendor  Product ID    Status    |
|-----|-----|-----|-----|-----|
| 1  sg12   0003:00:01.0/2:0:2:0    IBM    HUS151473VLS300  Active    |
| 1  sg13   0003:00:01.0/2:0:3:0    IBM    HUS151473VLS300  Active    |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
| e=Exit   q=Cancel  t=Toggle                                             |
+-----+
```

7. Optional: If fewer member disks are needed, for example when migrating from RAID 10 to RAID 0, the adapter will select the extra disks and take them out of the disk array automatically. No user intervention is needed.
8. Press **Enter** to confirm the RAID migration.
The migration between different RAID levels takes anywhere from several seconds to several hours. A progress screen displays the migration progress. Migration progress can also be found in the **Display Disk Array Status** screen and **Display Hardware Status** screen.

Managing the volatile adapter write cache

The 8247-21L, 8247-22L, and 8247-42L systems support the volatile adapter write cache feature with the 57D7 integrated SAS controllers or the 57B4 PCIe3 RAID SAS adapter, when configured in a single controller configuration. After the required level of adapter firmware and Linux software operating system are installed, this feature is enabled by default, but the feature can be disabled manually for each of the individual RAID arrays. This feature enables a volatile write cache in the SAS RAID adapter into which data is written. The data is then flushed to the underlying disks. This feature allows the RAID adapter to combine multiple host write operations into fewer write operations to the disks. It also enables the RAID adapter to issue stripe write operations to RAID 5 or 6 arrays, when a full parity stripe of data is available in the cache.

Before you enable or disable the volatile adapter write cache, ensure that the following requirements are met:

- The `iprutils` package must be at least version 2.4.10. For instruction about obtaining this version of the `iprutils` package for your Linux distribution, if it is not already installed, see [Updating the iprutils package](#).
- The SAS adapter firmware level must be at least version 14511800. For instructions about obtaining and installing the microcode, see [Updating the controller microcode](#).

To display and change the write cache mode, complete the following steps:

1. To list the current hardware RAID arrays that are configured on the system, run the `iprconfig` utility:
`iprconfig -c show-arrays`

Information similar to the following example is displayed:

```
+-----+
| # iprconfig -c show-arrays
| Name      PCI/SCSI Location      Description      Status          |
| -----  -
| sda       0003:70:00.0/0:2:0:0           RAID 0 Array    Optimized       |
| sdb       0003:70:00.0/0:2:1:0           RAID 10 Array   Optimized       |
| sdc       0003:70:00.0/0:2:2:0           RAID 5 Array    Optimized       |
+-----+
```

2. To view the current caching mode for a specific adapter, type the following command: `iprconfig -c query-write-cache-policy sda`

Information similar to the following example is displayed:

```
+-----+
| # iprconfig -c query-write-cache-policy sda
| write back
+-----+
```

- To disable the volatile adapter write cache feature for a specific adapter, type the following command:
`iprconfig -c set-write-cache-policy sda writethrough`
 Information similar to the following example is displayed indicating the change in the caching policy:

```
+-----+
| # iprconfig -c set-write-cache-policy sda writethrough           |
| # iprconfig -c query-write-cache-policy sda                     |
| write through                                                    |
+-----+
```

Hot spare disks

Use hot spare disks to automatically replace failed disks in a RAID environment. .

Upon a disk failure in a RAID array is detected, the RAID adapter will look into all disks it controls for a hot spare disk. If one is available, the RAID adapter will replace the role of the failed disk in the RAID array with the hot spare disk. The data and parity are rebuilt from the survival member disks automatically. The RAID array is changed back to optimized state when the rebuild is complete.

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.

For disk arrays with Easy Tier function, it is important to note that a hot spare disk will only replace a disk in the tier that has the similar performance characteristics as the hot spare. So, you need different hot spare disks to fully cover all tiers in a tiered RAID level. For example, an SSD hot spare and a HDD hot spare.

Creating hot spare disks

Create a hot spare disk to automatically replace a failed disk in a RAID environment.

Format a disk for advanced function if you want to use it as a hot spare disk.

Complete the following steps to create hot spare disks:

- Run the `iprconfig` utility by typing `iprconfig`.
- Select **Work with disk arrays**.
- Select **Create a hot spare**.
- Select the adapter for which you want to create hot spares and press Enter.
- 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.
- 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.

Related tasks:

“Formatting to advanced function” on page 52

Before a disk can be used in a disk array or as a hot spare disk, it must be formatted for advanced function. Advanced function disks are physical disks that are formatted to a block size that is compatible with SAS RAID. The RAID block size is larger than a JBOD block size due to the SCSI T10 standardized data integrity fields along with logically bad block checking stored on each block with the data. The SAS RAID adapters support disk blocks based on 512 Bytes of data or 4K Bytes of data. The RAID block size for the 512 disks is 528 Bytes per sector and the RAID block size for the 4K disks is 4224 bytes per sector.

Deleting hot spare disks

Delete a hot spare disk to free it for use in other disk arrays.

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. To reconfigure the disks, press Enter. If you do not want to reconfigure the disks as hot spares, type `q` to cancel.

Drive queue depth

For performance reasons, you might want to change the queue depth of the disk command. The disk queue depth limits the maximum number of commands that the Linux software can issue concurrently to that disk at any time.

Increasing the disk queue depth might improve the disk performance by increasing disk throughput (or I/O) but might also increase latency (response delay). Decreasing the disk queue depth might improve disk response time but decrease the overall throughput. The queue depth is viewed and changed on each individual disk. When changing the disk queue depth, the command elements and data transfer window on the parent adapter might also need to be changed.

Viewing the drive queue depth

To view the current queue depth on any disk (JBOD or RAID), use the `iprconfig -c query-qdepth sda` command from the `iprconfig` utility command line. Alternatively, you can use `iprconfig` graphical interface (`iprconfig > Work with disk configuration`).

The `queue_depth` attribute contains the current setting. The default value for the disk queue depth is determined by the adapter family.

Table 13. Drive queue depth for different adapter families

Queue depth	PCI-X and PCIe adapter family	PCIe2 adapter family
Default JBOD disk queue depth	16	16
Default RAID disk queue depth	4 times the number of disks in the RAID array	16 times the number of disks in the RAID array

Example

To list the current `queue_depth` attribute value for the `sda` disk, type the following command:

```
iprconfig -c query-qdepth sda
```

The system displays a message similar to the following: 16.

Changing the drive queue depth

You can change the drive queue depth from the command line by running the following command:

```
iprconfig -c set-qdepth sda 64
```

Linux command-line interface

You can perform tasks to manage the SAS RAID controllers by using the `iprconfig` utility command line instead of using the `iprconfig` graphical interface. The following table summarizes the commonly used command-line interface commands. For the full list of command line options, see the `iprconfig` man page.

Table 14. *iprconfig* utility commands

Task	Command
General help	<code>man iprconfig</code>
Viewing the disk array status	<code>iprconfig -c show-arrays</code>
Format device for advanced function	<code>iprconfig -c format-for-raid [disk]....[disk]</code>
Format device for JBOD function	<code>iprconfig -c format-for-jbod [disk]...[disk]</code>
Creating a SAS disk array	<code>iprconfig -c raid-create -r [raid-level] -s [strip_size] [devices]</code>
Deleting a SAS disk array	<code>iprconfig -c raid-delete [RAID device]</code>
Adding disks to an existing disk array	<code>iprconfig -c raid-include [array] [disk]...[disk]</code>
Creating hot spare disks	<code>iprconfig -c hot-spare-create [disk]</code>
Deleting hot spare disks	<code>iprconfig -c hot-spare-delete [disk]</code>
Displaying rechargeable battery information	<code>iprconfig -c show-battery-info [IOA]</code>

Considerations for solid-state drives (SSD)

The purpose of this section is to cover important controller functions when using solid-state drives (SSD).

Hard disk drives (HDD) use a spinning magnetic platter to store data in magnetic fields in nonvolatile storage. SSDs are storage devices that use nonvolatile solid-state memory (usually flash memory) to emulate HDDs. HDDs have an inherent latency and access time caused by mechanical delays in the spinning of the platter and movement of the head. SSDs greatly reduce the latency and time to access the stored data. The nature of solid-state memory is such that read operations can be performed faster than write operations and that write cycles are limited. Using techniques such as wear leveling and overprovisioning, enterprise class SSDs are designed to withstand many years of continuous use.

SSD usage specifications

While using SSDs, consider the following specifications:

- Intermixing of SSDs and HDDs within the same disk array is not allowed. A disk array must contain all SSDs or all HDDs.
- It is important to properly plan for hot-spare devices when using arrays of SSDs. An SSD hot-spare device is used to replace a failed device in an SSD disk array and an HDD hot spare is used for an HDD disk array
- Although SSDs can be used in a RAID 0 disk array, it is preferred that SSDs to be protected by RAID levels 5, 6, 10, 5T2, 6T2, or 10T2.
- See Installing and configuring Solid-state drives to identify specific configuration and placement requirements related to the SSD devices.
- Some adapters, known as RAID and SSD adapters, contain SSDs, which are integrated on the adapter. See the PCIe SAS RAID card comparison table for features and additional information for your specific adapter type.
- SSDs are supported only when formatted to a RAID block size and used as part of a RAID array.

RAID 0 automatically created disk array on a PCIe or PCIe2 SAS RAID Controller

During the controller boot process, any 528 bytes per sector (not 4224 bytes per sector) SSD array candidate attached to a PCIe or PCIe2 SAS RAID Controller that is not already part of a disk array is automatically created as a single-drive RAID 0 disk array. Following are the two options for changing the RAID 0 disk array to a protected RAID level:

- The RAID 0 disk array can be migrated to a RAID 10 disk array by using the technique described in Migrating an existing disk array to a new RAID level.

- The automatically created RAID 0 disk array can be deleted (see Deleting a disk array) and a new SSD disk array can be created with a different level of RAID protection (see Creating a disk array).

Adapter cache control

Adapter caching improves overall performance with disk drives. In some configurations, adapter caching might not improve performance when you use SSD disk arrays. In these situations, adapter caching can be disabled by using the Change Configuration of Adapter window.

To disable adapter caching, complete the following steps:

1. Run the `iprconfig` utility by typing **iprconfig**.
2. Select **Work with adapter configuration**.
3. Select the desired adapter by typing 1, and press Enter.
4. Select **IOA Caching Mode**.
5. Type `c`.
6. Select **Disabled** and press Enter. A display similar to the following example appears in the window.

```
+-----+
|                                     |
|               Change configuration of Adapter               |
|                                     |
| Current Adapter configurations are shown. To change setting hit 'c' for options |
| menu. Highlight desired option then hit Enter.             |
| c=Change Setting                                           |
|                                     |
| Adapter:  0001:00:01.0/33  IBM      574E001SISI0A           |
|                                     |
| Preferred Dual Adapter State . . . . . : Primary          |
| Active/Active Mode . . . . . : Enabled                    |
| IOA Caching Mode . . . . . : Default                      |
|                                     |
|                                     |
| Or leave blank and press Enter to cancel                    |
| e=exit  q=cancel                                           |
|                                     |
+-----+
```

Multi-initiator and high availability

You can increase availability using multi-initiator and high availability to connect multiple controllers to a common set of disk expansion drawers.

The terms *multi-initiator* and *high availability* (HA) in the SAS arena refer to connecting multiple controllers (typically two controllers) to a common set of disk expansion drawers for the purpose of increasing availability. IBM SAS RAID controllers support high availability with up to two controllers in one or two systems.

IBM SAS RAID controllers support the following configurations:

HA two-system configuration

An HA two-system configuration provides a high-availability environment for system storage by enabling two systems or partitions to have access to the same set of disks and disk arrays. This feature is typically used with high availability cluster applications such as the IBM PowerHA® SystemMirror®. The IBM HACMP™ software provides a commercial computing environment that ensures that mission-critical applications can recover quickly from hardware and software failures.

The HA two-system configuration is intended for using disk arrays. The disks must be formatted to RAID format. Any RAID level, or combination of RAID levels, can be used.

Use of disks without RAID (referred to as JBOD) is also possible. The disks must be formatted to JBOD format. This JBOD alternative is supported only on particular controllers and requires unique setup and cabling.

HA single-system configuration

An HA single-system configuration provides for redundant controllers from a single system to the same set of disks and disk arrays. This feature is typically used with Multi-Path I/O (MPIO). MPIO support in Linux is provided by Device-Mapper Multipath and can be used to provide a redundant IBM SAS RAID controller configuration with RAID-protected disks.

When using an HA single-system configuration, the disks must be formatted to RAID format and used in one or more disk array. Any RAID level, or combination of RAID levels, may be used. Disks formatted to JBOD format are not supported in an HA single-system configuration.

Disks formatted to 512 bytes per sector are allowed in an HA single-system configuration. These disks are detected through both SAS controllers with two different paths and assigned with two different /dev/sdX names, for example, /dev/sde and /dev/sdf. The disks can be accessed using either path. However, there is no additional coordination between the two SAS controllers to automatically switch between the two paths. Multipath IO software such as Device-Mapper Multipath needs to be configured to monitor and switch paths. Using Device-Mapper Multipath to provide HA function is a software-only solution that is independent of IBM SAS RAID controllers. Its use is beyond the scope of this document.

Not all controllers support all configurations. Refer to the comparison tables and look for controllers that have the entry **Yes** for **HA two-system RAID**, **HA two-system JBOD**, or **HA single-system RAID**.

Possible HA configurations


Compare RAID and JBOD features as used with single- and two-system HA configurations.

Table 15. SAS RAID and JBOD HA configurations

Multi-initiator configuration	HA two system (for example PowerHA for AIX)	HA single system (for example MPIO)
RAID (disks formatted RAID block size per sector)	<ul style="list-style-type: none"> • Maximum of two controllers • Both controllers must have same write cache capability and write cache sizes • Both controllers must support “HA two-system RAID” • Controllers are in different systems or partitions 	<ul style="list-style-type: none"> • Maximum of two controllers • Both controllers must have same write cache capability and write cache sizes • Both controllers must support “HA single-system RAID” • Controllers are in the same system or partition
JBOD (disks formatted JBOD block size per sector)	<ul style="list-style-type: none"> • Maximum of two controllers • Both controllers must support HA two-system JBOD • Controllers are in different systems or partitions • Requires unique setup and cabling 	<ul style="list-style-type: none"> • Not supported

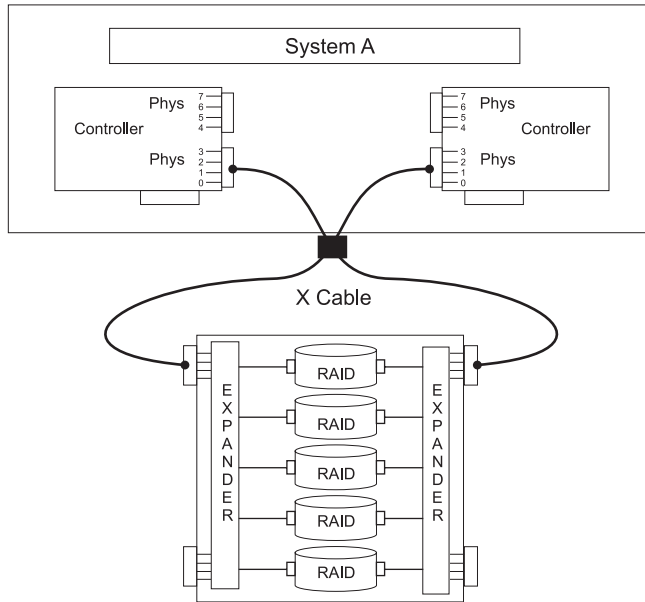
The following figures illustrate an example of each configuration.

HA Two System RAID



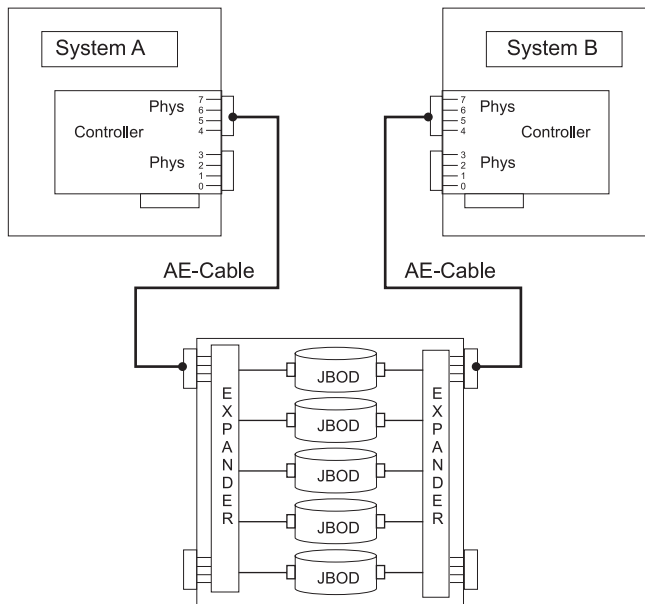
AREBJ514-0

HA Single System RAID



AREBJ515-0

HA Two System JBOD



AREBJ516-0

Controller functions

Consider these factors when using multi-initiator and HA functions.

Use of the multi-initiator and HA functions require controller and Linux software support. Controller support is shown in the “Feature comparison of SAS RAID cards” on page 2 table. Look for controllers

that have HA two-system RAID, HA two-system JBOD, or HA single-system RAID marked as Yes for the configuration that you want. The Linux software levels required for multi-initiator support are identified in Versions of iprconfig for SAS functions.

Specific controllers are intended only to be used in either an HA two-system RAID or HA single-system RAID configuration. Use the “Feature comparison of SAS RAID cards” on page 2 table to look for controllers that have Requires HA RAID configuration marked as Yes. This type of controller may not be used in an HA two-system JBOD or a stand-alone configuration.

Controllers connected in a RAID configuration must have the same write cache size (given they support write cache). A configuration error will be logged if the controllers' write caches are not the same size.

When reconfiguring a controller previously configured in a different HA configuration, it is recommended to configure the High-Availability Mode of the controllers to RAID or JBOD before attaching the SAS cables.

For all HA RAID configurations, one controller functions as the primary controller and manages the physical devices, such as creating a disk array or downloading disk microcode. The other controller functions as the secondary controller and is not capable of physical device management.

Note: On two-system configurations, the usage of the disk array may need to be discontinued from the secondary controller before some actions can be performed from the primary controller.

If the secondary controller detects the primary controller going offline, it will switch roles to become the primary controller. When the original primary controller comes back online, it will become the secondary controller. The exception to this case is if the original primary controller was previously designated as the “preferred” primary controller.

Both controllers are capable of performing direct I/O accesses to the disk arrays for purposes of read and write operations, but at any given time only one controller in the pair is “optimized” for the disk array. The controller that is optimized for a disk array is the one that directly accesses the physical devices for I/O operations. The controller that is non-optimized for a disk array will forward read and write requests through the SAS fabric to the optimized controller. See “HA asymmetric access optimization” on page 70 for more information on setting and viewing disk array optimization.

The primary controller logs most errors related to problems with a disk array. Some disk array errors may also be logged on the secondary if a disk array is optimized on the secondary, at the time the error occurred.

Typical reasons for the primary and secondary controllers to switch roles from what was expected or preferred are as follows:

- Controllers will switch roles for asymmetric reasons. For example, one controller detects more disk drives than the other. If the secondary controller is able to find devices that are not found by the primary controller, an automatic transition (failover) occurs. The controllers will communicate with each other, compare device information, and switch roles.
- Powering off the primary controller or the system that contains the primary controller causes an automatic transition (failover) to occur.
- Failure of primary controller or the system that contains the primary controller causes an automatic transition (failover) to occur.
- If the preferred primary controller is delayed in becoming active, the other controller assumes the role of primary controller. After the preferred primary controller becomes active, an automatic transition (failover) occurs.
- If the primary controller loses contact with the disks that are also accessible by the secondary controller, an automatic transition (failover) occurs.

- Downloading controller microcode might cause an automatic transition (failover) to occur. Such a transition is because the controller would reset itself to activate the new microcode. The controller will be temporarily offline until the reset is complete. Failover to another controller can prevent disk access disruption.

For all JBOD configurations, both controllers function only as stand-alone controllers and do not communicate directly with each other.

Users and their applications are responsible to ensure orderly read and write operations to the shared disks or disk arrays, for example, by using device reservation commands (persistent reservation is not supported).

Controller function attributes

Compare important attributes of controller functions.

Table 16. SAS controller functions.

Controller functions	HA two-system RAID configuration	HA two-system JBOD configuration	HA single-system RAID configuration
JBOD block size disks supported	No ¹	Yes	No ¹
RAID block size disks supported	Yes	No	Yes
Mirrored write cache between controllers (for controllers that have write cache)	Yes		Yes
Mirrored RAID parity footprints between controllers	Yes		Yes
Dual paths to disks	Yes	No	Yes
Target mode initiator device support	Yes	No	No
Only IBM qualified disk drives supported	Yes	Yes	Yes
Only IBM qualified disk expansion drawer supported	Yes	Yes	Yes
Tape or optical devices supported	No	No	No
Boot support	No	No	Yes
Operating mode ²	Primary Adapter or Secondary Adapter ³	StandaloneAdapter ³	Primary Adapter or Secondary Adapter ³
Preferred Dual Initiator Operating mode ²	None (no preference) or Primary ³	None (no preference) or Primary ³	None (no preference) or Primary ³
Dual Initiator Configuration ²	Default ³	JBOD HA Single Path ³	Default ³
High-Availability Mode	RAID ³	JBOD ³	RAID ³
Manage HA Access Characteristics ⁴	Yes	No	Yes

Table 16. SAS controller functions (continued).

Controller functions	HA two-system RAID configuration	HA two-system JBOD configuration	HA single-system RAID configuration
<ol style="list-style-type: none"> 1. JBOD block size (512 or 4096 bytes per sector) disks are not to be used functionally, but will be available to be formatted to RAID block size (528 or 4224 bytes per sector.) 2. Can be viewed by using the Change Adapter Configuration screen. 3. This option can be set by using the Change Adapter Configuration screen. 4. For information on managing the HA asymmetric access state for a disk array, see "HA asymmetric access optimization" on page 70. 			

Viewing HA controller attributes

You can view HA configuration-related information from the **Change Adapter Configuration** menu located under the `iprconfig` menu option titled **Work with adapter configuration**.

1. Run `iprconfig`.
2. Type 7 to select **Work with adapter configuration**.
3. Use the up/down arrow keys to move the cursor to the SAS controller and type 1; then, press Enter.
4. In the menu option titled **Change Configuration of Adapter**:
 - a. Configure **Preferred Dual Adapter State** to **None** or **Primary**.
 - b. Configure **High-Availability Mode** to **RAID** or **JBOD**.

The screen displayed will look similar to the following example.

```

+-----+
|                                         |
|                          Change Configuration of Adapter                          |
|                                         |
|                                         |
|Current Disk configurations are shown. To change setting hit 'c' for options |
|menu. Highlight desired option then hit Enter.                                |
|  c=Change Setting                                                              |
|                                         |
|Adapter: 0000:c8:01.0/0   IBM          572A001SISI0A                            |
|                                         |
|Preferred Dual Adapter State . . . . . : None                                  |
|High-Availability Mode . . . . . : Normal                                      |
|Active/Active Mode . . . . . : Disabled                                       |
|                                         |
|                                         |
|                                         |
|                                         |
|                                         |
|                                         |
|                                         |
|                                         |
|                                         |
|                                         |
|                                         |
+-----+

```

```
|
|
|
|
|
|
|
|Or leave blank and press Enter to cancel
|
|e=Exit  q=Cancel
+-----+
```

Note: For additional details on how to set up a configuration, see “Installing an HA single-system RAID configuration” on page 74, “Installing an HA two-system RAID configuration” on page 76 or “Installing an HA two-system JBOD configuration” on page 79.

HA cabling considerations

There are different types of cables to consider with high availability.

Correct cabling is one of the most important aspects of planning for a multi-initiator and HA configuration. For RAID configurations with disk expansion drawers, correct cabling is required to provide redundancy between each controller and the disk expansion drawer. For JBOD configurations, correct cabling is required, but generally provides much less redundancy between each controller and the disk expansion drawer. Thus there is better SAS fabric redundancy for RAID configurations versus JBOD configurations.

To see examples of how to cable HA configurations, see the section entitled Serial attached SCSI cable planning.

Note: Some systems have SAS RAID adapters integrated on the system boards. Separate SAS cables are not required to connect the two integrated SAS RAID adapters to each other.

HA performance considerations

Controller failures can impact performance.

The controller is designed to minimize performance impacts when running in an HA configuration. When using RAID 5, 6, 10, 5T2, 6T2, and 10T2 parity footprints are mirrored between the controller's nonvolatile memory, which causes only a slight impact to performance. For controllers with write cache, all cache data is mirrored between the controller's nonvolatile memories, which also causes only a slight impact to performance.

If one controller fails in an HA configuration, the remaining controller will disable write caching and begin to keep an additional copy of parity footprints on disk. This can significantly impact performance, particularly when using RAID 5, 6, 5T2, and 6T2.

Configuration and serviceability considerations for HA RAID configurations

There are configuration and serviceability differences between the primary and secondary controllers.

There are configuration and serviceability differences between the primary controller (that performs direct management of the physical devices) and the secondary controller (that runs as a client of the primary controller). This functional difference requires many of the configuration and serviceability functions to be performed on the primary controller because it is the only controller that can perform the commands.

Attention: Attempting these commands on a secondary controller might return unexpected results.

The following tasks are the common iprconfig tasks that must be performed from the primary controller:

- Under the iprconfig menu option titled **Work with disk arrays::**
 - Create a disk array
 - Delete a disk array

Note: On two-system configurations, the usage of the disk array may need to be discontinued from the secondary controller before some actions can be performed from the primary controller such as deleting the array.

- Add a device to a disk array
- Format device for RAID function
- Format device for JBOD function
- Work with hot spares
- Work with asymmetric access
- Force RAID Consistence Check
- Migrate disk array protection
- Under the iprconfig menu option titled **Work with Disk Unit Recovery::**
 - Concurrent add device
 - Concurrent remove device
 - Initialize and format disk
 - Rebuild disk unit data
 - Force RAID Consistency Check
- Under the iprconfig menu option titled **Work with disk configuration::**
 - Change Queue Depth
 - Change Format Timeout
- Under the iprconfig menu option titled **Download microcode::**
 - Adapter microcode download
 - Disk microcode download
 - SES device microcode download

HA asymmetric access optimization


HA access characteristics can balance the controller workload.

Important: Update each controller to the latest SAS controller microcode from the code download website. For instructions, see “Updating the controller microcode” on page 82. This update must be performed to obtain the latest critical fixes to ensure proper functioning.

By default, the primary adapter is optimized for all RAID arrays. The secondary adapter is non-optimized for all RAID arrays. The workload is not balanced between the two controllers. With either of the HA RAID configurations, maximum performance might be achieved by defining HA asymmetric access state for each disk array such that the workload is balanced between the two controllers. Setting the HA asymmetric access state for a disk array specifies which controller is preferred to be optimized for the disk array and perform direct reads and writes to the physical devices.

HA Access Optimization

The example assumes this setup:
 RAID Array "A" Optimized on Primary
 RAID Array "B" Optimized on Secondary



AREBU526-0

Figure 40. HA access optimization

Once the HA asymmetric access is enabled, HA asymmetric access states are displayed on the **Array Asymmetric Access** screen under the menu option titled **Work with Disk Arrays**, similar to the following screen; note that the status column shows the primary adapter view of the HA asymmetric access states of RAID arrays in HA single system RAID configuration:

```
+-----+
|                                     |
|               Array Asymmetric Access               |
|                                     |
|Select the disk array path.          |
|Type choice, press Enter.           |
| 1=change asymmetric access for a disk array         |
|OPT Name  PCI/SCSI Location          Description      Status  |
|-----|-----|-----|-----|
|  sdd     0002:00:01.0/1:255:0:0      RAID 5 Disk Array   Active  |
|  sde     0002:00:01.0/1:255:1:0      RAID 10 Disk Array  Optimized |
|  sdh     0002:00:01.0/1:255:2:0      RAID 0 Disk Array   Optimized |
|e=Exit  q=Cancel  r=Refresh  t=Toggle  |
+-----+

```

By using the up/down arrow keys to move the cursor to the desired RAID array, you can select it by pressing 1 followed by Enter. The HA asymmetric access state of the selected RAID array is displayed on the **Change Asymmetric Access Configuration of Array** screen, similar to the following:

```

+-----+
|           Change Asymmetric Access Configuration of Array           |
|                                                                       |
|Current array asymmetric access configuration is shown. To change setting hit |
| 'c' for options menu. Highlight desired option then hit Enter.      |
|  c=Change Setting                                                    |
|Array: /dev/sdh                                                       |
|Current asymmetric access state: Optimized                            |
|Saved asymmetric access state: Not Set                                |
|Preferred Asymmetric Access State. . . . : Not Set                    |
|Or leave blank and press Enter to cancel                              |
|e=Exit  q=Cancel                                                      |
+-----+

```

This display shows the HA asymmetric access state for the disk arrays selected. For each disk array listed, the current and preferred HA asymmetric access states are indicated. The *current* value shows how the disk array is currently accessed from the controller that was selected. The *preferred* value is the desired access state that is saved with the disk array configuration. Selecting the remote controller would show the opposite settings for the current and preferred access states.

There are three possible access state settings used for the HA asymmetric access state:

Optimized

The selected controller performs direct access for this disk array. This gives I/O operations performed on selected controller optimized performance compared to the remote controller. The selected disk array will be set to Non-Optimized on the remote adapter.

Non-Optimized

The selected controller performs indirect access for this disk array. This gives I/O operations performed on selected controller non-optimized performance compared to the remote controller. The selected disk array will be set to Optimized on the remote adapter.

Not Set

Neither an Optimized nor Non-Optimized access state has been set for this disk array. By default the disk array will be optimized on the primary controller.

The HA asymmetric access state can only be displayed and modified from the primary controller. The **Preferred Asymmetric Access** state can be modified when the disk array is selected from the primary controller. Changing the **Preferred Asymmetric Access** state from the primary controller stores the settings in the disk array and will automatically set the opposite settings on the secondary controller.

The controller will always try to switch the **Current Asymmetric Access** state of the disk array to match the **Preferred Asymmetric Access** state. This switching is done in the background by the controller; therefore, there may be delays between setting the **Preferred Asymmetric Access** state and seeing **Current Asymmetric Access** state switch. There are also situations where the controller will not switch HA asymmetric access state due to scenarios involving configuration errors, failed components, and certain RAID configuration activities.

By default all disk arrays are created with a **Preferred Asymmetric Access** state of Not Set. In order to maximize performance it is recommended that, when appropriate, multiple disk arrays be created and

optimized equally between the controller pair. This is accomplished by setting the **Preferred Asymmetric Access** to Optimized for half of the disk arrays and Non-Optimized to the other half.

Enabling asymmetric access

HA asymmetric access, an advanced function, is not enabled by default. Use the procedure here to enable it.

Perform the following steps on both the primary and the secondary controllers to enable HA asymmetric access:

1. Run the **iprconfig** command.
2. Type 7 to select **Work with adapter configuration**.
3. Use the up/down arrow keys to move the cursor to the SAS controller and type 1; then press Enter.
4. In the menu option titled **Change Configuration of Adapter**, configure **Active/Active Mode** to **Enabled**.

Asymmetric access status of disk arrays

The current asymmetric access status of your disk arrays can be easily accessed.

The current asymmetric access states of your disk arrays are displayed in the Status column of the **Display Hardware Status** screen (located under the main menu of iprconfig) and **Display Disk Array Status** screen (located under the menu option titled **Work with disk arrays**). An array can be in an Optimized state or an Active (non-optimized) state; see “HA asymmetric access optimization” on page 70.

IBM SAS RAID controllers are in a pair in HA single- and two-system RAID configurations. A RAID array that is Active (non-optimized) on the primary adapter is Optimized on the secondary adapter, and vice versa. RAID arrays are optimized on one of the two IBM SAS RAID controllers unless an error is detected.

The status of array members can be used to determine if the RAID array is on the primary or secondary adapter. When an array member has a status of Active, it indicates the RAID array is on the primary adapter. When the status is Remote, it indicates the RAID array is on the secondary adapter.

In the sample **Display Disk Array Status** screen of the HA single-system configuration shown in the following screen, the RAID 5 array is shown as `sdd` and is Active (non-optimized) on the primary adapter, while on the secondary adapter it is shown as `sdf` and Optimized. The RAID 10 array is shown as `sde` and Optimized on the primary adapter, and `sdg` and Active (non-optimized) on the secondary adapter.

```
+-----+
|                                     |
|               Display Disk Array Status               |
|                                     |
|Type option, press Enter.                |
| 1=Display hardware resource information details        |
|                                     |
|OPT Name  PCI/SCSI Location          Description        Status  |
|-----|-----|-----|-----|
|  sdd     0002:00:01.0/1:255:0:0      RAID 5 Disk Array  Active  |
|          0002:00:01.0/1:0:1:0        RAID 5 Array Member Active  |
|          0002:00:01.0/1:0:2:0        RAID 5 Array Member Active  |
|-----|-----|-----|-----|
```

	0002:00:01.0/1:0:3:0	RAID 5 Array Member	Active
sde	0002:00:01.0/1:255:1:0	RAID 10 Disk Array	Optimized
	0002:00:01.0/1:0:8:0	RAID 10 Array Member	Active
	0002:00:01.0/1:0:9:0	RAID 10 Array Member	Active
sdh	0002:00:01.0/1:255:2:0	RAID 0 Disk Array	Optimized
	0002:00:01.0/1:0:10:0	RAID 0 Array Member	Active
	0002:00:01.0/1:0:11:0	RAID 0 Array Member	Active
sdf	0003:00:01.0/2:255:0:0	RAID 5 Disk Array	Optimized
	0003:00:01.0/2:0:3:0	RAID 5 Array Member	Remote
	0003:00:01.0/2:0:1:0	RAID 5 Array Member	Remote
	0003:00:01.0/2:0:2:0	RAID 5 Array Member	Remote
sdg	0003:00:01.0/2:255:1:0	RAID 10 Disk Array	Active
	0003:00:01.0/2:0:9:0	RAID 10 Array Member	Remote
	0003:00:01.0/2:0:8:0	RAID 10 Array Member	Remote
sdi	0003:00:01.0/2:255:2:0	RAID 0 Disk Array	Active
	0003:00:01.0/2:0:10:0	RAID 0 Array Member	Remote
	0003:00:01.0/2:0:11:0	RAID 0 Array Member	Remote
e=Exit q=Cancel r=Refresh t=Toggle			

Installing high availability

Use the procedures in this section when performing HA installations.

Installation procedures are described for an HA two-system RAID configuration, an HA single-system RAID configuration, and an HA two-system JBOD configuration.

Installing an HA single-system RAID configuration

Use this procedure to help you to install an HA single-system RAID configuration.

To avoid problems during installation, follow the steps exactly as written.

Attention: Disk arrays can be created either before or after the HA RAID configuration is set up. See “Configuration and serviceability considerations for HA RAID configurations” on page 69 for important considerations.

1. Install and update the `iprutils` package on each system or partition. See “Updating the `iprutils` package” on page 44 for more information.
2. Power off the system or partition and install the SAS controllers into the system or partition. If you do not want to power off the machine or partition, Linux hot plug can be used to install the SAS controllers without powering off the system or partition.

Attention: Do not attach any cables to the SAS controllers at this time.

3. Update each controller to the latest SAS controller microcode from the code download Web site. See "Updating the controller microcode" on page 82.
4. To prevent errors while connecting the cables, configure the SAS controllers to the HA single-system RAID configuration before attaching any cables to the SAS controllers:
 - a. Run iprconfig
 - b. Type 7 to select Work with adapter configuration.
 - c. Use the up/down arrow keys to move the cursor to the SAS controller and type 1; then press Enter to select it.
 - d. In the menu option titled **Change Configuration of Adapter**, configure **Preferred Dual Adapter State** to None and **High-Availability Mode** to RAID.
5. Attach an X cable from the shared disk expansion drawer to the same SAS connector on each controller. To see examples of how to cable the HA configurations, see Serial attached SCSI cable planning. For examples of how to cable the HA configurations, see [Two SAS adapters to disk expansion drawer – RAID](#).
6. Power on the system or partition if it was powered off. Verify that the cabling and functioning of the controllers are correct by using the **Display Hardware Status** screen. Each SAS controller in dual-initiator configuration should show an optional block of Remote Adapter information for the other SAS controller.
 - a. Run iprconfig
 - b. Type 1 to select **Display Hardware Status**.
 - c. Use the up/down arrow keys to move the cursor to the SAS controller and type 1; then press Enter to select it. An optional block of Remote Adapter information for the other SAS controller in the remote system or partition is shown in the bottom of the **IOA Hardware Resource Information Details** screen, similar to the following:

```

+-----+
|                IOA Hardware Resource Information Details                |
|                                                                           |
|                                                                           |
|Manufacturer . . . . . : IBM                                             |
|Machine Type and Model . . . . . : 572B001SISIOA                       |
|Firmware Version . . . . . : 03200046                                   |
|Serial Number. . . . . : 07125793                                       |
|Part Number. . . . . : 0000042R4591                                     |
|Plant of Manufacturer. . . . . : 0022                                   |
|Cache Size . . . . . : 175 MB                                           |
|DRAM Size. . . . . : 0EE MB                                             |
|Resource Name. . . . . : /dev/sg20                                       |
|                                                                           |
|Physical location                                                       |
|PCI Address. . . . . : 0002:00:01.0                                       |
|SCSI Host Number . . . . . : 1                                           |
|                                                                           |

```

```

|Current Dual Adapter State . . . . . : Primary          |
|Preferred Dual Adapter State . . . . . : No Preference   |
|Remote Adapter Manufacturer. . . . . : IBM              |
|Remote Adapter Machine Type And Model. . : 572B001SISIOA  |
|Remote Adapter Serial Number . . . . . : 07125687       |
|                                                       |
|Current Asymmetric Access State. . . . . : Disabled     |
|                                                       |
|                                                       |
|Press Enter to Continue                               |
|                                                       |
|e=Exit  q=Cancel                                     |
+-----+

```

7. Optional: Configure one of the controllers in the HA single-system RAID configuration to be the Preferred Primary controller by following the steps below:

Note: Such configuration is often done for performance and usability reasons, such as disk configuration changes. If neither controller is configured as the Preferred Primary controller, the controllers will default to primary or secondary through a negotiation process during boot.

- a. Run `iprconfig`
- b. Type `7` to select **Work with adapter configuration**.
- c. Use the up/down arrow keys to move the cursor to the SAS controller and type `1`; then press Enter to select it.
- d. In the menu option titled **Change Configuration of Adapter**, configure **Preferred Dual Adapter State** to Primary.

When determining the Preferred Primary controller, consider the following:

- Because all disk array access must go through the primary controller, performance will be better for disk I/O operations from the system or partition containing the primary controller.
- All disk array configuration changes must be done on the system or partition containing the primary controller.
- Most disk service, including error log analysis, is performed from the system or partition containing the primary controller. However, errors might also be presented by the secondary controller that might require actions on the system or partition containing the secondary controller.

Installing an HA two-system RAID configuration

Use this procedure to help you to install an HA two-system RAID configuration.

To avoid problems during installation, follow the steps exactly as written.

Attention: Disk arrays can be created either before or after the HA RAID configuration is set up. See “Configuration and serviceability considerations for HA RAID configurations” on page 69 and “Functions requiring special attention in an HA two-system RAID configuration” on page 78 for important considerations.

1. Install and update the `iprutils` package on each system or partition. See “Updating the `iprutils` package” on page 44 for more information.

2. Power off the system or partition and install the SAS controllers into the system or partition. If you do not want to power off the machine or partition, Linux hot plug can be used to install the SAS controllers without powering off the system or partition.
Attention: Do not attach any cables to the SAS controllers at this time.
3. Update each controller to the latest SAS controller microcode from the code download website. See "Updating the controller microcode" on page 82.
4. To prevent errors while connecting the cables, configure the SAS controllers to the HA two-system RAID configuration before attaching any cables to the SAS controllers:
 - a. Run the iprconfig command.
 - b. Type 7 to select Work with adapter configuration.
 - c. Use the up/down arrow keys to move the cursor to the SAS controller and type 1; then press Enter to select it.
 - d. In the menu option titled **Change Configuration of Adapter**, configure **Preferred Dual Adapter State** to None and **High-Availability Mode** to RAID.
5. Attach an X cable from the shared disk expansion drawer to the same SAS connector on each controller. To see examples of how to cable the HA configurations, see Serial attached SCSI cable planning. For examples of how to cable the HA configurations, see [Two SAS adapters to disk expansion drawer – RAID](#).
6. Power on the system or partition if it was powered off. Verify that the cabling and functioning of the controllers are correct by using the **Display Hardware Status** screen. Each SAS controller in dual-initiator configuration should show an optional block of Remote Adapter information for the other SAS controller.
 - a. Run the iprconfig command.
 - b. Type 1 to select **Display Hardware Status**.
 - c. Use the up/down arrow keys to move the cursor to the SAS controller and type 1; then press Enter to select it. An optional block of Remote Adapter information for the other SAS controller in the remote system or partition is shown in the bottom of the **IOA Hardware Resource Information Details** screen, similar to the following:

```

+-----+
|                IOA Hardware Resource Information Details                |
|                                                                           |
|                                                                           |
|Manufacturer . . . . . : IBM                                             |
|Machine Type and Model . . . . . : 572B001SISIOA                         |
|Firmware Version . . . . . : 03200046                                    |
|Serial Number. . . . . : 07125687                                        |
|Part Number. . . . . : 0000042R4591                                       |
|Plant of Manufacturer. . . . . : 0022                                     |
|Cache Size . . . . . : 175 MB                                           |
|DRAM Size. . . . . : 0EE MB                                             |
|Resource Name. . . . . : /dev/sg26                                       |
|                                                                           |
|Physical location                                                         |

```

```

|PCI Address. . . . . : 0002:00:01.0 |
|SCSI Host Number . . . . . : 2 |
|
|Current Dual Adapter State . . . . . : Primary |
|Preferred Dual Adapter State . . . . . : No Preference |
|Remote Adapter Manufacturer. . . . . : IBM |
|Remote Adapter Machine Type And Model. . : 572B001SISIOA |
|Remote Adapter Serial Number . . . . . : 07125793 |
|
|Current Asymmetric Access State. . . . . : Disabled |
|
|
|Press Enter to Continue |
|
|e=Exit q=Cancel |
+-----+

```

7. Optional: Configure one of the controllers in the HA two-system RAID configuration to be the Preferred Primary controller by following the steps below:

Note: Such configuration is often done for performance and usability reasons, such as disk configuration changes. If neither controller is configured as the Preferred Primary controller, the controllers will default to primary or secondary through a negotiation process during boot.

- a. Run the `iprconfig` command.
- b. Type 7 to select **Work with adapter configuration**.
- c. Use the up/down arrow keys to move the cursor to the SAS controller and type 1; then press Enter to select it.
- d. In the menu option titled **Change Configuration of Adapter**, configure **Preferred Dual Adapter State** to Primary.

When determining the Preferred Primary controller, consider the following:

- Because all disk array access must go through the primary controller, performance will be better for disk I/O operations from the system or partition containing the primary controller.
- All disk array configuration changes must be done on the system or partition containing the primary controller.
- Most disk service, including error log analysis, is performed from the system or partition containing the primary controller. However, errors might also be presented by the secondary controller that might require actions on the system or partition containing the secondary controller.

Functions requiring special attention in an HA two-system RAID configuration

Manual intervention might be required on the system or partition containing secondary controller to get visibility to the new configuration.

Many configuration and serviceability functions need to be performed on the system or partition containing the primary controller. Any functions performed on the system or partition containing the

primary controller might also require some manual intervention on the system or partition containing secondary controller to get visibility to the new configuration.

The following table lists some of the common functions and the required steps to perform on the secondary controller:

Table 17. Configuration steps for the secondary controller.

Function performed on primary controller	Required configuration on secondary controller
Format device for advanced function	No configuration steps needed
Format device for JBOD function	No configuration steps needed
Create Disk Array	No configuration steps needed
Delete Disk Array ¹	Issue the SCSI Stop Unit command to the array <code>sg_start --stop /dev/sdX</code>
Add disks to Disk Array	No configuration steps needed
Reconstruct Disk Array	No configuration steps needed
Create/Delete Hot Spare disk	No configuration steps needed
Add disk (Hot Plug Manager)	No configuration steps needed
Remove disk (Hot Plug Manager)	No configuration steps needed
Reclaim Controller Cache Storage	No configuration steps needed
¹ The <code>sg_start</code> command is provided by the <code>sg3_utils</code> package.	

Installing an HA two-system JBOD configuration

Use this procedure to help you to install an HA configuration for two systems that both have JBOD-formatted disks.

To avoid problems during installation, follow the steps exactly as written.

Attention: Consider using an HA two-system RAID configuration rather than an HA two-system JBOD configuration due to the increased redundancy, performance, and reliability provided by the RAID subsystem.

Attention: Both controllers must have the High-Availability Mode option set to JBOD before being connected to the disk drives, all of the disk drives must be formatted to JBOD (JBOD block size per sector) format, and the correct cabling must be used. Remove any devices from the disk expansion drawer that are not set to JBOD or reformat them to JBOD block size per sector.

1. Install and update the `iprutils` package on each system or partition. See “Updating the `iprutils` package” on page 44 for more information.
2. Power off the system or partition and install the SAS controllers into the system or partition. If you do not want to power off the machine or partition, Linux hot plug can be used to install the SAS controllers without powering off the system or partition.

Attention: Do not attach any cables to the SAS controllers at this time.
3. Update each controller to the latest SAS controller microcode from the code download Web site. See “Updating the controller microcode” on page 82.
4. To prevent errors while connecting the cables, configure the SAS controllers to the HA two-system RAID configuration before attaching any cables to the SAS controllers:
 - a. Run the `iprconfig` command.
 - b. Type 7 to select Work with adapter configuration.
 - c. Use the up/down arrow keys to move the cursor to the SAS controller and type 1; then press Enter to select it.

- d. In the menu option titled **Change Configuration of Adapter**, configure **Preferred Dual Adapter State** to None and **High-Availability Mode** to JBOD.
5. Attach an AE cable from the shared disk expansion drawer to the same SAS connector on each controller. To see examples of how to cable the HA configurations, see Serial attached SCSI cable planning.
6. Power on the system or partition if it was powered off. Verify that the cabling and functioning of the controllers are correct by using the **Display Hardware Status** screen. Each SAS controller in dual-initiator configuration should show an optional block of Remote Adapter information for the other SAS controller.
 - a. Run the **iprconfig** command.
 - b. Type 1 to select **Display Hardware Status**.
 - c. Use the up/down arrow keys to move the cursor to the SAS controller and type 1; then press Enter to select it. An optional block of Remote Adapter information for the other SAS controller in the remote system or partition is shown in the bottom of the **IOA Hardware Resource Information Details** screen, similar to the following:

```

+-----+
|                IOA Hardware Resource Information Details                |
|                                                                           |
|                                                                           |
|Manufacturer . . . . . : IBM                                             |
|Machine Type and Model . . . . . : 572B001SISIOA                       |
|Firmware Version . . . . . : 03200046                                   |
|Serial Number. . . . . : 07125793                                       |
|Part Number. . . . . : 0000042R4591                                       |
|Plant of Manufacturer. . . . . : 0022                                       |
|Cache Size . . . . . : 175 MB                                           |
|DRAM Size. . . . . : 0EE MB                                             |
|Resource Name. . . . . : /dev/sg20                                       |
|                                                                           |
|Physical location                                                       |
|PCI Address. . . . . : 0002:00:01.0                                       |
|SCSI Host Number . . . . . : 1                                           |
|                                                                           |
|Current Dual Adapter State . . . . . : Primary                           |
|Preferred Dual Adapter State . . . . . : No Preference                   |
|Remote Adapter Manufacturer. . . . . : IBM                               |
|Remote Adapter Machine Type And Model. . : 572B001SISIOA               |
|Remote Adapter Serial Number . . . . . : 07125687                       |
|                                                                           |

```

```

|Current Asymmetric Access State. . . . . : Disabled          |
|
|
|Press Enter to Continue                                     |
|
|e=Exit   q=Cancel                                          |
+-----+

```

IBM SAS RAID controller maintenance

Maintenance procedures include microcode updates, failure recovery procedures, and working with the battery pack.

Follow the recommended maintenance procedures in this section to help maintain your RAID controller for Linux.

Usage tips

Follow recommended usage tips to help avoid controller and disk array problems.

The following usage tips will help you maintain your controller.

- 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 hot plug, PCI hot plug 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 a 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 "Concurrent device remove" screen to remove and replace a physical disk. For instructions on how to remove and replace a disk, see "Removing a failed disk" on page 95.
- 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 "Problem determination and recovery" on page 103.
- 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 might be a common cause.

Updating the controller microcode

Disk drives attached to this RAID adapter must contain the latest microcode level. Microcode can be updated using the iprconfig utility.

Download the latest microcode level for your drive model and controller from Fix Central. If you download new microcode from the website, follow the instructions in the readme file or Desc link on the website to install it. Use the installation procedure below only if there are no instructions available.

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.

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

1. Install the package by typing `rpm -ivh pci.101402BD.20-01200041-1.Linux.noarch.rpm`.
2. Update the microcode by typing `iprconfig`.
3. Select **Work with microcode updates** and press Enter.
4. Select **Download microcode** and press Enter.
5. Select the device or adapter you wish to update by typing a 1, and then press Enter. Multiple devices can be selected.
6. Select the microcode image to be downloaded to the device or adapter by typing a 1 and then press Enter.
7. Press Enter again to confirm the download.
8. If multiple devices were selected, repeat from Step 6 for the next device.

Rechargeable battery maintenance

Proper battery maintenance procedures can help to prevent data loss.

Perform battery maintenance tasks only when directed to do so by an isolation procedure or a maintenance analysis procedure (MAP).

Displaying rechargeable battery information

You can view information about the rechargeable battery by using the iprconfig utility.

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.
5. Press Enter to display battery information. The Battery Information screen displayed will look similar to the following:

```
+-----+
|                                     |
|                               Battery Information                               |
|                                     |
|                                     |
|                                     |
|Resource Name. . . . . : /dev/sg15                                         |
```


- 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.

Perform the following steps to force the battery error:

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. This forces the battery error.
Attention: Perform this procedure only under the guidance of IBM hardware service support. This procedure poses serious risks to performance until the battery pack is replaced.
5. 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 battery pack

You can replace battery packs that are maintainable both concurrently and nonconcurrently.

The battery pack replacement procedure you follow depends on whether your battery pack is concurrently maintainable or nonconcurrently maintainable.

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

CAUTION:

The battery is a lithium ion battery. To avoid possible explosion, do not burn. Exchange only 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: Data loss can occur; if the cache battery pack is not already in the error state, follow the steps described in “Forcing a rechargeable battery error” on page 83 before proceeding. After the battery is in the error state, replace it with a new battery by completing the appropriate replacement procedure.

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 a controller, 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.

Replacing a nonconcurrently maintainable battery pack

Use this procedure to replace a battery pack that is nonconcurrently maintainable or to replace a battery pack that is not concurrently maintainable or is concurrently maintainable but not accessible due to the structure of the system enclosure.

Determine whether it is safe to replace the cache battery pack by referring to “Displaying rechargeable battery information” on page 82. It is safe to replace the cache battery pack when the Battery Information screen displays Yes at the entry “Battery pack can be safely replaced.”

1. Remove the controller from the system. See your system documentation for instructions.
2. Place the controller on an ESD protected surface.




AREBJ518-1

Figure 41. Removing the cache battery

- A** Cache battery lever
- B** Cache battery pack
- C** Cache present LED

3. Move the cache battery lever **A** away from the connector to disengage the battery from the connector.
4. Continue to slide the cache battery pack out of the mounting guides and remove it from the controller.

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.



AREBJ519-0

Figure 42. Replacing the cache battery

- A** Cache battery lever
- B** Cache battery pack
- C** Cache battery connector

5. Before installing the new cache battery pack, move the lever to the unlatched position (away from the connector).
6. Slide the new cache battery pack into the mounting guides on the controller until it is seated in the battery connector.
7. Once the battery is seated in the connector, move the lever to the latched position to fully seat the battery into the connector.
8. Install the controller into the system. See your system documentation for instructions.


Replacing a concurrently maintainable battery pack

Use this procedure to replace a battery pack that is concurrently maintainable.

Note: If your system enclosure does not allow concurrent access to the battery, refer to “Replacing a nonconcurrently maintainable battery pack” on page 85 to replace the battery pack.

Determine whether it is safe to replace the cache battery pack by referring to “Displaying rechargeable battery information” on page 82. It is safe to replace the cache battery pack when the Battery Information screen displays Yes at the entry “Battery pack can be safely replaced.”

1. Remove the controller from the system. See your system documentation for instructions.
2. Place the controller on an ESD protected surface.




AREBJ518-1

Figure 43. Removing the cache battery

- A** Cache battery lever
- B** Cache battery pack
- C** Cache present LED

3. Move the cache battery lever **A** away from the connector to disengage the battery from the connector.
4. Continue to slide the cache battery pack out of the mounting guides and remove it from the controller.

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.



AREBJ519-0

Figure 44. Replacing the cache battery

- A** Cache battery lever
- B** Cache battery pack
- C** Cache battery connector

5. Before installing the new cache battery pack, move the lever to the unlatched position (away from the connector).
6. Slide the new cache battery pack into the mounting guides on the controller until it is seated in the battery connector.
7. Once the battery is seated in the connector, move the lever to the latched position to fully seat the battery into the connector.
8. Restart the adapter's write cache by completing the following steps:
 - a. Run the `iprconfig` utility by typing `iprconfig`.
 - b. Select **Work with disk unit recovery**.
 - c. Select **Work with resources containing cache battery packs**.
 - d. Start the IOA cache on the adapter whose battery you just replaced by typing `3`.
 - e. Press Enter.

Replacing the cache directory card

Replace the cache directory card only if you are directed to do so by an isolation procedure or a maintenance analysis procedure (MAP).


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

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 a controller, 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.

To replace the cache directory card, perform the following steps:

1. Remove the controller using the remove procedures for the model or expansion unit on which you are working.
2. Locate the cache directory card **A**.



3. Unseat the connector on the cache directory card by wiggling the two corners over the connector using a rocking motion. Then, lift the cache directory card off the connector and out of the guides on the cache battery rail.
4. Install the replacement card **A** by inserting it into the guides on the cache battery rail and then seating it on the connector.



AREBJ523-0

5. Install the controller using the installation procedures for the model or expansion unit on which you are working.

Replacing an SSD module on the PCIe RAID and SSD SAS adapter

Use this procedure to replace an integrated solid-state drive (SSD) on a PCIe serial-attached SCSI (SAS) RAID and SSD Adapter.

Complete the following steps to perform a nonconcurrent SSD replacement on a PCIe SAS RAID and SSD adapter:


Note: When an SSD on the PCIe adapter fails, the entire adapter must be removed from the system prior to replacing the individual SSD. See the documentation for your system for removing a PCI RAID and SSD SAS adapter from the system.

1. Remove the adapter from the system. See PCI adapters.

Important: Ensure to follow the concurrent or nonconcurrent replacement procedures depending on the type of data protection enabled:

- If the data protection is RAID, use the nonconcurrent procedure.
 - If the data protection is mirrored (card to card) and the SSD is located in the 5802 or 5803 expansion unit, use the concurrent procedure.
 - If the data protection is mirrored (card to card) but the SSD is not located in the 5802 or 5803 expansion unit, use the nonconcurrent procedure.
2. Place the adapter on a surface that is electrostatic-discharge protected.
 3. Lift the lever (A) for the SSD that you are replacing to a fully vertical position.


Note: Each of the levers (A) undocks two SSDs at a time.



P7EBJ511-0

Figure 45. Lifting the levers


4. With the lever (**A**) in the vertical position, firmly push the lever (**A**) away from the adapter tailstock to undock the two SSDs from their connectors.



P7EBJ512-0

Figure 46. Pushing the lever away from the adapter tailstock


5. Lift the single device retaining latch (**B**) for only the SSD that you are replacing by first moving it away from the center of the SSD divider and then lifting it to a fully vertical position.



P7EBJ513-0

Figure 47. Lifting the device retaining latch

- Using the device access finger openings (C), push the single SSD that you are replacing out of the device holder.



P7EBJ514-0

Figure 48. Pushing the SSD that is being replaced

- Grasp the SSD and continue to remove it from the adapter.
- Install the replacement SSD by performing steps 2 on page 91 to 7.

- Note:** Ensure that the device retaining latch and undocking lever are in the fully closed position.
- Reinstall the adapter in the system. See PCI adapters.

10. If you replaced the SSD as a part of another procedure, return to that procedure.

Replacing a disk in a SAS RAID adapter

The data from the failing disk can be reconstructed to a replacement hot-spare disk if the hot-spare disk is available during the failure. If the hot-spare is active and is available during the failure, the state of the affected disk array is either **Rebuilding** or **Optimal** because of the use of a hot-spare disk.

To complete a disk replacement on a SAS RAID adapter, complete the following steps:

1. If you want the new disk to be designated as a hot-spare disk, you must first prepare the disk to be used in the array. Complete the following steps to prepare the disk for use in a disk array. Otherwise, go to step 2.

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

- a. Run the `iprconfig` utility by typing **iprconfig**.
 - b. Select **Work with disk arrays > Format device for RAID function**.
 - c. Select the adapter to be formatted and press Enter.
 - d. Select the disks that must be prepared for use in disk arrays and formatted to 528 Byte Sectors from the list.
 - e. Return to the main menu.
 - f. Select **Work with disk arrays > Create a hot spare**.
 - g. Select the adapter for which you want to create the hot-spare disks and press Enter.
 - h. Select the disks that you want to designate as hot spares and press Enter.
 - i. Press Enter to configure the disks as hot-spare disks. If you do not want to configure the disks as hot-spare disks, type `q` to cancel.
 - j. Return to the system service procedure that sent you here.
2. If the state of the array is **Failed** or **Missing**, go to step 3. If the state of the array is **Degraded**, complete the following steps to change the state of the array to **Optimal**:
 - a. Run the `iprconfig` utility by typing **iprconfig**.
 - b. Select **Work with disk unit recovery > Rebuild disk unit data**.
 - c. Select the failed disks to rebuild.
 - d. Return to the system service procedure that sent you here.
 3. If the state of the array is **Failed** or **Missing**, delete and recreate the array, and then restore the data from the backup disk by completing the following steps:

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

- a. To delete a disk array, run the `iprconfig` utility by typing **iprconfig**.
- b. Select **Work with disk arrays > Delete a disk array**.
- c. From the listed disk arrays, select the array you want to delete, and then press Enter.
- d. To create a disk array, run the `iprconfig` utility by typing **iprconfig**.
- e. Select **Work with disk arrays > Create a disk array**.
- f. Select the controller under which you want to create a disk array, and then press Enter.
- g. Select the disk units to be included in the disk array, and then press Enter.

Information similar to the following example is displayed:

```
+-----+
|                                             |
|                                             |
|               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 . . . . . : 256 k
|Queue Depth (default = 12). . . . . : 12
|
|
|
|
|
|
|
|
|
|Press Enter to Continue
|
|e=Exit q=Cancel
+-----+

```

- h. Select the protection level for the array. For more information about selecting the appropriate RAID level, see “Supported RAID levels” on page 27.
- i. Select the stripe size, in kilobytes, for the array. For more information about selecting the appropriate stripe size, see “Stripe-unit size” on page 35.
- j. Press Enter to continue.
Data can be restored from a backup disk. The disk array can be added to a volume group. Logical volumes and file systems can also be created.
- k. Return to the system service procedure that sent you here.

Physical disks

Replace failed disks as soon as possible, even if a rebuild was initiated on a hot spare by the controller.

Use the **Concurrent remove device** option in the iprconfig utility to replace physical disks attached to the controller.

Removing a failed disk

Use the iprconfig utility to remove a failed disk.

1. Run the iprconfig utility by typing iprconfig.
2. Select **Work with disk unit recovery**.
3. Select **Concurrent Device Remove**. The resulting screen is similar to the following:

```

+-----+
|                                     |
|               Concurrent Device Remove               |
|                                     |
| Choose a single location for remove operations        |
| 1=Select                                             |
|                                     |
|OPT Name  Platform Location      Description      Status      |
|-----|-----|-----|-----|
|  sg1    U5888.001.6BAH025-P1-D1  Advanced Function Disk  Active      |
|  sg0    U5888.001.6BAH025-P1-D2  Advanced Function Disk  Active      |
|  sg2    U5888.001.6BAH025-P1-D5  Advanced Function Disk  Active      |
|                                     U5888.001.6BAH025-P1-D6      Empty      |
|                                     U5888.001.6BAH025-P1-D7      Empty      |
|                                     U5888.001.6BAH025-P1-F8      Empty      |
|                                     |
|                                     |
|e=Exit  q=Cancel  t=Toggle      |
+-----+

```

4. Select the device to be removed by typing 1, and press Enter.
5. Verify that the selected device is the device to be removed. The Identify indicator should now be set. Press Enter.

Attention: iprconfig allows 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.
7. Press Enter to complete removing the failed disk.

Installing a new disk

Use the iprconfig utility to install a new disk after the failed disk has been removed.

1. Run the iprconfig utility by typing iprconfig.
2. Select **Concurrent Device Add**. The resulting screen is similar to the following:

```

+-----+
|                                     |
|               Concurrent Device Add               |
|                                     |
| Choose a single location for add operations        |
| 1=Select                                             |
|                                     |
|OPT Name  Platform Location      Description      Status      |
|-----|-----|-----|-----|

```

```

|-----|
|      U5888.001.6BAH025-P1-D6      Empty      |
|      U5888.001.6BAH025-P1-D7      Empty      |
|      U5888.001.6BAH025-P1-F8      Empty      |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|                                     |
|e=Exit  q=Cancel  t=Toggle
+-----+

```

3. Select the location to install the device by typing 1; then, press Enter. If you do not see the location to install, press t to toggle the dual path locations.
4. Verify the location of the device to install. The identify indicator must 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.

Related tasks:

“Rebuilding a newly installed disk” on page 98

If a hot spare rebuild operation is not initiated by the controller when you have installed a new disk, you must initiate a rebuild operation before the disk is recognized by the system.

Rebuilding a newly installed disk:

If a hot spare rebuild operation is not initiated by the controller when you have installed a new disk, you must initiate a rebuild operation before the disk is recognized by the system.

1. Run the `iprconfig` utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select the disk you want to rebuild; then, press Enter.

Attention: Data currently on the disk will be overwritten.

4. To rebuild data, press Enter. If you do not want to rebuild data, type `q` to cancel.

Disk failure recovery

The IBM SAS RAID controller handles disk failures differently, depending on the RAID level involved.

The recovery procedure you use will depend on the level of the RAID adapter.

For a description of disk array and physical disk states, see “Disk array overview” on page 35. For a description of RAID levels, see “Supported RAID levels” on page 27.

RAID 0 failure

RAID 0 does not provide data protection. A single disk failure causes a RAID 0 array to transition to the Failed state.

If a single disk failure causes a RAID 0 array to transition to the Failed state, you must 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 5 disk recovery

RAID 5 protects data with parity information distributed on all member disks.

The data can be recovered from the remaining disks if one disk fails.

Recovering a RAID 5 single-disk failure:

RAID 5 protects the data with parity information distributed on all member disks. The data can be recovered from the remaining disks if one disk fails.

When a single disk in a RAID 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. If a hot spare disk is available, the controller can rebuild the data on the disk automatically. If a hot spare disk is not available, you must replace the failed disk and then initiate a rebuild.

Complete the following steps to initiate a rebuild:

1. Run the `iprconfig` utility by typing `iprconfig`.
2. Select **Work with disk unit recovery**.
3. Select **Rebuild disk unit data**.
4. Select the disks you want to rebuild, then press Enter.
5. 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.

RAID 5 multiple-disk failure:

If a second disk in a RAID 5 disk array fails, the array also fails and its data is not accessible.

If a second disk in a RAID level 5 disk array fails, you must 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 6 disk recovery

RAID 6 protects data with dual parity information distributed on all member disks.

The data is recoverable from the remaining disks if one or two disks fail. If a third disk in a RAID 6 disk array fails, the data is not recoverable.

Recovering a RAID 6 single-disk or dual-disk failure:

Whenever one or two disks in a RAID 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 is available, the controller can rebuild the data on the disk automatically. If a hot spare disk is not available, you will need to replace the failed disk and then initiate a rebuild.

Perform the following steps to initiate a rebuild:

1. Run the iprconfig utility by typing iprconfig.
2. Select **Work with disk unit recovery**.
3. Select **Rebuild disk unit data**.
4. Select the disks you want to rebuild, then press Enter.
5. 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.

RAID 6 failure of three or more disks:

If a third disk in a RAID 6 disk array fails, the array is marked as Failed and its data is not accessible.

If a third disk in a RAID 6 disk array fails, you must 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 10 disk recovery

RAID 10 protects the data on mirrored disk pairs.

Data can be recovered unless both disks of a mirrored disk pair fails.

Recovering a RAID 10 single-disk failure:

When a single disk in a RAID 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.

When ever a disk fails, replace it as soon as possible. If a hot spare disk is available, the controller can rebuild the data on the disk automatically. If a hot spare disk is not available, you will need to replace the failed disk and then initiate a rebuild.

Complete the following steps to initiate a rebuild:

1. Run the iprconfig utility by typing iprconfig.
2. Select **Work with disk unit recovery**.

3. Select **Rebuild disk unit data**.
4. Select the disks you want to rebuild, then press Enter.
5. 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.

RAID 10 multiple-disk failure:

The procedure for handling a RAID 10 multiple-disk failure depends on whether the failed disk is on the same mirrored pair or on a different mirrored pair.

If each failed disk belongs to a different mirrored pair, then the array is 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 that for a single-disk failure case where you 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 must 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.

Viewing SAS fabric path information

Use the iprconfig utility to view details of the SAS fabric information.

1. Type iprconfig and press Enter.
2. Select **Work with disk unit recovery**.
3. Select **Display SAS path status**.
4. Select the device by typing 1 in the OPT field. The resulting screen is similar to the following example:

```
+-----+
|                                     |
|                               Display SAS Path Details                       |
|                                     |
| Device.....: /dev/sg3              |
| Location.....: 0/00-04-05          |
| SAS Address      Description   Active  Status      Info      |
|-----|-----|-----|-----|-----|
| 00000005         Resource Path  Yes    Healthy      |
| 0000FF/500507604B809D41 IOA Port           Functional  6.0Gbps    |
| 00001E/500507604AC00DB0 Expander Port           Functional  6.0Gbps    |
| 000005/500507604AC00DB0 Expander Port           Functional  6.0Gbps    |
| 000001/500051610004B542 Device port           Functional  6.0Gbps    |
| 000001/500051610004B540 Device LUN           Functional  Enabled     |
| 00000405         Resource Path  Yes    Healthy      |
| 0004FF/500507604B809D45 IOA Port           Functional  6.0Gbps    |
| 000422/500507604AC00DB4 Expander Port           Functional  6.0Gbps    |
| 000405/500507604AC00DB4 Expander Port           Functional  6.0Gbps    |
```

```

|000400/500051610004B541 Device port          Functional      6.0Gbps      |
|000405/500051610004B540 Device LUN          Functional      Enabled       |
|e=Exit    q=Cancel                            |
+-----+

```

Table 18. Possible values for the SAS path status.

Status	Description
Functional	No problem detected
Degraded	The SAS node is degraded
Failed	The SAS node is failed
Suspect¹	The SAS node is suspected of contributing to a failure
Missing¹	The SAS node is no longer detected by the controller
Not valid	The SAS node is incorrectly connected
Unknown	Unknown or unexpected status

¹This status is an indication of a possible problem; however, the controller is not always able to determine the status of a node. The node can have this status even when the status of the node is not displayed.

Viewing SCSI location and physical resource path

Use the `iprconfig` utility to view the Serial Computer System Interface (SCSI) location and physical resource path of a device.

1. Type `iprconfig` and press Enter.
2. Select **Display hardware status**.
3. Select the device by typing 1 in the OPT field. The resulting screen is similar to the following example:

```

+-----+
|          Disk Unit Hardware Resource Information Details          |
|                                                                    |
|Manufacturer.....: IBM                                           |
|Product ID.....: V2-TX21B10400G                                  |
|Firmware Version.....: 30473330 (0G30)                            |
|Serial Number.....: 50401GDG                                       |
|Capacity.....: 387.96 GB                                           |
|                                                                    |
|Physical location                                                |
|PCI Address.....: 0000:01:00.0                                       |
|Resource Path.....: 00-04-05                                         |
|SCSI Host Number .....: 0                                           |
|SCSI Channel.....: 0                                               |
|SCSI Id.....: 3                                                    |

```

```
|SCSI Lun.....: 0 |
|Platform Location.....: U5888.001.6BAH025-P1-C1-P1-D6 |
|Capacity.....: 387.96 GB |
| | |
|Exteded Details |
|Press Enter to Continue |
|e=Exit  q=Cancel  f=PageDn  b=PageUp |
+-----+
```

Reclaiming IOA cache storage

A reclamation procedure for the I/O adapter (IOA) cache storage is sometimes used to resolve cache problems. Perform this procedure only if directed to do so by a maintenance analysis procedure (MAP).

Attention: Data may be lost. When an auxiliary cache adapter connected to the RAID controller logs a 9055 URC in the error log, the Reclaim procedure does not result in lost sectors. Otherwise, the Reclaim procedure will result in lost sectors.

Complete the following steps to reclaim IOA cache storage:

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. A screen display will then show the devices that might be affected.
5. To reclaim IOA cache storage, type c to confirm. If you do not want to reclaim IOA cache storage, type q to cancel. A screen display similar to one of the following will then be displayed.

```
+-----+
|
|
|
|
|
|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
|
+-----+

```

Problem determination and recovery

Several tools are available that the trained service representative can use to identify and recover from problems.

Trained service representatives use error logs, Unit Reference Code tables, and maintenance analysis procedures (MAPs) in problem determination and recovery procedures.

Note: The procedures contained in this section 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 topic might require involvement of the system administrator. For additional sources of related service procedures required to service this subsystem, see “Related information” on page 23.

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 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 MAP to further determine what actions should be taken to resolve the problem.

This topic collection provides many of these MAPs. The MAPs are intended to address only problems directly related to disk arrays and SAS cabling problem isolation. MAPs related to other device or adapter problems, when applicable, are located in other Linux or system documentation.

Analyzing error logs

The experienced user can search through this file with commands, such as `grep` and `sed`, but the `iprconfig` utility also offers a convenient menu for analyzing the error log.

Errors logged by the adapters and devices are stored in the `/var/log/messages` file. Options on the error log screen that are not mentioned here are used to 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.

1. Select **Analyze Log** from the main menu in `iprconfig`. A screen will display information 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                       |
+-----+
```

2. To change the directory where the error log is stored, use the **Set root kernel message log** directory option. Normally, this does not need to be used. It can be useful when booting from an alternate root file system if you want to access `/var/log/messages` on another file system.
3. To change the text editor used to view the error log from the default text editor, `vi`, choose the **Set default editor** option. Invoking `iprconfig` with the `-e` command line option will also change the default editor. Type the command `iprconfig -h` for more information.
4. 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.

Basic vi commands

There are several frequently-used **vi** commands.

Table 19. Frequently-used vi commands

Command	Action
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 <i>n</i> . If <i>n</i> is omitted, go to the last line in file.
/pattern	Search forward for pattern. After searching, type <i>n</i> to repeat the search.
?pattern	Search backward for pattern. After searching, type <i>n</i> to repeat the search.
n	Repeat the previous search
:q	Quit vi and return to the iprconfig utility

Searching logs

Use commands similar to these to perform a search on an error log.

The steps in this procedure are presented as a sample search for the URC "3400." Replace the variable "3400" with your own search term.

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.

Samples: Error logs

When viewed through iprconfig, ipr error messages are displayed in various formats.

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.

Sample: Generic IOA or device error log

Sample generic IOA or device error log

The generic IOA or device error log provides information in the following format:

```
+-----+
|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 |

```

+-----+

Sample: Device configuration error log

Sample device configuration error log

The device configuration error log provides information in the following format:

+-----+

```

|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 |
| WWN: 5000CCA00336F4AC |
|-----New Device Information----- |
|Vendor/Product ID: |
| Serial Number: 00000000 |
| WWN: 0000000000000000 |

```

```

|Cache Directory Card Information: |
|Vendor/Product ID: |
|  Serial Number: 00000000 |
|          WWN: 0000000000000000 |
|Adapter Card Information: |
|Vendor/Product ID: |
|  Serial Number: 00000000 |
|          WWN: 0000000000000000 |
|Additional IOA Data: 00000000 00000000 00000000 00000000 00000000 |

```

```

+-----+

```

Sample: Array error log

Sample array error log

The array error log provides information in the following format:

```

+-----+

```

```

|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 |

```

```

|           WWN: 5000CCA00336F4AC |

```

```

| Current Location: 2:0:4:0 |

```

```

| Expected Location: 2:0:4:0 |

```

```

|-----|

```

```

| Array Member 1: |

```

```

| Vendor/Product ID: IBM      IC35L036UCDY10-0 |

```

```

|   Serial Number: E3V0J55B |

```

```

|           WWN: 5005076C0400C703 |

```

```

| Current Location: 2:0:5:0 |

```

```

| Expected Location: 2:0:5:0 |

```

```

+-----+

```

Sample: Cache error log

Sample cache error log

The cache error log provides information in the following format:

```
+-----+
|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|
|           WWN: 5005076C03023F00|
|Adapter Card Information:|
|Vendor/Product ID: IBM      5703001|
|  Serial Number: 03060038|
|           WWN: 5005076C03023F00|
|-----Expected Configuration-----|
|Cache Directory Card Information:|
|Vendor/Product ID:|
|  Serial Number:|
|           WWN: 0000000000000000|
|Adapter Card Information:|
|Vendor/Product ID: IBM      5703001|
|  Serial Number: 03060038|
|           WWN: 5005076C03023F00|
|Additional IOA Data: 00000000 00000000 152101D0|
+-----+
```

Disk array problem identification

A disk array problem is uniquely identified by a Unit Reference Code (URC).

A Unit Reference Code (URC) is used to indicate the specific problem that has occurred. This code is required in order to determine which MAP to use.

A URC is provided in the ipr error log with each error. This URC describes the exact problem that has been detected and is 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 about this status screen, see “Viewing device status” on page 46.

When you have obtained the URC, you can then determine which Maintenance Analysis Procedure (MAP) to use.

Unit reference code tables

Use the information in this section, along with the unit reference code (URC) obtained from the ipr error log, to determine which maintenance analysis procedures (MAPs) to use.

Table 20. URC

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 131.	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 131.	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 131.	Disk drive
3020	Storage subsystem configuration error IOA detected a SAS fabric configuration error	Perform "MAP 3350" on page 129 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
3029	A device replacement has occurred	No action required.	
3100	Device bus error	A SAS fabric error occurred. Perform "MAP 3350" on page 129.	
3109	IOA timed out a device command	Perform "MAP 3350" on page 129.	
310D	Logical block reference tag error detected by the IOA Logical block guard error detected by the IOA	Controller T10 DIF host bus error. Perform "MAP 3350" on page 129.	
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 signal cables
4010	Incorrect connection between cascaded expanders	Perform "MAP 3342" on page 126 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
4020	Connections exceed IOA design limits	Perform "MAP 3343" on page 127 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	

Table 20. URC (continued)

URC	Description text	Service action	Failing items
4030	Incorrect multipath connection	Perform "MAP 3344" on page 127 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
4040	Incomplete multipath connection between IOA and enclosure	Perform "MAP 3344" on page 127 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
4041	Incomplete multipath connection between enclosure and device	Perform "MAP 3346" on page 128 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
4050	Enclosure does not support a required multipath function	Perform "MAP 3348" on page 128 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
4060	Multipath redundancy level got worse	Perform "MAP 3353" on page 133 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
4061	Multipath redundancy level got better	No action required.	
4080	IOA exceeded maximum operating temperature	Perform "MAP 3495" on page 150 for a PCIe2 or a PCIe3 controller.	
4085	Service required	Perform "MAP 3490" on page 150.	
4100	Hard device bus fabric error	Perform "MAP 3352" on page 132.	
4101	Soft device bus fabric error	Perform "MAP 3352" on page 132.	
4102	Device bus fabric performance degradation	Perform "MAP 3490" on page 150.	
4110	Unsupported enclosure function	Perform "MAP 3345" on page 127.	
4120	SAS cable VPD cannot be read	Perform "MAP 3461" on page 149 for a PCIe2 or a PCIe3 controller.	
4121	Configuration error, required cable is missing	Perform "MAP 3461" on page 149.	
4123	Configuration error, invalid cable vital product data	Perform "MAP 3461" on page 149.	

Table 20. URC (continued)

URC	Description text	Service action	Failing items
4170	Scatter list tag / sequence number error Logical block sequence number error on IOA to Host transfer	Controller recovered T10 DIF host bus error. Perform "MAP 3390" on page 134	
4171	Recovered scatter list tag / sequence number error Recovered logical block sequence number error on IOA to Host transfer	Controller T10 DIF host bus error. Perform "MAP 3390" on page 134	
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 131.	Disk drive
8008	Permanent cache battery pack failure Impending cache battery pack failure	Perform "MAP 3300" on page 116.	Cache battery pack
8009	Permanent cache battery pack failure Impending cache battery pack failure	Perform "MAP 3300" on page 116.	Cache battery pack
8150	Permanent IOA failure	Exchange the failing items in the Failing Items list one at a time.	I/O adapter
	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.	Any device on I/O signal cables
8151	IOA microcode error	Update adapter microcode. See "Updating the controller microcode" on page 82. 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

Table 20. URC (continued)

URC	Description text	Service action	Failing items
9000	IOA reserved area data check	Perform "MAP 3390" on page 134 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9001	IOA reserved area invalid data pattern	Perform "MAP 3337" on page 124 for a PCI-X or PCIe controller or MAP 3450 for a PCIe2 or a PCIe3 controller.	
9002	IOA reserved area LRC error	Perform "MAP 3390" on page 134 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9008	IOA does not support functions expected by devices	Perform "MAP 3330" on page 121 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9020	Array missing two or more devices with only one device present	Perform "MAP 3311" on page 118 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9021	Array missing two or more devices with two or more devices present	Perform "MAP 3311" on page 118 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9022	Exposed array is missing a required device	Perform "MAP 3311" on page 118 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9023	Array members not at required physical locations	Perform "MAP 3312" on page 118 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9024	Array not functional due to present hardware configuration	Perform "MAP 3390" on page 134.	
9025	Disk unit is not supported at its physical location	Perform "MAP 3310" on page 117.	
9026	Array not functional due to present hardware configuration	Perform "MAP 3390" on page 134.	
9027	Array is missing a device and parity is out of sync	Perform "MAP 3313" on page 119 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	

Table 20. URC (continued)

URC	Description text	Service action	Failing items
9028	Maximum number of arrays already exist	Perform "MAP 3390" on page 134 for a PCI-X or PCIe controller or "<"MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9029	Incorrect hardware configuration change has been detected	Perform "MAP 3390" on page 134 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9030	Array no longer protected due to missing or failed disk unit	Perform "MAP 3310" on page 117 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9031	Array protection temporarily suspended, protection resuming	Perform "MAP 3310" on page 117 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9032	Array exposed but still protected	Perform "MAP 3310" on page 117 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
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 errors corrected. Perform "MAP 3390" on page 134 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9042	Corrupt array parity detected on specified device	Perform "MAP 3390" on page 134 for a PCI-X or PCIe controller or MAP 3450 for a PCIe2 or a PCIe3 controller.	
9054	IOA resources not available due to previous problems	Perform "MAP 3321" on page 120 for a PCI-X or PCIe controller or MAP 3450 for a PCIe2 or a PCIe3 controller.	
9060	One or more disk pairs are missing from an array	Perform "MAP 3311" on page 118 for a PCI-X or PCIe controller or MAP 3450 for a PCIe2 or a PCIe3 controller.	

Table 20. URC (continued)

URC	Description text	Service action	Failing items
9061	One or more disks are missing from an array	Perform "MAP 3311" on page 118 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9062	One or more disks are missing from an array	Perform "MAP 3311" on page 118 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9063	Maximum number of functional arrays has been exceeded	Perform "MAP 3390" on page 134 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9073	Invalid multi-adapter configuration	Multiple controllers connected in an invalid configuration. Perform "MAP 3340" on page 125	
9074	Multiple controllers are not capable of similar functions or of controlling the same set of devices	Perform "MAP 3341" on page 126 for a PCIe2 or a PCIe3 controller.	
9075	Incomplete multipath connection between IOA and remote IOA	Perform "MAP 3349" on page 128 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9076	Configuration error, missing remote IOA	Missing remote controller. Perform "MAP 3347" on page 128	
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 signal cables
9090	Disk unit has been modified after the last known status	Perform "MAP 3333" on page 122 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
9091	Incorrect hardware configuration change has been detected	Perform "MAP 3333" on page 122 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	

Table 20. URC (continued)

URC	Description text	Service action	Failing items
9092	Disk unit requires initialization before use	Perform "MAP 3334" on page 122 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
FF3D	Soft IOA error recovered by the IOA Soft PCI bus 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 124 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	
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 131.	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" on page 131.	Disk drive
	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	No action required.	
	Media error recovered by device rewrite procedures		
FFF9	Soft media error. Sector reassignment recommended Device sector reassign successful	No action required.	

Table 20. URC (continued)

URC	Description text	Service action	Failing items
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 131.	Disk drive
FFFB	SCSI bus was reset	No action required.	
FFFC	Logical block guard error recovered by the device Logical block reference tag error recovered by the device	Device recovered T10 DIF device bus error. Perform "MAP 3350" on page 129.	
FFFD	Recovered logical block reference tag error detected by the IOA Logical block guard error recovered by the IOA	Controller recovered T10 DIF device bus error. Perform "MAP 3350" on page 129.	
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 129 for a PCI-X or PCIe controller or "MAP 3450" on page 145 for a PCIe2 or a PCIe3 controller.	

Maintenance analysis procedures

Use these procedures to resolve adapter, cache, or disk array problems associated with your controller.

For help in determining which MAP to use, see "Unit reference code tables" on page 109.

MAP 3300

Use this MAP to resolve the following problems:

- A permanent cache battery pack failure occurred (URC - 8008) for a PCI-X or PCIe controller
- Impending cache battery pack failure (URC - 9030) for a PCI-X or PCIe controller

Step 3300–1

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

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 "Forcing a rechargeable battery error" on page 83.

Go to "Step 3300-2."

Step 3300-2

Follow the actions recommended in Replacing a battery pack

When the problem is resolved, see the removal and replacement procedures topic for the system unit on which you are working, and do the "Verifying the repair" procedure.

MAP 3310

Use this MAP to resolve the following problems:

- Disk unit is not supported at its physical location (URC 9025) for a PCI-X or PCIe controller
- Array no longer protected due to missing or failed disk unit (URC - 9030) for a PCI-X or PCIe controller
- Array protection temporarily suspended (URC - 9031) for a PCI-X or PCIe controller
- Disk Array is degraded due to missing or failed disk (URC - 9032) for a PCI-X or PCIe controller

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 array status" on page 49.

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 details, see "Viewing array status" on page 49. 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 details, see "Physical disks" on page 95.
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 "RAID and JBOD formats" on page 52.
4. Assign the new disk as a hot spare for the disk. For details, see "Hot spare disks" on page 59.

When the problem is resolved, go to MAP 0410: Repair Checkout in *RS/6000 eServer 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 details, see "Viewing array status" on page 49. 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 details, see “Physical disks” on page 95.

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 eServer pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3311

Use this MAP to resolve the following problems:

- Array not functional due to present hardware configuration (URC 9020 / 9021 / 9022) for a PCI-X or PCIe controller
- One or more disk pairs are missing from an array (URC 9060) for a PCI-X or PCIe controller
- One or more disks are missing from an array (URC 9061 / 9062) for a PCI-X or PCIe controller

Step 3311–1

Identify the disks missing from the disk array by examining the error log and the **Display Hardware Status** screen. For details, see “Viewing array status” on page 49.

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 details, see “Deleting an IBM SAS RAID disk array” on page 54.

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. 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 eServer pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3312

Use this MAP to resolve the following problem:

- Array members not at required resource address (URC 9023) for a PCI-X or PCIe controller

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 details, see “Viewing device status” on page 46.

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 details, see “Deleting an IBM SAS RAID disk array” on page 54.
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, go to MAP 0410: Repair Checkout in *RS/6000 eServer pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3313

Use this MAP to resolve the following problem:

- Array not functional due to present hardware configuration (URC 9027) for a PCI-X or PCIe controller

Step 3313-1

Identify the adapter and disks related to the error by examining the error log and the **Display Hardware Status** screen. For details, see “Viewing device status” on page 46.

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” on page 120.

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 device status” on page 46.
 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 details, see “Deleting an IBM SAS RAID disk array” on page 54.
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. Select **Work with disk unit recovery**.
 2. Select **Initialize and format disk unit**.
 3. Select the drives you wish to format then press **Enter**.
 4. Run the **iprconfig** utility by typing `iprconfig`.

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

MAP 3320

Use this MAP to resolve the following problem:

- Cache data associated with attached disks cannot be found (URC 9010) for a PCI-X or PCIe controller

Step 3320–1

Contact your service support organization.

MAP 3321

Use this MAP to resolve the following problems:

- IOA resources not available due to previous problems (URC 9054) with a PCI-X or PCIe controller

The possible causes follow:

- The adapter or disks have been physically moved or changed such that the adapter does not support a function that the disks require.
- The disks were last used under the IBM i operating system.
- The disks were moved from a PCIe2 controller to a PCI-X or PCIe controller.

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 eServer pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3330

Use this MAP to resolve the following problem:

- 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 details, see “Viewing device status” on page 46.

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: 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 eServer pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3331

Use this MAP to resolve the following problem:

- Required cache data cannot be located for one or more disks (URC 9050) for a PCI-X or PCIe controller

Step 3331–1

Contact your service support organization.

MAP 3332

Use this MAP to resolve the following problem:

- Cache data exists for one or more missing or failed disks (URC 9051) for a PCI-X or PCIe controller

Step 3332-1

Contact your service support organization.

MAP 3333

Use this MAP to resolve the following problem:

- Disk unit has been modified after the last known status (URC 9090) for a PCI-X or PCIe controller
- Incorrect hardware configuration change has been detected (URC 9091) for a PCI-X or PCIe controller

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 eServer pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3334

Use this MAP to resolve the following problem:

- Disk unit requires initialization before use (URC 9092) for a PCI-X or PCIe controller

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 “Physical disks” on page 95) 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 details, see “Viewing device status” on page 46.

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” on page 123.

YES Go to “Step 3334-3” on page 123.

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 eServer 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 device status" on page 46.
 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 eServer 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 an IBM SAS RAID disk array” on page 54.

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 eServer pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3335

Use this MAP to resolve the following problem:

- Disk media format bad (URC FFF3) for a PCI-X or PCIe controller

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 details, see “Viewing device status” on page 46.

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 eServer pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3337

Use this MAP to resolve the following problem:

- 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” on page 125.

YES Go to “Step 3337-3” on page 125.

Step 3337-2

Contact your service support organization.

Exit this procedure.

Step 3337-3

Identify the affected disks by examining the error log and the **Display Hardware Status** screen. For further details, see "Viewing array status" on page 49.

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 arrays. For further details, see "Deleting an IBM SAS RAID disk array" on page 54.

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 an IBM SAS RAID disk array" on page 53.

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

MAP 3340

Use this MAP to resolve the following problem:

- Multiple controllers connected in an invalid configuration (URC 9073) for a PCI-X or PCIe controller

The possible causes follow:

- Incompatible adapters are connected to each other. Such an incompatibility includes invalid adapter combinations such as the following situations. See the feature comparison tables for PCIe and PCI-X cards for a list of the supported adapters and their attributes.
 - Adapters have different write cache sizes.
 - One adapter is not supported by the Linux operating system.
 - An adapter that does not support auxiliary cache is connected to an auxiliary cache adapter.
 - An adapter that supports multi-initiator and high availability is connected to another adapter that does not have the same support.
 - Adapters connected for multi-initiator and high availability are not operating in the same dual initiator configuration. For example, both are not set to the default value or both are not set to the JBOD HA single path value.
 - More than two adapters are connected for multi-initiator and high availability.
 - Adapter microcode levels are not up to date or are not at the same level of function.
 - One adapter, of a connected pair of adapters, is not operating under the Linux operating system. Connected adapters must both be controlled by the Linux operating system. Additionally, both adapters must be in the same logical partition if one adapter is an auxiliary cache adapter.
- An adapter is CCIN 572A but has a part number of either 44V4266 or 44V4404 (feature code 5900) that does not support multi-initiator and high availability.
- Adapters connected for multi-initiator and high availability are not cabled correctly. Each type of high availability configuration requires specific cables to be used in a supported manner.

Step 3340–1

Determine which of the possible causes apply to the current configuration and take the appropriate actions to correct the problem. If this action does not correct the error, contact your hardware service provider.

When the problem is resolved, see the removal and replacement procedures topic for the system unit on which you are working and complete the Verifying a repair procedure.

Related concepts:

“Feature comparison of SAS RAID cards” on page 2

Compare the main features of PCI-X, PCI Express (PCIe), PCIe2, and PCIe3 SAS RAID cards.

MAP 3341

Use this MAP to resolve the following problem:

Multiple controllers not capable of similar functions or not controlling the same set of devices (URC 9074) for a PCI-X or PCIe controller

Step 3341–1

Contact your service support organization.

MAP 3342

Use this MAP to resolve the following problem:

- Configuration error; incorrect connection between cascaded enclosure (URC 4010) for a PCI-X or PCIe controller

Step 3342–1

Incorrect SAS fabric connections. Contact your service provider.

MAP 3343

Use this MAP to resolve the following problem:

- Configuration error; connections exceed controller design limits (URC 4020) for a PCI-X or PCIe controller

Step 3343–1

Incorrect SAS fabric connections. Contact your service provider.

MAP 3344

Use this MAP to resolve the following problems:

- Configuration error, incorrect multipath connection (URC 4030) for a PCI-X or PCIe controller
- Configuration error, incorrect multipath connection between controller and enclosure detected (URC 4030) for a PCI-X or PCIe controller

The possible causes follow:

- Incorrect cabling to device enclosure.
Attention: Ensure that the YO-cable, YI-cable, or X-cable is routed along the right side of the rack frame (as viewed from the rear) when connecting the cable to a disk expansion drawer. Examine the device enclosure cabling, and correct the cabling as required.
- A failed connection caused by a failing component in the SAS fabric between, and including, the adapter and device enclosure.

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 the disk enclosure or removable media enclosure integrated in the system with no cables. For these configurations the SAS connections are integrated onto the system boards and a failed connection can be the result of a failed system board or integrated device enclosure.

Attention: Replacing RAID adapters is not recommended without assistance from your hardware service support organization when SAS fabric problems exist. The adapter might contain nonvolatile write cache data and configuration data for the attached disk arrays. Additional problems can be created by replacing an adapter when SAS fabric problems exist.

Step 3344–1

Incorrect SAS fabric connections. Contact your service provider.

MAP 3345

Use this MAP to resolve the following problem:

- Unsupported enclosure function detected (URC 4110) for a PCI-X or PCIe controller

The possible causes follow:

- Device enclosure or adapter microcode levels are not up to date
- Unsupported type of device enclosure or device

Step 3345–1

Ensure device enclosure or adapter microcode levels are up to date. If this does not correct the problem, contact your service provider.

MAP 3346

Use this MAP to resolve the following problem:

- Configuration error; incomplete multipath connection between enclosures and device detected (URC 4041) for a PCI-X or PCIe controller

The possible cause follows:

- A failed connection caused by a failing component within the device enclosure, including the device itself.

Note: The adapter is not a likely cause of this problem.

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 the disk enclosure or removable media enclosure integrated in the system with no cables. For these configurations the SAS connections are integrated onto the system boards and a failed connection can be the result of a failed system board or integrated device enclosure.

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

Step 3346–1

Ensure device enclosure and SAS disks are properly plugged in their slots. If this does not correct the problem, contact your service provider.

MAP 3347

Use this MAP to resolve the following problem:

- Multiple controllers not capable of similar functions or not controlling the same set of devices (URC 9076) for a PCI-X or PCIe controller

Step 3347–1

Contact your service support organization.

MAP 3348

Use this MAP to resolve the following problem:

- Attached enclosure does not support required multipath function (URC 4050) for a PCI-X or PCIe controller

Step 3348–1

Contact your service provider.

MAP 3349

Use this map to resolve the following problem:

- Incomplete multipath connection between controller and remote controller (URC 9075) for a PCI-X or PCIe controller

Step 3349–1

Contact your service provider.

MAP 3350

Use the following to perform SAS fabric problem isolation for a PCI-X or PCIe controller.

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 SAS 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 SAS/PCI-X buses. For these configurations, replacement of the RAID Enablement Card is unlikely to solve a SAS fabric-related problem because the SAS interface logic is on the system board.

Attention: Replacing RAID adapters is not recommended without assistance from your service support organization when SAS fabric 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 SAS fabric problems exist.

Attention: Removing functioning disks in a disk array is not recommended without assistance from your service support organization. A disk array may become *Degraded* or *Failed* if functioning disks are removed. Also, additional problems may be created.

Identify the SAS fabric 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 SAS configuration recently?

NO Go to “Step 3350–5” on page 130.

YES Go to “Step 3350–3.”

Step 3350–3

Check for the following problems:

- Cabling problems such as configurations that exceed the maximum cable lengths
- Ensure the SAS configuration 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 SAS cabling, refer to *RS/6000 pSeries Adapters, Devices, and Cable Information for Multiple Bus Systems*.

Did you find a problem?

NO Go to “Step 3350–5” on page 130.

YES Go to “Step 3350–4” on page 130.

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 SAS fabric-related failure occur?

NO Go to “Step 3350–14” on page 131.

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 details, see “Viewing array status” on page 49.

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 disks 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 disks from each *Degraded* disk array. For details, see “Physical disks” on page 95.
3. Reboot the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SAS fabric related failure occur?

NO Go to “Step 3350–14” on page 131.

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 SAS fabric?

NO Go to “Step 3350–10” on page 131.

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 SAS fabric related failure occur?

NO Go to “Step 3350–9” on page 131.

YES Go to “Step 3350–7” on page 130.

Step 3350–9

The last removable media device removed from the SAS fabric may be the cause of the SAS fabric 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 or 4096 byte per sector *JBOD* disks, *Hot Spare* disks, or *Advanced Function* disks) on the SAS fabric?

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 details, see “Physical disks” on page 95.
2. Reboot the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SAS fabric related failure occur?

NO Go to “Step 3350–12.”

YES Go to “Step 3350–10.”

Step 3350–12

The last disk removed from the SAS fabric may be the cause of the SAS fabric problems. Follow the repair procedures for that device.

Go to “Step 3350–14.”

Step 3350–13

Contact your service support organization.

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-SAS fabric related errors if they exist.

MAP 3351

Use the following procedure to determine what other field replaceable units (FRUs) besides a disk might need to be replaced in order to solve a problem.

You are sent here when an Advanced Function disk, a physical disk in 528 or 4224 bytes per 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 might 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 SAS 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 SAS/PCI-X busses. For these configurations, replacement of the RAID Enablement Card is unlikely to solve a SAS fabric related problem since the SAS interface logic is on the system board.

Attention: Replacing RAID adapters is not recommended without assistance from your service support organization when SAS fabric 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 SAS fabric problems exist.

Attention: Removing functioning disks in a disk array is not recommended without assistance from your service support organization. A disk array may become *Degraded* or *Failed* if functioning disks are removed. Also, additional problems may be created.

Step 3351-1

Identify the SAS fabric 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 SAS fabric in the following order.

1. Cable (if present)
2. Adapter (if SAS interface logic is on the adapter) or system board (if SAS 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.
4. 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 pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3352

Use this MAP to resolve the following problems:

- Device bus fabric error (URC 4100) for a PCI-X or PCIe controller
- Temporary device bus fabric error (URC 4101) for a PCI-X or PCIe controller

The possible causes follow:

- A failed connection caused by a failing component in the SAS fabric between, and including, the adapter and device enclosure.
- A failed connection caused by a failing component within the device enclosure, including the device itself.

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 SAS and PCI-X/PCIe bus interface logic integrated onto the system boards and use a pluggable RAID Enablement Card (a non-PCI form factor card) for these SAS/PCI-X/PCIe buses. For these configurations, replacement of the RAID Enablement Card is unlikely to solve a SAS related problem because the SAS interface logic is on the system board.
- Some systems have the disk enclosure or removable media enclosure integrated in the system with no cables. For these configurations the SAS connections are integrated onto the system boards and a failed connection can be the result of a failed system board or integrated device enclosure.

Attention: Replacing RAID adapters is not recommended without assistance from your service provider when SAS fabric 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 SAS fabric problems exist.

Attention: Removing functioning disks in a disk array is not recommended without assistance from your service provider. A disk array may become *Degraded* or *Failed* if functioning disks are removed and additional problems may be created.

Step 3352-1

Contact your service provider.

MAP 3353

Use this map to resolve the following problem:

- Multipath redundancy level got worse (URC 4060) for a PCI-X or PCIe controller

The possible causes follow:

- A failed connection caused by a failing component in the SAS fabric between, and including, the adapter and device enclosure.
- A failed connection caused by a failing component within the device enclosure, including the device itself.

Note: The failed connection was previously working, and may have already recovered.

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 SAS and PCI-X/PCIe bus interface logic integrated onto the system boards and use a pluggable RAID Enablement Card (a non-PCI form factor card) for these SAS/PCI-X/PCIe buses. For these configurations, replacement of the RAID Enablement Card is unlikely to solve a SAS related problem because the SAS interface logic is on the system board.
- Some systems have the disk enclosure or removable media enclosure integrated in the system with no cables. For these configurations the SAS connections are integrated onto the system boards and a failed connection can be the result of a failed system board or integrated device enclosure.

Attention: Replacing RAID adapters is not recommended without assistance from your service provider when SAS fabric 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 SAS fabric problems exist.

Attention: Removing functioning disks in a disk array is not recommended without assistance from your service provider. A disk array may become *Degraded* or *Failed* if functioning disks are removed and additional problems may be created.

Step 3353-1

Contact your service provider.

MAP 3390

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

The possible causes for 9002 are:

- One or more SAS devices were moved from a PCIe2 controller to a PCI-X or PCIe controller. If the device was moved from a PCIe2 controller to a PCI-X or PCIe controller, the **Detail Data** section of the hardware error log contains a reason for failure of Payload CRC Error. For this case, the error can be ignored and the problem is resolved if the devices are moved back to a PCIe2 controller or if the devices are formatted on the PCI-X or PCIe controller.

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 details, see "Viewing device status" on page 46.

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.

MAP 3410

Use this MAP to resolve the following problems:

- Disk unit not supported at its physical location (URC 9025) for a PCIe2 controller
- Array no longer protected due to missing or failed disk unit (URC - 9030) for a PCIe2 controller
- Array protection temporarily suspended (URC - 9031) for a PCIe2 controller
- Disk array is degraded due to missing or failed disk (URC - 9032) for a PCIe2 controller

Step 3410-1

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 array status" on page 49.

Go to “Step 3410–2”

Step 3410–2

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

NO Go to “Step 3410–3.”

YES Go to “Step 3410–4.”

Step 3410–3

The affected disk array must 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 completing the following steps:

1. Identify the failed disk by using the **Display Hardware Status** screen. For details, see “Viewing array status” on page 49. The failed disk has a *Failed* status.
2. Remove the failed disk and replace it with a new disk that is used as the hot spare. For details, see “Physical disks” on page 95.
3. The new disk must be formatted for advanced function to be used as a hot spare. If the new disk needs to be formatted, see “RAID and JBOD formats” on page 52.
4. Assign the new disk as a hot spare for the disk. For details, see “Hot spare disks” on page 59.

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

Step 3410–4

The failed disk must be replaced by completing the following steps:

1. Identify the failed disk by using the Display Hardware Status screen. For details, see “Viewing array status” on page 49. The failed disk has a *Failed* status.
2. Remove the failed disk, replace it with a new disk, and rebuild the disk unit data. For details, see “Physical disks” on page 95.

Note: The replacement disk must 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 eServer pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3411

Use this MAP to resolve the following problems:

- Array not functional due to present hardware configuration (URC 9020 / 9021 / 9022) for a PCIe2 controller
- One or more disk pairs are missing from an array (URC 9060) for a PCIe2 controller
- One or more disks are missing from an array (URC 9061 / 9062) for a PCIe2 controller

Step 3411–1

Identify the disks missing from the disk array by examining the error log and the **Display Hardware Status** screen. For details, see “Viewing array status” on page 49.

Go to “Step 3411–2” on page 136.

Step 3411–2

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 might be necessary to boot into Rescue mode), or reset the adapter by performing the following steps:
 1. Find the SCSI host number associated with the adapter by using the **Display Hardware Status** screen. For more details, see “Viewing array status” on page 49.
 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 details, see “Deleting an IBM SAS RAID disk array” on page 54.
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 want to format then press **Enter**.

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

MAP 3412

Use this MAP to resolve the following problem:

- Array members not at required physical location (URC 9023) for a PCIe2 controller

Step 3412–1

Identify the disks that are not at their required physical locations by using the error log entry and the **Display Hardware Status** screen. For details, see “Viewing device status” on page 46.

Go to “Step 3412–2.”

Step 3412–2

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 might be necessary to boot into *Rescue* mode), or reset the adapter by performing the following steps:
 1. Find the scsi host number associated with the adapter by using the **Display Hardware Status** screen. For more details, see “Viewing device status” on page 46.
 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 details, see “Deleting an IBM SAS RAID disk array” on page 54.
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 want to format then press **Enter**.

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

MAP 3413

Use this MAP to resolve the following problem:

- Array not functional due to present hardware configuration (URC 9027) for a PCIe2 controller

Step 3413-1

Identify the adapter and disks related to the error by examining the error log and the **Display Hardware Status** screen. For details, see “Viewing device status” on page 46.

Go to “Step 3413-2.”

Step 3413-2

Have the disks or adapter card been physically moved recently?

NO Contact your service support organization.

YES Go to “Step 3413-3.”

Step 3413-3

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 might be necessary to boot into *Rescue* mode), or reset the adapter by performing the following steps:
 1. Find the SCSI host number associated with the adapter by using the **Display Hardware Status** screen. For more details, see “Viewing device status” on page 46.
 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 details, see “Deleting an IBM SAS RAID disk array” on page 54.

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

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

MAP 3420

Use this MAP to resolve the following problem:

- Cache data associated with attached disks cannot be found (URC 9010) for a PCIe2 controller

Step 3420–1

Contact your service support organization.

MAP 3421

Use this MAP to resolve the following problems:

- IOA resources not available due to previous problems (URC 9054) for a PCIe2 controller

Step 3421–1

Perform the following steps:

1. Remove any new or replacement disks which are attached to the adapter.
2. Resolve the other errors which occurred at the same time as this error.

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

MAP 3430

Use this MAP to resolve the following problem:

- IOA does not support functions expected by devices (URC 9008) for a PCIe2 controller

The possible causes follow:

- The adapter or disks are physically moved or changed such that the adapter does not support a function which the disks require.
- The disks were last used under the IBM i operating system.
- The disks were moved from a PCI-X or PCIe controller to a PCIe2 controller and the disks had one of the following attributes which is not supported by a PCIe2 controller:
 - The disks were used in a disk array with a stripe-unit size of 16 KB, 64 KB, or 512 KB (a PCIe2 controller supports only a stripe-unit size of 256 KB).
 - The disks were used in a RAID level 5 or 6 disk array which had disks added to it after it was initially created (a PCIe2 controller does not support adding disks to a previously created RAID level 5 or 6 disk array).

Step 3430–1

Identify the affected disks by examining the error log and the **Display Hardware Status** screen. For details, see “Viewing device status” on page 46.

Go to “Step 3430–2.”

Step 3430–2

Have the adapter card or disks been physically moved recently?

NO Contact your service support organization.

YES Go to “Step 3430–3.”

Step 3430–3

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 to their original configuration. Then, either IPL the system or logical partition (it might be necessary to boot into *Rescue* mode) or reset the adapter by performing the following steps:
 1. Find the SCSI host number associated with the adapter by using the **Display Hardware Status** screen. For more details, see “Viewing device status” on page 46.
 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 want to format then press **Enter**.

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

MAP 3431

Use this MAP to resolve the following problem:

- Required cache data cannot be located for one or more disks (URC 9050) for a PCIe2 controller

Step 3431–1

Contact your service support organization.

MAP 3432

Use this MAP to resolve the following problem:

Cache data exists for one or more missing or failed disks (URC 9051) for a PCIe2 controller

Step 3432–1

Contact your service support organization.

MAP 3433

Use this MAP to resolve the following problems:

- Disk unit is modified after the last known status (URC 9090) for a PCIe2 controller
- Incorrect hardware configuration change is detected (URC 9091) for a PCIe2 controller

Step 3433–1

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

1. Find the SCSI host number associated with the adapter by using the **Display Hardware Status** screen. For details, see “Viewing device status” on page 46.
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.

Resolve any new errors which surface.

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

MAP 3434

Use this MAP to resolve the following problem:

- Disk unit requires initialization before use (URC 9092) for a PCIe2 controller

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. For example, not using the **Device Concurrent Maintenance** screen in **iprconfig** when concurrently removing and installing disks (see “Physical disks” on page 95) or not performing a normal power down of the system before 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 3434-1

Identify the affected disks by examining the error log and the **Display Hardware Status** screen. For details, see “Viewing device status” on page 46.

Go to “Step 3434-2.”

Step 3434-2

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

NO Go to “Step 3434-4.”

YES Go to “Step 3434-3.”

Step 3434-3

Resolve the other errors that occurred at the same time as this error.

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

Step 3434-4

Have the disks or adapter card been physically moved recently?

NO Go to “Step 3434-5.”

YES Go to “Step 3434-6” on page 141.

Step 3434-5

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

NO Go to “Step 3434-7” on page 141.

YES Go to “Step 3434-6” on page 141.

Step 3434–6

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

- IPL the system or logical partition (it might be necessary to boot into *Rescue* mode), or reset the adapter by performing the following steps:
 1. Find the SCSI host number associated with the adapter by using the **Display Hardware Status** screen. For more details, see “Viewing device status” on page 46.
 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 might be necessary to boot into *Rescue* mode), or reset the adapter by performing the following steps:
 1. Find the SCSI host number associated with the adapter by using the Display Hardware Status screen. For more details, see “Viewing device status” on page 46.
 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 eServer pSeries Diagnostic Information for Multiple Bus Systems*.

Step 3434–7

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 want to format then press **Enter**.
- If the disks are members of a disk array, delete the disk array. See “Deleting an IBM SAS RAID disk array” on page 54.

Note: In some rare scenarios, deleting the disk array will not have any 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 eServer pSeries Diagnostic Information for Multiple Bus Systems*.

MAP 3435

Use this MAP to resolve the following problem:

- Disk media format bad (URC FFF3) for a PCIe2 controller

The possible causes follow:

- Disk was being formatted and was powered off during this process.
- Disk was being formatted and was reset during this process.

Step 3435–1

Identify the affected disk by examining the error log and the **Display Hardware Status** screen. For details, see “Viewing device status” on page 46.

Go to "Step 3435-2."

Step 3435-2

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 want to format then press **Enter**.

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

MAP 3440

Use this MAP to resolve the following problem:

- Multiple controllers connected in an invalid configuration (URC 9073) for a PCIe2 controller

The possible causes follow:

- Incompatible adapters are connected to each other. Such an incompatibility includes invalid adapter combinations such as the following situations. See the feature comparison tables for PCIe and PCI-X cards for a list of the supported adapters and their attributes.
 - Adapters have different write cache sizes.
 - One adapter is not supported by the Linux operating system.
 - An adapter that does not support auxiliary cache is connected to an auxiliary cache adapter.
 - An adapter that supports multi-initiator and high availability is connected to another adapter that does not have the same support.
 - Adapters connected for multi-initiator and high availability are not operating in the same dual initiator configuration. For example, both are not set to the default value or both are not set to the JBOD HA single path value.
 - More than two adapters are connected for multi-initiator and high availability.
 - Adapter microcode levels are not up to date or are not at the same level of function.
 - One adapter, of a connected pair of adapters, is not operating under the Linux operating system. Connected adapters must both be controlled by the Linux operating system. Additionally, both adapters must be in the same logical partition if one adapter is an auxiliary cache adapter.
- An adapter is CCIN 572A but has a part number of either 44V4266 or 44V4404 (feature code 5900) that does not support multi-initiator and high availability.
- Adapters connected for multi-initiator and high availability are not cabled correctly. Each type of high availability configuration requires specific cables to be used in a supported manner.

Step 3440-1

Determine which of the possible causes applies to the current configuration and take the appropriate actions to correct the problem. If this action does not correct the error, contact your hardware service provider.

When the problem is resolved, see the removal and replacement procedures topic for the system unit on which you are working and complete the Verifying a repair procedure.

Related concepts:

“Feature comparison of SAS RAID cards” on page 2

Compare the main features of PCI-X, PCI Express (PCIe), PCIe2, and PCIe3 SAS RAID cards.

MAP 3441

Use this MAP to resolve the following problem:

- Multiple controllers not capable of similar functions or controlling the same set of devices (URC 9074) for a PCIe2 controller

Step 3441–1

Contact your service support organization.

MAP 3442

Use this MAP to resolve the following problem:

- Configuration error; incorrect connection between cascaded enclosure (URC 4010) for a PCIe2 controller

Step 3442–1

Incorrect SAS fabric connections. Contact your service provider.

MAP 3443

Use this MAP to resolve the following problem:

- Configuration error; connections exceed controller design limits (URC 4020) for a PCIe2 controller

Step 3443–1

Incorrect SAS fabric connections. Contact your service provider.

MAP 3444

Use this MAP to resolve the following problems:

- Configuration error, incorrect multipath connection (URC 4030) for a PCIe2 controller
- Configuration error, incorrect multipath connection between controller and enclosure detected (URC 4030) for a PCIe2 controller

The possible causes follow:

- Incorrect cabling to device enclosure.

Attention: Ensure that a YO-cable, YI-cable, or X-cable is routed along the right side of the rack frame (as viewed from the rear) when connecting the cable to a disk expansion drawer. Review the device enclosure cabling and correct the cabling as required.

- A failed connection caused by a failing component in the SAS fabric between, and including, the adapter and device enclosure.

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 the disk enclosure or removable media enclosure integrated in the system with no cables. For these configurations, the SAS connections are integrated onto the system boards and a failed connection can be the result of a failed system board or integrated device enclosure.

Attention: Do not replace the RAID adapters without assistance from your hardware service support organization when SAS fabric problems exist. The adapter might contain nonvolatile write cache data and configuration data for the attached disk arrays. Additional problems can be created by replacing an adapter when SAS fabric problems exist.

Step 3444–1

Incorrect SAS fabric connections. Contact your service provider.

MAP 3445

Use this MAP to resolve the following problem:

- Unsupported enclosure function detected (URC 4110) for a PCIe2 controller

The possible causes follow:

- Device enclosure or adapter microcode levels are not up to date
- Unsupported type of device enclosure or device

Step 3445–1

Ensure that device enclosure or adapter microcode levels are up to date. If this action does not correct the problem, contact your service provider.

MAP 3446

Use this MAP to resolve the following problem:

- Configuration error; incomplete multipath connection between enclosures and device detected (URC 4041) for a PCIe2 controller

The possible cause follows:

- A failed connection caused by a failing component within the device enclosure, including the device itself.

Note: The adapter is not a likely cause of this problem.

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 the disk enclosure or removable media enclosure integrated in the system with no cables. For these configurations, the SAS connections are integrated onto the system boards and a failed connection can be the result of a failed system board or integrated device enclosure.

Attention: Do not remove the functioning disks in a disk array without assistance from your Hardware Service Support organization. A disk array might become degraded or failed if functioning disks are removed and additional problems might be created.

Step 3446–1

Ensure that device enclosure and SAS disks are properly plugged in their slots. If this action does not correct the problem, contact your service provider.

MAP 3447

Use this MAP to resolve the following problem:

- Multiple controllers not capable of similar functions or controlling the same set of devices (URC 9076) for a PCIe2 controller

Step 3447-1

Contact your service support organization.

MAP 3448

Use this MAP to resolve the following problem:

- Attached enclosure does not support required multipath function (URC 4050) for a PCIe2 controller

Step 3448-1

Contact your service provider.

MAP 3449

Use this map to resolve the following problem:

- Incomplete multipath connection between controller and remote controller (URC 9075) for a PCIe2 controller

Note: This problem is not expected to occur for a PCIe2 controller.

Step 3449-1

Proceed to “MAP 3490” on page 150

MAP 3450

Use the following to perform SAS fabric problem isolation for a PCIe2 controller.

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 the disk enclosure or removable media enclosure integrated in the system with no cables. For these configurations, the SAS connections are integrated onto the system boards and a failed connection can be the result of a failed system board or integrated device enclosure.

Attention:

- Do not replace the RAID adapters without assistance from your service support organization when SAS fabric problems exist. Because the adapter may contain non-volatile write cache data and configuration data for the attached disk arrays, additional problems might be created by replacing an adapter when SAS fabric problems exist.
- Do not remove the functioning disks in a disk array without assistance from your service support organization. A disk array might become *Degraded* or *Failed* if functioning disks are removed. Also, additional problems might be created.

Identify the SAS fabric on which the problem is occurring by examining the error log entry.

Go to “Step 3450-2.”

Step 3450-2

Have changes been made to the SAS configuration recently?

NO Go to “Step 3450-5” on page 146.

YES Go to “Step 3450–3.”

Step 3450–3

Check for the following problems:

- Cabling problems such as configurations that exceed the maximum cable lengths
- Ensure that the SAS configuration 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 details about supported SAS cabling, see *RS/6000 pSeries Adapters, Devices, and Cable Information for Multiple Bus Systems*.

Did you find a problem?

NO Go to “Step 3450–5.”

YES Go to “Step 3450–4.”

Step 3450–4

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 SAS fabric-related failure occur?

NO Go to “Step 3450–14” on page 148.

YES Go to “Step 3450–5.”

Step 3450–5

Determine if any of the disk arrays on the adapter are in a *Degraded* state. For details, see “Viewing array status” on page 49.

Does any disk array have a *Degraded* state?

NO Go to “Step 3450–7” on page 147.

YES Go to “Step 3450–6.”

Step 3450–6

1. Identify the failed disks by first finding disk arrays with a *Degraded* state, then finding disks on those arrays with a state of *Failed*.
2. Remove the failed disks from each *Degraded* disk array. For details, see “Physical disks” on page 95.
3. Reboot the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SAS fabric-related failure occur?

NO Go to “Step 3450–14” on page 148.

YES Go to "Step 3450-7."

Step 3450-7

Are there any nonessential removable media devices (such as Tape, CD, and DVD) on the SAS fabric?

NO Go to "Step 3450-10."

YES Go to "Step 3450-8."

Step 3450-8

1. Power off the system or logical partition.
2. Remove one of the nonessential 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 SAS fabric-related failure occur?

NO Go to "Step 3450-9."

YES Go to "Step 3450-7."

Step 3450-9

The last removable media device removed from the SAS fabric might be the cause of the SAS fabric problems. Follow the repair procedures for that device.

Go to "Step 3450-14" on page 148.

Step 3450-10

Are there any non essential disks that are not disk array members (such as 512 or 4096 byte per sector *JBOD* disks, *Hot Spare* disks, or *Advanced Function* disks) on the SAS fabric?

NO Go to "Step 3450-13" on page 148.

YES Go to "Step 3450-11."

Step 3450-11

1. Remove one of the nonessential disk devices. For details, see "Physical disks" on page 95.
2. Reboot the system or logical partition. If you cannot power on normally, boot to *Rescue* mode. Examine the error log.

Did a SAS fabric-related failure occur?

NO Go to "Step 3450-12."

YES Go to "Step 3450-10."

Step 3450-12

The last disk removed from the SAS fabric might be the cause of the SAS fabric problems. Follow the repair procedures for that device.

Go to "Step 3450-14" on page 148.

Step 3450–13

Contact your service support organization.

Step 3450–14

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

MAP 3452

Use this MAP to resolve the following problems:

- Device bus fabric error (URC 4100) for a PCIe2 controller
- Temporary device bus fabric error (URC 4101) for a PCIe2 controller

The possible causes follow:

- A failed connection caused by a failing component in the SAS fabric between, and including, the adapter and device enclosure.
- A failed connection caused by a failing component within the device enclosure, including the device itself.

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 the disk enclosure or removable media enclosure integrated in the system with no cables. For these configurations, the SAS connections are integrated onto the system boards and a failed connection can be the result of a failed system board or integrated device enclosure.

Attention:

- Do not replace the RAID adapters without assistance from your service provider when SAS fabric problems exist. Because the adapter might contain non-volatile write cache data and configuration data for the attached disk arrays, additional problems can be created by replacing an adapter when SAS fabric problems exist.
- Do not remove the functioning disks in a disk array without assistance from your service provider. A disk array might become *Degraded* or *Failed* if functioning disks are removed and additional problems might be created.

Step 3453–1

Contact your service provider.

MAP 3453

Use this map to resolve the following problem:

- Multipath redundancy level got worse (URC 4060) for a PCIe2 controller

The possible causes follow:

- A failed connection caused by a failing component in the SAS fabric between, and including, the adapter and device enclosure.
- A failed connection caused by a failing component within the device enclosure, including the device itself.

- A failed connection caused by a failing component between two SAS adapters, including the AA-cable or the SAS adapters themselves.

Note: The failed connection was previously working, and might already be recovered.

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 the disk enclosure or removable media enclosure integrated in the system with no cables. For these configurations, the SAS connections are integrated onto the system boards and a failed connection can be the result of a failed system board or integrated device enclosure.

Attention:

- Do not replace the RAID adapters without assistance from your service provider when SAS fabric problems exist. Because the adapter might contain non-volatile write cache data and configuration data for the attached disk arrays, additional problems can be created by replacing an adapter when SAS fabric problems exist.
- Do not remove the functioning disks in a disk array without assistance from your service provider. A disk array might become *Degraded* or *Failed* if functioning disks are removed and additional problems might be created.

Step 3453–1

Contact your service provider.

MAP 3454

Use this map to resolve the following problem:

Device bus fabric performance degradation (URC 4102) for a PCIe2 controller

Note: This problem is not common for a PCIe2 controller.

Step 3454–1

Proceed to “MAP 3490” on page 150

MAP 3460

Use this map to resolve the following problems:

- Scatter list tag / sequence number error (URC 4170) for a PCIe2 controller
- Recovered scatter list tag / sequence number error (URC 4171) for a PCIe2 controller

Note: This problem is not common for a PCIe2 controller.

Step 3460–1

Proceed to “MAP 3490” on page 150

MAP 3461

Use this map to resolve the following problems:

- Configuration error, cable VPD cannot be read (URC 4120) for a PCIe2 controller
- Configuration error, required cable is missing (URC 4121) for a PCIe2 controller

Note: This problem is not common for a PCIe2 controller.

Step 3461–1

Proceed to “MAP 3490”

MAP 3490

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

Step 3490–1

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

Go to “Step 3490–2.”

Step 3490–2

Collect the current disk array configuration. For details, see “Viewing device status” on page 46

Go to “Step 3490–3.”

Step 3490–3

Collect any `ipr` dump files that might be applicable to the problem. The files are located at `/var/log/`.

Go to “Step 3490–4.”

Step 3490–4

Contact your service support organization.

MAP 3495

Use this MAP to resolve the following problem:

- I/O adapter (IOA) exceeded maximum operating temperature (URC 4080) for a PCIe2 controller.

Step 3495–1

Determine which of the following items are the cause for the exceeded maximum operating temperature and take the appropriate actions. If this action does not correct the error, contact your hardware service provider.

- The adapter is installed in an unsupported system. Verify that the adapter is supported on this system by checking the PCI adapter information by feature type.
- The adapter is installed in an unsupported slot location within the system unit or an I/O enclosure. Verify that the adapter is located in a supported slot location. See the PCI adapter placement information for the machine-type model (MTM) where the adapter is located.
- The adapter is installed in a supported system, but the system is not operating in the required airflow mode. For example, the adapter is in an 8202-E4B or 8205-E6B system that is running in acoustic mode. Verify any system specific requirements for this adapter by checking the PCI adapter information by feature type.
- Fan failures or obstructions affecting the cooling of the adapter.

Note: The adapter that is logging this error continues to log this error while it remains at a temperature higher than the maximum operating temperature or every time it exceeds the maximum operating temperature.

When the problem is resolved, see the removal and replacement procedures topic for the system unit on which you are working and verify the repair. For instructions, see Verifying a repair.

Notices

This information was developed for products and services offered in the US.

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

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

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions 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. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

The drawings and specifications contained herein shall not be reproduced in whole or in part without the written permission of IBM.

IBM has prepared this information for use with the specific machines indicated. IBM makes no representations that it is suitable for any other purpose.

IBM's computer systems contain mechanisms designed to reduce the possibility of undetected data corruption or loss. This risk, however, cannot be eliminated. Users who experience unplanned outages, system failures, power fluctuations or outages, or component failures must verify the accuracy of operations performed and data saved or transmitted by the system at or near the time of the outage or failure. In addition, users must establish procedures to ensure that there is independent data verification before relying on such data in sensitive or critical operations. Users should periodically check IBM's support websites for updated information and fixes applicable to the system and related software.

Homologation statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

Accessibility features for IBM Power Systems servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power Systems servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power Systems servers use the latest W3C Standard, WAI-ARIA 1.0 (www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and Web Content Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power Systems servers.

The IBM Power Systems servers online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the Accessibility section of the IBM Knowledge Center help (www.ibm.com/support/knowledgcenter/doc/kc_help.html#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power Systems servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power Systems servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power Systems servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power Systems servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see IBM Accessibility (www.ibm.com/able).

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as the customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM’s Privacy Policy at <http://www.ibm.com/privacy> and IBM’s Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at <http://www.ibm.com/software/info/product-privacy>.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

Electronic emission notices

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Class A Notices

The following Class A statements apply to the IBM servers that contain the POWER8 processor and its features unless designated as electromagnetic compatibility (EMC) Class B in the feature information.

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. IBM is not 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.

Industry Canada Compliance Statement

CAN ICES-3 (A)/NMB-3(A)

European Community Compliance Statement

This product is in conformity with the protection requirements of EU Council Directive 2014/30/EU on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

European Community contact:
IBM Deutschland GmbH
Technical Regulations, Abteilung M456

IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 800 225 5426
email: halloibm@de.ibm.com

Warning: 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.

VCCI Statement - Japan

この装置は、クラスA 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。 VCCI-A

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

This is a Class A product based on the standard of the VCCI Council. If this equipment is used in a domestic environment, radio interference may occur, in which case, the user may be required to take corrective actions.

Japan Electronics and Information Technology Industries Association Statement

This statement explains the Japan JIS C 61000-3-2 product wattage compliance.

(一社) 電子情報技術産業協会 高調波電流抑制対策実施
要領に基づく定格入力電力値： Knowledge Centerの各製品の
仕様ページ参照

This statement explains the Japan Electronics and Information Technology Industries Association (JEITA) statement for products less than or equal to 20 A per phase.

高調波電流規格 JIS C 61000-3-2 適合品

This statement explains the JEITA statement for products greater than 20 A, single phase.

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 : 6 (単相、PFC回路付)
- 換算係数 : 0

This statement explains the JEITA statement for products greater than 20 A per phase, three-phase.

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 : 5 (3相、PFC回路付)
- 換算係数 : 0

Electromagnetic Interference (EMI) Statement - People's Republic of China

声 明

此为 A 级产品,在生活环境中,该产品可能会造成无线电干扰。在这种情况下,可能需要用户对其干扰采取切实可行的措施。

Declaration: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may need to perform practical action.

Electromagnetic Interference (EMI) Statement - Taiwan

警告使用者：

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

The following is a summary of the EMI Taiwan statement above.

Warning: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user will be required to take adequate measures.

IBM Taiwan Contact Information:

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

Electromagnetic Interference (EMI) Statement - Korea

이 기기는 업무용 환경에서 사용할 목적으로 적합성평가를 받은 기기로서 가정용 환경에서 사용하는 경우 전파간섭의 우려가 있습니다.

Germany Compliance Statement

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55022 / EN 55032 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

EN 55022 / EN 55032 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden: "Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem "Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)". Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse A

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road

Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 / EN 55032 Klasse A.

Electromagnetic Interference (EMI) Statement - Russia

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать
радиопомехи, для снижения которых необходимы
дополнительные меры

Class B Notices

The following Class B statements apply to features designated as electromagnetic compatibility (EMC) Class B in the feature installation information.

Federal Communications Commission (FCC) Statement

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation.

This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation.

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

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

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for any radio or television interference caused by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate this 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.

Industry Canada Compliance Statement

CAN ICES-3 (B)/NMB-3(B)

European Community Compliance Statement

This product is in conformity with the protection requirements of EU Council Directive 2014/30/EU on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

European Community contact:
IBM Deutschland GmbH
Technical Regulations, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 800 225 5426
email: halloibm@de.ibm.com

VCCI Statement - Japan

この装置は、クラスB情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。

取扱説明書に従って正しい取り扱いをして下さい。 VCCI-B

Japan Electronics and Information Technology Industries Association Statement

This statement explains the Japan JIS C 61000-3-2 product wattage compliance.

(一社) 電子情報技術産業協会 高調波電流抑制対策実施
要領に基づく定格入力電力値 : Knowledge Centerの各製品の
仕様ページ参照

This statement explains the Japan Electronics and Information Technology Industries Association (JEITA) statement for products less than or equal to 20 A per phase.

高調波電流規格 JIS C 61000-3-2 適合品

This statement explains the JEITA statement for products greater than 20 A, single phase.

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 : 6 (単相、PFC回路付)
- 換算係数 : 0

This statement explains the JEITA statement for products greater than 20 A per phase, three-phase.

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- 回路分類 : 5 (3相、PFC回路付)
- 換算係数 : 0

IBM Taiwan Contact Information

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

Germany Compliance Statement

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse B EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55022/ EN 55032 Klasse B ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem "Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)". Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse B

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 (0) 800 225 5426
email: HalloIBM@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022/ EN 55032 Klasse B.

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

