

Power Systems

Enclosures and expansion units

IBM

Power Systems

Enclosures and expansion units

IBM

Note

Before using this information and the product it supports, read the information in “Safety notices” on page vii, “Notices” on page 143, 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 POWER7 processor and to all associated models.

© Copyright IBM Corporation 2010, 2013.

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

Contents

Safety notices	vii
-----------------------	------------

I/O expansion units	1
----------------------------	----------

Overview for I/O expansion units	1
Adding I/O expansion units	2
Preparing to add an I/O expansion unit	2
Where to install a new expansion unit	3
Identifying 12X and SPCN cables	3
Expansion unit configuration rules	4
Adding an I/O expansion unit with the system powered on	7
Adding an I/O expansion unit with the system powered off	13
Connecting expansion units with 12X cables	16
Adding expansion units to a new 12X loop	16
Adding expansion units to an existing 12X loop	17
Examples: 12X connections	18
Connecting expansion units with SPCN cables	19
Examples: SPCN connections	20
Verifying that the new configuration is functioning	21
Verifying the 12X loops with an HMC	22
Verifying the 12X loops without an HMC	23
Verifying the system power control network	25
Setting I/O enclosure configuration ID and MTMS value	29
Using the ASMI to verify and set the configuration ID and MTMS value	29
Updating the SPCN firmware	31
Configuring I/O expansion units	32
Configuring the disk subsystem in the 5802 expansion unit	32
Cabling the 5802 disk subsystem	35
Removing I/O expansion units	41
Removing an I/O expansion unit with the system powered on	41
Removing an I/O expansion unit with the system powered off	44
Removing an enclosure from the rack	46

Disk drive enclosures	51
------------------------------	-----------

Overview for disk drive enclosures	51
5786, 5787, 7031-D24, or 7031-T24 SCSI disk drive enclosure	51
Connecting the 5786, 5787, 7031-D24, or 7031-T24 to a system running AIX	52
Connecting and configuring the 5786, 5787, 7031-D24, or 7031-T24 SCSI disk drive enclosure in an AIX clustered environment	55
Connecting and configuring the disk drive enclosure in a system running Linux	57
Connecting and configuring the enclosure in a system running IBM i	61
SCSI repeater card	64
5886 SAS disk drive enclosure	67
Connecting the SAS adapter to the 5886 disk drive enclosure	68
5887 SAS disk drive enclosure	72
Connecting the SAS adapter to the 5887 disk drive enclosure	73
Removing your disk drive enclosures	85

PCIe storage enclosures	87
--------------------------------	-----------

Overview for PCIe storage enclosures	87
Adding PCIe storage enclosures	87
Preparing to add a PCIe storage enclosure	88
Where to install a PCIe storage enclosure	88
Identifying PCIe cables	88
PCIe storage enclosure configuration rules	90
Adding a PCIe storage enclosure with the system powered on	92

Adding a PCIe storage enclosure with the system powered off	98
Connecting PCIe storage enclosures with PCIe cables	99
Verifying the PCIe hardware topology	102
Verifying the PCIe hardware topology with an HMC	102
Verifying the PCIe hardware topology without an HMC	103
Verifying that the PCIe storage enclosure is functioning	105
Checking for new serviceable events	106
Using identify indicators to locate hardware components	107
Configuring PCIe storage enclosures	108
Configuring the 5888 PCIe storage enclosure	108
Configuring the EDR1 PCIe storage enclosure	109
Removing PCIe storage enclosures	110
Removing a PCIe storage enclosure with the system powered on	110
Removing a PCIe storage enclosure with the system powered off	116
Reference information	119
Common procedures	119
Starting the system or logical partition	119
Starting a system that is not managed by an HMC or an SDMC	119
Starting a system or logical partition by using the HMC	120
Stopping a system or logical partition	121
Stopping a system that is not managed by an HMC or an SDMC	121
Stopping a system by using the HMC	122
Powering on a system to firmware standby state	123
Removing an enclosure from the rack	124
Connector locations	126
Server connectors	127
Model 8202-E4B connector locations	127
Model 8202-E4C connector locations	127
Model 8202-E4D connector locations	128
Model 8205-E6B connector locations	129
Model 8205-E6C connector locations	130
Model 8205-E6D connector locations	131
Model 8231-E1C connector locations	132
Model 8231-E1D or 8268-E1D connector locations	132
Model 8231-E2B connector locations	133
Model 8231-E2C connector locations	133
Model 8231-E2D connector locations	134
Model 8233-E8B connector locations	135
Model 8246-L1S connector locations	136
Model 8246-L1T connector locations	136
Model 8246-L2S connector locations	136
Model 8246-L2T connector locations	137
Model 8248-L4T, 8408-E8D, or 9109-RMD connector locations	137
Model 9117-MMB or 9179-MHB connector locations	138
Model 9117-MMC or 9179-MHC connector locations	139
Model 9117-MMD or 9179-MHD connector locations	139
Enclosure connectors	140
Model 5796 or 7314-G30 connector locations	140
Model 5802 and 5877 connector locations	141
Model 5886 connector locations	141
Model 5887 connector locations	142
Model 5888 connector locations	142
Model EDR1 connector locations	142
Notices	143
Trademarks	144
Electronic emission notices	144
Class A Notices	144
Class B Notices	148

Terms and conditions. 151

Safety notices

Safety notices may be printed throughout this guide:

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

World Trade safety information

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

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

German safety information

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

Laser safety information

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

Laser compliance

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

DANGER

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

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

- Connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

To Disconnect:

1. Turn off everything (unless instructed otherwise).
2. Remove the power cords from the outlets.
3. Remove the signal cables from the connectors.
4. Remove all cables from the devices.

To Connect:

1. Turn off everything (unless instructed otherwise).
2. Attach all cables to the devices.
3. Attach the signal cables to the connectors.
4. Attach the power cords to the outlets.
5. Turn on the devices.

(D005)

DANGER

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

- Heavy equipment—personal injury or equipment damage might result if mishandled.
- Always lower the leveling pads on the rack cabinet.
- Always install stabilizer brackets on the rack cabinet.
- To avoid hazardous conditions due to uneven mechanical loading, always install the heaviest devices in the bottom of the rack cabinet. Always install servers and optional devices starting from the bottom of the rack cabinet.
- Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices.



- Each rack cabinet might have more than one power cord. Be sure to disconnect all power cords in the rack cabinet when directed to disconnect power during servicing.
- Connect all devices installed in a rack cabinet to power devices installed in the same rack cabinet. Do not plug a power cord from a device installed in one rack cabinet into a power device installed in a different rack cabinet.
- An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock.

CAUTION

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

(R001)

CAUTION:

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


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

(R002)


(L001)



(L002)



(L003)



or



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)

CAUTION:

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

Do Not:

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

Exchange only with the IBM-approved part. Recycle or discard the battery as instructed by local regulations. In the United States, IBM has a process for the collection of this battery. For information, call 1-800-426-4333. Have the IBM part number for the battery unit available when you call. (C003)

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 metallically 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 metallically 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.

I/O expansion units

Use this information to learn about connecting and configuring input/output (I/O) expansion units to system units.

Overview for I/O expansion units

Input/output expansion units provide additional I/O slots. Some I/O expansion units also provide additional disk and removable media slots.

Input/output (I/O) expansion units are connected to system units with 12X cables and system power control network (SPCN) cables. The 12X cables can transport data and control information. The SPCN cables can only transport control information.

The 12X cables from an I/O expansion unit connect to a 12X cable port pair (also called a connector pair) at the system unit. The 12X port pair is either on a pluggable card that is installed in the system unit or mounted directly on the system unit chassis. In either case, the connectors are wired to an I/O hub chip that drives the 12X protocol. The pluggable cards are referred to as GX adapters or GX cards.

The 12X interconnection between I/O expansion units and system units can be single data rate (SDR) or double data rate (DDR). Some I/O expansion units are capable of SDR only. Others are capable of both SDR and DDR. Similarly, some I/O hub chips and 12X cables are capable of SDR only while others are capable of both SDR and DDR. In order for a 12X interconnection between a system unit and I/O expansion unit to run at DDR, the I/O hub chip, 12X cable, and expansion unit must all be DDR capable. While components that are strictly SDR capable cannot run at DDR, components that are DDR capable can be configured by the system firmware to run at SDR if one or more of the other components in the interconnection are strictly SDR capable.

For example, a DDR-capable I/O expansion unit that is connected to a GX adapter that has a strictly SDR-capable I/O hub chip will be configured to run at SDR. Similarly, a DDR-capable GX adapter that is connected to a strictly SDR-capable I/O expansion unit will be configured to run at SDR.

Notes:

- Except where the distinction is important, a 12X cable refers to either a 12X SDR cable or a 12X DDR cable.
- Except where the distinction is important, a GX adapter or GX card refers to either an SDR or a DDR version.
- The terms I/O expansion unit and expansion unit are used interchangeably.

Important:

- You can add your expansion units concurrently (with the server powered on and partitions active) if you have one of the following configurations:
 - Your system is managed by an IBM Hardware Management Console (HMC).
 - Your system is not managed by an HMC, but has only one partition, and that partition is running the IBM i operating system.

If your configuration does not permit adding expansion units concurrently, you must power off the server to add expansion units.

- If additional GX adapters are needed to accommodate the new I/O expansion units, refer to the instructions that came with the GX adapters for information about how to install them. Installing GX

adapters must be done separately from installing expansion units. Depending on your configuration, it might be necessary to power off the system to install the GX adapters.

- Expansion units cannot be moved or relocated concurrently. If existing expansion units must be relocated in the configuration to add the new expansion units, complete the following tasks:
 1. Power off the system.
 2. Physically relocate the existing expansion units.
 3. Power on the system to firmware standby or firmware running state. The relocation of the existing expansion units is complete.
 4. Add the new expansion units.
- Before performing any procedure that involves changes to the 12X cabling or configuration for a system that is managed by an HMC, obtain a record that identifies for each expansion unit which I/O buses exist in that expansion unit. To acquire this information, perform the following steps:
 1. On the HMC, select the managed system and then select the Properties task.
 2. Enter the following from the HMC command line:

```
lshwres -r io --subtype bus -m managed_system
```
- Improper completion of the procedure to connect expansion units can result in unintended I/O configuration changes, including a change to the bus numbers assigned to existing expansion units. When bus numbers change, partition profiles cannot find existing I/O resources. If this procedure completes with unexpected results, contact your next level of support. Your next level of support can attempt to restore bus numbers, provided that you obtained a record of the original configuration.

Adding I/O expansion units

Learn how to add an input/output (I/O) expansion unit to the system and verify that it is functioning.

You might be able to add an I/O expansion unit to your system with the system powered on and logical partitions active. You can add an I/O expansion unit with the system powered on and logical partitions active if your configuration meets one of the following requirements:

- Your system is managed by an IBM Hardware Management Console (HMC).
- Your system is not managed by an HMC, but has only one partition, and that partition is running the IBM i operating system.

If your configuration does not meet any of the requirements specified above, you must power off the system to add the I/O expansion unit.

Preparing to add an I/O expansion unit

Use this information to plan the installation and configuration of an I/O expansion unit.

This topic collection provides information about connecting an I/O expansion unit with 12X and system power control network (SPCN) cables. This task is a customer task. You can perform this task yourself or contact a service provider for this service.

To prepare to connect the I/O expansion unit, complete the following tasks:

1. Decide where you want to install the new expansion units. For details, see “Where to install a new expansion unit” on page 3.
2. Unpack the expansion units by using the unpacking instructions.
3. Identify the cables. For details, see “Identifying 12X and SPCN cables” on page 3.
4. Plan for cable layout. When you decide where to place the cables, follow your site plan and keep the following items in mind:
 - Avoid creating a safety hazard.
 - Avoid damaging the cables.

- Avoid placing cables parallel to high-voltage lines.
5. Continue with “Adding I/O expansion units” on page 2.

Where to install a new expansion unit

Before you begin the installation process, plan where to install the new expansion units.

When deciding where to install a new expansion unit, consider several elements that include size, security, and environmental factors. For more information, see Site preparation and physical planning (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ebe/p7ebegenconsiderations.htm>).

Note: Two expansion units, the 5796 and 7314-G30, can be ordered with a 12X adapter that has a repeater to support longer 12X cable lengths, or a 12X adapter without a repeater that does not support longer cable lengths. If you are installing a 5796 or 7314-G30 expansion unit with a Short Run Dual-Port Channel Attach 12X adapter (FC 6446), the expansion unit must be placed in the same rack as the system unit. If you are installing a 5796 or 7314-G30 expansion unit with a Long Run Dual-Port Channel Attach 12X adapter (FC 6457), the expansion unit can be placed in the same rack as the system unit, or a different rack.

Identifying 12X and SPCN cables

Use this information to help identify the 12X cables and system power control network (SPCN) cables for the expansion unit.

The system uses 12X SDR or 12X DDR cables to send and receive customer data and auxiliary control information to and from the expansion unit, and to download expansion unit firmware. The SPCN is used to control power to the expansion units and is used as a backup to the 12X cables for auxiliary control and firmware download.

12X SDR cables

Observe the following guidelines when you are determining whether to use a 12X SDR cable:

- A 12X SDR cable can be used only with 5796, 5797, 5798, and 7314-G30 expansion units.
- A 12X SDR cable cannot be used in place of a 12X DDR cable.

Table 1. 12X SDR cables

Cable feature	Length	Part number	Additional information
1829	0.6 m (1.9 ft.)	39J2064	Withdrawn from new feature sale
1830	1.5 m (4.9 ft.)	97P3650	Withdrawn from new feature sale
1831	2.5 m (8.2 ft.)	45D4787 or 42V2131	Withdrawn from new feature sale
1834	8.0 m (26.2 ft.)	45D4788 or 42R6160	Withdrawn from new feature sale
1840	3.0 m (9.8 ft.)	45D5271 or 42V2132	Withdrawn from new feature sale

12X DDR cables

Observe the following guidelines when you are determining whether to use a 12X DDR cable:

- Any 12X DDR cable can be used in any 12X application, as long as the specified length restrictions are observed.

- A 12X DDR cable can be used instead of a 12X SDR cable, as long as the specified length restrictions are observed.

Table 2. 12X DDR cables

Cable feature	Length	Part number	Additional information
1828	1.5 m (4.9 ft.)	45D2235	12X to 4X channel conversion adapter used for clustering
1841	3.0 m (9.8 ft.)	45D2236	12X to 4X channel conversion cable used for clustering
1842	10.0 m (32.8 ft.)	45D2237	12X to 4X channel conversion cable used for clustering
1861	0.6 m (1.9 ft.)	45D4785	
1862	1.5 m (4.9 ft.)	45D4786	
1863	2.5 m (8.2 ft.)	45D4787	
1864	8.0 m (26.2 ft.)	45D4788	
1865	3.0 m (9.8 ft.)	45D5271	

SPCN cables

Table 3. SPCN cables

Cable feature	Length	Part number
6001	2.0 m (6.6 ft.)	22R5217
6006	3.0 m (9.8 ft.)	22R5239
6007	15.0 m (49.2 ft.)	22R5221
6008	6.0 m (19.6 ft.)	22R5219
6029	30.0 m (98.4 ft.)	22R5222

Expansion unit configuration rules

Learn about important configuration rules that must be followed when connecting expansion units to system units.

When you set up an expansion unit, observe the configuration rules in the following list:

- Expansion units must be connected to system units with 12X cables in a loop topology. Doing so produces redundancy in the 12X connections so that a single cable failure does not result in loss of communication between the system unit and one or more expansion units.
- Expansion units must be connected to system units with SPCN cables in a loop topology. Doing so produces redundancy in the SPCN connections so that a single cable failure does not result in loss of communication between the system unit and one or more expansion units.
- Expansion units with 12X SDR capability cannot be in the same 12X loop as expansion units with 12X DDR capability. The following expansion units have 12X SDR capability:
 - 7314-G30
 - 5796
 - 5797
 - 5798

The following expansion units have 12X DDR capability:

- 5802
- 5803
- 5873
- 5877

The following table specifies configuration limitations related to 12X loops.

Table 4. 12X loop configuration limitations

Machine type and model	Supported expansion units	Maximum number of loops	Maximum number of expansion units per loop	Maximum number of supported expansion units
8202-E4B 8202-E4C 8202-E4D	7314-G30 5796 5802 5877	Maximum number of loops: 1 loop	7314-G30 = 4 5796 = 4 5802 = 2 5877 = 2	7314-G30 = 4 5796 = 4 5802 = 2 5877 = 2
8205-E6B 8205-E6C 8205-E6D	7314-G30 5796 5802 5877	Maximum number of loops: 2 loops	7314-G30 = 4 5796 = 4 5802 = 2 5877 = 2	7314-G30 = 8 5796 = 8 5802 = 4 5877 = 4
8231-E2C 8231-E2D 8246-L1S 8246-L2T	5802 5877	Maximum number of loops: 1 loop	5802 = 2 5877 = 2	5802 = 2 5877 = 2
8233-E8B	7314-G30 5796 5802 5877	The maximum number of loops varies with the number of processor cards, with a maximum of four processor cards: 1 processor card = 1 loop 2 processor cards = 2 loops 3 processor cards = 2 loops 4 processor cards = 2 loops	7314-G30 = 4 5796 = 4 5802 = 2 5877 = 2	7314-G30 = 8 5796 = 8 5802 = 4 5877 = 4

Table 4. 12X loop configuration limitations (continued)

Machine type and model	Supported expansion units	Maximum number of loops	Maximum number of expansion units per loop	Maximum number of supported expansion units
<p>8248-L4T Note: The 16-way, 24-way, and 32-way configurations support expansion units. The 8-way configuration does not support expansion units.</p> <p>8408-E8D Note: The 16-way, 24-way, and 32-way configurations support expansion units. The 8-way configuration does not support expansion units.</p> <p>9109-RMD Note: The 24-way, 36-way, and 48-way configurations support expansion units. The 12-way configuration does not support expansion units.</p>	<p>5802</p> <p>5877</p>	<p>Maximum number of loops: 2 loops</p>	<p>5802 = 2</p> <p>5877 = 2</p>	<p>5802 = 4</p> <p>5877 = 4</p>
<p>8412-EAD</p> <p>9117-MMB</p> <p>9117-MMC</p> <p>9117-MMD</p> <p>9179-MHB</p> <p>9179-MHC</p> <p>9179-MHD</p>	<p>7314-G30</p> <p>5796</p> <p>5802</p> <p>5877</p>	<p>The maximum number of loops is equal to two loops per processor drawer in the system configuration:</p> <p>1 processor drawer = 2 loops</p> <p>2 processor drawers = 4 loops</p> <p>3 processor drawers = 6 loops</p> <p>4 processor drawers = 8 loops</p>	<p>7314-G30 = 4</p> <p>5796 = 4</p> <p>5802 = 2</p> <p>5877 = 2</p>	<p>7314-G30 = 32</p> <p>5796 = 32</p> <p>5802 = 16</p> <p>5877 = 16</p>

Table 4. 12X loop configuration limitations (continued)

Machine type and model	Supported expansion units	Maximum number of loops	Maximum number of expansion units per loop	Maximum number of supported expansion units
9119-FHB	5797 5798 5803 5873	The maximum number of loops is equal to two loops per processor drawer in the configuration: 1 processor book = 4 loops 2 processor books = 8 loops 3 processor books = 12 loops 4 processor books = 16 loops 5 processor books = 20 loops 6 processor books = 24 loops 7 processor books = 28 loops 8 processor books = 32 loops	5797 = 1 5798 = 1 5803 = 1 5873 = 1	5797 = 32 5798 = 32 5803 = 32 5873 = 32
8231-E1C 8231-E1D 8231-E2B 8236-E8C 8246-L1C 8246-L1D 8246-L1T 8246-L2C 8246-L2D 8246-L2S 8268-E1D	System does not support expansion units.	Not applicable.	Not applicable.	Not applicable.

Adding an I/O expansion unit with the system powered on

Learn about adding an input/output (I/O) expansion unit to the system with the system powered on.

Important: To use the following procedure to add I/O expansion units to the system, your configuration must meet one of the following requirements:

- Your system is managed by an IBM Hardware Management Console (HMC).
- Your system is not managed by an HMC, but has only one partition, and that partition is running the IBM i operating system.

If your configuration does not meet any of these requirements, you must add your I/O expansion unit by completing the procedure at “Adding an I/O expansion unit with the system powered off” on page 13.

Prerequisites: If you did not review the information in “Overview for I/O expansion units” on page 1 and “Preparing to add an I/O expansion unit” on page 2, do so now.

Note: Steps in this topic require you to have connector information for the hardware in your configuration. For details about expansion unit or system connector locations, see “Connector locations” on page 126.

To add an I/O expansion unit to the system with the system powered on, complete the following steps:

1. Verify the existing 12X configuration by choosing one of the following options.

Important: Any problems that you discover with the 12X loops in this step must be corrected before you continue with the procedure.

- If your system is managed by an HMC, see Verifying the 12X loops with an HMC (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_loopverifyconsole.htm)

Tip: This task includes a step to verify the status of the links in the 12X loops. When you are instructed to use the **View Hardware Topology** task, you can click **Save** to save the current hardware topology. Use this action to compare the topology that existed before the expansion unit was added with the topology that was created by adding the expansion unit.

- If your system is not managed by an HMC, see Verifying the 12X loops without an HMC (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_loopverifynoconsole.htm).
2. Verify the system power control network (SPCN). For instructions, see Verifying the system power control network (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_spcnverify_5802.htm).

Important: Any problems that you discover with the SPCN in this step must be corrected before you continue with the procedure.

3. Set the SPCN Firmware Update Policy to prevent firmware updates over the serial SPCN interface by completing the following steps.

Note: The SPCN Firmware Update Policy controls when and how SPCN firmware on expansion units is updated. The default setting for the SPCN Firmware Update Policy is Enabled. This setting allows for firmware updates to be completed over the 12X interface whenever an update is required and does not allow updates over the slower serial SPCN interface.

- a. Access the Advanced System Management Interface (ASMI) by using an authority level of administrator or authorized service provider. For instructions, see Managing the Advanced System Management Interface (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hby/asmi.htm>).
- b. In the ASMI navigation area, expand **System Configuration** and select **Configure I/O Enclosures**.
 - 1) If the SPCN Firmware Update Policy shows Expanded, record the current setting so that you can restore it later, and then, change the setting to Enabled.
 - 2) If the SPCN Firmware Update Policy shows Enabled or Disabled, continue without changing the setting.
- c. Check to see whether an SPCN firmware update is in progress and choose one of the following options.

Note: An SPCN firmware update is in progress if the Power Control Network Firmware Update Status column shows In Progress and displays the percentage that is completed.

- If an SPCN firmware update is in progress, continue with step 4.
 - If an SPCN firmware update is not in progress, continue with step 5.
4. Check the Power Control Network Firmware Update Status column to determine the type of update that is in progress, and then, choose one of the following options:
 - If the SPCN firmware update is a High Speed Link (HSL) update, wait for the update to complete before you continue. An HSL update completes quickly. To determine when the update is complete, select **Configure I/O Enclosures** every 15 – 30 seconds until Not Required is displayed in the Power Control Network Firmware Update Status column. Do not click **Back** or **Refresh** on the browser to monitor the status.
 - If the SPCN firmware update is a serial update, choose one of the following options:
 - Wait for the update to complete. A serial update can take 20 minutes or more depending on how many expansion units are in the configuration. To determine when the update is complete, select **Configure I/O Enclosures** every 15 – 30 seconds until Not Required is displayed in the Power Control Network Firmware Update Status column. Do not click **Back** or **Refresh** on the browser to monitor the status.
 - Stop the update by completing the following steps:

Note: The update is restarted from the beginning later in the procedure.

- a. Click **Stop SPCN Firmware Update** in the Configure I/O Enclosures window.
 - b. Click **Configure I/O Enclosures** every 15 - 30 seconds until Pending is displayed in the Power Control Network Firmware Update Status column. Do not use click the **Back** or **Refresh** on the browser to monitor the status.
5. If the system has a rear cover, remove or open it.
 6. Record the current date and time for use later in the procedure when you check for serviceable events.
 7. Connect the expansion unit with 12X cables. For instructions, see Connecting expansion units with 12X cables (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_conncable12x_5802.htm).

Note: When you connect an expansion unit in an existing 12X loop with the system powered on, serviceable events and informational logs might be generated when you disconnect and reconnect cables. This behavior is normal and is covered later in the procedure.

8. Connect the expansion unit with SPCN cables. For instructions, see Connecting expansion units with SPCN cables (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_conncablespcn_5802.htm).

Note: When you connect an expansion unit in an existing SPCN loop with the system powered on, serviceable events and informational logs might be generated when you disconnect and reconnect cables. This behavior is normal and is covered later in the procedure.

9. Connect the power cables for the expansion unit that you are installing. The expansion unit powers on automatically when the power cables are connected.

Important: Ensure that you plug the power cable into the power source before you connect it into the power supply of the expansion unit.

10. Wait for the green power indicator on the control panel of the expansion unit to light solid. This action might take 1 - 10 minutes. During this time, a yellow attention indicator might be lit on the control panel of the expansion unit. After the green power indicator is lit on the control panel, wait 10 extra minutes for hardware initialization to complete before you proceed with the installation.
11. Verify the SPCN configuration by completing the following steps.

- a. Read and understand the following information about reference codes that might be generated when an expansion unit is added while the system is powered on to firmware standby state.

Note: When an expansion unit is added while the system is powered on, it is normal for error logs and event logs to be created with the following SPCN-related reference codes. The number and combination of these logs is dependent on the configuration.

Logs with the following reference codes can be ignored during this verification step if their time stamps indicate that they occurred during the add procedure and the resources that they refer to were involved in the procedure.

- 10009133: The 12X loop was disrupted. This reference code appears in an informational log.
- 10009135: The SPCN loop was opened. This reference code might appear in a serviceable event or an informational log.
- 10009136: An SPCN firmware update was stopped. This reference code appears in an informational log.
- 10009137: The 12X loop was disrupted. This reference code appears in an informational log.
- 10009139: The SPCN loop was closed. This reference code appears in an informational log.
- 1000913B: An SPCN firmware update was required, but not automatically started. This reference code appears in a serviceable event.
- 1000910A: An SPCN firmware update was initiated. This reference code appears in an informational log.
- 100091DE: An SPCN firmware update has completed. This reference code appears in an informational log.

- b. Complete the verification task that is specified in *Verifying the system power control network* (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_spcnverify_5802.htm).

12. Verify that the Power Control Network Identifier value is set correctly for the expansion unit that you are adding by completing the following steps:

- a. Access the ASMI by using an authority level of administrator or authorized service provider. For instructions, see *Managing the Advanced System Management Interface* (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hby/asmi.htm>).
- b. In the ASMI navigation area, expand **System Configuration** and select **Configure I/O Enclosures**.
- c. Review the Enclosure Configuration table to locate the expansion unit that you added.
- d. Use the following information to verify that the expansion unit that you added has the correct Power Control Network Identifier:
 - Identifier 0x8D is for 5796 and 7314-G30 expansion units.
 - Identifier 0x8E is for 5802 and 5877 expansion units.
- e. If an incorrect Power Control Network Identifier is specified for the expansion unit that you added, contact your service provider for assistance before you continue with this procedure.

13. Verify the 12X configuration by completing the following steps:

- a. Read and understand the following information about reference codes that might be generated when an expansion unit is added while the system is powered on.

Note: When an expansion unit is added while the system is powered on, it is normal for error logs and event logs to be created with the following 12X-related reference codes. The number and combination of these logs is dependent on the configuration.

Logs with the following reference codes can be ignored during this verification step if their time stamps indicate that they occurred during the add procedure and the resources that they refer to were involved in the procedure.

- B7006907: An expansion unit was added that was previously configured on a different system. This reference code appears in an informational log.
- B7006982: A 12X connection failure was experienced. This reference code appears in an informational log.
- B7006984: A 12X loop was opened. This reference code appears in an informational log.
- B7006985: A 12X loop was closed. This reference code appears in an informational log.
- B70069E6: A 12X link was lost. This reference code appears in an informational log.
- B70069E7: A 12X link was restored. This reference code appears in an informational log.

b. Choose one of the following options.

- If your system is managed by an HMC, see *Verifying the 12X loops with an HMC* (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_loopverifyconsole.htm).

Tip: If you saved the original hardware topology in step 1 on page 8 before you added the expansion unit, the saved topology now appears in the last Valid Hardware Topology section of the **View Hardware Topology** output. Use this operation to compare the topology that was created by adding the expansion unit to the topology that existed before the expansion unit was added.

- If your system is not managed by an HMC, see *Verifying the 12X loops without an HMC* (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_loopverifynoconsole.htm).

14. Check for new serviceable events that were generated during the procedure by choosing one of the following options:

Important: Multiple options might apply to your configuration. Choose only the first option that applies.

a. If your system is managed by an HMC, complete the following steps:

- 1) In the navigation area, select **Serviceability > Manage Serviceable Events**.
- 2) Specify event criteria that includes events that were generated during the procedure. In particular, specify the managed system that the expansion unit was added to, and specify a date and time range that includes the data and time that were recorded in step 6 on page 9.
- 3) Continue with step 15 on page 12.

b. If your system is not managed by an HMC and has a partition that is running the Virtual I/O Server (VIOS) operating system, complete the following steps:

- 1) Log in as the padmin user. If you need assistance, contact the system administrator.
- 2) Type `diagmenu -d sysplanar0 -E xx` and press Enter, where `xx` specifies the number of days in the range 1 - 60. This command returns results for the most recent number of days specified. Specify a range that includes the date and time that was recorded in step 6 on page 9.
- 3) On the DIAGNOSTIC MODE SELECTION display, highlight the **Problem Determination** option, and press Enter.
- 4) Continue with step 15 on page 12.

c. If your system is not managed by an HMC and has a partition that is running the AIX® operating system, complete the following steps:

- 1) Log in as the root user, or use the CE login. If you need assistance, contact the system administrator.
- 2) Type `diag -d sysplanar0 -E xx` and press Enter, where `xx` specifies the number of days in the range 1 - 60. This command returns results for the most recent number of days specified. Specify a range that includes the date and time that was recorded in step 6 on page 9.

- 3) On the DIAGNOSTIC MODE SELECTION display, highlight the **Problem Determination** option, and press Enter.
- 4) Continue with step 15.
- d. If your system is not managed by an HMC and has a partition that is running the IBM i operating system, complete the following steps:
 - 1) Sign on with at least service level authority.
 - 2) At the command line of the IBM i session, type strsst and press Enter.
 - 3) Type your service tools user ID and service tools password on the System Service Tools (SST) Sign On display, and press Enter. The service tools password is case-sensitive.
 - 4) Click **Start a Service Tool > Hardware Service Manager > Work with Service Action Event Log**.
 - 5) On the **Select Timeframe** display, change the **From Date and Time** field to the wanted date and time range. Specify a range that includes the date and time that was recorded in step 6 on page 9.
 - 6) Continue with step 15.
- e. If your system is not managed by an HMC and has a partition that is running the Linux operating system, complete the following steps:
 - 1) Log in as the root user. If you need assistance, contact the system administrator.
 - 2) Type `servicelog -query='serviceable=1 AND closed=0 AND time_event>="yyyy-mm-dd"'` and press Enter, where yyyy-mm-dd is the day on which the procedure occurred.
 - 3) Continue with step 15.
15. Handle serviceable events that were generated during the procedure by completing the following steps:
 - a. Close any serviceable events that were generated during the procedure that includes one of the reference codes that are listed in step 11 on page 9 or step 13 on page 10.
 - b. Perform problem analysis on any remaining open serviceable events.
16. Start necessary SPCN firmware updates.

If you found the reference code 1000913B in step 11a on page 10, an SPCN firmware update is needed. However, because of the current SPCN Firmware Update Policy, the update could not be initiated automatically. To initiate the SPCN firmware update manually, complete the following steps:

 - a. To initiate the SPCN firmware updates, see [Updating the SPCN firmware \(http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_spcnupdate.htm\)](http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_spcnupdate.htm).
 - b. Wait for the SPCN firmware updates to complete before you continue to the next step. This update can take 20 minutes to several hours, depending on the number of expansion units on the SPCN loop. The amount of time can also be influenced by the current level of the SPCN firmware in the expansion units that were added.
17. If you changed the SPCN Firmware Update Policy in step 3b1 on page 8, change the policy back to the original setting now by completing the following steps:
 - a. Access the ASMI by using an authority level of administrator or authorized service provider. For instructions, see [Managing the Advanced System Management Interface \(http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hby/asmi.htm\)](http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hby/asmi.htm).
 - b. In the ASMI navigation area, expand **System Configuration** and select **Configure I/O Enclosures**.
 - c. Change the **SPCN Firmware Update Policy** setting to the setting recorded in step 3b1 on page 8.

Note: The default setting of **Enabled** is recommended.
18. Verify the new configuration. For instructions, see [Verifying that the new configuration is functioning \(http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_connverify_5802.htm\)](http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_connverify_5802.htm).

19. If the system has logical partitions, you can now assign I/O slots in the expansion units that were added to logical partitions.

You have completed the addition of your expansion unit.

Adding an I/O expansion unit with the system powered off

Learn about adding an input/output (I/O) expansion unit to the system with the system powered off.

Important: If your configuration meets any of the following requirements, you might be able to add I/O expansion units to the system with the system powered on:

- Your system is managed by an IBM Hardware Management Console (HMC).
- Your system is not managed by an HMC, but has only one partition, and that partition is running the IBM i operating system.

If your configuration does not meet any of the requirements that are specified above, you must use this procedure to add the I/O expansion unit. Otherwise, you can use the following procedure, or you can use “Adding an I/O expansion unit with the system powered on” on page 7 instead of this procedure.

Prerequisites: If you have not reviewed the information in “Overview for I/O expansion units” on page 1 and “Preparing to add an I/O expansion unit” on page 2, do so now.

Note: Steps in this topic require you to have connector information for the hardware in your configuration. For details about expansion unit or system connector locations, see “Connector locations” on page 126.

To add an I/O expansion unit to the system with the system powered off, complete the following steps:

1. If the system is not already powered off, remove power from the system. For instructions, see Stopping a system or logical partition (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7haj/crustopsys.htm>).
2. If the system has a rear cover, remove or open it.
3. Record the current date and time for use later in the procedure when you are checking for serviceable events.
4. Connect the expansion unit with SPCN cables. For instructions, see Connecting expansion units with SPCN cables (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_conncablespcn_5802.htm).
5. Power on the system to firmware standby state. For instructions, see .Powering on a system to firmware standby state (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/pxham_firmware_standby.htm).

Note: Connecting expansion units with 12X cables while the system is powered off might cause bus numbers assigned to existing expansion units to change. The circumstances that cause this change to occur are not common. One example is the presence of failed links elsewhere in the loop to which you are adding an expansion unit. Powering the system on to firmware standby state eliminates the risk of bus numbers that are assigned to existing expansion units from changing.

6. Connect the expansion unit with 12X cables. For instructions, see Connecting expansion units with 12X cables (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_conncable12x_5802.htm).

Note: When you connect an expansion unit in an existing 12X loop with the system powered on, serviceable events and informational logs might be generated when you disconnect and reconnect cables. This behavior is normal and is covered later in the procedure.

7. Connect the power cables for the expansion unit that you are installing.

Important: Ensure that you plug the power cable into the power source before you connect it into the power supply of the expansion unit.

8. Verify the system power control network (SPCN). For instructions, see Verifying the system power control network (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_spcnverify_5802.htm).

Important: Any problems that you discover with the SPCN in this step must be corrected before you continue with the procedure.

9. Verify that the Power Control Network Identifier value is set correctly for the expansion unit that you are adding by completing the following steps:
 - a. Access the ASMI by using an authority level of administrator or authorized service provider. For instructions, see Managing the Advanced System Management Interface (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hby/asmi.htm>).
 - b. In the ASMI navigation area, expand **System Configuration** and select **Configure I/O Enclosures**.
 - c. Review the Enclosure Configuration table to locate the expansion unit that you added.
 - d. Use the following information to verify that the expansion unit that you added has the correct Power Control Network Identifier:
 - Identifier 0x8D is for 5796 and 7314-G30 expansion units.
 - Identifier 0x8E is for 5802 and 5877 expansion units.
 - e. If an incorrect Power Control Network Identifier is specified for the expansion unit that you have added, contact your service provider for assistance before you continue with this procedure.
10. Verify the 12X configuration by completing the following steps:
 - a. Read and understand the following information about reference codes that might be generated when an expansion unit is added while the system is powered on.

Note: When an expansion unit is connected with 12X cables while the system is powered on to standby state, it is normal for error logs and event logs to be created. The number and combinations of these logs is dependent on the configuration.

Logs with the following reference codes can be ignored during this verification step if their time stamps indicate that they occurred during the add procedure and the resources that they refer to were involved in the procedure.

- B7006907: An expansion unit was added that was previously configured on a different system. This reference code appears in an informational log.
 - B7006982: A 12X connection failure was experienced. This reference code appears in an informational log.
 - B7006984: A 12X loop was opened. This reference code appears in an informational log.
 - B7006985: A 12X loop was closed. This reference code appears in an informational log.
 - B70069E6: A 12X link was lost. This reference code appears in an informational log.
 - B70069E7: A 12X link was restored. This reference code appears in an informational log.
- b. Choose one of the following options:
 - If your system is managed by an HMC, see Verifying the 12X loops with an HMC (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_loopverifyconsole.htm).
 - If your system is not managed by an HMC, see Verifying the 12X loops without an HMC (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_loopverifynoconsole.htm).
11. Check for new serviceable events that were generated during the procedure by choosing one of the following options:

Important: Multiple options might apply to your configuration. Choose only the first option that applies.

- a. If your system is managed by an HMC, complete the following steps:
 - 1) In the navigation area, select **Serviceability > Manage Serviceable Events**.
 - 2) Specify event criteria that includes events that were generated during the procedure. In particular, specify the managed system that the expansion unit was added to, and specify a date and time range that includes the data and time that were recorded in step 3 on page 13.
 - 3) Continue with step 12.
 - b. If your system is not managed by an HMC and has a partition that is running the Virtual I/O Server (VIOS) operating system, complete the following steps:
 - 1) Log in as the padmin user. If you need assistance, contact the system administrator.
 - 2) Type `diagmenu -d sysplanar0 -E xx` and press Enter, where xx specifies the number of days in the range 1 - 60. This command returns results for the most recent number of days specified. Specify a range that includes the date and time that was recorded in step 3 on page 13.
 - 3) On the DIAGNOSTIC MODE SELECTION display, highlight the **Problem Determination** option, and press Enter.
 - 4) Continue with step 12.
 - c. If your system is not managed by an HMC and has a partition that is running the AIX operating system, complete the following steps:
 - 1) Log in as the root user, or use the CE login. If you need assistance, contact the system administrator.
 - 2) Type `diag -d sysplanar0 -E xx` and press Enter, where xx specifies the number of days in the range 1 - 60. This command returns results for the most recent number of days specified. Specify a range that includes the date and time that was recorded in step 3 on page 13.
 - 3) On the DIAGNOSTIC MODE SELECTION display, highlight the **Problem Determination** option, and press Enter.
 - 4) Continue with step 12.
 - d. If your system is not managed by an HMC and has a partition that is running the IBM i operating system, the following steps:
 - 1) Sign on with at least service level authority.
 - 2) At the command line of the IBM i session, type `strsst` and press Enter.
 - 3) Type your service tools user ID and service tools password on the System Service Tools (SST) Sign On display, and press Enter. The service tools password is case-sensitive.
 - 4) Click **Start a Service Tool > Hardware Service Manager > Work with Service Action Event Log**.
 - 5) On the **Select Timeframe** display, change the **From Date and Time** field to the wanted date and time range. Specify a range that includes the date and time that was recorded in step 3 on page 13
 - 6) Continue with step 12.
 - e. If your system is not managed by an HMC and has a partition that is running the Linux operating system, the following steps:
 - 1) Log in as the root user. If you need assistance, contact the system administrator.
 - 2) Type `servicelog -query='serviceable=1 AND closed=0 AND time_event>="yyyy-mm-dd"'` and press Enter, where yyyy-mm-dd is the day on which the procedure occurred.
 - 3) Continue with step 12.
12. Handle serviceable events that were generated during the procedure, by completing the following steps:
- a. Specify event criteria that includes events that were generated during the procedure.

- b. Close any serviceable events that are generated during the procedure that include one of the reference codes that are listed in step 10 on page 14.
 - c. Perform problem analysis on any remaining open serviceable events.
13. Power on the system to running state by choosing one of the following options:
- Power on one or more logical partitions. For instructions, see Starting a system or logical partition (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7haj/crustartsys.htm>).
- Note:** Powering on a partition while the system is in standby state changes the system out of standby state into operating or running state.
- Power off the system and then power it back on.
 - For details about powering off the system, see Stopping a system or logical partition (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7haj/crustopsys.htm>).
 - For details about powering on the system, see Starting a system or logical partition (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7haj/crustartsys.htm>).
14. If the system has logical partitions, you can now assign I/O slots in the expansion units that were added to logical partitions.
15. You have completed the addition of the expansion unit.

Connecting expansion units with 12X cables

These instructions explain how to connect expansion units to a system unit with 12X cables in either a new or existing loop.

Tip: If multiple expansion units are to be added with the system power on, work with one 12X loop at a time. For example, if the plan is to add three expansion units to loop X and two to loop Y, follow the complete procedure to simultaneously add the three units to loop X. When the first procedure has been completed, then follow the procedure to simultaneously add the two units to loop Y. It is not necessary to complete the procedure for each expansion unit separately.

After reviewing “Examples: 12X connections” on page 18, select one of the following options:

- If you are creating a new 12X loop with the expansion units you are adding, see “Adding expansion units to a new 12X loop.”
- If you are adding expansion units to an existing 12X loop, see “Adding expansion units to an existing 12X loop” on page 17.

Adding expansion units to a new 12X loop

Use this procedure when adding new expansion units creates a new 12X loop, as opposed to being added to an existing 12X loop.

Unused 12X connectors might be covered with metal clips. Remove the clips before you install the 12X cables.

Complete the following steps to add the expansion units to the 12X configuration:

1. On the system unit where you are connecting new expansion units, identify connectors 0 and 1 of the 12X connector pair.
2. Find the 12X cables that were included with the expansion units.
3. Connect a 12X cable to connector 0 on one of the expansion units to be added.
4. Connect another 12X cable to connector 1 on the expansion unit added in step 3.
5. If there is another expansion unit to add, do the following:
 - a. Choose the cable that you just connected to connector 1 on the expansion unit that you just added. Connect the other end of this cable to connector 0 on the next expansion unit to be added.
 - b. Connect another cable to connector 1 on the expansion unit added in step 5a.

6. Repeat step 5 on page 16 for each additional expansion unit.
7. Connect the other end of the cable from connector 0 on the first expansion unit added to connector 0 on your system unit.
8. Connect the other end of the cable from connector 1 on the last expansion unit added to connector 1 on your system unit.
9. Return to the procedure that sent you here.

Adding expansion units to an existing 12X loop

Use this procedure to add expansion units to an existing 12X loop.

Unused 12X connectors might be covered with metal clips. Remove the clips before you install the 12X cables.

To add expansion units to a 12X configuration, complete the following steps:

1. Identify the 12X loop to which you want to add the expansion units.
2. Identify connectors 0 and 1 of the 12X connector pair on the system unit that is associated with the loop identified in step 1.
3. Find the 12X cables that were included with the expansion units.
4. Choose from the following options:
 - To add the expansion units to the beginning of the loop, go to step 5.
 - To add the expansion units to the middle of the loop, go to step 6.
 - To add the expansion units to the end of the loop, go to step 7.
5. To add the expansion units to the beginning of the loop, complete the following steps:
 - a. Find connector 0 on the first expansion unit on the loop (the one connected to connector 0 on the system unit). Move the 12X cable from connector 0 on this unit to connector 0 on an expansion unit to be added.
 - b. Connect another 12X cable to connector 1 on the expansion unit added in step 5a.
 - c. If you are adding another expansion unit, complete the following steps:
 - 1) Find the other end of the cable that is connected to connector 1 on the expansion unit just added. Connect that end to connector 0 on the next expansion unit to be added.
 - 2) Connect another cable to connector 1 on the expansion unit added in step 5c1.
 - d. Repeat step 5c for each additional expansion unit.
 - e. Find the other end of the cable from connector 1 on the last expansion unit added. Connect that end to connector 0 on the expansion unit that was disconnected in step 5a.
6. To add the expansion units to the middle of the loop, complete the following steps:
 - a. Find connector 0 on the expansion unit that will follow the expansion units that are being added. Move the 12X cable from connector 0 on that unit to connector 0 on an expansion unit to be added.
 - b. Connect another 12X cable to connector 1 on the expansion unit added in step 6a.
 - c. If you are adding another expansion unit, complete the following steps:
 - 1) Find the other end of the cable that is connected to connector 1 on the expansion unit just added. Connect that end to connector 0 on the next expansion unit to be added.
 - 2) Connect another cable to connector 1 on the expansion unit added in step 6c1.
 - d. Repeat step 6c for each additional expansion unit.
 - e. Find the other end of the cable from connector 1 on the last expansion unit added. Connect that end to connector 0 on the expansion unit that was disconnected in step 6a.
7. To add the expansion units to the end of the loop, complete the following steps:
 - a. Move the 12X cable from connector 1 on the system unit to connector 0 on one of the expansion units to be added.

- b. Connect another 12X cable to connector 1 on the expansion unit added in step 7a on page 17.
 - c. If you are adding another expansion unit, complete the following steps:
 - 1) Find the other end of the cable from connector 1 on the expansion unit just added. Connect that end to connector 0 on the next expansion unit to be added.
 - 2) Connect another cable to connector 1 on the expansion unit added in step 7c1.
 - d. Repeat step 7c for each additional expansion unit.
 - e. Connect the other end of the cable from connector 1 on the last expansion unit added to connector 1 on the system unit.
8. Return to the procedure that sent you here.

Examples: 12X connections

See how 12X cables are connected to expansion units and server units to form 12X loops.

Note: The physical connectors on the server units and expansion units might not have labels, or might have labels that are different from what is shown in the examples. In all cases, the top or left 12X connector corresponds to connector 0 in the examples, and the bottom or right 12X connector corresponds to connector 1 in the examples.

1. Connect a server that has one GX adapter to two expansion units, as shown in the following figure.

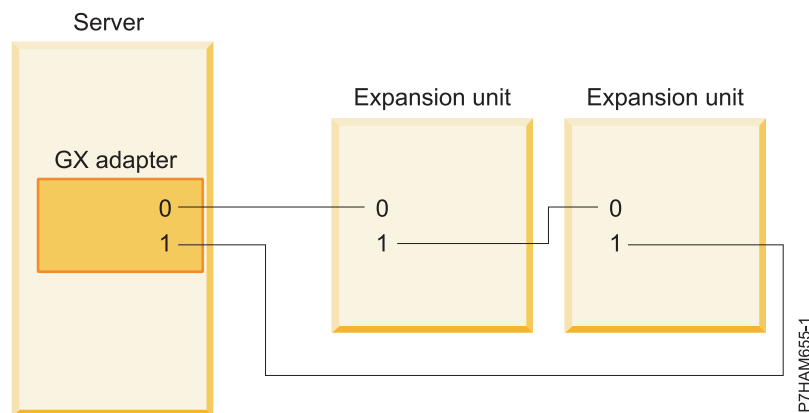


Figure 1. Example of a server that has one GX adapter connected to two expansion units

2. Connect a server to four expansion units using two GX adapters, where each adapter connects to two expansion units, as shown in the following figure.




Figure 2. Example of a server that has two GX adapters, each connected to two expansion units

Connecting expansion units with SPCN cables

Use these instructions to learn how to connect expansion units into the system power control network (SPCN).

- Find the SPCN cables that were included with your expansion units.
- Review “Examples: SPCN connections” on page 20 and then return here and continue with the instructions.

Important: As noted in “Overview for I/O expansion units” on page 1, there is only one SPCN loop per system to which all expansion units are connected.

1. Choose from the following options:
 - To add expansion units when there is not already an SPCN loop, go to step 2.
 - To add expansion units to the beginning of an existing SPCN loop, go to step 3 on page 20.
 - To add expansion units to the middle of an existing SPCN loop, go to step 4 on page 20.
 - To add expansion units to the end of an existing SPCN loop, go to step 5 on page 20.
2. To add expansion units when there is not already an SPCN loop, do the following:
 - a. Connect an SPCN cable to connector SPCN0 on your system unit.
 - b. Connect the other end of the SPCN cable to connector SPCN 0 on one of the expansion units to be added.
 - c. Connect another cable to connector SPCN 1 on the expansion unit added in step 2b.
 - d. If there is another expansion unit to add, do the following:
 - 1) Connect the other end of the cable from connector SPCN 1 on the expansion unit just added to connector SPCN 0 on the next expansion unit.
 - 2) Connect another cable to connector SPCN 1 on the expansion unit added in step 2d1.

- e. Repeat step 2d on page 19 for each additional expansion unit to be added.
 - f. Connect the other end of the cable from connector SPCN 1 on the last expansion unit added to connector SPCN1 on your system unit.
 - g. Return to the procedure that sent you here.
3. To add expansion units to the beginning of an existing SPCN loop, do the following:
 - a. Remove the SPCN cable end from connector SPCN 0 on the first expansion unit (the one connected to connector SPCN0 on the system unit). Reconnect this cable end to connector SPCN 0 on one of the expansion units to be added.
 - b. Connect another SPCN cable to connector SPCN 1 on the expansion unit added in step 3a.
 - c. If there is another expansion unit to add, do the following:
 - 1) Find the other end of the cable attached to connector SPCN 1 on the expansion unit just added. Connect this end to connector SPCN 0 on the next expansion unit to be added.
 - 2) Connect another cable to connector SPCN 1 on the expansion unit added in step 3c1.
 - d. Repeat step 3c for each additional expansion unit.
 - e. Find the other end of the cable attached to connector SPCN 1 on the last expansion unit added. Connect this end to connector SPCN 0 on the expansion unit that was disconnected in step 3a.
 - f. Return to the procedure that sent you here.
 4. To add expansion units to the middle of an existing SPCN loop, do the following:
 - a. Remove the SPCN cable end from connector SPCN 0 on the expansion unit that will follow the expansion units that are being added. Reconnect this cable to connector SPCN 0 on one of the expansion units to be added.
 - b. Connect another SPCN cable to connector SPCN 1 on the expansion unit added in step 4a.
 - c. If there is another expansion unit to add, do the following:
 - 1) Find the other end of the cable attached to connector SPCN 1 on the expansion unit just added. Connect this end to connector SPCN 0 on the next expansion unit to be added.
 - 2) Connect another cable to connector SPCN 1 on the expansion unit added in step 4c1.
 - d. Repeat step 4c for each additional expansion unit.
 - e. Find the other end of the cable attached to connector SPCN 1 on the last expansion unit added. Connect this end to connector SPCN 0 on the expansion unit that was disconnected in step 4a.
 - f. Return to the procedure that sent you here.
 5. To add expansion units to the end of an existing SPCN loop, do the following:
 - a. Move the SPCN cable from connector SPCN 1 on the system unit to connector SPCN 0 on one of the expansion units to be added.
 - b. Connect another SPCN cable to connector SPCN 1 on the expansion unit added in step 5a.
 - c. If there is another expansion unit to add, do the following:
 - 1) Find the other end of the cable attached to connector SPCN 1 on the expansion unit just added. Connect this end to connector SPCN 0 on the next expansion unit to be added.
 - 2) Connect another SPCN cable to connector SPCN 1 on the expansion unit added in step 5c1.
 - d. Repeat step 5c for each additional expansion unit.
 - e. Connect the other end of the cable from connector SPCN 1 on the last expansion unit added to connector SPCN 1 on the system unit.
 - f. Return to the procedure that sent you here.

Examples: SPCN connections

Review examples that show how SPCN cables are connected to expansion units and system units for form SPCN loops.

Note: The physical connectors on the system units and expansion units might not have labels, or might have labels that are different from what is shown in the examples. In all cases, the top or left SPCN connector corresponds to SPCN 0 in the examples, and the bottom or right SPCN connector corresponds to SPCN 1 in the examples.

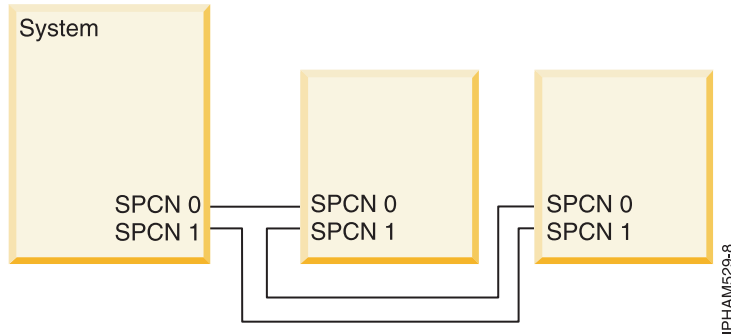


Figure 3. Example of SPCN connections

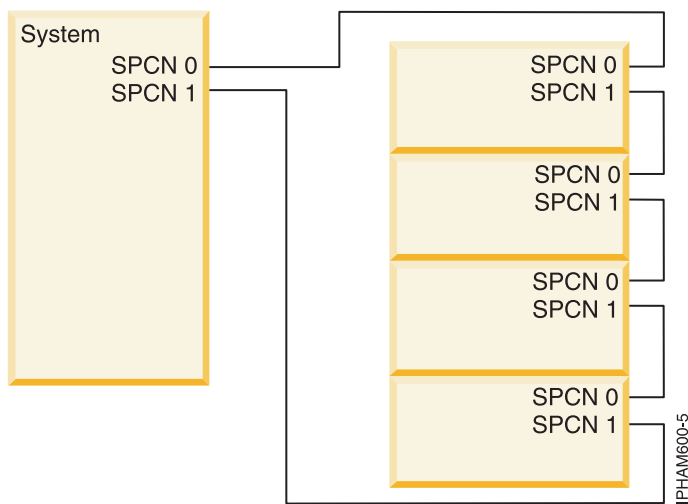


Figure 4. Example of SPCN connections

Verifying that the new configuration is functioning

Use this procedure to verify that your system lists the expansion units in your configuration and that the expansion units are functioning correctly.

1. Choose from the following options:
 - If your system is managed by an HMC, continue with step 2.
 - If your system is not managed by an HMC, see [Verifying the installed part \(http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7haj/p7hajhsmverify.htm\)](http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7haj/p7hajhsmverify.htm). When you are finished, continue with step 4 on page 22.
2. To view more information about the managed system, perform the following steps:
 - a. From the HMC navigation area, expand **Systems Management**.
 - b. Click **Servers**.
 - c. Select the system with which you want to work. Information about the system is shown under the Tasks section in the content area.
3. To verify the new expansion units, complete the following steps:

- a. Select the server for which you want to verify the configuration.
- b. Click **Tasks > Properties**.
- c. Click the **I/O** tab.
- d. Verify that the new expansion units are displayed in the list. It might take several minutes for the new expansion units to be displayed in the list.

Tip: To refresh the list of expansion units, repeat steps 3b and 3c. It is not necessary to exit the **Properties** task before repeating these steps, but doing so reduces the number of active tasks.

4. If the I/O Properties window does not display the new expansion units after 10 minutes, verify the installation by completing the following steps:
 - a. Ensure that the expansion units are powered on and that the power cables are connected to the power source and the expansion units.
 - b. Ensure that the 12X cables are connected correctly. For details, see [Connecting expansion units with 12X cables \(http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_conncable12x_5802.htm\)](http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_conncable12x_5802.htm).
 - c. Ensure that the SPCN cables are connected correctly. For details, see [Connecting expansion units with SPCN cables \(http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_conncablespcn_5802.htm\)](http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_expunit_conncablespcn_5802.htm).
 - d. If the expansion units still are not displayed, contact your next level of support for assistance.
5. Return to the procedure that sent you here.

Verifying the 12X loops with an HMC

If your system is managed by a Hardware Management Console (HMC), learn how to verify that the 12X loops have no failed links.

To verify that the 12X loops do not have failed links on a system that is managed by an HMC, complete the following steps with the system powered on:

1. Verify that no serviceable events include B70069xx reference codes (where x is any character 0 - 9 or A - F) by performing the following steps:
 - a. In the navigation area of the HMC, select **Service Management**.
 - b. In the content area, select **Manage Serviceable Events**.
 - c. In the Manage Serviceable Events window, perform the following steps:
 - 1) Click **Open** for the **Serviceable Event Status** field.
 - 2) Select the machine type, model, and serial number (MTMS) of the server you are working with for the **Reporting MTMS** field.
 - 3) Click **ALL** for all other fields, and click **OK**.
 - d. Scan for any serviceable events that include B70069xx reference codes, and perform one of the following actions:
 - If there are no serviceable events that include B70069xx reference codes, close the Reference Code History window.
 - If there are serviceable events that include B70069xx reference codes, there might be problems with the 12X loops. Perform problem analysis on the serviceable events to correct the problems before continuing. If you need assistance, contact your service provider.
2. Verify the status of the links in the 12X loops by performing the following steps:
 - a. From the navigation bar, expand **Systems Management**.
 - b. Click **Servers**.
 - c. Select the server you are working with.
 - d. From the Tasks area, complete the following steps:
 - 1) Expand **Hardware Information**.
 - 2) Click **View Hardware Topology**.

- e. Scan the Current Hardware Topology section of the Hardware Topology data to identify rows with links that must be repaired. A row must be repaired if either the Leading Port Status column or Trailing Port Status column contains a value other than Operational.

Exception: Exclude any rows that represent nodes belonging to a loop with no I/O expansion units connected. The Hardware Topology data displays a loop with no I/O expansion units as a bus with only one associated row. That row contains a Node Type of Local NIC, a Leading Port Status of Open/Failed, and Trailing Port Status of Open/Failed.

- f. If you identified links in step 2e with a status other than Operational, the links must be repaired before you continue. To repair the links, repeat step 1 on page 22. If you already performed step 1 on page 22 or need assistance, contact your next level of support.
3. Return to the procedure that sent you here.

Verifying the 12X loops without an HMC

If your system is not managed by a Hardware Management Console (HMC), learn how to use the IBM i, AIX, Virtual I/O Server (VIOS), or Linux operating system to verify that the 12X loops have no failed links.

To verify that the 12X loops do not have failed links on a system that is not managed by an HMC, complete the following steps with the system powered on:

1. Choose the operating system that best applies to your configuration.

Important: Multiple options might apply to your configuration. Choose only the first option that applies.

- If your system has a partition running VIOS operating system, continue with step 2.
 - If your system has a partition running the AIX operating system, continue with step 3.
 - If your system has a partition running the IBM i operating system, continue with step 4 on page 24.
 - If your system has a partition running the Linux operating system, continue with step 6 on page 25.
2. Use the VIOS error log to verify that there are no B70069xx errors (where *x* is any character 0 - 9 or A - F) by completing the following steps.
 - a. Log in as the padmin user. If you need assistance, contact the system administrator.
 - b. Type `diagmenu -d sysplanar0 -E xx` and press Enter, where *xx* is any number 1 - 60 that specifies a number of days. This command returns results for the most recent number of days specified. The suggested range is 30 days.
 - c. Highlight the **Problem Determination** option in the DIAGNOSTIC MODE SELECTION display. Press Enter.
 - d. Search the results that are displayed to identify problems with B70069xx reference codes. You might need to scroll down in the display to see all of the results.

Note: If the Problem Determination results contain serviceable events that have already been reported, a PREVIOUS DIAGNOSTICS RESULTS display appears. Respond to the **Do you want to review these previously reported errors?** prompt by highlighting the YES response, and then press Enter.

- e. Based on the results of your search, perform one of the following options:
 - If there are no problems with B70069xx reference codes, press Enter to return to the command line.
 - If there are problems with B70069xx reference codes, there might be problems with the 12X loops. Perform problem analysis on the problems to correct them before continuing. If you need assistance, contact your service provider.
 - f. Go to step 7 on page 25.
3. Use the AIX error log to verify that there are no B70069xx errors (where *x* is any character 0 - 9 or A - F) by completing the following steps.

- a. Log in as the padmin user. If you need assistance, contact the system administrator.
- b. Type `diag -d sysplanar0 -E xx` and press Enter, where `xx` is any number 1 - 60 that specifies a number of days. This command returns results for the most recent number of days specified. The suggested range is 30 days.
- c. Highlight the **Problem Determination** option in the DIAGNOSTIC MODE SELECTION display. Press Enter.
- d. Search the results that are displayed to identify problems with B70069xx reference codes. You might need to scroll down in the display to see all of the results.

Note: If the Problem Determination results contain serviceable events that have already been reported, a PREVIOUS DIAGNOSTICS RESULTS display appears. Respond to the **Do you want to review these previously reported errors?** prompt by highlighting the YES response, and then press Enter.

- e. Based on the results of your search, perform one of the following options:
 - If there are no problems with B70069xx reference codes, press Enter to return to the command line.
 - If there are problems with B70069xx reference codes, there might be problems with the 12X loops. Perform problem analysis on the problems to correct them before continuing. If you need assistance, contact your service provider.
 - f. Go to step 7 on page 25.
4. Use the IBM i service action event log to verify that there are no B70069xx errors (where `x` is any character 0 - 9 or A - F) by completing the following steps.
 - a. Sign on with at least service level authority.
 - b. On the command line of the IBM i session, type `strsst` and press Enter.
 - c. Type your service tools user ID and service tools password on the System Service Tools (SST) Sign On display. Press Enter.

Tip: The service tools password is case-sensitive.

- d. Click **Start a Service Tool > Hardware Service Manager > Work with Service Action Event Log**.
 - e. On the **Select Timeframe** display, change the **From: Date and Time** to the desired date and time range. The suggested range is 30 days.
 - f. Search for serviceable events with B70069xx reference codes, and perform one of the following actions:
 - If there are no serviceable events with B70069xx reference codes, press F3 (Exit) to return to the Hardware Service Manager display. Go to step 5
 - If there are serviceable events with B70069xx reference codes, there might be problems with the 12X loops. Perform problem analysis on the serviceable events to correct the problems before continuing. If you need assistance, contact your service provider. Go to step 7 on page 25
5. Use the Hardware Service Manager to verify the statuses of the links in the 12X loops by completing the following steps:
 - a. Type your service tools user ID and service tools password on the System Service Tools (SST) Sign On display. Press Enter. The service tools password is case-sensitive.
 - b. Select **Logical hardware resources** from the Hardware Service Manager display. Press Enter.
 - c. Select **High-speed link resources** from the Logical Hardware Resources display. Press Enter.
 - d. Select **Display port information** for the first 12X loop. Press Enter.
 - e. Scan the port information screen to identify rows with links that must be repaired. A row must be repaired if the connection status is a value other than Operational.

Exception: Do not include any rows belonging to a loop with no I/O expansion units connected. The port information display shows a loop with no I/O expansion units as a single row that contains a connection status of Failed.

- f. Choose one of the following options:
 - 1) If you completed step 5e on page 24 without identifying any links to be repaired, press F3 to exit.
 - 2) If you completed step 5e on page 24 with one or more links that are not Operational, then the rows that you identified contain failed or degraded links in the 12X configuration. The links must be repaired before you continue. To repair the links, repeat step 5e on page 24. If you already performed step 5e on page 24 or need assistance, contact your next level of support.
- g. Go to step 7.
6. Use the Linux error log to verify that there are no B70069xx errors (where x is any character 0 - 9 or A - F) by completing the following steps.
 - a. Log in as the root user. If you need assistance, contact the system administrator.
 - b. To list every platform error with a reference code that starts with B70069, type `service log --query='refcode like "B70069%" AND serviceable=1 AND closed=0'` and press Enter.
 - c. Search the results that are displayed for problems with B70069xx reference codes and for statuses that are open. If there are problems with B70069xx reference codes, there might be problems with the 12X loops. Perform problem analysis on the problems to correct them before continuing. If you need assistance, contact your service provider.

Note: You might need to scroll down in the display to see all of the results.
- d. Go to step 7.
7. Return to the procedure that sent you here.

Verifying the system power control network

Check for system power control network (SPCN) errors and inactive enclosures, and fix machine type, model, and serial number (MTMS) problems.

1. Choose one the following options to verify the SPCN:
 - If your system is not managed by a Hardware Management Console (HMC), continue with step 2.
 - If your system is managed by an HMC, continue with step 3.
2. Choose one the following options:

Important: Multiple options might apply to your configuration. Choose only the first option that applies.

- If your system is running the Virtual I/O Server (VIOS) operating system, continue with step 4 on page 26.
 - If your system is running the AIX operating system, continue with step 5 on page 26.
 - If your system is running the IBM i operating system, continue with step 6 on page 27.
 - If your system is running the Linux operating system, continue with step 7 on page 27.
3. Use the HMC to verify that no serviceable events include 10009xxx reference codes (where x is any character 0 - 9 or A - F) by completing the following steps:
 - a. In the navigation area, select **Service Management**.
 - b. In the content area, select **Manage Serviceable Events**.
 - c. In the Manage Serviceable Events window, perform the following steps:
 - 1) Click **Open** for the **Serviceable Event Status** field.
 - 2) Select the machine type, model, and serial number (MTMS) of the server you are working with for the **Reporting MTMS** field.
 - 3) Click **ALL** for all other fields and click **OK**.

- d. Scan for any serviceable events that include 10009xxx reference codes, and perform one of the following actions:
 - If there are no serviceable events with 10009xxx reference codes, click **Cancel** twice to exit Manage Serviceable Events.
 - If there are serviceable events with 10009xxx reference codes, there might be problems with the SPCN loop. Perform problem analysis on the serviceable events to correct the problems before continuing. If you need assistance, contact your service provider.
- e. Go to step 8 on page 27.
4. Use the VIOS error log to verify that there are no 10009xxx errors (where *x* is any character 0 - 9 or A - F) by completing the following steps.
 - a. Log in as the padmin user. If you need assistance, contact the system administrator.
 - b. Type `diagmenu -d sysplanar0 -E xx` and press Enter, where *xx* is any number 1 - 60 that specifies a number of days. This command returns all 10009xxx error log entries for the most recent number of days specified. The suggested range is 30 days.
 - c. Highlight the **Problem Determination** option in the DIAGNOSTIC MODE SELECTION display. Press Enter.
 - d. Search the results that are displayed to identify problems with 10009xxx reference codes. You might need to scroll down in the display to see all of the results.

Note: If the Problem Determination results contain serviceable events that have already been reported, the PREVIOUS DIAGNOSTICS RESULTS display is shown. Respond to the **Do you want to review these previously reported errors?** prompt by highlighting the YES response, and then press Enter.

- e. Based on the results of your search, perform one of the following options:
 - If there are no problems with 10009xxx reference codes, press Enter to return to the command line.
 - If there are problems with 10009xxx reference codes, there might be problems with the SPCN loop. Perform problem analysis on the problems to correct them before continuing. If you need assistance, contact your service provider.
- f. Go to step 8 on page 27.
5. Use the AIX error log to verify that there are no 10009xxx errors (where *x* is any character 0 - 9 or A - F) by completing the following steps.
 - a. Log in as the root user, or use the CE login. If you need assistance, contact the system administrator.
 - b. Type `diag -d sysplanar0 -E xx` and press Enter, where *xx* is any number 1 - 60 that specifies a number of days. This command returns all 10009xxx error log entries for the most recent number of days specified. The suggested range is 30 days.
 - c. Highlight the **Problem Determination** option in the DIAGNOSTIC MODE SELECTION display. Press Enter.
 - d. Search the results that are displayed to identify problems with 10009xxx reference codes. You might need to scroll down in the display to see all of the results.

Note: If the Problem Determination results contain serviceable events that have already been reported, the PREVIOUS DIAGNOSTICS RESULTS display is shown. Respond to the **Do you want to review these previously reported errors?** prompt by highlighting the YES response, and then press Enter.

- e. Based on the results of your search, perform one of the following options:
 - If there are no problems with 10009xxx reference codes, press Enter to return to the command line.

- If there are problems with 10009xxx reference codes, there might be problems with the SPCN loop. Perform problem analysis on the problems to correct them before continuing. If you need assistance, contact your service provider.
- f. Go to step 8.
6. Use the IBM i service action event log to verify that there are no 10009xxx errors (where *x* is any character 0 - 9 or A - F) by completing the following steps.
 - a. Sign on with at least service level authority.
 - b. At the command line of the IBM i session, type `strsst` and press Enter.
 - c. Type your service tools user ID and service tools password on the System Service Tools (SST) Sign On display. Press Enter. The service tools password is case-sensitive.
 - d. Select **Start a Service Tool > Hardware Service Manager > Work with Service Action Event Log**.
 - e. On the **Select Timeframe** display, change the **From: Date and Time** to the desired date and time range. The suggested range is 30 days.
 - f. Search for serviceable events with 10009xxx reference codes, and perform one of the following actions:
 - If there are no serviceable events with 10009xxx reference codes, press F3 (Exit) to return to the Hardware Service Manager display.
 - If there are serviceable events with 10009xxx reference codes, there might be problems with the SPCN loop. Perform problem analysis on the serviceable events to correct the problems before continuing. If you need assistance, contact your service provider.
 - g. Go to step 8.
 7. Use the Linux error log to verify that there are no 10009xxx errors (where *x* is any character 0 - 9 or A - F) by completing the following steps.
 - a. Log in as the root user. If you need assistance, contact the system administrator.
 - b. Type `service log --query='refcode like "10009%" AND serviceable=1 AND closed=0'` and press Enter
 - c. Search the results that are displayed for problems with 10009xxx reference codes and statuses that are open. If there are problems with 10009xxx reference codes, there might be problems with the SPCN loop. Perform problem analysis on the problems to correct them before continuing. If you need assistance, contact your service provider.

Note: You might need to scroll down in the display to see all of the results.

- d. Go to step 8.
8. Verify that no inactive expansion units are in the SPCN by performing the following steps:
 - a. Access the ASMI by using an authority level of administrator or authorized service provider. For details about using the ASMI, see *Managing the Advanced System Management Interface* (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hby/asmi.htm>).
 - b. In the ASMI navigation area, expand **System Configuration** and click **Configure I/O Enclosures**.
 - c. Scan for inactive expansion units by looking for rows that show *Inactive* in the Status column. If inactive expansion units exist, perform the following steps:

Important: If you are here because the procedure to concurrently add expansion units instructed you to perform this task following steps to connect the power cables and power on the expansion units, do not perform the following steps to clear the inactive enclosures. Contact your service provider for assistance.

- 1) Click **Clear inactive enclosures**.
 - 2) Wait for the completion indication, and repeat steps 8b and 8c. If inactive expansion units remain, do not continue. Contact your service provider for assistance.
9. Check for temporary machine type, model, and serial number (MTMS) values assigned to the expansion units connected to the system by completing the following steps:

- a. Access the ASMI by using an authority level of administrator or authorized service provider. For details about using the ASMI, see *Managing the Advanced System Management Interface* (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hby/asmi.htm>).
 - b. In the ASMI navigation area, expand **System Configuration** and click **Configure I/O Enclosures**.
 - c. Review the display to determine if there are any expansion units that require correction. If you find one or more rows that contain a Location code value of UTMPx.xxx.xxxxxxx, where x is any character 0 - 9 or A - Z, continue to step 10 to provide the necessary corrections. If there are no rows with Location codes that require correction, continue to step 11 on page 29.
10. For each expansion unit found in your review, complete the following steps to correct the machine type, model, and serial number (MTMS) values. After you have completed the corrections required for all expansion units identified in your review, continue to step 11 on page 29.

Important: If your expansion unit contains one or more I/O slots that are owned by active partitions running the VIOS, AIX, or Linux operating system, these partitions will not automatically recognize the MTMS value corrections that you are about to make. For these partitions to recognize the new values, you will be required to take additional action that is potentially disruptive. Before proceeding with this step, review step 10e to determine whether it is possible to take the required actions at this time. If it is not, continue to step 11 on page 29 without performing this step.

- a. Access the ASMI by using an authority level of administrator or authorized service provider. For details about using the ASMI, see *Managing the Advanced System Management Interface* (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hby/asmi.htm>).
- b. In the ASMI navigation area, expand **System Configuration** and click **Configure I/O Enclosures**.
- c. Select the row associated with the expansion unit to be corrected and click **Change settings**. Follow these steps to update the information:
 - 1) Enter the correct type and model in the **Type-Model** field.
 - 2) Enter the correct serial number in the **Serial number** field.
 - 3) Do not update the **Power Control Network Identifier** field.
 - 4) Click **Save settings**.
- d. In the ASMI navigation area, expand **System Configuration** and click **Configure I/O Enclosures** again. Verify that the values you entered are displayed in the information for the expansion unit that you selected in step 10c. If the values that you entered are not displayed, wait for a few minutes and repeat this step. If the problem persists, contact your service provider for assistance.

Important: Do not click **Back** for the purpose of verifying your values.

- e. If the expansion unit that you corrected contains one or more I/O slots that are owned by active partitions running the AIX, VIOS, or Linux operating system, the corrections will not be recognized by those partitions until you take one of the following actions:
 - Power off and power on the expansion unit. To perform this task, select one of the following options:
 - If the system is managed by an HMC, click **Serviceability > Hardware > Power On/Off Unit** to perform this task.
 - If the system is not managed by an HMC, you must choose one of the other options in this list.
 - Power off and power on active partitions that own I/O slots in the expansion unit that you corrected and that are running the AIX, VIOS, or Linux operating system.
 - Dynamically remove I/O slots in the expansion unit that you corrected from active partitions that are running the AIX, VIOS, or Linux operating system, and then dynamically add the slots back to those partitions. To perform this task, select one of the following options:
 - If the system is managed by an HMC, click **Dynamic Logical Partitioning > Physical Adapters** to perform this task.

- If the system is not managed by an HMC, you must choose one of the other options in this list.
 - Power off and power on the server.
11. Return to the procedure that sent you here.

Setting I/O enclosure configuration ID and MTMS value

The configuration ID and the machine type, model, and serial number (MTMS) values are used by the system to identify characteristics of an expansion unit. If these values are not set correctly, the system might not manage the expansion unit properly.

To set the configuration ID, use the Advanced System Management Interface (ASMI). However if the ASMI is not available you can use the physical control panel to set the configuration ID. To set the MTMS value you must use the ASMI, not the control panel. However, if you do not have access to the ASMI, the system still operates without updating this information. The MTMS value should match the original value of the enclosure, which can be found on a label affixed to the enclosure.

Using the ASMI to verify and set the configuration ID and MTMS value:

You can use the Advanced System Management Interface (ASMI) to verify and set certain values for an expansion unit. These values include the machine type, model, and serial number (MTMS) and the configuration ID.

To perform this operation, verify that the following prerequisites have been met:

- The system must be powered on to firmware standby or firmware running state.
- The expansion unit must have power and be correctly installed in the system power control network.
- Your authority level must be one of the following:
 - Administrator
 - Authorized service provider

Complete the following steps to verify and set values for the configuration ID and MTMS.

1. Log in to ASMI.
2. Expand **System Configuration**.
3. Click **Configure I/O Enclosures > Clear Inactive Enclosures**.
4. If you were directed here from a part replacement procedure, you must restore information that was contained in the nonvolatile storage of the original part. The expansion unit MTMS must be restored, and you might also be required to set or change the configuration ID (power control network identifier).

Note: In the replacement part, the nonvolatile storage that contains the expansion unit MTMS value is uninitialized. Following the Clear Inactive Enclosures task, the system detected the uninitialized value and replaced it with an obvious, unique value. The replacement MTMS value is formatted as `TMPx.xxx.xxxxxxx`, where `x` is any character 0 - 9 or A - Z. The expansion unit location code, which you must use to select the expansion unit for operations in this task, became `UTMPx.xxx.xxxxxxx`.

5. From the ASMI utility, expand **System Configuration**.
6. Click **Configure I/O Enclosures**.
7. Verify the configuration ID and MTMS data by completing the following steps:
 - a. Compare the Power Control Network Identifier value shown for the expansion unit you are working with to the following Power Control Network Identifier (configuration ID) values:
 - 0x84 for 5096 and 5296 expansion units
 - 0x88 for 7311-D11 and 5790 expansion units
 - 0x89 for 5088 and 0588 expansion units

- 0x8A for 5094, 5294, 5096, and 5296 expansion units
 - 0x8B for 5095 and 0595 expansion units
 - 0x8C for 7311-D20 expansion units
 - 0x8D for 7314-G30 and 5796 expansion units
 - 0x8E for 5802 and 5877 expansion units
- b. Compare the Type-Model and Serial number column values for the expansion unit you are working with to the type, model, and serial values on the expansion unit label. The expansion unit location code displayed in the Location code column appears in one of the following formats:
- *UTMPx.xxx.xxxxxxx*, where *x* is any character 0 - 9 or A - Z.
 - *TTTT.MMM.SSSSSS*, where *TTTT* is the enclosure type, *MMM* is the model, and *SSSSSS* is the serial number.

Note: A 10009132 informational log is produced during an initial program load (IPL), or if a temporary MTMS was assigned to an expansion unit that was added concurrently. The temporary MTMS causes the expansion unit to have a location code of *UTMPx.xxx.xxxxxxx*, where *x* is any character 0 - 9 or A - Z.

- c. If any changes must be made, go to step 8. Otherwise, go to step 18 on page 31.
8. If the system is powered on to the firmware running state, go to step 9. If the system is powered on to the firmware standby state, go to step 10.
9. Choose from the following options:
- If the configuration ID must be changed, perform the following steps:
 - a. Read the following notes.
 - b. Power off the I/O expansion unit.
 - c. Go to step 10.
 - If the MTMS data must be changed and any expansion unit resources are owned by active system partitions that are running the AIX or Linux operating system, perform the following steps:
 - a. Read the following notes.
 - b. Power off the I/O enclosure.
 - c. Go to step 10.
 - If none of the previous options apply, go to step 10.

Notes:

- Do not disconnect the power cables after powering off the expansion unit.
 - If the expansion unit does not display immediately on the service utility that is used to power off the expansion unit, refresh the utility periodically for up to 10 minutes. If the expansion unit does not display during this timeframe, go to step 1 on page 29 and repeat this procedure.
 - If you were directed here from a part replacement procedure, remember to use the new *UTMPx.xxx.xxxxxxx* location code when you select the expansion unit to power off.
 - If the procedure that directed you here involved powering off the expansion unit and while powering off the expansion unit you were instructed to use panel function 69 to power on the expansion unit, perform panel function 69 now (with the control panel set to manual mode) from the control panel of the logical partition where you were instructed to use panel function 69. Do this even though the expansion unit is already on.
10. From the ASMI utility, expand **System Configuration**.
11. Click **Configure I/O Enclosure > the expansion you are working with > Change settings**.
12. If in step 7 on page 29 you determined that the Power Control Network Identifier (configuration ID) value was not correct, enter the correct value now.

13. If in step 7 on page 29 you determined that the Type-Model and Serial number values were not correct, enter the correct values now.

Note: Serial numbers are case-sensitive. All alphabetic characters contained in the serial number must be entered as a capital letter.

14. Click **Save settings** to complete the operation.
15. Verify that the values you just entered are reflected in the Power Control Network Identifier, Type-Model, Serial number, and Location code columns for the expansion unit you are working with by expanding **System Configuration > Configure I/O Enclosures**. Do not use the **Back** button to do this.
16. If the system is powered on to firmware standby state and you entered a new power control network identifier (configuration ID) in step 12 on page 30, the expansion unit will power off and back on automatically. If this is the case, go to step 18. Otherwise, go to step 17.
17. If you powered off the expansion unit in step 9 on page 30, read the following notes to power on the expansion unit.

Notes:

- If the system is managed by an HMC, power on the expansion unit by using the **Serviceability > Hardware > Power On/Off Unit** task on the HMC. If the values you just entered are not immediately reflected in the location code of the expansion unit in the Power On/Off Unit task, restart the task periodically for up to 10 minutes until the values you entered are shown.
 - If the system is not managed by an HMC, disconnect all power to the expansion unit by disconnecting the cables from the power supplies on the expansion unit. Wait for the power indicator light on the expansion unit control panel to go off, wait an additional 30 seconds, and then reconnect the power cables. The expansion unit will power on automatically.
18. Log off and close the ASMI.
 19. Return to the procedure that sent you here.

Updating the SPCN firmware

If a system power control network (SPCN) firmware update is required but is not started automatically due to the current SPCN firmware update policy, the SPCN firmware update must be started manually.

Notes:

- Do not perform maintenance on an expansion unit or change the SPCN network while the SPCN firmware update is being performed.
- If you power off the system while performing firmware updates, the SPCN firmware update is interrupted and must be restarted.
- This procedure can take several hours depending on the number of expansion units on the SPCN loop and their SPCN firmware levels.

To change the SPCN firmware update policy to expanded and start an SPCN firmware update, perform the following steps:

1. Access the Advanced System Management Interface (ASMI).
2. Click **System Configuration > Configure I/O Enclosures**.
3. Record the current SPCN Firmware Update Policy setting so that it can be restored later.
4. Change the **SPCN Firmware Update Policy** setting to **Expanded**, and click **Save Policy Setting** to allow SPCN firmware updates to be performed over the serial SPCN interfaces.
5. Click **Start SPCN Firmware Update**. The SPCN firmware is then downloaded to the expansion units that require an update.
6. Change the SPCN Firmware Update Policy setting back to the setting that was recorded in step 3, and click **Save Policy Setting**.

Notes:

- To monitor the progress of the SPCN firmware update, click **Configure I/O Enclosures**. The window is refreshed. Do not use the browser **Back** or **Refresh** button to monitor the update progress. The Power Control Network Firmware Update Status column shows the percentage completed, and In Progress is displayed while the download operation is in progress. Not Required is displayed when the download process is completed.
- To stop the SPCN firmware update (not recommended), click **Stop SPCN Firmware Update**.

Configuring I/O expansion units

Read information that has been provided for configuring input/output (I/O) expansion unit models.

Configuring the disk subsystem in the 5802 expansion unit

Learn how to configure the SAS disk subsystem in the 5802 expansion unit.

Note: For 5802 expansion unit connector locations referred to in this procedure, see “Model 5802 and 5877 connector locations” on page 141.

1. The SAS disk subsystem in the 5802 expansion unit can hold up to 18 disk drives. The disks in this subsystem can be organized in several different configurations depending on the operating system used, the type of SAS adapter, and the position of the hard-disk-drive partitions switch.

Notes:

- The hard-disk-drive partitions switch is located at the rear of the 5802 expansion unit, just below the SAS connectors.
- If the hard-disk-drive partitions switch position is changed, you must power the 5802 expansion unit off and on in order for the 5802 to sense the new position of the hard-disk-drive partitions switch.

The SAS disk subsystem in the 5802 expansion unit is supported by the following operating systems:

- AIX
- IBM i
- Linux
- VIOS

To determine the level of software you need to support the SAS disk subsystem in the 5802 expansion unit, see IBM Prerequisite (http://www-912.ibm.com/e_dir/eServerPrereq.nsf).

Table 5. Positions of the hard-disk-drive toggle switch

Operating system	Position 1	Position 2	Position 4
AIX	<p>Options:</p> <ul style="list-style-type: none"> • One FC 5901 adapter (PCIe Dual - x4 SAS) • One FC ESA1 adapter (PCIe2 RAID SAS Dual-port 6 Gb) • Two FC 5903/5805 adapters (PCIe Dual - x4 3 Gb SAS RAID) • Two FC ESA1 adapters (PCIe2 RAID SAS Dual-port 6 Gb) • Two FC 5913 or two FC ESA3 adapters (PCIe2 1.8 GB cache RAID SAS Tri-port 6 Gb) <p>Note: The top ports of the FC 5913 or the FC ESA3 pair must attach by a high-density (HD) SAS AA cable.</p>	<p>Options:</p> <ul style="list-style-type: none"> • Two FC 5901 adapters (PCIe Dual - x4 SAS) • Two FC 5903/5805 adapters (PCIe Dual - x4 3 Gb SAS RAID) • Two FC ESA1 adapters (PCIe2 RAID SAS Dual-port 6 Gb) • Two FC 5913 or two FC ESA3 adapters (PCIe2 1.8 GB cache RAID SAS Tri-port 6 Gb) <p>Note: The top ports of the FC 5913 or the FC ESA3 pair must attach by an HD SAS AA cable.</p>	<p>Options:</p> <ul style="list-style-type: none"> • Four single FC 5901 adapters (PCIe Dual - x4 SAS) • Four single FC ESA1 adapters (PCIe2 RAID SAS Dual-port 6 Gb)
IBM i	Not supported	<p>Options:</p> <ul style="list-style-type: none"> • Two single FC 5901 adapters (PCIe Dual - x4 SAS) • Two FC 5903/5805 adapters (PCIe Dual - x4 3 Gb SAS RAID) • Two FC ESA1 adapters (PCIe2 RAID SAS Dual-port 6 Gb) • Two FC 5913 or two FC ESA3 adapters (PCIe2 1.8 GB cache RAID SAS Tri-port 6 Gb) <p>Note: The top ports of the FC 5913 or the FC ESA3 pair must attach by an HD SAS AA cable.</p>	Not supported

Table 5. Positions of the hard-disk-drive toggle switch (continued)

Operating system	Position 1	Position 2	Position 4
Linux	<p>Options:</p> <ul style="list-style-type: none"> • One FC 5901 adapter (PCIe Dual - x4 SAS) • One FC ESA1 adapter (PCIe2 RAID SAS Dual-port 6 Gb) • Two FC 5903/5805 adapters (PCIe Dual - x4 3 Gb SAS RAID) • Two FC ESA1 adapters (PCIe2 RAID SAS Dual-port 6 Gb) • Two FC 5913 or two FC ESA3 adapters (PCIe2 1.8 GB cache RAID SAS Tri-port 6 Gb) <p>Note: The top ports of the FC 5913 or the FC ESA3 pair must attach by an HD SAS AA cable.</p>	<p>Options:</p> <ul style="list-style-type: none"> • Two FC 5901 adapters (PCIe Dual - x4 SAS) • Two FC 5903/5805 adapters (PCIe Dual - x4 3 Gb SAS RAID) • Two FC ESA1 adapters (PCIe2 RAID SAS Dual-port 6 Gb) • Two FC 5913 or two FC ESA3 adapters (PCIe2 1.8 GB cache RAID SAS Tri-port 6 Gb) <p>Note: The top ports of the FC 5913 or the FC ESA3 pair must attach by an HD SAS AA cable.</p>	<p>Options:</p> <ul style="list-style-type: none"> • Four single FC 5901 adapters (PCIe Dual - x4 SAS) • Four single FC ESA1 adapters (PCIe2 RAID SAS Dual-port 6 Gb)
<p>Note: FC means feature code.</p>			

Tip: If you are planning to configure RAID arrays, ensure that you have the minimum number of available disks for each RAID level:

RAID 0

Minimum of one drive per array

RAID 5

Minimum of three drives per array

RAID 6

Minimum of four drives per array

RAID 10

Minimum of two drives per array

For additional SAS RAID controller information for the AIX operating system, see SAS RAID controller for AIX (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ebj/p7ebjkickoff.htm>).

For additional SAS RAID controller information for the IBM i operating system, see SAS RAID controllers for IBM i (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ed5/p7ed5kickoff.htm>).

For additional SAS RAID controller information for the Linux operating system, see SAS RAID controller for Linux (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ebk/p7ebkkickoff.htm>).

For information related to device parity protection, see Device parity protection (<http://publib.boulder.ibm.com/infocenter/systems/scope/i5os/topic/rzaly/rzalydpy.htm>).

Cabling the 5802 disk subsystem

Learn how to connect the disk subsystem in the 5802 expansion unit.

To cable the disk subsystem in the 5802 expansion unit, complete the following steps.

For more information related to SAS cabling and cabling configurations, see Serial-attached SCSI cable planning (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7had/p7hadsascabbling.htm>).

1. Select one of the following configurations to cable the disk subsystem in the 5802 expansion unit.

Note: When you connect any multi-initiator configuration that includes dual SAS adapters, ensure that all of the same disk devices are attached to both adapters and that the disk devices are cabled to the same numbered ports on each adapter.

- To cable a configuration that includes two hard-disk-drive partitions, nine disks per hard-disk-drive partition, and a dual SAS adapter pair with connectivity to all drives, go to step 2 on page 36.

Notes:

- This is a multi-initiator high-availability configuration.
- This configuration can be used with the following operating systems: AIX, IBM i, Linux
- To cable a configuration that includes one hard-disk-drive partition, 18 disks, and one SAS adapter, go to step 3 on page 36.

Note: This configuration can be used with the following operating systems: AIX, Linux

- To cable a configuration that includes two hard-disk-drive partitions, nine disks per hard-disk-drive partition, and one SAS adapter, go to step 4 on page 37.

Note: This configuration can be used with the following operating systems: AIX, IBM i, Linux

- To cable a configuration that includes two hard-disk-drive partitions, nine disks per hard-disk-drive partition, and two SAS adapters, go to step 5 on page 38.

Note: This configuration can be used with the following operating systems: AIX, IBM i, Linux

- To cable a configuration that includes two hard-disk-drive partitions, nine disks per hard-disk-drive partition, a dual SAS adapter pair per hard-disk-drive partition that controls the disks in the partition, and dual SAS connectivity to all drives, go to step 6 on page 39.

Notes:

- This is a multi-initiator high-availability configuration.
- This configuration can be used with the following operating systems: AIX, IBM i, Linux
- To cable a configuration that includes one hard-disk-drive partition, 18 disks, and a dual SAS adapter pair with connectivity to all drives, go to step 7 on page 39.

Notes:

- This is a multi-initiator high-availability configuration.
- This configuration can be used with the following operating systems: AIX, Linux
- To cable a configuration that includes four hard-disk-drive partitions, 18 disks, and one SAS adapter per hard-disk-drive partition, go to step 8 on page 40.

Notes:

- The 18 disks in this configuration are divided among the four hard-disk-drive partitions as follows:
 - Partition 1: Disk drives 01 - 05.
 - Partition 2: Disk drives 06 - 09.

- Partition 3: Disk drives 10 - 14.
 - Partition 4: Disk drives 15 - 18.
 - This configuration is optimized for boot partitions.
 - This configuration can be used with the following operating systems: AIX, Linux
2. Complete the following steps to cable a configuration that includes two hard-disk-drive partitions, nine disks per hard-disk-drive partition, and a dual SAS adapter pair with connectivity to all drives.

Notes:

- For the following operating systems, this must be a RAID configuration: AIX, Linux
 - You must have a minimum of one disk drive per hard-disk-drive partition.
- a. Set the hard-disk-drive partitions switch to position 2. The hard-disk-drive partitions switch is located at the rear of the 5802 expansion unit, just below the SAS connectors.
- b. Choose one of the following SAS adapter and cable combination options:
- Two FC ESA1 adapters (PCIe2 RAID SAS Dual-port 6 Gb) cabled with four FC 3689 cables (AT 0.6 meter - HD SAS)
 - Two FC 5805 or FC 5903 adapters (PCIe Dual - x4 3 Gb SAS RAID) cabled with four FC 3688 cables (AT 0.6 meter - Mini SAS)
 - Two FC 5901 adapters (PCIe Dual - x4 SAS) cabled with four FC 3688 cables (AT 0.6 meter - Mini SAS)

Note: Dual FC 5901 adapters are not supported in IBM i.

- Two FC 5913 or two FC ESA3 adapters (PCIe2 1.8 GB cache RAID SAS Tri-port 6 Gb) cabled with four FC 3689 cables (AT 0.6 meter - HD SAS)

Note: The top ports of the FC 5913 or the FC ESA3 pair must attach by an HD SAS AA cable.

- c. Attach an AT cable (**B**) from the same numbered port on each SAS adapter (**A**) of a dual SAS adapter pair to SAS connectors 1 and 2 on the 5802 SAS conduit card (**C**), as shown in the following figure.




Figure 5. Configuration with two hard-disk-drive partitions that have dual SAS connectivity in a multi-initiator high-availability configuration

- d. Attach an AT cable (**B**) from other same numbered ports on the same adapter pair (**A**) to SAS connectors 3 and 4 on the 5802 SAS conduit card (**C**).
- e. Go to step 9 on page 41.
3. Complete the following steps to cable a configuration that includes one hard-disk-drive partition, 18 disks, and one SAS adapter.

- a. Set the hard-disk-drive partitions switch to position 1. The hard-disk-drive partitions switch is located at the rear of the 5802 expansion unit, just below the SAS connectors.
- b. Choose one of the following SAS adapter and cable combination options:
 - One FC ESA1 adapter (PCIe2 RAID SAS Dual-port 6 Gb) cabled with one FC 3689 cable (AT 0.6 meter - HD SAS)
 - One FC 5901 adapter (PCIe Dual - x4 SAS) cabled with one FC 3688 cable (AT 0.6 meter - Mini SAS)
- c. Attach an AT cable (**B**) from a port on a SAS adapter (**A**) to SAS connector 1 on the 5802 SAS conduit card (**C**), as shown in the following figure.




Figure 6. Configuration with one hard-disk-drive partition and one SAS adapter

- d. Optional: Use another port on the same SAS adapter to connect a 5886 disk drive enclosure or 5887 disk drive enclosure by using a SAS YO cable (**D**).
- e. Go to step 9 on page 41.
4. Complete the following steps to cable a configuration that includes two hard-disk-drive partitions, nine disks per hard-disk-drive partition, and one SAS adapter.

Note: You must have a minimum of one disk drive per hard-disk-drive partition.

- a. Set the hard-disk-drive partitions switch to position 2. The hard-disk-drive partitions switch is located at the rear of the 5802 expansion unit, just below the SAS connectors.
- b. Choose one of the following SAS adapter and cable combination options:
 - One FC ESA1 adapter (PCIe2 RAID SAS Dual-port 6 Gb) cabled with two FC 3689 cables (AT 0.6 meter - HD SAS)
 - One FC 5901 adapter (PCIe Dual - x4 SAS) cabled with two FC 3688 cables (AT 0.6 meter - Mini SAS)
- c. Attach an AT cable (**B**) from one port on the SAS adapter (**A**) to SAS connector 1 on the 5802 SAS conduit card (**C**), as shown in the following figure.




Figure 7. Configuration with two hard-disk-drive partitions and one SAS adapter

- d. Attach an AT cable (**B**) from another port on the SAS adapter (**A**) to SAS connector 1 on the 5802 SAS conduit card (**C**).
- e. Go to step 9 on page 41.
5. Complete the following steps to cable a configuration that includes two hard-disk-drive partitions, nine disks per hard-disk-drive partition, and two SAS adapters.

Note: You must have a minimum of one disk drive per hard-disk-drive partition.

- a. Set the hard-disk-drive partitions switch to position 2. The hard-disk-drive partitions switch is located at the rear of the 5802 expansion unit, just below the SAS connectors.
- b. Choose one of the following SAS adapter and cable combination options:
 - Two FC ESA1 adapters (PCIe2 RAID SAS Dual-port 6 Gb) cabled with two FC 3689 cables (AT 0.6 meter - HD SAS)
 - Two FC 5901 adapters (PCIe Dual - x4 SAS) cabled with two FC 3688 cables (AT 0.6 meter - Mini SAS)
- c. Attach an AT cable (**B**) from one port on one SAS adapter (**A**) to SAS connector 1 on the 5802 SAS conduit card (**C**), as shown in the following figure.




Figure 8. Configuration with two hard-disk-drive partitions and two SAS adapters

- d. Attach an AT cable (**B**) from one port on the other SAS adapter (**A**) to SAS connector 3 on the 5802 SAS conduit card (**C**).
- e. Optional: Use other ports on the same SAS adapters to connect a 5886 disk drive enclosure or 5887 disk drive enclosure by using SAS YO cables (**D**).
- f. Go to step 9 on page 41.

6. Complete the following steps to cable a configuration that includes two hard-disk-drive partitions, nine disks per hard-disk-drive partition, a dual SAS adapter pair per hard-disk-drive partition that controls the disks in the partition, and dual SAS connectivity to all drives.

Notes:

- For the following operating systems, this configuration must be a RAID configuration: AIX, Linux
 - You must have a minimum of one disk drive per hard-disk-drive partition.
- a. Set the hard-disk-drive partitions switch to position 2. The hard-disk-drive partitions switch is located at the rear of the 5802 expansion unit, just below the SAS connectors.
 - b. Choose one of the following SAS adapter and cable combination options:
 - Four FC ESA1 adapters (PCIe2 RAID SAS Dual-port 6 Gb) cabled with four FC 3689 cables (AT 0.6 meter - HD SAS)
 - Four FC 5805 or FC 5903 adapters (PCIe Dual - x4 3 Gb SAS RAID) cabled with four FC 3688 cables (AT 0.6 meter - Mini SAS)
 - Four FC 5901 adapters (PCIe Dual - x4 SAS) cabled with four FC 3688 cables (AT 0.6 meter - Mini SAS)

Note: Dual FC 5901 adapters are not supported in IBM i.

- Four FC 5913 or four FC ESA3 adapters (PCIe2 1.8 GB cache RAID SAS Tri-port 6 Gb) cabled with four FC 3689 cables (AT 0.6 meter - HD SAS)

Note: The top ports of the FC 5913 or the FC ESA3 pair must attach by an HD SAS AA cable.

- c. Attach an AT cable (**B**) from the same numbered port on each SAS adapter (**A**) of the first dual SAS adapter pair to SAS connectors 1 and 2 on the 5802 SAS conduit card (**C**), as shown in the following figure.




Figure 9. Configuration with two hard-disk-drive partitions that have dual SAS connectivity in a multi-initiator high-availability configuration

- d. Attach an AT cable (**B**) from the same numbered port on each SAS adapter of the second dual SAS adapter pair to SAS connectors 3 and 4 on the 5802 SAS conduit card (**C**).
 - e. Optional: Use other ports on the same SAS adapter pair to connect a 5886 disk drive enclosure by using SAS X cables or to connect a 5887 disk drive enclosure by using SAS YO cables (**D**). Both ports from the same adapter pair must be connected to the same disk drive enclosure.
 - f. Go to step 9 on page 41.
7. Complete the following steps to cable a configuration that includes one hard-disk-drive partition, 18 disks, and a dual SAS adapter pair with connectivity to all drives.

Note: For the following operating systems, this must be a RAID configuration: AIX, Linux

- a. Set the hard-disk-drive partitions switch to position 1. The hard-disk-drive partitions switch is located at the rear of the 5802 expansion unit, just below the SAS connectors.
- b. Choose one of the following SAS adapter and cable combination options:
 - Two FC ESA1 adapters (PCIe2 RAID SAS Dual-port 6 Gb) cabled with two FC 3689 cables (AT 0.6 meter - HD SAS)
 - Two FC 5805 or FC 5903 adapters (PCIe Dual - x4 3 Gb SAS RAID) cabled with two FC 3688 cables (AT 0.6 meter - Mini SAS)
 - Two FC 5901 adapters (PCIe Dual - x4 SAS) cabled with two FC 3688 cables (AT 0.6 meter - Mini SAS)
 - Two FC 5913 or two FC ESA3 adapters (PCIe2 1.8 GB cache RAID SAS Tri-port 6 Gb) cabled with two FC 3689 cables (AT 0.6 meter - HD SAS)

Note: The top ports of the FC 5913 or the FC ESA3 pair must attach by an HD SAS AA cable.

- c. Attach an AT cable (**B**) from the same numbered port on each SAS adapter (**A**) of a dual SAS adapter pair to SAS connectors 1 and 2 on the 5802 SAS conduit card (**C**), as shown in the following figure.




Figure 10. Configuration with a hard-disk-drive partition that has dual SAS connectivity in a multi-initiator high-availability configuration

- d. Optional: Use other ports on the same SAS adapter pair to connect a 5886 disk drive enclosure by using a SAS X cable or to connect a 5887 disk drive enclosure by using SAS YO cables (**D**). Both ports from the same adapter pair must be connected to the same disk drive enclosure.
 - e. Go to step 9 on page 41.
8. Complete the following steps to cable a configuration that includes four hard-disk-drive partitions, 18 disks, and one SAS adapter per hard-disk-drive partition.

Note: The 18 disks in this configuration are divided among the four hard-disk-drive partitions as follows:

- Partition 1: Disk drives 01 - 05.
- Partition 2: Disk drives 06 - 09.
- Partition 3: Disk drives 10 - 14.
- Partition 4: Disk drives 15 - 18.

- a. Set the hard-disk-drive partitions switch to position 4. The hard-disk-drive partitions switch is located at the rear of the 5802 expansion unit, just below the SAS connectors.
- b. Choose one of the following SAS adapter and cable combination options:

Note: You must provide an adapter and cable combination for each of the four disk partitions that has an installed disk drive.

- One FC ESA1 adapter (PCIe2 RAID SAS Dual-port 6 Gb) cabled with one FC 3689 cable (AT 0.6 meter - HD SAS)
 - One FC 5901 adapter (PCIe Dual - x4 SAS) cabled with one FC 3688 cable (AT 0.6 meter - Mini SAS)
- c. Attach an AT cable (**B**) from one port on each SAS adapter (**A**) to a SAS connector on the 5802 SAS conduit card (**C**), as shown in the following figure.




Figure 11. Configuration with four hard-disk-drive partitions that have four SAS adapters

- d. Optional: Use other ports on the SAS adapters to connect a 5886 disk drive enclosure or a 5887 disk drive enclosure by using SAS YO cables (**D**).
 - e. Go to step 9.
9. If the configuration that you chose required the hard-disk-drive partitions switch to be set to a different position, the new setting will be active when the 5802 is powered on. Otherwise, activate the new setting by powering off and powering on the 5802 expansion unit.

Removing I/O expansion units

Learn how to remove an input/output (I/O) expansion unit from the system and remove it from a rack.

Follow the procedures to remove an I/O expansion unit from a system with the system powered on or with the system powered off.

Note: For information about physically removing an I/O expansion unit from a rack, see “Removing an enclosure from the rack” on page 46.

Removing an I/O expansion unit with the system powered on

Learn about removing an input/output (I/O) expansion unit from the system with the system powered on.

Perform the following tasks before beginning this procedure:

- Record the time that you start the procedure. The time is referenced later during error log analysis.
- If you are performing a procedure that changes the 12X cabling or configuration for a system that is managed by a management console, you must first obtain a record that identifies for each expansion unit which I/O buses exist in that expansion unit. To acquire this information, perform the following steps:
 1. Select the managed system and then select the **Properties** task.
 2. From the HMC command line, enter the command `lshwres -r io --subtype bus -m managed system`.

Important: Be aware of the following items when you are removing an expansion unit:

- Concurrent I/O expansion unit removal is supported only on IBM Power[®] 770, IBM Power 780, and IBM Power 795 systems.
- Concurrent I/O expansion unit removal is supported only for 12X-attached I/O expansion units.
- Concurrent I/O expansion unit removal is supported only on systems that are managed by a Hardware Management Console (HMC).
- An I/O expansion unit can be concurrently removed from a system using these procedures and then later added back to the system configuration either concurrently or nonconcurrently. When the expansion unit is added back to the configuration, it is assigned new bus numbers.
- Concurrent removal of I/O expansion units must be done one expansion unit at a time. If multiple expansion units are to be removed concurrently, the procedure must be done separately for each one.

To remove an I/O expansion unit from the system configuration, complete the following steps:

1. Verify the existing 12X configuration. For details, see “Verifying the 12X loops with an HMC” on page 22. If problems exist with the 12X configuration, they must be corrected before continuing with the expansion unit removal procedure.
2. Verify the existing system power control network (SPCN) configuration by using the Advanced System Management Interface (ASMI). For details, see “Verifying the system power control network” on page 25. If problems exist with the SPCN configuration, they must be corrected before continuing with the expansion unit removal procedure.
3. Optional: Identify the expansion unit that you want to remove. Select the **Operations > LED Status > Identify LED** task for the managed system from which the expansion unit is being removed in order to activate and deactivate identify indicators.
4. Select from the following options:
 - If the system is not in manufacturing default configuration (MDC), continue with step 5
 - If the system is in manufacturing default configuration (MDC), continue with step 7.
5. If the system is not in manufacturing default configuration (MDC), use one of the following methods to ensure that none of the I/O resources in the expansion unit to be removed are owned by active partitions:
 - Use dynamic logical partitioning to dynamically remove I/O resources in the expansion unit that are owned by active partitions.
 - Power off logical partitions that own I/O resources in the expansion unit.
6. If the system is not in manufacturing default configuration (MDC), remove all I/O resources in the expansion unit to be removed from logical partition profiles.

Note: Failure to complete this step can result in messages when a profile is activated or edited that includes resources in an expansion unit that was removed. When such a profile is activated or edited, the HMC presents a message about missing resources. At that time, the option to have the HMC automatically remove the missing resources from the profile is given.

7. Power off the expansion unit that is to be removed by performing the following steps:
 - a. Select *the managed system*, **Serviceability > Hardware > Power On/Off Unit**.
 - b. In the Power On/Off Unit window, expand the managed system and select the expansion unit that is to be removed.
 - c. Click **Power Off**, and follow the instructions on the HMC until the expansion unit is powered off.
8. Remove ac power from the expansion unit that is to be removed by disconnecting the ac cables from the power supplies on the expansion unit.
9. Familiarize yourself with the SPCN cabling rules. For SPCN cabling examples, see “Examples: SPCN connections” on page 20.
10. Perform the following steps:

- a. Remove the SPCN cable between the SPCN 0 connector on the expansion unit that is to be removed and the corresponding connector on the preceding unit on the SPCN loop.

Note: The preceding unit on the SPCN loop might be another expansion unit or a system unit.

- b. Disconnect the SPCN cable from the SPCN 1 connector on the expansion unit that is to be removed. Attach that end to the connector on the preceding unit on the SPCN loop that was opened in the previous step.
11. Familiarize yourself with the 12X cabling rules. For 12X cabling examples, see “Examples: 12X connections” on page 18.
 12. Choose from the following options for removing an expansion unit from a 12X loop.

Note: To review the connector locations for your expansion unit, see “Connector locations” on page 126.

- To remove an expansion unit from a 12X loop that has only one expansion unit, continue with step 13.
 - To remove an expansion unit from a 12X loop that has more than one expansion unit, continue with step 14.
13. To remove an expansion unit from a 12X loop with only one expansion unit, complete the following steps:
 - a. Remove the 12X cable between connector T1 on the system unit and connector T1 on the expansion unit that is to be removed.
 - b. Disconnect the 12X cable from connector T2 on the expansion unit that is to be removed and attach it to connector T1 on the system unit.

Important: This step creates a 12X loop with no expansion units on it. The 12X cable runs between connectors T1 and T2 on the system unit. Do not skip this step. It is necessary for the system firmware to complete the removal process.

- c. Wait 2 minutes.
 - d. Remove the 12X cable between connectors T1 and T2 on the system unit.
 - e. Continue with step 15.
14. Perform the following steps to remove an expansion unit from a 12X loop that contains multiple expansion units.
 - a. Remove the 12X cable between connector 0 on the expansion unit to be removed and the corresponding connector on the preceding unit on the 12X loop.

Note: The preceding unit on the 12X loop might be another expansion unit or a system unit.

- b. Disconnect the 12X cable from connector 1 on the expansion unit to be removed. Attach that end to the connector on the preceding unit on the 12X loop that was opened in step 14a.
 - c. Continue with step 15.
15. Clear the inactive enclosure record for the expansion unit that was removed.
 - a. Access the ASMI from the HMC
 - 1) In the navigation pane, select **System Management > Servers**.
 - 2) In the content pane, select the system that has the expansion unit you want to remove.
 - 3) Select **Tasks > Operations > Launch Advanced Systems Management (ASM)**.
 - b. Log in to ASMI with administrator or authorized service provider authority.
 - c. Expand **System Configuration**.
 - d. Click **Configure I/O Enclosures > Clear inactive enclosures**.
 16. Repeat step 1 on page 42 to verify the 12X configuration.

Note: While concurrently removing an expansion unit, it is normal for error logs and event logs to be created with the following reference codes. The number and combination of these logs is dependent on the configuration. The logs can be ignored during this verification step if their timestamps indicate that they occurred during the removal procedure and if the resources that they reference were involved in the procedure.

- B7006981: This permanent reference code indicates a 12X adapter failure.
- B7006982: This permanent reference code indicates 12X link failure.
- B7006984: This informational reference code indicates an open 12X loop.
- B7006985: This informational reference code indicates the 12X loop was closed.
- B70069E6: This informational reference code indicates a 12X link stopped.
- B70069E7: This informational reference code indicates a 12X link started.

17. Repeat step 2 on page 42 to verify the SPCN configuration.

Note: While concurrently removing an expansion unit, it is normal for error logs and event logs to be created with the following reference codes. The number and combination of these logs is dependent on the configuration. The logs can be ignored during this verification step if their timestamps indicate that they occurred during the removal procedure and if the resources that they reference were involved in the procedure.

- 100090F0: This permanent reference code indicates that an expansion unit was dropped from the power control network.
- 10009135: This informational or permanent reference code indicates an open SPCN loop.
- 10009137: This permanent reference code indicates disruption of the 12X loop.
- 10009139: This informational reference code indicates that the SPCN loop was closed.

18. Check for new serviceable events that were generated during the removal procedure by performing the following steps:

- a. Select the managed system, and then click **Serviceability > Manage Serviceable Events**.
- b. In the Manage Serviceable Events window, specify event criteria to include any events that were generated during the procedure. Optionally, you can use the default criteria.
- c. Click **OK**. A table appears with the serviceable events that match your criteria.
- d. Close any serviceable events generated during the procedure that include one of the reference codes identified in step 16 on page 43 or 17.
- e. Perform standard problem analysis on any remaining open serviceable events generated during the procedure.

The expansion unit is now removed from the system configuration.

19. Physically remove the expansion unit from the rack. To learn about how to remove an enclosure from a rack, see “Removing an enclosure from the rack” on page 46.

Removing an I/O expansion unit with the system powered off

Learn about removing an input/output (I/O) expansion unit from the system with the system powered off.

Perform the following tasks before beginning this procedure:

- If you are performing a procedure that changes the 12X cabling or configuration for a system that is managed by a management console, you must first obtain a record that identifies for each expansion unit which I/O buses exist in that expansion unit. To acquire this information, perform the following steps:
 1. Select the managed system and then select the **Properties** task.
 2. From the HMC command line, enter the command `lshwres -r io --subtype bus -m managed system`.

To remove an I/O expansion unit from the system configuration that has the system powered off, complete the following steps:

1. If the system is powered on, power it off.
2. Remove ac power from the expansion unit that is to be removed by disconnecting the ac cables from the power supplies on the expansion unit.
3. Familiarize yourself with the SPCN cabling rules. For SPCN cabling examples, see “Examples: SPCN connections” on page 20.
4. Perform the following steps:
 - a. Remove the SPCN cable between the SPCN 0 connector on the expansion unit that is to be removed and the corresponding connector on the preceding unit on the SPCN loop.

Note: The preceding unit on the SPCN loop might be another expansion unit or a system unit.

- b. Disconnect the SPCN cable from the SPCN 1 connector on the expansion unit that is to be removed. Attach that end to the connector on the preceding unit on the SPCN loop that was opened in the previous step.
5. Familiarize yourself with the 12X cabling rules. For 12X examples, see “Examples: 12X connections” on page 18.
6. Choose from the following options for removing an expansion unit from a 12X loop.

Note: To review the connector locations for your expansion unit, see “Connector locations” on page 126.

- To remove an expansion unit from a 12X loop that has only one expansion unit, continue with step 7.
 - To remove an expansion unit from a 12X loop that has more than one expansion unit, continue with step 8.
7. To remove an expansion unit from a 12X loop with only one expansion unit, complete the following steps:
 - a. Remove the 12X cable between connector T1 on the system unit and connector T1 on the expansion unit that is to be removed.
 - b. Remove the 12X cable between connector T2 on the system unit and connector T2 on the expansion unit that is to be removed.
 - c. Continue with step 9.
 8. To remove an expansion unit from a 12X loop that contains multiple expansion units, complete the following steps:
 - a. Remove the 12X cable between connector 0 on the expansion unit to be removed and the corresponding connector on the preceding unit on the 12X loop.

Note: The preceding unit on the 12X loop might be another expansion unit or a system unit.

- b. Disconnect the 12X cable from connector 1 on the expansion unit to be removed. Attach that end to the connector on the preceding unit on the 12X loop that was opened in step 8a.
 - c. Continue with step 9.
9. Physically remove the expansion unit from the rack. To learn how to remove an enclosure from a rack, see “Removing an enclosure from the rack” on page 46.
10. Power on the system but do not start the partitions.
11. Check for new serviceable events that were created since starting the system:
 - If the system is managed by an HMC, complete the following steps to check for new serviceable events:
 - a. Select *the managed system* **Serviceability > Manage Serviceable Events**.
 - b. In the Manage Serviceable Events window, click **OK** to accept the default criteria. A table is shown with the serviceable events that match the criteria.

- c. Continue with step 12.
 - If the system is not managed by an HMC, use the serviceable event log tool in the operating system to check for new serviceable events. Continue with step 12.
12. Perform standard problem analysis on any open serviceable events that were created since starting the system.
 13. Start the logical partitions.

Removing an enclosure from the rack

This procedure provides you with general steps to help you physically remove an enclosure that is mounted in a rack. Because methods for the attachment of enclosures to a rack vary by model, this procedure provides only conceptual removal steps.

To perform this task, you need the following items:

- A flat-head screwdriver
- A Philips-head screwdriver
- A Torx screwdriver
- Up to three people to physically lift and move the enclosure
- A clear space to place the enclosure and mounting hardware you are removing

Important: Supplement each step of this procedure with detailed information about the installation of the enclosure, when possible. Before you begin this procedure, check to see if the following information is available about the enclosure that you are removing:

- The inventory list of parts provided with the delivery of your enclosure.
- Installation documentation about your enclosure model provided online or with the delivery of your enclosure.

Tip: Online installation documentation for your enclosure might not be present in the Power Systems™ hardware level that you are currently searching. You might need to search earlier Power Systems hardware levels to locate installation documentation for the enclosure that you are removing.

To remove an enclosure from the rack, complete the following steps:

1. If the enclosure is not in a powered-off state, power off the enclosure.

Important: If the enclosure that you are removing is an expansion unit, the expansion unit must be powered off and removed from the system. Perform the removal of the expansion unit with the system powered on and running (concurrently) or with the system powered off (nonconcurrently).

- To learn about removing an expansion unit concurrently, see “Removing an I/O expansion unit with the system powered on” on page 41.
 - To learn about removing an expansion unit nonconcurrently, see “Removing an I/O expansion unit with the system powered off” on page 44.
2. If a service position exists for your enclosure, place the enclosure in its service position.
 3. If the enclosure was installed with shipping or bracing brackets, complete the following steps:
 - a. Remove any screws that attach the brackets to the rack.
 - b. Remove any screws that attach the brackets to the enclosure.
 - c. Carefully lift the brackets from the enclosure and remove them from the rack.

CAUTION:

Shipping and bracing brackets can be heavier than expected. Use more than one person to lift and move the brackets, if required.

4. Disconnect the enclosure from all devices and power as follows:

- a. At the rear of the enclosure, disconnect power cords, any I/O bus cables connecting the enclosure to adapters, and other cables that might be present.
 - b. At the front of the enclosure, disconnect any UPIC cables and other cables that might be present.
5. Note the location of all screws that provide for the attachment of the enclosure, the enclosure mounting hardware, and the rack to each other. If enclosure bezels cover screws on the left and right ends of the front enclosure panel, remove those bezels.
 6. Identify which screws secure the enclosure in position, which screws are weight-bearing screws, and which screws serve additional purposes. To remove the enclosure safely, you must understand the functions of all screws that are used to install the enclosure. Evaluate the screws in the following order:
 - a. Make a note of which screws are securing screws. You can identify a securing screw as any screw that fastens the enclosure chassis to any other fixture, such as the rack or a rail. Even if a screw attaches the chassis to another fixture that bears weight, it remains a securing screw and not a weight-bearing screw. If a screw secures only a mounting hardware item to the rack, it is not a securing screw. Check both the front and rear of the enclosure for securing screws. During the enclosure installation task, screws to secure the enclosure were the last screws to be installed. Therefore, they are the first screws that you are directed to remove later in this task to enable the detachment of the enclosure.

Note: The enclosure might be installed in a shell that is, in turn, installed to a rack or rails. If so, then any screw that fastens the enclosure to the shell is a securing screw. Also, any screw that fastens the shell to the rack or rails is a securing screw.

Exception: If the enclosure is installed on slide-out rails, the securing screws that attach the enclosure to the rails might also be weight-bearing screws. These models were installed in a two-step process that used rail-mounting guides. In the first step, the guides were temporarily attached to the sides of the enclosure. In the second step, the guides were removed after the enclosure was lifted onto the rails and securing screws were installed. For the purposes of this procedure, identify these screws as the weight-bearing screws that you are directed to remove later in this task.

- b. Make a note of which screws are enclosure weight-bearing screws. You can identify a weight-bearing screw as any screw that fastens a rail to the rack, unless the screw is already identified as a securing screw. The type of rail, such as a stationary rail or a slide-out rail, is not a factor in identifying weight-bearing screws. During the enclosure installation task, screws to attach weight-bearing mounting hardware to the rack were installed before the enclosure was lifted onto the rails and secured. Therefore, you will be directed to remove them later in this task only after the enclosure has been safely removed.
 - c. Make a note of any remaining screws used in the installation of the enclosure. These screws serve other purposes and are the last screws that you are directed to remove later in this task.
7. If the enclosure is installed on slide-out rails and you identified the securing screws as weight-bearing screws in step 6a, review the following options for the removal of the enclosure. Then complete the option that is best for your situation:
 - a. If you have the original rail-mounting guides and screws that were used during the installation of the enclosure, complete the following steps:
 - 1) Attach the rail-mounting guides to the sides of the enclosure using the screws provided for the original installation.
 - 2) Remove the screws that you identified in the Exception note of step 6a as weight-bearing screws.
 - 3) Using three people, lift the enclosure from the rails and place it carefully in the space that you cleared. If the enclosure cannot be lifted because it is still attached to the rails, return to step 6 to identify additional screws that might need to be removed.
 - 4) Continue with step 12 on page 48.

- b. If you do not have the original rail-mounting guides and screws that were used during the installation of the enclosure, complete the following steps:
 - 1) Position three people at the front and sides of the enclosure to bear its weight and hold it steady during the removal of weight-bearing screws.
 - 2) Remove the screws that you identified in the Exception note of step 6a on page 47 as weight-bearing screws in the slide-out rails. Order the removal of screws so that the last two screws to be removed are ones installed at diagonal points on opposite rails.
 - 3) Using the same three people holding the enclosure, lift the enclosure from the rails and place it carefully in the space that you cleared. If the enclosure cannot be lifted because it is still attached to the rails, return to step 6 on page 47 to identify additional screws that might need to be removed.
 - 4) Continue with step 12.

8. If the enclosure is installed in a shell, complete the following steps:
 - a. If a cable retention bracket is present, remove the thumb screw that attaches the bracket to the shell, and remove the bracket.
 - b. Remove any securing screws that fasten the enclosure to the shell.
 - c. Slide the enclosure as required into a position where it can be removed. Using one or two people to handle the enclosure safely, slide the enclosure out from the shell.
 - d. If additional enclosures must be removed from the shell, repeat steps 8a through 8c as required to remove the enclosures.
 - e. Remove all remaining securing screws that fasten the shell to the rack or rails.
 - f. Lift the shell out from the rack and place it carefully in the space that you cleared.
 - g. Continue with step 12.
9. Remove all screws that you identified in step 6a on page 47 as securing screws. When this step is complete, the enclosure should no longer be secured to the rack or to any mounting hardware.

CAUTION:

Be careful in this step to remove only the screws that prevent the enclosure from being detached from the rack or mounting hardware. If you believe that a particular screw might be weight-bearing, do not remove it.

10. Perform a test lift of the enclosure, using three people, to determine how many people are required to handle the enclosure safely.

CAUTION:

The heaviest of enclosures requires three people for safe lifting. Using fewer people than are required to lift an enclosure can result in injury.

11. Slide the enclosure as required into a position where it can be removed. Lift the enclosure out from the rack and place it carefully in the space that you cleared. If the enclosure is still attached, return to step 6 on page 47 to identify additional screws that might need to be removed.

Note: Depending on the method of attachment, you might be required to open latches to fully detach the enclosure from the rack and mounting hardware.

12. If you are required to remove the rails, you can remove them safely now. The procedure to remove rails includes some of the following common steps at one or both ends of the rail:
 - a. Fold down any hinge brackets located where the rail attaches to the rack.
 - b. Remove all screws that you identified in step 6b on page 47 as weight-bearing screws that attach the rail to the rack.
 - c. Push any spring-loaded seating pins in the ends of the rails back through their rack holes.
 - d. Open latches or loosen fasteners on the rail as required to detach the rail.
 - e. Perform any additional action required to retract and remove the rail.
13. Remove any remaining screws that you identified in step 6c on page 47, and remove any remaining enclosure mounting hardware from the rack.

14. Store all enclosure mounting hardware, brackets, bezels, and the shell for future use.

Disk drive enclosures

Use this information to learn about the cabling, SCSI addressing, and postconnection requirements for the disk drive enclosure.

Overview for disk drive enclosures

The disk drive subsystem is the portion of the enclosure that houses and controls the disk drives.

The following connector information pertains to the systems and disk drive enclosures that are configured in this topic:

1. Adapter cables attach to ports at the rear of the systems in which the adapters are installed. To identify port locations for the system in your configuration, see "Server connectors" on page 127.
2. Adapters are cabled to the disk drive enclosures by using ports at the rear of the enclosure. To identify port locations for the disk drive enclosures in your configuration, see "Enclosure connectors" on page 140.

5786, 5787, 7031-D24, or 7031-T24 SCSI disk drive enclosure

Learn about the SCSI disk drive enclosure and how to assemble it.

This SCSI disk drive enclosure can hold up 24 disk drives. The enclosure is organized into four groups, with each group containing six disk drive slots. The enclosure can be used as a stand-alone enclosure or a rack-mounted drawer.

When connecting with your server, you can connect your SCSI cables to either a single-initiator repeater card or to a dual-initiator repeater card on your SCSI disk drive enclosure. The dual-initiator repeater card allows for high availability and combining groups of disk drive slots.

To determine the level of software that you need to support the SCSI disk drive enclosure, see IBM Prerequisite.

Notes: Before you connect the enclosure, read the following information.

- You must install the SCSI adapter that is used to connect the enclosure to your system. To view the PDF file of PCI adapters, approximately 40 MB in size, see <http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hak/p7hak.pdf>.
- For systems or partitions with an AIX or Linux operating system, you must have the SCSI adapter information for your operating system. This is critical information for configuring the disks after you have attached the enclosure. Print the information before you begin the procedure. This information is available from the SCSI PCI Adapters (http://publib.boulder.ibm.com/infocenter/pseries/v5r3/index.jsp?topic=/com.ibm.pseries.doc/hardware_docs/scsipciadapters.htm) website.
- If you are going to configure RAID arrays, ensure that you have the following number of available disks for each RAID level:

RAID 0 or 1

Two drives per array

RAID 5

At least three drives per array

RAID 6

At least four drives per array

Connecting the 5786, 5787, 7031-D24, or 7031-T24 to a system running AIX

Use the information in this section to perform the cabling, SCSI addressing, and postconnection requirements for the disk drive enclosure.

To connect the 5786, 5787, 7031-D24, or 7031-T24 to a system that has an AIX operating system, perform the following steps:

1. Install the initiator repeater cards into the disk drive enclosure. For instructions, see “SCSI repeater card” on page 64.
2. Connect the SCSI cable to your repeater card. Choose one of the following methods based on your situation:
 - **Single repeater cards or dual repeater cards:** If you are using only single repeater cards or only dual repeater cards, install your SCSI cable (A) to the initiator repeater card (B).
 - **Single SCSI adapter connected to single repeater card:** If, in this case, you are not sharing resources (disk drives), simply attach each SCSI cable to each repeater card.




Figure 12. SCSI cable to single repeater card

- **Two SCSI adapters connected to dual repeater cards:** If, in this case, you are sharing resources (disk drives) between systems or logical partitions, ensure that you understand the SCSI addressing and cabling order. For more information, see “Connecting and configuring the 5786, 5787, 7031-D24, or 7031-T24 SCSI disk drive enclosure in an AIX clustered environment” on page 55.




Figure 13. SCSI cable to a dual repeater card that is connected with two SCSI cards

- **Four SCSI adapters connected to dual repeater cards:** If, in this case, you are sharing resources (disk drives) between systems or logical partitions, ensure that you understand the SCSI addressing and cabling order. For more information, see “Connecting and configuring the 5786, 5787, 7031-D24, or 7031-T24 SCSI disk drive enclosure in an AIX clustered environment” on page 55.




Figure 14. SCSI cable to a dual repeater card that is connected with four SCSI cards

- **Single SCSI adapter connected to a combination of single and dual repeater cards:** If, in this case, you are sharing resources (disk drives) between systems or logical partitions, ensure that you

understand the SCSI addressing and cabling order. For more information, see “Connecting and configuring the 5786, 5787, 7031-D24, or 7031-T24 SCSI disk drive enclosure in an AIX clustered environment” on page 55.




Figure 15. SCSI cable to a dual repeater card that is connected to a single repeater card

3. Use the information that you printed for the SCSI adapter to complete the configuration of your disk drives. For more information, see the SCSI PCI Adapters (http://publib.boulder.ibm.com/infocenter/pseries/v5r3/index.jsp?topic=/com.ibm.pseries.doc/hardware_docs/scsipciadapters.htm) Web page and locate the file for the adapter that you are using to attach the enclosure to your system.

Important: If you plan on using the disk enclosure in an AIX environment, it is critical that you set up an AIX CRON job so that if errors occur, they are reported to the attached system. The AIX CRON job must be set up during the initial installation.

Note: Errors on disks in the enclosure are found without this CRON job (script) through normal automatic error log analysis. However, other errors, such as fan and power supply problems in the enclosure, are found only when you run diagnostics manually. You can run diagnostics manually by using the following scripts.

To collect enclosure errors, add the `run_ses_healthcheck` job to the system CRON table. Edit the system CRON by using the `crontab -e` command. At the bottom of that file, enter:

```
* 3 * * * /usr/lpp/diagnostics/bin/run_ses_healthcheck 1>/dev/null 2>/dev/null
```

This CRON runs the script `run_ses_healthcheck` daily at 3 a.m. The contents of the script depend on the configuration of the system to which the enclosure is attached.

Example 1

If the system to which the enclosure is attached has a Hardware Management Console (HMC) or is running the Electronic Service Agent™ application, create a file named `run_ses_healthcheck` in the `/usr/lpp/diagnostics/bin` directory, with the following contents:

```
-----
#!/bin/ksh
#Name:run_ses_healthcheck
#Location:/usr/lpp/diagnostics/bin
#Function: SCSI SES hourly heathcheck
for i in `lsdev -Cc container -t ses -s scsi -F name -S available`
do
    diag -cd $i > /dev/null
done
#any registered "external notification" will be notified of errors
#(such as HMC or SDMC or Electronic Service Agent)
-----
```

Note: On a logically partitioned system, the CRON job and script only need to run on one partition that is associated with the enclosure. However, it is acceptable to have the same CRON job and scripts on any partition that is associated with the enclosure.

Example 2

If the system to which the enclosure is attached has no HMC and is not running the Electronic Service Agent application, create a file named `run_ses_healthcheck` in the `/usr/lpp/diagnostics/bin` directory, with the following contents:

```
-----
#!/bin/ksh
#Name:run_ses_healthcheck
#Location:/usr/lpp/diagnostics/bin
#Function: SCSI SES hourly heathcheck
for i in `lsdev -Cc container -t ses -s scsi -F name -S available`
do
    diag -cd $i > /dev/null
    if [ $? -ne 0 ]
    then
        /usr/lpp/diagnostics/bin/diagrpt -o >/tmp/ses.health.output
        #you might want to process the output prior to placing it in
        #a file
        #somhow notify the user of the error. A sample is shown
        #below.
        mail -s "7031 Health Check" root</tmp/ses.health.output
        rm /tmp/ses.health.output
    fi
done
-----
```

Note: You can customize the actual notification in the script based on your preferences. If errors need attention, these scripts send mail to the root user. You can modify the script to inform certain users about the errors.

In both examples, you have created the `run_ses_healthcheck` file. To run this file, type the following command at the AIX command prompt:

```
chmod 544 /usr/lpp/diagnostics/bin/run_ses_healthcheck
```

Connecting and configuring the 5786, 5787, 7031-D24, or 7031-T24 SCSI disk drive enclosure in an AIX clustered environment

Special considerations exist for connecting the disk drive enclosure in a clustered environment. The information in this section can help you meet these considerations.

Gather and have available any planning and cabling documentation that you have for the system to which you are connecting in a clustered environment so that you can refer to it during this procedure.

To assure that the SCSI disk drive enclosure is ready to be connected to the cluster, complete the following steps:

1. Ensure that each SCSI device that is connected to the shared SCSI bus has a unique ID. A common configuration is to set the SCSI ID of the adapters on the nodes to be higher than the SCSI IDs of the shared devices. (Devices with higher IDs take precedence in a SCSI bus contention.)
 - Use the command `lscfg | grep scsi` to determine and record the logical name of each adapter. In the command output, the first column lists the logical name of the SCSI adapter, such as `+ SCSI0`.
 - Use the `lscfg -vp1` command and logical name of the adapter to record the I/O slot (physical slot) that each SCSI adapter uses, for example, `lscfg -vp1 scsi0`, where `scsi0` is the logical name of the adapter.
 - Use the `lsattr` command, as in the following example, to find the ID of the adapter `scsi0`: `lsattr -E -l scsi0 | grep id`

Note: Do not use wildcard characters or full path names on the command line for the device name designation.

In the resulting output, the first column lists the attribute names. The integer to the right of the ID attribute is the adapter SCSI ID.

2. If a SCSI adapter address conflict exists, go to the next step. If no conflict exists, return to the SCSI connection method overview that referred you to this procedure.
3. Examine the cabling to see if the correct ports on the system are connected to the correct ports on the repeater cards on the disk drive enclosure. For example, if you are connecting the system using two host adapter cards and a dual repeater card as shown in the following figure, connect the systems as in the instructions that follow.




Figure 16. SCSI cable to a dual repeater card that is connected with two SCSI cards

4. Cable the first SCSI bus or set of shared disks by performing the following steps:
 - a. Connect the bottom port of the adapter on Host A (SCSI address 7) to the exterior port on the repeater card in slot C2 on the enclosure.
 - b. Connect the top port of the adapter on Host B (SCSI address 6) to the exterior port on the repeater card in slot C3 on the enclosure.

This connection creates a set of shared disks and a shared SCSI bus (such as scsi0) between Host A and Host B. The important thing to consider at this point is that you connect the cables in descending order of the SCSI address on the host adapter. Connecting the cables in descending order of the host SCSI address reduces the potential for conflicting addresses.

5. Cable the second SCSI bus or set of shared disks by performing the following steps:
 - a. Connect the bottom port of the adapter on Host B (SCSI address 7) to the interior port on the repeater card in slot C5 on the enclosure.
 - b. Connect the top port of the adapter on Host A (SCSI address 6) to the interior port on the repeater card in slot C4 on the enclosure.

This connection creates a set of shared disks and a shared SCSI bus (such as scsi1) between Host A and Host B. The important thing to consider at this point is that this is an independent SCSI bus or set of shared disks from the first set. Ensuring that you do not mistakenly connect one of the cables from either SCSI bus to another SCSI bus reduces the potential for conflicting SCSI addresses.

Important: When adding, removing, or replacing SCSI adapters in host systems, you might want to disconnect the cabling at the enclosure first and reconnect the cabling to the enclosure last. This is because the new adapter SCSI ID default might be 7. You should ensure that the correct SCSI ID is set for the port you are cabling.

Connecting and configuring the disk drive enclosure in a system running Linux

Use the information in this section to perform the cabling, SCSI addressing, and postconnection requirements for the disk drive enclosure.

To connect and configure the disk drive enclosure in a system that has a Linux operating system, perform the following steps:

1. Install the initiator repeater cards into the disk drive enclosure. For instructions, see SCSI repeater card(<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/scsirepeatcard.htm>).
2. Connect the SCSI cable to your repeater card. Choose one of the following options based on your situation:
 - If you are using only single repeater cards or only dual repeater cards, connect your SCSI cable (A) to the initiator repeater card (B). For more details, see Figure 17 on page 58, Figure 18 on page 58, or Figure 19 on page 59, depending on your situation.




Figure 17. SCSI cable to a single repeater card





Figure 18. SCSI cable to a dual repeater card that is connected with two SCSI cards



IPHAL981-2

Figure 19. SCSI cable to a dual repeater card that is connected with four SCSI cards

- If you are using a combination of single and dual repeater cards, connect a SCSI cable to the dual-initiator repeater card (A). Then, connect the dual repeater card (A) to the single repeater card (B) by using a different SCSI cable. For more information, see Figure 20 on page 60.




Figure 20. SCSI cable to a dual repeater card that is connected to a single repeater card

3. Complete the SCSI cable connection process in this step.

Important: The SCSI adapter or the system or partition must be powered off before making the connection in this step.

Connect the other end of the SCSI cable (C) to the SCSI adapter (D) on the server. For more details, see Figure 21.




Figure 21. SCSI cable to a SCSI interface card

4. Add the disk drives to the operating system that you are using. See Disk drives and solid state drives (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hal/p7halkickoff.htm>).

Important: If you plan on using the disk enclosure in a Linux environment, it is critical that you set up a Linux CRON job so that if errors exist, they are reported to the attached system. The Linux CRON job must be set up during the initial installation.

5. Set up the Linux CRON job during the initial installation of your SCSI disk drive enclosure. Perform the following tasks to ensure that your Linux CRON job is set up:
 - a. Install the following RPMs from Service and productivity tools (<http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html>):
 - librtas
 - lsvpd
 - ppc64-utils
 - diagela
 - b. If your logical partitions are managed by an HMC, install the `src`, `rsct.core.utils`, `rsct.core`, `csn.core`, `csn.client`, and `devices.chrp.base.ServiceRM` packages in order to report errors to the HMC. Instructions for installing the packages are located at Service and productivity tools (<http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html>).

Note: Errors on hard disks within the enclosure are not found using this CRON job. This diagnostic test is limited to reporting errors with the fans, power supplies, Vital Product Data (VPD) card, and repeater in the enclosure.

- c. To run enclosure diagnostics, add the `diag_encl` job to the system CRON table. As root user, edit the system CRON by using the `crontab -e` command. At the bottom of that file, enter:

```
* 3 * * * /usr/sbin/diag_encl -s 1>/dev/null 2>&1
```

The contents of the script depend on the configuration of the system to which the enclosure is attached.

This CRON runs the diagnostic application `diag_encl` daily at 3 a.m. The application reads the contents of the `/etc/diagela/diagela.config` file to determine which applications have been registered to be notified of errors. To add additional notification methods or for more information, see that file. By default, the following notifications occur:

- If the system is managed by an HMC, notifications are shown in the Manage Serviceable Events task on the HMC. If the system is not managed by an HMC, the root group is notified.
- If the system has the Electronic Service Agent application installed, it is notified.
- The error details are printed to the console, to the end of the `/var/log/platform` log file, and to the `syslog` file in `/var/log/messages`.

Notes:

- 1) On a logically partitioned system, it is only necessary to have the CRON job on one partition that is associated with the enclosure. However, it is acceptable to have the same CRON job on all partitions.
- 2) The `diag_encl` application might run at any time to diagnose enclosures that are attached to the system.

Connecting and configuring the enclosure in a system running IBM i

If you are running the IBM i operating system, use this information to perform the cabling, SCSI addressing, and postconnection requirements for the disk drive enclosure.

To connect and configure the disk drive enclosure in a system that has an IBM i operating system, perform the following steps:

1. Install the initiator repeater cards into the disk drive enclosure. For instructions, see “SCSI repeater card” on page 64.
2. Connect the SCSI cable to your repeater card. Choose one of the following options based on your situation:

Tip: Ensure that the cable length for connections across the enclosure are long enough to allow concurrent maintenance of disk drives.

- If you are using only single repeater cards or only dual repeater cards, connect your SCSI cable (A) to the initiator repeater card (B).




Figure 22. SCSI cable to a single repeater card

- If you are using a combination of single and dual repeater cards, connect a SCSI cable to the dual-initiator repeater card (A). Then, connect the dual repeater card (A) to the single repeater card (B) with a different SCSI cable.




Figure 23. SCSI cable to a dual repeater card that is connected to a single repeater card

3. Complete the SCSI cable connection process in this step.

Important: The SCSI adapter or the system or partition must be powered off before making the connection in this step.

Connect the other end of the SCSI cable (C) to the SCSI adapter (D) on the server. For more details, see Figure 24.




Figure 24. SCSI cable to a SCSI interface card

4. Add the disk drives to the operating system.

SCSI repeater card

Learn about where to place and how to install the SCSI repeater card.

Use the following figures and instructions to place your repeater cards.




Figure 25. Repeater card location of the 5786 and 7031-D24 SCSI disk drive enclosure (drawer models)




Figure 26. Repeater card location of the 5787 and 7031-T24 SCSI disk drive enclosure (deskside models)

If you only have dual repeater cards, place them as follows:

Important: On deskside models, the locations C3 and C5 are located on top, and C2 and C4 are located on the bottom. Follow the location codes when placing repeater cards.

Table 6. Repeater card placement for single and dual style repeater cards

Type and number of repeater cards	Single repeater card placement	Dual repeater card placement
1 single repeater card and 0 dual repeater cards	C5	
2 single repeater cards and 0 dual repeater cards	Place in order C5 and C2	
3 single repeater cards and 0 dual repeater cards	Place in order C5, C2, and C4	
4 single repeater cards and 0 dual repeater cards	Place in order C5, C2, C4, and C3	
1 single repeater card and 1 dual repeater card	Place in C5 first, and then place dual repeater card	C4
1 single repeater card and 2 dual repeater cards	Place in C5 first, and then place dual repeater card	Place in order C4 and C3

Table 6. Repeater card placement for single and dual style repeater cards (continued)


Type and number of repeater cards	Single repeater card placement	Dual repeater card placement
2 single repeater cards and 1 dual repeater cards	Place in C5, C2, and then place dual repeater card	C4
2 single repeater cards and 2 dual repeater cards	Place in C5, C2, and then place dual repeater card	Place in order C4 and C3
0 single repeater cards and 1 dual repeater card		C4
0 single repeater cards and 2 dual repeater cards		Place in order C4 and C5
0 single repeater cards and 3 dual repeater cards		Place in order C4, C5, and C3
0 single repeater cards and 4 dual repeater cards Restriction: This configuration is for systems that are running the AIX or Linux operating system only.		Place in order C4, C5, C3, and C2

The following figures show the locations of the disk drives for the front and back of the 5786, 5787, 7031-D24, or 7031-T24 SCSI disk drive enclosure.



IPHAL975-0

Figure 27. Front view of 5786, 5787, 7031-D24, or 7031-T24 SCSI disk drive enclosure



IPHAL976-0

Figure 28. Back view of 5786, 5787, 7031-D24, or 7031-T24 SCSI disk drive enclosure

The following list describes how the repeater card locations affect the disk drive slots.

- The top-left repeater card (C2) drives the front right six slots (D07 to D12). See Figure 27 on page 66.
- The top-right repeater card (C3) drives the front left six slots (D01 to D06). See Figure 27 on page 66.
- The bottom-left repeater card (C4) drives the back left six slots (D13 to D18). See Figure 28.
- The bottom-right repeater card (C5) drives the back right six slots (D19 to D24). See Figure 28.

To install your repeater card, complete the following steps:

1. Remove the filler from the repeater card slot.
2. Install your new repeater card.

5886 SAS disk drive enclosure

Learn about the SAS disk drive enclosure and how to assemble it.

1. Review the following information:

The SAS disk drive enclosure can hold up to 12 disk drives. The enclosure cannot be split into two independent groups.

The SAS disk drive enclosure supports the following operating systems:

- AIX
- IBM i
- Linux
- VIOS

To determine the level of software that you need to support the SAS disk drive enclosure, see IBM Prerequisite (http://www-912.ibm.com/e_dir/eServerPrereq.nsf).

Tip: If you are planning to configure RAID arrays, ensure that you have the minimum number of available disks for each RAID level:

RAID 0

Minimum of one drive per array

RAID 5

Minimum of three drives per array

RAID 6

Minimum of four drives per array

RAID 10

Minimum of two drives per array

Connecting the SAS adapter to the 5886 disk drive enclosure

Learn how to connect the 5886 disk drive enclosure

To connect the 5886 disk drive enclosure to one or more Serial-attached SCSI (SAS) adapters or to a system with support for a SAS drawer, complete the following steps.

For additional information related to SAS cabling and cabling configurations, see Serial-attached SCSI cable planning (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7had/p7hadsascabing.htm>).

1. Confirm that all adapters that you need to connect to the 5886 disk drive enclosure have been installed in the system. If the adapters have not been installed, complete the adapter installation procedure for your system before continuing with this task.
2. If the server requires an internal cable to be installed to produce an external SAS port for connection with the 5886 disk drive enclosure, confirm that the installation has been performed. Review the following list of systems that require the installation of an external SAS port, and then complete the action for your system, if required:

Remember: When installing or confirming the use of an external SAS port, record the location of the external SAS port on the server. Later in this procedure, you are instructed to install the external SAS cable to this server connector location.

- For installing the external SAS port on the IBM Power 710 Express (8231-E2B and 8231-E1C) or IBM Power 730 Express (8231-E2B and 8231-E2C), see Installing the external SAS port in the 8231-E2B, 8231-E1C, 8231-E1D, 8231-E2C, or 8231-E2D (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hal/installsasport_71x_73x.htm).
 - For installing the external SAS port on the IBM Power 720 Express (8202-E4B, 8202-E4C, and 8202-E4D) or IBM Power 740 Express (8205-E6B, 8205-E6C, and 8205-E6D), see Installing the external SAS port in the 8202-E4B, 8202-E4C, 8202-E4D, 8205-E6B, 8205-E6C, or 8205-E6D system (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hal/installsasport_72x_74x.htm).
 - For installing the external SAS port on the IBM Power 750 Express (8233-E8B) or IBM Power 755 (8236-E8C), see Installing the external SAS port on an 8233-E8B or 8236-E8C (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hal/dasd550backplanesasfeature.htm>).
 - For installing the external SAS port on the IBM Power 770 (9117-MMB, 9117-MMC, and 9117-MMD) or IBM Power 780 (9179-MHB, 9179-MHC, and 9179-MHD), see Installing the external SAS port in the 9117-MMB, 9117-MMC, 9117-MMD, 9179-MHB, 9179-MHC, or 9179-MHD (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hal/9117installsasport.htm>).
3. Locate the connection for each adapter for which you are using an external SAS cable to connect the 5886.

Note: The configuration figures in step 4 use adapter cards to represent the external server connections. The adapter card can represent one of the following connection types:

- The external port of an adapter that you confirmed in step 1.
 - The external port of an internal adapter cable that you confirmed in step 2.
4. Choose one of the following options to cable the SAS adapter:
 - To perform a connection of one 5886 disk drive enclosure to a single SAS adapter by using a YO cable, go to step 5 on page 69.

- To perform a connection of two 5886 disk drive enclosures to a single SAS adapter by using YO cables, where each enclosure can be attached to an additional 5886 disk drive enclosure by using an EE cable, go to step 6.
- To perform a connection of two 5886 disk drive enclosures to a SAS adapter pair by using X cables, where each enclosure can be attached to an additional 5886 disk drive enclosure by using an EE cable, go to step 7 on page 70.
- To perform a connection of one 5886 disk drive enclosure by using a YI cable to a system with support for a single SAS drawer, go to step 8 on page 71.

If your SAS configuration requirements are not supported by any of these options, go to step 9 on page 72.

5. Perform the connection of one 5886 disk drive enclosure (**A**) to a single SAS adapter (**C**) by using a YO cable (**B**) as shown in Figure 29.




Figure 29. Connection of one 5886 disk drive enclosure to a single SAS adapter by using a YO cable

Go to step 9 on page 72.

6. Perform the connection of two 5886 disk drive enclosures (**A and C**) to a single SAS adapter (**I**) by using YO cables (**G and H**), where each enclosure can be attached to an additional 5886 disk drive enclosure (**B and D**) by using two EE cables (**E and F**) as shown in Figure 30 on page 70.




Figure 30. Connection of two 5886 disk drive enclosures to a single SAS adapter by using YO cables, where each enclosure can be attached to an additional 5886 disk drive enclosure by using two EE cables

Go to step 9 on page 72.

7. Perform the connection of two 5886 disk drive enclosures (**A and C**) to a SAS adapter pair (**I**) by using X cables (**G and H**), where each enclosure can be attached to an additional 5886 disk drive enclosure (**B and D**) by using two EE cables (**E and F**) as shown in Figure 31 on page 71.




Figure 31. Connection of two 5886 disk drive enclosures to a SAS adapter pair by using X cables, where each enclosure can be attached to an additional 5886 disk drive enclosure by using two EE cables

Go to step 9 on page 72.

8. Perform the connection of one 5886 disk drive enclosure (A) by using a YI cable (B) to a system with support for a single SAS drawer (C) as shown in Figure 32 on page 72.




Figure 32. Connection of one 5886 disk drive enclosure by using a YI cable to a system with support for a single SAS drawer

Go to step 9.

- For additional information to plan or complete your SAS cable installation, see Serial-attached SCSI cable planning (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7had/p7hadsascabbling.htm>).

Remember:

- Connect power cables and apply power to the 5886 disk drive enclosure.
- Add the disk drives to the operating system you are using.
- Verify the new configuration. To verify that the system or logical partition recognizes the disk drive enclosure, see Verifying the installed part (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7haj/p7hajhsmverify.htm>).

5887 SAS disk drive enclosure

Learn about the SAS disk drive enclosure and how to assemble it.

- Review the following information:

The SAS disk drive enclosure can hold up to 24 disk drives. The enclosure can be split logically into one, two, or four independent groups.

The SAS disk drive enclosure supports the following operating systems:

- AIX
- IBM i
- Linux
- VIOS

To determine the level of software that you need to support the SAS disk drive enclosure, see IBM Prerequisite (http://www-912.ibm.com/e_dir/eServerPrereq.nsf).

Tip: If you are planning to configure RAID arrays, ensure that you have the minimum number of available disks for each RAID level:

RAID 0

Minimum of one drive per array

RAID 5

Minimum of three drives per array

RAID 6

Minimum of four drives per array

RAID 10

Minimum of two drives per array

Connecting the SAS adapter to the 5887 disk drive enclosure

Learn how to connect the 5887 disk drive enclosure.

Complete the following steps to connect the 5887 disk drive enclosure to one or more Serial-attached SCSI (SAS) adapters, to a system with support for a SAS disk drawer, or to a PCIe storage enclosure with support for a SAS disk drawer.

For additional information related to SAS cabling and cabling configurations, see Serial-attached SCSI cable planning (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7had/p7hadsascabing.htm>).

1. If the system is powered on, you must do one of the following actions, depending on the supported functions of your operating system:
 - Unconfigure the adapters to which you are connecting the 5887 disk drive enclosure.
 - Power off the adapters to which you are connecting the 5887 disk drive enclosure.
 - Power off the logical partitions or systems that own the adapters to which you are connecting the 5887 disk drive enclosure.

To do one of these required actions, choose one of the following options.

Important: If you are connecting to two SAS adapters or to the two internal RAID controllers of a PCIe storage enclosure, both of the SAS adapters or RAID controllers must be powered off or unconfigured.

- Power off the system or the logical partitions that own the adapters or PCIe storage enclosures by completing the steps in Stopping a system or logical partition (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7haj/crustopsys.htm>).

Note: Use this option if one of the following conditions apply:

- Your system model does not support slot power control.
- Your adapters are not located in an I/O enclosure that supports slot power control.
- You are unable to tolerate a temporary loss of access to other disk devices that might already exist on the same adapters.
- For all other situations, choose one of the following options:
 - If you are connecting to SAS adapters, complete the following steps:
 - a. Unconfigure the two SAS adapters.
 - b. Connect the SAS cables from the 5887 disk drive enclosure to the SAS adapters.
 - c. Reconfigure the SAS adapters.
 - If you are connecting to a PCIe storage enclosure, complete the following steps:
 - a. Unconfigure the two internal RAID controllers in the PCIe storage enclosure.

- b. Connect the SAS cables from the the 5887 disk drive enclosure to the PCIe storage enclosure.
- c. Reconfigure the internal RAID controllers in the PCIe storage enclosure.
- If you are unable to use any of the previous options, choose one of the following options:
 - If you are connecting to SAS adapters, complete the following steps:
 - a. Power off the two SAS adapters.
 - b. Connect the SAS cables from the 5887 disk drive enclosure to the SAS adapters.
 - c. Power on the SAS adapters.
 - d. Configure the SAS adapters and devices.
 - If you are connecting to a PCIe storage enclosure, complete the following steps:
 - a. Power off the two internal RAID controllers in the PCIe storage enclosure.
 - b. Connect the SAS cables from the 5887 disk drive enclosure to the PCIe storage enclosure.
 - c. Power on the internal RAID controllers in the PCIe storage enclosure.
 - d. Configure the internal RAID controllers in the PCIe storage enclosure and devices.
- 2. Confirm the factory-set mode of the 5887 disk drive enclosure by using information printed on stickers at the rear of the enclosure. Stickers are attached to the lower-left shelf of the chassis (**A**) and the center support between the enclosure services manager (ESM) modules (**B**). The stickers indicate whether the enclosure is set to mode 1, mode 2, or mode 4, as shown in Figure 33.




Figure 33. Mode sticker locations at the rear of the 5887 disk drive enclosure

3. Confirm that all adapters that you need to connect to the 5887 disk drive enclosure have been installed in the system. If the adapters have not been installed, complete the adapter installation procedure for your system before continuing with this task.
4. If the server requires an internal cable to be installed to produce an external SAS port for connection with the 5887 disk drive enclosure, confirm that the installation has been performed. Review the following list of systems that require the installation of an external SAS port, and then complete the action for your system, if required:

Remember: When installing or confirming the use of an external SAS port, record the location of the external SAS port on the server. Later in this procedure, you are instructed to install the external SAS cable to this server connector location.

- For installing the external SAS port on the IBM Power 710 Express (8231-E2B, 8231-E1C, and 8231-E1D) or IBM Power 730 Express (8231-E2B, 8231-E2C, and 8231-E2D), see Installing the external SAS port in the 8231-E2B, 8231-E1C, 8231-E1D, 8231-E2C, or 8231-E2D (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hal/installsasport_71x_73x.htm).
 - For installing the external SAS port on the IBM Power 720 Express (8202-E4B, 8202-E4C, and 8202-E4D) or IBM Power 740 Express (8205-E6B, 8205-E6C, and 8205-E6D), see Installing the external SAS port in the 8202-E4B, 8202-E4C, 8202-E4D, 8205-E6B, 8205-E6C, or 8205-E6D system (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hal/installsasport_72x_74x.htm).
 - For installing the external SAS port on the IBM Power 750 Express (8233-E8B) or IBM Power 755 (8236-E8C), see Installing the external SAS port on an 8233-E8B or 8236-E8C (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hal/dasd550backplanesasfeature.htm>).
 - For installing the external SAS port on the IBM Power 770 (9117-MMB, 9117-MMC, and 9117-MMD) or IBM Power 780 (9179-MHB, 9179-MHC, and 9179-MHD), see Installing the external SAS port in the 9117-MMB, 9117-MMC, 9117-MMD, 9179-MHB, 9179-MHC, or 9179-MHD (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hal/9117installsasport.htm>).
5. Confirm the configuration that you will use to cable the SAS adapter to the 5887 disk drive enclosure as follows:
 - See the common SAS adapter configurations that are listed in step 7 to determine if one of them applies to your requirements.
 - To learn about other specific SAS adapter configurations for the 5887 disk drive enclosure, see SAS cabling for the 5887 drawer (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7had/p7hadsascabling5887.htm>).
 6. Locate the connection for each adapter for which you are using an external SAS cable to connect the 5887.

Note: The configuration figures in step 7 use adapter cards to represent the external server connections. The adapter card can represent one of the following connection types:

- The external port of an adapter that you confirmed in step 3 on page 74.
 - The external port of an internal adapter cable that you confirmed in step 4 on page 74.
7. Choose one of the following options to cable the SAS adapter:
 - To perform a mode 1 connection of one 5887 disk drive enclosure by using a YI cable to a system with support for a single SAS disk drawer, go to step 8 on page 76.
 - To perform a mode 1 connection of one 5887 disk drive enclosure by using a YO cable to a single SAS adapter, go to step 9 on page 76.
 - To perform a mode 1 connection of two 5887 disk drive enclosures by using YO cables to a single SAS adapter, go to step 10 on page 77.
 - To perform a mode 1 connection of one 5887 disk drive enclosure by using YO cables to a SAS adapter pair, go to step 11 on page 78.
 - To perform a mode 1 connection of two 5887 disk drive enclosures by using YO cables to a SAS adapter pair, go to step 12 on page 79.
 - To perform a mode 1 connection of one 5887 disk drive enclosure by using EX cables to a PCIe storage enclosure with support for SAS disk drawers, go to step 13 on page 80.
 - To perform a mode 1 connection of two 5887 disk drive enclosures by using EX cables to a PCIe storage enclosure with support for SAS disk drawers, go to step 14 on page 81.

Note: A mode 2 or mode 4 connection of 5887 disk drive enclosures to a PCIe storage enclosure is not supported.

- To perform a mode 2 connection of one 5887 disk drive enclosure by using YO cables to two independent SAS adapters, go to step 15 on page 81.
- To perform a mode 2 connection of one 5887 disk drive enclosure by using X cables to two SAS adapter pairs, go to step 16 on page 82.
- To perform a mode 4 connection of one 5887 disk drive enclosure by using X cables to four independent SAS adapters, go to step 17 on page 83.

If your SAS configuration requirements are not supported by any of these options, go to step 18 on page 85.

8. Perform the mode 1 connection of one 5887 disk drive enclosure (A) by using a YI cable (B) to a system with support for a single SAS drawer (C) as shown in Figure 34.

Note:

- The system with support for a single SAS drawer (C) has access to all 24 drive bays.




Figure 34. Mode 1 connection of one 5887 disk drive enclosure by using a YI cable to a system with support for a single SAS drawer

Go to step 18 on page 85.

9. Perform the mode 1 connection of one 5887 disk drive enclosure (A) by using a YO cable (B) to a single SAS adapter (C) as shown in Figure 35 on page 77.

Note:

- The single SAS adapter (C) has access to all 24 drive bays.




Figure 35. Mode 1 connection of one 5887 disk drive enclosure by using a YO cable to a single SAS adapter

Go to step 18 on page 85.

10. Perform the mode 1 connection of two 5887 disk drive enclosures (**A and B**) by using YO cables (**C and D**) to a single SAS adapter (**E**) as shown in Figure 36 on page 78.

Note:

- The single SAS adapter (**E**) has access to all 48 drive bays.




Figure 36. Mode 1 connection of two 5887 disk drive enclosures by using YO cables to a single SAS adapter

Go to step 18 on page 85.

11. Perform the mode 1 connection of one 5887 disk drive enclosure (A) by using YO cables (B) to a SAS adapter pair (C) as shown in Figure 37 on page 79.

Notes:

- Each adapter in the SAS adapter pair (C) has access to the other adapter and to all 24 drive bays.
- If 6 Gb SAS adapters are being used, they must be cabled by using 6 Gb SAS cables.
- For SAS adapter pairs, you can attach the cables to any adapter port as long as you use the same port in both adapters.




Figure 37. Mode 1 connection of one 5887 disk drive enclosure by using YO cables to a SAS adapter pair

Go to step 18 on page 85.

12. Perform the mode 1 connection of two 5887 disk drive enclosures (**A and B**) by using YO cables (**C and D**) to a SAS adapter pair (**E**) as shown in Figure 38 on page 80.

Notes:

- Each adapter in the SAS adapter pair (**C**) has access to the other adapter and to all 48 drive bays.
- If 6 Gb SAS adapters are being used, they must be cabled by using 6 Gb SAS cables.
- For SAS adapter pairs, you can attach the cables to any adapter port as long as you use the same port in both adapters.




Figure 38. Mode 1 connection of two 5887 disk drive enclosures by using YO cables to a SAS adapter pair

Go to step 18 on page 85.

13. Perform the mode 1 connection of one 5887 disk drive enclosure (B) by using EX cables (C) to a PCIe storage enclosure with support for SAS disk drawers (A) as shown in Figure 39.

Notes:

- Each internal adapter in the PCIe storage enclosure (A) has access to the other adapter and to all disk drive bays.
- Both EX cables (C) from the same 5887 disk drive enclosure must be attached to the same numbered SAS port on the PCIe storage enclosure as shown in Figure 39.




Figure 39. Mode 1 connection of one 5887 disk drive enclosure by using EX cables to a PCIe storage enclosure

Go to step 18 on page 85.

14. Perform the mode 1 connection of two 5887 disk drive enclosures (**B**) by using EX cables (**C**) to a PCIe storage enclosure with support for SAS disk drawers (**A**) as shown in Figure 40.

Notes:

- The AIX and Linux operating systems support the attachment of up to two 5887 disk drive enclosures. The IBM i operating system supports the attachment of one 5887 disk drive enclosure.
- Each internal adapter in the PCIe storage enclosure (**A**) has access to the other adapter and to all disk drive bays.
- Both EX cables (**C**) from the same 5887 disk drive enclosure must be attached to the same numbered SAS port on the PCIe storage enclosure as shown in Figure 40.




Figure 40. Mode 1 connection of two 5887 disk drive enclosures by using EX cables to a PCIe storage enclosure

Go to step 18 on page 85.

15. Perform the mode 2 connection of one 5887 disk drive enclosure (**A**) by using YO cables (**B**) to two independent SAS adapters as shown in Figure 41 on page 82.

Notes:

- Independent SAS adapter 1 (**C**) does not have access to the other independent adapter and has access only to drive bays D1 - D12.
- Independent SAS adapter 2 (**D**) does not have access to the other independent adapter and has access only to drive bays D13 - D24.




Figure 41. Mode 2 connection of one 5887 disk drive enclosure by using YO cables to two independent SAS adapters

Go to step 18 on page 85.

16. Perform the mode 2 connection of one 5887 disk drive enclosure (A) by using X cables (B) to two SAS adapter pairs as shown in Figure 42 on page 83.

Notes:

- Each adapter in the SAS adapter pair 1 (C) has access to the other adapter and to drive bays D1 - D12.
- Each adapter in the SAS adapter pair 2 (D) has access to the other adapter and to drive bays D13 - D24.
- If 6 Gb SAS adapters are being used, they must be cabled by using 6 Gb SAS cables.
- For SAS adapter pairs, you can attach the cables to any adapter port as long as you use the same port in both adapters.




Figure 42. Mode 2 connection of one 5887 disk drive enclosure by using X cables to two SAS adapter pairs

Go to step 18 on page 85.

17. Perform the mode 4 connection of one 5887 disk drive enclosure (A) by using X cables (B) to four independent SAS adapters as shown in Figure 43 on page 84.

Notes:

- The cable that plugs into independent SAS adapter 1 (C) contains a label with the P1 identifier (G), as shown in Figure 44 on page 84. This adapter does not have access to any other independent adapter and has access only to drive bays D1 - D6.
- The cable that plugs into independent SAS adapter 2 (D) contains a label with the P2 identifier (G), as shown in Figure 44 on page 84. This adapter does not have access to any other independent adapter and has access only to drive bays D7 - D12.
- The cable that plugs into independent SAS adapter 3 (E) contains a label with the P1 identifier (G), as shown in Figure 44 on page 84. This adapter does not have access to any other independent adapter and has access only to drive bays D13 - D18.
- The cable that plugs into independent SAS adapter 4 (F) contains a label with the P2 identifier (G), as shown in Figure 44 on page 84. This adapter does not have access to any other independent adapter and has access only to drive bays D19 - D24.




Figure 43. Mode 4 connection of one 5887 disk drive enclosure by using X cables to four independent SAS adapters




Figure 44. Labels for SAS adapter cables showing P1 and P2 identifiers

Go to step 18 on page 85.

18. For additional information to plan or complete your SAS cable installation, see Serial-attached SCSI cable planning (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7had/p7hadsascablng.htm>).

Remember:

- Connect power cables and apply power to the 5887 disk drive enclosure.
- Power on the system, partition, or reconfigure the adapters depending what option you chose in the beginning of this procedure.
- Add the disk drives to the operating system you are using.
- Verify the new configuration. To verify that the system or logical partition recognizes the disk drive enclosure, see Verifying the installed part (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7haj/p7hajsmverify.htm>).

Removing your disk drive enclosures

Learn about removing a disk drive enclosure from a rack.

For information about physically removing a disk drive enclosure from a rack, see “Removing an enclosure from the rack” on page 46.

PCIe storage enclosures

Use this information to learn about connecting and configuring PCI Express (PCIe) storage enclosures to system units.

Overview for PCIe storage enclosures

PCI Express (PCIe) storage enclosures provide additional storage capacity to the system in the form of hard disk drives (HDDs) or solid-state drives (SSDs).

PCIe storage enclosures are connected to the system with PCIe cables. The PCIe cables extend the PCIe buses that are generated by PCI host bridges in the system, and transport PCIe protocol information and data between the system and the storage controllers in the PCIe storage enclosure.

Each PCIe cable connects to a PCIe connector on the system unit. The connector can be on a pluggable card, such as a PCIe GX++ adapter, or can be mounted directly on the system unit chassis. In either case, the connector is wired to a PCI host bridge that controls the PCIe protocol. The other end of the cable connects to a PCIe connector on the PCIe storage enclosure that is wired to a storage controller in the PCIe storage enclosure.

The PCIe interconnection between a system and a PCIe storage enclosure is point to point. Although a PCIe storage enclosure is connected to the system unit with two PCIe cables, there is no redundancy or cable failover capability in the PCIe links. However, access to all drives in the PCIe storage enclosure is maintained in the event of a single PCIe link failure because of the dual controllers and fully redundant SAS fabric within the PCIe storage enclosure.

Furthermore, there are multiple ways to connect the two PCIe cables from a PCIe storage enclosure to the system, some of which result in a better high availability characteristics than others. For example, connecting the two PCIe cables from a PCIe storage enclosure to different PCIe GX++ adapters in the system creates a more highly available configuration than connecting the cables to the same PCIe GX++ adapter. If the PCIe GX++ adapters are in different processor drawers, such as is possible with multi-drawer mid-range systems, such as 9117-MMC and 9179-MHC, the configuration is even more highly available.

With proper software support in the operating systems, it is even possible to connect a PCIe storage enclosure to two different systems for maximum high availability.

Important:

- If additional PCIe GX++ adapters are needed to accommodate new PCIe storage enclosures, see the instructions that came with the PCIe GX++ adapters for information about how to install them. Installing PCIe GX++ adapters must be done separately from installing PCIe storage enclosures.
- Before you change the system configuration, such as adding or removing PCIe storage enclosures, it is a good practice to resolve all open serviceable events. If that is not possible, evaluate the open serviceable events that cannot be resolved to determine whether or not they might interfere with the procedure, and record them as evidence that they were not caused by the procedure.

Related reference:

“Connecting PCIe storage enclosures with PCIe cables” on page 99
Learn how PCIe cables are used to connect PCIe storage enclosures to servers.

Adding PCIe storage enclosures

Learn how to add a PCIe storage enclosure to the system and verify that it is functioning.

Depending on the feature of the PCIe storage enclosure that you are working with, you might be able to add it to the system with the system powered on and the logical partitions active.

- If you are adding a 5888 PCIe storage enclosure, the system must be powered off to complete the addition.
- If you are adding any other PCIe storage enclosure, you can add it either with the system powered on and partitions active or with the system powered off.

Preparing to add a PCIe storage enclosure

Use this information to plan the installation and configuration of a PCI Express (PCIe) storage enclosure.

This topic provides information about adding a PCIe storage enclosure with PCIe cables. You can perform this task yourself or contact a service provider for this service.

To prepare to add the PCIe storage enclosure, complete the following steps:

1. Decide where you want to install the new PCIe storage enclosures. For details, see “Where to install a PCIe storage enclosure.”
2. Unpack the PCIe storage enclosures by using the unpacking instructions.
3. Identify the cables. For details, see “Identifying PCIe cables.”
4. Plan for cable layout. When you decide where to place the cables, follow your site plan and keep the following items in mind:
 - Avoid creating a safety hazard.
 - Avoid damaging the cables.
 - Avoid placing cables parallel to high-voltage lines.
5. Continue with “Adding PCIe storage enclosures” on page 87.

Where to install a PCIe storage enclosure

Before you begin the installation process, plan where you want to install the PCIe storage enclosures.

When deciding where to install a PCIe storage enclosure, consider several elements that include size, security, and environmental factors. For more information, see Site preparation and physical planning (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ebe/p7ebegenconsiderations.htm>).

Identifying PCIe cables

Use this information to help identify the PCIe cables for the PCIe storage enclosure.

The system uses PCIe cables to transmit data and control information to and from the PCIe storage enclosure.

Table 7. PCIe cables

Cable feature	Part number	Length	PCIe storage enclosures supported	Systems supported
EN05	46K3769	1.5 m (4.9 ft.)	5888	8231-E1C 8231-E2C 8202-E4C 8205-E6C
			EDR1	8202-E4D 8205-E6D 8231-E1D 8231-E2D 8246-L2T 8248-L4T 8268-E1D 8408-E8D 8412-EAD 9109-RMD 9117-MMD 9179-MHD
EN07	46K3770	3.0 m (9.8 ft.)	5888	8231-E1C 8231-E2C 8202-E4C 8205-E6C
			EDR1	8202-E4D 8205-E6D 8231-E1D 8231-E2D 8246-L2T 8248-L4T 8268-E1D 8408-E8D 8412-EAD 9109-RMD 9117-MMD 9179-MHD

Table 7. PCIe cables (continued)

Cable feature	Part number	Length	PCIe storage enclosures supported	Systems supported
EN08	41U8581	8.0 m (26.2 ft.)	EDR1	8202-E4D 8205-E6D 8231-E1D 8231-E2D 8246-L2T 8248-L4T 8268-E1D 8408-E8D 8412-EAD 9109-RMD 9117-MMD 9179-MHD

PCIe storage enclosure configuration rules

Learn about important configuration rules that must be followed when connecting PCIe storage enclosures to system units.

When connecting a PCIe storage enclosure, observe the configuration rules in the following list:

- PCIe storage enclosures are connected to system units in a point-to-point topology.
- Only one PCIe storage enclosure can be connected to a system-side PCIe connector. Daisy-chain cabling of PCIe storage enclosures is not supported.
- PCIe storage enclosures can be connected to two PCIe connectors on the same PCIe GX++ adapter, or split across two different PCIe GX++ adapters. Connecting the two PCIe cables to different PCIe GX++ adapters can improve system serviceability and availability.

The following table specifies PCIe storage enclosure limits.

Table 8. PCIe storage enclosures limits

Machine type and model	Supported PCIe storage enclosures	Maximum number of PCIe GX++ adapters	Maximum number of PCIe storage enclosures
8202-E4C	5888	1	1
8231-E1C	5888	1	1 Note: The 8231-E1C must be one of two systems in a high-availability configuration in order to support a PCIe storage enclosure.
8231-E2C	5888	2	1
8205-E6C	5888	2	2

Table 8. PCIe storage enclosures limits (continued)

Machine type and model	Supported PCIe storage enclosures	Maximum number of PCIe GX++ adapters	Maximum number of PCIe storage enclosures
8202-E4D 8246-L2T	EDR1	1	1
8231-E1D	EDR1	1	1 Note: The 8231-E1D must be one of two systems in a high-availability configuration in order to support a PCIe storage enclosure.
8268-E1D	EDR1	1	1 Note: The 8268-E1D must be one of two systems in a high-availability configuration in order to support a PCIe storage enclosure.
8231-E2D	EDR1	2	1
8205-E6D 8248-L4T Note: The 16-way, 24-way, and 32-way configurations support PCIe storage enclosures. The 8-way configuration does not support PCIe storage enclosures. 8408-E8D Note: The 16-way, 24-way, and 32-way configurations support PCIe storage enclosures. The 8-way configuration does not support PCIe storage enclosures. 9109-RMD Note: The 24-way, 36-way, and 48-way configurations support PCIe storage enclosures. The 12-way configuration does not support PCIe storage enclosures.	EDR1	2	2

Table 8. PCIe storage enclosures limits (continued)

Machine type and model	Supported PCIe storage enclosures	Maximum number of PCIe GX++ adapters	Maximum number of PCIe storage enclosures
8412-EAD 9117-MMD 9179-MHD	EDR1	The maximum number of PCIe GX++ adapters is equal to two PCIe GX++ adapters per processor drawer in the system configuration: 1 processor drawer = 2 PCIe GX++ adapters 2 processor drawers = 4 PCIe GX++ adapters 3 processor drawers = 6 PCIe GX++ adapters 4 processor drawers = 8 PCIe GX++ adapters	The maximum number of PCIe storage enclosures is equal to two PCIe storage enclosures per processor drawer in the system configuration: 1 processor drawer = 2 PCIe storage enclosures 2 processor drawers = 4 PCIe storage enclosures 3 processor drawers = 6 PCIe storage enclosures 4 processor drawers = 8 PCIe storage enclosures

Adding a PCIe storage enclosure with the system powered on

Learn how to add a PCIe storage enclosure to the system with the system powered on.

Important: To use this procedure to add a PCIe storage enclosure to the system, the enclosure must be a model that can be added with the system powered on. Review the following information to learn which options you have for adding your enclosure model to the system:

- If you are adding a 5888 PCIe storage enclosure, you must use “Adding a PCIe storage enclosure with the system powered off” on page 98 instead of this procedure.
- If you are adding any other PCIe storage enclosure, you can add it with the system powered on using this procedure.

Prerequisites: If you have not reviewed the information in “Overview for PCIe storage enclosures” on page 87, do so now.

Notes:

- To learn about the different options for cabling PCIe storage enclosures to the system, see “Connecting PCIe storage enclosures with PCIe cables” on page 99.
- Steps in this topic require you to have connector information for the hardware in your configuration. For details about PCIe storage enclosure or system connector locations, see “Connector locations” on page 126.

To add a PCIe storage enclosure to the system with the system powered on, complete the following steps:

1. Record the current date and time for use later in the procedure when checking for serviceable events.
2. Verify the existing PCIe hardware topology by completing the steps in the applicable procedure:
 - If you have a 5888 PCIe storage enclosure, see “Verifying the PCIe hardware topology” on page 102.
 - If you have an EDR1 PCIe storage enclosure, see “Verifying the PCIe hardware topology” on page 102.

All PCIe hardware topology problems that are discovered in this step must be corrected before you continue with this procedure.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

3. If the system has a rear cover, remove or open it.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

4. Locate the PCIe GX++ adapter or adapters to which the PCIe storage enclosure will be connected.

Note: This procedure assumes that the PCIe GX++ adapters are already installed. If the adapters are not installed, stop now and follow the instructions that came with the adapters to install them. After the PCIe GX++ adapters are installed, continue with the steps in this procedure to connect the PCIe storage enclosure.

5. Using PCIe cables, connect the PCIe storage enclosure to the GX adapters that were identified in step 4.

Note: If you are connecting a PCIe cable to an unused PCIe connector that is covered with a plastic or metal clip, remove the clip before installing the PCIe cable.

6. Identify the PCIe link information that is associated with the connections to the PCIe storage enclosure by completing the following steps.

Note: Two PCIe links connect a PCIe storage enclosure.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system. Each system will host one of the PCIe links that connect the PCIe storage enclosure.

- a. Determine the location codes of the PCIe connectors on the system to which you connected the PCIe storage enclosure.

- b. Choose one of the following options:

- If your system is managed by a Hardware Management Console (HMC), complete the following steps:

- 1) From the navigation bar, expand **Systems Management**.
- 2) Click **Servers**.
- 3) Select the server you are working with.
- 4) In the Tasks area, expand **Hardware Information**.
- 5) Click **PCIe Hardware Topology**.

- If your system is not managed by an HMC, complete the following steps:

- 1) Access the Advanced System Management Interface (ASMI) by using an authority level of administrator or authorized service provider. For details about using the ASMI, see *Managing the Advanced System Management Interface* (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hby/asmi.htm>).
- 2) In the ASMI navigation area, expand **System Configuration**.
- 3) Click **PCIe Hardware Topology**.

- c. Scan the PCIe Hardware Topology data for entries with Host Port column values that match the location codes of the PCIe GX++ adapter connectors that you connected PCIe cables to in step 5. For each match, complete the following steps:

- 1) Select the entry and click **Identify Indicators**.
- 2) On the Identify Indicators display, select the entry with a Location column value that matches the Host Port column value from the PCIe Hardware Topology entry that you selected in step 6c1.
- 3) Click **Activate LED**.
- 4) Record the link identifier and the I/O slot identifier string for the PCIe link entry that you identified in step 6c2. The link identifier is in the Link ID column, and the I/O slot identifier string is in the I/O Slots column.

- d. Visually inspect the identify indicators for the PCIe connectors to which you connected the PCIe storage enclosure to verify that they are active.
- e. For each identify indicator that you recorded in step 6c4 on page 93 that is active, complete the following steps:
 - 1) On the Identify Indicators display, select the entry with a Location column value that matches the Host Port column value from the PCIe Hardware Topology entry or entries that you selected in step 6c1 on page 93.
 - 2) Click **Deactivate LED**.
- f. For each identify indicator that you recorded in step 6c4 on page 93 that is not active, complete the following steps.

Note: Identify indicators are not active if the PCIe cables are connected to the wrong PCIe connectors or the PCIe connector location codes were calculated incorrectly.

- 1) On the Identify Indicators display, select the entry with a Location column value that matches the Host Port column value from the PCIe Hardware Topology entry or entries that you selected in step 6c1 on page 93.
 - 2) Click **Deactivate LED**.
 - 3) Restart the procedure beginning at step 4 on page 93. If the problem persists with the entry or entries that you selected in step 6c1 on page 93, contact your next level of support for assistance.
7. If the system has a rear cover, close it.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

8. If the power cords for the PCIe storage enclosure are not connected to the power source, connect them now. Then connect the power cords to the power supplies on the PCIe storage enclosure.
9. During the time that the PCIe storage enclosure is powering on and initializing, complete the following steps:
 - a. Verify that the green chassis power light-emitting diode (LED) comes on solid. This LED is located on the front of the enclosure, and it has an overlay with a lightning bolt symbol.
 - b. Verify that the green **AC good** LED and **DC good** LED on each of the two power supplies come on solid. These LEDs are located on the power supplies near the connectors for the power cords, and are clearly labeled **AC** and **DC**.
 - c. Verify that the amber fault LED on each power supply is off. This LED is labeled with an exclamation point (!).

Important: The PCIe storage enclosure usually requires no more than 5 minutes to power on and initialize. If any of the LED states previously described cannot be verified after 5 minutes have elapsed, there could be problems with the PCIe storage enclosure. Verify that the power cords are securely connected to the power supplies and the power source. If the problem persists, contact your next level of support for assistance.

10. For each of the two input/output (I/O) slots that you recorded in step 6c4 on page 93, power on the I/O slot by choosing one of the following options.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system. Each system will host one of the I/O slots that connect the PCIe storage enclosure.

- If the I/O slot is assigned to a partition that is powered on, choose one of the following options.

Note: If your system is not managed by an HMC and is powered on to the firmware running state, choose one of the following options. The system is in firmware running state if there are active partitions.

- If the I/O slot is assigned to a partition that is powered on and running the Virtual I/O Server (VIOS) or AIX operating system, continue with step 11.
- If the I/O slot is assigned to a partition that is powered on and running the IBM i operating system, continue with step 12 on page 96.
- If the I/O slot is assigned to a partition that is powered on and running the Linux operating system, continue with step 13 on page 96.
- If the I/O slot is not assigned to a partition that is powered on, continue with step 14 on page 96.

Note: If your system is powered on to firmware standby state, choose this option. The system is in firmware standby state if there are no active partitions.

11. If you are using AIX, complete the following steps:

- a. Log in to the operating system with admin or service level authority. If you need assistance, contact the system administrator.
- b. Choose one of the following options:
 - If the partition is running the VIOS operating system, type `diagmenu` at the VIOS command prompt, and press Enter.
 - If the partition is running the AIX operating system, type `diag` at the AIX command prompt, and press Enter.
- c. On the DIAGNOSTIC OPERATING INSTRUCTIONS display, press Enter to continue.
- d. On the FUNCTION SELECTION display, use the arrow keys to select **Task Selection (Diagnostics, Advanced Diagnostics, Service Aids, etc.)**, and press Enter.
- e. On the TASKS SELECTION LIST display, use the arrow keys to select **Hot Plug Task**, and press Enter.
- f. On the Hot Plug Task display, use the arrow keys to select **PCI Hot Plug Manager**, and press Enter.
- g. On the PCI Hot Plug Manager display, use the arrow keys to select **Add a PCI Hot Plug Adapter**, and press Enter.
- h. On the Add a PCI Hot Plug Adapter display, use the arrow keys to select the entry that has the identifier string of the I/O slot in the Slot column, and press Enter.
- i. Follow the online instructions to complete the add operation. However, do not perform the steps to physically add the adapter.
- j. Press **F3** (Exit) to return to the PCI Hot Plug Manager display.

Note: If your terminal emulation does not support **F3** to exit, press **ESC**, then type 3 to exit.

- k. On the PCI Hot Plug Manager display, use the arrow keys to select **Install/Configure Devices Added After IPL**, and press Enter.
- l. On the Install/Configure Devices Added After IPL display, verify that **INPUT device/directory for software** is set to none, and press Enter.
- m. On the COMMAND STATUS display, wait for the command status to change from running to OK.
- n. Choose one of the following options:
 - If you have another I/O slot to power on, complete the following steps:
 - 1) Press **F3** (Exit) twice to return to the PCI Hot Plug Manager display.

Note: If your terminal emulation does not support **F3** to exit, press **ESC**, type 3, press **ESC**, then type 3 to exit.

- 2) Return to step 10 on page 94.
- If you do not have another I/O slot to power on, complete the following steps:
 - 1) Press **F10** twice to return to the operating system command prompt.
 - 2) Continue with step 15 on page 97.

12. If you have IBM i, complete the following steps:
 - a. Log in to the operating system with admin or service level authority. If you need assistance, contact the system administrator.
 - b. At the IBM i command prompt, type `strsst` and press Enter.
 - c. Type the service tools user ID and service tools password on the System Service Tools (SST) Sign On display, and press Enter.
 - d. Select the **Start a service tool** option, and press Enter.
 - e. Select the **Hardware service manager** option, and press Enter.
 - f. Select the **Packaging hardware resources (systems, frames, cards)** option, and press Enter.
 - g. On the Packaging Hardware Resources display, locate the entry that contains the I/O slot. To make this identification, scan the display for an entry that has a Description value of System Unit and a Unit ID that matches the U label portion of the I/O slot identifier string. The U label portion is the substring that begins with U and continues until the first dash (-) character.
 - h. Type 9 in the Opt column for the system unit entry that you located in step 12g, and press Enter.

Note: Specifying option 9 starts the Hardware contained within package view for the selected packaging hardware resource.

- i. On the Packaging Hardware Resources display, locate the entry for the I/O slot. To make this identification, scan the display for an entry that has a Location column value that, when appended to the Unit ID value in the upper right corner, forms the complete I/O slot identifier string.
- j. Type 3 in the Opt column for the entry that you located in step 12i, and press Enter.

Note: Specifying option 3 starts the Hardware Resource Concurrent Maintenance menu for the selected packaging hardware resource.

- k. On the Hardware Resource Concurrent Maintenance display, start the Power off domain function by pressing **F9**.
- l. Follow the online instructions to complete the power-off operation.
- m. On the Hardware Resource Concurrent Maintenance display, start the Power on domain function by pressing **F10**.
- n. Follow the online instructions to complete the power-on operation.
- o. If you have another I/O slot to power on, return to step 10 on page 94. Otherwise, continue with step 15 on page 97.

13. If you are using Linux, complete the following steps:

- a. Log in to the operating system with admin or service level authority. If you need assistance, contact the system administrator.
- b. At the Linux command prompt, type `drmgr -c pci -a -s io_slot_identifier` (where *io_slot_identifier* is the identifier string of the I/O slot), and press Enter.
- c. Follow the online instructions to complete the add operation. However, do not perform the steps to physically add the adapter.
- d. At the Linux command prompt, type `vpdupdate`, and press Enter.
- e. If you have another I/O slot to power on, return to step 10 on page 94. Otherwise, continue with step 15 on page 97.

14. Complete the following steps:

- a. Choose one of the following options:
 - If your system is managed by an HMC, complete the following steps:
 - 1) From the navigation bar, expand **Systems Management**.
 - 2) Click **Servers**.
 - 3) Select the server you are working with.

- 4) In the Tasks area, expand **Hardware Information**.
- 5) Click **PCIe Hardware Topology**.
- If your system is not managed by an HMC, complete the following steps:
 - 1) Access the ASMI by using an authority level of administrator or authorized service provider. For details about using the ASMI, see Managing the Advanced System Management Interface (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hby/asmi.htm>).
 - 2) In the ASMI navigation area, expand **System Configuration**.
 - 3) Click **PCIe Hardware Topology**.
- b. Scan the PCIe Hardware Topology data for a value in the I/O Slots column that matches the identifier string of the I/O slot, and select the entry.
- c. Click **Probe Link**.

Important: If the Probe Link function fails, verify that the I/O slot is not owned by a partition that is powered on. Then choose from the following options:

- If you determine that the I/O slot is owned by a partition that is powered on, return to step 10 on page 94 and choose the alternative option that supports an I/O slot assigned to a partition that is powered on.
 - If you are certain the I/O slot is not owned by a partition that is powered on, contact your next level of support for assistance.
- d. If you have another I/O slot to power on, return to step 10 on page 94. Otherwise, continue with step 15.
15. Verify the PCIe hardware topology by completing the steps in the applicable procedure:
 - If you have a 5888 PCIe storage enclosure, see “Verifying the PCIe hardware topology” on page 102.
 - If you have an EDR1 PCIe storage enclosure, see “Verifying the PCIe hardware topology” on page 102.

While you are completing the verification steps of that procedure, confirm that the PCIe links that connect the PCIe storage enclosure being added have Link Status column values of `Operational` and valid I/O Enclosure Port column values. The I/O Enclosure Port column value looks similar to `Uxxxx.001.yyyyyyy-P1-Cz-T3`, where `xxxx` is the feature code of the PCIe storage enclosure connected by the PCIe link, `yyyyyyy` is the serial number of the PCIe storage enclosure connected by the PCIe link, and `z` is either 1 or 2. An example of an actual I/O Enclosure Port column value is `UEDR1.001.6BAH004-P1-C1-T3`.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

Requirement: All PCIe hardware topology problems that are discovered in this step must be corrected before continuing with this procedure.

16. Check for and handle new serviceable events that were generated during the procedure by completing the steps in Checking for new serviceable events (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_genfile_checkevents.htm). Use the date and time information that you recorded in step 1 on page 92 as the start time for searching for new serviceable events.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

17. Verify the new configuration. For details, see Verifying that the new configuration is functioning (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_pcistor_connverify.htm).

18. If the system is managed by an HMC, is not in the manufacturing default configuration, and the I/O slots associated with the connection to the PCIe storage enclosure are not assigned to logical partitions, you can now assign them. For information about adding I/O slots to partitions dynamically, see *Adding physical I/O devices and slots dynamically* (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hat/iphbldlpariopaddp6.htm>).

Adding a PCIe storage enclosure with the system powered off

Learn how to add a PCIe storage enclosure to the system with the system powered off.

Prerequisites: If you have not reviewed the information in “Overview for PCIe storage enclosures” on page 87, do so now.

Notes:

- Depending on the feature of PCIe storage enclosure that you are working with, you might be able to add it to the system with the system powered on and the logical partitions active. To learn which options you have for adding your PCIe storage enclosure to the system, read the following information:
 - If you are adding a 5888 PCIe storage enclosure, you must use this procedure to add the enclosure to the system with the system powered off.
 - If you are adding any other PCIe storage enclosure, you can use this procedure to add the PCIe storage enclosure to the system with the system powered off, or you can use “Adding a PCIe storage enclosure with the system powered on” on page 92 instead of this procedure to add it with the system powered on.
- To learn about the different options for cabling PCIe storage enclosures to the system, see “Connecting PCIe storage enclosures with PCIe cables” on page 99.
- Steps in this topic require you to have connector information for the hardware in your configuration. For details about PCIe storage enclosure or system connector locations, see “Connector locations” on page 126.

To add a PCIe storage enclosure to the system with the system powered off, complete the following steps:

1. If the system is not already powered off, remove power from the system. For details, see *Stopping a system or logical partition* (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7haj/crustopsys.htm>)

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

2. If the system has a rear cover, remove or open it.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

3. Locate the PCIe GX++ adapter or adapters to which the PCIe storage enclosure will be connected.

Note: This procedure assumes that the PCIe GX++ adapters are already installed. If the adapters are not installed, stop now and follow the instructions that came with the GX adapters to install them. After the PCIe GX++ adapters are installed, continue with the steps in this procedure to add the PCIe storage enclosure.

4. Using PCIe cables, connect the PCIe storage enclosure to the PCIe GX++ adapters that you identified in step 3.

Note: If you are connecting a PCIe cable to an unused PCIe connector that is covered with a plastic or metal clip, remove the clip before installing the PCIe cable.

5. If the system has a rear cover, close it.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

6. If the power cords for the PCIe storage enclosure are not connected to the power source, connect them now. Then connect the power cords to the power supplies on the PCIe storage enclosure.
7. Record the current date and time for use later in the procedure when checking for serviceable events.
8. Power on the system. For details, see Starting the system or logical partition (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7haj/crustartsys.htm>).
9. Verify the PCIe hardware topology by completing the steps in the applicable procedure:
 - If you have a 5888 PCIe storage enclosure, see “Verifying the PCIe hardware topology” on page 102.
 - If you have an EDR1 PCIe storage enclosure, see “Verifying the PCIe hardware topology” on page 102.

While you are completing the verification steps of that procedure, confirm that the PCIe links that connect the PCIe storage enclosure being added have Link Status column values of `Operational` and valid I/O Enclosure Port column values. The I/O Enclosure Port column value looks similar to `Uxxxx.001.yyyyyyy-P1-Cz-T3`, where `xxxx` is the feature code of the PCIe storage enclosure connected by the PCIe link, `yyyyyyy` is the serial number of the PCIe storage enclosure connected by the PCIe link, and `z` is either 1 or 2. An example of an actual I/O Enclosure Port column value is `UEDR1.001.6BAH004-P1-C1-T3`.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

Requirement: All PCIe hardware topology problems that are discovered in this step must be corrected before continuing with this procedure.

10. Check for and handle new serviceable events that were generated during the procedure by completing the steps in Checking for new serviceable events (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_genfile_checkevents.htm). Use the date and time information that you recorded in step 7 as the start time for searching for new serviceable events.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

11. Verify the new configuration. For details, see Verifying that the new configuration is functioning (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_pcistor_connverify.htm).
12. If the system is managed by an HMC, is not in the manufacturing default configuration, and the I/O slots associated with the connection to the PCIe storage enclosure are not assigned to logical partitions, you can now assign them. For information about adding I/O slots to partitions dynamically, see Adding physical I/O devices and slots dynamically (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/iphath/iphbldlpariopaddp6.htm>).

Connecting PCIe storage enclosures with PCIe cables

Learn how PCIe cables are used to connect PCIe storage enclosures to servers.

The following examples provide possible configurations for connecting PCIe storage enclosures to servers by using PCIe cables:

1. Connect a server that has one PCIe GX++ adapter to a single storage enclosure, as shown in the following figure.




Figure 45. Example of a server that has one PCIe GX++ adapter connected to a single storage enclosure

2. Connect a server that has two PCIe GX++ adapters to two storage enclosures, as shown in the following figure.




Figure 46. Example of a server that has two PCIe GX++ adapters connected to two storage enclosures

3. Connect a server that has two PCIe GX++ adapters to a single storage enclosure, as shown in the following figure.




Figure 47. Example of a server that has two PCIe GX++ adapters connected to a single storage enclosure

4. Connect a server that has two PCIe GX++ adapters to two storage enclosures for high availability, as shown in the following figure.




Figure 48. Example of a server that has two PCIe GX++ adapters connected to two storage enclosures for high availability

5. Connect two servers, each with a PCIe GX++ adapter, to a single storage enclosure for high availability, as shown in the following figure.

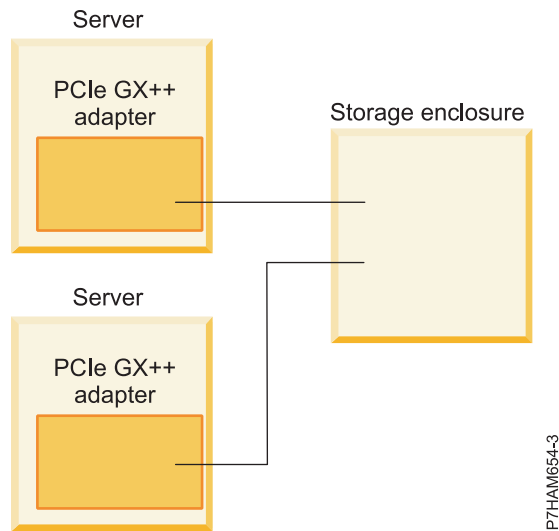


Figure 49. Example of two servers, each with a PCIe GX++ adapter, connected to a single storage enclosure for high availability

6. Connect two processor drawers, each with a PCIe GX++ adapter, to a single storage enclosure for high availability, as shown in the following figure.

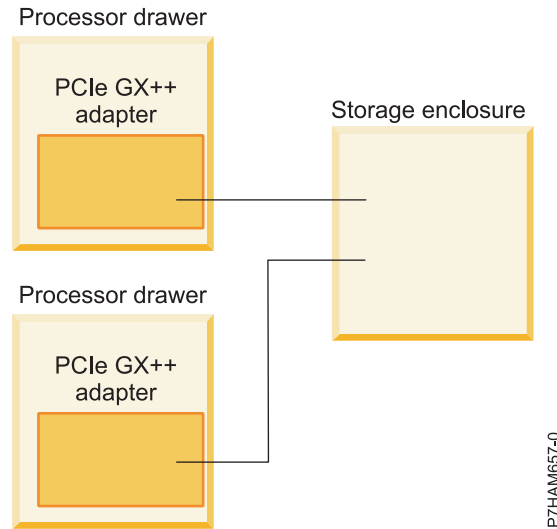


Figure 50. Example of two processor drawers, each with a PCIe GX++ adapter, connected to a single storage enclosure for high availability

Verifying the PCIe hardware topology

Learn how to verify that the PCIe links in the system are functioning properly.

The functioning of PCIe links can be verified in systems that are managed by a Hardware Management Console (HMC) and in systems that are not managed by an HMC.

Verifying the PCIe hardware topology with an HMC

Learn how to verify that there are no failed or degraded PCIe links by using a Hardware Management Console (HMC).

To use an HMC to verify that there are no failed or degraded PCIe links, complete the following steps with the system powered on:

1. Verify that no serviceable events include B7006Axx reference codes (where *x* is any character 0 - 9 or A - F) by completing the following steps:
 - a. In the navigation area of the HMC, select **Service Management**.
 - b. In the content area, select **Manage Serviceable Events**.
 - c. In the Manage Serviceable Events window, perform the following steps:
 - 1) Click **Open** for the **Serviceable Event Status** field.
 - 2) Select the machine type, model, and serial number (MTMS) of the server you are working with for the **Reporting MTMS** field.
 - 3) Click **ALL** for all other fields, and click **OK**.
 - d. Scan for any serviceable events that include B7006Axx reference codes, and choose one of the following options:
 - If no serviceable events include B7006Axx reference codes, close the Reference Code History window.
 - If any serviceable events include B7006Axx reference codes, there might be problems with the PCIe hardware topology. Perform problem analysis on the serviceable events to correct the problems before continuing. If you need assistance, contact your service provider.
2. Verify the status of the PCIe links by completing the following steps:
 - a. From the navigation bar, expand **Systems Management**.

- b. Click **Servers**.
- c. Select the server you are working with.
- d. From the Tasks area, expand **Hardware Information**.
- e. Click **PCIe Hardware Topology**.
- f. Scan the PCIe Hardware Topology data to identify rows that have a link status of Failed or Degraded.
- g. If you identified links in step 2f that have a link status of Failed or Degraded, the links must be repaired before you continue. To repair the links, repeat step 1 on page 102. If you already performed step 1 on page 102 or need assistance, contact your next level of support.

Verifying the PCIe hardware topology without an HMC

Learn how to verify that there are no failed or degraded PCIe links by using the AIX, Virtual I/O Server (VIOS), IBM i, or Linux operating system.

To verify that there are no failed or degraded PCIe links by using the AIX, Virtual I/O Server (VIOS), IBM i, or Linux operating system, complete the following steps with the system powered on:

1. Choose the operating system that best applies to your configuration.

Important: Multiple options might apply to your configuration. Choose only the first option that applies.

 - If your system has a partition running the VIOS operating system, continue with step 2.
 - If your system has a partition running the AIX operating system, continue with step 3.
 - If your system has a partition running the IBM i operating system, continue with step 4 on page 104.
 - If your system has a partition running the Linux operating system, continue with step 5 on page 104.
2. Use the VIOS error log to verify that there are no B7006Axx errors (where x is any character 0 - 9 or A - F) by completing the following steps.
 - a. Log in to the operating system with admin or service level authority. If you need assistance, contact the system administrator.
 - b. At the VIOS command prompt, type `diagmenu -d sysplanar0 -E xx` and press Enter, where xx is any number 1 - 60 that specifies the most recent number of days for which results are returned. The suggested range is 30 days.
 - c. Highlight the **Problem Determination** option in the DIAGNOSTIC MODE SELECTION display. Press Enter.
 - d. Search the results that are displayed to identify problems with B7006Axx reference codes. You might need to scroll down in the display to see all of the results.

Note: If the Problem Determination results contain serviceable events that have already been reported, a PREVIOUS DIAGNOSTICS RESULTS display is shown. Respond to the **Do you want to review these previously reported errors?** prompt by highlighting the YES response, and then press Enter.
 - e. Based on the results of your search, choose one of the following options:
 - If there are no problems that have B7006Axx reference codes, press Enter to return to the command line.
 - If there are problems that have B7006Axx reference codes, there might be problems with the PCIe hardware topology. Perform problem analysis to correct the problems before continuing. If you need assistance, contact your service provider.
 - f. Go to step 6 on page 105.
3. Use the AIX error log to verify that there are no B7006Axx errors (where x is any character 0 - 9 or A - F) by completing the following steps.

- a. Log in to the operating system with admin or service level authority. If you need assistance, contact the system administrator.
- b. At the AIX command prompt, type `diag -d sysplanar0 -E xx` and press Enter, where `xx` is any number 1 - 60 that specifies the number of days. This command returns results for the most recent number of days specified. The suggested range is 30 days.
- c. Highlight the **Problem Determination** option in the DIAGNOSTIC MODE SELECTION display. Press Enter.
- d. Search the results that are displayed to identify problems with B7006Axx reference codes. You might need to scroll down in the display to see all of the results.

Note: If the Problem Determination results contain serviceable events that have already been reported, a PREVIOUS DIAGNOSTICS RESULTS display is shown. Respond to the **Do you want to review these previously reported errors?** prompt by highlighting the YES response, and then press Enter.

- e. Based on the results of your search, choose one of the following options:
 - If there are no problems that have B7006Axx reference codes, press Enter to return to the command line.
 - If there are problems that have B7006Axx reference codes, there might be problems with the PCIe hardware topology. Perform problem analysis to correct the problems before continuing. If you need assistance, contact your service provider.
 - f. Go to step 6 on page 105.
4. Use the IBM i service action event log to verify that there are no B7006Axx errors (where `x` is any character 0 - 9 or A - F) by completing the following steps:
 - a. Log in to the operating system with admin or service level authority. If you need assistance, contact the system administrator.
 - b. At the IBM i command prompt, type `strsst` and press Enter.
 - c. Type your service tools user ID and service tools password on the System Service Tools (SST) Sign On display. Press Enter.
 - d. Click **Start a service tool** and press Enter.
 - e. Click **Hardware service manager** and press Enter.
 - f. Click **Work with service action log** and press Enter.
 - g. On the **Select Timeframe** display, change the **From: Date and Time** to the desired date and time range. The suggested range is 30 days.
 - h. Search for serviceable events with B7006Axx reference codes, and choose one of the following options:
 - If there are no serviceable events that have B7006Axx reference codes, press F3 (Exit) to return to the Hardware Service Manager display.
 - If there are serviceable events that have B7006Axx reference codes, there might be problems with the PCIe hardware topology. Perform problem analysis to correct the problems before continuing. If you need assistance, contact your service provider.
 - i. Go to step 6 on page 105.
 5. Use the Linux error log to verify that there are no B7006Axx errors (where `x` is any character 0 - 9 or A - F) by completing the following steps.
 - a. Log in as the root user. If you need assistance, contact the system administrator.
 - b. At the Linux command prompt, type `service log --query='refcode like "B7006A%" AND serviceable=1 AND closed=0'` and press Enter.
 - c. Search the results that are displayed for problems that have B7006Axx reference codes and for statuses that are open. If there are problems that have B7006Axx reference codes, there might be problems with the PCIe hardware topology. Perform problem analysis to correct the problems before continuing. If you need assistance, contact your service provider.

Note: You might need to scroll down in the display to see all of the results.

- d. Go to step 6.
6. Check for failed or degraded PCIe links by completing the following steps:
 - a. Access the Advanced System Management Interface (ASMI) by using an authority level of administrator or authorized service provider. For details about using the ASMI, see *Managing the Advanced System Management Interface* (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hby/asmi.htm>).
 - b. In the ASMI navigation area, expand **System Configuration**.
 - c. Click **PCIe Hardware Topology**.
 - d. Scan the PCIe Hardware Topology data to identify rows that have a link status of **Failed** or **Degraded**.
 - e. If you identified links in step 6d that have a link status of **Failed** or **Degraded**, the links must be repaired before you continue. To repair the links, repeat step 1 on page 103. If you already completed step 1 on page 103 or need assistance, contact your next level of support.

Verifying that the PCIe storage enclosure is functioning

Use this procedure to verify that a PCIe storage enclosure is functioning and is recognized by the system.

Prerequisite: To use this procedure, the PCIe storage enclosure must have power, must be connected to the system, and the system must be powered on.

1. Choose from the following options:
 - If your system is managed by an Hardware Management Console (HMC), continue with step 2.
 - If your system is not managed by an HMC, see *Verifying the installed part* (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7haj/p7hajhsmverify.htm>). When you are finished, continue with step 4 on page 106.
2. If your system is managed by an HMC, complete the following steps:
 - a. From the navigation area, expand **Systems Management**.
 - b. Click **Servers**.
 - c. From the content area, select the system with which you want to work.
 - d. Click **Tasks > Properties**.
 - e. Click the **I/O** tab.
3. Verify that the system recognizes the storage controllers in the PCIe storage enclosure by completing the following steps:
 - a. Scan the Slot column values to identify the entries that correspond to the connections to the PCIe storage enclosure. The following tips can be helpful in identifying these entries:
 - When an entry corresponds to a PCIe storage enclosure connection, the Slot column contains an identifier that resembles `Uwww.uuu.xxxxxx-P1-Cy-Tz-L1`, where *www* is the feature code of the system, *uuu* is a set of three numbers 0-9 or letters A-F, *xxxxxx* is the serial number of the system, *y* is a number greater than 0, and *z* is 1 or 2. An example of an actual Slot column value is `U8205.E6C.DQDMLTM-P1-C1-T1-L1`. The exact value of the identifier varies depending on the type of server and the location of the PCIe connector. Identifiers for slots that connect a PCIe storage enclosure terminate with `-T1-L1` or `-T2-L1`.
 - You can use identify indicators to visually confirm that an entry represents the connection to the storage enclosure being added. Record the PCIe connector location code, which is the slot identifier value minus the `-L1` at the end. Then, complete the procedure in *Using identify indicators to locate hardware components* (http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ham/p7ham_genfile_setidindicator.htm) to illuminate the identify indicator for the PCIe connector that is associated with that location code.

- b. For each entry that you identified in step 3a on page 105, verify that the Description column correctly identifies the type of adapter in the PCIe storage enclosure that is connected to that slot. Use the following information to verify whether the system correctly identifies the adapter:
 - 1) When the system recognizes the PCIe storage enclosure, the Description column contains a value that resembles PCIe2 3.1GB Cache RAID SAS Enclosure 6Gb x8. The exact string is dependent on the type of enclosure.
 - 2) When the system does not recognize the PCIe storage enclosure, the value in the Description column is Empty slot, Unknown, or PCI-to-PCI bridge. To ensure that the PCIe storage enclosure has been connected correctly, complete the following steps:
 - a) Verify that the PCIe cables are connected to the system and to the PCIe storage enclosure. If you find an incorrect connection at either end, connect the cables correctly.
 - b) Verify that the power cables are connected to the power source and to the power supplies on the PCIe storage enclosure, and verify that the green enclosure power light-emitting diode (LED) is on solid. If you find an incorrect connection at either end, connect the cables correctly.
 - c) If you needed to reconnect the PCIe cables or power cables, power the system off and back on. Then repeat step 3 on page 105.

If you detect no physical connection problems, or if the problem persists after a system reboot was required, contact your next level of support for assistance.

4. Return to the procedure that sent you here.

Checking for new serviceable events

Learn how to check for new serviceable events in various console and operating system environments.

To check for new serviceable events, complete the following steps:

1. Choose one the following options.

Important: Multiple options might apply to your configuration. Choose only the first option that applies.

- If your system is managed by a Hardware Management Console (HMC), complete the following steps:
 - a. In the navigation area, click **Serviceability > Manage Serviceable Events**.
 - b. Continue with step 2 on page 107.
- If your system is not managed by an HMC and has a partition running the Virtual I/O Server (VIOS) operating system, complete the following steps:
 - a. Log in to the operating system with admin or service level authority. If you need assistance, contact the system administrator.
 - b. Type `diagmenu -d sysplanar0 -E xx`, where *xx* is a number 1 - 60 that specifies the number of days, and press Enter. This command returns results for the most recent number of days specified. Specify a value that includes the time period during which the procedure was performed.
 - c. On the DIAGNOSTIC MODE SELECTION display, highlight the **Problem Determination** option, and press Enter.
 - d. Continue with step 2 on page 107.
- If your system is not managed by an HMC and has a partition running the AIX operating system, complete the following steps:
 - a. Log in to the operating system with admin or service level authority. If you need assistance, contact the system administrator.

- b. Type `diag -d sysplanar0 -E xx`, where `xx` is a number 1 - 60 that specifies the number of days, and press Enter. This command returns results for the most recent number of days specified. Specify a value that includes the time period during which the procedure was performed.
 - c. On the DIAGNOSTIC MODE SELECTION display, highlight the **Problem Determination** option, and press Enter.
 - d. Continue with step 2.
- If your system is not managed by an HMC and has a partition running the IBM i operating system, complete the following steps:
 - a. Log in to the operating system with admin or service level authority. If you need assistance, contact the system administrator.
 - b. At the command line of the IBM i session, type `strsst` and press Enter.
 - c. Type your service tools user ID and service tools password on the System Service Tools (SST) Sign On display, and press Enter. The service tools password is case-sensitive.
 - d. Click **Start a Service Tool > Hardware Service Manager > Work with Service Action Event Log**.
 - e. On the **Select Timeframe** display, change the **From Date and Time** field to the desired date and time range. Specify a value that includes the time period during which the procedure was performed.
 - f. Continue with step 2.
 - If your system is not managed by an HMC and has a partition running the Linux operating system, complete the following steps:
 - a. Log in as the root user. If you need assistance, contact the system administrator.
 - b. Type `servicelog -query='serviceable=1 AND closed=0 AND time_event>="yyyy-mm-dd"'` (where `yyyy` is the year, `mm` is the month, and `dd` is the day on which the procedure occurred). Press Enter.
 - c. Continue with step 2.
2. Perform problem analysis on open serviceable events. If you need assistance, contact your service provider.

Using identify indicators to locate hardware components

Many hardware components, such as PCI adapters, power supplies, and cable connectors, are equipped with small amber lights that are called identify indicators. The identify indicator, typically a light-emitting diode (LED) device, can be illuminated manually to assist you in making positive visual identification of a component. Use this procedure to activate and deactivate an identify indicator when you need to locate a hardware component.

To activate an identify indicator to locate a hardware component, complete the following steps:

1. Access the Advanced System Management Interface (ASMI) by using an authority level of administrator or authorized service provider. For details about using the ASMI, see *Managing the Advanced System Management Interface* (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hby/asmi.htm>).
2. In the ASMI navigation area, expand **System Configuration > Service Indicators**.
3. Click **Enclosure Indicators**.
4. From the list of enclosures, select the enclosure that contains the component to be identified. Click **Continue**.
5. Scan the list of locations for the enclosure that you selected, and identify the location that corresponds to the component to be identified.
6. Activate the identify indicator by changing the value from Off to Identify. Click **Save settings**. The setting is saved when the message *Operation completed successfully* is displayed.

7. Physically locate the component to visually confirm that its identify indicator is illuminated. If you identified the component to perform a specific task, complete that task now. When you are ready to deactivate the identify indicator, continue with the next step.
8. Repeat steps 2 on page 107 - 5 on page 107, choosing the same enclosure and location as you did previously.

Restriction: Do not use the **Back** button on your browser to get back to the page where you selected the component.

9. Deactivate the identity indicator by changing the value from **Identify** to **Off**. Click **Save settings**. The setting is saved when the message **Operation completed successfully** is displayed.

Configuring PCIe storage enclosures

Read information that has been provided for configuring PCIe storage enclosure models.

Configuring the 5888 PCIe storage enclosure

Learn how to configure the 5888 PCIe storage enclosure.

1. Review the following information:

The 5888 PCIe storage enclosure can hold up to 30 solid state drives (SSDs). The SSDs are physically organized in a single device partition within the enclosure. The devices cannot be split into separate physical groups. The SSDs are controlled by the dual PCIe2 SAS RAID controllers contained within the PCIe storage enclosure.

The 5888 PCIe storage enclosure is supported by the following operating systems:

- AIX
- Linux
- VIOS

Note: VIOS does not support the 5888 PCIe storage enclosure for an IBM i client.

To determine the level of software that you need to support the PCIe storage enclosure feature, see **IBM Prerequisite** (http://www-912.ibm.com/e_dir/eServerPrereq.nsf).

Tip: If you are planning to configure RAID arrays, ensure that you have the minimum number of available disks for each RAID level:

RAID 0
Minimum of one drive per array

RAID 5
Minimum of three drives per array

RAID 6
Minimum of four drives per array

RAID 10
Minimum of two drives per array

For additional SAS RAID controller information for the AIX operating system, see **SAS RAID controller for AIX** (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ebj/p7ebjkickoff.htm>).

For additional SAS RAID controller information for the Linux operating system, see **SAS RAID controller for Linux** (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ebk/p7ebkkickoff.htm>).

Configuring the EDR1 PCIe storage enclosure

Learn how to configure the EDR1 PCIe storage enclosure.

1. Review the following information:

The EDR1 PCIe storage enclosure can hold up to 30 solid-state drives (SSDs). The SSDs are physically organized in a single device partition within the enclosure. The devices cannot be split into separate physical groups.

As shown in the following figure, the EDR1 PCIe storage enclosure (A) supports the attachment of up to two optional 5887 disk drive enclosures (B) by using EX cables (C) for increased storage capacity.

Note: The AIX and Linux operating systems support the attachment of up to two 5887 disk drive enclosures. The IBM i operating system supports the attachment of one 5887 disk drive enclosure.




Figure 51. Connection of two 5887 disk drive enclosures by using EX cables to a PCIe storage enclosure

For detailed information about cabling a 5887 disk drive enclosure to the EDR1 PCIe storage enclosure, see Serial-attached SCSI cable planning (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7had/p7hadsascabling.htm>).

The 5887 disk drive enclosures must contain only hard disk drives (HDDs) and be factory-set to mode 1. The internal SSDs and any HDDs from attached 5887 disk drive enclosures are controlled by the dual PCIe2 SAS RAID controllers contained within the PCIe storage enclosure. SSDs and HDDs cannot be mixed in the same RAID array.

The EDR1 PCIe storage enclosure is supported by the following operating systems:

- AIX
- IBM i
- Linux
- VIOS

Note: VIOS does not support the EDR1 PCIe storage enclosure for an IBM i client.

To determine the level of software that you need to support the PCIe storage enclosure feature, see IBM Prerequisite (http://www-912.ibm.com/e_dir/eServerPrereq.nsf).

Tip: If you are planning to configure RAID arrays, ensure that you have the minimum number of available disks for each RAID level:

RAID 0

Minimum of one drive per array

RAID 5

Minimum of three drives per array

RAID 6

Minimum of four drives per array

RAID 10

Minimum of two drives per array

Best performance is achieved when the 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.

For additional SAS RAID controller information for the AIX operating system, see SAS RAID controller for AIX (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ebj/p7ebjkickoff.htm>).

For additional SAS RAID controller information for the IBM i operating system, see SAS RAID controllers for IBM i (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ed5/p7ed5kickoff.htm>).

For additional SAS RAID controller information for the Linux operating system, see SAS RAID controller for Linux (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7ebk/p7ebkkickoff.htm>).

Removing PCIe storage enclosures

Learn how to remove a PCIe storage enclosure from the system and remove it from the rack.

Depending on the feature of PCIe storage enclosure that you are working with, you might be able to remove it from the system with the system powered on and the logical partitions active.

- If you are removing a 5888 PCIe storage enclosure, the system must be powered off to complete the removal.
- If you are removing any other PCIe storage enclosure, you can remove it either with the system powered on and the logical partitions active or with the system powered off.

Related tasks:

“Removing an enclosure from the rack” on page 46

This procedure provides you with general steps to help you physically remove an enclosure that is mounted in a rack. Because methods for the attachment of enclosures to a rack vary by model, this procedure provides only conceptual removal steps.

Removing a PCIe storage enclosure with the system powered on

Learn about removing a PCIe storage enclosure from a system with the system powered on.

Important: To use this procedure to remove a PCIe storage enclosure from the system, the enclosure must be a model that can be removed with the system powered on. Review the following information to learn which options you have for removing your enclosure from the system:

- If you are removing a 5888 PCIe storage enclosure, you must use “Removing a PCIe storage enclosure with the system powered off” on page 116 instead of this procedure.
- If you are removing any other PCIe storage enclosure, you can remove it with the system powered on by using this procedure.

Prerequisites: Locate the hardware and hardware components that pertain to this procedure as follows:

- Identify the PCIe storage enclosure that is being removed.
- Identify the system or systems to which the PCIe storage enclosure is connected.

To remove a PCIe storage enclosure from the system with the system powered on, complete the following steps:

1. Record the feature code and serial number of the PCIe storage enclosure that you are removing. The feature code and serial number appear on the front panel of the PCIe storage enclosure.
2. Record the current date and time for use later in the procedure when you check for serviceable events.
3. Verify the existing PCIe hardware topology by completing the steps the applicable procedure:
 - If you have a 5888 PCIe storage enclosure, see “Verifying the PCIe hardware topology” on page 102.
 - If you have an EDR1 PCIe storage enclosure, see “Verifying the PCIe hardware topology” on page 102.

All PCIe hardware topology problems that are discovered in this step must be corrected before you continue with this procedure.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

4. Determine the identifier strings and link identifiers for the two I/O slots that are associated with the connections to the PCIe storage enclosure that is being removed by completing the following steps.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system. Each system hosts one of the I/O slots that connect the PCIe storage enclosure.

a. Choose one of the following options:

- If your system is managed by a Hardware Management Console (HMC), complete the following steps:

- 1) From the navigation bar, expand **Systems Management**.
- 2) Click **Servers**.
- 3) Select the server that you are working with.
- 4) In the Tasks area, expand **Hardware Information**.
- 5) Click **PCIe Hardware Topology**.

- If your system is not managed by an HMC, complete the following steps:

- 1) Access the ASMI by using an authority level of administrator or authorized service provider. For details about using the ASMI, see *Managing the Advanced System Management Interface* (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hby/asmi.htm>).
- 2) In the ASMI navigation area, expand **System Configuration**.
- 3) Click **PCIe Hardware Topology**.

- b. Review the PCIe Hardware Topology data to identify all PCIe links that are associated with the PCIe storage enclosure that you are removing. To make this identification, scan the I/O Enclosure Port column for values that contain the feature code and serial number that you recorded in step 1.

Note: The I/O Enclosure Port column value looks similar to Uxxxx.001.yyyyyyy-P1-Cz-T3, where xxxx is the feature code of the PCIe storage enclosure that is connected by the PCIe link, yyyyyyy is the serial number of the PCIe storage enclosure that is connected by the PCIe link, and z is either 1 or 2. An example of an actual I/O Enclosure Port column value is UEDR1.001.6BAH004-P1-C1-T3.

- c. For each PCIe link that you identified in step 4b on page 111, record the I/O slot identifier string from the I/O Slots column and the link identifier from the Link ID column.
5. For each I/O slot that you recorded in step 4c, choose one of the following options:
 - If your system is managed by an HMC and is not in the manufacturing default configuration, continue with step 6.
 - If your system is managed by an HMC and is in the manufacturing default configuration, or if your system is not managed by an HMC, choose one of the following options:
 - If the I/O slot is assigned to a logical partition that is powered on and running the Virtual I/O Server (VIOS) or AIX operating system, continue with step 7.
 - If the I/O slot is assigned to a logical partition that is powered on and running the IBM i operating system, continue with step 8 on page 113.
 - If the I/O slot is assigned to a logical partition that is powered on and running the Linux operating system, continue with step 9 on page 114.
6. Complete the following steps:
 - a. If the I/O slot is assigned to a logical partition that is powered on, complete the following steps to dynamically remove the I/O slot from the running logical partition:
 - 1) From the navigation bar, expand **Systems Management**.
 - 2) Click **Servers**.
 - 3) Select the server that you are working with.
 - 4) Scan the list of logical partitions that are defined on the selected server, and select the logical partition that owns the I/O slot.
 - 5) In the Tasks area, expand **Dynamic Logical Partitioning**.
 - 6) Expand **Physical Adapters**.
 - 7) Click **Move or Remove**.
 - 8) On the Remove I/O Adapters display, select the entry that contains the identifier string of the I/O slot in the Location Code column.
 - 9) In the **Move to partition** field, specify **None**.
 - 10) Click **OK**.
 - b. For each logical partition that includes the I/O slot in one or more partition profiles, complete the following steps to remove the I/O slot from partition profiles that include it:
 - 1) From the navigation bar, expand **Systems Management**.
 - 2) Click **Servers**.
 - 3) Select the server that you are working with.
 - 4) Scan the list of logical partitions that are defined on the selected server, and select the logical partition that includes the I/O slot in one or more of its profiles.
 - 5) In the Tasks area, expand **Configuration**.
 - 6) Click **Manage Profiles**.
 - 7) For each profile for the selected logical partition that includes the I/O slot, complete the following steps:
 - a) On the Manage Profiles display, click the name of the profile that includes the I/O slot.
 - b) On the Logical Partition Profile Properties display, click the **I/O** tab.
 - c) Select the entry that has the identifier string of the I/O slot in the Location Code column.
 - d) Click **Remove**.
 - e) Click **OK**.
 - c. If there is another I/O slot to process, return to step 5. Otherwise, continue with step 10 on page 114.
7. Complete the following steps:

- a. Log in to the operating system with admin or service level authority. If you need assistance, contact the system administrator.
 - b. Choose one of the following options:
 - If the partition is running the VIOS operating system, type `diagmenu` at the VIOS command prompt, and press Enter.
 - If the partition is running the AIX operating system, type `diag` at the AIX command prompt, and press Enter.
 - c. On the DIAGNOSTIC OPERATING INSTRUCTIONS display, press Enter to continue.
 - d. On the FUNCTION SELECTION display, use the arrow keys to select **Task Selection (Diagnostics, Advanced Diagnostics, Service Aids, etc.)**, and press Enter.
 - e. On the TASKS SELECTION LIST display, use the arrow keys to select **Hot Plug Task**, and press Enter.
 - f. On the Hot Plug Task display, use the arrow keys to select **PCI Hot Plug Manager**, and press Enter.
 - g. On the PCI Hot Plug Manager display, use the arrow keys to select **List PCI Hot Plug Slots**, and press Enter.
 - h. On the COMMAND STATUS display, locate the entry with the identifier string of the I/O slot in the Slot column, and record the logical device name in the Devices column.
 - i. Press **F3** (Exit).

Note: If your terminal emulation does not support **F3** to exit, press **ESC** then type 3 to exit.
 - j. On the PCI Hot Plug Manager display, use the arrow keys to select **Unconfigure a device**, and press Enter.
 - k. On the Unconfigure a device display, complete the following steps:
 - 1) Enter the logical device name that was recorded in step 7h for the **Device Name** field.
 - 2) Use the Tab key to select yes for the **Unconfigure any Child Devices** field.
 - 3) Use the Tab key to select no for the **KEEP definition in database** field.
 - 4) Press Enter.
 - l. On the ARE YOU SURE dialog, press Enter.
 - m. On the COMMAND STATUS display, wait for processing to complete and then press **F3** twice to return to the PCI Hot Plug Manager display.
 - n. On the PCI Hot Plug Manager display, use the arrow keys to select **Replace/Remove a PCI Hot Plug Adapter**, and press Enter.
 - o. On the Replace/Remove a PCI Hot Plug Adapter dialog, use the arrow keys to select the entry with the identifier string of the I/O slot in the Slot column, and press Enter.
 - p. On the Replace/Remove a PCI Hot Plug Adapter display, use the Tab key to select the **Remove** operation, and press Enter.
 - q. Follow the online instructions to complete the remove operation. However, do not complete the steps to physically remove the adapter.
 - r. If there is another I/O slot to process, return to step 5 on page 112. Otherwise, continue with step 10 on page 114.
8. Complete the following steps:
 - a. Log in to the operating system with admin or service level authority. If you need assistance, contact the system administrator.
 - b. At the IBM i command prompt, type `strsst` and press Enter.
 - c. Type the service tools user ID and service tools password on the System Service Tools (SST) Sign On display, and press Enter.
 - d. Select the **Start a service tool** option, and press Enter.
 - e. Select the **Hardware service manager** option, and press Enter.

- f. Select the **Packaging hardware resources (systems, frames, cards)** option, and press Enter.
- g. On the Packaging Hardware Resources display, locate the entry that contains the I/O slot. To make this identification, scan the display for an entry that has a Description column value of System Unit and Unit ID column value that matches the U label portion of the I/O slot identifier string. The U label portion is the substring that begins with U and continues until the first dash (-) character.
- h. Type 9 in the Opt column for the system unit entry that you located in step 8g, and press Enter.

Note: Specifying option 9 starts the Hardware that is contained within package view for the selected packaging hardware resource.

- i. On the Packaging Hardware Resources display, locate the entry for the I/O slot. To make this identification, scan the display for an entry with a Location column value that, when appended to the Unit ID value in the upper right corner, forms the complete I/O slot identifier string.
- j. Type 3 in the Opt column for the entry that you located in step 8i, and press Enter.

Note: Specifying option 3 starts the Hardware Resource Concurrent Maintenance menu for the selected packaging hardware resource.

- k. On the Hardware Resource Concurrent Maintenance display, start the Power off domain function by pressing F9.
- l. Follow the online instructions to complete the power-off operation.
- m. If there is another I/O slot to process, return to step 5 on page 112. Otherwise, continue with step 10.

9. Complete the following steps:

- a. Log in to the operating system with admin or service level authority. If you need assistance, contact the system administrator.
- b. At the Linux command prompt, type `drmgr -c pci -r -s io_slot_identifier` (where *io_slot_identifier* is the identifier string of the I/O slot), and press Enter.
- c. Follow the online instructions to complete the remove operation. However, do not complete the steps to physically remove the adapter.
- d. If there is another I/O slot to process, return to step 5 on page 112. Otherwise, continue with step 10.

10. If the system has a rear cover, remove or open it.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

11. Disconnect the PCIe cables that connect the PCIe storage enclosure by completing the following steps:

- a. For each PCIe link that you recorded in step 4 on page 111, complete the following steps:
 - 1) Choose one of the following options:
 - If your system is managed by an HMC, complete the following steps:
 - a) From the navigation bar, expand **Systems Management**.
 - b) Click **Servers**.
 - c) Select the server that you are working with.
 - d) In the Tasks area, expand **Hardware Information**.
 - e) Click **PCIe Hardware Topology**.
 - If your system is not managed by an HMC, complete the following steps:
 - a) Access the ASMI by using an authority level of administrator or authorized service provider. For details about using the ASMI, see Managing the Advanced System Management Interface (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hby/asmi.htm>).

- b) In the ASMI navigation area, expand **System Configuration**.
 - c) Click **PCIe Hardware Topology**.
- 2) Select the PCIe link and click **Identify Indicators**.
 - 3) On the Identify Indicators display, select both locations and click **Activate LED**.
- b. Locate the PCIe storage enclosure that you are removing. Disconnect the PCIe cables from the Enclosure RAID Modules (ERMs) in the PCIe storage enclosure. Use the identify indicators that are activated in step 11a on page 114 to assist in locating the PCIe storage enclosure and ERMs.
 - c. Locate the PCIe connectors on the systems that connect to the PCIe storage enclosure. Disconnect the PCIe cables from the systems. Use the identify indicators that are activated in step 11a on page 114 to assist in locating the PCIe connectors.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

- d. Deactivate the identify indicators by repeating step 11a on page 114, except to click **Deactivate LED** in step 11a3 rather than **Activate LED**.
12. Remove power from the PCIe storage enclosure by disconnecting the power cables from the power supplies on the PCIe storage enclosure.
 13. Remove the PCIe storage enclosure from the rack. For instructions, see “Removing an enclosure from the rack” on page 46.
 14. If the system has a rear cover, close it.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

15. Clear the PCIe links that are associated with the PCIe storage enclosure that you are removing by completing the following steps for each PCIe link that you recorded in step 4 on page 111:
 - a. Choose one of the following options:
 - If your system is managed by an HMC, complete the following steps:
 - 1) From the navigation bar, expand **Systems Management**.
 - 2) Click **Servers**.
 - 3) Select the server that you are working with.
 - 4) From the Tasks area, expand **Hardware Information**.
 - 5) Click **PCIe Hardware Topology**.
 - If your system is not managed by an HMC, complete the following steps:
 - 1) Access the ASMI by using an authority level of administrator or authorized service provider. For details about using the ASMI, see Managing the Advanced System Management Interface (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7hby/asm.htm>).
 - 2) In the ASMI navigation area, expand **System Configuration**.
 - 3) Click **PCIe Hardware Topology**.
 - b. Select the PCIe link.

Note: The link status for each link that is associated with the PCIe storage enclosure that you are removing should be Failed.

- c. Click **Clear Link**.

Note: The link status changes to Open.

16. Verify the PCIe hardware topology by completing the steps in the applicable procedure:
 - If you have a 5888 PCIe storage enclosure, see “Verifying the PCIe hardware topology” on page 102.

- If you have an EDR1 PCIe storage enclosure, see “Verifying the PCIe hardware topology” on page 102.

All PCIe hardware topology problems that are discovered in this step must be corrected before you continue with this procedure.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

17. Check for and handle new serviceable events that were generated during the procedure by completing the steps in “Checking for new serviceable events” on page 106. Use the date and time information that you recorded in step 2 on page 111 as the start time for searching for new serviceable events.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

Removing a PCIe storage enclosure with the system powered off

Learn about removing a PCIe storage enclosure from a system with the system powered off.

Note: Depending on the feature of the PCIe storage enclosure that you are working with, you might be able to remove it from the system with the system powered on and the logical partitions active. To learn which options you have for removing your PCIe storage enclosure from the system, review the following information:

- If you are removing a 5888 PCIe storage enclosure, you must use this procedure to remove the enclosure from the system with the system powered off.
- If you are removing any other PCIe storage enclosure, you can use this procedure to remove the enclosure from the system with the system powered off, or you can use “Removing a PCIe storage enclosure with the system powered on” on page 110 instead of this procedure to remove it with the system powered on.

Prerequisites: Locate the hardware and hardware components that pertain to this procedure as follows:

- Identify the PCIe storage enclosure that is being removed.
- Identify the system or systems to which the PCIe storage enclosure is connected.

To remove a PCIe storage enclosure from the system with the system powered off, complete the following steps:

1. Record the feature code and serial number of the PCIe storage enclosure that you are removing. The feature code and serial number appear on the front panel of the PCIe storage enclosure.
2. Optional: You can use identify indicators to assist in locating the PCIe cables and connectors that are associated with the PCIe storage enclosure you are removing. To illuminate the identify indicators for the two connections to the PCIe storage enclosure that is being removed, complete the following steps.

Important:

- To complete this step, your system must be running firmware at the 01Ax760 level, or later.
- The system must be powered on to do this step. After you complete this step for locating the PCIe connectors, the system is powered off.
- If the PCIe storage enclosure is connected to two systems, complete this step for each system.

Note: The identify indicators that are activated in this step provide support for steps 5 on page 117 - 6 on page 117.

- a. Choose one of the following options:

- If your system is managed by a Hardware Management Console (HMC), complete the following steps:
 - 1) From the navigation bar, expand **Systems Management**.
 - 2) Click **Servers**.
 - 3) Select the server that you are working with.
 - 4) From the Tasks area, expand **Hardware Information**.
 - 5) Click **PCIe Hardware Topology**.
 - If your system is not managed by an HMC, complete the following steps:
 - 1) Access the ASMI by using an authority level of administrator or authorized service provider. For details about using the ASMI, see Managing the Advanced System Management Interface (<http://www-01.ibm.com/support/knowledgecenter/api/redirect/systems/scope/hw/topic/p7hby/asmi.htm>).
 - 2) In the ASMI navigation area, expand **System Configuration**.
 - 3) Click **PCIe Hardware Topology**.
 - b. Review the PCIe Hardware Topology data to identify all PCIe links that are associated with the PCIe storage enclosure that you are removing. To make this identification, scan the I/O Enclosure Port column for values that contain the feature code and serial number that you recorded in step 1 on page 116.
- Note:** The I/O Enclosure Port column value looks similar to Uxxx.001.yyyyyyy-P1-Cz-T3, where xxx is the feature code of the PCIe storage enclosure that is connected by the PCIe link, yyyyyyy is the serial number of the PCIe storage enclosure that is connected by the PCIe link, and z is either 1 or 2. An example of an actual I/O Enclosure Port column value is UEDR1.001.6BAH004-P1-C1-T3.
- c. For each PCIe link that you identified in the previous step, complete the following steps:
 - 1) Select the link and click **Identify Indicators**.
 - 2) On the Identify Indicators display, select both locations and click **Activate LED**.
3. If the system is not already powered off, remove power from the system. For details, see Stopping a system or logical partition (<http://www-01.ibm.com/support/knowledgecenter/api/redirect/systems/scope/hw/topic/p7haj/crustopsys.htm>).

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

4. If the system has a rear cover, remove or open it.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

5. Disconnect the PCIe cables from the PCIe storage enclosure.

Note: If you chose to activate identify indicators in step 2 on page 116, they can be used now to assist in locating the PCIe connectors from which to disconnect the cables.

6. Disconnect the PCIe cables from the PCIe connectors on the system.

Note: If you chose to activate identify indicators in step 2 on page 116, they can be used now to assist in locating the PCIe connectors from which to disconnect the cables.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

7. Remove power from the PCIe storage enclosure by disconnecting the power cables from the power supplies on the PCIe storage enclosure.

8. Remove the PCIe storage enclosure from the rack. For instructions, see “Removing an enclosure from the rack” on page 46.
9. If the system has a rear cover, close it.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

10. Record the current date and time for use when you check for serviceable events.
11. Power on the system. For details, see Starting the system or logical partition (<http://publib.boulder.ibm.com/infocenter/systems/scope/hw/topic/p7haj/crustartsys.htm>).

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

12. Verify the PCIe hardware topology by completing the steps in the applicable procedure:
 - If you have a 5888 PCIe storage enclosure, see “Verifying the PCIe hardware topology” on page 102.
 - If you have an EDR1 PCIe storage enclosure, see “Verifying the PCIe hardware topology” on page 102.

All PCIe hardware topology problems that are discovered in this step must be corrected before you continue with this procedure.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

13. Check for and handle new serviceable events that were generated during the procedure by completing the steps in “Checking for new serviceable events” on page 106. Use the date and time information that you recorded in step 10 as the start time for searching for new serviceable events.

Important: If the PCIe storage enclosure is connected to two systems, complete this step for each system.

Reference information

Use the information in this section as needed to complete enclosure installation and configuration tasks.

Common procedures

These procedures are used throughout the topic.

Starting the system or logical partition


Learn how to start a system or logical partition after performing a service action or system upgrade.

Starting a system that is not managed by an HMC or an SDMC

You can use the power button or the Advanced System Management Interface (ASMI) to start a system that is not managed by a Hardware Management Console (HMC) or an IBM Systems Director Management Console (SDMC).

To start a system that is not managed by a HMC or SDMC, follow these steps:

1. Open the front rack door, if necessary.
2. Before you press the power button on the control panel, ensure that power is connected to the system unit as follows:
 - All system power cables are connected to a power source.
 - The power LED, as shown in the following figure, is slowly flashing.
 - The top of the display, as shown in the following figure, shows 01 V=F.
3. Press the power button (A), as shown in the following figure, on the control panel.



- A: Power-on button

- **B:** Power LED
 - A constant light indicates full system power to the unit.
 - A flashing light indicates standby power to the unit.

Note: There is approximately a 30-second transition period from the time the power-on button is pressed to when the power LED goes from flashing to solid. During the transition period, the LED might flash faster.

- **C:** Enclosure identify light
 - A constant light indicates the identify state for the enclosure or for a resource within the enclosure.
 - No light indicates that no resources in the enclosure are being identified.
- **D:** Attention light
 - No light indicates that the system is operating normally.
 - A solid light indicates that the system requires attention.
- **E:** USB port
- **F:** Enclosure fault roll-up light
 - A constant light indicates a fault indicator active in the system.
 - No light indicates that the system is operating normally.
- **G:** Function/Data display
- **H:** Decrement button
- **I:** Enter button
- **J:** Increment button
- **K:** Pinhole reset button

4. Observe the following aspects after pressing the power button:

- The power-on light begins to flash faster.
- The system cooling fans are activated after approximately 30 seconds and begin to accelerate to operating speed.
- Progress indicators, also referred to as checkpoints, appear on the control panel display while the system is being started. The power-on light on the control panel stops flashing and remains on, indicating that the system power is on.

Tip: If pressing the power button does not start the system, do the following steps to start the system by using the Advanced System Management Interface (ASMI):

1. Access the ASMI. For instructions, see [Accessing the ASMI without an HMC](#).
2. Start the system by using the ASMI. For instructions, see [Powering the system on and off](#).

Starting a system or logical partition by using the HMC

You can use the Hardware Management Console (HMC) to start the system or logical partition after the required cables are installed and the power cables are connected to a power source.

For instructions on working with the HMC, see [Managing the Hardware Management Console](#). For instructions on starting a logical partition, see [Logical partitioning](#). For instructions on starting the system, see [Powering on the managed system](#).

Progress indicators, also referred to as checkpoints, appear on the control panel display while the system is being started. When the power-on light on the control panel stops blinking and remains on, the system power is on.

Stopping a system or logical partition

Learn how to stop a system or logical partition as a part of a system upgrade or service action.

Attention: Using either the power-on button on the control panel or entering commands at the Hardware Management Console (HMC) to stop the system can cause unpredictable results in the data files. Also, the next time you start the system, it might take longer if all applications are not ended before stopping the system.

To stop the system or logical partition, select the appropriate procedure.

Stopping a system that is not managed by an HMC or an SDMC

You might need to stop the system to perform another task. If your system is not managed by the Hardware Management Console (HMC) or the IBM Systems Director Management Console (SDMC), use these instructions to stop the system by using the power button or the Advanced System Management Interface (ASMI).

Before you stop the system, follow these steps:

1. If an Integrated xSeries Adapter (IXA) is present on the system, shut it down using IBM i options.
2. Ensure that all jobs are completed and end all applications.
3. Ensure that the operating system is stopped.
Attention: Failure to do so can result in the loss of data.
4. If a Virtual I/O Server (VIOS) logical partition is running, ensure that all clients are shut down or that the clients have access to their devices by using an alternate method.

The following procedure describes how to stop a system that is not managed by the HMC or the SDMC.


1. Log in to the system as a user with the authority to run the **shutdown** or **pwrdownsys** (Power Down System) command.
2. At the command line, enter one of the following commands:
 - If your system is running the AIX operating system, type **shutdown**.
 - If your system is running the Linux operating system, type **shutdown -h now**.
 - If your system is running the IBM i operating system, type **PWRDOWNSYS**. If your system is partitioned, use the **PWRDOWNSYS** command to power down each of the secondary partitions. Then, use the **PWRDOWNSYS** command to power down the primary partition.

The command stops the operating system. The system power turns off, the power-on light begins to slowly flash, and the system goes into a standby state.

3. At the Linux command line, type **shutdown -h now**.
The command stops the operating system. The system power turns off, the power-on light begins to slowly flash, and the system goes into a standby state.
4. Record the IPL type and the IPL mode from the control panel display to help you return the system to this state when the installation or replacement procedure is completed.
5. Set the power switches of any devices connected to the system to off.
6. Unplug any power cables that are attached to the peripheral devices, such as printers and expansion units.

Important: The system may be equipped with a second power supply. Before continuing with this procedure, ensure that all power sources to the system have been disconnected.

(L003)



or



Stopping a system by using the HMC

You can use the Hardware Management Console (HMC) to stop the system or a logical partition.

By default, the managed system is set to power off automatically when you shut down the last running logical partition on the managed system. If you set the managed system properties on the HMC so that the managed system does not power off automatically, you must use this procedure to power off your managed system.

Attention: If possible, shut down the running logical partitions on the managed system before powering off the managed system. Powering off the managed system without shutting down the logical partitions first causes the logical partitions to shut down abnormally and can cause data loss. If you use a Virtual I/O Server (VIOS) logical partition, ensure that all clients are shut down or that the clients have access to their devices using an alternate method.

To power off a managed system, you must be a member of one of the following roles:

- Super administrator
- Service representative
- Operator
- Product engineer

Use the following steps to stop the system by using the HMC:

1. In the navigation area, expand the **Systems Management** folder.
2. Click the **Servers** icon.
3. In the Contents area, select the managed system.
4. Select **Tasks > Operations > Power Off**.
5. Select the appropriate power-off mode and click **OK**.

Related information:

Shutting down and restarting logical partitions

Powering on a system to firmware standby state

When a system is in the firmware standby state, the system firmware is active but all partitions remain powered off.

Complete the following steps to power on the system to firmware standby state:

1. Select one of the following options:
 - If your system is managed by a Hardware Management Console (HMC), continue with step 2.
 - If your system is not managed by an HMC, continue with step 9.
2. In the HMC navigation area, expand **Systems Management** and click **Servers**.
3. Select the check box next to the name of the desired server to enable the tasks for that server.
4. Complete the following steps to enable the server to power on to the firmware standby state:
 - a. From the **Tasks** menu, click **Properties**.
 - b. Click the **Power-On Parameters** tab.
 - c. Record the current setting for the **Partition start policy** field so that you can restore the setting later in this procedure.
 - d. Set the **Partition start policy** field to User-Initiated. Click **OK**.
5. From the **Tasks** menu, click **Operations > Power on**. Click **OK**.
6. In the content area, observe the server that you are working with until the value in the Status column changes to Standby.
7. To reset the **Partition start policy** field to its original value, complete the following steps:
 - a. From the **Tasks** menu, click **Properties**.
 - b. Click the **Power-On Parameters** tab.
 - c. Set the **Partition start policy** field to the value that you recorded in step 4c.
 - d. Click **OK**.
8. Return to the procedure that sent you here.
9. Access the Advanced System Management Interface (ASMI) by using an authority level of administrator or authorized service provider.
10. Complete the following steps to enable the server to power on to the firmware standby state:
 - a. In the ASMI navigation area, expand **Power/Restart Control**.
 - b. Click **Power On/Off System**.
 - c. Record the current setting for the **Server firmware start policy** field so that you can restore the setting later in this procedure.
 - d. Set the **Server firmware start policy** field to Standby (User-Initiated).
11. Click **Save settings and power on**.
12. Periodically, click **Power On/Off System** again to refresh the information on the display until the value of the **Current system server firmware state** field changes to Standby.
13. To reset the **Server firmware start policy** field to its original value, complete the following steps:
 - a. In the ASMI navigation area, expand **Power/Restart Control**.
 - b. Click **Power On/Off System**.
 - c. Set the **Server firmware start policy** field to the value that you recorded in step 10c.
14. Return to the procedure that sent you here.

Removing an enclosure from the rack

This procedure provides you with general steps to help you physically remove an enclosure that is mounted in a rack. Because methods for the attachment of enclosures to a rack vary by model, this procedure provides only conceptual removal steps.

To perform this task, you need the following items:

- A flat-head screwdriver
- A Philips-head screwdriver
- A Torx screwdriver
- Up to three people to physically lift and move the enclosure
- A clear space to place the enclosure and mounting hardware you are removing

Important: Supplement each step of this procedure with detailed information about the installation of the enclosure, when possible. Before you begin this procedure, check to see if the following information is available about the enclosure that you are removing:

- The inventory list of parts provided with the delivery of your enclosure.
- Installation documentation about your enclosure model provided online or with the delivery of your enclosure.

Tip: Online installation documentation for your enclosure might not be present in the Power Systems hardware level that you are currently searching. You might need to search earlier Power Systems hardware levels to locate installation documentation for the enclosure that you are removing.

To remove an enclosure from the rack, complete the following steps:

1. If the enclosure is not in a powered-off state, power off the enclosure.

Important: If the enclosure that you are removing is an expansion unit, the expansion unit must be powered off and removed from the system. Perform the removal of the expansion unit with the system powered on and running (concurrently) or with the system powered off (nonconcurrently).

- To learn about removing an expansion unit concurrently, see “Removing an I/O expansion unit with the system powered on” on page 41.
 - To learn about removing an expansion unit nonconcurrently, see “Removing an I/O expansion unit with the system powered off” on page 44.
2. If a service position exists for your enclosure, place the enclosure in its service position.
 3. If the enclosure was installed with shipping or bracing brackets, complete the following steps:
 - a. Remove any screws that attach the brackets to the rack.
 - b. Remove any screws that attach the brackets to the enclosure.
 - c. Carefully lift the brackets from the enclosure and remove them from the rack.

CAUTION:

Shipping and bracing brackets can be heavier than expected. Use more than one person to lift and move the brackets, if required.

4. Disconnect the enclosure from all devices and power as follows:
 - a. At the rear of the enclosure, disconnect power cords, any I/O bus cables connecting the enclosure to adapters, and other cables that might be present.
 - b. At the front of the enclosure, disconnect any UPIC cables and other cables that might be present.
5. Note the location of all screws that provide for the attachment of the enclosure, the enclosure mounting hardware, and the rack to each other. If enclosure bezels cover screws on the left and right ends of the front enclosure panel, remove those bezels.

6. Identify which screws secure the enclosure in position, which screws are weight-bearing screws, and which screws serve additional purposes. To remove the enclosure safely, you must understand the functions of all screws that are used to install the enclosure. Evaluate the screws in the following order:
 - a. Make a note of which screws are securing screws. You can identify a securing screw as any screw that fastens the enclosure chassis to any other fixture, such as the rack or a rail. Even if a screw attaches the chassis to another fixture that bears weight, it remains a securing screw and not a weight-bearing screw. If a screw secures only a mounting hardware item to the rack, it is not a securing screw. Check both the front and rear of the enclosure for securing screws. During the enclosure installation task, screws to secure the enclosure were the last screws to be installed. Therefore, they are the first screws that you are directed to remove later in this task to enable the detachment of the enclosure.

Note: The enclosure might be installed in a shell that is, in turn, installed to a rack or rails. If so, then any screw that fastens the enclosure to the shell is a securing screw. Also, any screw that fastens the shell to the rack or rails is a securing screw.

Exception: If the enclosure is installed on slide-out rails, the securing screws that attach the enclosure to the rails might also be weight-bearing screws. These models were installed in a two-step process that used rail-mounting guides. In the first step, the guides were temporarily attached to the sides of the enclosure. In the second step, the guides were removed after the enclosure was lifted onto the rails and securing screws were installed. For the purposes of this procedure, identify these screws as the weight-bearing screws that you are directed to remove later in this task.

- b. Make a note of which screws are enclosure weight-bearing screws. You can identify a weight-bearing screw as any screw that fastens a rail to the rack, unless the screw is already identified as a securing screw. The type of rail, such as a stationary rail or a slide-out rail, is not a factor in identifying weight-bearing screws. During the enclosure installation task, screws to attach weight-bearing mounting hardware to the rack were installed before the enclosure was lifted onto the rails and secured. Therefore, you will be directed to remove them later in this task only after the enclosure has been safely removed.
 - c. Make a note of any remaining screws used in the installation of the enclosure. These screws serve other purposes and are the last screws that you are directed to remove later in this task.
7. If the enclosure is installed on slide-out rails and you identified the securing screws as weight-bearing screws in step 6a on page 47, review the following options for the removal of the enclosure. Then complete the option that is best for your situation:
 - a. If you have the original rail-mounting guides and screws that were used during the installation of the enclosure, complete the following steps:
 - 1) Attach the rail-mounting guides to the sides of the enclosure using the screws provided for the original installation.
 - 2) Remove the screws that you identified in the Exception note of step 6a on page 47 as weight-bearing screws.
 - 3) Using three people, lift the enclosure from the rails and place it carefully in the space that you cleared. If the enclosure cannot be lifted because it is still attached to the rails, return to step 6 on page 47 to identify additional screws that might need to be removed.
 - 4) Continue with step 12 on page 48.
 - b. If you do not have the original rail-mounting guides and screws that were used during the installation of the enclosure, complete the following steps:
 - 1) Position three people at the front and sides of the enclosure to bear its weight and hold it steady during the removal of weight-bearing screws.
 - 2) Remove the screws that you identified in the Exception note of step 6a on page 47 as weight-bearing screws in the slide-out rails. Order the removal of screws so that the last two screws to be removed are ones installed at diagonal points on opposite rails.

- 3) Using the same three people holding the enclosure, lift the enclosure from the rails and place it carefully in the space that you cleared. If the enclosure cannot be lifted because it is still attached to the rails, return to step 6 on page 47 to identify additional screws that might need to be removed.
 - 4) Continue with step 12 on page 48.
8. If the enclosure is installed in a shell, complete the following steps:
 - a. If a cable retention bracket is present, remove the thumb screw that attaches the bracket to the shell, and remove the bracket.
 - b. Remove any securing screws that fasten the enclosure to the shell.
 - c. Slide the enclosure as required into a position where it can be removed. Using one or two people to handle the enclosure safely, slide the enclosure out from the shell.
 - d. If additional enclosures must be removed from the shell, repeat steps 8a on page 48 through 8c on page 48 as required to remove the enclosures.
 - e. Remove all remaining securing screws that fasten the shell to the rack or rails.
 - f. Lift the shell out from the rack and place it carefully in the space that you cleared.
 - g. Continue with step 12 on page 48.
 9. Remove all screws that you identified in step 6a on page 47 as securing screws. When this step is complete, the enclosure should no longer be secured to the rack or to any mounting hardware.

CAUTION:
Be careful in this step to remove only the screws that prevent the enclosure from being detached from the rack or mounting hardware. If you believe that a particular screw might be weight-bearing, do not remove it.
 10. Perform a test lift of the enclosure, using three people, to determine how many people are required to handle the enclosure safely.

CAUTION:
The heaviest of enclosures requires three people for safe lifting. Using fewer people than are required to lift an enclosure can result in injury.
 11. Slide the enclosure as required into a position where it can be removed. Lift the enclosure out from the rack and place it carefully in the space that you cleared. If the enclosure is still attached, return to step 6 on page 47 to identify additional screws that might need to be removed.

Note: Depending on the method of attachment, you might be required to open latches to fully detach the enclosure from the rack and mounting hardware.
 12. If you are required to remove the rails, you can remove them safely now. The procedure to remove rails includes some of the following common steps at one or both ends of the rail:
 - a. Fold down any hinge brackets located where the rail attaches to the rack.
 - b. Remove all screws that you identified in step 6b on page 47 as weight-bearing screws that attach the rail to the rack.
 - c. Push any spring-loaded seating pins in the ends of the rails back through their rack holes.
 - d. Open latches or loosen fasteners on the rail as required to detach the rail.
 - e. Perform any additional action required to retract and remove the rail.
 13. Remove any remaining screws that you identified in step 6c on page 47, and remove any remaining enclosure mounting hardware from the rack.
 14. Store all enclosure mounting hardware, brackets, bezels, and the shell for future use.

Connector locations

Learn about 12X, system power control network (SPCN), and PCIe connector locations.

Server connectors

Learn about connector locations for servers.

Model 8202-E4B connector locations

Learn about connector locations on rack-mounted and stand-alone models.




Figure 53. Model 8202-E4B connector locations on a rack-mounted model




Figure 54. Model 8202-E4B connector locations on a stand-alone model

Model 8202-E4C connector locations

Learn about connector locations on rack-mounted models.

The 8202-E4C server provides cable connector locations for the following enclosures:

- Support for expansion units, as shown in the following figure, is as follows:

- The card installed in the C1 slot can contain two 12X double data rate (DDR) cable ports. However, the C1 slot cannot provide both 12X DDR and PCIe cable ports.
- The server contains a serial-attached SCSI (SAS) port.
- Support for disk drive enclosures, as shown in the following figure, is as follows:
 - The server contains a SAS port.
- Support for PCIe storage enclosures, as shown in the following figure, is as follows:
 - The card installed in the C1 slot can contain two PCIe cable ports. However, the C1 slot cannot provide both PCIe and 12X DDR cable ports.
 - The server contains a SAS port.




Figure 55. Model 8202-E4C connections for expansion units, disk drive enclosures, and PCIe storage enclosures

Model 8202-E4D connector locations

Learn about connector locations on rack-mounted models.

The 8202-E4D server provides cable connector locations for the following enclosures:

- Support for expansion units, as shown in the following figure, is as follows:
 - The card installed in the C1 slot can contain two 12X double data rate (DDR) cable ports. However, the C1 slot cannot provide both 12X DDR and PCIe cable ports.
 - The server contains a serial-attached SCSI (SAS) port.
- Support for disk drive enclosures, as shown in the following figure, is as follows:
 - The server contains a SAS port.
- Support for PCIe storage enclosures, as shown in the following figure, is as follows:
 - The card installed in the C1 slot can contain two PCIe cable ports. However, the C1 slot cannot provide both PCIe and 12X DDR cable ports.
 - The server contains a SAS port.




Figure 56. Model 8202-E4D connections for expansion units, disk drive enclosures, and PCIe storage enclosures

Model 8205-E6B connector locations

Learn about connector locations on rack-mounted and stand-alone models.




Figure 57. Model 8205-E6B connector locations on a rack-mounted model




Figure 58. Model 8205-E6B connector locations on a stand-alone model

Model 8205-E6C connector locations

Learn about connector locations on rack-mounted models.

The 8205-E6C server provides cable connector locations for the following enclosures:

- Support for expansion units, as shown in the following figure, is as follows:
 - The card installed in the C1 slot can contain two 12X double data rate (DDR) cable ports. However, the C1 slot cannot provide both 12X DDR and PCIe cable ports.
 - The card installed in the C8 slot can contain two 12X DDR cable ports. However, the C8 slot cannot provide both 12X DDR and PCIe cable ports.
 - The port card type installed in the C1 slot can differ from the port card type installed in the C8 slot.
 - The server contains a serial-attached SCSI (SAS) port.
- Support for disk drive enclosures, as shown in the following figure, is as follows:
 - The server contains a SAS port.
- Support for PCIe storage enclosures, as shown in the following figure, is as follows:
 - The card installed in the C1 slot can contain two PCIe cable ports. However, the C1 slot cannot provide both PCIe and 12X DDR cable ports.
 - The card installed in the C8 slot can contain two PCIe cable ports. However, the C8 slot cannot provide both PCIe and 12X DDR cable ports.
 - The port card type installed in the C1 slot can differ from the port card type installed in the C8 slot.
 - The server contains a SAS port.




Figure 59. Model 8205-E6C connections for expansion units, disk drive enclosures, and PCIe storage enclosures

Model 8205-E6D connector locations

Learn about connector locations on rack-mounted models.

The 8205-E6D server provides cable connector locations for the following enclosures:

- Support for expansion units, as shown in the following figure, is as follows:
 - The card installed in the C1 slot can contain two 12X double data rate (DDR) cable ports. However, the C1 slot cannot provide both 12X DDR and PCIe cable ports.
 - The card installed in the C8 slot can contain two 12X DDR cable ports. However, the C8 slot cannot provide both 12X DDR and PCIe cable ports.
 - The port card type installed in the C1 slot can differ from the port card type installed in the C8 slot.
 - The server contains a serial-attached SCSI (SAS) port.
- Support for disk drive enclosures, as shown in the following figure, is as follows:
 - The server contains a SAS port.
- Support for PCIe storage enclosures, as shown in the following figure, is as follows:
 - The card installed in the C1 slot can contain two PCIe cable ports. However, the C1 slot cannot provide both PCIe and 12X DDR cable ports.
 - The card installed in the C8 slot can contain two PCIe cable ports. However, the C8 slot cannot provide both PCIe and 12X DDR cable ports.
 - The port card type installed in the C1 slot can differ from the port card type installed in the C8 slot.
 - The server contains a SAS port.




Figure 60. Model 8205-E6D connections for expansion units, disk drive enclosures, and PCIe storage enclosures

Model 8231-E1C connector locations

Learn about connector locations.

The 8231-E1C server provides cable connector locations for the following enclosures:

- Support for disk drive enclosures, as shown in the following figure, is as follows:
 - The server contains a serial-attached SCSI (SAS) port.
- Support for PCIe storage enclosures, as shown in the following figure, is as follows:
 - The card installed in the C1 slot can contain one PCIe cable port.
 - The server contains a SAS port.




Figure 61. Model 8231-E1C connections for disk drive enclosures and PCIe storage enclosures

Model 8231-E1D or 8268-E1D connector locations

Learn about connector locations.

The 8231-E1D or 8268-E1D servers provides cable connector locations for the following enclosures:

- Support for disk drive enclosures, as shown in the following figure, is as follows:
 - The server contains a serial-attached SCSI (SAS) port.
- Support for PCIe storage enclosures, as shown in the following figure, is as follows:
 - The card installed in the C1 slot can contain one PCIe cable port.
 - The server contains a SAS port.




Figure 62. Model 8231-E1D or 8268-E1D connections for disk drive enclosures and PCIe storage enclosures

Model 8231-E2B connector locations

Learn about connector locations on rack-mounted models.

The 8231-E2B server provides cable connector locations for the following enclosures:

- Support for disk drive enclosures, as shown in the following figure, is as follows:
 - The server contains a SAS port.




Figure 63. Model 8231-E2B connections for disk drive enclosures

- The 8231-E2B server provides no support for expansion units. The 12X DDR cable ports are for high speed system-to-system interconnection only.

Model 8231-E2C connector locations

Learn about connector locations.

The 8231-E2C server provides cable connector locations for the following enclosures:

- Support for expansion units, as shown in the following figure, is as follows:
 - The C1 slot must contain the system power control network (SPCN) card, and the C8 slot must contain the card with two 12X double data rate (DDR) cable ports.

Note: Installing the 12X DDR cable ports at this location causes both the C8 and C6 slot space to be occupied.

- The server cannot support both 12X DDR cable ports and PCIe cable ports at the same time.
- The server contains a serial-attached SCSI (SAS) port.
- Support for disk drive enclosures, as shown in the following figure, is as follows:
 - The server contains a SAS port.




Figure 64. Model 8231-E2C connections for expansion units and disk drive enclosures

- Support for PCIe storage enclosures, as shown in the following figure, is as follows:
 - The card installed in the C1 slot can contain one PCIe cable port, and the card installed in the C8 slot can contain one PCIe cable port.
 - The server cannot support both PCIe cable ports and 12X DDR cable ports at the same time.
 - The server contains a SAS port.




Figure 65. Model 8231-E2C connections for PCIe storage enclosures

Model 8231-E2D connector locations

Learn about connector locations.

The 8231-E2D server provides cable connector locations for the following enclosures:

- Support for expansion units, as shown in the following figure, is as follows:
 - The C1 slot must contain the system power control network (SPCN) card, and the C8 slot must contain the card with two 12X double data rate (DDR) cable ports.

Note: Installing the 12X DDR cable ports at this location causes both the C8 and C6 slot space to be occupied.

- The server cannot support both 12X DDR cable ports and PCIe cable ports at the same time.
- The server contains a serial-attached SCSI (SAS) port.
- Support for disk drive enclosures, as shown in the following figure, is as follows:
 - The server contains a SAS port.




Figure 66. Model 8231-E2D connections for expansion units and disk drive enclosures

- Support for PCIe storage enclosures, as shown in the following figure, is as follows:
 - The card installed in the C1 slot can contain one PCIe cable port, and the card installed in the C8 slot can contain one PCIe cable port.
 - The server cannot support both PCIe cable ports and 12X DDR cable ports at the same time.
 - The server contains a SAS port.




Figure 67. Model 8231-E2D connections for PCIe storage enclosures

Model 8233-E8B connector locations

Learn about connector locations on rack-mounted models.




Figure 68. Model 8233-E8B connector locations on a rack-mounted model

Model 8246-L1S connector locations

Learn about connector locations.




Figure 69. Model 8246-L1S connector locations

Model 8246-L1T connector locations

Learn about connector locations.




Figure 70. Model 8246-L1T connector locations

Model 8246-L2S connector locations

Learn about connector locations.




Figure 71. Model 8246-L2S connector locations

Model 8246-L2T connector locations

Learn about connector locations.




Figure 72. Model 8246-L2T connector locations

Model 8248-L4T, 8408-E8D, or 9109-RMD connector locations

Learn about connector locations.

The 8248-L4T, 8408-E8D, and 9109-RMD servers provide cable connector locations for the following enclosures:

- Support for expansion units, as shown in the following figure, is as follows:
 - The card installed in the C2 slot can contain two 12X double data rate (DDR) cable ports. However, the C2 slot cannot provide both 12X DDR and PCIe cable ports.
 - The card installed in the C3 slot can contain two 12X DDR cable ports. However, the C3 slot cannot provide both 12X DDR and PCIe cable ports.
- Support for disk drive enclosures, as shown in the following figure, is as follows:
 - The server contains a SAS port.
- Support for PCIe storage enclosures, as shown in the following figure, is as follows:
 - The card installed in the C2 slot can contain two PCIe cable ports. However, the C2 slot cannot provide both PCIe and 12X DDR cable ports.
 - The card installed in the C3 slot can contain two PCIe cable ports. However, the C3 slot cannot provide both PCIe and 12X DDR cable ports.

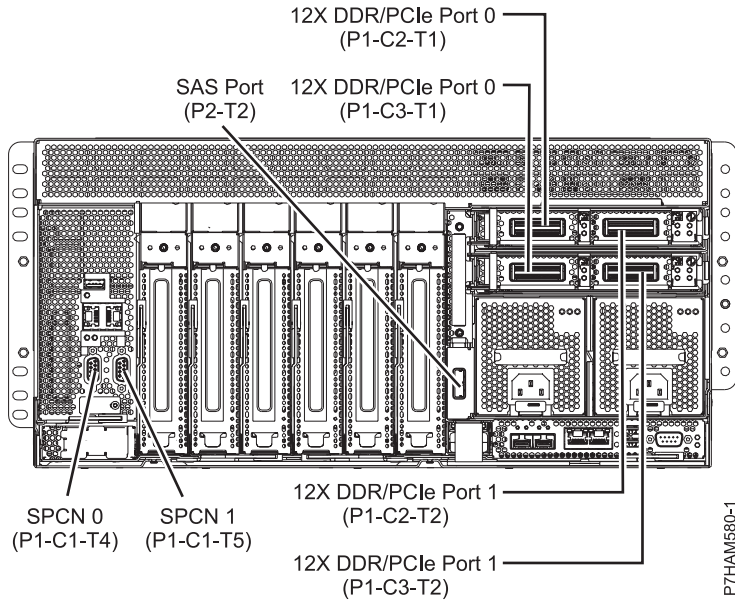


Figure 73. Model 8248-L4T, 8408-E8D, or 9109-RMD connector locations for expansion units, disk drive enclosures, and PCIe storage enclosures

Model 9117-MMB or 9179-MHB connector locations

Learn about connector locations.




Figure 74. Model 9117-MMB or 9179-MHB connector locations

Model 9117-MMC or 9179-MHC connector locations

Learn about connector locations.




Figure 75. Model 9117-MMC or 9179-MHC connector locations

Model 9117-MMD or 9179-MHD connector locations

Learn about connector locations.

The 9117-MMD and 9179-MHD servers provide cable connector locations for the following enclosures:

- Support for expansion units, as shown in the following figure, is as follows:
 - The card installed in the C2 slot can contain two 12X double data rate (DDR) cable ports. However, the C2 slot cannot provide both 12X DDR and PCIe cable ports.
 - The card installed in the C3 slot can contain two 12X DDR cable ports. However, the C3 slot cannot provide both 12X DDR and PCIe cable ports.
- Support for disk drive enclosures, as shown in the following figure, is as follows:
 - The server contains a SAS port.
- Support for PCIe storage enclosures, as shown in the following figure, is as follows:
 - The card installed in the C2 slot can contain two PCIe cable ports. However, the C2 slot cannot provide both PCIe and 12X DDR cable ports.
 - The card installed in the C3 slot can contain two PCIe cable ports. However, the C3 slot cannot provide both PCIe and 12X DDR cable ports.




Figure 76. Model 9117-MMD or 9179-MHD connector locations for expansion units, disk drive enclosures, and PCIe storage enclosures

Enclosure connectors

Learn about connector locations for I/O enclosures.

Model 5796 or 7314-G30 connector locations

Learn about connector locations.




Figure 77. Model 5796 or 7314-G30 connector locations

Model 5802 and 5877 connector locations

Learn about connector locations.




Figure 78. Model 5802 and 5877 connector locations

Model 5886 connector locations

Learn about connector locations.




Figure 79. Model 5886 connector locations

Model 5887 connector locations

Learn about connector locations.




Figure 80. Model 5887 connector locations

Model 5888 connector locations

Learn about connector locations.




Figure 81. Model 5888 connector locations

Model EDR1 connector locations

Learn about connector locations.




Figure 82. Model EDR1 connector locations

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 PUBLICATION 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 websites not owned by the manufacturer are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this product and use of those websites 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.

Homologation statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Electronic emission notices

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices supplied with the monitor.

Class A Notices

The following Class A statements apply to the IBM servers that contain the POWER7[®] processor and its features unless designated as electromagnetic compatibility (EMC) Class B in the feature information.

Federal Communications Commission (FCC) statement

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with

the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Industry Canada Compliance Statement

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 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 Deutschland GmbH
Technical Regulations, Department M372
IBM-Allee 1, 71139 Ehningen, Germany
Tele: +49 7032 15 2941
email: lugi@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 업무용(A급)으로 전자파적합기기로서 판매자 또는 사용자는 이 점을 주의하시기 바라며, 가정외의 지역에서 사용하는 것을 목적으로 합니다.

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 von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von 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 Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Regulations, Abteilung M372
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 7032 15 2941
email: lugi@de.ibm.com

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

Class B Notices

The following Class B statements apply to features designated as electromagnetic compatibility (EMC) Class B in the feature installation information.

Federal Communications Commission (FCC) statement

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation.

This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation.

If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult an IBM-authorized dealer or service representative for help.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM-authorized dealers. IBM is not responsible for

any radio or television interference caused by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate this equipment.

This device complies with Part 15 of the FCC rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Industry Canada Compliance Statement

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

Avis de conformité à la réglementation d'Industrie Canada

Cet appareil numérique de la classe B 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 B Information Technology Equipment according to European Standard EN 55022. The limits for Class B equipment were derived for typical residential environments to provide reasonable protection against interference with licensed communication equipment.

European Community contact:
IBM Deutschland GmbH
Technical Regulations, Department M372
IBM-Allee 1, 71139 Ehningen, Germany
Tele: +49 7032 15 2941
email: lugi@de.ibm.com

VCCI Statement - Japan

この装置は、クラスB情報技術装置です。この装置は、家庭環境で使用することを目的としていますが、この装置がラジオやテレビジョン受信機に近接して使用されると、受信障害を引き起こすことがあります。

取扱説明書に従って正しい取り扱いをして下さい。 VCCI-B

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

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 가정용(B급)으로 전자파적합기기로서 주로 가정에서 사용하는 것을 목적으로 하며, 모든 지역에서 사용할 수 있습니다.

Germany Compliance Statement

**Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse B EU-Richtlinie zur
Elektromagnetischen Verträglichkeit**

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2004/108/EG zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55022 Klasse B ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem "Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)". Dies ist die Umsetzung der EU-Richtlinie 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 B**

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

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:
International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:
IBM Deutschland GmbH
Technical Regulations, Abteilung M372
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 7032 15 2941
email: lugi@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 Klasse B.

Terms and conditions

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

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

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

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.



Printed in USA