



Power Systems

SAS RAID controllers for IBM i







Power Systems

SAS RAID controllers for IBM i

**Note**

Before using this information and the product it supports, read the information in “Notices,” on page 57, “Safety notices” on page v, 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 POWER6® processor and to all associated models.

© Copyright International Business Machines Corporation 2009.

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

---

# Contents

<b>Safety notices</b>	<b>v</b>
<b>Chapter 1. What's new in SAS RAID controllers for IBM i</b>	<b>1</b>
<b>Chapter 2. SAS RAID controllers for IBM i</b>	<b>3</b>
Feature comparison of SAS RAID cards	3
PCI-X SAS RAID card comparison	3
PCIe SAS RAID card comparison	6
SAS architecture	8
Disk arrays	9
Supported RAID levels	10
RAID 5	10
RAID 6	11
System mirroring	11
Disk array capacities	11
RAID level summary	12
Valid states for disk arrays and physical disk units	12
States for disk arrays	12
States for disk units	13
Auxiliary write cache	13
Auxiliary write cache adapter	14
Viewing link status information	15
<b>Chapter 3. Controller software</b>	<b>17</b>
Verifying the controller software	17
<b>Chapter 4. Common controller and disk array management tasks</b>	<b>19</b>
Viewing IBM SAS disk information	19
Considerations for solid-state drives	19
<b>Chapter 5. Dual storage IOA configurations</b>	<b>21</b>
Possible disk storage IOA configurations	21
Dual storage IOA functions	22
Dual storage IOA function attributes	24
Viewing dual storage IOA attributes	24
SAS cabling considerations	25
Performance considerations	26
Dual storage IOA access optimization	26
Installing dual storage IOA configurations	28
<b>Chapter 6. SAS RAID controller maintenance</b>	<b>31</b>
Rechargeable battery maintenance	31
Displaying rechargeable battery information	31
Error state	33
Forcing a rechargeable battery error	34
Replacing a battery pack	34
Replacing a 572B nonconcurrent maintainable battery pack	35
Replacing a 57B7 concurrent maintainable battery pack	37
Replacing a 574E concurrent maintainable battery pack	38
Replacing a 572F/575C card set concurrent maintainable battery pack	39
Separating the 572F/575C card set and moving the cache directory card	40
Viewing SAS fabric path information	45
Example: Using SAS fabric path information	49

<b>Chapter 7. SAS address and physical location information . . . . .</b>	<b>55</b>
<b>Appendix. Notices . . . . .</b>	<b>57</b>
Trademarks . . . . .	58
Electronic emission notices . . . . .	58
Class A Notices . . . . .	58
Terms and conditions . . . . .	62

---

## 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, a safety information booklet is included in the publications package shipped with the product. The booklet 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 in the booklet. You should also refer to the booklet any time you do not clearly understand any safety information in the U.S. English publications.

### 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

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. (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)

**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.



---

## **Chapter 1. What's new in SAS RAID controllers for IBM i**

This topic collection is new for the October 2009 release.



---

## Chapter 2. SAS RAID controllers for IBM i

Find usage and maintenance information regarding controllers for the serial-attached SCSI (SAS) Redundant Array of Independent Disks (RAID) for IBM i. Use this information with your specific system unit and operating system documentation. General information is intended for all users of this product. Service information is intended for a service representative trained on the system unit and the subsystem being serviced.

The SAS RAID controllers for IBM i have the following features:

- PCI-X 266 system interface or PCI Express (PCIe) system interface.
- Physical link speed of 3 Gbps SAS that supports transfer rates of 300 MB per second.
- Support for SAS devices and nondisk Serial Advanced Technology Attachment (SATA) devices.
- Optimized for SAS disk configurations that use dual paths through dual expanders for redundancy and reliability.
- Controller managed path redundancy and path switching for multiported SAS devices.
- Embedded PowerPC® Reduced Instruction Set Computer (RISC) processor, hardware XOR Direct Memory Access engine, and hardware finite field multiplier (FFM) DMA engine for RAID 6.
- Some adapters that support nonvolatile write cache.
- Support for RAID 5 and RAID 6 disk arrays and system mirroring.
- 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 and 6 disk arrays and system mirroring.
  - 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 bytes per sector, providing cyclical redundancy checking (CRC) and logically bad-block checking.
  - Optimized hardware for RAID 5 and 6 sequential write workloads.
  - Optimized skip read-and-write disk support for transaction workloads.
- Supports a maximum of 64 advanced function disks with a maximum of 255 devices.

**Note:** The number of all physical SAS and SATA devices plus the number of logical RAID disk arrays must be less than 255 per controller.

---

### Feature comparison of SAS RAID cards

Compare the main features of PCI-X and PCI Express (PCIe) SAS RAID cards for IBM i.

The tables in this section provide a breakdown of the main features of the SAS RAID PCI-X and PCIe controller cards.

#### PCI-X SAS RAID card comparison

Use the table in this topic to compare the features of PCI-X SAS RAID cards for IBM i. There are also images of adapters for you to view.

Table 1. PCI-X SAS RAID controller card comparison

Features	572A	572C	572F and 575C	57B8
Description	PCI-X 266 Ext Dual-x4 3 Gb SAS 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	Low profile 64 bit PCI-X	Planar integrated	Long 64 bit PCI-X, double-wide card set	Planar RAID enablement
Physical links	8 (two mini SAS 4x connectors)	8 <sup>1</sup>	12 (bottom 3 mini SAS 4x connectors) and 2 (top mini SAS 4x connector for high availability only)	8 <sup>1</sup>
RAID levels supported	RAID 5 <sup>3</sup> , RAID 6 <sup>3</sup> , system mirroring	System mirroring	RAID 5, RAID 6, system mirroring	RAID 5, RAID 6, system mirroring
Write cache size			Up to 1.5 Gb (compressed)	175 MB
Read cache size			Up to 1.6 Gb (compressed)	
Cache battery pack technology			Lithium ion	Not applicable <sup>2</sup>
Cache battery concurrent maintenance	No	No	Yes <sup>4</sup>	Not applicable <sup>2</sup>
Cache data present LED	No	No	No	No
Removable cache card	No	No	No	No
Auxiliary write cache (AWC) support	No	No	Yes	Yes
Dual storage IOA configuration	No	No	Yes	No
Requires dual storage IOA configuration	No	No	No	No
<ol style="list-style-type: none"> <li>1. Some systems provide an external mini SAS 4x connector from the integrated backplane controller.</li> <li>2. The controller contains battery-backed cache, but the battery power is supplied by the 57B7 controller through the backplane connections.</li> <li>3. 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.</li> <li>4. 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.</li> </ol>				

## Adapter graphics

View the SAS RAID controllers.




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




Figure 2. CCIN 57B8 planar RAID enablement card




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

## Related concepts

Chapter 5, “Dual storage IOA configurations,” on page 21

You can increase availability using a dual storage I/O adapter (IOA) configuration to connect multiple controllers to a common set of disk expansion drawers and the included disks and disk arrays.

“Dual storage IOA functions” on page 22

Consider these factors when using dual storage I/O adapter (IOA) functions.

## PCIe SAS RAID card comparison

Use the table in this topic to compare the features of a PCI Express (PCIe) SAS RAID cards for IBM i. There are also images of adapters for you to view.

Table 2. PCIe SAS RAID controller card comparison

Features	57B7	57B3	574E
Description	PCIe x1 Auxiliary Cache adapter	PCIe x8 Ext Dual-x4 3 Gb SAS adapter	PCIe x8 Ext Dual-x4 3 Gb SAS RAID adapter
Form factor	Planar Auxiliary Cache	PCIe x8	PCIe x8
Physical links	2	8 (two mini SAS 4x connectors)	8 (two mini SAS 4x connectors)
RAID levels supported		RAID 5 <sup>1</sup> , RAID 6 <sup>1</sup> , system mirroring	RAID 5, RAID 6, system mirroring
Write cache size	175 MB		380 MB
Read cache size			
Cache battery pack technology	Lithium ion		Lithium ion
Cache battery concurrent maintenance	Yes	No	Yes
Cache data present LED	Yes	No	Yes
Removable cache card	No	No	Yes

Table 2. PCIe SAS RAID controller card comparison (continued)

Features	57B7	57B3	574E
Auxiliary write cache (AWC) support	Yes	No	No
Dual storage IOA configuration	No	No	Yes
Requires dual storage IOA configuration	No	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.			

## Adapter graphics

View the SAS RAID controllers.




Figure 4. CCIN 57B7 Planar Auxiliary Cache




Figure 5. CCIN 57B3 PCIe x8 Ext Dual-x4 3 Gb SAS adapter




Figure 6. CCIN 574E PCIe x8 Ext Dual-x4 3 Gb SAS RAID adapter

### Related concepts

Chapter 5, “Dual storage IOA configurations,” on page 21

You can increase availability using a dual storage I/O adapter (IOA) configuration to connect multiple controllers to a common set of disk expansion drawers and the included disks and disk arrays.

“Dual storage IOA functions” on page 22

Consider these factors when using dual storage I/O adapter (IOA) functions.

## SAS architecture

Serial-attached SCSI (SAS) architecture defines a serial device interconnect and transport protocol that 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 (phys) 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. Phys are contained in SAS ports which contain one or more phys. A port is a wide port if there are more than one phy in the port. If there is only one phy in the port, it is a narrow port. A port is identified by a unique SAS worldwide name (also called 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 by using either the SCSI command set or the ATA/ATAPI command set depending on the device type.

A SAS expander enables connections between a controller port and multiple I/O device ports by routing connections between the expander ports. Only a single connection through an expander can exist 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, more than one path to the device can exist when there are expander devices included in the path.

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

The following example SAS subsystem shows some of the concepts described in this SAS overview. A controller is shown with eight SAS phys. Four of those phys are connected into two different wide ports. One connector contains four phys grouped into two ports. The connectors have no significance in SAS other than causing a physical wire connection. The four-phy connector can contain between one and four ports depending on the type of cabling that is used. The uppermost port in the figure shows a controller-wide port number 6 that consists of phy numbers 6 and 7. Port 6 connects to an expander, which attaches to one of the dual ports of the I/O devices. The dashed red line indicates a path between the controller and an I/O device. Another path runs 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.




Figure 7. Example SAS Subsystem

## Disk arrays

*Disk arrays* are groups of disks that work together with a specialized array controller to take advantage of potentially higher data transfer rates and data redundancy.

Disk arrays use RAID technology to offer data redundancy and to provide improved data transfer rates over single large disks. If a disk failure occurs, the disk can typically be replaced without interrupting normal system operation.

## Data redundancy

The disk array controller tracks how the data is distributed across the disks. RAID 5 and RAID 6 disk arrays provide data redundancy, ensuring that data is not lost if a disk in the array fails. If a disk failure occurs, the disk can typically be replaced without interrupting normal system operations. System mirroring provides data redundancy by mirroring the same data across pairs of disks.

## Supported RAID levels

The RAID level of a disk array determines how data is stored on the disk array and the level of protection that is provided.

If a part of the RAID system fails, different RAID levels help to recover lost data in different ways. If a single drive fails within an array, the array controller can reconstruct the data for the failed disk by using the data stored on other hard drives within the array. This data reconstruction has little or no impact to current system programs and users. The controller supports RAID levels 5 and 6 as well as system mirroring. Not all controllers support all RAID levels. Each RAID level supported by the controller has its own attributes and uses a different method of writing data. The following information provides details for each supported RAID level.

### Related concepts

“PCI-X SAS RAID card comparison” on page 3

Use the table in this topic to compare the features of PCI-X SAS RAID cards for IBM i. There are also images of adapters for you to view.

“PCIe SAS RAID card comparison” on page 6

Use the table in this topic to compare the features of a PCI Express (PCIe) SAS RAID cards for IBM i. There are also images of adapters for you to view.

### Related information

 Device parity protection concepts

## RAID 5

Learn how data is written to a RAID 5 array.

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




Figure 8. 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 information

 RAID 5 concepts

## RAID 6

Learn how data is written to a RAID 6 array.

RAID 6 stripes data across all disks in the array. RAID level 6 also writes array "P" and "Q" parity data. The P and Q parity data, 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 9. 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 a one or two Failed disks is said to be operating in Degraded mode. Whenever data is read from a Degraded disk array, the array controller recalculates the data on the failed 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 information

 RAID 6 concepts

## System mirroring

Mirrored protection is beneficial if you have a multibus system or a system with a large single bus. A greater number of disk units provides more opportunity for failure and increased recovery time.

Refer to Mirrored protection for more information.

## Disk array capacities

These guidelines will help you calculate the capacity of a disk array.

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

### RAID 5

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

## RAID 6

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

## System mirroring

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

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

## 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 3. RAID level summary*

RAID level	Data redundancy	Usable disk capacity	Read performance	Write performance	Devices per array
RAID 5	Very good	67% - 94%	Very good	Good	Minimum: 3 Maximum: 18
RAID 6	Excellent	50% - 89%	Very good	Fair to good	Minimum: 4 Maximum: 18
System mirroring	Excellent	50%	Excellent	Very good	Not applicable

## RAID 5

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

## System Mirroring

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.

---

## Valid states for disk arrays and physical disk units

Disk arrays and physical disks have several operational states.

## States for disk arrays

There are four valid states for disk arrays.

View disk arrays by completing the following steps:

1. Select Work with disk units on the Use System Service Tools (SST) menu.
2. Select Display disk configuration from the Work with Disk Configuration display.
3. Select Display device parity status on the Display Disk Configuration display.

The valid four states for IBM SAS Disk Arrays are RAID 5, RAID 6, RAID 5/Availability, and RAID 6/Availability.

## RAID 5

Indicates that this parity set is configured as a RAID 5 parity set. If any one unit in the set should fail, the other units in the set would be able to sustain the functions of the failed unit. RAID 5 requires one drive's worth of capacity per array to hold the parity data.

## RAID 6

Indicates that this parity set is configured as a RAID 6 parity set. If one or two units in the set should fail, the other units in the set would be able to sustain the functions of the failed units. RAID 6 requires two drives' worth of capacity per array to hold the parity data. If one unit in a RAID 6 parity set should fail, the parity set would have protection equivalent to a RAID 5 parity set.

## RAID 5/Availability

Indicates that this parity set is configured as a RAID 5 parity set which is optimized for Availability. There is at most one disk per I/O bus.

## RAID 6/Availability

Indicates that this parity set is configured as a RAID 6 parity set which is optimized for Availability. There are at most two disks per I/O bus.

## States for disk units

There are seven valid states for physical disk units.

View disk arrays by completing the following steps.

1. Select **Work with disk units** on the Use System Service Tools (SST) menu.
2. Select **Display disk configuration** from the Work with Disk Configuration display.
3. Select **Display device parity status** on the Display Disk Configuration display.

The valid states for disk units are Active, Read/Write Protected, Failed, Rebuilt/Rebuilding, Unprotected, Power Loss, and Unknown.

**Active** The disk is functioning correctly.

### Read/Write Protected

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

**Failed** The controller cannot communicate with the disk unit, or the disk unit the cause of the disk array being in a Degraded (exposed) state.

### Rebuilt/Rebuilding

The data on one of the units in the parity set is being rebuilt from other units in the set. All the units in the set will have this status. Some subsystems will report their progress to the system. If they do, it will be displayed in the status field. If another unit in the disk unit subsystem fails, data could be lost.

### Unprotected

This unit is operational. However, another unit in the disk unit subsystem has failed. If another unit in the disk unit subsystem fails, data could be lost.

### Power Loss

This unit has lost power. The unit was previously connected to the controller but is no longer detected.

### Unknown

The state of the disk could not be determined.

---

## Auxiliary write cache

A duplicate, nonvolatile copy of write cache data can be preserved.

## 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 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 dual storage IOA connections and AWC connections. Connecting controllers in a dual storage IOA 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.

**Important:** If a failure of either the RAID controller or the Auxiliary Cache occurs, the Isolation and Recovery procedures for the System Reference Codes (SRCs) in the Service Action Log (SAL) or Product Activity Log (PAL) must be followed precisely.

The RAID controller and the AWC adapter each require a PCI bus connection and 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, write out 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.

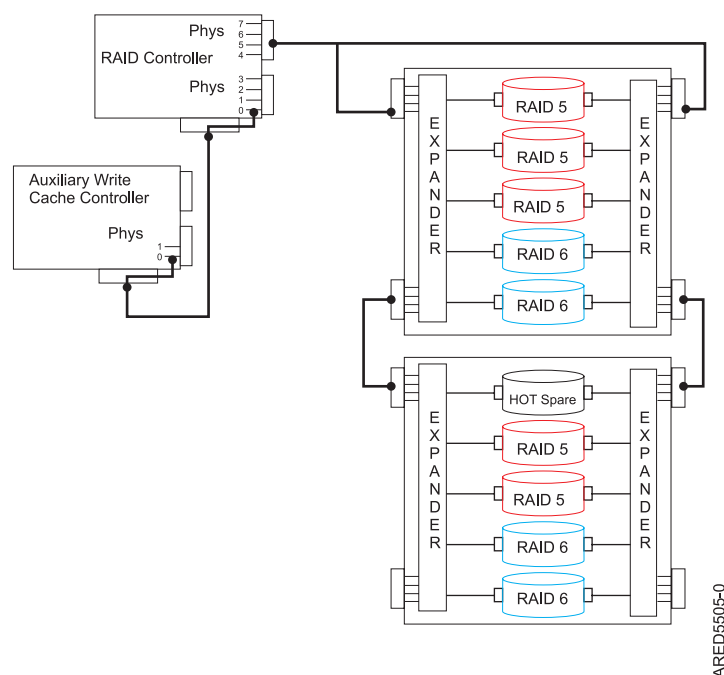


Figure 10. Example RAID and AWC controller configuration

## Viewing link status information

You can view more detailed link status information in the Hardware Service Manager.

1. Select **Start a service tool** on the Use System Service Tools (SST) menu.
2. Select **Hardware service manager** on Start a Service Tool menu.
3. Select **Logical hardware resources (buses, IOPs, controllers,...)**.
4. Select **System bus resources**.
5. Select the IBM Dual Storage IOA by placing a **9 (Resources associated with IOP)** in front of the desired adapter.
6. Enter **5 (Display detail)** in front of **Storage IOA** for the desired adapter. The screen displayed will look similar to the following example:

#### Auxiliary Storage Hardware Resource Detail

Description . . . . . : Storage IOA  
Type-model . . . . . : 572F-001  
Status . . . . . : Operational  
Serial number . . . . . : YL3028269C6B  
Part number . . . . . : 0000044V4193  
Resource name . . . . . : DC10  
Cache size (MB) . . . . . : 390  
PCI bus . . . . . :  
    System bus . . . . . : 517  
    System board . . . . . : 0  
    System card . . . . . : 0  
Storage . . . . . :  
    I/O adapter . . . . . : Not used  
    I/O bus . . . . . : 127  
    Controller . . . . . :  
    Device . . . . . :  
Operating mode . . . . . : Primary Storage IOA

More...

F3=Exit      F5=Refresh      F6=Print

F9=Change detail   F12=Cancel   F14=Dual Storage IOA Configuration

#### 7. Press **Page Down** to display **Link Status** on second page.

#### Auxiliary Storage Hardware Resource Detail

Remote storage IOA resource name . : DC01  
Remote storage IOA serial number . : YL3229021013  
Remote storage IOA link status . . : Operational  
Remote storage IOA type-model . . . : 572F-001  
Attached auxiliary IOA resource name: DC05  
Attached auxiliary IOA serial number: YL3229020FF9  
Attached auxiliary IOA link status : Operational  
Attached auxiliary IOA type-model . : 575C-001

Bottom

F3=Exit      F5=Refresh      F6=Print

F9=Change detail   F12=Cancel   F14=Dual Storage IOA Configuration



---

## Chapter 3. Controller software

For the adapter to be identified and configured by IBM i, the requisite software support must be installed. The requisite software for the adapter is often preinstalled during IBM i installation.

It might be necessary to perform operations related to the installation, verification, and maintenance of the IBM i software support for the adapter.

Software for the adapter is packaged in Program Temporary Fix (PTF) format and distributed as part of the base IBM i installation media, cumulative package media, and through the Web-based Fix Delivery Center for IBM i. This information is an overview of the IBM i software support that is required for the adapter. For complete information related to the installation and maintenance of IBM i, see the IBM System i® and IBM i Information Center Web site.

The adapter runs onboard microcode. Although a version of adapter microcode might be distributed along with the IBM i, this does not necessarily represent the most recent version of microcode that is available for the adapter. Newer PTFs might be available for the most current level of adapter microcode. Contact your technical support to verify the latest PTFs available for your specific adapters.

For the latest PTF group, HIPER (High Impact PERvasive) PTF, and cumulative PTF packages for your release, see Fix Central.

For the latest fixes and updates, go to the Support & downloads Web site, and search by entering your system type and controller type.

---

### Verifying the controller software

Verify the minimum software support that is required for your specific controller.

Support for the controller is contained in the Licensed Internal Code of IBM i.

Each controller requires a supported release of IBM i. Verify other possible code prerequisites that are described in the following table and at IBM Prerequisite.

**Attention:** Ensure that the adapters have the latest adapter microcode PTFs as part of the initial installation.

*Table 4. CCIN and version and release data*

CCIN (Custom card identification number)	Minimum required IBM i version and release
572A	IBM i V5R4M5 or later <sup>1</sup>
572C	IBM i V5R4M5 or later
572F and 575C	IBM i V5R4M5 or later <sup>1</sup>
574E	IBM i 6.1.1 or later <sup>1</sup>
57B3	IBM i 6.1 or later <sup>1</sup>
57B7	IBM i V5R4M5 or later
57B8	IBM i V5R4M5 or later
1. Refer to the PCI adapter information by feature type for the minimum IBM i level requirements.	

It might become necessary to install software updates so that you have the latest available level of adapter software support. Updates to the adapter software support are packaged, distributed, and installed through the same mechanisms that are used for other portions of the IBM i Licensed Internal Code. The standard IBM i technical support procedures can be used to determine the latest available level of adapter software support.

**Related concepts**

“Dual storage IOA functions” on page 22

Consider these factors when using dual storage I/O adapter (IOA) functions.

**Related information**

PCI adapter information by feature type

---

## Chapter 4. Common controller and disk array management tasks

You can perform various tasks to manage SAS RAID disk arrays.

Use the information in this section to manage your SAS RAID disk arrays.

- Device parity protection  
The topic describes the use of device parity protection on IBM i.
- Managing disk arrays  
Refer to this topic to see the interface for performing various tasks with disk arrays.
- Creating a disk array  
Use this procedure to start device parity protection.
- Using hot spare disks  
Hot spare disks are used to automatically replace a disk that has failed in a RAID environment.
- Disk unit management  
This procedure allows viewing disk status and disk unit details.

---

### Viewing IBM SAS disk information

This procedure enables you to view SAS disk information, status, and details.

To view the SAS disk information and status, see the following:

- IBM i Service Functions
- IBM i Dedicated Service Tools (DST) options

**Note:** The disk unit information options can also be accessed through System Service Tools. The system does not need to be in dedicated service mode to display disk information. Some disk configuration functions do require dedicated service mode.

- Work with disk units
- Display disk configuration

**Note:** This display shows disk unit details such as type, model, serial number, operating status, capacity, and protection status.

---

### Considerations for solid-state drives

Use this information to understand the importance of controller functions when you use solid-state drives (SSD).

Hard-disk drives (HDD) use a spinning magnetic platter to store nonvolatile data in magnetic fields. SSDs are a storage device using nonvolatile solid-state memory, typically 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 read/write 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 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 and HDD use

Follow these guidelines when using SSDs and HDDs.

- Do not mix SSDs and HDDs within the same disk array. A disk array must only contain SSDs or HDDs.
- Do not mix SSDs and HDDs with system mirroring in the same mirrored pair. A mirrored pair must only contain SSDs or HDDs.
- It is important to plan for hot-spare devices when you use arrays of SSDs. An SSD hot-spare device replaces a failed device in an SSD disk array. An HDD hot-spare device replaces a failed device for an HDD disk array.
- It is recommended that SSDs be protected by RAID 5, RAID 6, or by system mirroring.
- See Installing and configuring Solid State Drives to identify specific configuration and placement requirements related to the SSD devices.

### Related information



Installing and configuring Solid State Drives

---

## Chapter 5. Dual storage IOA configurations

You can increase availability using a dual storage I/O adapter (IOA) configuration to connect multiple controllers to a common set of disk expansion drawers and the included disks and disk arrays.

**Note:** Not all controllers support all configurations. See the PCI-X SAS RAID cards or the PCIe SAS RAID card comparison tables to look for controllers that have dual storage IOA configurations.

### Related concepts

“PCI-X SAS RAID card comparison” on page 3

Use the table in this topic to compare the features of PCI-X SAS RAID cards for IBM i. There are also images of adapters for you to view.

“PCIe SAS RAID card comparison” on page 6

Use the table in this topic to compare the features of a PCI Express (PCIe) SAS RAID cards for IBM i. There are also images of adapters for you to view.

---

## Possible disk storage IOA configurations

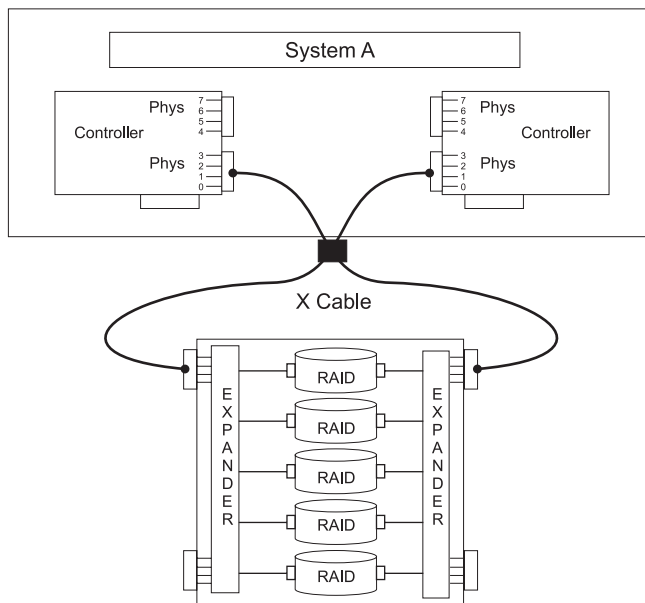
This topic shows a table illustrating what is needed to have dual storage IOA configurations with RAID or system mirroring and images of dual storage IOA configurations.

*Table 5. Disk protection with dual storage.* This table describes what is needed to have dual storage with different kinds of disk protection.

Multi-initiator configuration	Dual storage IOA
RAID	<ul style="list-style-type: none"><li>• Two controllers</li><li>• Both controllers must have the same write cache capability and write cache sizes</li><li>• Both controllers must support dual storage IOA configuration</li><li>• Controllers are in the same system or partition</li></ul>
System mirroring	<ul style="list-style-type: none"><li>• Four controllers (two pairs of controllers)</li><li>• Each pair of controllers must have the same write cache capability and write cache sizes</li><li>• Each pair of controllers must support dual storage IOA configuration</li><li>• Controllers are in the same system or partition</li></ul>

The following figure illustrates an example of a Dual storage IOA configuration with RAID.

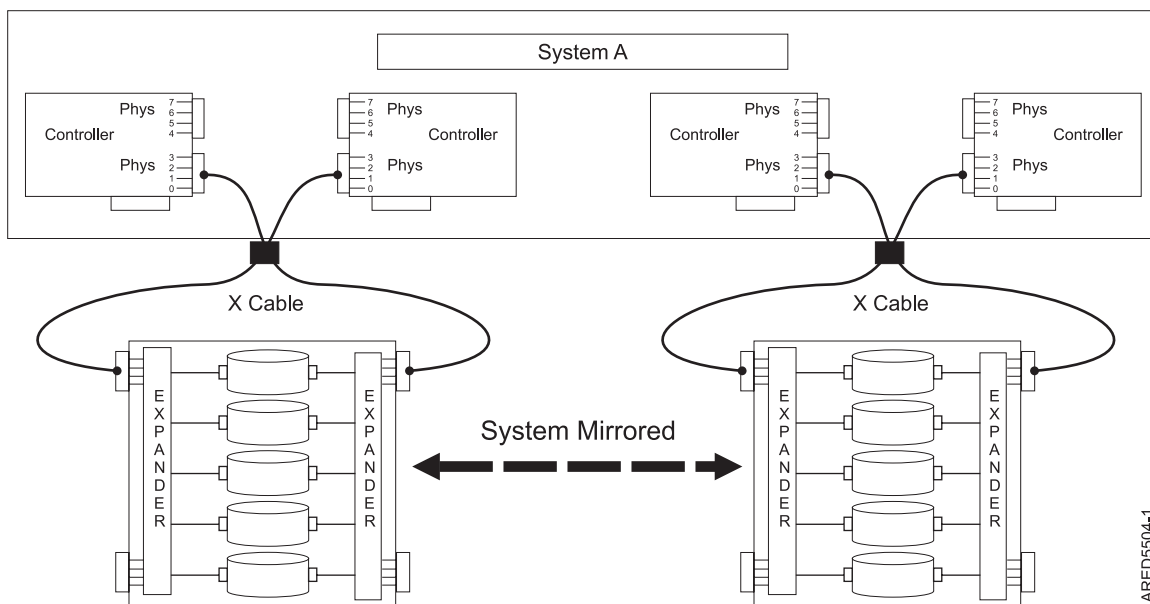
## Dual Storage IOA's with RAID



ARED5500-0

Figure 11. Dual storage IOA RAID configuration

## Dual Storage IOA's with System Mirroring



ARED5504-1

Figure 12. Dual storage IOA system mirroring configuration

## Dual storage IOA functions

Consider these factors when using dual storage I/O adapter (IOA) functions.

Use of the dual storage IOA function requires controller and IBM i software support. Controller support is shown in the feature comparison tables for PCIe and PCI-X cards. Look for controllers that have *Dual Storage IOA configuration* marked as Yes. The IBM i software levels that are required for multi-initiator support are identified in the Controller software verification topic.

Controllers connected in a dual storage IOA configuration must have the same write cache size (assuming that they support write cache). A configuration error is logged if the write caches for the controllers are not the same size.

When you configure a controller for a dual storage IOA configuration, no mode jumpers or special configuration settings are needed.

For all dual storage IOA configurations, one controller functions as the primary controller. Primary controllers perform management of the physical devices, such as creating a disk array. The other controller functions as the secondary controller and is not capable of physical device management.

If the secondary controller detects that the primary controller is going offline, it switches roles to become the primary controller. When the original primary controller comes back online, it becomes the secondary controller.

Both controllers are capable of performing direct I/O accesses (read and write operations) to the disk arrays. At any given time, only one controller in the pair is optimized for the disk array. The controller optimized for a disk array is the one that directly accesses the physical devices for I/O operations. The controller that is not optimized for a disk array forwards read and write requests, through the SAS fabric, to the optimized controller.

The primary controller logs most errors that are related to problems with a disk array. Disk array errors might also be logged on the secondary controller if a disk array is optimized on the secondary controller at the time the error occurred.

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

- Controllers 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 communicate with each other, compare device information, and switch roles.
- Powering off the primary controller causes an automatic transition (failover) to occur.
- Failure of the primary controller causes an automatic transition (failover) to occur.
- 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.

## Related concepts

“Verifying the controller software” on page 17

Verify the minimum software support that is required for your specific controller.

“PCI-X SAS RAID card comparison” on page 3

Use the table in this topic to compare the features of PCI-X SAS RAID cards for IBM i. There are also images of adapters for you to view.

“PCIe SAS RAID card comparison” on page 6

Use the table in this topic to compare the features of a PCI Express (PCIe) SAS RAID cards for IBM i. There are also images of adapters for you to view.

“Dual storage IOA access optimization” on page 26

View the active or passive path of your disk units and controller.

---

## Dual storage IOA function attributes

Find out which controller functions are supported with dual storage IOA configurations.

*Table 6. SAS controller functions.* This table describes controller functions that are supported with dual storage IOA configurations.

Controller functions	Dual storage IOA configurations
Disks formatted to 512 bytes per sector	No <sup>1</sup>
Disks formatted to 528 bytes per sector	Yes
Mirrored write cache between controllers that have write cache	Yes
Mirrored RAID parity footprints between controllers	Yes
Dual paths to disks	Yes
System-level mirroring	Yes
IBM-qualified disk drives	Yes
IBM-qualified disk expansion drawers	Yes
Tape or optical devices	No
Load source capable	Yes
Operating mode <sup>2</sup>	Primary or secondary adapter
1. Disks formatted to 512 bytes per sector are not to be used functionally, but these disks can be formatted to 528 bytes per sector.	
2. The operating mode can be viewed by using the Auxiliary Storage Hardware Resource Detail display.	

---

## Viewing dual storage IOA attributes

This topic collection provides the details for using the Auxiliary Storage Hardware Resource Detail display to obtain dual storage I/O adapter (IOA) configuration information.

Perform the following steps to view details about your adapters.

1. Select **Start a service tool** on the Use System Service Tools (SST) menu.
2. Select **Hardware service manager** on the Start a Service Tool menu.
3. Select **Logical hardware resources (buses, IOPs, controllers)** on the Hardware Service Manager menu.
4. Select **System bus resources on the Logical Hardware Resources** on the System Bus menu.
5. Type 9 (Resources associated with IOP) in front of the adapter that you want.
6. Type 5 (Display detail) in front of **Storage IOA** to get details about the storage IOA. This is an example of the display:



#### Auxiliary Storage Hardware Resource Detail

```

Description . . . . . : Storage IOA
Type-model . . . . . : 574E-001
Status . . . . . : Operational
Serial number . . . . . : YL3028269C6B
Part number . . . . . : 0000044V4198
Resource name . . . . . : DC10
Cache size (MB) . . . . . : 380
PCI bus . . . . . :
  System bus . . . . . : 517
  System board . . . . . : 0
  System card . . . . . : 0
Storage . . . . . :
  I/O adapter . . . . . : Not used
  I/O bus . . . . . : 127
  Controller . . . . . :
  Device . . . . . :
Operating mode . . . . . : Primary Storage IOA

```

F3=Exit      F5=Refresh      F6=Print  
 F9=Change detail    F12=Cancel    F14=Dual Storage IOA Configuration

- Press F14 (Dual Storage IOA Configuration) to view a list of both adapters in the dual storage IOA pair. This is an example of the display:

#### Dual Storage IOA Configuration

Type options, press Enter.  
 2=Change detail    5=Display detail    6=I/O debug  
 8=Associated packaging resource(s) 9=Resources associated with controlling IOP

Resource Opt Name	Type- Model	Status	Serial Number	Operating Mode
DC10	574E-001	Operational	YL3028269C6B	Primary Storage IOA
DC09	574E-001	Operational	YL3028270DA0	Secondary Storage IOA

F3=Exit      F5=Refresh      F6=Print      F12=Cancel

- To see details about each individual adapter, type 5 (Display detail) in front of the adapter that you want.

For additional details on how to set up a configuration, see [Installing dual storage IOA configurations](#).

#### Related concepts

“Installing dual storage IOA configurations” on page 28

Use this procedure to help you to install a dual storage IOA configuration.

## SAS cabling considerations

Cabling your system correctly is one of the most important aspects of planning for a dual storage I/O adapter (IOA) configuration.

Follow these guidelines when you cable your system.

- For RAID configurations on a 5886 EXP 12S disk expansion drawer, X cables provide redundancy for two wide SAS ports between each controller and disk expansion drawer, and it also provides redundancy for two narrow SAS ports for each disk drive.
- For RAID configurations with a 5802 or 5803 PCIe 12X I/O drawer, AT cables are used. SAS topology is incorporated with in the IO drawer wiring. This provides redundancy similar to X cables.

- For RAID configurations with internal SAS disk slots, YR cables provide redundancy for two narrow SAS ports between each controller and internal disk enclosure, and it also provides redundancy for two narrow SAS ports for each disk drive.

To see examples of how to cable dual storage IOA configurations, see Serial attached SCSI cable planning. Look for the following sections for examples of how to cable dual storage IOA configurations:

- Two SAS adapters to disk expansion drawer – multi-initiator HA RAID configuration
- PCIe SAS adapter in PCIe 12X I/O drawer to the internal SAS disk slots
- Two SAS adapters to internal SAS disk slots in models 9117-MMA and 9406-MMA

#### **Related concepts**

“Installing dual storage IOA configurations” on page 28

Use this procedure to help you to install a dual storage IOA configuration.

#### **Related information**

 Serial attached SCSI cable planning

---

## **Performance considerations**

Controller failures can affect performance.

The controller is designed to minimize performance impacts when running in a dual storage IOA configuration. When using RAID 5 and RAID 6, 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 a dual storage IOA configuration, the remaining controller disables write caching (if auxiliary cache is not also provided by the controllers) and begins to keep an additional copy of parity footprints on disk. This can significantly affect performance, particularly when using RAID 5 and RAID 6.

---

## **Dual storage IOA access optimization**

View the active or passive path of your disk units and controller.

Dual storage IOA access characteristics can balance the controller workload. The dual storage IOA access characteristics for a disk array, parity set, specifies which controller is preferred to be optimized for the disk array. It performs direct read and write operations to the physical devices. The controller that is preferred to be optimized for the disk array, contains the active path to the disk units in the disk array. The other controller contains the passive path. The system only sends read and write operations down the active path. The passive path is only used if the active path fails.


Best performance is achieved when the dual storage IOA access characteristics on each disk array have a balanced workload. This happens when the two controllers have an equal number of disk arrays with active paths to the disk units.

The system selects the disk units and dual storage IOA access characteristics for each disk array. When creating disk arrays, set the parity set optimization to Performance. This attribute sets an even number of disk arrays (for example, 2, 4, 6, and so forth) to be created. It also enables the system to optimize disk arrays on each controller. As a result, the two controllers will have an equal number of disk units with an active path.

To change the parity set optimization, see Changing parity set optimization.

## Dual Storage IOA Optimization

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



ARE5503-1

Figure 13. Dual storage IOA optimization. This figure shows RAID arrays with primary and secondary adapters.

## Viewing the active or passive path of disk units

To view the active or passive path of the disk units, complete the following steps.

1. Select **Work with disk units** on the Use System Service Tools (SST) menu.
2. Select **Display disk configuration** from the Work with Disk Configuration display.
3. Select **Display path status** on the Display Disk Configuration display.

### Display Disk Path Status

ASP Unit	Serial Number	Type	Model	Resource Name	Path Status
* *	Y6800024F78E	433C	099	DMP001	Active
				DMP002	Passive
* *	Y680002AEB3D	433C	099	DMP003	Active
				DMP004	Passive
* *	Y6800024F754	433C	099	DMP005	Active
				DMP006	Passive
* *	Y6800024F771	433C	099	DMP007	Active
				DMP008	Passive
* *	Y68000268517	433C	099	DMP009	Active
				DMP010	Passive
* *	Y680002B31DD	433C	099	DMP011	Active
				DMP012	Passive
* *	Y6800024F74D	433C	099	DMP013	Active
				DMP014	Passive

More...

Press Enter to continue.

F3=Exit      F5=Refresh      F9=Display disk unit details  
F11=Display encryption status      F12=Cancel

## Viewing the active or passive path role for a controller

To view the active or passive path role for a controller, complete the following steps.

1. Select **Start a service tool** on the Use System Service Tools (SST) menu.
2. Select **Hardware service manager** on the Start a Service Tool menu.
3. Select **Logical hardware resources (buses, IOPs, controllers)** on the Hardware Service Manager menu.
4. Select **System bus resources** on the Logical Hardware Resources on the System Bus menu.
5. Select the **Virtual IOP** by typing a 9 (Resources associated with IOP) in front of the desired IBM dual storage IOA.
6. Press F11 (function key) until Path Role is shown.

Logical Hardware Resources Associated with IOP				
Type options, press Enter.				
2=Change detail    4=Remove    5=Display detail    6=I/O debug				
7=Verify    8=Associated packaging resource(s)				
Opt	Description	Type-Model	Path Role	Resource Name
	Virtual IOP	* 572F-001		CMB01
	Storage IOA	572F-001		DC02
	Disk Unit	433B-099	Active	DMP002
	Disk Unit	433B-099	Passive	DMP004
	Disk Unit	433B-099	Active	DMP006
	Disk Unit	433B-099	Passive	DMP008
	Disk Unit	433B-099	Active	DMP010
	Disk Unit	433C-099	Active	DMP012
	Disk Unit	433C-099	Active	DMP014
	Disk Unit	433C-099	Passive	DMP016
	Disk Unit	433B-099	Passive	DMP018
				More...
F3=Exit    F5=Refresh    F6=Print    F8=Include non-reporting resources				
F9=Failed resources    F10=Non-reporting resources				
F11=Display status/resource name    F12=Cancel				

### Related concepts

“Dual storage IOA functions” on page 22

Consider these factors when using dual storage I/O adapter (IOA) functions.

“Installing dual storage IOA configurations”

Use this procedure to help you to install a dual storage IOA configuration.

### Related information

 Changing parity set optimization

## Installing dual storage IOA configurations

Use this procedure to help you to install a dual storage IOA configuration.

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

**Attention:** Disk Arrays can be created either before or after the Dual Storage IOA configuration is set up.

Each controller requires a supported release of IBM i and verify other possible code pre-requisites. Refer to the Feature Pre-requisite web site.

**Attention:** Ensure the adapters are updated with the latest adapter microcode PTF's as part of the initial installation.

Perform the following steps to install a dual storage IOA configuration

1. Verify that all pre-requisites are permanently applied.
2. Install the SAS controllers into the system or partition.

**Note:** Do not attach any cables to the SAS controllers.

3. To prevent errors while connecting the cables, perform a normal shutdown of the system or partition before you attach the cables.
4. Attach the necessary SAS cables from the shared disk enclosure to the same SAS connector on each controller. To see examples of how to cable the Dual storage IOA configurations, see SAS cabling considerations.
5. Power on your system or partition.
6. Verify that the cabling and functioning of the controllers are correct by using the Dual storage IOA configuration screen, see Viewing dual storage IOA attributes.
7. Best performance is achieved when dual storage IOA access characteristics for each disk array is such that the workload is balanced between the two controllers. Refer to Dual storage IOA access optimization and create or change the RAID configuration as necessary.

For additional details on how to set up a configuration, see Installing dual storage IOA configuration.

#### **Related concepts**

“SAS cabling considerations” on page 25

Cabling your system correctly is one of the most important aspects of planning for a dual storage I/O adapter (IOA) configuration.

“Dual storage IOA access optimization” on page 26

View the active or passive path of your disk units and controller.

“Installing dual storage IOA configurations” on page 28

Use this procedure to help you to install a dual storage IOA configuration.

#### **Related tasks**

“Viewing dual storage IOA attributes” on page 24

This topic collection provides the details for using the Auxiliary Storage Hardware Resource Detail display to obtain dual storage I/O adapter (IOA) configuration information.



---

## Chapter 6. SAS RAID controller maintenance

Ensure optimal performance of your controller by using these maintenance procedures.

To help avoid controller and disk array problems, use the following tips:

- Perform a normal system shutdown before physically replacing or moving the RAID controller or members of disk arrays. A normal shutdown of the system will flush the controller's write cache and remove dependencies between the controller and the disks. Powering off the PCI slot using concurrent maintenance options in Hardware Service Manager (HSM) has the same effect as it would on a single controller when the PWRDWN SYS command is used.

**Note:** Disks that are a Failed member of an Unprotected (exposed) disk array can be replaced and the disk unit data rebuilt while the system continues to run. No system shutdown is required.

- You can physically move disks from one controller to another. However, if the disks are members of a disk array, be sure to move all the disks in the array as a group. Prior to attempting a disk movement, ensure that the disk array is not in an Unprotected state because of a disk failure. The system/partition must be powered off normally before disks are moved.
- When physically removing disks that are members of a disk array, remove the disks from the Auxiliary Storage Pool (ASP) and then stop RAID on the disk array before removing the disks. This action avoids loss of data and disk-array-related problems the next time that these disks are used. The system/partition must be powered off normally before disks are physically removed.
- Always use the Device Concurrent Maintenance option to remove and replace a disk.
- If the load source disk is part of a disk array and the system fails to IPL because of a suspected disk array problem, IPL the system/partition using D-IPL media (CD/DVD or SAVESYS media). Error Log Analysis, and other tools are available on the Dedicated Service Tools menu to help determine and resolve the problem with the disk array.
- Do not attempt to correct problems by swapping controllers and disks unless you are directed to do so by the service procedures. Use Error Log Analysis to determine what actions to perform, and when appropriate, follow the appropriate Isolation Procedures for problem determination. If multiple errors occur at approximately the same time, look at them as a whole to determine if there is a common cause.
- Do not confuse the cache directory card, which is a small rectangular card with round, button-shaped batteries, for a removable cache card. The nonvolatile write cache memory is integrated into the controller. The write cache memory itself is battery-backed by the large, rechargeable cache battery pack. The cache directory card contains only a secondary copy of the write cache directory and no cache data. Do not remove this card except under very specific recovery cases as described in the Isolation Procedures.
- Do not unplug or exchange a cache battery pack without following the procedures as outlined in this section or in the Isolation Procedures. Failure to follow these procedures might result in data loss.

---

### Rechargeable battery maintenance

Rechargeable battery maintenance tasks include displaying rechargeable battery information, forcing a rechargeable battery error, and replacing the rechargeable cache battery pack.

### Displaying rechargeable battery information

Use this procedure to display information about the controller's rechargeable battery using the Hardware Service Manager in the IBM® i operating system.

1. Sign on to the system with at least service level authority.
2. On the command line, type `strsst`. Press Enter.

3. On the System Service Tools (SST) Sign On display, type your service tools user ID and service tools password. Press Enter.
4. On the System Service Tools (SST) display, select Start a Service Tool. Press Enter.
  - a. On the Start a Service Tool display, select Hardware Service Manager. Press Enter.
  - b. On the hardware Service Manager display, select Work with resources containing cache battery packs. Press Enter.
  - c. Select Display battery information.

#### Battery Information

```

Resource name . . . . . : DC01
Serial number . . . . . : YL3229021013
Actual type-model . . . . . : 572F-001
Unit ID . . . . . : U5094.001.10XS187
Planar ID . . . . . : CB1
Card . . . . . : C01
Battery type . . . . . : Lithium Ion (LiIon)
Battery state . . . . . : No battery warning/error
Power-on time (days) . . . . . : 215
Adjusted power-on time (days) . . . : 236
Estimated time to warning (days) . : 673
Estimated time to error (days) . . : 756
Concurrently maintainable      :
    battery pack . . . . . : Yes
Battery pack can be safely replaced : No
  
```

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

#### **Resource name**

The resource name of the selected controller.

#### **Serial number**

Serial number of the selected controller

#### **Actual type-model**

CCIN of the selected controller

#### **Unit ID**

Enclosure feature identifier containing the selected controller

#### **Planar ID**

Planar identifier containing the selected controller

**Card** Physical card slot identifier containing the selected controller

#### **Battery Type**

The type of rechargeable cache battery pack.

#### **Battery State**

Indicates if an error condition currently exists related to the rechargeable cache battery pack. The



possible values for this field are:

*Table 7. Battery state*

No battery warning/error	Warning condition	Error condition	Unknown
No warning or error condition currently exists.	A warning condition currently exists and an error has been logged.	An error condition currently exists and an error has been logged.	Information is not available to determine whether a warning or error condition currently exists.

**Power-on time (days)**

Indicates the raw power-on time, in units of days, of the rechargeable cache battery pack.

**Adjusted power-on time (days)**

Indicates the adjusted (prorated) power-on time, in units of days, of the rechargeable cache battery pack.

**Note:** Some rechargeable cache battery packs are negatively affected by higher temperatures and thus are prorated based on the amount of time that they spend at various ambient temperatures.

**Estimated time to warning (days)**

Estimated time, in units of days, until a message is issued indicating that the replacement of the rechargeable cache battery pack should be scheduled.

**Estimated time to error (days)**

Estimated time, in units of days, until an error is reported indicating that the rechargeable cache battery pack must be replaced.

**Concurrently maintainable battery pack**

Indicates if the rechargeable cache battery pack can be replaced while the controller continues to operate.

**Battery pack can be safely replaced**

Indicates if the controller's write cache has been disabled and the rechargeable cache battery pack can be safely replaced.

## Error state

The cache battery pack should be in an error state before you replace it.

To prevent possible data loss, ensure that the cache battery pack is in an error state before replacing it. This will ensure all cache data is written to disk before battery replacement. Forcing the battery error will result in the following:

- The system logs an error.
- Data caching becomes disabled on the selected controller.
- System performance could become significantly degraded until the cache battery pack is replaced and charging of the new battery pack has completed. The new battery pack may take several hours to charge.
- The battery pack can be safely replaced field on the controller rechargeable battery information screen indicates Yes.
- Cache data present LED stops flashing. See the feature descriptions and the figures in the Replacing a battery pack section to determine if your adapter has a cache data present LED and the location of the LED.

This error state requires replacement of the cache battery. Ensure that you have the correct type and quantity of cache battery packs to do the replacement. To resume normal operations, replace the cache battery pack.

The cache battery pack for the 572F storage I/O adapter and the 575C auxiliary cache adapter is contained in a single battery field replacement unit (FRU) that is physically located on the 575C auxiliary cache adapter. The functions of forcing a battery pack error and starting the adapter cache on either adapter in the card set results in the same function automatically being performed on the other adapter in the card set.

## Forcing a rechargeable battery error

Use this procedure to place the controller's rechargeable battery into an error state using the Hardware Service Manager in the IBM i operating system.

This topic applies to the IBM i operating system. For information about maintaining the rechargeable battery using the AIX® or Linux® operating systems, see SAS RAID controller for AIX or SAS RAID controller for Linux.

To force the cache battery pack into an error state, do the following steps on the system or partition that is using the adapter.

1. Sign on to the system with at least service level authority.
2. On the command line, type `strsst`. Press Enter.
3. On the System Service Tools (SST) Sign On display, type your service tools user ID and service tools password. Press Enter.
4. On the System Service Tools (SST) display, select **Start a Service Tool**. Press Enter.
  - a. On the Start a Service Tool display, select **Hardware Service Manager**. Press Enter.
  - b. On the hardware Service Manager display, select **Work with resources containing cache battery packs**. Press Enter.
  - c. On the Work with Resources containing Cache Battery Packs display, select **Force battery pack into error state for the I/O card**. Press Enter.
  - d. On the **Force Battery Packs Into Error State** display, verify that the correct I/O adapter has been selected, and press the function key that confirms your choice.
  - e. Return to the **Work with Resources containing Cache Battery Packs** display and select **Display battery information** and verify that the **Battery pack can be safely replaced** field indicates yes. If it does not indicate yes, contact your next level of support before continuing this procedure.
5. Verify that the cache data present light emitting diode (LED) is no longer flashing before replacing the cache battery pack. See the feature descriptions and the figures in the "Replacing a battery pack" section to determine if your adapter has a cache data present LED and the location of the LED.
6. Replace the cache battery pack using the procedure that sent you here. For a list of replacement procedures, see "Replacing a battery pack."
7. Restart the adapter's write cache by doing the following:
  - a. Return to the **Work with Resources containing Cache Battery Packs display** and select the **Start IOA cache**. Press Enter.
  - b. Ensure that you get the message Cache was started.

---

## Replacing a battery pack

Follow these guidelines before replacing your battery pack.

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

**Note:** 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:** To prevent data loss, if the cache battery pack is not already in the error state, follow the steps described in “Forcing a rechargeable battery error” on page 34 before proceeding. If the cache data present LED is flashing, do not replace the cache battery pack or data will be lost. See the feature descriptions and the figures in the following sections to determine if your adapter has a cache data present LED and the location of the LED.

**Attention:** Static electricity can damage this device and your system unit. To avoid damage, keep this device in its antistatic protective bag until you are ready to install it. To reduce the possibility of electrostatic discharge, read the following precautions:

- 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 antistatic package, touch it to an unpainted metal part of the system unit for at least 2 seconds. (This duration 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 572B nonconcurrent maintainable battery pack

Use this procedure to replace the nonconcurrent maintainable battery pack on adapter type CCIN 572B.

**Attention:** Before continuing with this procedure, determine that it is safe to replace the cache battery pack. Refer to Displaying rechargeable battery information. It is safe to replace the cache battery pack when **Yes** is displayed next to **Battery pack can be safely replaced**.

Complete the following steps to replace a nonconcurrent maintainable battery pack.

1. Remove the controller from the system. See your system documentation for instructions.
2. Place the controller on a surface that is electrostatic-discharge protected.
3. Unplug the battery connector (**B**) from its connector on the adapter, squeezing the retaining latch while gently pulling on the plug. The plug connects to the board in only one way, so it cannot be inserted incorrectly during the replacement procedure.

**Note:** Ensure that the cache battery pack is disconnected for at least 60 seconds before connecting the new battery. This duration is the minimum amount of time needed for the adapter to recognize that the battery has been replaced.




Figure 14. Removing the cache battery

(A) Cache battery pack

(B) Battery connector

(C) Plastic rivet

(D) Plastic pin

4. Locate the two plastic rivets (C) that hold the cache battery pack in place. From the back of the adapter, remove the two pins (D) that are inserted inside of the rivets.
5. Release the rivets (C) that secure the battery assembly to the adapter. Press the rivets through the back of the adapter and remove the battery pack (A) from the adapter. If the rivets (C) cannot be pressed through the back of the adapter, follow these steps to press out the rivets with a ballpoint pen:
  - a. Locate a retractable ballpoint pen.

**Note:** A medium-sized retractable ballpoint pen is preferred, or an equivalent item with a small opening can be used. The small opening must be large enough so that the pen (or equivalent) can go around the tip of the rivet, but small enough that it does not slide over the rivet and contact the battery assembly bracket.

  - b. Slide the card off the edge of the work area enough so rivet (C) can be pressed out of the back of the adapter.
  - c. Hold the pen with the ballpoint retracted, place the pen on top of rivet (C), and gently press straight down until rivet (C) presses out.
  - d. Repeat steps 5b and 5c for the other rivet (C).
  - e. Remove the cache battery pack (A) from the adapter.
  - f. Turn the adapter over and press the rivets (C) back into the adapter.
6. Install the new battery pack. (A) onto the press rivets (C) of the adapter.
7. Reinsert the pins (D) into the rivets from the back of the adapter.
8. Connect the cache battery pack connector (B) to the adapter. The plug connects to the adapter in only one way, so it cannot be inserted incorrectly.

9. Reinstall the adapter.


## Replacing a 57B7 concurrent maintainable battery pack

Use this procedure to replace the concurrent maintainable battery pack on adapter type CCIN 57B7.

**Attention:** Before continuing with this procedure, determine that it is safe to replace the cache battery pack. Refer to Displaying rechargeable battery information. It is safe to replace the cache battery pack when **Yes** is displayed next to **Battery pack can be safely replaced**. If the cache data present LED is flashing, do not replace the cache battery pack or data will be lost.

Complete the following steps to replace a 57B7 concurrent maintainable battery pack.

1. Using the following illustration to locate the battery components, verify that the cache data present LED (C) is not flashing. If it is flashing, do not continue; return to "Forcing a rechargeable battery error" on page 34.



AREBJ518-0


- (A) Cache battery lever
- (B) Cache battery pack
- (C) Cache present LED

Figure 15. Removing the 57B7 cache battery

2. Move the cache battery lever (A) away from the connector to disengage the battery from the connector.
3. 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 duration is the minimum amount of time needed for the card to recognize that the battery has been replaced.

4. Using the following illustration to locate the battery components, move the lever to the unlatched position (away from the connector).



AREBJ519-0

Figure 16. Replacing the 57B7 cache battery

- (A) Cache battery lever
- (B) Cache battery pack
- (C) Cache battery connector

5. Slide the new cache battery pack into the mounting guides on the controller until it is seated in the battery connector.
6. After the battery is seated in the connector, move the lever to the latched position to fully seat the battery into the connector.
7. Restart the adapter's write cache by doing the following:
  - a. Return to the **Work with Resources containing Cache Battery Packs display** and select the **Start IOA cache**. Press Enter.
  - b. Ensure that you get the message Cache was started.

## Replacing a 574E concurrent maintainable battery pack

Use this procedure to replace the concurrent maintainable battery pack on adapter type CCIN 574E.

**Attention:** Before continuing with this procedure, determine that it is safe to replace the cache battery pack. See *Displaying rechargeable battery information*. It is safe to replace the cache battery pack when **Yes** is displayed next to **Battery pack can be safely replaced**. If the cache data present LED is flashing, do not replace the cache battery pack or data will be lost.

Complete the following steps to replace a 574E concurrent maintainable battery pack.

1. Using the following illustration to locate the battery components, verify that the cache data present LED (C) is not flashing. If it is flashing, do not continue; return to "Forcing a rechargeable battery

error” on page 34.




Figure 17. Replacing the 574E cache battery

- (B) Cache battery pack
- (C) Cache data present LED
- (D) Cache battery tab
- (E) Cache battery tab

2. Squeeze tab (D) against tab (E) to disengage the battery retaining tab, pull out the cache battery pack (B), 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 duration is the minimum amount of time needed for the card to recognize that the battery has been replaced.

3. Install the new cache battery pack by reversing this procedure. Ensure that the replacement cache battery pack is fully seated.
4. Restart the adapter's write cache by doing the following:
  - a. Return to the **Work with Resources containing Cache Battery Packs display** and select the **Start IOA cache**. Press Enter.
  - b. Ensure that you get the message Cache was started.

## Replacing a 572F/575C card set concurrent maintainable battery pack

Use this procedure to replace the concurrent maintainable battery pack on adapter type CCIN 572F/575C card set.

**Attention:** Before continuing with this procedure, determine that it is safe to replace the cache battery pack. See Displaying rechargeable battery information. It is safe to replace the cache battery pack when **Yes** is displayed next to **Battery pack can be safely replaced**.

Complete the following steps to replace a 572F/575C concurrent maintainable battery pack.

1. Using the following illustration to locate the battery's components, locate the metal cover (A) that holds the battery pack. Pull out on the push-rivet (B) to release the metal cover (A).




Figure 18. Replacing the 572F/575C cache battery

- (A) Metal cover
- (B) Push-rivet
- (C) Tab

2. Remove the battery unit by pulling on tab (C).

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

3. Install the new cache battery pack by reversing this procedure. Ensure that the replacement cache battery pack is fully seated.
4. Restart the adapter's write cache by doing the following:
  - a. Return to the **Work with Resources containing Cache Battery Packs display** and select the **Start IOA cache**. Press Enter.
  - b. Ensure that you get the message Cache was started.

## Separating the 572F/575C card set and moving the cache directory card

When the maintenance procedures direct you to separate the 572F/575C card set and move the cache directory card on a 572F controller for recovery purposes, carefully follow this procedure.

**Important:** To avoid loss of cache data, do not remove the cache battery during this procedure.

### Notes:

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



- If you are removing the adapter from a double-wide cassette, go to the procedures in your system unit's service information for removing a double-wide adapter from a double-wide cassette.

**Attention:** All cards are sensitive to electrostatic discharge. See Handling static-sensitive devices before beginning this procedure.

To separate the 572F/575C card set and move the cache directory card, complete the following steps.

1. Label both sides of the card before separating them.
2. Place the 572F/575C card set adapter on an ESD protective surface and orient it as shown in Figure 19.




Figure 19. 572F/575C card set adapter

3. To prevent possible card damage, loosen all five retaining screws **C** before removing any of them. After all five retaining screws have been loosened, remove the screws **C** from the 572F storage adapter.

**Important:** Failure to loosen all five retaining screws prior to removing any of the screws can result in damage to the card.




Figure 20. Location of screws on the 572F/575C card set adapter

**C** Screws

4. Grasp the 572F and 575C adapters close to the interconnect connector **A**, as shown in the following figure, and carefully pull the connector apart; then, set the adapters on the ESD protective surface.




Figure 21. Location of interconnect connector on the 572F/575C card set adapter

**A** Interconnect connector

5. Turn the 572F storage adapter over so the components are facing up. Locate the cache directory card **D** on the 572F storage adapter. The cache directory card is the small rectangular card mounted on

the I/O card.




Figure 22. Cache directory card

**D** Cache directory card

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




Figure 23. Unseating the connector

7. Move the cache directory card to the replacement 572F storage adapter and seat it on the connector and mounting pegs.
8. To reassemble the cards, perform the preceding procedure in reverse order. When connecting the two adapters together, carefully align guide pins **B** on each side of the interconnect connector **A**. After the connector is seated correctly, apply pressure to completely squeeze the connector together. To prevent possible card damage, insert all five screws **C** before tightening any of them.




Figure 24. Reassembling the cards

- A** Interconnect connector
- B** Guide pins
- C** Screws

9. Cassette installations only: If you are installing the 572F/575C card set adapter into a cassette, perform the following steps:
  - a. Remove the adapter handle **B** as shown in Figure 25 on page 45.




Figure 25. Cassette adapter handle attachment

**A** Push-rivets

**B** Adapter handle

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

10. Return to the procedure that sent you here. **This ends this procedure.**

## Viewing SAS fabric path information

Use the Hardware Service Manager to view details of the SAS fabric information.

Perform the following steps to view the SAS fabric information:

1. Start the Hardware Service manager and page forward to the second selection screen.
2. Select **SAS Resource Path Information**.
3. Enter **1 (Display resource path information)** in front of the desired adapter resource. The screen displayed will look similar to the following example:

### SAS Resource Path Information

Type option, press Enter.

1=Display SAS fabric path graphical view

Opt	Resource Name	Path 1 Status	Path 2 Status
	DMP001	Operational	Operational
	DMP003	Operational	Operational
	DMP015	Operational	Operational
	DMP005	Operational	Operational
	DMP009	Operational	Operational
	DMP017	Operational	Operational
	DMP019	Operational	Operational
	DMP021	Operational	Operational
	DMP007	Operational	Operational
	DMP011	Operational	Operational
	DMP013	Operational	Operational
	DMP024	Operational	Operational
	DMP026	Operational	Operational
	DMP027	Operational	Operational
	DMP029	Operational	Operational

More...

F3=Exit

F5=Refresh

F6=Print

F12=Cancel

Selecting a device will display the details of all the nodes on each path between the controller and the device. Following is an example for Display SAS fabric path graphical view.

# SAS Fabric Path Graphical View

		Adapter DC03			
Path Active	: Yes			Path Active	: Yes
Path State	: Operational			Path State	: Operational
+-----+-----+					
SAS Address	: 5005076C07377C01			SAS Address	: 5005076C07377C01
Port	: 0			Port	: 2
Status	: Operational			Status	: Operational
Info	: 3.0 GBPS			Info	: 3.0 GBPS
+-----++-----+					
+-----++-----+					
Expander : 1				Expander : 1	
+-----++-----+					
SAS Address	: 500A0B8370F9D000			SAS Address	: 500A0B82FC269000
Port	: 22			Port	: 22
Status	: Operational			Status	: Operational
+-----++-----+					
SAS Address	: 500A0B8370F9D000			SAS Address	: 500A0B82FC269000
Port	: 5			Port	: 5
Status	: Operational			Status	: Operational
Info	: 3.0 GBPS			Info	: 3.0 GBPS
+-----++-----+					
+-----++-----+					
Device				Device	
+-----++-----+					
SAS Address	: 5000CCA00357B5CF			SAS Address	: 5000CCA00397B5CF
Port	: 0			Port	: 1
Status	: Operational			Status	: Operational
Info	: 3.0 GBPS			Info	: 3.0 GBPS
+-----++-----+					
SAS Address	: 5000CCA00317B5CF			SAS Address	: 5000CCA00317B5CF
Status	: Operational			Status	: Operational
+-----++-----+					
Device Lun					
DMP003					
+-----++-----+					

Bottom

F3=Exit

F5=Refresh

F6=Print

F12=Cancel

F11=SAS Fabric Path Data View

An alternative view can be displayed by selecting F11 (SAS Fabric Path Data View). Following is an example of SAS Fabric Path Data View.

SAS Fabric Path Data View

Adapter	Adapter Port	Path Active	Path State	Device
DC03	0	Yes	Operational	DMP003
DC03	2	Yes	Operational	DMP003

Node	SAS Address	Port Type	Port	Status	Info
1	5005076C07377C01	Adapter	0	Operational	3.0 GBPS
2	500A0B8370F9D000	Expander	22	Operational	3.0 GBPS
3	500A0B8370F9D000	Expander	5	Operational	3.0 GBPS
4	5000CCA00357B5CF	Device	0	Operational	3.0 GBPS
5	5000CCA00317B5CF	Device Lun	5	Operational	Status 0
1	5005076C07377C01	Adapter	2	Operational	3.0 GBPS
2	500A0B82FC269000	Expander	22	Operational	3.0 GBPS
3	500A0B82FC269000	Expander	5	Operational	3.0 GBPS
4	5000CCA00397B5CF	Device	1	Operational	3.0 GBPS
5	5000CCA00317B5CF	Device Lun	5	Operational	Status 0

F3=Exit

F5=Refresh

F6=Print

F12=Cancel

The possible status values for the SAS Fabric Path Data View and the SAS Fabric Path Graphical View follow.

Table 8.

Status	Description
Operational	No problem detected
Degraded	The SAS node is degraded
Failed	The SAS node is failed
Suspect	The SAS node is suspect of contributing to a failure



Table 8. (continued)

Status	Description
Missing	The SAS node is no longer detected by controller
Not valid	The SAS node is incorrectly connected
Unknown	Unknown or unexpected status

## Example: Using SAS fabric path information

This data becomes helpful in determining the cause of configuration or SAS fabric problems.

The following example assumes a cascaded disk enclosure with a broken connection on one path between the cascaded enclosures.

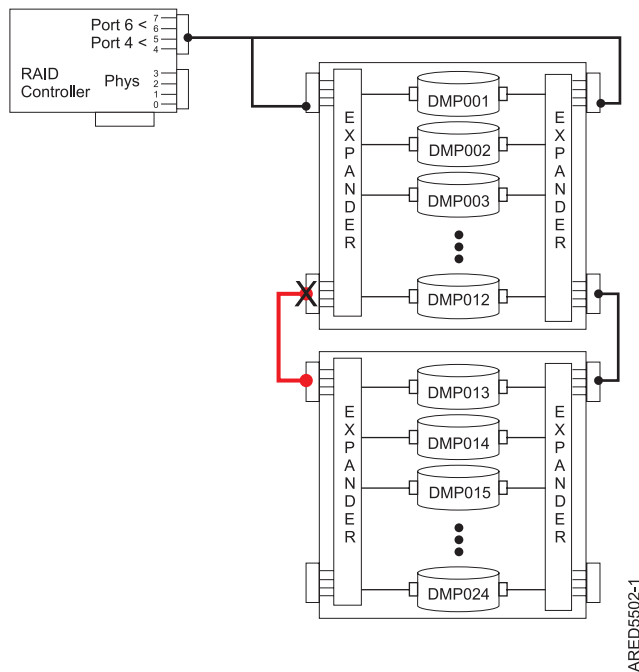


Figure 26. Cascaded disk enclosure

The state of all paths to all devices displays information similar to the following.

#### SAS Resource Path Information

Type option, press Enter.

1=Display SAS fabric path graphical view

Opt	Resource Name	Path 1 Status	Path 2 Status
	DMP001	Operational	Operational
	DMP002	Operational	Operational
	DMP003	Operational	Operational
	DMP012	Operational	Operational
	DMP013	Failed	Operational
	DMP014	Failed	Operational
	DMP015	Failed	Operational
	DMP024	Failed	Operational
	D01	Operational	
	D02	Operational	
	D03	Failed	
	D04	Operational	

For Display SAS fabric path graphical view, choosing one of the devices with a Failed path will display information similar to the following.

Chapter 6. SAS RAID controller maintenance **51**

		Info	: 3.0 GBPS	
		+-----+		
		+-----+		
			Device	
		+-----+		
		SAS Address	: 5000CCA003A63941	
		Port	: 1	
		Status	: Operational	
		Info	: 3.0 GBPS	
+-----+				
	SAS Address	: 5000CCA003263941	SAS Address	: 5000CCA003263941
	Status	: Missing	Status	: Operational
+-----+				
		Device Lun		
		DMP013		
+-----+				

For SAS Fabric Path Data View, pressing the F11 key will display information similar to the following.

#### SAS Fabric Path Data View

Adapter	Adapter Port	Path Active	Path State	Device
DC01	4	No	Failed	DMP013
DC01	6	Yes	Operational	DMP013

Node	SAS Address	Port Type	Port	Status	Info
1	5005076C07434609	Adapter	4	Operational	3.0 GBPS
2	500A0B81E1AA9000	Expander	20	Operational	3.0 GBPS
3	500A0B81E1AA9000	Expander	FF	Missing	Status 0
4	5000CCA003263941	Device Lun	1	Missing	Status 0
1	5005076C07434609	Adapter	6	Operational	3.0 GBPS
2	500A0B81E1B07000	Expander	20	Operational	3.0 GBPS
3	500A0B81E1B07000	Expander	16	Operational	3.0 GBPS
4	500A0B8245C4A000	Expander	16	Operational	3.0 GBPS
5	500A0B8245C4A000	Expander	1	Operational	3.0 GBPS
6	5000CCA003A63941	Device	1	Operational	3.0 GBPS
7	5000CCA003263941	Device Lun	1	Operational	Status 0



---

## Chapter 7. SAS address and physical location information

Many hardware error logs identify the location of a physical device, such as a SAS disk, using what is called a unit address.

The resource format is: cceellFF where:

- cc identifies the controller's port to which the device, or device enclosure, is attached.
- ee is the expander's port to which the device is attached. When a device is not connected to a SAS expander, for example, the device is directly connected, the expander port is set to zero.

Typically, the expander port will be in a range of 00 to 3F hex. A value greater than 3F indicates there are two expanders (for example, cascaded expanders) between the controller and device. For example, a device connected through a single expander might show an expander port of 1A, while a device connected through a cascaded expander might show an expander port of 5A (that is, a value of 40 hex added to the expander port indicates the presence of a cascaded expander), but in both cases, the device is connected off port 1A of the expander.

A value of FF indicates the expander port is not known.

- ll is the logical unit number (LUN) of the device.

A value of FF indicates the LUN is not known.

A resource can identify a physical device or it can identify other SAS components. For example:

- FFFFFFFF indicates the identity of the device is not known or indicates a SAS RAID controller.
- ccFFFFFF identifies only a controller's SAS port.
- cceellFF identifies the controller port, expander port, and LUN of an attached device.

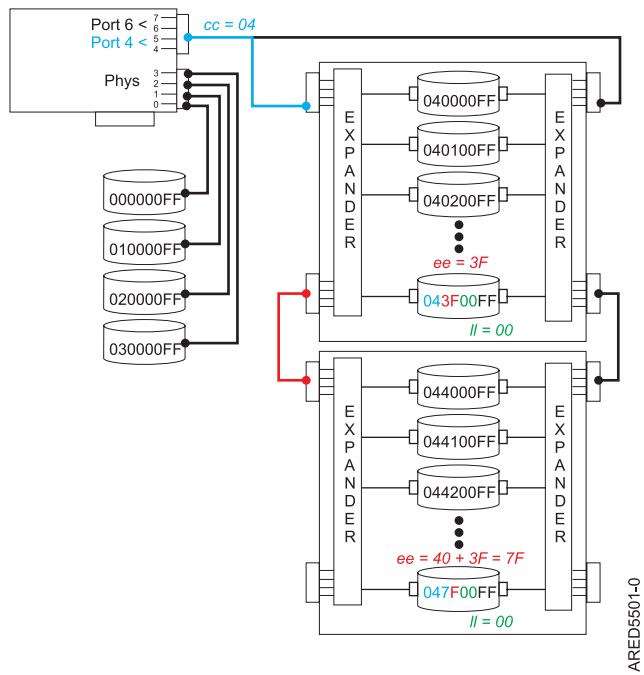


Figure 27. Example of SAS subsystem unit addresses



---

## Appendix. Notices

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

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

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

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** THIS INFORMATION IS PROVIDED "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 states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

Any references in this information to Web sites not owned by the manufacturer are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this product and use of those Web sites is at your own risk.

The manufacturer may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

All statements regarding the manufacturer's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The manufacturer's prices shown are the manufacturer'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 the names and addresses used by an actual business enterprise 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 the manufacturer.

The manufacturer has prepared this information for use with the specific machines indicated. The manufacturer makes no representations that it is suitable for any other purpose.

The manufacturer'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 the manufacturer's support websites for updated information and fixes applicable to the system and related software.

---

## Trademarks

IBM, the IBM logo, and [ibm.com](http://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](http://www.ibm.com/legal/copytrade.shtml).

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

---

## Electronic emission notices

### Class A Notices

The following Class A statements apply to the IBM servers that contain the POWER6 processor.

### 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**

This Class A digital apparatus complies with Canadian ICES-003.

### **Avis de conformité à la réglementation d'Industrie Canada**

Cet appareil numérique de la classe A respecte est conforme à la norme NMB-003 du Canada.

### **European Community Compliance Statement**

This product is in conformity with the protection requirements of EU Council Directive 2004/108/EC 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.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to European Standard EN 55022. The limits for Class A equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

European Community contact:  
IBM Technical Regulations  
Pascalstr. 100, Stuttgart, Germany 70569  
Tele: 0049 (0)711 785 1176  
Fax: 0049 (0)711 785 1283  
E-mail: tjahn@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.

**Japanese Electronics and Information Technology Industries Association (JEITA)  
Confirmed Harmonics Guideline (products less than or equal to 20 A per phase)**

高調波ガイドライン適合品

**Japanese Electronics and Information Technology Industries Association (JEITA)  
Confirmed Harmonics Guideline with Modifications (products greater than 20 A per phase)**

高調波ガイドライン準用品

**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 업무용으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이점을 주의하시기 바라며, 만약 잘못 판매 또는 구입하였을 때에는 가정용으로 교환하시기 바랍니다.

Please note that this equipment has obtained EMC registration for commercial use. In the event that it has been mistakenly sold or purchased, please exchange it for equipment certified for home use.

## 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 2004/108/EG zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55022 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 der IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung der IBM gesteckt/eingebaut werden.

EN 55022 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 2004/108/EG in der Bundesrepublik Deutschland.

### Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC EG Richtlinie 2004/108/EG) für Geräte der Klasse A.

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Konformitätserklärung nach des EMVG ist die IBM Deutschland GmbH, 70548 Stuttgart.

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 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

---

## Terms and conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

**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 the manufacturer.

**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 the manufacturer.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any data, software or other intellectual property contained therein.

The manufacturer 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 the manufacturer, 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.

THE MANUFACTURER MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THESE 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.





Printed in USA