IBM Storwize V7000

Troubleshooting, Recovery, and Maintenance Guide

IBM

> **Note**
>
> Before using this information and the product it supports, read the following information:
>
> - The general information in "Notices" on page 395
> - The information in the "Safety and environmental notices" on page xi
> - The information in the *IBM Environmental Notices and User Guide* (provided on a DVD)

Contents

Figures

Tables

ix

Safety and environmental notices

Review the safety notices, environmental notices, and electronic emission notices for IBM® Storwize® V7000 before you install and use the product.

Suitability for telecommunication environment: This product is not intended to connect directly or indirectly by any means whatsoever to interfaces of public telecommunications networks.

To find the translated text for a caution or danger notice, complete the following steps.

1. Look for the identification number at the end of each caution notice or each danger notice. In the following examples, the numbers (C001) and (D002) are the identification numbers.

   CAUTION:

   A caution notice indicates the presence of a hazard that has the potential of causing moderate or minor personal injury. (C001)

   DANGER

   A danger notice indicates the presence of a hazard that has the potential of causing death or serious personal injury. (D002)

2. Locate the *IBM Storwize V7000 Safety Notices* with the user publications that were provided with the Storwize V7000 hardware.

3. Find the matching identification number in the *IBM Storwize V7000 Safety Notices*. Then, review the topics about the safety notices to ensure that you are in compliance.

4. (Optional) Read the multilingual safety instructions on the Storwize V7000 website.

   a. Go to www.ibm.com/support

   b. Search for "Storwize V7000"

   c. Click the documentation link

Safety notices and labels

Review the safety notices and safety information labels before using this product.

To view a PDF file, you need Adobe Acrobat Reader. You can download it at no charge from the Adobe website:

www.adobe.com/support/downloads/main.html

IBM Systems Safety Notices

This publication contains the safety notices for the IBM Systems products in English and other languages. Anyone who plans, installs, operates, or services the system must be familiar with and understand the safety notices. Read the related safety notices before you begin work.

Note: The *IBM System Safety Notices* document is organized into two sections. The danger and caution notices without labels are organized alphabetically by language

in the "Danger and caution notices by language" section. The danger and caution notices that are accompanied with a label are organized by label reference number in the "Labels" section.

Note: You can find and download the current *IBM System Safety Notices* by searching for Publication number **G229-9054** in the IBM Publications Center.

The following notices and statements are used in IBM documents. They are listed in order of decreasing severity of potential hazards.

Danger notice definition
>   A special note that emphasizes a situation that is potentially lethal or extremely hazardous to people.

Caution notice definition
>   A special note that emphasizes a situation that is potentially hazardous to people because of some existing condition, or to a potentially dangerous situation that might develop because of some unsafe practice.

Note: In addition to these notices, labels might be attached to the product to warn of potential hazards.

Finding translated notices

Each safety notice contains an identification number. You can use this identification number to check the safety notice in each language.

To find the translated text for a caution or danger notice:

1. In the product documentation, look for the identification number at the end of each caution notice or each danger notice. In the following examples, the numbers (D002) and (C001) are the identification numbers.

   DANGER

   > **A danger notice indicates the presence of a hazard that has the potential of causing death or serious personal injury. (D002)**

   CAUTION:
   A caution notice indicates the presence of a hazard that has the potential of causing moderate or minor personal injury. (C001)

2. After you download the *IBM System Safety Notices* document, open it.
3. Under the language, find the matching identification number. Review the topics about the safety notices to ensure that you are in compliance.

Note: This product was designed, tested, and manufactured to comply with IEC 60950-1, and where required, to relevant national standards that are based on IEC 60950-1.

Caution notices for the Storwize V7000

Ensure that you understand the caution notices for Storwize V7000.

Use the reference numbers in parentheses at the end of each notice (for example, D005) to find the matching translated notice in *IBM Storwize V7000 Safety Notices*.

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. (C003)

CAUTION:

33.6-46.3 kg (74-102 lbs)	46.3-61.7 kg (102-136 lbs)	≥61.7-100 kg (136-220 lbs)

svc01053

The weight of this part or unit is more than 55 kg (121.2 lb). It takes specially trained persons, a lifting device, or both to safely lift this part or unit. (C011)

CAUTION:
To avoid personal injury, before lifting this unit, remove all appropriate subassemblies per instructions to reduce the system weight. (C012)

CAUTION:
Electrical current from power, telephone, and communication cables can be hazardous. To avoid personal injury or equipment damage, disconnect the attached power cords, telecommunication systems, networks, and modems before you open the machine covers, unless instructed otherwise in the installation and configuration procedures. (26)

CAUTION:

CAUTION regarding IBM provided VENDOR LIFT TOOL:

- Operation of LIFT TOOL by authorized personnel only

- LIFT TOOL intended for use to assist, lift, install, remove units (load) up into rack elevations. It is not to be used loaded transporting over major ramps nor as a replacement for such designated tools like pallet jacks, walkies, fork trucks and such related relocation practices. When this is not practicable, specially trained persons or services must be used (for instance, riggers or movers). Read and completely understand the contents of LIFT TOOL operator's manual before using.

- Read and completely understand the contents of LIFT TOOL operator's manual before using. Failure to read, understand, obey safety rules, and follow instructions may result in property damage and/or personal injury. If there are questions, contact the vendor's service and support. Local paper manual must remain with machine in provided storage sleeve area. Latest revision manual available on vendor's website.

- Test verify stabilizer brake function before each use. Do not over-force moving or rolling the LIFT TOOL with stabilizer brake engaged.

- Do not raise, lower or slide platform load shelf unless stabilizer (brake pedal jack) is fully engaged. Keep stabilizer brake engaged when not in use or motion.

- Do not move LIFT TOOL while platform is raised, except for minor positioning.

- Do not exceed rated load capacity. See LOAD CAPACITY CHART regarding maximum loads at center versus edge of extended platform.

- Only raise load if properly centered on platform. Do not place more than 200 lb (91 kg) on edge of sliding platform shelf also considering the load's center of mass/gravity (CoG).

- Do not corner load the platform tilt riser accessory option. Secure platform riser tilt option to main shelf in all four (4x) locations with provided hardware only, prior to use. Load objects are designed to slide on/off smooth platforms without appreciable force, so take care not to push or lean. Keep riser tilt option flat at all times except for final minor adjustment when needed.

- Do not stand under overhanging load.

- Do not use on uneven surface, incline or decline (major ramps).

- Do not stack loads. (C048, part 1 of 2)

- Do not operate while under the influence of drugs or alcohol.
- Do not support ladder against LIFT TOOL.
- Tipping hazard. Do not push or lean against load with raised platform.
- Do not use as a personnel lifting platform or step. No riders.
- Do not stand on any part of lift. Not a step.
- Do not climb on mast.
- Do not operate a damaged or malfunctioning LIFT TOOL machine.
- Crush and pinch point hazard below platform. Only lower load in areas clear of personnel and obstructions. Keep hands and feet clear during operation.
- No Forks. Never lift or move bare LIFT TOOL MACHINE with pallet truck, jack or fork lift.
- Mast extends higher than platform. Be aware of ceiling height, cable trays, sprinklers, lights, and other overhead objects.
- Do not leave LIFT TOOL machine unattended with an elevated load.
- Watch and keep hands, fingers, and clothing clear when equipment is in motion.
- Turn Winch with hand power only. If winch handle cannot be cranked easily with one hand, it is probably over-loaded. Do not continue to turn winch past top or bottom of platform travel. Excessive unwinding will detach handle and damage cable. Always hold handle when lowering, unwinding. Always assure self that winch is holding load before releasing winch handle.
- A winch accident could cause serious injury. Not for moving humans. Make certain clicking sound is heard as the equipment is being raised. Be sure winch is locked in position before releasing handle. Read instruction page before operating this winch. Never allow winch to unwind freely. Freewheeling will cause uneven cable wrapping around winch drum, damage cable, and may cause serious injury. (C048, part 2 of 2)

CAUTION:
- Do not install a unit in a rack where the internal rack ambient temperatures will exceed the manufacturer's recommended ambient temperature for all your rack-mounted devices.
- Do not install a unit in a rack where the air flow is compromised. Ensure that air flow is not blocked or reduced on any side, front, or back of a unit used for air flow through the unit.
- Consideration should be given to the connection of the equipment to the supply circuit so that overloading of the circuits does not compromise the supply wiring or overcurrent protection. To provide the correct power connection to a rack, refer to the rating labels located on the equipment in the rack to determine the total power requirement of the supply circuit.
- (For sliding drawers) Do not pull out or install any drawer or feature if the rack stabilizer brackets are not attached to the rack. Do not pull out more than one drawer at a time. The rack might become unstable if you pull out more than one drawer at a time.
- (For fixed drawers) This drawer is a fixed drawer and must not be moved for servicing unless specified by the manufacturer. Attempting to move the drawer partially or completely out of the rack might cause the rack to become unstable or cause the drawer to fall out of the rack. (R001 part 2 of 2)

CAUTION:
Removing components from the upper positions in the rack cabinet improves rack stability during a 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.**
- **If the rack cabinet you are relocating was supplied with removable outriggers they must be reinstalled before the cabinet is relocated.**
- **Inspect the route that you plan to take to eliminate potential hazards.**
- **Verify that the route that you choose can support the weight of the loaded rack cabinet. Refer to the documentation that comes with your rack cabinet for the weight of a loaded rack cabinet.**
- **Verify that all door openings are at least 760 x 230 mm (30 x 80 in.).**
- **Ensure that all devices, shelves, drawers, doors, and cables are secure.**
- **Ensure that the four leveling pads are raised to their highest position.**
- **Ensure that there is no stabilizer bracket installed on the rack cabinet during movement.**
- **Do not use a ramp inclined at more than 10 degrees.**
- **When the rack cabinet is in the new location, complete the following steps:**
  - **Lower the four leveling pads.**
  - **Install stabilizer brackets on the rack cabinet.**
  - **If you removed any devices from the rack cabinet, repopulate the rack cabinet from the lowest position to the highest position.**
- **If a long-distance relocation is required, restore the rack cabinet to the configuration of the rack cabinet as you received it. Pack the rack cabinet in the original packaging material, or equivalent. Also lower the leveling pads to raise the casters off the pallet and bolt the rack cabinet to the pallet. (R002)**

Danger notices for Storwize V7000

Ensure that you are familiar with the danger notices for Storwize V7000.

Use the reference numbers in parentheses at the end of each notice (for example, D005) to find the matching translated notice in *IBM Storwize V7000 Safety Notices*.

DANGER

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

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

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- 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.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

DANGER

Heavy equipment–personal injury or equipment damage might result if mishandled. (D006)

DANGER

> **DANGER: Serious injury or death can occur if loaded lift tool falls over or if a heavy load falls off the lift tool. Always completely lower the lift tool load plate and properly secure the load on the lift tool before moving or using the lift tool to lift or move an object. (D010)**

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.**
>
> f2c00064
>
> - **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. (R001 part 1 of 2)**

DANGER

> **Racks with a total weight of > 227 kg (500 lb.), Use Only Professional Movers! (R003)**

DANGER

> Do not transport the rack via fork truck unless it is properly packaged, secured on top of the supplied pallet. (R004)

DANGER:

⏚

Main Protective Earth (Ground):

This symbol is marked on the frame of the rack.
The PROTECTIVE EARTHING CONDUCTORS should be terminated at that point. A recognized or certified closed loop connector (ring terminal) should be used and secured to the frame with a lock washer using a bolt or stud. The connector should be properly sized to be suitable for the bolt or stud, the locking washer, the rating for the conducting wire used, and the considered rating of the breaker. The intent is to ensure the frame is electrically bonded to the PROTECTIVE EARTHING CONDUCTORS. The hole that the bolt or stud goes into where the terminal conductor and the lock washer contact should be free of any non-conductive material to allow for metal to metal contact. All PROTECTIVE EARTHING CONDUCTORS should terminate at this main protective earthing terminal or at points marked with ⏚ . (R010)

Special caution and safety notices

This information describes special safety notices that apply to the Storwize V7000. These notices are in addition to the standard safety notices supplied and address specific issues relevant to the equipment provided.

General safety

When you service the Storwize V7000, follow general safety guidelines.

Use the following general rules to ensure safety to yourself and others.

- Observe good housekeeping in the area where the devices are kept during and after maintenance.
- Follow the guidelines when lifting any heavy object:
    1. Ensure that you can stand safely without slipping.
    2. Distribute the weight of the object equally between your feet.
    3. Use a slow lifting force. Never move suddenly or twist when you attempt to lift.
    4. Lift by standing or by pushing up with your leg muscles; this action removes the strain from the muscles in your back. *Do not attempt to lift any objects that weigh more than 18 kg (40 lb) or objects that you think are too heavy for you.*
- Do not perform any action that causes a hazard or makes the equipment unsafe.
- Before you start the device, ensure that other personnel are not in a hazardous position.
- Place removed covers and other parts in a safe place, away from all personnel, while you are servicing the unit.
- Keep your tool case away from walk areas so that other people cannot trip over it.
- Do not wear loose clothing that can be trapped in the moving parts of a device. Ensure that your sleeves are fastened or rolled up above your elbows. If your hair is long, fasten it.

- Insert the ends of your necktie or scarf inside clothing or fasten it with a nonconducting clip, approximately 8 cm (3 in.) from the end.
- Do not wear jewelry, chains, metal-frame eyeglasses, or metal fasteners for your clothing.

  Remember: Metal objects are good electrical conductors.
- Wear safety glasses when you are hammering, drilling, soldering, cutting wire, attaching springs, using solvents, or working in any other conditions that might be hazardous to your eyes.
- After service, reinstall all safety shields, guards, labels, and ground wires. Replace any safety device that is worn or defective.
- Reinstall all covers correctly after you have finished servicing the unit.

Handling static-sensitive devices

Ensure that you understand how to handle devices that are sensitive to static electricity.

Attention: Static electricity can damage electronic devices and your system. To avoid damage, keep static-sensitive devices in their static-protective bags until you are ready to install them.

To reduce the possibility of electrostatic discharge, observe the following precautions:
- Limit your movement. Movement can cause static electricity to build up around you.
- Handle the device carefully, holding it by its edges or frame.
- Do not touch solder joints, pins, or exposed printed circuitry.
- Do not leave the device where others can handle and possibly damage the device.
- While the device is still in its antistatic bag, touch it to an unpainted metal part of the system unit for at least two seconds. (This action removes static electricity from the package and from your body.)
- Remove the device from its package and install it directly into your Storwize V7000, without putting it down. If it is necessary to put the device down, place it onto its static-protective bag. (If your device is an adapter, place it component-side up.) Do not place the device onto the cover of the Storwize V7000 or onto a metal table.
- Take additional care when you handle devices during cold weather. Indoor humidity tends to decrease in cold weather, causing an increase in static electricity.

Sound pressure

Attention: Depending on local conditions, the sound pressure can exceed 85 dB(A) during service operations. In such cases, wear appropriate hearing protection.

Environmental notices

The *IBM Systems Environmental Notices* contains all of the required environmental notices for IBM Systems products in English and other languages.

The *IBM Systems Environmental Notices* (`http://ibm.co/1fBgWFI`) includes statements on limitations, product information, product recycling and disposal, battery information, flat panel display, refrigeration and water-cooling systems, external power supplies, and safety data sheets.

About this guide

This guide describes how to service, maintain, and troubleshoot the IBM Storwize V7000.

The chapters that follow introduce you to the hardware components and to the tools that assist you in troubleshooting and servicing the Storwize V7000, such as the management GUI and the service assistant.

The troubleshooting procedures can help you analyze failures that occur in a Storwize V7000 system. With these procedures, you can isolate the components that fail.

You are also provided with step-by-step procedures to remove and replace parts.

Who should use this guide

This guide is intended for system administrators who use and diagnose problems with the Storwize V7000.

Storwize V7000 library and related publications

Product manuals, other publications, and websites contain information that relates to Storwize V7000.

IBM Knowledge Center for Storwize V7000

The information collection in the IBM Knowledge Center contains all of the information that is required to install, configure, and manage the system. The information collection in the IBM Knowledge Center is updated between product releases to provide the most current documentation. The information collection is available at the following website:

http://www.ibm.com/support/knowledgecenter/STPVGU

Storwize V7000 library

Unless otherwise noted, the publications in the library are available in Adobe portable document format (PDF) from a website.

ibm.com/shop/publications/order

Click **Search for publications** to find the online publications you are interested in, and then view or download the publication by clicking the appropriate item.

Table 1 lists websites where you can find help, services, and more information.

Table 1. IBM websites for help, services, and information

Website	Address
Directory of worldwide contacts	http://www.ibm.com/planetwide
Support for Storwize V7000 (2076)	www.ibm.com/support

Table 1. IBM websites for help, services, and information (continued)

Website	Address
Support for IBM System Storage® and IBM TotalStorage products	www.ibm.com/support/

Each PDF publication in the Table 2 library is also available in the IBM Knowledge Center by clicking the number in the "Order number" column:

Table 2. Storwize V7000 library

Title	Description	Order number
IBM Storwize V7000 Gen2 and Gen2+ Quick Installation Guide	The guide provides detailed instructions for unpacking your order and installing your system. The first chapter describes verifying your order, becoming familiar with the hardware components, and meeting environmental requirements. The second chapter describes installing the hardware and attaching data cables and power cords. The last chapter describes accessing the management GUI to initially configure your system.	GC27-6500
IBM Storwize V7000 Troubleshooting, Recovery, and Maintenance Guide	The guide describes how to service, maintain, and troubleshoot the Storwize V7000 system.	GC27-2291
IBM Spectrum Virtualize for SAN Volume Controller and Storwize Family Command-Line Interface User's Guide	The guide describes the commands that you can use from the Storwize V7000 command-line interface (CLI).	GC27-2287

IBM documentation and related websites

Table 3 lists websites that provide publications and other information about the Storwize V7000 or related products or technologies. The IBM Redbooks® publications provide positioning and value guidance, installation and implementation experiences, solution scenarios, and step-by-step procedures for various products.

Table 3. IBM documentation and related websites

Website	Address
IBM Publications Center	ibm.com/shop/publications/order
IBM Redbooks publications	www.redbooks.ibm.com/

Related accessibility information

To view a PDF file, you need Adobe Reader, which can be downloaded from the Adobe website:

www.adobe.com/support/downloads/main.html

IBM Publications Center

The IBM Publications Center is a worldwide central repository for IBM product publications and marketing material.

The IBM Publications Center website offers customized search functions to help you find the publications that you need. You can view or download publications at no charge. Access the IBM Publications Center through the following website:

ibm.com/shop/publications/order

Related websites

The following websites provide information about Storwize V7000 or related products or technologies:

Type of information	Website
Storwize V7000 support	www.ibm.com/support
Technical support for IBM storage products	www.ibm.com/support/
IBM Electronic Support registration	www-01.ibm.com/support/electronicsupport/

Sending comments

Your feedback is important in helping to provide the most accurate and highest quality information.

Procedure

To submit any comments about this publication or any other IBM storage product documentation:

Send your comments by email to starpubs@us.ibm.com. Be sure to include the following information:
- Exact publication title and version
- Publication form number (for example, GA32-1234-00)
- Page, table, or illustration numbers that you are commenting on
- A detailed description of any information that should be changed

How to get information, help, and technical assistance

If you need help, service, technical assistance, or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you.

Information

IBM maintains pages on the web where you can get information about IBM products and fee services, product implementation and usage assistance, break and fix service support, and the latest technical information. For more information, refer to Table 4 on page xxvi.

Table 4. IBM websites for help, services, and information

Website	Address
Directory of worldwide contacts	http://www.ibm.com/planetwide
Support for Storwize V7000 (2076)	www.ibm.com/support
Support for IBM System Storage and IBM TotalStorage products	www.ibm.com/support/

Note: Available services, telephone numbers, and web links are subject to change without notice.

Help and service

Before calling for support, be sure to have your IBM Customer Number available. If you are in the US or Canada, you can call 1 (800) IBM SERV for help and service. From other parts of the world, see http://www.ibm.com/planetwide for the number that you can call.

When calling from the US or Canada, choose the **storage** option. The agent decides where to route your call, to either storage software or storage hardware, depending on the nature of your problem.

If you call from somewhere other than the US or Canada, you must choose the **software** or**hardware** option when calling for assistance. Choose the **software** option if you are uncertain if the problem involves the Storwize V7000 software or hardware. Choose the **hardware** option only if you are certain the problem solely involves the Storwize V7000 hardware.When calling IBM for service regarding the product, follow these guidelines for the **software** and **hardware** options:

Software option
> Identify the Storwize V7000 product as your product and supply your customer number as proof of purchase. The customer number is a 7-digit number (0000000 - 9999999) assigned by IBM when the product is purchased. Your customer number should be on the customer information worksheet or on the invoice from your storage purchase. If asked for an operating system, use **Storage**.

Hardware option
> Provide the serial number and appropriate 4-digit machine type. For Storwize V7000, the machine type is 2076.

In the US and Canada, hardware service and support can be extended to 24x7 on the same day. The base warranty is 9x5 on the next business day.

Getting help online

You can find information about products, solutions, partners, and support on the IBM website.

To find up-to-date information about products, services, and partners, visit the IBM website at www.ibm.com/support.

Before you call

Make sure that you have taken steps to try to solve the problem yourself before you call.

Some suggestions for resolving the problem before calling IBM Support include:

- Check all cables to make sure that they are connected.
- Check all power switches to make sure that the system and optional devices are turned on.
- Use the troubleshooting information in your system documentation. The troubleshooting section of the knowledge center contains procedures to help you diagnose problems.
- Go to the IBM Support website at www.ibm.com/support to check for technical information, hints, tips, and new device drivers or to submit a request for information.

Using the documentation

Information about your IBM storage system is available in the documentation that comes with the product.

That documentation includes printed documents, online documents, readme files, and help files in addition to the knowledge center. See the troubleshooting information for diagnostic instructions. The troubleshooting procedure might require you to download updated device drivers or software. IBM maintains pages on the web where you can get the latest technical information and download device drivers and updates. To access these pages, go to www.ibm.com/support and follow the instructions. Also, some documents are available through the IBM Publications Center.

Sign up for the Support Line Offering

If you have questions about how to use and configure the machine, sign up for the IBM Support Line offering to get a professional answer.

The maintenance that is supplied with the system provides support when there is a problem with a hardware component or a fault in the system machine code. At times, you might need expert advice about using a function that is provided by the system or about how to configure the system. Purchasing the IBM Support Line offering gives you access to this professional advice while deploying your system, and in the future.

Contact your local IBM sales representative or your support group for availability and purchase information.

What's new

New and updated information was included in this version of the book as a result of usability testing and other feedback. Read all of the steps no matter how familiar you are with the installation.

Chapter 1. Storwize V7000 hardware components

A Storwize V7000 system consists of one or more machine type 2076 rack-mounted enclosures. Control enclosures contain the node canisters that manage the system operation and provide the host interfaces. Expansion enclosures provide more extra drives that can be managed by the system. Enclosures can support 2.5 inch (6.35 cm) small form factor drives or 3.5 inch (8.89 cm) large form factor drives.

There are several model types. The main differences among the model types are the following items:

- Whether the model is Gen1 or the newer Gen2

  Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 5. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 6. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

- The number of drives that an enclosure can hold

  Drives are on the front of the enclosure. An enclosure can hold up to twelve 3.5-inch drives or up to twenty-four 2.5-inch drives.

- Whether the model is a control enclosure or an expansion enclosure

  Control enclosures contain the main processing units that control the whole system. They are where external systems, such as host application servers, other storage systems, and management workstations are connected through the

Ethernet ports or Fibre Channel ports. Control enclosures can also be connected to expansion enclosures through the serial-attached SCSI (SAS) ports.

Expansion enclosures contain more storage capacity. Expansion enclosures connect either to control enclosures or to other expansion enclosures through the SAS ports.

- If the control enclosure has either 1 Gbps Ethernet capability or 10 Gbps Ethernet capability

The machine type and model (MTM) are shown on labels on the front and the rear of each enclosure:

- The left end cap label on the front of the enclosure indicates whether the enclosure is a control enclosure or an expansion enclosure.
- The label on the rear of the left enclosure flange.

Note: Labels also show the enclosure serial number. You must know the serial number when you contact IBM support.

You must be able to distinguish between control enclosures and expansion enclosures before you service the system. Be aware of the following differences:

- The model type that is shown on the labels.
- The model description that is shown on the left end cap.
- The number of ports at the rear of the enclosure. Control enclosures have Ethernet ports, Fibre Channel ports, and USB ports. Expansion enclosures do not have any of these ports.

Components in the front of the enclosure

The front of each control enclosure features several different components.

Drives for control enclosures

Storwize V7000 enclosures use different drives, depending on the generation of your control enclosure model. An enclosure can hold up to 12 3.5 in. (8.89 cm) drives or up to 24 2.5 in. (6.35 cm) drives.

Storwize V7000 Gen1

The drives are in the front of the enclosure. The 12 large form factor (LFF) 3.5 in. (8.89 cm) drives are mounted horizontally in four columns with three rows. The 24 small form factor (SFF) 2.5 in. (6.35 cm) drives are mounted vertically in one row.

Important: Drive slots cannot be empty. Install a drive assembly or blank carrier in each slot.

Note: Drives that are sold as Storwize V7000 options are the only drives that are supported. For more information, see the Support website.

Figure 1 on page 3 shows 12 LFF drives, as viewed from the front of the enclosure.

Figure 1. Storwize V7000 Gen1 Drives on a 12-drive enclosure

Figure 2 shows a large form factor 3.5 in. (8.89 cm) drive with the latching mechanism open.
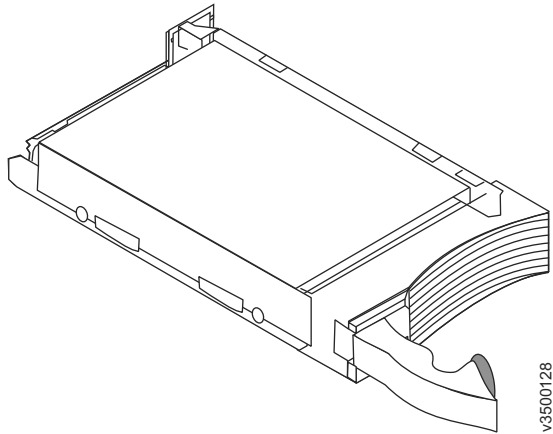
Figure 2. Storwize V7000 Gen2 Large form factor horizontal drive

Figure 3 shows 24 SFF drives, as viewed from the front of the enclosure.

Figure 3. Storwize V7000 Gen1 Drives on a 24-drive enclosure

Storwize V7000 Gen2

Figure 4 shows a small form factor 2.5 in. (6.35 cm) drive with the latching mechanism open.
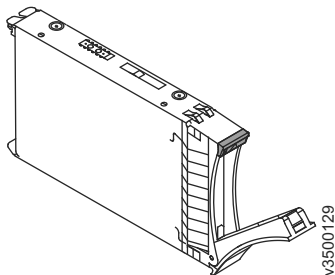
Figure 4. Storwize V7000 Gen2 Small form factor vertical drive

Drive indicators for control enclosures

Storwize V7000 enclosures use different drive indicators, depending on the generation of your control enclosure model. Drives have two light-emitting diode (LED) indicators each; they have no controls or connectors.

Storwize V7000 Gen1

The LED color is the same for both drives. The LEDs for the 3.5-inch drives are placed vertically above and below each other. Figure 5 displays the LED indicators on a single 3.5-inch drive. The LEDs for the 2.5-inch drives are placed next to each other at the bottom. Figure 6 displays the LED indicators on a single 2.5-inch drive.
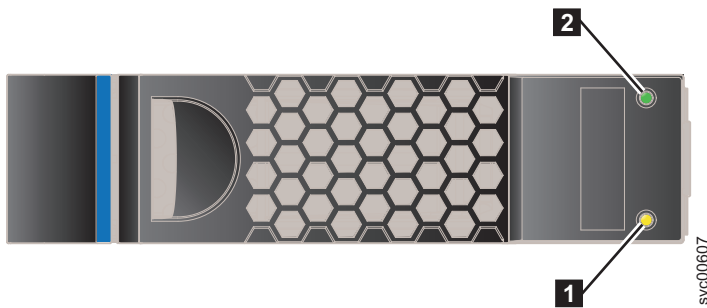
Figure 5. LED indicators on a single 3.5 inch drive

Figure 6. LED indicators on a single 2.5 inch drive

1 Fault LED

2 Activity LED

Table 7 on page 5 shows the status descriptions for the two LEDs.

Table 7. Drive LEDs

Name	Description	Color
Activity	Indicates whether the drive is ready or active. • If the LED is on, the drive is ready to be used. • If the LED is off, the drive is not ready. • If the LED is flashing, the drive is ready, and there is activity.	Green
Fault	Indicates a fault or identifies a drive. • If the LED is on, a fault exists on the drive. • If the LED is off, no known fault exists on the drive. • If the LED is flashing, the drive is being identified. A fault might or might not exist.	Amber

Storwize V7000 Gen2

The drive indicator LEDs on all systems use the same color and flashing patterns for both sizes of drives. Figure 7 shows the location of the LEDs on a 3.5-inch drive.

Figure 7. Storwize V7000 LED indicators on a horizontal 3.5 in. (8.89 cm) drive

Figure 8 shows the location of the LEDs on a 2.5-inch drive.
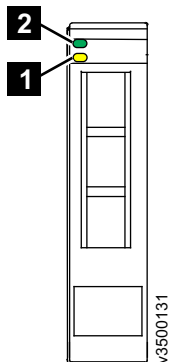
Figure 8. Storwize V7000 LED indicators on a vertical 2.5 in. (6.35 cm) drive

Table 8 on page 6 lists the status descriptions for the two LEDs.

Table 8. Storwize V7000 drive LED status

LED name	Callout	Color	State	Meaning
Fault	**1**	Amber	OFF	No known fault exists.
			FLASHING	The drive is being identified; a fault might or might not exist.
			ON	A fault exists on the drive.
Activity	**2**	Green	OFF	The drive is not ready for use.
			FLASHING	The drive is ready. Activity is in progress.
			ON	The drive is ready. No activity is in progress.

Enclosure end cap indicators

Storwize V7000 enclosure end cap indicators vary, depending on the generation of your control enclosure model. The enclosure indicators are on the left end cap of the enclosure and provide an overview of the enclosure status.

Storwize V7000 Gen1

Figure 9 shows where the end caps are on the front of an enclosure with 12 drives. The end caps are in the same position for an enclosure with 24 drives.

1 Left end cap
2 Drives
3 Right end cap

The left enclosure end caps for both enclosures are identical and contain only indicators. The left enclosure end cap contains no controls or connectors. The right enclosure end cap for both enclosures has no controls, indicators, or connectors.
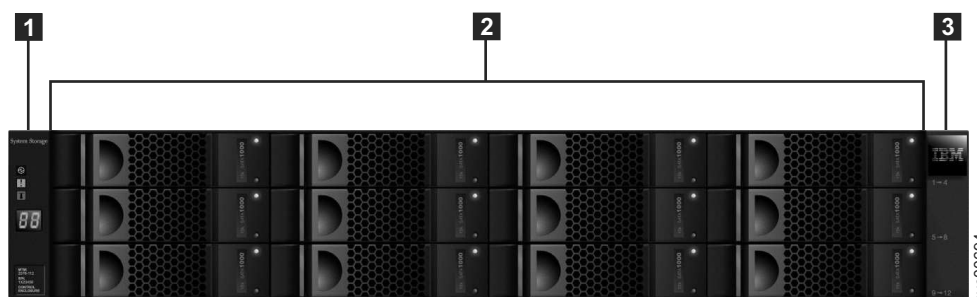
Figure 9. Storwize V7000 Gen1 12 drives and two end caps

Figure 10 on page 7 shows the indicators on the front of the enclosure end cap.
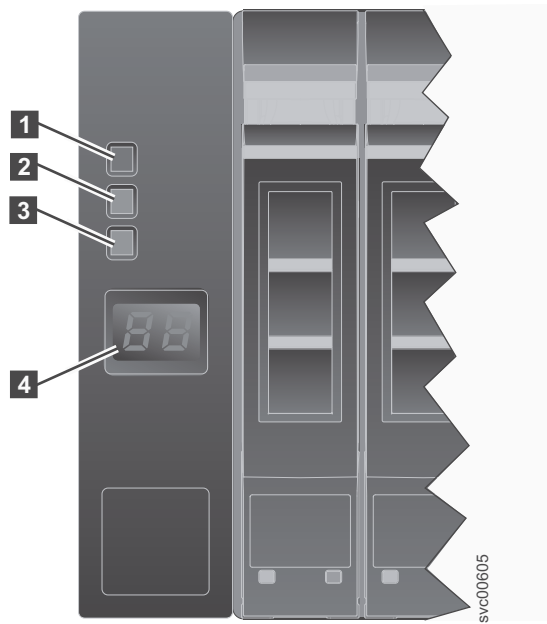
Figure 10. Storwize V7000 Gen1 Left enclosure end cap

Table 9. Storwize V7000 Gen1 LED descriptions for left enclosure end cap

Name	Description	Color	Symbol
Power	**1** The power LED is the upper LED. When the green LED is lit, it indicates that the main power is available to the enclosure	Green	⌾
Fault	**2** The fault LED is the middle LED. When the amber LED is lit, it indicates one of the enclosure components has a hardware fault.	Amber	!
Identify	**3** The identify LED is the lower LED. When the blue LED is lit, it identifies the enclosure.	Blue	((¡))
N/A	**4** The two-character LCD display shows the enclosure ID.	N/A	N/A

Storwize V7000 Gen2

Labels and indicators on the enclosure left end cap provide information about the enclosure and the enclosure status. The product name, machine type-model, serial number, and type of enclosure (control or expansion) are shown on the Storwize V7000 left end cap. When multiple enclosures are in a machine room or rack, the serial number is used to identify the enclosure being referenced.
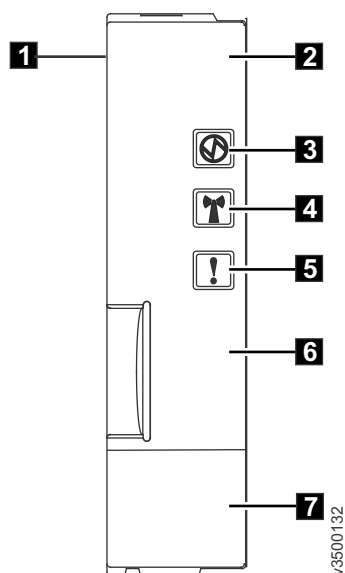
Figure 11. Storwize V7000 Gen2 Left end cap items

1	Enclosure left end cap
2	Storwize V7000 product name
3	Enclosure Power LED
4	Identify (Locator) LED
5	Enclosure Fault LED
6	Enclosure type label (control or expansion)
7	Product Asset Tag label (displays machine type, serial number, part number, and 3D bar code)

The LEDs on the left end cap display the status of the enclosure. Table 10 on page 9 summarizes the meaning of the enclosure LED activity.

Table 10. Storwize V7000 Gen2 Enclosure LED descriptions

LED name	Call out	Symbol	Color	State	Meaning
Enclosure Operational	**3**	⊘	Green	OFF	The enclosure does not have any power or it is running on battery power.
				SLOW BLINK 500 ms ON, 500 ms OFF	All node canisters are powered, but they are in stand-by mode. **Note:** Not used on expansion enclosures.
				ON	At least one canister in the enclosure is operational.
Identify (Locator)	**4**	(•)	Blue	OFF	The enclosure is not being identified.
				ON	The enclosure has been identified. You can use the `chenclosure` command to turn this LED on or off.

Table 10. Storwize V7000 Gen2 Enclosure LED descriptions (continued)

LED name	Call out	Symbol	Color	State	Meaning
Enclosure Fault	**5**	!	Amber	OFF	No faults are identified within any hardware component of the enclosure.
				ON	When ON, this LED indicates one the following: • The enclosure is still in the process of coming up. • An error is logged against the enclosure backplane, a canister, a drive or a PSU within the enclosure. • The enclosure contains a failed drive.

Components in the rear of the enclosure

This topic describes the hardware components in the rear of the enclosure.

Two canisters are located in the middle of each enclosure. The power supply units are located on the left and right of the canisters. The left slot is power supply 1 (**1**), and the right slot is power supply 2 (**2**). Power supply 1 is top side up, and power supply 2 is inverted. The upper slot is canister 1 (**3**), and the lower slot is canister 2 (**4**). Canister 1 is top side up, and canister 2 is inverted.

Figure 12 on page 11 shows the rear view of a model 2076-112 or a model 2076-124 control enclosure. Figure 13 on page 11 shows the rear view of a model 2076-312 or a model 2076-324 control enclosure with the 10 Gbps Ethernet port (**5**). Figure 14 on page 11 shows the rear of an expansion enclosure.
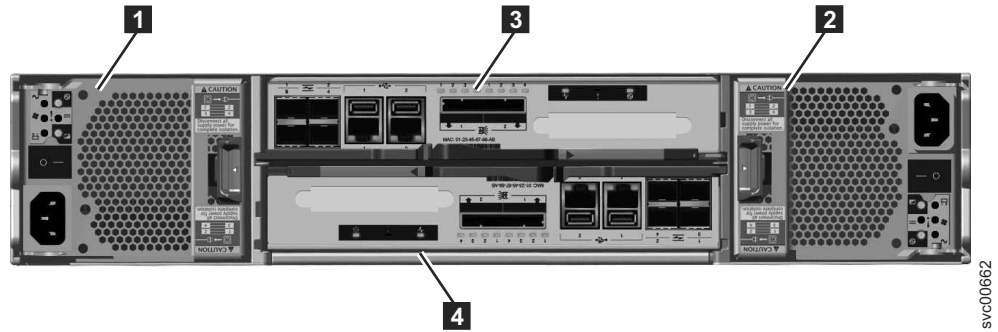
Figure 12. Rear view of a model 2076-112 or a model 2076-124 control enclosure
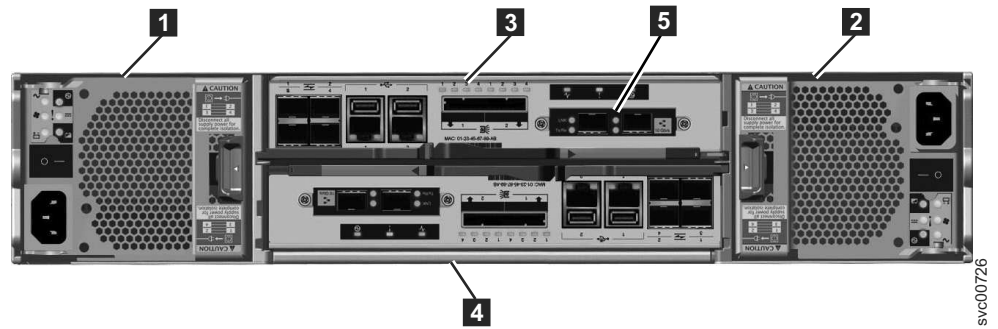
Figure 13. Rear view of a model 2076-312 or a model 2076-324 control enclosure
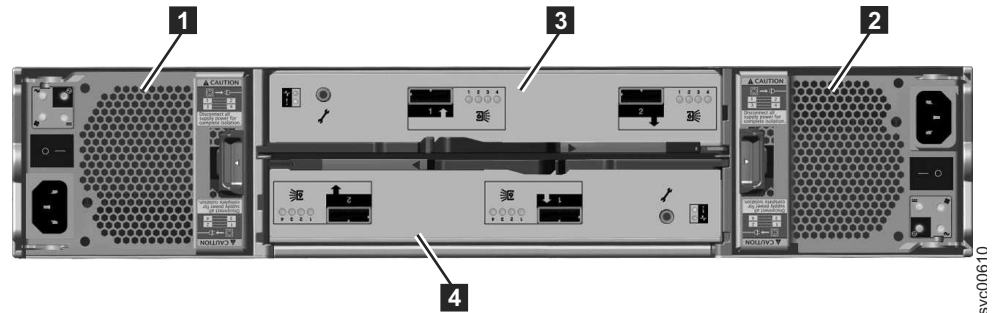
Figure 14. Rear view of a model 2076-212 or a model 2076-224 expansion enclosure

1 Power supply unit 1
2 Power supply unit 2
3 Canister 1
4 Canister 2
1 Node canisters
2 Power supply units

Power supply units for control enclosures

Storwize V7000 enclosures use different power supply units, depending on the generation of your control enclosure model.

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 11. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 12. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Storwize V7000 Gen2 power supply units

Each Storwize V7000 Gen2 enclosure contains two power supply units. Each power supply unit can provide power to the whole enclosure.

Note: The power supply has no power switch. A power supply is active when a power cord is connected to the power connector and to a power source.

Figure 15 on page 13 shows the rear view of a control enclosure and identifies the location of the power supply units and node canisters.
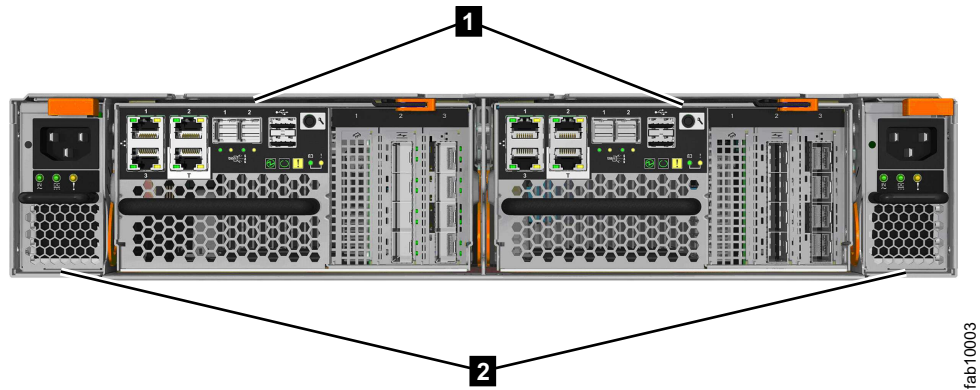
Figure 15. Rear view of a Storwize V7000 Gen2 control enclosure

1 Node canisters
2 Power supply units

Each power supply also contains a fan that cools the power supply unit itself. Cool air is drawn in and passes over each power supply. The warmed air is ejected through the rear of each power supply. For optimal cooling, do not obstruct this airflow. Also, ensure that all enclosure components or fillers are installed while the system is operational.

Storwize V7000 Gen1 Power supply unit and battery for the control enclosure

The Storwize V7000 Gen1 control enclosure contains two power supply units, each with an integrated battery. Each power supply unit can provide power to the whole enclosure.

The two power supply units in the enclosure are installed with one unit top side up and the other inverted. The power supply unit for the control enclosure has six LEDs.

Each power supply unit has a power switch. The switch must be on for the power supply unit to be operational. If the power switches are turned off, or the main power is removed, the integrated batteries temporarily continue to supply power to the node canisters. As a result, the canisters can store configuration data and cached data to their internal drives. Battery power is required only if both power supply units stop operating.

Figure 16 on page 14 shows the location of the LEDs **1** in the rear of the power supply unit.
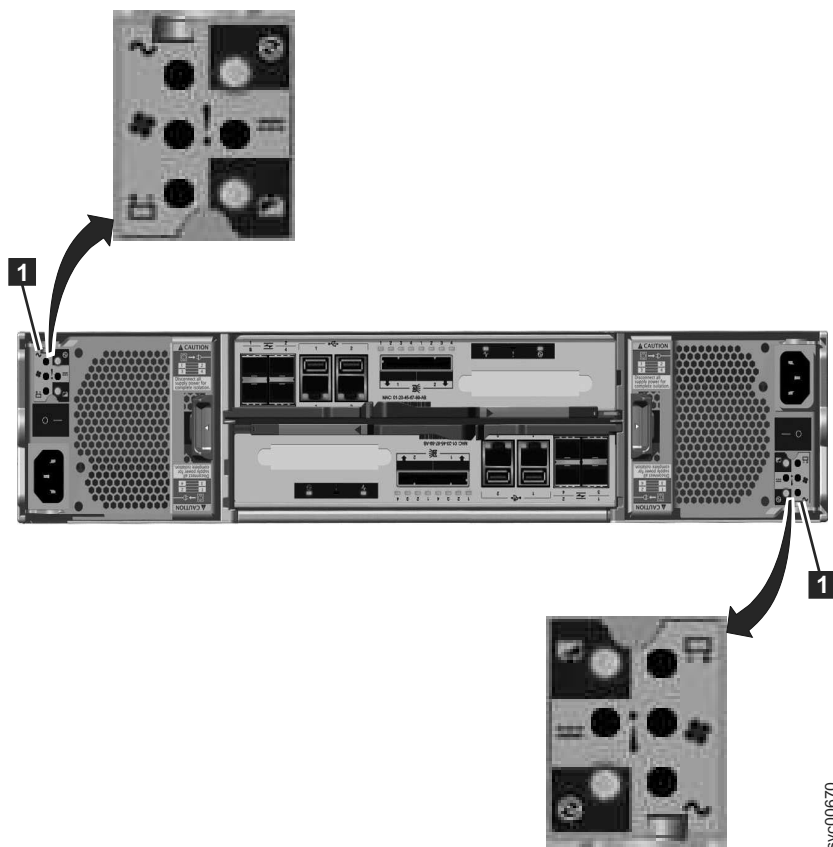
Figure 16. LEDs on the power supply units of the control enclosure

Table 13 identifies the LEDs in the rear of the control enclosure. Each power supply also contains fans that cool the enclosure. Cool air is drawn in through the front of the enclosure. The air passes over the drives, node canisters, and power supplies. The warmed air is ejected through the rear of each power supply. For optimal cooling, do not obstruct the airflow and ensure that all enclosure components or fillers are installed while the system is operational.

Table 13. Power supply unit LEDs in the rear of the control enclosure

Name	Color	Symbol
AC power failure	Amber	∿
Power supply OK	Green	⏻
Fan failure	Amber	✲
DC power failure	Amber	▭
Battery failure	Amber	🔋
Battery state	Green	🔋

See "Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs" on page 112 for help in diagnosing a particular failure.

Power supply units for expansion enclosures

Storwize V7000 enclosures use different power supply units, depending on the generation of your expansion enclosure model.

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 14. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 15. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Storwize V7000 Gen2 power supply units for expansion enclosures

The Storwize V7000 Gen2 expansion enclosure contains two power supply units (PSU).

Figure 17 on page 16 shows the locations of the expansion canisters and the two power supply units in the rear of the expansion enclosure.
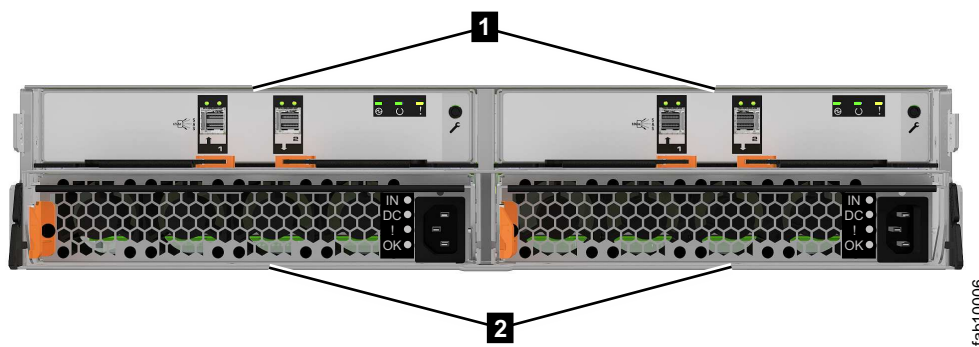
Figure 17. Rear view of a Storwize V7000 Gen2 expansion enclosure

1 Expansion canisters
2 Power supply units

Each power supply unit has four LED indicators (Table 16).

Table 16. Power supply LEDs

Name	Label	Color	Description	
Input status	IN $\approx$	Green	**Off**	No input power detected
			On	Direct current input power detected
Output status	DC $=$	Green	**Off**	PSU is not providing dc output power
			On	PSU is providing dc output power
Fault	!	Amber	**Off**	No fault detected
			On	PSU fault has been detected
			Flash	PSU is being identified. A fault might have been detected.
(None)	OK	Blue	Not used	

See "Procedure: Understanding the Storwize V7000 Gen2 system status from the LEDs" on page 107 for help in diagnosing a particular failure.

Storwize V7000 Gen1 power supply units for expansion enclosures

The Storwize V7000 Gen1 expansion enclosure contains two power supply units (PSU).

The two power supply units in the enclosure are installed with one unit top side up and the other inverted. The power supply unit for the expansion enclosure has four LEDs, two less than the power supply for the control enclosure.

There is a power switch on each of the power supply units. The switch must be on for the power supply unit to be operational. If the power switches are turned off, the power supply units stop providing power to the system.

Table 17 on page 17 shows the locations of the LEDs **1** in the rear of the power supply unit.
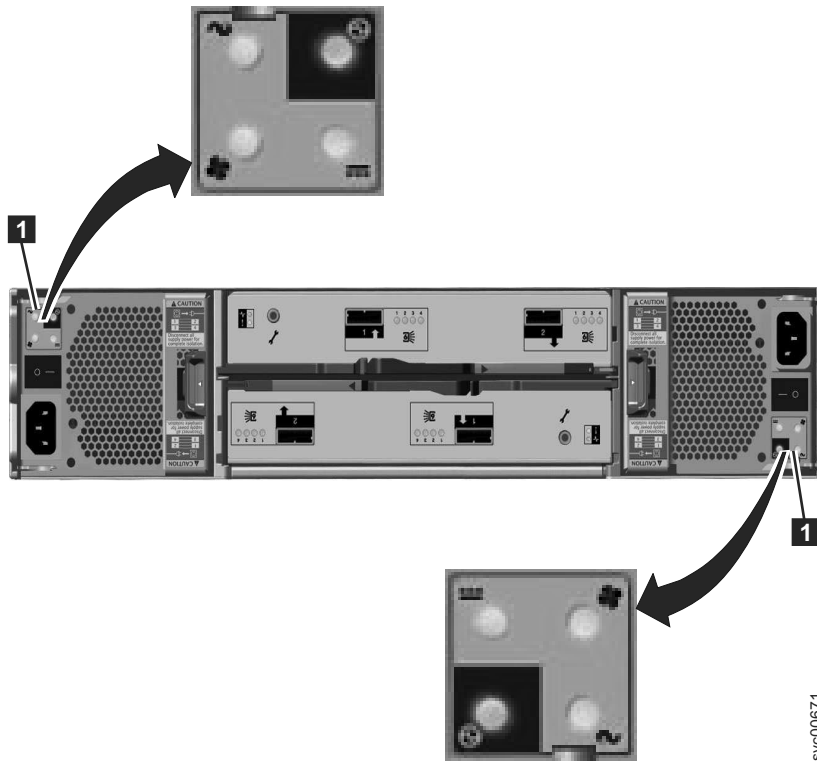
Figure 18. LEDs on the power supply units of the expansion enclosure

svc00671

Table 17 identifies the LEDs in the rear of the expansion enclosure.

Table 17. Power supply unit LEDs in the rear of the expansion enclosure

Name	Color	Symbol
ac power failure	Amber	∿
Power supply OK	Green	⊘
Fan failure	Amber	✿
dc power failure	Amber	⎓

See "Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs" on page 112 for help in diagnosing a particular failure.

Storwize V7000 2076-524 node canister ports and indicators

The node canister has indicators and ports but no controls.

A Fibre Channel over Ethernet (FCoE)/Internet Small Computer System Interface (iSCSI) host interface adapter may be installed in a node canister.

A node canister contains a battery that provides power to the canister as it stores cache and system data to an internal drive in the event of a power failure. This process is known as a *fire hose dump*.

Storwize V7000 Gen2 node canister ports

Each Storwize V7000 2076-524 node canister has ports for connecting Ethernet, iSCSI, and USB peripherals, and optional expansion enclosures.

Chapter 1. Storwize V7000 hardware components **17**

Figure 19 illustrates the location of the ports.

Figure 19. Storwize V7000 2076-524 node canister ports

- **1** USB ports. Each canister has two USB ports. One port is used during installation.
- **2** Ethernet ports. Each canister has two 1 Gbps Ethernet ports.

  Port 1 Must be connected for system management. Can optionally be used for iSCSI host connectivity.

  Port 2 Optional. Can be used for iSCSI host connectivity or to provide an alternative (redundant) management address.

- **3** Serial-attached SCSI (SAS) ports. Each canister has two SAS ports. Ports 1 and 2 are exclusively used for connecting to expansion enclosures.

Storwize V7000 2076-524 node canister indicators

Each Storwize V7000 2076-524 node canister has indicator LEDs that provide status information about the canister.

Using the callout numbers in Figure 20, refer to the tables for a listing of the Storwize V7000 2076-524 node canister LEDs and a description of the meaning of the LED activity.
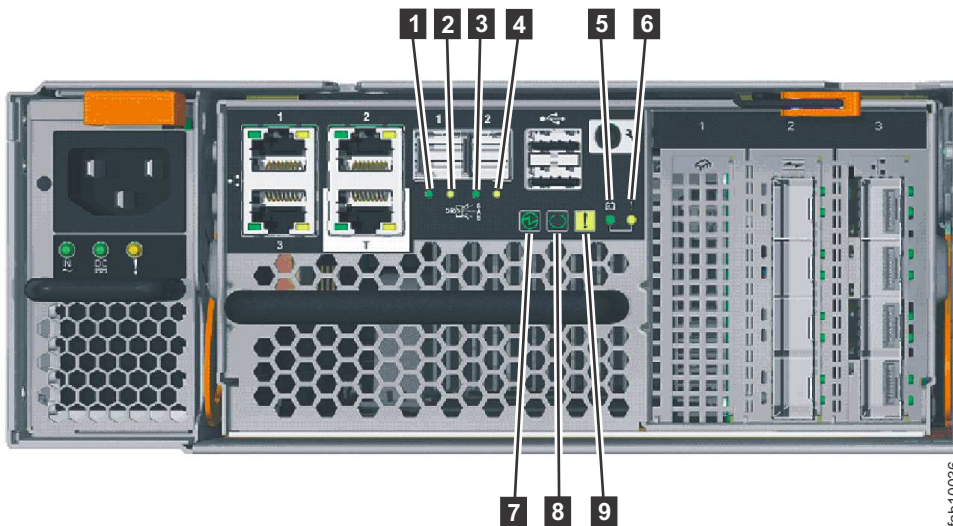
Figure 20. Storwize V7000 2076-524 node canister indicators

Storwize V7000 2076-524 node canister SAS port LEDs

Table 18 depict the status LEDs for SAS ports 1 and 2, and their location in Figure 20 on page 18.

Table 18. Storwize V7000 2076-524 SAS ports 1 and 2 LEDs. Storwize V7000 2076-524 SAS ports 1 and 2 LED legend.

Name	Call out	Symbol	Color	State	Meaning
SAS Port 1 Link	**1**	None	Green	OFF	No link connection on any phys (lanes). The connection is down.
				ON	The port is connected to at least one phy. At least one of the phys to that connector is up.
SAS Port 1 Fault	**2**	None	Amber	OFF	No fault. All four phys have a link connection.
				ON	This status can indicate several different error conditions: • One or more, but not all, of the 4 phys are connected. • Not all 4 phys are at the same speed. • One or more of the connected phys are attached to an address different from the others. • An unsupported device is plugged in to this SAS port.

Table 18. Storwize V7000 2076-524 SAS ports 1 and 2 LEDs (continued). Storwize V7000 2076-524 SAS ports 1 and 2 LED legend.

Name	Call out	Symbol	Color	State	Meaning
SAS Port 2 Link	**3**	None	Green	OFF	No link connection on any phys (lanes). The connection is down.
				ON	The port is connected to at least one phy. At least one of the lanes to that connector is up.
SAS Port 2 Fault	**4**	None	Amber	OFF	No fault. All four phys have a link connection.
				ON	This status can indicate several different error conditions: • One or more, but not all, of the 4 phys are connected. • Not all 4 phys are at the same speed. • One or more of the connected phys are attached to an address different from the others. • An unsupported device is plugged in to this SAS port.

Storwize V7000 2076-524 node canister battery status LEDs

Table 19 on page 21 show battery status LEDs and their location in Figure 20 on page 18.

Table 19. Storwize V7000 2076-524 battery status LEDs

Name	Call out	Color	State	Meaning
Battery status	**5**	Green	OFF	Indicates that the battery is not available for use. The battery might be missing or a battery fault was detected.
			FAST BLINK	The battery has insufficient charge to complete a "fire hose" dump.
			BLINK	The battery has sufficient charge to complete a single "fire hose" dump.
			ON	The battery has sufficient charge to complete at least two "fire hose" dumps.
Battery fault	**6**	Amber	OFF	No fault. An exception to this would be where a battery has insufficient charge to complete a single "fire hose" dump. Refer to the documentation for the Battery status LED.
			ON	A battery fault was detected.

Storwize V7000 2076-524 node canister system status LEDs

Table 20 on page 22 show system status LEDs and their location in Figure 20 on page 18.

Table 20. Storwize V7000 2076-524 node canister system status LEDs

Name	Call out	Color	State	Meaning
Power	**7**	Green	OFF	No power is available or power is coming from the battery.
			SLOW BLINK	Power is available but the main processor is not running; this state is called *standby mode*.
			FAST BLINK	In self-test.
			ON	Power is available and the system code is running.
Status	**8**	Green	OFF	The system code has not started. The system is off, in standby, or in self-test.
			BLINK	The canister is in candidate or service state. It is not completing I/O operations. It is safe to remove the node.
			FAST BLINK	The canister is active, able to complete I/O operations, or starting.
			ON	The canister is active, able to complete I/O operations, or starting. The node is part of a cluster.

Name	Call out	Color	State	Meaning
Canister fault	**9**	Amber	OFF	The canister can function as an active member of the system. If the node canister has a problem, it is not severe enough to stop the node canister from completing I/O operations.
			BLINK	The canister is being identified. There might or might not be a fault condition.
			ON	The node is in service state or an error exists that might be stopping the system code from starting. The node canister cannot become active in the system until the problem is resolved. You must determine the cause of the error before you replace the node canister. The error might be due to insufficient battery charge. To resolve this error, wait for the battery to charge.

USB ports in the Storwize V7000 Gen2 node canister

Two USB ports are located on each Storwize V7000 Gen2 node canister.

The USB ports are numbered 1 on top and 2 on the bottom as shown in Figure 21 on page 24. One port is used during installation.
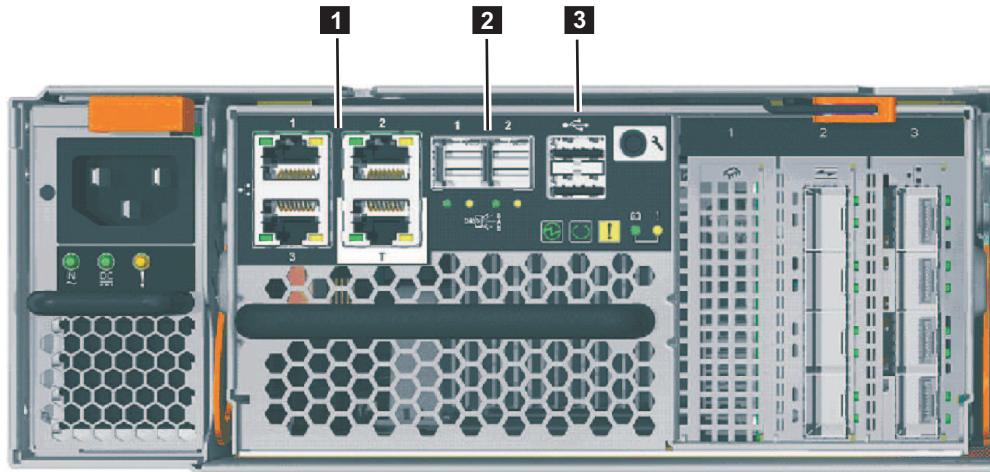
Figure 21. USB ports on the Storwize V7000 Gen2 node canister

1 Ethernet ports. The technician port is labeled **T**. It is used during installation.

2 Serial-attached SCSI (SAS) ports.

3 USB ports. Each canister has two USB ports.

The USB ports have no indicators.

Storwize V7000 2076-524 host interface adapter ports and indicators

An 8 Gbps or 10 Gbps host interface adapter can be installed in each node canister.

Storwize V7000 2076-524 Fibre Channel host interface adapter ports and indicators:

If you specified Fibre Channel host interface adapters for your Storwize V7000 2076-524 system, an adapter is preinstalled in each node canister.

Each 8 Gbps Fibre Channel 4-port host interface adapter (feature code ACHK) can have from two to four short wave (SW) small form-factor pluggable (SFP) transceivers installed. Cap unused ports with safety caps.

Fibre Channel host interface adapter ports

Fibre Channel ports **1** are in 1 - 4 order, starting at the top. Ports and their indicators are shown in Figure 22 on page 25.

Each port can have up to an 8 Gbps SW SFP transceiver installed. Each transceiver connects to a host or Fibre Channel switch with an LC-to-LC Fibre Channel cable.
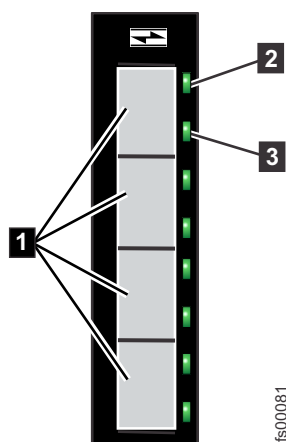
Figure 22. Fibre Channel ports and indicators

1 Fibre Channel 8 Gbps ports (x4)

2 Link-state LED (x4 - one for each port)

3 Speed-state LED (x4 - one for each port)

Fibre Channel host interface adapter indicators

Each Fibre Channel port has two green LED indicators. The link-state LED **2** is above the speed-state LED **3** for each port. Consider the LEDs as a pair to determine the overall link state, which is decoded in Table 21.

Table 21. Fibre Channel host interface adapter port-state LEDs

Link-state LED	Speed-state LED	Link state
OFF	OFF	Inactive
ON or FLASHING	OFF	Active low speed (2 Gbps)
ON or FLASHING	FLASHING	Active medium speed (4 Gbps)
ON or FLASHING	ON	Active high speed (8 Gbps)

One or two Fibre Channel interface adapters can be installed in each node canister. They can be installed in slots 2 and 3 of the node canister. When a single interface adapter is installed in either slot 2 or slot 3, the Fibre Channel ports on the adapter are numbered 1, 2, 3, and 4.

Storwize V7000 2076-524 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapter ports and indicators:

You can install 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapters for your Storwize V7000 2076-524 system, the adapter is preinstalled in each node canister.

The 4 port FCoE/iSCSI host interface adapter is used for Fibre Channel over Ethernet (FCoE) or internet Small Computer System Interface (iSCSI) connections to host systems or for Fibre Channel over Ethernet connections to host system or storage systems. Each port can support simultaneous FCoE and iSCSI connections. The Small Form-factor Pluggable (SFP) transceivers that are installed on the adapter support data transfer speeds of 10 Gbps.

Note: This adapter can be installed only in slots 2 and 3. Figure 23 shows two 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapters, both installed in slot 3.
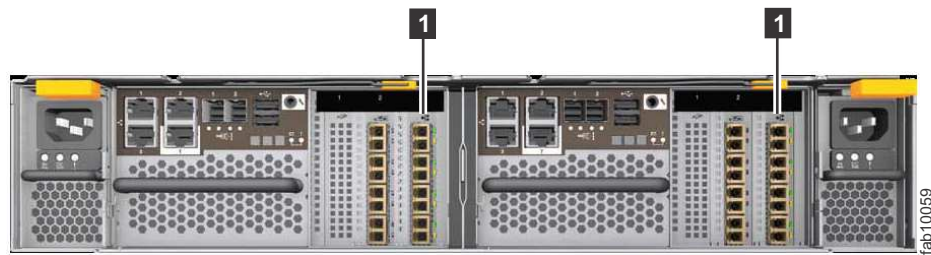
1 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapter

Figure 23. Example of installed 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapters

Storwize V7000 2076-524 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapter ports

The adapter has four Ethernet ports, none of which are used for system management. The ports are named 1, 2, 3 and 4 (Figure 24) when installed in a slot.

1 Port 1
2 Port 2
3 Port 3
4 Port 4

Figure 24. 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapter ports

Storwize V7000 2076-524 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapter indicators

Each port has two LED indicators, one green and one amber (see Figure 25 on page 27).

1 Green LED

2 Amber LED
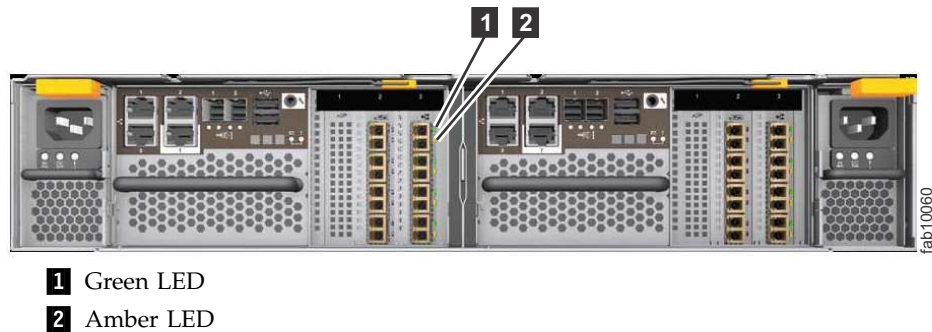
Figure 25. 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapter indicator LEDs

The LED states and their meanings are explained in Table 22.

Table 22. Storwize V7000 2076-524 host interface adapter LED states and meanings

Green LED	Amber LED	Meaning
OFF	OFF	The port is not configured in flex hardware and the port is not active in the current profile. For example, in the 2-by-16 Gbps profile, two ports are not active.
OFF	ON	The port is configured, but is not connected or the negotiation of the link failed (the link is not detected at the transport layer).
ON	OFF	The link is up and is running at the configured speed. **Note:** This code does not indicate logical connectivity, such as the completion of FLOGI (Fabric login) or FIP (Fibre Channel over Ethernet Initialization Protocol).
ON	ON	The link is up and is running at less than the configured (degraded) speed.

Node canister ports and indicators

The node canister has indicators and ports but no controls.

Fibre Channel ports and indicators

The Fibre Channel port LEDs show the speed of the Fibre Channel ports and activity level.

Each node canister has four Fibre Channel ports located on the left side of the canister as shown in Figure 26 on page 28. The ports are in two rows of two ports. The ports are numbered 1 - 4 from left to right and top to bottom.

Note: The reference to the left and right locations applies to canister 1, which is the upper canister. The port locations are inverted for canister 2, which is the lower canister.
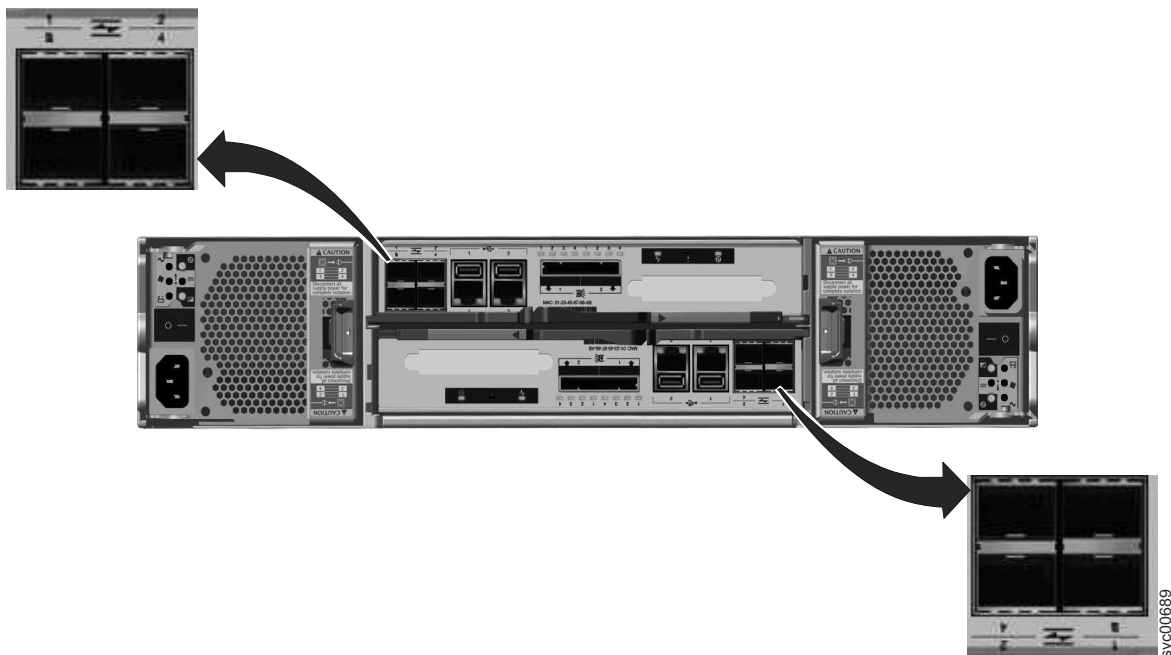
Figure 26. Fibre Channel ports on the node canisters

There are two green LEDs associated with each port: the speed LED and the link activity LED. These LEDs are in the shape of a triangle. The LEDs are located in between the two rows of the ports as shown in Figure 27. Figure 27 shows the LEDs for the Fibre Channel ports on canister 1. Each LED points to the associated port. The first and second LEDs in each set show the speed state, and the third and fourth LEDs show the link state.
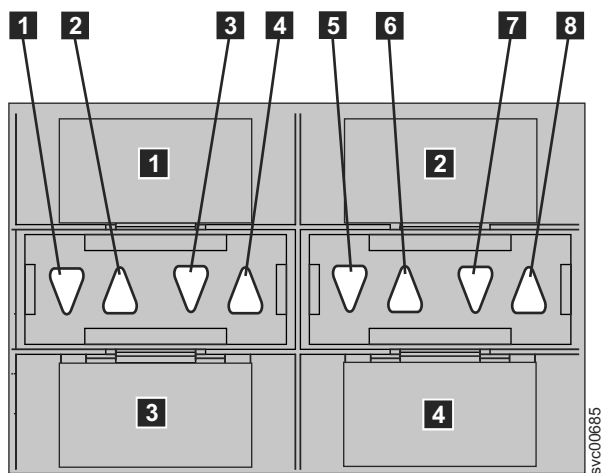
Figure 27. LEDs on the Fibre Channel ports

Table 23. Fibre Channel port LED locations on canister 1

Associated port	LED location	LED status
Port 3 **3**	First LED between ports 1 and 3 **1**	Speed
Port 1 **1**	Second LED between ports 1 and 3 **2**	Speed

Table 23. Fibre Channel port LED locations on canister 1 (continued)

Associated port	LED location	LED status
Port 3 **3**	Third LED between ports 1 and 3 **3**	Link
Port 1 **1**	Fourth LED between ports 1 and 3 **4**	Link
Port 4 **4**	First LED between ports 2 and 4 **5**	Speed
Port 2 **2**	Second LED between ports 2 and 4 **6**	Speed
Port 4 **4**	Third LED between ports 2 and 4 **7**	Link
Port 2 **2**	Fourth LED between ports 2 and 4 **8**	Link

Table 24 provides the status descriptions for the LEDs on the Fibre Channel ports.

Table 24. Fibre Channel port LED status descriptions

Speed state LED	Link state LED	Link state
Off	Off	Inactive
Off	On or flashing	Active low speed (2 Gbps)
Flashing	On or flashing	Active medium speed (4 Gbps)
On	On or flashing	Active high speed (8 Gbps)

USB ports in Storwize V7000 Gen1 node canisters

Two USB ports are located side by side on each Storwize V7000 Gen1 node canister.

The USB ports are numbered 1 on the left and 2 on the right as shown in Figure 28 on page 30. One port is used during installation.

Note: The reference to the left and right locations applies to canister 1, which is the upper canister. The port locations are inverted for canister 2, which is the lower canister.
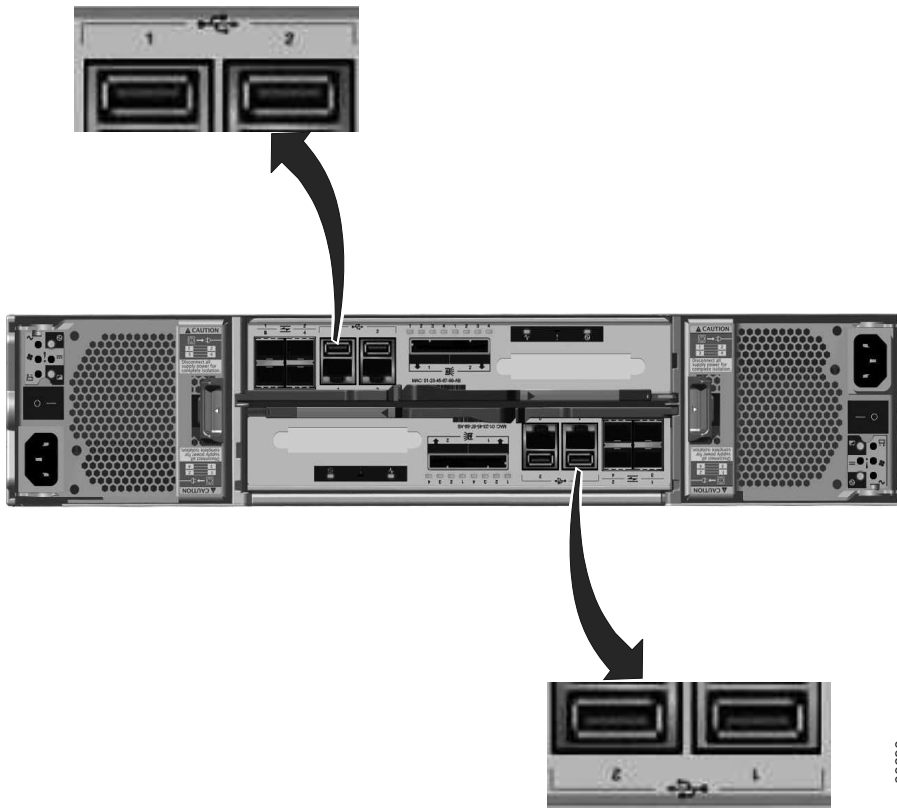
Figure 28. USB ports on the node canisters

The USB ports have no indicators.

Ethernet ports and indicators

Ethernet ports are located side by side on the rear of the node canister. All control enclosure models have two 1 Gbps Ethernet ports per node canister. Model 2076-312 and model 2076-324 also have two 10 Gbps Ethernet ports per node canister.

For the 1 Gbps support, the Ethernet ports are numbered 1 on the left and 2 on the right as shown in Figure 29 on page 31. Port 1 must be connected; the use of port 2 is optional. Two LEDs are associated with each port.

Note: The reference to the left and right locations applies to canister 1, which is the upper canister. The port locations are inverted for canister 2, which is the lower canister.
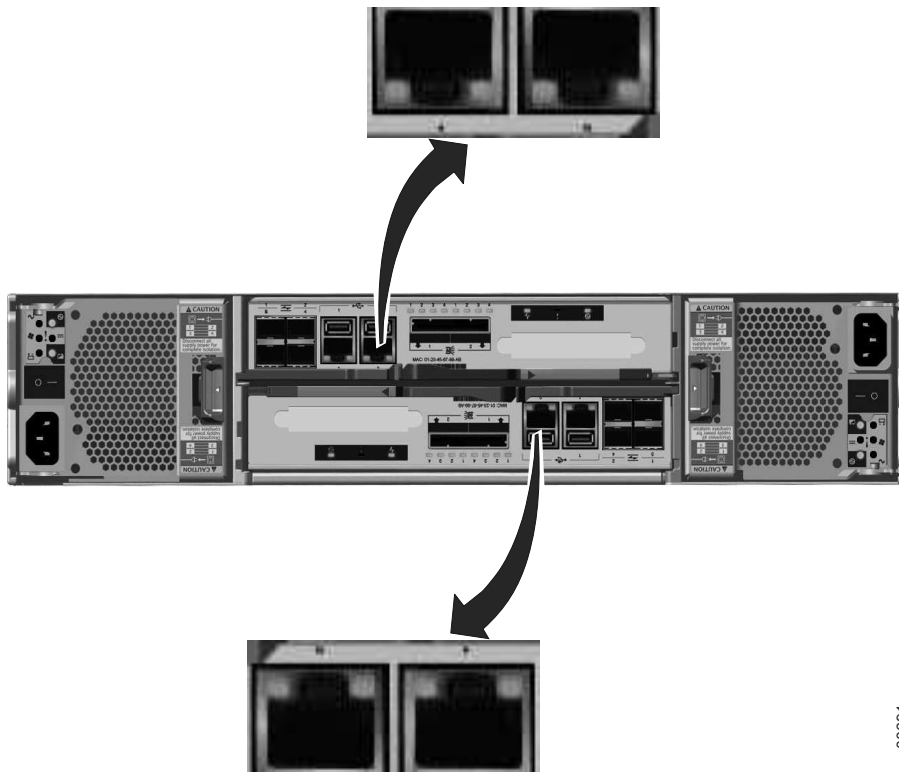
Figure 29. Ethernet ports on the 2076-112 and 2076-124 node canisters

Table 25 provides a description of the two LEDs.

Table 25. 1 Gbps Ethernet port LEDs

Name	Description	Color
Link speed (LED on right of upper canister)	The LED is on when there is a link connection; otherwise, the LED is off.	Green
Activity (LED on left of upper canister)	The LED is flashing when there is activity on the link; otherwise, the LED is off.	Yellow

Figure 30 shows the location of the 10 Gbps Ethernet ports.
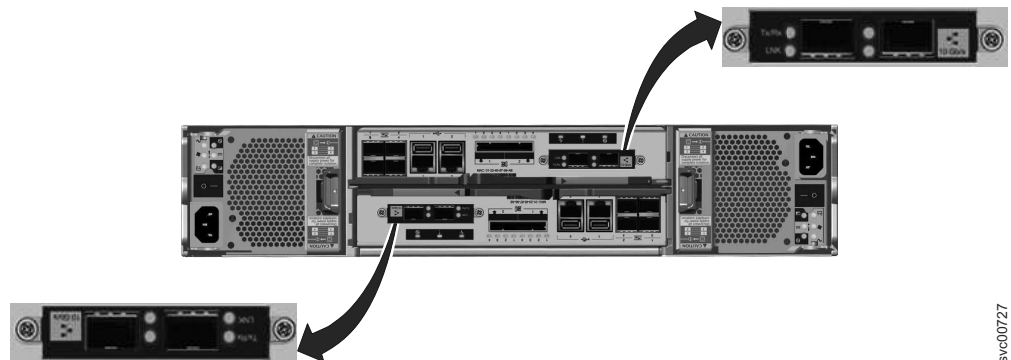
Figure 30. 10 Gbps Ethernet ports on the 2076-312 and 2076-324 node canisters

Table 26 on page 32 provides a description of the LEDs.

Table 26. 10 Gbps Ethernet port LEDs

Name	Symbol	Description	Color
Activity	Tx/Rx	The LED is flashing when there is activity on the link; otherwise, the LED is off.	Green
Link	LNK	The LED is on when there is a link connection; otherwise, the LED is off.	Amber

Node canister SAS ports and indicators

Two serial-attached SCSI (SAS) ports are located side by side in the rear of the node canister.

The SAS ports are numbered 1 on the left and 2 on the right as shown in Figure 31. Port 1 is used if you add one expansion enclosure. Port 2 is used if you add a second expansion enclosure. Each port provides four data channels.

Note: The reference to the left and right locations applies to canister 1, which is the upper canister. The port locations are inverted for canister 2, which is the lower canister.
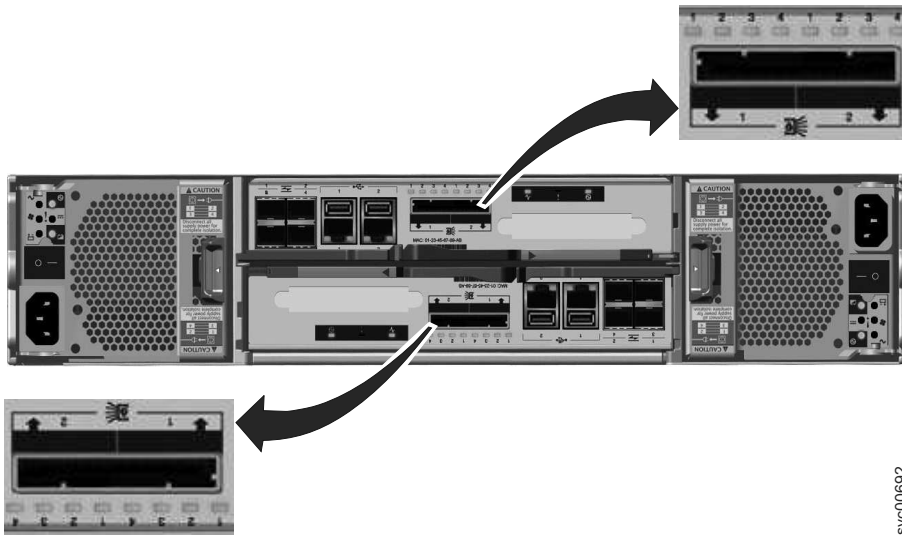
Figure 31. SAS ports on the node canisters.

SAS ports must be connected to Storwize V7000 enclosures only. See "Problem: Storwize V7000 Gen1 SAS cabling not valid" on page 97 for help in attaching the SAS cables.

Four LEDs are located with each port. Each LED describes the status of one data channel within the port. The data channel number is shown with the LED.

Table 27. SAS port LEDs on the node canister

LED state	Description
Off	No link is connected.
Flashing	The link is connected and has activity.
On	The link is connected.

Node canister LEDs

Each node canister has three LEDs that provide status and identification for the node canister.

The three LEDs are located in a horizontal row near the upper right of the canister **1** . Figure 32 shows the rear view of the node canister LEDs.
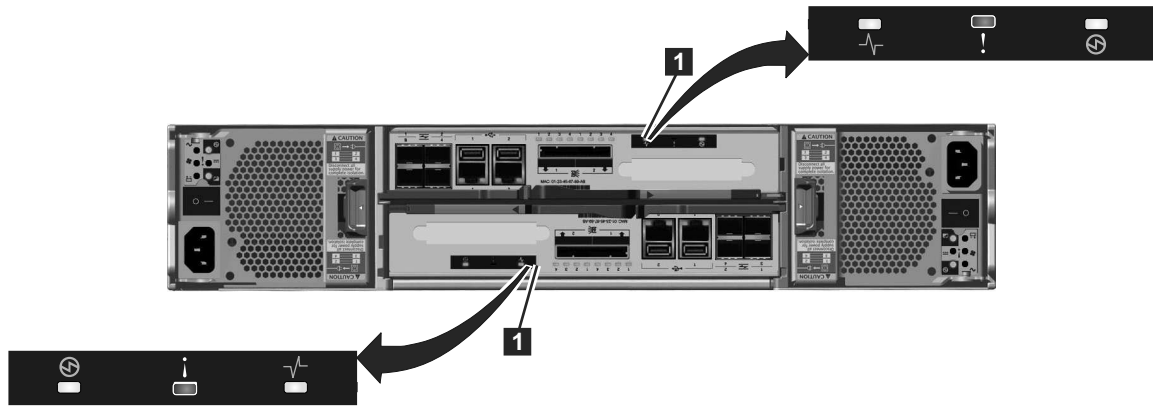
Figure 32. LEDs on the node canisters

Note: The reference to the left and right locations applies to canister 1, which is the upper canister. The port locations are inverted for canister 2, which is the lower canister.

Table 28. Node canister LEDs

Name	Description	Color	Symbol
System status	Indicates the status of the node. • The on status indicates that the node is active, that is, it is an active member of a clustered system. When the node is active, do not remove it. • The off state indicates there is no power to the canister or the canister is in standby mode. These conditions can cause the off state: – The main processor is off and only the service processor is active. – A power-on self-test (POST) is running on the canister. – The operating system is loading. • The flashing status indicates that the node is in candidate state or service state. It is not able to perform I/O in a system. When the node is in either of these states, it can be removed. Do not remove the canister unless directed by a service procedure.	Green	-\/-

Table 28. Node canister LEDs (continued)

Name	Description	Color	Symbol
Fault	Indicates if a fault is present and identifies which canister. • The on status indicates that the node is in service state or an error exists that might be preventing the code from starting. Do not assume that this status indicates a hardware error. Further investigation is required before replacing the node canister. • The off status indicates that the node is a candidate or is active. This status does not mean that there is not a hardware error on the node. Any error that was detected is not severe enough to stop the node from participating in a system. • The flashing status indicates that the canister is being identified. This status might or might not be a fault.	Amber	!
Power	Indicates if power is available and the boot status of the canister. • The on status indicates that the canister is powered on and that the main processor or processors are running. • The off status indicates that no power is available. • The slow flashing (1 Hz) status indicates that power is available and that the canister is in standby mode. The main processor or processors are off and only the service processor is active. • The fast flashing (2 Hz) indicates that the canister is running the power-on self-test (POST).	Green	⏻

Notes:

1. If the system status LED is on and the fault LED is off, the node canister is an active member of a system.

2. If the system status LED is on and the fault LED is on, there is a problem establishing a system.

For a more complete identification of the system LEDs, go to "Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs" on page 112.

Expansion canister ports and indicators

An expansion canister is one of two canisters that is located in the rear of a SAS expansion enclosure. The expansion canister has no controls.

There is a diagnostic port on the right of the expansion canister. There are no indicators that are associated with the port. There are no defined procedures that use the port.

Storwize V7000 Gen2 expansion canister SAS ports and indicators

Two SAS ports are located in the rear of the Storwize V7000 Gen2 expansion canister.

SAS ports are numbered at the bottom of the port, with 1 on the left and 2 on the right, as shown in Figure 33 on page 35. Use of port 1 is required. Use of port 2 is optional. Each port connects four data channels.
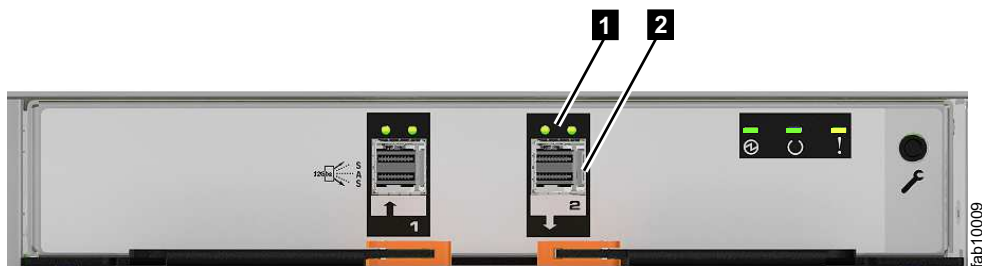
Figure 33. SAS ports and LEDs at rear of expansion canister

Figure 33 has callouts to show the location of the LEDs and the port for SAS port 2:

1 Port 2 LEDs

2 Port 2 12 Gbps SAS port

Table 29 describes LED states for each of the two LEDs per SAS port. The link LED is on the left of each set of ports.

Table 29. SAS port LEDs on the expansion canister

Name	Color	State	Meaning
SAS Port 1 Link	Green	OFF	No link connection on any phys (lanes). The connection is down.
		ON	There is a connection on at least one phy. At least one of the phys to that connector is up.
SAS Port 1 Fault	Amber	OFF	No fault. All four phys have a link connection.
		ON	This can indicate a number of different error conditions: • One or more, but not all, of the 4 phys are connected. • Not all 4 phys are at the same speed. • One or more of the connected phys are attached to an address different from the others
SAS Port 2 Link	Green	OFF	No link connection on any phys (lanes). The connection is down.
		ON	There is a connection on at least one phy. At least one of the lanes to that connector is up.
SAS Port 2 Fault	Amber	OFF	No fault. All four phys have a link connection.
		ON	This can indicate a number of different error conditions: • One or more, but not all, of the 4 phys are connected. • Not all 4 phys are at the same speed. • One or more of the connected phys are attached to an address different from the others

Storwize V7000 Gen1 expansion canister SAS ports and indicators

Two SAS ports are located in the rear of the Storwize V7000 Gen1 expansion canister.

SAS ports are numbered 1 on the left and 2 on the right as shown in Figure 34. Use of port 1 is required. Use of port 2 is optional. Each port connects four data channels.

Note: The reference to the left and right locations applies to canister 1, which is the upper canister. The port locations are inverted for canister 2, which is the lower canister.
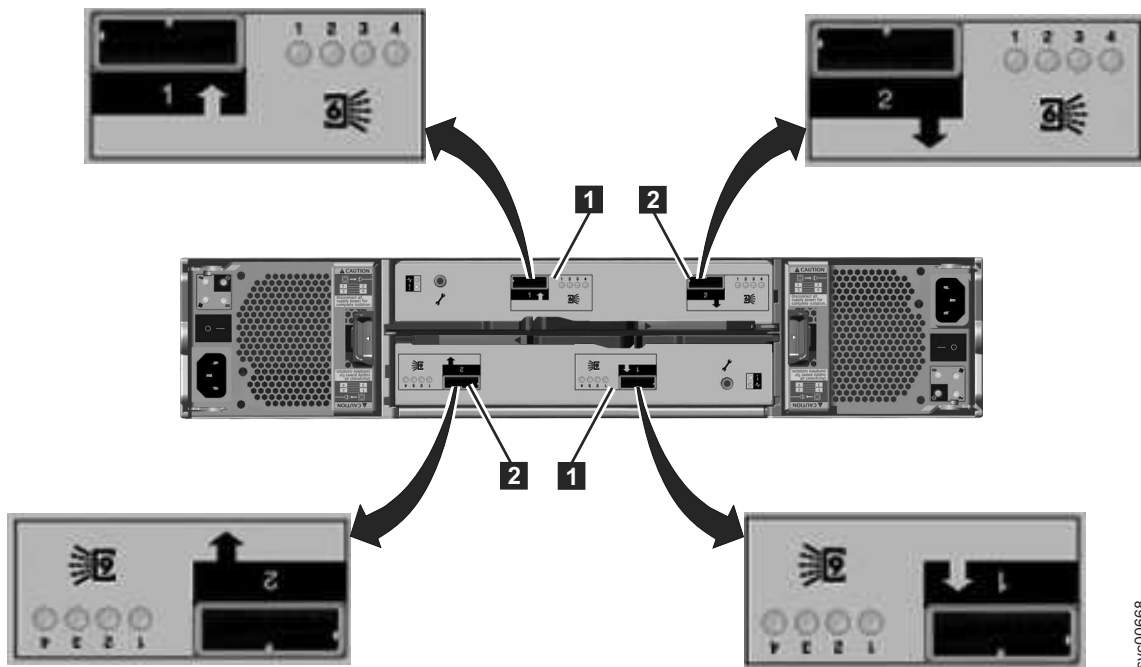
Figure 34. SAS ports and LEDs in rear of expansion enclosure

- **1** Port 1, 6 Gbps SAS port and LEDs
- **2** Port 2, 6 Gbps SAS port and LEDs

Four LEDs are located with each port. Each LED describes the status of one data channel within the port. The data channel is shown with the LED. Table 30 indicates the status of the SAS port LEDs on the expansion canister.

Table 30. SAS port LEDs on the expansion canister

LED state	Description
Off	No link is connected.
Flashing	The link is connected and has activity.
On	The link is connected.

Storwize V7000 Gen2 expansion canister LEDs

Each Storwize V7000 Gen2 expansion canister has three LEDs that provide status and identification for the expansion canister.

Three LEDs are located in a horizontal row on the right side (when viewed from the rear) of the expansion canister. Figure 35 shows the expansion canister LEDs, and Table 31 describes the LEDs.
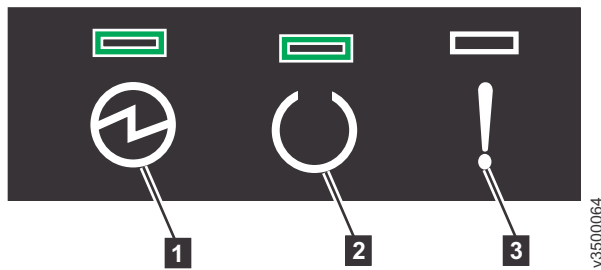
Figure 35. Expansion canister LEDs

Table 31. Expansion canister LED descriptions

Name	Description	Color	Symbol
1 Power	Indicates whether the expansion canister has power. • If the LED is on, the canister has power. • If the LED is off, the canister does not have power.	Green	
2 Status	Indicates whether the expansion canister is active. • If the LED is on, the canister is active. • If the LED is off, the canister is not active. • If the LED is flashing, there is a vital product data (VPD) error.	Green	
3 Fault	Indicates whether a fault is present and identifies the expansion canister. • If the LED is on, a fault exists. • If the LED is off, no fault exists. • If the LED is flashing, the expansion canister is being identified. This status might or might not be a fault.	Amber	

Expansion canister LEDs

Each expansion canister has two LEDs that provide status and identification for the expansion canister.

The two LEDs are located in a vertical row on the left side of the canister. Figure 36 on page 38 shows the LEDs (**1**) in the rear of the expansion canister.
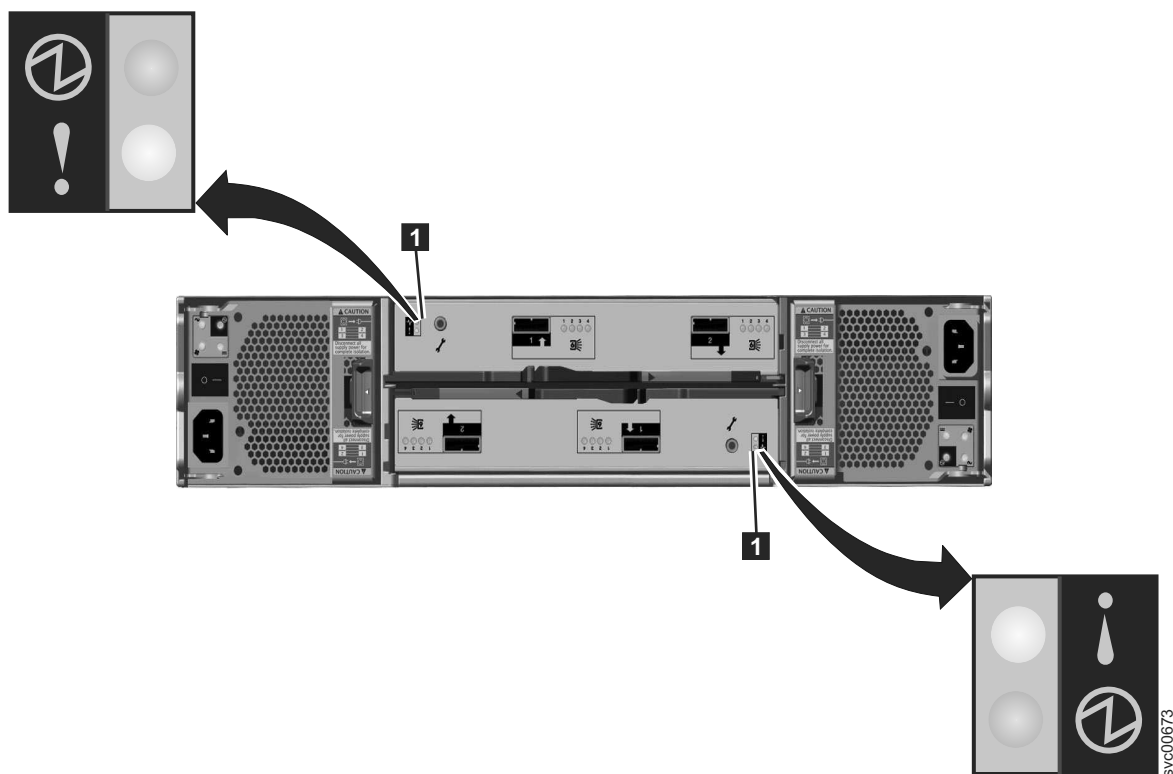
Figure 36. LEDs on the expansion canisters

Table 32 describes the expansion canister LEDs.

Table 32. Expansion canister LEDs

Name	Description	Color	Symbol
Status	Indicates if the canister is active. • If the LED is on, the canister is active. • If the LED is off, the canister is not active. • If the LED is flashing, there is a vital product data (VPD) error.	Green	⊕
Fault	Indicates if a fault is present and identifies the canister. • If the LED is on, a fault exists. • If the LED is off, no fault exists. • If the LED is flashing, the canister is being identified. This status might or might not be a fault.	Amber	!

Chapter 2. Best practices for troubleshooting

Taking advantage of certain configuration options, and ensuring vital system access information has been recorded, makes the process of troubleshooting easier.

Starting statistics collection

The system collects statistics over an interval and creates files that can be viewed.

Introduction

For each collection interval, the management GUI creates four statistics files: one for managed disks (MDisks), named **Nm_stat**; one for volumes and volume copies, named **Nv_stat**; one for nodes, named **Nn_stat**; and one for SAS drives, named **Nd_stat**. The files are written to the /dumps/iostats directory on the node. To retrieve the statistics files from the non-configuration nodes onto the configuration node, **svctask cpdumps** command must be used.

A maximum of 16 files of each type can be created for the node. When the 17th file is created, the oldest file for the node is overwritten.

Fields

The following fields are available for user definition:

Interval
> Specify the interval in minutes between the collection of statistics. You can specify 1 - 60 minutes in increments of 1 minute.

Tables

The following tables describe the information that is reported for individual nodes and volumes.

Table 33 describes the statistics collection for MDisks, for individual nodes.

Table 33. Statistics collection for individual nodes

Statistic name	Description
id	Indicates the name of the MDisk for which the statistics apply.
idx	Indicates the identifier of the MDisk for which the statistics apply.
rb	Indicates the cumulative number of blocks of data that is read (since the node has been running).
re	Indicates the cumulative read external response time in milliseconds for each MDisk. The cumulative response time for disk reads is calculated by starting a timer when a **SCSI read** command is issued and stopped when the command completes successfully. The elapsed time is added to the cumulative counter.
ro	Indicates the cumulative number of MDisk read operations that are processed (since the node is running).

Table 33. Statistics collection for individual nodes (continued)

rq	Indicates the cumulative read queued response time in milliseconds for each MDisk. This response is measured from above the queue of commands to be sent to an MDisk because the queue depth is already full. This calculation includes the elapsed time that is taken for **read** commands to complete from the time they join the queue.
wb	Indicates the cumulative number of blocks of data written (since the node is running).
we	Indicates the cumulative write external response time in milliseconds for each MDisk. The cumulative response time for disk writes is calculated by starting a timer when a **SCSI write** command is issued and stopped when the command completes successfully. The elapsed time is added to the cumulative counter.
wo	Indicates the cumulative number of MDisk write operations processed (since the node is running).
wq	Indicates the cumulative write queued response time in milliseconds for each MDisk. This time is measured from above the queue of commands to be sent to an MDisk because the queue depth is already full. This calculation includes the elapsed time taken for **write** commands to complete from the time they join the queue.

Table 34 describes the VDisk (volume) information that is reported for individual nodes.

Note: MDisk statistics files for nodes are written to the /dumps/iostats directory on the individual node.

Table 34. Statistic collection for volumes for individual nodes

Statistic name	Description
id	Indicates the volume name for which the statistics apply.
idx	Indicates the volume for which the statistics apply.
rb	Indicates the cumulative number of blocks of data read (since the node is running).
rl	Indicates the cumulative read response time in milliseconds for each volume. The cumulative response time for volume reads is calculated by starting a timer when a **SCSI read** command is received and stopped when the command completes successfully. The elapsed time is added to the cumulative counter.
rlw	Indicates the worst read response time in microseconds for each volume since the last time statistics were collected. This value is reset to zero after each statistics collection sample.
ro	Indicates the cumulative number of volume read operations processed (since the node has been running).
wb	Indicates the cumulative number of blocks of data written (since the node is running).
wl	Indicates the cumulative write response time in milliseconds for each volume. The cumulative response time for volume writes is calculated by starting a timer when a **SCSI write** command is received and stopped when the command completes successfully. The elapsed time is added to the cumulative counter.

Table 34. Statistic collection for volumes for individual nodes (continued)

wlw	Indicates the worst write response time in microseconds for each volume since the last time statistics were collected. This value is reset to zero after each statistics collection sample.
wo	Indicates the cumulative number of volume write operations processed (since the node is running).
wou	Indicates the cumulative number of volume write operations that are not aligned on a 4K boundary.
xl	Indicates the cumulative read and write data transfer response time in milliseconds for each volume since the last time the node was reset. When this statistic is viewed for multiple volumes and with other statistics, it can indicate whether the latency is caused by the host, fabric, or the Storwize V7000.

Table 35 describes the VDisk information that is related to Metro Mirror or Global Mirror relationships that is reported for individual nodes.

Table 35. Statistic collection for volumes that are used in Metro Mirror and Global Mirror relationships for individual nodes

Statistic name	Description
gwl	Indicates cumulative secondary write latency in milliseconds. This statistic accumulates the cumulative secondary write latency for each volume. You can calculate the amount of time to recovery from a failure based on this statistic and the **gws** statistics.
gwo	Indicates the total number of overlapping volume writes. An overlapping write is when the logical block address (LBA) range of write request collides with another outstanding request to the same LBA range and the write request is still outstanding to the secondary site.
gwot	Indicates the total number of fixed or unfixed overlapping writes. When all nodes in all clusters are running Storwize V7000 version 4.3.1, this records the total number of write I/O requests received by the Global Mirror feature on the primary that overlapped. When any nodes in either cluster are running Storwize V7000 versions earlier than 4.3.1, this value does not increment.
gws	Indicates the total number of write requests issued to the secondary site.

Table 36 describes the port information that is reported for individual nodes

Table 36. Statistic collection for node ports

Statistic name	Description
bbcz	Indicates the total time in microseconds for which the buffer credit counter was at zero. Note that this statistic is only reported by 8 Gbps Fibre Channel ports. For other port types, this is 0.
cbr	Indicates the bytes received from controllers.
cbt	Indicates the bytes transmitted to disk controllers.
cer	Indicates the commands received from disk controllers.
cet	Indicates the commands initiated to disk controllers.
dtdc	Indicates the number of transfers that experienced excessive data transmission delay.

Table 36. Statistic collection for node ports (continued)

dtdm	Indicates the number of transfers that had their data transmission delay measured.
dtdt	Indicates the total time in microseconds for which data transmission was excessively delayed.
hbr	Indicates the bytes received from hosts.
hbt	Indicates the bytes transmitted to hosts.
her	Indicates the commands received from hosts.
het	Indicates the commands initiated to hosts.
icrc	Indicates the number of CRC that are not valid.
id	Indicates the port identifier for the node.
itw	Indicates the number of transmission word counts that are not valid.
lf	Indicates a link failure count.
lnbr	Indicates the bytes received to other nodes in the same cluster.
lnbt	Indicates the bytes transmitted to other nodes in the same cluster.
lner	Indicates the commands received from other nodes in the same cluster.
lnet	Indicates the commands initiated to other nodes in the same cluster.
lsi	Indicates the lost-of-signal count.
lsy	Indicates the loss-of-synchronization count.
pspe	Indicates the primitive sequence-protocol error count.
rmbr	Indicates the bytes received to other nodes in the other clusters.
rmbt	Indicates the bytes transmitted to other nodes in the other clusters.
rmer	Indicates the commands received from other nodes in the other clusters.
rmet	Indicates the commands initiated to other nodes in the other clusters.
wwpn	Indicates the worldwide port name for the node.

Table 37 describes the node information that is reported for each node.

Table 37. Statistic collection for nodes

Statistic name	Description
cluster_id	Indicates the name of the cluster.
cluster	Indicates the name of the cluster.
cpu	**busy** - Indicates the total CPU average core busy milliseconds since the node was reset. This statistic reports the amount of the time the processor has spent polling while waiting for work versus actually doing work. This statistic accumulates from zero.
	comp - Indicates the total CPU average core busy milliseconds for compression process cores since the node was reset.
	system - Indicates the total CPU average core busy milliseconds since the node was reset. This statistic reports the amount of the time the processor has spent polling while waiting for work versus actually doing work. This statistic accumulates from zero. This is the same information as the information provided with the cpu busy statistic and will eventually replace the cpu busy statistic.

Table 37. Statistic collection for nodes (continued)

cpu_core	**id** - Indicates the CPU core id.
	comp - Indicates the per-core CPU average core busy milliseconds for compression process cores since node was reset.
	system - Indicates the per-core CPU average core busy milliseconds for system process cores since node was reset.
id	Indicates the name of the node.
node_id	Indicates the unique identifier for the node.
rb	Indicates the number of bytes received.
re	Indicates the accumulated receive latency, excluding inbound queue time. This statistic is the latency that is experienced by the node communication layer from the time that an I/O is queued to cache until the time that the cache gives completion for it.
ro	Indicates the number of messages or bulk data received.
rq	Indicates the accumulated receive latency, including inbound queue time. This statistic is the latency from the time that a command arrives at the node communication layer to the time that the cache completes the command.
wb	Indicates the bytes sent.
we	Indicates the accumulated send latency, excluding outbound queue time. This statistic is the time from when the node communication layer issues a message out onto the Fibre Channel until the node communication layer receives notification that the message arrived.
wo	Indicates the number of messages or bulk data sent.
wq	Indicates the accumulated send latency, including outbound queue time. This statistic includes the entire time that data is sent. This time includes the time from when the node communication layer receives a message and waits for resources, the time to send the message to the remote node, and the time taken for the remote node to respond.

Table 38 describes the statistics collection for volumes.

Table 38. Cache statistics collection for volumes and volume copies

Statistic	Acronym	Statistics for volume cache	Statistics for volume copy cache	Statistics for volume cache partition	Statistics for volume copy cache partition	Statistics for the Node Overall Cache	Cache statistics for mdisks	Units and state
read ios	ri	Yes	Yes					ios, cumulative
write ios	wi	Yes	Yes					ios, cumulative
read misses	r	Yes	Yes					sectors, cumulative
read hits	rh	Yes	Yes					sectors, cumulative
flush_through writes	ft	Yes	Yes					sectors, cumulative
fast_write writes	fw	Yes	Yes					sectors, cumulative
write_through writes	wt	Yes	Yes					sectors, cumulative
write hits	wh	Yes	Yes					sectors, cumulative

Table 38. Cache statistics collection for volumes and volume copies (continued)

Statistic	Acronym	Statistics for volume cache	Statistics for volume copy cache	Statistics for volume cache partition	Statistics for volume copy cache partition	Statistics for the Node Overall Cache	Cache statistics for mdisks	Units and state
prefetches	p		Yes					sectors, cumulative
prefetch hits (prefetch data that is read)	ph		Yes					sectors, cumulative
prefetch misses (prefetch pages that are discarded without any sectors read)	pm		Yes					pages, cumulative
modified data	m	Yes	Yes					sectors, snapshot, non-cumulative
read and write cache data	v	Yes	Yes					sectors snapshot, non-cumulative
destages	d	Yes	Yes					sectors, cumulative
fullness Average	fav			Yes	Yes			%, non-cumulative
fullness Max	fmx			Yes	Yes			%, non-cumulative
fullness Min	fmn			Yes	Yes			%, non-cumulative
Destage Target Average	dtav				Yes		Yes	IOs capped 9999, non-cumulative
Destage Target Max	dtmx				Yes			IOs, non-cumulative
Destage Target Min	dtmn				Yes			IOs, non-cumulative
Destage In Flight Average	dfav				Yes		Yes	IOs capped 9999, non-cumulative
Destage In Flight Max	dfmx				Yes			IOs, non-cumulative
Destage In Flight Min	dfmn				Yes			IOs, non-cumulative
destage latency average	dav	Yes	Yes	Yes	Yes	Yes	Yes	µs capped 9999999, non-cumulative

Table 38. Cache statistics collection for volumes and volume copies (continued)

Statistic	Acronym	Statistics for volume cache	Statistics for volume copy cache	Statistics for volume cache partition	Statistics for volume copy cache partition	Statistics for the Node Overall Cache	Cache statistics for mdisks	Units and state
destage latency max	dmx			Yes	Yes	Yes		µs capped 9999999, non-cumulative
destage latency min	dmn			Yes	Yes	Yes		µs capped 9999999, non-cumulative
destage count	dcn	Yes	Yes	Yes	Yes	Yes		ios, non-cumulative
stage latency average	sav	Yes	Yes			Yes		µs capped 9999999, non-cumulative
stage latency max	smx					Yes		µs capped 9999999, non-cumulative
stage latency min	smn					Yes		µs capped 9999999, non-cumulative
stage count	scn	Yes	Yes			Yes		ios, non-cumulative
prestage latency average	pav		Yes			Yes		µs capped 9999999, non-cumulative
prestage latency max	pmx					Yes		µs capped 9999999, non-cumulative
prestage latency min	pmn					Yes		µs capped 9999999, non-cumulative
prestage count	pcn		Yes			Yes		ios, non-cumulative
Write Cache Fullness Average	wfav					Yes		%, non-cumulative
Write Cache Fullness Max	wfmx					Yes		%, non-cumulative
Write Cache Fullness Min	wfmn					Yes		%, non-cumulative
Read Cache Fullness Average	rfav					Yes		%, non-cumulative

Table 38. Cache statistics collection for volumes and volume copies (continued)

Statistic	Acronym	Statistics for volume cache	Statistics for volume copy cache	Statistics for volume cache partition	Statistics for volume copy cache partition	Statistics for the Node Overall Cache	Cache statistics for mdisks	Units and state
Read Cache Fullness Max	rfmx					Yes		%, non-cumulative
Read Cache Fullness Min	rfmn					Yes		%, non-cumulative
Pinned Percent	pp	Yes	Yes	Yes	Yes	Yes		% of total cache snapshot, non-cumulative
data transfer latency average	tav	Yes	Yes					µs capped 9999999, non-cumulative
Track Lock Latency (Exclusive) Average	teav	Yes	Yes					µs capped 9999999, non-cumulative
Track Lock Latency (Shared) Average	tsav	Yes	Yes					µs capped 9999999, non-cumulative
Cache I/O Control Block Queue Time	hpt					Yes		Average µs, non-cumulative
Cache Track Control Block Queue Time	ppt					Yes		Average µs, non-cumulative
Owner Remote Credit Queue Time	opt					Yes		Average µs, non-cumulative
Non-Owner Remote Credit Queue Time	npt					Yes		Average µs, non-cumulative
Admin Remote Credit Queue Time	apt					Yes		Average µs, non-cumulative
Cdcb Queue Time	cpt					Yes		Average µs, non-cumulative
Buffer Queue Time	bpt					Yes		Average µs, non-cumulative
Hardening Rights Queue Time	hrpt					Yes		Average µs, non-cumulative

Note: Any statistic with a name **av**, **mx**, **mn**, and **cn** is not cumulative. These statistics reset every statistics interval. For example, if the statistic does not have a name with name **av**, **mx**, **mn**, and **cn**, and it is an Ios or count, it will be a field containing a total number.

- The term *pages* means in units of 4096 bytes per page.

- The term *sectors* means in units of 512 bytes per sector.
- The term *µs* means microseconds.
- Non-cumulative means totals since the previous statistics collection interval.
- Snapshot means the value at the end of the statistics interval (rather than an average across the interval or a peak within the interval).

Table 39 describes the statistic collection for volume cache per individual nodes.

Table 39. Statistic collection for volume cache per individual nodes. This table describes the volume cache information that is reported for individual nodes.

Statistic name	Description
cm	Indicates the number of sectors of modified or dirty data that are held in the cache.
ctd	Indicates the total number of cache destages that were initiated writes, submitted to other components as a result of a volume cache flush or destage operation.
ctds	Indicates the total number of sectors that are written for cache-initiated track writes.
ctp	Indicates the number of track stages that are initiated by the cache that are prestage reads.
ctps	Indicates the total number of staged sectors that are initiated by the cache.
ctrh	Indicates the number of total track read-cache hits on prestage or non-prestage data. For example, a single read that spans two tracks where only one of the tracks obtained a total cache hit, is counted as one track read-cache hit.
ctrhp	Indicates the number of track reads received from other components, treated as cache hits on any prestaged data. For example, if a single read spans two tracks where only one of the tracks obtained a total cache hit on prestaged data, it is counted as one track read for the prestaged data. A cache hit that obtains a partial hit on prestage and non-prestage data still contributes to this value.
ctrhps	Indicates the total number of sectors that are read for reads received from other components that obtained cache hits on any prestaged data.
ctrhs	Indicates the total number of sectors that are read for reads received from other components that obtained total cache hits on prestage or non-prestage data.
ctr	Indicates the total number of track reads received. For example, if a single read spans two tracks, it is counted as two total track reads.
ctrs	Indicates the total number of sectors that are read for reads received.
ctwft	Indicates the number of track writes received from other components and processed in flush through write mode.
ctwfts	Indicates the total number of sectors that are written for writes that are received from other components and processed in flush through write mode.
ctwfw	Indicates the number of track writes received from other components and processed in fast-write mode.
ctwfwsh	Indicates the track writes in fast-write mode that were written in write-through mode because of the lack of memory.
ctwfwshs	Indicates the track writes in fast-write mode that were written in write through due to the lack of memory.

ctwfws	Indicates the total number of sectors that are written for writes that are received from other components and processed in fast-write mode.
ctwh	Indicates the number of track writes received from other components where every sector in the track obtained a write hit on already dirty data in the cache. For a write to count as a total cache hit, the entire track write data must already be marked in the write cache as dirty.
ctwhs	Indicates the total number of sectors that are received from other components where every sector in the track obtained a write hit on already dirty data in the cache.
ctw	Indicates the total number of track writes received. For example, if a single write spans two tracks, it is counted as two total track writes.
ctws	Indicates the total number of sectors that are written for writes that are received from components.
ctwwt	Indicates the number of track writes received from other components and processed in write through write mode.
ctwwts	Indicates the total number of sectors that are written for writes that are received from other components and processed in write through write mode.
cv	Indicates the number of sectors of read and write cache data that is held in the cache.

Table 40 describes the XML statistics specific to an IP Partnership port.

Table 40. XML statistics for an IP Partnership port

Statistic name	**Description**
ipbz	Indicates the average size (in bytes) of data that is being submitted to the IP partnership driver since the last statistics collection period.
iprc	Indicates the total bytes that are received before any decompression takes place.
ipre	Indicates the bytes retransmitted to other nodes in other clusters by the IP partnership driver.
iprt	Indicates the average round-trip time in microseconds for the IP partnership link since the last statistics collection period.
iprx	Indicates the bytes received from other nodes in other clusters by the IP partnership driver.
ipsz	Indicates the average size (in bytes) of data that is being transmitted by the IP partnership driver since the last statistics collection period.
iptc	Indicates the total bytes that are transmitted after any compression (if active) has taken place.
iptx	Indicates the bytes transmitted to other nodes in other clusters by the IP partnership driver.

Table 41 on page 49 describes the offload data transfer (ODX) Vdisk and node level I/O statistics.

Table 41. ODX VDisk and node level statistics

Statistic name	Acronym	Description
Read cumulative ODX I/O latency	orl	Cumulative total read latency of ODX I/O per VDisk. The unit type is micro-seconds (US).
Write cumulative ODX I/O latency	owl	Cumulative total write latency of ODX I/O per VDisk. The unit type is micro-seconds (US).
Total transferred ODX I/O read blocks	oro	Cumulative total number of blocks read and successfully reported to the host, by ODX WUT command per VDisk. It is represented in blocks unit type.
Total transferred ODX I/O write blocks	owo	Cumulative total number of blocks written and successfully reported to the host, by ODX WUT command per VDisk. It is represented in blocks unit type.
Wasted ODX I/Os	oiowp	Cumulative total number of wasted blocks written by ODX WUT command per node. It is represented in blocks unit type.
WUT failure count	otrec	Cumulative total number of failed ODX WUT commands per node. It includes WUT failures due to a token getting revoked and expired.

Table 42 describes the statistics collection for cloud per cloud account id.

Table 42. Statistics collection for cloud per cloud account id

Statistic name	Acronym	Description
id	id	Cloud account id
Total Successful Puts	puts	Total number of successful PUT operations
Total Successful Gets	gets	Total number of successful GET operations
Bytes Up	bup	Total number of bytes successful transferred to the cloud
Bytes Down	bdown	Total number of bytes successful downloaded/read from the cloud
Up Latency	uplt	Total time taken to transfer the data to the cloud

Statistic name	Acronym	Description
Down Latency	dwlt	Total time taken to download the data from the cloud
Down Error Latency	dwerlt	Time taken for the GET errors
Part Error Latency	pterlt	Total time taken for part errors
Persisted Bytes Down	prbdw	Total number of bytes successfully downloaded from the cloud and persisted on the local storage that were part of successful GET operation
Persisted Bytes Up	prbup	Total number of bytes successfully transferred to the cloud and persisted on the cloud that were part of successful PUT operation. The difference is that we may have a 100 bytes file, of which we successfully had 80 bytes sent to the cloud through a PUT operation, but the last data transfer cycle carrying 20 bytes errored out, and the entire request failed. In that case the statistics indicates: BYTES_UP = 80 and PERSISTED_BYTES_UP = 0
Persisted Down Latency	prdwlt	Total time taken to download the data from the cloud that were part of successful GET operation
Persisted Up Latency	pruplt	Total time taken to transfer the data to the cloud that were part of successful PUT operation
Failed Gets	flgt	Total number of failed GET operations
Failed Puts	flpt	Total number of failed PUT operations
Get Errors	gter	Total number of times a read from the cloud failed (including the last retry that failed the GET request)
Get Retries	gtrt	Total number of GET retries
Part Errors	pter	Total number of part errors. It is the count in case of multi part upload. The *part* refers to the multi-part upload scenario.

Table 42. Statistics collection for cloud per cloud account id (continued)

Statistic name	Acronym	Description
Parts Put	ptpt	Total number of parts that are successfully transferred to the cloud
Persisted parts	prpt	Total number parts successfully persisted on the cloud that were part of successful put operation
Put retries	ptrt	Total number of PUT retries
Throttle upload latency	tuplt	Average delay introduced due to setting upload bandwidth limit
Throttle download latency	tdwlt	Average delay introduced due to setting download bandwidth limit
Throttle upload bandwidth utilization percentage	tupbwpc	Bandwidth utilization in percentage of configured upload bandwidth limit
Throttle download bandwidth utilization percentage	tdwbwpc	Bandwidth utilization in percentage of configured download bandwidth limit

Table 43 describes the statistics collection for cloud per VDisk.

Table 43. Statistics collection for cloud per VDisk

SNo	Statistic name	Acronym	Description
1	blocks up	bup	Number of blocks uploaded in cloud
2	blocks down	bdn	Number of block downloaded from cloud

Note: A block is 512 bytes.

Actions

The following actions are available to the user:

OK   Click this button to change statistic collection.

Cancel
        Click this button to exit the panel without changing statistic collection.

XML formatting information

The XML is more complicated now, as seen in this raw XML from the volume (Nv_statistics) statistics. Notice how the names are similar but because they are in a different section of the XML, they refer to a different part of the VDisk.

```
<vdsk idx="0"
ctrs="213694394" ctps="0" ctrhs="2416029" ctrhps="0"
ctds="152474234" ctwfts="9635" ctwwts="0" ctwfws="152468611"
ctwhs="9117" ctws="152478246" ctr="1628296" ctw="3241448"
```

```
                 ctp="0" ctrh="123056" ctrhp="0" ctd="1172772"
                 ctwft="200" ctwwt="0" ctwfw="3241248" ctwfwsh="0"
                 ctwfwshs="0" ctwh="538" cm="13768758912876544" cv="13874234719731712"
                 gwot="0" gwo="0" gws="0" gwl="0"

                 id="Master_iogrp0_1"
                 ro="0" wo="0" rb="0" wb="0"
                 rl="0" wl="0" rlw="0" wlw="0" xl="0">
                 Vdisk/Volume statistics
                 <ca r="0" rh="0" d="0" ft="0"
                 wt="0" fw="0" wh="0" ri="0"
                 wi="0" dav="0" dcn="0" pav="0" pcn="0" teav="0"  tsav="0"  tav="0"
                 pp="0"/>

                 <cpy idx="0">

                 volume copy statistics
                 <ca r="0" p="0" rh="0" ph="0"
                 d="0" ft="0" wt="0" fw="0"
                 wh="0" pm="0" ri="0" wi="0"
                 dav="0" dcn="0" sav="0" scn="0"
                 pav="0" pcn="0" teav="0"  tsav="0"
                 tav="0"  pp="0"/>

                 </cpy>
                 <vdsk>
```

The `<cpy idx="0">` means its in the volume copy section of the VDisk, whereas the
statistics shown under `Vdisk/Volume statistics` are outside of the cpy idx section
and therefore refer to a VDisk/volume.

Similarly for the volume cache statistics for node and partitions:

```
                 <uca><ca dav="18726" dcn="1502531" dmx="749846" dmn="89"
                 sav="20868" scn="2833391" smx="980941" smn="3"
                 pav="0" pcn="0" pmx="0" pmn="0"
                 wfav="0" wfmx="2" wfmn="0"
                 rfav="0" rfmx="1" rfmn="0"
                 pp="0"
                 hpt="0" ppt="0" opt="0" npt="0"
                 apt="0" cpt="0" bpt="0" hrpt="0"
                 /><partition id="0"><ca dav="18726" dcn="1502531" dmx="749846" dmn="89"
                 fav="0" fmx="2" fmn="0"
                 dfav="0" dfmx="0" dfmn="0"
                 dtav="0" dtmx="0" dtmn="0"
                 pp="0"/></partition>
```

This output describes the volume cache node statistics where `<partition id="0">`
the statistics are described for partition 0.

Replacing `<uca>` with `<lca>` means that the statistics are for volume copy cache
partition 0.

Record the access information

It is important that anyone who has responsibility for managing the system know
how to connect to and log on to the system. Give attention to those times when the
normal system administrators are not available because of vacation or illness.

Record the following information in Table 44 on page 53 and ensure that
authorized people know how to access the information.
• The management IP addresses. This address connects to the system using the
  management GUI or starts a session that runs the command-line interface (CLI)
  commands. The system has two Ethernet ports. Each port can have either an

IPv4 address or an IPv6 address or both. Record this address and any limitations regarding where it can be accessed from within your Ethernet network.

- The service IP addresses for the control enclosure canister. These addresses are normally not needed. You might need a service IP address to access the service assistant during some recovery procedures. Use this address if the control enclosure CLI is not working. These addresses are not set during the installation of a Storwize V7000 system, but you can set these IP addresses later by using the management GUI or the `chserviceip` CLI command.

- The service IP address of the node canisters on the control enclosures is used only in certain circumstances. The service IP address connects to a node canister in the control enclosure. Access to the address is sometimes required if the canister has a fault that stops it from becoming an active member of the system. Each of the two node canisters can have a service IP address that is specified for Ethernet port 1. Each address can have either an IPv4 address or an IPv6 address or both. Ensure that the address specified for each node canister is different.

- The system password for user `superuser`. The password is required to access the system through the service IP address. The authentication of superuser is always local; therefore, the user ID can be used when a remote authentication server that is used for other users is not available.

Table 44. Access information for your system

Item	Value	Notes
The management IP address for the management GUI and CLI		
The management user ID (the default is `admin`)		
The management user ID password (the default is `admin001`)		
The additional management user IDs and passwords that you create on your system		
The control enclosure superuser IP address (not applicable to Storwize V7000 Gen2)		
Control enclosure service IP address: node canister 1		
Control enclosure service IP address: node canister 2		
The control enclosure superuser password (the default is `passw0rd`)		

Follow proper power management procedures

Follow the proper power management procedures for your system. Power management procedures differ, depending on the generation of your control enclosure model.

About this task

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 45. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 46. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Follow proper Storwize V7000 Gen2 power management procedures

You might lose access to volume data if you incorrectly power off all or part of a Storwize V7000 Gen2 system.

Always use the management GUI function to power off the system.

Only power off or remove a node canister if instructed to do so in a service action.

If you power off an expansion enclosure, you cannot read or write to the drives in that enclosure or to any other expansion enclosure that is attached to it from the SAS ports. Powering off an expansion enclosure can also prevent the control enclosure from flushing all the data that it cached to the RAID arrays.

Follow proper Storwize V7000 Gen1 power management procedures

Access to your volume data can be lost if you incorrectly power off all or part of a system.

Use the management GUI or the CLI commands to power off a system. Using either of these methods ensures that the data that is cached in the node canister memory is correctly flushed to the RAID arrays.

Do not power off an enclosure unless instructed to do so. If you power off an expansion enclosure, you cannot read or write to the drives in that enclosure or to any other expansion enclosure that is attached to it from the SAS ports. Powering off an expansion enclosure can prevent the control enclosure from flushing all the data that it has cached to the RAID arrays.

Remove a node canister only when directed to do so by a service action. Physically removing an active node canister means that it is unable to write any configuration data or volume data that it has cached to its internal disk and the data is lost. If both node canisters in a control enclosure are removed in quick succession, run recovery actions, which might include restoring your volume data from a backup.

Set up event notifications

Configure your system to send notifications when a new event is reported.

Correct any issues reported by your system as soon as possible. To avoid monitoring for new events by constantly monitoring the management GUI, configure your system to send notifications when a new event is reported. Select the type of event that you want to be notified about. For example, restrict notifications to just events that require immediate action. Several event notification mechanisms exist:

- Email. An event notification can be sent to one or more email addresses. This mechanism notifies individuals of problems. Individuals can receive notifications wherever they have email access which includes mobile devices.
- Simple Network Management Protocol (SNMP). An SNMP trap report can be sent to a data-center management system, such as IBM Systems Director, that consolidates SNMP reports from multiple systems. Using this mechanism, you can monitor your data center from a single workstation.
- Syslog. A syslog report can be sent to a data-center management system that consolidates syslog reports from multiple systems. Using this mechanism, you can monitor your data center from a single workstation.
- Call Home. If your system is within warranty, or you have a hardware maintenance agreement, configure your system to send email events to IBM if an issue that requires hardware replacement is detected. This mechanism is called Call Home. When the event is received, IBM automatically opens a problem report, and if appropriate, contacts you to verify if replacement parts are required. If you set up Call Home to IBM, ensure that the contact details that you configure are correct and kept up to date as personnel change.

Set up inventory reporting

Inventory reporting is an extension to the Call Home email.

Rather than reporting a problem, an email is sent to IBM that describes your system hardware and critical configuration information. Object names and other information, such as IP addresses, are not sent. The inventory email is sent on a regular basis. Based on the information that is received, IBM can inform you if the hardware or software that you are using requires an update because of a known issue.

Back up your data

Back up your system configuration data and volume data.

The storage system backs up your control enclosure configuration data to a file every day. This data is replicated on each control node canister in the system. Download this file regularly to your management workstation to protect the data. This file must be used if there is a serious failure that requires you to restore your system configuration. It is important to back up this file after modifying your system configuration.

Your volume data is susceptible to failures in your host application or your Storwize V7000 system. Follow a backup and archive policy that is appropriate to the data that you have for storing the volume data on a different system.

Manage your spare and failed drives

Your RAID arrays that are created from drives consist of drives that are active members and drives that are spares.

The spare drives are used automatically if a member drive fails. If you have sufficient spare drives, you do not have to replace them immediately when they fail. However, monitoring the number, size, and technology of your spare drives, ensures that you have sufficient drives for your requirements. Ensure that there are sufficient spare drives available so that your RAID arrays are always online.

Resolve alerts in a timely manner

Your system reports an alert when there is an issue or a potential issue that requires user attention. The Storwize V7000 helps resolve these problems through the **Recommended actions only** option from the Events panel.

Complete the recommended actions as quickly as possible after the problem is reported. Your system is designed to be resilient to most single hardware failures. However, if you operate for any period of time with a hardware failure, the possibility increases that a second hardware failure can result in some volume data that is unavailable.

If there are a number of unfixed alerts, fixing any one alert might become more difficult because of the effects of the other alerts.

Keep your software up to date

Check for new code releases and update your code on a regular basis.

This can be done using the management GUI, or by checking the IBM support website to see if new code releases are available:

www.ibm.com/support

The release notes provide information about new function in a release plus any issues that have been resolved. Update your code regularly if the release notes indicate a potential issue.

Keep your records up to date

Follow the proper record keeping procedures for your system. Record management procedures differ, depending on the generation of your control enclosure model.

About this task

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 47. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 48. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Keep your Storwize V7000 Gen2 records up to date

Keep good records of the system location, names, and management addresses. Record the location information for your enclosures.

If you have only one system, it is relatively easy to identify the enclosures that make up the system. Identification becomes more difficult when you have multiple systems in your data center and multiple systems in the same rack.

Record the MT-M and serial number of each Storwize V7000 Gen2 enclosure. The information can be found on the IBM Standard Asset Tag attached to the left enclosure bezel, which includes a machine-readable data matrix to ISO/IEC 15434.

For each system, record the location of the control enclosure and the location of any expansion enclosures. It is useful to label the enclosures themselves with the system name and the management IP addresses.

Keep your Storwize V7000 Gen1 records up to date

Record the location information for your enclosures.

If you have only one system, it is relatively easy to identify the enclosures that make up the system. Identification becomes more difficult when you have multiple systems in your data center and multiple systems in the same rack.

The enclosure identifier that is displayed on the front of the display is unique within a system. However, the identifiers can be repeated between different systems. Do not rely solely on this identifier.

For each system, record the location of the control enclosure and the location of any expansion enclosures. It is useful to label the enclosures themselves with the system name and the management IP addresses.

Subscribe to support notifications

Subscribe to support notifications so that you are aware of best practices and issues that might affect your system.

Subscribe to support notifications by visiting the IBM support page on the IBM website:

www.ibm.com/support

By subscribing, you are informed of new and updated support site information, such as publications, hints and tips, technical notes, product flashes (alerts), and downloads.

Know your IBM warranty and maintenance agreement details

If you have a warranty or maintenance agreement with IBM, know the details that must be supplied when you call for support.

Have the phone number of the support center available. When you call support, provide the machine type (always 2076) and the serial number of the enclosure that has the problem. If the problem does not relate to a specific enclosure, provide the control enclosure serial number. The serial numbers are on the labels on the enclosures.

Support personnel also ask for your customer number, machine location, contact details, and the details of the problem.

How to get information, help, and technical assistance

If you need help, service, technical assistance, or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you.

Information

IBM maintains pages on the web where you can get information about IBM products and fee services, product implementation and usage assistance, break and fix service support, and the latest technical information. For more information, refer to Table 4 on page xxvi.

Table 49. IBM websites for help, services, and information

Website	Address
Directory of worldwide contacts	http://www.ibm.com/planetwide
Support for Storwize V7000 (2076)	www.ibm.com/support
Support for IBM System Storage and IBM TotalStorage products	www.ibm.com/support/

Note: Available services, telephone numbers, and web links are subject to change without notice.

Help and service

Before calling for support, be sure to have your IBM Customer Number available. If you are in the US or Canada, you can call 1 (800) IBM SERV for help and service. From other parts of the world, see http://www.ibm.com/planetwide for the number that you can call.

When calling from the US or Canada, choose the **storage** option. The agent decides where to route your call, to either storage software or storage hardware, depending on the nature of your problem.

If you call from somewhere other than the US or Canada, you must choose the **software** or**hardware** option when calling for assistance. Choose the **software** option if you are uncertain if the problem involves the Storwize V7000 software or hardware. Choose the **hardware** option only if you are certain the problem solely involves the Storwize V7000 hardware.When calling IBM for service regarding the product, follow these guidelines for the **software** and **hardware** options:

Software option
> Identify the Storwize V7000 product as your product and supply your customer number as proof of purchase. The customer number is a 7-digit number (0000000 - 9999999) assigned by IBM when the product is purchased. Your customer number should be on the customer information worksheet or on the invoice from your storage purchase. If asked for an operating system, use **Storage**.

Hardware option
> Provide the serial number and appropriate 4-digit machine type. For Storwize V7000, the machine type is 2076.

In the US and Canada, hardware service and support can be extended to 24x7 on the same day. The base warranty is 9x5 on the next business day.

Getting help online

You can find information about products, solutions, partners, and support on the IBM website.

To find up-to-date information about products, services, and partners, visit the IBM website at www.ibm.com/support.

Before you call

Make sure that you have taken steps to try to solve the problem yourself before you call.

Some suggestions for resolving the problem before calling IBM Support include:
- Check all cables to make sure that they are connected.
- Check all power switches to make sure that the system and optional devices are turned on.
- Use the troubleshooting information in your system documentation. The troubleshooting section of the knowledge center contains procedures to help you diagnose problems.
- Go to the IBM Support website at www.ibm.com/support to check for technical information, hints, tips, and new device drivers or to submit a request for information.

Using the documentation

Information about your IBM storage system is available in the documentation that comes with the product.

That documentation includes printed documents, online documents, readme files, and help files in addition to the knowledge center. See the troubleshooting information for diagnostic instructions. The troubleshooting procedure might require you to download updated device drivers or software. IBM maintains pages on the web where you can get the latest technical information and download device drivers and updates. To access these pages, go to www.ibm.com/support and follow the instructions. Also, some documents are available through the IBM Publications Center.

Sign up for the Support Line Offering

If you have questions about how to use and configure the machine, sign up for the IBM Support Line offering to get a professional answer.

The maintenance that is supplied with the system provides support when there is a problem with a hardware component or a fault in the system machine code. At times, you might need expert advice about using a function that is provided by the system or about how to configure the system. Purchasing the IBM Support Line offering gives you access to this professional advice while deploying your system, and in the future.

Contact your local IBM sales representative or your support group for availability and purchase information.

Chapter 3. Understanding battery operations

Understanding batter operations can extend the life of the battery.

Battery operation for the control enclosure

Storwize V7000 node canisters cache volume data and hold state information in volatile memory. Battery operations differ, depending on the generation of your control enclosure model.

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 50. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 51. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Battery operation for Storwize V7000 Gen2 control enclosures

Each node canister in the control enclosure caches critical data and holds state information in volatile memory.

If power to a node canister fails, the node canister uses battery power to write cache and state data to its boot drive.

Note: Storwize V7000 Gen2 expansion canisters do not cache volume data or store state information in volatile memory. Therefore, expansion canisters do not require battery power. If ac power to both power supplies in an expansion enclosure fails, the enclosure powers off. When ac power is restored to at least one power supply, the enclosure restarts without operator intervention.

The battery is maintained in a fully charged state by the battery subsystem. At maximum power, the battery can save critical data and state information in two back-to-back power failures. If power to a node canister is lost, saving critical data starts after a five-second wait. (If the outage is shorter than five seconds, the battery continues to support the node and critical data is not saved.) The node canister stops handling I/O requests from host applications. The saving of critical data runs to completion, even if power is restored during this time. The loss of power might be because the input power to the enclosure is lost, or because the node canister is removed from the enclosure.

When power is restored to the node canister, the system restarts without operator intervention. How quickly it restarts depends on whether there is a history of previous power failures. The system restarts only when the battery has sufficient charge for the node canister to save the cache and state data again. A node canister with multiple power failures might not have sufficient charge to save critical data. In such a case, the system starts in service state and waits to start I/O operations until the battery has sufficient charge.

Two light-emitting diode (LED) indicators indicate the state of the battery:
- Status LED - Green
- Fault LED - Amber

See "Procedure: Understanding the system status from the Storwize V7000 Gen2 LEDs" for a complete description of the battery LEDs.

Important: Although Storwize V7000 Gen2 is resilient to power failures and brown outs, always install Storwize V7000 Gen2 in an environment with reliable, consistent, and required ac power. Consider uninterruptible power supply units to avoid extended interruptions to data access.

Design parameters

Consider the following important design parameters:
- The design life of the battery in the Storwize V7000 Gen2 is five years service after one year on the shelf.
- Each battery is automatically reconditioned every three months to measure the battery capacity. Batteries in the same enclosure are not reconditioned within two days of each other. If a battery has a lower capacity than required (below the planned threshold), it is marked as "End Of Life" and should be replaced.
- Each battery provides power only for the canister in which it is installed. If a battery fails, the canister goes offline and reports a node error. The single running canister destages its cache and runs the I/O group in "write-through" mode until its partner canister is repaired and online.

Reconditioning the Storwize V7000 Gen2 battery
Reconditioning the battery ensures that the system can accurately determine the charge in the battery.

As a battery ages, it loses capacity. When a battery no longer has capacity to protect against two power loss events, it reports the battery end of life event and it should be replaced.

A reconditioning cycle is automatically scheduled to occur approximately once every three months, but reconditioning is rescheduled or canceled if the system loses redundancy. In addition, a two-day delay is imposed between the recondition cycles of the two batteries in one enclosure.

Battery operation for Storwize V7000 Gen1 control enclosures

Storwize V7000 node canisters cache volume data and hold state information in volatile memory.

If the power fails, cache and state data is written to a local flash drive in the canister. The batteries within the control enclosure provide the power to write the cache and state data to a local drive.

Note: Storwize V7000 expansion canisters do not cache volume data or store state information in volatile memory. Therefore, expansion canisters do not require battery power. If ac power to both power supplies in an expansion enclosure fails, the enclosure powers off. When ac power is restored to at least one power supply, the expansion enclosure restarts without operator intervention.

Two power supply units exist in the control enclosure. Each one contains an integrated battery. Both power supply units and batteries provide power to both control canisters. Each battery has sufficient charge to save critical data from both node canisters to the local drive. A fully redundant system has two batteries and two canisters. In such a system, the batteries can save critical data and state information from both canisters to a local drive twice. In a system with a failed battery, there is enough charge in the remaining battery to support saving critical data from both canisters to a local drive once.

Without ac power, a canister starts saving critical data to a local drive approximately 10 seconds after it detects the loss. If power is restored within the 10 seconds, the system continues to operate. This loss in power is called a *brown out*. When it saves critical data, the system stops handling I/O requests from host applications. At the same time, Metro Mirror and Global Mirror relationships go offline. The system powers off after it saves the critical data.

If both node canisters shut down without writing the cache and state data to the local drive, the system is unable to restart without an extended service action. The system configuration must be restored. If any cache write data is lost, volumes must be restored from a backup. Therefore, it is important not to remove the canisters or the power supply units from the control enclosures unless directed to do so by the service procedures. Removing either of these components might prevent the node canister from writing its cache and state data to the local drive.

When the ac power is restored to the control enclosure, the system restarts without operator intervention. How quickly it restarts depends on whether there is a history of previous power failures.

When ac power is restored after both canisters save critical data, the system restarts without operator intervention. How quickly it restarts depends on whether there is a history of previous power failures. The system restarts only when the batteries have sufficient charge to save critical data for both node canisters again. A

fully redundant system can restart as soon as power is restored after it saves critical data once. If a second ac power outage occurs before batteries charge completely, the system starts in service state. The system does not begin I/O operations until the batteries are half charged. Recharging takes approximately 30 minutes.

In a system with a failed battery, an ac power failure causes both canisters to save critical data and completely discharges the remaining battery. When the ac power is restored, the system starts in service state and does not start I/O operations until the remaining battery is fully charged. The recharging takes approximately 1 hour.

A battery is considered failed for the following conditions:
- The system can communicate with the battery, but the battery reports an error.
- The system is unable to communicate with the battery. The failure exists because the power supply is absent or because the power supply failed without communicating with the system.

Other conditions can cause critical data to be saved and the nodes to go into service state and stop I/O operations. Each node canister saves critical data if it detects there is no longer sufficient battery charge to support saving critical data. This situation happens, for example, when both batteries have two-thirds of a charge. The total charge is sufficient to support saving critical data once. Therefore, both canisters are in active state and I/O operations occur. If one battery fails, the remaining battery has only two-thirds of a charge. The total charge in the enclosure is now insufficient to save critical data if the ac power fails. Data protection is not guaranteed in this case. When the battery has sufficient charge, the system automatically restarts.

The node canister has two LEDs indicators that indicate the state of the battery:
- Status
- Fault

See "Procedure: Understanding the system status by using the LEDs" for a complete description of the battery LEDs.

Important: Although Storwize V7000 is resilient to power failures and brown outs, always install in an environment with reliable, consistent, and required AC power. Consider uninterruptible power supply units to avoid extended interruptions to data access.

Design parameters

Consider the following important design parameters:
- The design life of the battery in the Storwize V7000 is five years service after one year on the shelf.
- No periodic *learning mode* or reconditioning cycle occurs in the battery of this product.
- Each battery provides power only for the canister in which it is installed. If a battery fails, the canister goes offline and reports a node error. The single running canister destages its cache and runs the I/O group in "write-through" mode until its partner is repaired and online.

Maintenance discharge cycles

Maintenance discharge cycles extend the lifetime of the batteries and ensure that the system can accurately measure the charge in the batteries. Discharge cycles guarantee that the batteries have sufficient charge to protect the Storwize V7000 system.

Maintenance discharge cycles are scheduled automatically by the system and involve fully discharging a battery and then recharging it again. Maintenance discharges are normally scheduled only when the system has two fully charged batteries. This condition ensures that for the duration of the maintenance cycle, the system still has sufficient charge to complete a save of the critical data if the ac power fails. This condition also ensures that I/O operations continue while the maintenance cycle is completed. It is usual for both batteries to require a maintenance discharge at the same time. In these circumstances, the system automatically schedules the maintenance of one battery. When the maintenance on that battery completes, the maintenance on the other battery starts.

Maintenance discharges are scheduled for the following situations:
- A battery has been powered on for three months without a maintenance discharge.
- A battery has provided protection for saving critical data at least twice.
- A battery has provided protection for at least 10 brown outs, which lasted up to 10 seconds each.

A maintenance discharge takes approximately 10 hours to complete. If the ac power outage occurs during the maintenance cycle, the cycle must be restarted. The cycle is scheduled automatically when the battery is fully charged.

Under the following conditions, a battery is not considered when calculating whether there is sufficient charge to protect the system. This condition persists until a maintenance discharge cycle is completed.
- A battery is completing a maintenance discharge.
- A battery has provided protection for saving critical data at least four times without any intervening maintenance discharge.
- A battery has provided protection for at least 20 brown outs, which lasted up to 10 seconds each.
- A battery must restart a maintenance discharge because the previous maintenance cycle was disrupted by an ac power outage.

If a system suffers repeated ac power failures without a sufficient time interval in between the ac failures to complete battery conditioning, then neither battery is considered when calculating whether there is sufficient charge to protect the system. In these circumstances, the system enters service state and does not permit I/O operations to be restarted until the batteries have charged and one of the batteries has completed a maintenance discharge. This activity takes approximately 10 hours.

If one of the batteries in a system fails and is not replaced, it prevents the other battery from completing a maintenance discharge. Not only does this condition reduce the lifetime of the remaining battery, but it also prevents a maintenance discharge cycle from occurring after the battery has provided protection for at least 2 critical saves or 10 brown outs. Preventing this maintenance cycle from occurring increases the risk that the system accumulates a sufficient number of power outages to cause the remaining battery to be discounted when calculating whether

there is sufficient charge to protect the system. This condition results in the system entering service state while the one remaining battery completes a maintenance discharge. I/O operations are not permitted during this process. This activity takes approximately 10 hours.

Chapter 4. Understanding the medium errors and bad blocks

A storage system returns a medium error response to a host when it is unable to successfully read a block. The Storwize V7000 response to a host read follows this behavior.

The volume virtualization that is provided extends the time when a medium error is returned to a host. Because of this difference to non-virtualized systems, the Storwize V7000 system uses the term *bad blocks* rather than medium errors.

The system allocates volumes from the extents that are on the managed disks (MDisks). The MDisk can be a volume on an external storage controller or a RAID array that is created from internal drives. In either case, depending on the RAID level used, there is normally protection against a read error on a single drive. However, it is still possible to get a medium error on a read request if multiple drives have errors or if the drives are rebuilding or are offline due to other issues.

The system provides migration facilities to move a volume from one underlying set of physical storage to another or to replicate a volume that uses FlashCopy or Metro Mirror or Global Mirror. In all these cases, the migrated volume or the replicated volume returns a medium error to the host when the logical block address on the original volume is read. The system maintains tables of bad blocks to record where the logical block addresses that cannot be read are. These tables are associated with the MDisks that are providing storage for the volumes.

The **dumpmdiskbadblocks** command and the **dumpallmdiskbadblocks** command are available to query the location of bad blocks.

Important: The **dumpmdiskbadblocks** only outputs the virtual medium errors that have been created, and not a list of the actual medium errors on MDisks or drives.

It is possible that the tables that are used to record bad block locations can fill up. The table can fill either on an MDisk or on the system as a whole. If a table does fill up, the migration or replication that was creating the bad block fails because it was not possible to create an exact image of the source volume.

The system creates alerts in the event log for the following situations:
- When it detects medium errors and creates a bad block
- When the bad block tables fill up

Table 52 lists the bad block error codes.

Table 52. Bad block errors

Error code	Description
1840	The managed disk has bad blocks. On an external controller, this can only be a copied medium error.
1226	The system has failed to create a bad block because the MDisk already has the maximum number of allowed bad blocks.

Table 52. Bad block errors (continued)

Error code	Description
1225	The system has failed to create a bad block because the system already has the maximum number of allowed bad blocks.

The recommended actions for these alerts guide you in correcting the situation.

Clear bad blocks by deallocating the volume disk extent, by deleting the volume or by issuing write I/O to the block. It is good practice to correct bad blocks as soon as they are detected. This action prevents the bad block from being propagated when the volume is replicated or migrated. It is possible, however, for the bad block to be on part of the volume that is not used by the application. For example, it can be in part of a database that has not been initialized. These bad blocks are corrected when the application writes data to these areas. Before the correction happens, the bad block records continue to use up the available bad block space.

Chapter 5. Storwize V7000 user interfaces for servicing your system

The Storwize V7000 provides a number of user interfaces to troubleshoot, recover, or maintain your system. The interfaces provide various sets of facilities to help resolve situations that you might encounter.

The interfaces for servicing your system connect through the 1 Gbps Ethernet ports that are accessible from port 1 of each canister or by directly connecting to the technician port of a node canister. You cannot manage a system by using the 10 Gbps Ethernet ports.

- Use the initialization tool to do the initial setup of your system.
- Use the management GUI to monitor and maintain the configuration of storage that is associated with your clustered systems.
- Complete service procedures from the service assistant.
- Use the command-line interface (CLI) to manage your system. The front panel on the node provides an alternative service interface.

Management GUI interface

The management GUI is a browser-based GUI for configuring and managing all aspects of your system. It provides extensive facilities to help troubleshoot and correct problems.

About this task

You use the management GUI to manage and service your system. The **Monitoring > Events** panel provides access to problems that must be fixed and maintenance procedures that step you through the process of correcting the problem.

The information on the Events panel can be filtered three ways:

Recommended action (default)
Shows only the alerts that require attention and have an associated fix procedure. Alerts are listed in priority order and should be fixed sequentially by using the available fix procedures. For each problem that is selected, you can:

- Run a fix procedure.
- View the properties.

Unfixed messages and alerts
Displays only the alerts and messages that are not fixed. For each entry that is selected, you can:

- Run a fix procedure on any alert with an error code.
- Mark an event as fixed.
- Filter the entries to show them by specific minutes, hours, or dates.
- Reset the date filter.
- View the properties.

Show all

Displays all event types whether they are fixed or unfixed. For each entry that is selected, you can:

- Run a fix procedure on any alert with an error code.
- Mark an event as fixed.
- Filter the entries to show them by specific minutes, hours, or dates.
- Reset the date filter.
- View the properties.

Some events require a certain number of occurrences in 25 hours before they are displayed as unfixed. If they do not reach this threshold in 25 hours, they are flagged as expired. Monitoring events are below the coalesce threshold and are usually transient.

You can also sort events by time or error code. When you sort by error code, the most serious events, those with the lowest numbers, are displayed first. You can select any event that is listed and select **Actions** > **Properties** to view details about the event.

- Recommended Actions. For each problem that is selected, you can:
  - Run a fix procedure.
  - View the properties.
- Event log. For each entry that is selected, you can:
  - Run a fix procedure.
  - Mark an event as fixed.
  - Filter the entries to show them by specific minutes, hours, or dates.
  - Reset the date filter.
  - View the properties.

When to use the management GUI

The management GUI is the primary tool that is used to service your system.

Regularly monitor the status of the system using the management GUI. If you suspect a problem, use the management GUI first to diagnose and resolve the problem.

Use the views that are available in the management GUI to verify the status of the system, the hardware devices, the physical storage, and the available volumes. The **Monitoring** > **Events** panel provides access to all problems that exist on the system. Use the **Recommended Actions** filter to display the most important events that need to be resolved.

If there is a service error code for the alert, you can run a fix procedure that assists you in resolving the problem. These fix procedures analyze the system and provide more information about the problem. They suggest actions to take and step you through the actions that automatically manage the system where necessary. Finally, they check that the problem is resolved.

If there is an error that is reported, always use the fix procedures within the management GUI to resolve the problem. Always use the fix procedures for both system configuration problems and hardware failures. The fix procedures analyze the system to ensure that the required changes do not cause volumes to be

inaccessible to the hosts. The fix procedures automatically perform configuration changes that are required to return the system to its optimum state.

Accessing the management GUI

To view events, you must access the management GUI.

About this task

You must use a supported web browser. For a list of supported browsers, refer to the "Web browser requirements to access the management GUI" topic.

You can use the management GUI to manage your system as soon as you have created a clustered system.

Procedure

1. Start a supported web browser and point the browser to the management IP address of your system.

   The management IP address is set when the clustered system is created. Up to four addresses can be configured for your use. There are two addresses for IPv4 access and two addresses for IPv6 access. When the connection is successful, you will see a login panel.
2. Log on by using your user name and password.
3. When you have logged on, select **Monitoring** > **Events**.
4. Ensure that the events log is filtered using **Recommended actions**.
5. Select the recommended action and run the fix procedure.
6. Continue to work through the alerts in the order suggested, if possible.

Results

After all the alerts are fixed, check the status of your system to ensure that it is operating as intended.

If you encounter problems logging on the management GUI or connecting to the management GUI, see "Problem: Unable to log on to the management GUI" on page 89 or "Problem: Unable to connect to the Storwize V7000 Gen1 management GUI" on page 88.

Diagnosing and resolving problems with fix procedures

You can use fix procedures to diagnose and resolve problems with the system.

About this task

For example, to repair a Storwize V7000 system, you might complete the following tasks:
- Analyze the event log (if it is available, or view node errors)
- Replace failed components
- Verify the status of a repaired device
- Restore a device to an operational state in the system
- Mark the error as fixed in the event log

Fix procedures help simplify this process by automating as many of the tasks as possible.

The following example uses the management GUI to repair a Storwize V7000 system.

Procedure

Complete the following steps to start the fix procedure.

1. Click **Monitoring** > **Events** and ensure that you are filtering the event log to display **Recommended actions**.

   The list might indicate that many errors must be repaired. If the list contains several errors, the error at the top of the list has the highest priority and must always be fixed first. If you do not fix the higher priority errors first, you might not be able to fix the lower priority errors.

2. Select the error at the top of the list or select the **Next recommended action**.

3. Click **Run Fix Procedure**.

   The pane displays the error code and provides a description of the condition.

4. Click **Next** to go forward or **Cancel** to return to the previous pane. One or more panes might be displayed with instructions for you to replace parts or complete other repair activity.

5. If you are not able to complete the actions now, click **Cancel** until you return to the previous pane. Click **Cancel** until you are returned to the **Next recommended action** pane. When you return to the fix procedures, the repair can be restarted from step 1. After you complete all the instructions, click **OK**. When the last repair action is completed, the procedures might attempt to restore failed devices to the system.

6. After you complete the fix, you see the statement `Click OK to mark the error as fixed`. Click **OK**. This action marks the error as fixed in the event log and prevents this instance of the error from being listed again.

7. When you see the statement `The repair has been completed.`, click **Exit**. If other errors must be fixed, those errors are displayed and the fix procedures continue.

   If no errors remain, the following statement is displayed: `There are no unfixed errors in the event log.`

Results

While you are fixing hardware faults, the fix procedures might direct you to complete hardware actions that look like an error to the system (for example, replacing a drive). In these situations, the fix procedures enter maintenance mode automatically. New events are entered into the event log when they occur. However, a specific set of events are not notified unless they remain unfixed when exiting maintenance mode. The events that were recorded in maintenance mode are fixed automatically when the issue is resolved. Maintenance mode prevents unnecessary messages from being sent.

Service assistant interface

The service assistant interface is a browser-based GUI that is used to service individual node canisters in the control enclosures.

You connect to the service assistant on one node canister through the service IP address. If there is a working communications path between the node canisters, you can view status information and perform service tasks on the other node canister by making the other node canister the current node. You do not have to

reconnect to the other node. For Storwize V7000 Gen2 systems, you can also access the service assistant interface by using the technician port.

When to use the service assistant

The primary use of the service assistant is when a node canister in the control enclosure is in service state. The node canister cannot be active as part of a system while it is in service state.

Attention: Complete service actions on node canisters only when directed to do so by the fix procedures. If used inappropriately, the service actions that are available through the service assistant can cause loss of access to data or even data loss.

The node canister might be in a service state because it has a hardware issue, has corrupted data, or has lost its configuration data.

Use the service assistant in the following situations:
- When you cannot access the system from the management GUI and you cannot access the Storwize V7000 to run the recommended actions
- When the recommended action directs you to use the service assistant.

The storage system management GUI operates only when there is an online system. Use the service assistant if you are unable to create a system or if both node canisters in a control enclosure are in service state.

The service assistant does not provide any facilities to help you service expansion enclosures. Always service the expansion enclosures by using the management GUI.

The service assistant provides detailed status and error summaries, and the ability to modify the World Wide Node Name (WWNN) for each node.

You can also complete the following service-related actions:
- Collect logs to create and download a package of files to send to support personnel.
- Remove the data for the system from a node.
- Recover a system if it fails.
- Install a code package from the support site or rescue the code from another node.
- Update code on node canisters manually versus completing a standard update procedure.
- Configure a control enclosure chassis after replacement.
- Change the service IP address that is assigned to Ethernet port 1 for the current node canister.
- Install a temporary SSH key if a key is not installed and CLI access is required.
- Restart the services used by the system.

The service assistant completes a number of tasks that cause the node canister to restart. It is not possible to maintain the service assistant connection to the node canister when it restarts. If the current node canister on which the tasks are completed is also the node canister that the browser is connected to and you lose your connection, reconnect and log on to the service assistant again after running the tasks.

Accessing the service assistant

The service assistant is a web application that helps troubleshoot and resolve problems on a node canister in a control enclosure. The service assistant can also be accessed for the Storwize V7000 Gen2 by using the technician port.

About this task

You must use a supported web browser. For a list of supported browsers, refer to the topic Web browser requirements to access the management GUI.

Procedure

To start the application, complete the following steps.

1. Start a supported web browser and point your web browser to *serviceaddress*/service for the node canister that you want to work on.

   For example, if you set a service address of 11.22.33.44 for a node canister, point your browser to 11.22.33.44/service. If you are unable to connect to the service assistant, see "Problem: Cannot connect to the service assistant" on page 94.

2. Log on to the service assistant using the superuser password.

   If you are accessing a new node canister, the default password is passw0rd. If the node canister is a member of a system or was a member of a system, use the password for the superuser password.

   If you do not know the current superuser password, try to find out. If you cannot find out what the password is, reset the password. Go to "Procedure: Resetting the Storwize V7000 Gen1 superuser password" on page 100.

Results

Complete the service assistant actions on the correct node canister. If you did not connect to the node canister that you wanted to work on, access the **Change Node** panel from the home page to select a different current node.

Commands are run on the current node. The current node might not be the node canister that you connected to. The current node identification is shown on the left at the top of the service assistant screen. The identification includes the enclosure serial number, the slot location, and if it has one, the node name of the current node.

The Storwize V7000 Gen2 technician port

The technician port provides a convenient, direct connection to a node canister for servicing.

On uninitialized systems, the technician port provides access to the system initialization wizard instead of the service assistant. An uninitialized system is one where all node canisters have the green power LED on, the green status LED blinking, and amber fault LED is off.

Once a system has been initialized, the technician port provides access to:
- The service assistant
- The password reset facility (if enabled)

Command-line interface

Use the command-line interface (CLI) to manage a system with task commands and information commands.

For a full description of the commands and how to start an SSH command-line session, see the "Command-line interface" section of the Storwize V7000 Information Center.

When to use the CLI

The system command-line interface is intended for use by advanced users who are confident at using a CLI.

Nearly all of the flexibility that is offered by the CLI is available through the management GUI. However, the CLI does not provide the fix procedures that are available in the management GUI. Therefore, use the fix procedures in the management GUI to resolve the problems. Use the CLI when you require a configuration setting that is unavailable in the management GUI.

You might also find it useful to create command scripts that use CLI commands to monitor certain conditions or to automate configuration changes that you make regularly.

Accessing the system CLI

Follow the steps that are described in the Command-line interface section of the Storwize V7000 Knowledge Center to initialize and use a CLI session.

Service command-line interface

Use the service command-line interface (CLI) to manage a node canister in a control enclosure by using the task commands and information commands.

Note: The service command line interface can also be accessed by using the technician port.

For a full description of the commands and how to start an SSH command line session, see Command-line interface.

When to use the service CLI

The service CLI is intended for use by advanced users who are confident at using a command-line interface.

To access a node canister directly, it is normally easier to use the service assistant with its graphical interface and extensive help facilities.

Accessing the service CLI

To initialize and use a CLI session, complete the steps in the Command-line interface topic of this product information.

satask.txt commands

If you are creating the `satask.txt` command file by using a text editor, the file must contain a single command on a single line in the file.

The commands that you use are the same as the service CLI commands except where noted. Not all service CLI commands can be run from the USB flash drive. The **satask.txt** commands always run on the node that the USB flash drive is plugged into.

Reset service IP address and superuser password command:

Use this command to obtain service assistant access to a node canister even if the current state of the node canister is unknown. The physical access to the node canister is required and is used to authenticate the action.

Syntax

```
►►── satask ── chserviceip ── ──serviceip─ipv4─┬───────────┬──┬───────────┬──┬─────────────────┬──►◄
                                               └─gw─ipv4─┘    └─mask─ipv4─┘   └─-resetpassword─┘

►►── satask ── chserviceip ── ──serviceip_6─ipv6─┬──────────────┬──┬───────────────┬───────────►
                                                 └─gw_6─ipv6─┘      └─prefix_6─int─┘
►─┬───────────────────┬───────────────────────────────────────────────────────────────────────►◄
  └─-resetpassword─┘

►►── satask ── chserviceip ── ──default─┬───────────────────┬──────────────────────────────────►◄
                                        └─-resetpassword─┘
```

Parameters

-serviceip *ipv4*
    (Optional) The IPv4 address for the service assistant.

-gw *ipv4*
    (Optional) The IPv4 gateway for the service assistant.

-mask *ipv4*
    (Optional) The IPv4 subnet for the service assistant.

-serviceip_6 *ipv6*
    (Optional) The IPv6 address for the service assistant.

-gw_6 *ipv6*
    (Optional) The IPv6 gateway for the service assistant.

-default
    (Optional) Resets to the default IPv4 address.

-prefix_6 *int*
    (Optional) The IPv6 prefix for the service assistant.

-resetpassword
    (Optional) Sets the service assistant password to the default value.

Description

This command resets the service assistant IP address to the default value. If the command is run on the upper canister, the default value is 192.168.70.121 subnet mask: 255.255.255.0. If the command is run on the lower canister, the default value is 192.168.70.122 subnet mask: 255.255.255.0. If the node canister is active in a system, the superuser password for the system is reset; otherwise, the superuser password is reset on the node canister.

If the node canister becomes active in a system, the superuser password is reset to that of the system. You can configure the system to disable resetting the superuser password. If you disable that function, this action fails.

This action calls the **satask chserviceip** command and the **satask resetpassword** command.

Reset service assistant password command:

Use this command when you are unable to logon to the system because you have forgotten the superuser password, and you wish to reset it.

Syntax

```
►►── satask ── resetpassword ── ─────────────────────────────────────────────────── ►◄
```

Parameters

None.

Description

This command resets the service assistant password to the default value `passw0rd`. If the node canister is active in a system, the superuser password for the system is reset; otherwise, the superuser password is reset on the node canister.

If the node canister becomes active in a system, the superuser password is reset to that of the system. You can configure the system to disable resetting the superuser password. If you disable that function, this action fails.

This command calls the **satask resetpassword** command.

Snap command:

Use the **snap** command to collect diagnostic information from the node canister and to write the output to a USB flash drive.

Syntax

```
►►── satask ── snap ──┬─────────┬──┬─────────┬──┬──────────────┬── ►◄
                      └─ -dump ─┘  └─ -noimm ─┘  └─ panel_name ─┘
```

Parameters

-dump
    (Optional) Indicates the most recent dump file in the output.

-noimm
    (Optional) Indicates the `/dumps/imm.ffdc` file should not be included in the output.

panel_name
    (Optional) Indicates the node on which to execute the **snap** command.

Description

This command moves a snap file to a USB flash drive.

This command calls the **satask snap** command.

If collected, the IMM FFDC file is present in the **snap** archive in
/dumps/imm.ffdc.<node.dumpname>.<date>.<time>.tgz. The system waits for up to
five minutes for the IMM to generate its FFDC. The status of the IMM FFDC is
located in the **snap** archive in /dumps/imm.ffdc.log. These two files are not left on
the node.

An invocation example

```
satask snap -dump 111584
```

The resulting output:

```
No feedback
```

Install software command:

Use this command to install a specific update package on the node canister.

Syntax

```
►►── satask ── installsoftware ── ── -file ──filename── ──┬────────────┬── ──────────►◄
                                                           ├─ -ignore ──┤
                                                           └─ -pacedccu ─┘
```

Parameters

-file *filename*
>   (Required) The *filename* designates the name of the update package .

-ignore | -pacedccu
>   (Optional) Overrides prerequisite checking and forces installation of the update
>   package.

Description

This command copies the file from the USB flash drive to the update directory on
the node canister, and then installs the update package.

This command calls the **satask installsoftware** command.

Create system command:

Use this command to create a storage system.

Syntax

```
►►── satask ── mkcluster ── ── -clusterip ──ipv4── ──┬──────────┬── ──┬───────────┬── ──┬─────────────────────┬── ──────────►◄
                                                     └─ -gw ──ipv4─┘    └─ -mask ──ipv4─┘    └─ -name ──cluster_name─┘


►►── satask ── mkcluster ── ── -clusterip_6 ──ipv6── ──┬────────────┬── ──┬──────────────┬── ──┬─────────────────────┬── ──────►◄
                                                       └─ -gw_6 ──ipv6─┘    └─ -prefix_6 ──int─┘    └─ -name ──cluster_name─┘
```

Parameters

-clusterip *ipv4*
(Optional) The IPv4 address for Ethernet port 1 on the system.

-gw *ipv4*
(Optional) The IPv4 gateway for Ethernet port 1 on the system.

-mask *ipv4*
(Optional) The IPv4 subnet for Ethernet port 1 on the system.

-clusterip_6 *ipv6*
(Optional) The IPv6 address for Ethernet port 1 on the system.

-gw_6 *ipv6*
(Optional) The IPv6 gateway for Ethernet port 1 on the system.

-prefix_6 *int*
(Optional) The IPv6 prefix for Ethernet port 1 on the system.

-name *cluster_name*
(Optional) The name of the new system.

Description

This command creates a storage system.

This command calls the **satask mkcluster** command.

Query status command:

Use this command to determine the current service state of the node canister.

Syntax

```
►►── sainfo ── getstatus ── ──────────────────────────────────────────── ──◄
```

Parameters

None.

Description

This command writes the output from each node canister to the USB flash drive.

This command calls the **sainfo lsservicenodes** command, the **sainfo lsservicestatus** command, and the **sainfo lsservicerecommendation** command.

Initialization tool interface

Use the Initialization tool interface to initialize a system and to service the node canisters in a control enclosure. Although the Initialization tool wizard interface is similar, accessing the wizard differs, depending on the generation of your control enclosure model.

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 53. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 54. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Technician port for Storwize V7000 Gen2

A technician port is available on the rear of the Storwize V7000 Gen2 system for service support.

Use the technician port to access the initialization tool for a Storwize V7000 Gen2 control enclosure. You do not use the USB flash drive for Storwize V7000 Gen2 initialization.

(Storwize V7000 Gen2 control enclosures can use a Storwize V7000 Gen1 USB flash drive for other purposes but not for initialization.)

Verify that you are using a supported operating system. The initialization tool is valid for the following operating systems:

- Microsoft Windows 8.1 (64-bit), or Microsoft Windows 7 (64-bit)
- Apple MacOS X 10.7[1]
- Red Hat Enterprise Server 5 and 6
- Ubuntu desktop 11.04 and 13.10

---

1. Apple OS X 10.9 is undergoing testing and so, currently remains unclaimed.

USB flash drive and the Storwize V7000 Gen1 Initialization tool interface

Use a USB flash drive to initialize a system and also to help service the node canisters in a control enclosure. For Storwize V7000 Gen2 use "Procedure: Initializing the Storwize V7000 Gen2 system using the technician port" on page 124 to initialize and service the Storwize V7000 Gen2 system.

The initialization tool is a Windows application. Use the initialization tool to set up the USB flash drive to perform the most common tasks.

When a USB flash drive is inserted into one of the USB ports on a node canister in a control enclosure, the software searches for a control file on the USB flash drive and runs the command that is specified in the file. When the command completes, the command results and node status information are written to the USB flash drive.

Using the initialization tool on a Storwize V7000 Gen1 control enclosure

The initialization tool is a graphical user interface (GUI) wizard that creates the initial configuration on the control enclosure.

Before you begin

Access the initialization tool through a USB flash drive for a Storwize V7000 Gen1 control enclosure.

Note: Access the initialization tool for a Storwize V7000 Gen2 control enclosure through the technician port. You do not use the USB flash drive for Storwize V7000 Gen2 initialization.

Verify that you are using a supported operating system. The initialization tool is valid for the following operating systems:

- Microsoft Windows 8.1 (64-bit), or Microsoft Windows 7 (64-bit)
- Apple MacOS X 10.7[2]
- Red Hat Enterprise Server 5 and 6
- Ubuntu desktop 11.04 and 13.10

About this task

By using the initialization tool, you can set the USB flash drive to run one of the following tasks:

- Initialize a new system.
- Reset the superuser password.
- Set or reset the service assistant IP address on a node canister on the control enclosure.
- Set the management IP addresses.

For any other tasks that you want to complete on a node canister on the control enclosure, you must create the satask.txt file by using a text editor.

---

2. Apple OS X 10.9 is undergoing testing and so, currently remains unclaimed.

The initialization tool is available on the USB flash drive that is shipped with the control enclosures. The name of the file is `InitTool`. If you cannot locate the USB flash drive, you can download the application from the following support website (search for **initialization tool** or **init tool**).

`www.ibm.com/support`

Procedure

To use the initialization tool, complete the following steps.

1. If you downloaded the initialization tool, copy the file onto the USB flash drive that you are going to use.
2. To start the initialization tool, insert the USB flash drive that contains the program into a USB slot on a suitable personal computer.
3. Run the `InitTool` program from the USB drive.
   - **Windows:** Open the USB flash drive and double-click `InitTool.bat`.
   - **Apple Macitosh:** Locate the root directory of the USB flash drive (usually located in the `/Volumes/` directory). Type `sh InitTool.sh`.
   - **Linux:** Locate the root directory of the USB flash drive. (It is usually located in the `/media/` directory. If an automatic mount system is used, the root directory can be located by typing the mount command.) Type `sh InitTool.sh`.

   The initialization tool prompts you for the task that you want to complete and for the parameters that are relevant to that task. It prompts you when to put it in the node canister on the control enclosure.
4. After the `satask.txt` file is created, follow the instructions in "Using a USB flash drive" to run the commands on the node.
5. When the commands are completed, return the USB flash drive to your personal computer and start the tool again to see the results.

When to use the USB flash drive

The USB flash drive is normally used to initialize the configuration after installing a new system. In the case of Storwize V7000 Gen2, the USB flash drive is not required to initialize the configuration; however, it can be used for other functions.

Using the USB flash drive is required in the following situations:
- When you cannot connect to a node canister in a control enclosure using the service assistant and you want to see the status of the node.
- When you do not know, or cannot use, the service IP address for the node canister in the control enclosure and must set the address.
- When you have forgotten the superuser password and must reset the password.

Using a USB flash drive

Use any USB flash drive that is formatted with a FAT32 file system on its first partition.

About this task

When a USB flash drive is plugged into a node canister, the nodecanister code searches for a text file that is named `satask.txt` in the root directory. If the code finds the file, it attempts to run a command that is specified in the file. When the command completes, a file that is called `satask_result.html` is written to the root directory of the USB flash drive. If this file does not exist, it is created. If it exists, the data is inserted at the start of the file. The file contains the details and results

of the command that was run and the status and the configuration information from the node canister. The status and configuration information matches the detail that is shown on the service assistant home page panels.

The `satask.txt` file can be created on any workstation by using a text editor. If a Microsoft Windows workstation is being used, the initialization tool can be used to create the commands that are most often used.

The fault light-emitting diode (LED) on the nodecanister flashes when the USB service action is being completed. When the fault LED stops flashing, it is safe to remove the USB flash drive.

Results

The USB flash drive can then be plugged into a workstation, and the `satask_result.html` file can be viewed in a web browser.

To protect from accidentally running the same command again, the `satask.txt` file is deleted after it is read.

If no `satask.txt` file is found on the USB flash drive, the result file is still created, if necessary, and the status and configuration data is written to it.

Chapter 6. Resolving a problem

Described here are some procedures to help resolve fault conditions that might exist on your system and which assume a basic understanding of the Storwize V7000 system concepts.

The following procedures are often used to find and resolve problems:
- Procedures that involve data collection and system configuration
- Procedures that are used for hardware replacement.

Always use the recommended actions on the Events panel of the management GUI as the starting point to diagnose and resolve a problem.

The following topics describe a type of problem that you might experience, that is not resolved by using the management GUI. In those situations, review the symptoms and follow the actions that are provided here.

The "Start here: Use the management GUI recommended actions" topic gives the starting point for any service action. The situations covered in this section are the cases where you cannot start the management GUI or the node canisters in the control enclosure are unable to run the system software.

Note: After you have created your clustered system, remove hardware components only when directed to do so by the fix procedures. Failure to follow the procedures can result in loss of access to data or loss of data. Follow the fix procedures when servicing a control enclosure.

Start here: Use the management GUI recommended actions

The management GUI provides extensive facilities to help you troubleshoot and correct problems on your system.

You can connect to and manage a Storwize V7000 system using the management GUI as soon as you have created a clustered system. If you cannot create a clustered system, see the problem that contains information about what to do if you cannot create one. Go to "Problem: Cannot initialize or create a Storwize V7000 Gen1 clustered system" on page 91.

To run the management GUI, start a supported web browser and point it to the management IP address of your system. Up to four addresses can be configured for your use. There are two addresses for IPv4 access, and two addresses for IPv6 access. If you do not know the system management IP address, go to "Problem: Management IP address unknown" on page 86. After the connection is successful, you see a login panel. If you are unable to access the login panel, go to "Problem: Unable to connect to the Storwize V7000 Gen1 management GUI" on page 88.

Log on using your user name and password. If you are unable to log on, go to "Problem: Unable to log on to the management GUI" on page 89.

When you have logged on, select **Monitoring** > **Events**. Depending on how you choose to filter alerts, you might see only the alerts that require attention, alerts and messages that are not fixed, or all event types whether they are fixed or unfixed.

Select the recommended alert, or any other alert, and run the fix procedure. The fix procedure steps you through the process of troubleshooting and correcting the problem. The fix procedure displays information that is relevant to the problem and provides various options to correct the problem. Where it is possible, the fix procedure runs the commands that are required to reconfigure the system.

Always use the recommended action for an alert because these actions ensure that all required steps are taken. Use the recommended actions even in cases where the service action seems obvious, such as a drive showing a fault. In this case, the drive must be replaced and reconfiguration must be performed. The fix procedure performs the reconfiguration for you.

The fix procedure also checks that another existing problem does not result in a fix procedure that causes volume data to be lost. For example, if a power supply unit in a node enclosure must be replaced, the fix procedure checks and warns you if the integrated battery in the other power supply unit is not sufficiently charged to protect the system.

If possible, fix the alerts in the order shown to resolve the most serious issues first. Often, other alerts are fixed automatically because they were the result of a more serious issue.

After all the alerts are fixed, go to "Procedure: Checking the status of your system" on page 104.

Problem: Management IP address unknown

If you are not able to run the management GUI because you do not know the management IP address, you can retrieve the management IP address. For Storwize V7000 Gen2 follow the procedure in "Procedure: Initializing the Storwize V7000 Gen2 system using the technician port" on page 124.

The management IP address is set when the clustered system is created. An address for port 2 can be added after the clustered system is created.

The management IP address is part of the data that is shown in the service assistant home panel or the data that is returned by the USB flash drive.

Problem: Unable to connect to the management GUI

If you are unable to connect to the management GUI from your web browser and received a `Page not found` or similar error, this information might help you resolve the issue. The connection information differs, depending on the generation of your control enclosure model.

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 55. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives

Table 55. Storwize V7000 Gen1 model numbers (continued)

Machine type/model	Description
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 56. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Problem: Unable to connect to the Storwize V7000 Gen2 management GUI

If you are unable to connect to the management GUI from your web browser and received a `Page not found` or similar error, this information might help you resolve the issue.

If you are unable to connect to the management GUI from your web browser and received a `Certificate expired` or similar error, see Resolving a problem with SSL certificates.

If you are unable to connect to the management GUI from your web browser and received a `cipher error`, `SSL error`, `TLS error`, or `handshake error` or similar error, see Resolving a problem with SSL/TLS clients.

Consider the following possibilities if you are unable to connect to the Storwize V7000 Gen2 management GUI:

- You cannot connect if the system is not operational with at least one node online. If you know the service address of a node canister, either use the service assistant to verify that the state of at least one node canister is active, or if the node canister is not active, use the LEDs to see if any node canister state is active.

  If there is not a node canister with a state of active, resolve the reason why it is not in active state. If the state of all node canisters is candidate, then there is not

a clustered system to connect to. If the node state is service, see the topic that contains information about fixing node errors.

- Ensure that you are using the correct system IP address. If you can access the service assistant using the service address or the technician port of a node canister, log in to find the node and system addresses on the **Access** tab of the Node Detail table.

- Ensure that all node canisters have an Ethernet cable that is connected to port 1 and that the port is working. Use the Ethernet LEDs to understand the port status.

- Ping the management address to see if the Ethernet network permits the connection. If the ping fails, check the Ethernet network configuration to see if there is a routing or a firewall issue. Ensure that the Ethernet network configuration is compatible with the gateway and subnet or prefix settings. Ensure that you did not use the Ethernet address of another device as the management address. If necessary, modify your network settings to establish a connection.

- If the system IP address settings are incorrect for your environment, take these steps:

  1. You can determine this if you can access the service assistant on any node canister. Access the service assistant using the technician port on the rear of a node canister if it cannot be accessed over your network. Alternatively use the summary data returned, when a USB flash drive is plugged into a node canister.

  2. You can temporarily run the management GUI on the service address of the configuration node. Point your browser to service address/gui. For example, if the service address of the configuration node is 11.22.33.44, point your browser to 11.22.33.44/gui.

  3. In the Management GUI, use the options in the **Settings** > **Network** to change the management IP settings.

  4. As an alternative to using the management GUI, you can use the `chsystemip` CLI command to correct the system IP address settings by using ssh to the service IP of the configuration node.

Problem: Unable to connect to the Storwize V7000 Gen1 management GUI

If you are unable to connect to the management GUI from your web browser and received a `Page not found` or similar error, this information might help you resolve the issue.

If you are unable to connect to the management GUI from your web browser and received a `Certificate expired` or similar error, see Resolving a problem with SSL certificates.

If you are unable to connect to the management GUI from your web browser and received a `cipher error`, `SSL error`, `TLS error`, or `handshake error` or similar error, see Resolving a problem with SSL/TLS clients.

Consider the following possibilities if you are unable to connect to the management GUI:

- You cannot connect if the system is not operational with at least one node online. If you know the service address of a node canister, go to "Procedure: Getting node canister and system information using the service assistant" on page 105; otherwise, go to "Procedure: Getting node canister and system

information by using a USB flash drive" on page 105 and obtain the state of each of the node canisters from the data that is returned. If there is not a node canister with a state of active, resolve the reason why it is not in active state. If the state of all node canisters is candidate, then there is not a clustered system to connect to. If all nodes are in a service state, go to "Procedure: Fixing node errors" on page 121.

- Ensure that you are using the correct system IP address. If you know the service address of a node canister, go to "Procedure: Getting node canister and system information using the service assistant" on page 105; otherwise, go to "Procedure: Getting node canister and system information by using a USB flash drive" on page 105 and obtain the management IP address from the data that is returned.
- Ensure that all node canisters have an Ethernet cable that is connected to port 1 and that the port is working. To understand the port status, go to "Procedure: Finding the status of Storwize V7000 Gen1 Ethernet connections" on page 118.
- Ping the management address to see if the Ethernet network permits the connection. If the ping fails, check the Ethernet network configuration to see if there is a routing or a firewall issue. Ensure that the Ethernet network configuration is compatible with the gateway and subnet or prefix settings. Ensure that you have not used the Ethernet address of another device as the management address. If necessary, modify your network settings to establish a connection.
- If the system IP address settings are incorrect for your environment, take these steps:
  1. Determine the service address of the configuration node canister. You can determine this if you can access the service assistant on any node canister, alternatively use the summary data returned, when a USB flash drive is plugged into a node canister.
  2. You can temporarily run the management GUI on the service address of the configuration node. Point your browser to *service address*/gui. For example, if the service address of the configuration node is 11.22.33.44, point your browser to 11.22.33.44/gui.
  3. Use the options in the **Settings** > **Network** panel to change the management IP settings.
  4. As an alternative to using the management GUI, you can use the **chsystemip** CLI command to correct the system IP address settings by using ssh to the service IP of the configuration node.

Problem: Unable to log on to the management GUI

If you can see the management GUI login screen but cannot log on, you have several options for correcting the problem.

Log on using your user name and password. Complete the suggested actions when you encounter a specific situation.

- If you are not logging on as superuser, contact your system administrator to verify your user name and reset your account password.
- If the user name that you are using is authenticated through a remote authentication server, verify that the server is available. If the authentication server is unavailable, you can log on as user name superuser. This user is always authenticated locally.
- If you do not know the password for superuser, go to "Procedure: Resetting superuser password" on page 99.

Problem: Cannot initialize or create a clustered system

Use this information if your attempt to create a clustered system has failed. This information varies depending on the generation of your control enclosure model.

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 57. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 58. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Problem: Cannot initialize or create a Storwize V7000 Gen2 clustered system

Use this information if your attempt to create a Storwize V7000 Gen2 clustered system failed.

The failure is reported regardless of the method that you used to create a system:

- System initialization wizard (using the technician port of a node canister)
- USB flash drive
- Management console
- Service assistant
- Service command line

To prevent accidental loss of volumes, a system cannot be initialized on an enclosure that is already configured in an existing system. Check the following details using the service assistant to confirm that the enclosure is not already configured in a system:

- The node canister that you are attempting to create a clustered system on must be in candidate state.
- The partner node canister in the control enclosure must not be in active state.
- The latest system ID of the control enclosure must be 0. If the partner node is in active state, see "Start here: Use the management GUI recommended actions" on page 85. If the partner code is not in the active state and the node canister on which you are attempting to create the system is in service state, see "Procedure: Fixing node errors" on page 121.

Problem: Cannot initialize or create a Storwize V7000 Gen1 clustered system

If your attempt to create a system has failed, there are several possible solutions depending on the current state of the storage system.

The failure is reported regardless of the method that you used to create a clustered storage system:

- USB flash drive
- Management console
- Service assistant
- Service command line

The create clustered-system function protects the system from loss of volume data. If you create a clustered system on a control enclosure that was previously used, you lose all of the volumes that you previously had. To determine whether there is an existing system, use data that is returned by "Procedure: Getting node canister and system information using the service assistant" on page 105 or "Procedure: Getting node canister and system information by using a USB flash drive" on page 105.

- The node canister that you are attempting to create a clustered system on is in candidate state. The node canister is in candidate state if it is a new canister.
- The partner node canister in the control enclosure is not in active state.
- The latest system ID of the control enclosure is 0.

If the create function failed because there is an existing system, fix the existing clustered system; do not re-create a new clustered system. If you want to create a clustered system and do not want to use any data from the volumes that were used in the previous clustered system, go to "Procedure: Deleting a system completely" on page 120, and then run the create function again.

You might not be able to create a cluster if the node canister (the one on which you are attempting to create the clustered system) is in service state. Check whether the node canister is in service state by using the data that was returned by "Procedure: Getting node canister and system information using the service assistant" on page 105 or "Procedure: Getting node canister and system information by using a USB flash drive" on page 105. If the node is in service state, fix the reported node errors. For more information, go to "Procedure: Fixing node errors" on page 121. After the node error is corrected, attempt to create a clustered storage system again.

Problem: Node canister service IP address unknown

You can use several methods to determine the service address of a node canister. The methods of determining the service address differ, depending on the generation of your control enclosure model.

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 59. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 60. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Problem: Storwize V7000 Gen2 node canister service IP address unknown

You can use several methods to determine the service address of a node canister.

A default service address is initially assigned to each node canister, as shown in Table 61 on page 93. Try using these addresses if the node has not been reconfigured, and the addresses are valid on your network.

If you are able to access the management GUI, the service IP addresses of the node canisters are shown by selecting a node and port at **Settings** > **Network** > **Service IP Addresses**.

If you are unable to access the management GUI but you know the management IP address of the system, you can use the address to log in to the service assistant that is running on the configuration node:

1. Point your browser at the /service directory of the management IP address of the system. If your management IP address is 11.22.33.44, point your web browser to 11.22.33.44/service.
2. Log in to the service assistant.
3. The service assistant home page lists the node canister that can communicate with the node.
4. If the service address of the node canister that you are looking for is listed in the Change Node window, make the node the current node. Its service address is listed under the Access tab of the node details.

If you know the service IP address of any node canister in the system, you can log in to the service assistant of that node. Follow the previous instructions for using the service assistant, but at step 1, point your browser at the /service directory of the service IP address you know. If you know that a service IP address is 11.22.33.56, point your web browser to 11.22.33.56/service

Some types of errors can prevent nodes from communicating with each other; in that event, it might be necessary to point your browser directly at the service assistant of the node that requires administering, rather than change the current node in the service assistant.

If you cannot find the service address of the Storwize V7000 Gen2 node using the management GUI or service assistant of a different node, two options remain:
- You can connect directly to the service assistant of a node using the technician port of the node, as described in "Procedure: Initializing the Storwize V7000 Gen2 system using the technician port" on page 124.
- You can also use a USB flash drive to find the service address of the node. For more information, see "Procedure: Getting node canister and system information by using a USB flash drive" on page 105.

Table 61. Default service IP addresses

Canister and port	IPv4 address	IPV4 subnet mask
Canister 1 (left) port 1 (left)	192.168.70.121	255.255.255.0
Canister 2 (right) port 1 (left)	192.168.70.122	255.255.255.0

Problem: node canister service IP address unknown

You can use several methods to determine the service address of a node canister.

A default service address is initially assigned to each node canister, as shown in Table 62 on page 94. Try using these addresses if the node has not been reconfigured, and the addresses are valid on your network.

If you are able to access the management GUI, the service IP addresses of the node canisters are shown by selecting a node and port at **Settings** > **Network** > **Service IP Addresses**.

If you are unable to access the management GUI but know the management IP address of the system, use the IP address to log in to the service assistant that is running on the configuration node.
1. Point your browser at the /service directory of the management IP address of the system. If your management IP address is 11.22.33.44, point your web browser to 11.22.33.44/service.

2. Log in to the service assistant.

    The service assistant home page lists the node canister that can communicate with the node.

3. If the service address of the node canister that you are looking for is listed in the Change Node window, make the node the current node. Its service address is listed under the Access tab of the node details.

If you know the service IP address of any node canister in the system, you can log in to the service assistant of that node. Follow the previous instructions for using the service assistant, but at step 1, point your browser at the /service directory of the service IP address you know. For example, if you know the service IP address is 11.22.33.56, point your web browser to 11.22.33.56/service.

Some types of errors can prevent nodes from communicating with each other; in that event, it might be necessary to point your browser directly at the service assistant of the node that requires administering, rather than change the current node in the service assistant.

If you are unable to find the service address of the node by using the management GUI or service assistant, you can also use a USB flash drive to find it. For more information, see "Procedure: Getting node canister and system information by using a USB flash drive" on page 105.

Table 62. Default service IP addresses

Canister	IPv4 address	IPV4 subnet mask
Canister 1 (left)	192.168.70.121	255.255.255.0
Canister 2 (right)	192.168.70.122	255.255.255.0

Problem: Cannot connect to the service assistant

Use this information if you are unable to display the service assistant on your browser via the service IP address.

You might encounter a number of situations when you cannot connect to the service assistant.

- Check that you have entered the /service path after the service IP address. Point your web browser to *service IP address*/service for the node that you want to work on. For example, if you set a service address of 11.22.33.44 for a node canister, point your browser to 11.22.33.44/service.
- Check that you are using the correct service address for the node canister. To find the IPv4 and IPv6 addresses that are configured on the node, go to "Problem: node canister service IP address unknown" on page 93. Try accessing the service assistant through these addresses. Verify that the IP address, subnet, and gateway are specified correctly for IPv4 addresses. Verify that the IP address, prefix, and gateway are specified for the IPv6 addresses. If any of the values are incorrect, see "Procedure: Changing the service IP address of a node canister" on page 121.
- You cannot connect to the service assistant if the node canister is not able to start the Storwize V7000 code. To verify that the LEDs indicate that the code is active, see "Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs" on page 112.
- The service assistant is configured on Ethernet port 1 of a node canister. Verify that an Ethernet cable is connected to this port and to an active port on your

Ethernet network. See "Procedure: Finding the status of Storwize V7000 Gen1 Ethernet connections" on page 118 for details.

- Ping the service address to see if the Ethernet network permits the connection. If the ping fails, check the Ethernet network configuration to see if there is a routing or a firewall issue. Check that the Ethernet network configuration is compatible with the gateway and subnet or prefix settings. Check that you have not used an address that is used by another device on your Ethernet network. If necessary, change the network configuration or see "Procedure: Changing the service IP address of a node canister" on page 121 to change the service IP address of a node.

- A default service address is initially assigned to each node canister. The service IP address 192.168.70.121 subnet mask 255.255.255.0 is preconfigured on Ethernet port 1 of the upper canister, canister 1. The service IP address 192.168.70.122 subnet mask 255.255.255.0 is preconfigured on Ethernet port 1 of the lower canister, canister 2.

  You might not be able to access these addresses because of the following conditions:

  – These addresses are the same as the addresses that are used by other devices on the network.

  – These addresses cannot be accessed on your network.

  – There are other reasons why they are not suitable for use on your network.

  If the previous conditions apply, see "Procedure: Changing the service IP address of a node canister" on page 121 to change the service IP address to one that works in your environment.

  If you are unable to change the service address, for example, because you cannot use a USB flash drive in the environment, see "Procedure: Accessing aStorwize V7000 Gen1 canister by using a directly attached Ethernet cable" on page 126.

Problem: Management GUI or service assistant does not display correctly

If the Management GUI or the service assistant does not display correctly, verify that you are using a supported web browser.

For a list of supported browsers, see http://pic.dhe.ibm.com/infocenter/storwize/ ic/topic/com.ibm.storwize.v7000.730.doc/svc_configuringbrowser_1obg15.html.

Problem: A node canister has a location node error

The node error listed on the service assistant home page or in the event log can indicate a location error.

A location error means that the node canister or the enclosure midplane has been moved or changed. This is normally due to a service action not being completed or not being implemented correctly.

A number of different conditions are reported as location errors. Each condition is indicated by different node error. To find out how to resolve the node error, go to "Procedure: Fixing node errors" on page 121.

Be aware that after a node canister has been used in a system, the node canister must not be moved to a different location, either within the same enclosure or in a different enclosure because this might compromise its access to storage, or a host

application's access to volumes. Do not move the canister from its original location unless directed to do so by a service action.

Problem: SAS cabling not valid

Use this procedure if you receive errors to determine if your SAS cabling is valid. The procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 63. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 64. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Problem: Storwize V7000 Gen2 SAS cabling not valid

This topic provides information to be aware of if you receive errors that indicate the SAS cabling is not valid.

Check the following items:
- No more than 10 expansion enclosures can be chained to port 1 (below the control enclosure). The connecting sequence from SAS port 1 of the node canister is called chain 1.

- No more than 10 expansion enclosures can be chained to port 2 (above the control enclosure). The connecting sequence from SAS port 2 of the node canister is called chain 2.
- Do not connect a SAS cable between a port on a left canister and a port on a right canister.
- In any enclosure, the same ports must be used on both canisters.
- Do not connect a SAS cable between ports in the same enclosure.
- In each enclosure, where a cable connects from SAS port 1 of the left canister to another enclosure, a cable must also connect from SAS port 1 of the right canister to the other enclosure. Similarly, where a cable connects from SAS port 2 of the left canister to another enclosure, a cable must also connect from SAS port 2 of the right canister to the other enclosure.
- For cables connected between expansion enclosures, one end is connected to port 1 while the other end is connected to port 2.
- For cables that are connected between a control enclosure and expansion enclosures, port 1 must be used on the expansion enclosures.
- The last enclosure in a chain must not have cables in port 2 of the left canister, nor port 2 of the right canister.
- Ensure that each SAS cable is fully inserted.

See the topic about installing SAS cables in the *IBM Storwize V7000 Gen2 Quick Installation Guide*.

Problem: Storwize V7000 Gen1 SAS cabling not valid

This topic provides information to be aware of if you receive errors that indicate the SAS cabling is not valid.

Check the following items:
- No more than five expansion enclosures can be chained to port 1 (below the control enclosure). The connecting sequence from port 1 of the node canister is called chain 1.
- No more than four expansion enclosures can be chained to port 2 (above the control enclosure). The connecting sequence from port 2 of the node canister is called chain 2.
- Do not connect a SAS cable between a port on an upper canister and a port on a lower canister.
- In any enclosure, the same ports must be used on both canisters.
- No SAS cable can be connected between ports in the same enclosure.
- For any enclosure, the cables that are connected to SAS port 1 on each canister must attach to the same enclosure. Similarly, for any enclosure, the cables that are connected to SAS port 2 on each canister must attach to the same enclosure. Cable attachments for SAS port 1 and cable attachments for SAS port 2 do not go to the same enclosure.
- For cables connected between expansion enclosures, one end is connected to port 1 while the other end is connected to port 2.
- For cables that are connected between a control enclosure and expansion enclosures, port 1 must be used on the expansion enclosures.
- The last enclosure in a chain must not have cables in port 2 of canister 1 and port 2 of canister 2.
- Ensure that each SAS cable is fully inserted.

Problem: New expansion enclosure not detected

Determine why a newly installed expansion enclosure was not detected by the system.

When you install a new expansion enclosure, follow the management GUI Add Enclosure wizard. Select **Monitoring** > **System**. From the **Actions** menu, select **Add Enclosures**.

If the expansion enclosure is not detected, complete the following verifications:

- Verify the status of the LEDs at the back of the expansion enclosure. At least one power supply unit must be on with no faults shown. At least one canister must be active, with no fault LED on. Serial-attached SCSI (SAS) ports of Storwize V7000 Gen1 each have four green link-status LEDS, which must be on for any port in use.

  Storwize V7000 Gen2 has two LEDs per SAS port, one green link-status LED and one amber fault LED. The link status LED of the ports that are in use is on while the fault LED is off. For details about LED status, see "Procedure: Understanding the system status using the LEDs" on page 106.

- Verify that the SAS cabling to the expansion enclosure is correctly installed. To review the requirements, see "Problem: SAS cabling not valid" on page 96.

Problem: Control enclosure is not detected

If a control enclosure is not detected by the system, this procedure can help you resolve the problem.

When you install a new control enclosure, use the **Add Enclosures** wizard in the management GUI. To access this wizard, select **Monitoring** > **System**. On the **Systems** page, select **Actions** > **Add Enclosures**.

If the control enclosure is not detected, check the following items:

- The enclosure is powered on.
- The enclosure is not part of another system.
- At least one node is in candidate state.
- The Fibre Channel cables are connected and zoning is set up according to the zoning rules defined in the *SAN configuration and zoning rules summary* topic. There must be a zone that includes all ports from all node canisters.
- The existing system and the nodes in the enclosure that are not detected have version 6.2 or later installed.

Problem: Mirrored volume copies no longer identical

The management GUI provides options to either check copies that are identical or to check that the copies are identical and to process any differences that are found.

To confirm that the two copies of a mirrored volume are still identical, choose the volume view that works best for you. Select one of the volume copies in the volume that you want to check. From the **Actions** menu, select the **Validate Volume Copies** option.

You have the following choices:

- Validate that the volume copies are identical.

- Validate that the volume copies are identical, mark, and repair any differences that are found.

If you want to resolve any differences, you have the following options:

- Consider that the primary volume copy is correct and make the other volume copy match the primary volume copy if any differences are found. The primary volume copy is the copy that is considered correct.
- Do not assume that either volume copy is correct. If a difference is found, the sector is marked. A media error is returned if the volume is read by a host application.

Problem: Command file not processed from USB flash drive

This information assists you in determining why the command file is not being processed, when using a USB flash drive.

You might encounter this problem during initial setup (the USB flash drive is not used for initial setup of the Storwize V7000 Gen2) or when running commands if you are using your own USB flash drive rather than the USB flash drive that was packaged with your order.

If you encounter this situation, verify the following items:

- That an `satask_result.html` file is in the root directory on the USB flash drive. If the file does not exist, then the following problems are possible:
  - The USB flash drive is not formatted with the correct file system type. Use any USB flash drive that is formatted with FAT32 file system on its first partition; for example, NTFS is not a supported type. Reformat the key or use a different key.
  - The USB port is not working. Try the key in the other USB port.
  - The node is not operational. Check the status using the light-emitting diodes (LEDs). See "Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs" on page 112.
- If there is a `satask_result.html` file, check the first entry in the file. If there is no entry that matches the time the USB flash drive was used, it is possible that the USB port is not working or the node is not operational. Check the node status using the LEDs. See "Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs" on page 112.
- If there is a status output for the time the USB flash drive was used, then the `satask.txt` file was not found. Check that the file was named correctly. The `satask.txt` file is automatically deleted after it has been processed.

Procedure: Resetting superuser password

You can reset the superuser password to the default password of `passw0rd` by using a special command action. The password procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 65. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 66. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Procedure: Resetting the superuser password for Storwize V7000 Gen2

The primary method for resetting the superuser password is to change the password as you log in, with the link on the log-in page. You can also access the service assistant from the technician port to change the password.

If the password reset function is enabled, the log-in page displays a link for resetting the password. You can also use the technician port to access the Storwize V7000 Gen2 service assistant.

If the password reset function is not enabled, there is no work-around. You must contact the person who knows the password. The USB flash drive interface also supports resetting the password.

Procedure: Resetting the Storwize V7000 Gen1 superuser password

You can reset the superuser password to the default password of `passw0rd` by using a USB flash drive command action.

About this task

You can use this procedure to reset the superuser password if you have forgotten the password. This command runs differently depending on whether you run it on a node canister that is active in a clustered system.

Note: If a node canister is not in active state, the superuser password is still required to log on to the service assistant.

It is possible to configure your system so that resetting the superuser password with the USB flash drive command action is not permitted. If your system is configured this way, there is no work-around. Contact the person who knows the password.

To use a USB flash drive to reset the superuser password, see "USB flash drive and the Storwize V7000 Gen1 Initialization tool interface" on page 81.

See also "Problem: Unable to log on to the management GUI" on page 89.

Results

If the node canister is active in a clustered system, the password for superuser is changed on the clustered system. If the node canister is not in active state, the superuser password for the node canister is changed. If the node canister joins a clustered system later, the superuser password is reset to that of the clustered system.

Procedure: Identifying which enclosure or canister to service

Use this procedure to identify which enclosure or canister must be serviced. Identification procedures differ, depending on the generation of your control enclosure model.

About this task

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 67. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 68. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Procedure: Identifying which Storwize V7000 Gen2 enclosure or canister to service

Use this procedure to identify which Storwize V7000 Gen2 enclosure or canister must be serviced.

About this task

To prevent loss of access to data or loss of data when you service your system, it is important to be able to identify the correct enclosure or canister when you complete a service action. Each enclosure is identified by its model type and serial number. Model type and serial number are indicated by labels on the enclosure front and rear.

Each canister is identified by its enclosure and slot location. Viewing an enclosure from the rear, slot 1 is on the left and slot 2 is on the right. There are physical differences between Storwize V7000 Gen2 control enclosures and expansion enclosures.

Looking at the front of a rack:
- The type of the enclosure, either Control or Expansion, is labeled on the left bezel.
- The model type and serial number of the enclosure are found at the bottom of the left bezel.

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 69. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives

Table 69. Storwize V7000 Gen2 model numbers (continued)

Machine type/model	Description
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Looking at the rear of a rack:

- Control enclosures contain tall and narrow power supplies at the far left and right of the enclosure, with the node canisters installed in between, side by side.
- Expansion enclosures contain short node canisters installed side by side, above short power supply units also installed side by side.
- Control enclosures contain node canisters that have Ethernet ports, and USB ports. Expansion enclosures do not have any of these ports. The model type is shown on a label.
- Each canister is identified by its slot location. Viewing an enclosure from the rear, slot 1 is on the left and slot 2 is on the right.

A canister is uniquely identified by the enclosure that it is in and its slot location. This ID is shown as E-C or E|C where E is the enclosure ID and C is the canister location. On the service assistant, the ID is known as the Panel.

Notes:

- When a node canister is added to a clustered system as a node, it is given a node name and a node ID. The default node name is nodeN, where N is the node ID. The node ID does not represent the slot location of the node. On the **Monitoring** > **System** page use the dynamic graphic to show the back of the system. Hover over the canister to display the node name and canister location. The service assistant home page also shows both the node name and the canister location. If you have only the node name, use these panels to determine the node canister location.
- It is good practice to use this procedure to identify which enclosure or canister must be serviced, as completing a service action on the wrong canister can lead to loss of access to data or loss of data.

To control the identify LED of an enclosure or online canister, use the management GUI:

1. Log in to the management GUI for the system.
2. Select the **Monitoring** > **System** panel.
3. Select the canister or enclosure to be identified.
4. Select **Action** > **Identify** to control the identify LEDs for the component.

Alternatively, if a node canister is not online to the system, use the service assistant to control the identify LED.

1. Log in to the service assistant of the node canister to be identified.
2. Click **Identify** at the upper left of the page to control the identify LEDs.

Procedure: Identifying which Storwize V7000 Gen1 enclosure or canister to service

Before you begin service actions on an enclosure or canister, verify that you are working on the correct enclosure or canister.

About this task

Procedure

Use the following options to identify an enclosure. An enclosure is identified by its ID and serial number.

- The ID is shown on the LCD panel on the front left of the enclosure. The serial number is also found on the front left end cap of the enclosure and is repeated on the rear left flange of the enclosure. The enclosure ID is unique within a system. However, if you have more than one system, the same ID can be used within more than one system. The serial number is always unique.

  Note: To change the ID of an enclosure, select **Monitoring** > **System** in the management GUI. Use the dynamic image to display the enclosure to rename. Right-click the enclosure and select **Modify ID**. Use this option to set a unique ID on all your enclosures.

- Within an enclosure, a canister is identified by its slot location. Slot 1 is the upper canister. Slot 2 is the lower canister. A canister is uniquely identified by the enclosure that it is in and the slot location. The ID can be shown as E-C or E|C where *E* is the enclosure ID and *C* is the canister location. On the service assistant, the ID is known as the *Panel Name*.

  Note: When a node canister is added to a clustered system as a node, it is given a node name and a node ID. The default node name is node*N*, where *N* is an integer number. This number does not represent the slot location of the node. Similarly, the node ID does not indicate the slot location. To display the slot location, select **Monitoring** > **System** in the management GUI. Use the dynamic image to display the enclosure that contains the canister that you want to view. The enlarged view represents the actual location of the node. Hover over the node to display the node ID and other details. The service assistant home page also shows both the node name and the canister location. If you have only the node name, use these panels to determine the node canister location.

- Use the service assistant to identify a node canister by turning on the identify LED of the containing enclosure. This option is at the upper left of the service assistant page. It is a good practice to identify a node in this way before you perform any service action. Performing a service action on the wrong canister can lead to loss of access to data or loss of data.

Related tasks:

"Accessing the service assistant" on page 74
The service assistant is a web application that helps troubleshoot and resolve problems on a node canister in a control enclosure. The service assistant can also be accessed for the Storwize V7000 Gen2 by using the technician port.

Procedure: Checking the status of your system

Use this procedure to verify the status of objects in your system using the management GUI. If the status of the object is not online, view the alerts and run the recommended fix procedures.

About this task

Volumes normally show offline because another object is offline. A volume is offline if one of the MDisks that makes up the storage pool that it is in is offline. You do not see an alert that relates to the volume; instead, the alert relates to the

MDisk. Performing the fix procedures for the MDisk enables the volume to go online.

Procedure

Use the following management GUI functions to find a more detailed status:
- **Monitoring** > **System**
- **Pools** > **MDisks by Pools**
- **Volumes** > **Volumes**
- **Monitoring** > **Events**, and then use the filtering options to display alerts, messages, or event types.

Procedure: Getting node canister and system information using the service assistant

This procedure explains how to view information about the node canisters and system using the service assistant.

About this task

To obtain the information:
1. Log on to the service assistant.
2. View the information about the node canister to which you connected or the other node canisters in the enclosure. To change which node's information is shown, select the node in the **Change Node** table of the Home page.

The Home page shows a table of node errors that exist on the node canister and a table of node details for the current node. The node errors are shown in priority order.

The node details are divided into several sections. Each section has a tab. Examine the data that is reported in each tab for the information that you want.
- The Node tab shows general information about the node that includes the node state and whether it is a configuration node.
- The Hardware tab shows information about the hardware.
- The Access tab shows the management IP addresses and the service addresses for this node.
- The Location tab identifies the enclosure in which the node canister is located.
- The Ports tab shows information about the I/O ports.

Procedure: Getting node canister and system information by using a USB flash drive

You can view information about the node canister and system by using a USB flash drive.

About this task

Use any USB flash drive with a FAT32 file system on its first partition.
1. Ensure that the USB flash drive does not contain a file that is named `satask.txt` in the root directory.

If `satask.txt` does exist in the directory, the node attempts to run the command that is specified in the file. The information that is returned is appended to the `satask_result.html` file. Delete this file if you no longer want the previous output.

Procedure

1. Insert the USB flash drive in the USB port of the node from which you want to collect data. The node fault light-emitting diode (LED) flashes while information is collected and written to the USB flash drive.
2. Wait until the LED stops flashing before you remove the USB flash drive. Because the LED is a fault indicator, it might remain permanently on or off.
3. View the results in file `satask_result.html` in a web browser. The file contains the details and results of the command that was run and the status and the configuration information from the node canister.

Procedure: Understanding the system status using the LEDs

Use this procedure to determine the system status using the LED indicators on the system. The procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 70. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 71. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Procedure: Understanding the Storwize V7000 Gen2 system status from the LEDs

To determine the Storwize V7000 2076-524 system status using the LED indicators on a control enclosure, use this procedure.

About this task

To understand the status of the I/O port at the rear of a control enclosure, refer to the topic about Storwize V7000 2076-524 node canister ports and indicators that is linked at the end of this topic.

Status indicators on the front of a control enclosure are described in the topic about components in the front of the enclosure that is linked at the end of this topic.

A detailed view of the system state is provided in the Monitoring sections of the management GUI and by the service assistant. If neither the management GUI nor the service assistant is accessible, use this procedure to determine the system status using the LED indicators on the control enclosures.

The system status LEDs visible at the rear of each control enclosure can show one of several states, as described in Table 72.

Table 72. LED state descriptions used in the Storwize V7000 2076-524 enclosure

State description	Detail
Off	The LED is continuously not lit.
Flashing slowly	The LED turns on and off at a frequency of 1 Hz: It is on for 500 ms, then off for 500 ms, then repeats.
Flashing	The LED turns on and off at a frequency of 2 Hz: It is on for 250 ms, then off for 250 ms, then repeats.
Flashing fast	The LED turns on and off at a frequency of 4 Hz: It is on for 125 ms, then off for 125 ms, then repeats.
On	The LED is continuously lit.
Flashing	The LED is lit to indicate some activity, then turns off. The rate and duration that the LED is lit depends on the rate and duration of the activity.

Procedure

Complete the following steps to understand when a node canister is not participating in the system, and what remedial action to take.

1. Identify the control enclosures for the system you are troubleshooting.

   To understand which are the control enclosures, refer to the topic about identifying the Storwize V7000 2076-524 enclosure or canister to service.

2. Use Table 73 on page 108 to check the status of each power supply unit (PSU) in a control enclosure.

3. If at least one PSU is providing power to a control enclosure, use Table 74 on page 108 to check the status of each node canister in that enclosure.

Table 73. Understanding the power supply unit LEDs

LED state			Action
! Fault (amber)	**IN ≈ ac power (green)**	**DC === dc power (green)**	
On	(any)	(any)	Replace the power supply unit, as described in "Replacing a Storwize V7000 Gen2 power supply unit for a control enclosure" on page 188.
Off	On	On	The power supply is functioning normally, providing power to the enclosure.
		Off	Replace the power supply unit, as described in "Replacing a Storwize V7000 Gen2 power supply unit for a control enclosure" on page 188.
	Off	On	Replace the power supply unit, as described in "Replacing a Storwize V7000 Gen2 power supply unit for a control enclosure" on page 188.
		Off	The power supply is not receiving power from the power outlet through the power cord. To power on the system, connect the power supply to an outlet in use for the power cord and turn on the power outlet.
Note: The fourth status LED on the power supply unit is not used.			

Table 74. Understanding the node canister status LEDs

LED state			Description
⊘ Power (green)	**! Fault (amber)**	**⏻ System status (green)**	
Off	(any)	(any)	Power is not being provided by any power supply unit (PSU) in the enclosure, or power is coming from the battery in the canister. If at least one PSU is powering the enclosure, the canister or the enclosure midplane might be faulty. If both node canisters in an enclosure indicate this state at the same time, it is more likely that the enclosure midplane is faulty. Reseat any node canister in this state, as described in "Procedure: Reseating a Storwize V7000 Gen2 node canister" on page 128. If the condition persists, replace the node canister. If just one canister is affected, see "Replacing a Storwize V7000 Gen2 node canister" on page 177. If both canisters are affected, see "Replacing a Storwize V7000 Gen2 control enclosure midplane assembly" on page 235.
Flashing slowly			Power is available, but the canister is powered off. (The main CPU is not running.) Restart the node canister, as described in "Procedure: Reseating a Storwize V7000 Gen2 node canister" on page 128.
Flashing fast			The node canister is doing a self test during start-up. Wait for the canister to complete its start-up sequence.

Table 74. Understanding the node canister status LEDs (continued)

LED state			Description
⊕ Power (green)	! Fault (amber)	⏻ System status (green)	
On	Off	Off	The node canister is in standby mode. (The Storwize V7000 software is not running.) It is safe to remove or reseat the canister. Restart the node canister, as described in "Procedure: Reseating a Storwize V7000 Gen2 node canister" on page 128.
		Flashing	The node canister is in candidate state. The node is not doing I/O in the system. Unless indicated by the battery status LED, it is safe to remove or reseat the canister. See Table 75 on page 111. If the node state is candidate and the system is running on the other node canister, the candidate node is automatically added to the system. If both node canisters are in candidate state, determine whether you must recover the system or create a new system.
		Flashing fast	The node is doing an emergency shutdown operation, using the battery for power, after detecting a loss of power from the power supply units. Wait for the emergency shutdown operation to complete. If the partner node has also done an emergency shutdown operation, there was most likely a loss of input power to both enclosure power supply units and the system restarts when the input power is restored. Otherwise, there might be a fault the node canister, enclosure midplane, or power supply units.
		On	The Storwize V7000 software is running, and the node canister is participating in the system. The canister must not be removed. This is the normal operational LED state.
On	Flashing	Off	The node canister is in standby mode. (The Storwize V7000 software is not running.) It is safe to remove or reseat the canister.
		Flashing	The Identify function for this canister has been activated. To determine if it is safe to remove the canister, use the management GUI or service assistant to turn off the Identify function, then check the node canister status LEDs again.

Table 74. Understanding the node canister status LEDs (continued)

LED state			Description
⊕ Power (green)	! Fault (amber)	⏻ System status (green)	
On	Flashing	Flashing fast	The Identify function for this canister has been activated. The node is doing an emergency shutdown operation, using the battery for power, after detecting a loss of power from the power supply units. Wait for the emergency shutdown operation to complete. If the partner node has also done an emergency shutdown operation, there was most likely a loss of input power to both enclosure power supply units and the system restarts when the input power is restored. Otherwise, there might be a fault in the power supply units or the node canister.
On	Flashing	On	The Identify function for this canister was activated. To determine if it is safe to remove the canister, use the management GUI or service assistant to turn off the Identify function, then check the node canister status LEDs, again.
On	On	Off	The Storwize V7000 software is not running. The BIOS might have detected a fault. It is safe to remove or reseat the canister. Try reseating the canister, as described in "Procedure: Reseating a Storwize V7000 Gen2 node canister" on page 128. If the canister still shows this fault, replace the node canister, as described in "Replacing a Storwize V7000 Gen2 node canister" on page 177.
On	On	Flashing	The node is in service state. It is safe to remove or reseat the canister. See Table 75 on page 111 to determine whether the battery charge is insufficient. Use the service assistant to identify any node errors or to determine what to do, see "Procedure: Fixing node errors" on page 121. If the node is able to communicate with the configuration node, there might also be an error alert in the system event log, in which case you should run the associated fix procedure.
On	On	Flashing fast	The node is in service state. There is a code update in progress. The canister must not be removed. No action is required until the code update completes.

Table 74. Understanding the node canister status LEDs (continued)

LED state			Description
⊕ Power (green)	! Fault (amber)	⏻ System status (green)	
On	On	On	The Storwize V7000 software is running but there might be an error alert in the event log, such as error code 550. The canister must not be removed. If possible, go to the management GUI and run the fix procedure for the error alerts listed there. If this is not possible, refer to the service actions for node error 550.

Table 75. Understanding the node canister battery status LEDs

LED state		Description
! Fault (amber)	▱ Status (green)	
On	On	A battery fault exists that caused the node canister to do an emergency shutdown operation (storing its cache data to the system disk). The node canister is in service state or is going into service state. The service assistant shows a node error. Replace the battery in the node canister, as described in "Replacing the battery in a Storwize V7000 Gen2 node canister" on page 198.
On	Off	A battery fault exists that caused the node canister to shut down to service state. There was insufficient power in the battery or the supply for the node canister to do an emergency shutdown operation. The service assistant shows a node error. Replace the battery in the node canister, as described in "Replacing the battery in a Storwize V7000 Gen2 node canister" on page 198.
Off	Flashing fast	The battery is charging and has insufficient charge to protect against a single ac loss. The node is held in service state with a fatal node error until the battery has charged.
Off	Flashing	The battery has sufficient charge to complete one emergency shutdown operation. This state does not prevent the canister from participating in the system. The battery continues to charge until it is able to complete two emergency shutdown operations. No action is necessary.
Off	On	The battery is optimally charged, such that the node canister is able to complete two emergency shutdown operations. This state does not prevent the canister from participating in the system. The canister continues to manage the battery charge. No action is necessary.

4. Repeat steps 2 on page 107 and 3 on page 107 for each control enclosure in the system.

Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs

Use this procedure to determine the Storwize V7000 Gen1 system status using the LED indicators on the system.

About this task

The LEDs provide a general idea of the system status. You can obtain more detail from the management GUI and the service assistant. Examine the LEDs when you are not able to access the management GUI or the service assistant, or when the system is not showing any information about a device.

The procedure shows the status for the enclosure chassis, power supply units and batteries, and canisters. It does not show the status for the drives.

The first step is to determine the state of the control enclosure, which includes its power supply units, batteries, and node canisters. Your control enclosure is operational if you can manage the system using the management GUI. You might also want to view the status of the individual power supply units, batteries, or node canisters.

Find the control enclosure for the system that you are troubleshooting. There is one control enclosure in a system. If you are unsure which one is the control enclosure, go to "Procedure: Identifying which Storwize V7000 Gen1 enclosure or canister to service" on page 103.

Procedure

1. Use the state of the ac power failure, power supply OK, fan failure, and dc power failure LEDs on each power supply unit in the enclosure to determine if there is power to the system, or if there are power problems. Figure 37 on page 113 shows the LEDs on the power supply unit for the 2076-112 or 2076-124. The LEDs on the power supply units for the 2076-312 and 2076-324 are similar, but they are not shown here.
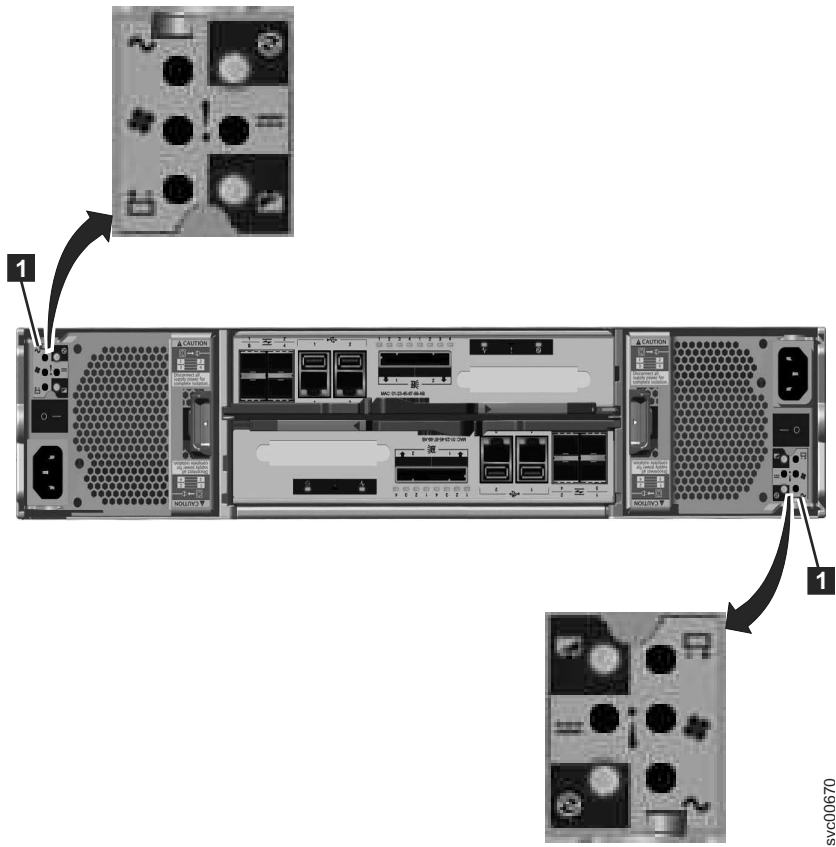
Figure 37. LEDs on the power supply units of the control enclosure

svc00670

Table 76. Power supply unit LEDs

Power supply OK	AC failure	Fan failure	DC failure	Status	Action
On	On	On	On	Communication failure between the power supply unit and the enclosure chassis	Replace the power supply unit. If failure is still present, replace the enclosure chassis.
Off	Off	Off	Off	No ac power to the enclosure.	Turn on power.
Off	Off	Off	On	The ac power is on but power supply unit is not seated correctly in the enclosure.	Seat the power supply unit correctly in the enclosure.

Table 76. Power supply unit LEDs (continued)

Power supply OK ⏻	AC failure ∿	Fan failure ✿	DC failure ⚏	Status	Action
Off	On	Off	On	No ac power to this power supply	1. Check that the switch on the power supply unit is on. 2. Check that the ac power is on. 3. Reseat and replace the power cable.
On	Off	Off	Off	Power supply is on and operational.	No actions
Off	Off	On	Off	Fan failure	Replace the power supply unit.
Off	On	On	On	Communication failure and power supply problem	Replace the power supply unit. If replacing the power supply unit does not fix the problem, replace the enclosure chassis.
Flashing	X	X	X	No canister is operational.	Both canisters are either off or not seated correctly. Turn off the switch on both power supply units and then turn on both switches. If this action does not resolve the problem, remove both canisters slightly and then push the canisters back in.
Off	Flashing	Flashing	Flashing	Firmware is downloading.	No actions. Do not remove ac power. **Note:** In this case, if there is a battery in a power supply unit, its LEDs also flash.

2. At least one power supply in the enclosure must indicate `Power supply OK` or `Power supply firmware downloading` for the node canisters to operate. For this situation, review the three canister status LEDs on each of the node canisters. Start with the power LED.

Table 77. Power LEDs

Power LED status ⏻	Description
Off	There is no power to the canister. Try reseating the canister. Go to "Procedure: Reseating a Storwize V7000 Gen1 node canister" on page 129. If the state persists, follow the hardware replacement procedures for the parts in the following order: node canister, enclosure chassis.

Table 77. Power LEDs (continued)

Power LED status ⏻	Description
Slow flashing (1 Hz)	Power is available, but the canister is in standby mode. Try to start the node canister by reseating it. Go to "Procedure: Reseating a Storwize V7000 Gen1 node canister" on page 129.
Fast flashing (2 Hz)	The canister is running its power-on self-test (POST). Wait for the test to complete. If the canister remains in this state for more than 10 minutes, try reseating the canister. Go to "Procedure: Reseating a Storwize V7000 Gen1 node canister" on page 129. If the state persists, follow the hardware replacement procedure for the node canister.

Figure 38 shows the LEDs on the node canister.
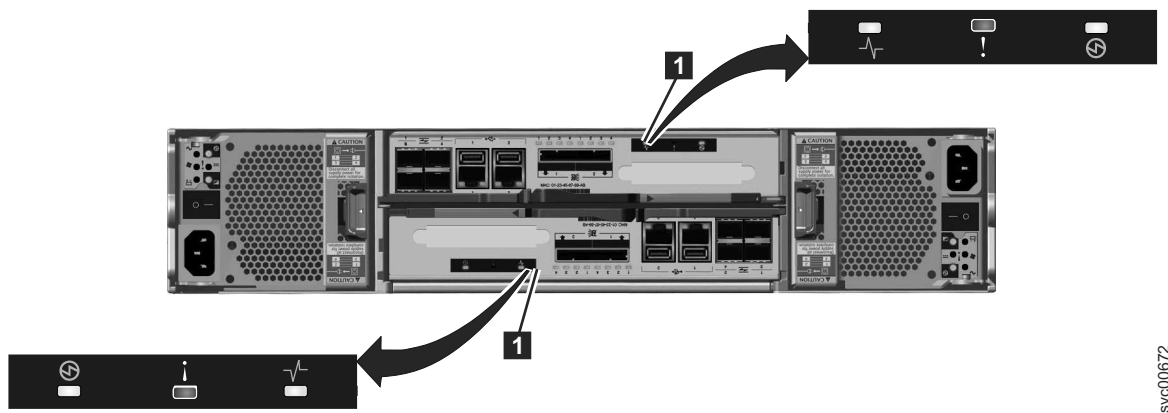
Figure 38. LEDs on the node canisters

3. If the power LED is on, consider the states of the clustered-system status and fault LEDs.

Table 78. System status and fault LEDs

System status LED ⎍	Fault LED ❗	Status ⏻	Action
Off	Off	Code is not active.	• Follow procedures for reviewing power LEDs. • If the power LEDs show green, reseat the node canister. See "Procedure: Reseating a Storwize V7000 Gen1 node canister" on page 129. If the LED status does not change, see "Replacing a Storwize V7000 Gen1 node canister" on page 178.
Off	On	Code is not active. The BIOS or the service processor has detected a hardware fault.	Follow the hardware replacement procedures for the node canister.
On	Off	Code is active. Node state is active.	No action. The node canister is part of a system and can be managed by the management GUI.

Table 78. System status and fault LEDs (continued)

System status LED ⌁	Fault LED ❗	Status ⓘ	Action
On	On	Code is active and is in starting state. However, it does not have enough resources to form the system.	The node canister cannot become active in a system. There are no detected problems on the node canister itself. However, it cannot connect to enough resources to safely form a system. Follow the procedure to fix the node errors. Go to "Procedure: Fixing node errors" on page 121.
Flashing	Off	Code is active. Node state is candidate.	Create a system on the node canister, or add the node canister to the system. If the other node canister in the enclosure is in active state, it automatically adds this node canister into the system. A node canister in this state can be managed using the service assistant.
Flashing	On	Code is active. Node state is service.	The node canister cannot become active in a system. Several problems can exist: hardware problem, a problem with the environment or its location, or problems with the code or data on the canister. Follow the procedure to fix the node errors. Go to "Procedure: Fixing node errors" on page 121.
Any	Flashing	The node canister is being identified so that you can locate it.	The fix procedures in the management GUI might have identified the component because it requires servicing. Continue to follow the fix procedures. The service assistant has a function to identify node canisters. If the identification LED is on in error, use the service assistant node actions to turn off the LED.

Results

To review the status of the control enclosure batteries, see Table 79.

Table 79. Control enclosure battery LEDs

Battery Good 🔋	Battery Fault 🔋	Description	Action
On	Off	Battery is good and fully charged.	None
Flashing	off	Battery is good but not fully charged. The battery is either charging or a maintenance discharge is in process.	None
Off	On	Nonrecoverable battery fault.	Replace the battery. If replacing the battery does not fix the issue, replace the power supply unit.

Battery Good	Battery Fault	Description	Action
Off	Flashing	Recoverable battery fault.	None
Flashing	Flashing	The battery cannot be used because the firmware for the power supply unit is being downloaded.	None

Procedure: Finding the status of Ethernet connections

Use this procedure to find the status of Ethernet connections when you cannot connect. This procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 80. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 81. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Procedure: Finding the status of Storwize V7000 Gen2 Ethernet connections

Use this procedure to find the status of Ethernet connections in the Storwize V7000 Gen2 system when you cannot connect.

About this task

Ethernet port 1 of each node canister must be connected to an active port on your Ethernet network. You can determine the status of the Ethernet ports by using any of the following methods:

- Connect a personal computer directly to the node by following "Procedure: Accessing the Storwize V7000 Gen2 service assistant from the technician port" on page 126t. In the service assistant, the status, speed, and MAC address for each port are shown in the **Ports** tab of the **Node Details** table. Any node errors are shown in the **Node Errors** table.
- Use a USB flash drive to obtain node configuration information (see "Procedure: Getting node canister and system information by using a USB flash drive" on page 105). The results file contains the status, speed, and MAC address for each port; whether the node is the configuration node and if any node errors are being reported.
- Examine the LEDs of the Ethernet ports. The Ethernet ports on the left end of the rear of each node canister are 1 Gbps Ethernet ports. For these ports, the link state LED is ON if the link is connected.

Procedure

1. Verify that each end of the cable is securely connected.
2. Verify that the port on the Ethernet switch or hub is configured correctly.
3. Connect the cable to a different port on your Ethernet network.
4. Replace the Ethernet cable.
5. Review any node errors that are reported in the service assistant or on the USB flash drive. Follow fixing node errors for each node error that is reported.
6. Follow the hardware replacement procedure for a node canister.

Procedure: Finding the status of Storwize V7000 Gen1 Ethernet connections

This procedure explains how to find the status of the Ethernet connections when you cannot connect.

About this task

Ethernet port 1 must be connected to an active port on your Ethernet network. Determine the state of the Ethernet LEDs by using the following methods:

- If the node software is active on the node, use the USB flash drive to obtain the most comprehensive information for the node status. Go to "Procedure: Getting node canister and system information by using a USB flash drive" on page 105. The status, speed, and MAC address are returned for each port. Information is returned that identifies whether the node is the configuration node and whether any node errors were reported.
- Examine the LEDs of the Ethernet ports. For the status of the LEDs, go to "Ethernet ports and indicators" on page 30.

Procedure

If your link is not connected, complete the following actions to check the port status each time until it is corrected or connected.

1. Verify that each end of the cable is securely connected.
2. Verify that the port on the Ethernet switch or hub is configured correctly.
3. Connect the cable to a different port on your Ethernet network.
4. If the status is obtained using the USB flash drive, review all the node errors that are reported.
5. Replace the Ethernet cable.
6. If replacing the Ethernet cable does not fix the problem, contact IBM support.

Procedure: Finding the status of Storwize V7000 Gen2 SAS connections

Find the status of the SAS connections between Storwize V7000 Gen2 canisters in different enclosures.

About this task

Ensure that the Storwize V7000 machine code is active on the node before you begin this procedure. To determine if the machine code is active, see "Procedure: Understanding the Storwize V7000 Gen2 system status from the LEDs" on page 107.

Procedure

Determine the state of the SAS ports by using one of the following methods:

- Go to **Monitoring** > **System** in the management GUI. Use the dynamic image to display the rear of the system. Hover over each of the SAS ports on the canisters to display the status. A port with offline status indicates that its link is not connected.
  - It is normal for port 2 of the canisters in the expansion enclosure at the end of a SAS chain to be offline.
  - If no expansion enclosures are connected to a system, it is normal for port 4 of each canister in the control enclosure to be offline.

  Attention: The system can identify some SAS cabling errors and log an event to alert you to the error. Go to the **Monitoring** > **Events** page of the management GUI and identify alerts concerning hardware errors and SAS cabling errors. Run fix procedures in the recommended order.
- Determine the meaning of the LEDs of the SAS ports, as described in "Expansion canister ports and indicators" on page 34.
  - If the link LED is off, the link is not connected.
  - If the fault LED is on, the link is partially operational, with reduced performance. Consider the state of any other link between the two enclosures before servicing this link.
- To connect a link that is not connected, complete the following actions while checking the link status after each step until the link is connected:
  1. Ensure that both ends of each SAS cable are correctly inserted into their correct ports, as described in "Problem: Storwize V7000 Gen2 SAS cabling not valid" on page 96.

2. Replace the SAS cable.
3. Replace the expansion canister at one end of the connection.
4. Replace the canister at the other end of the connection. If it is a node canister, see "Replacing a Storwize V7000 Gen2 node canister" on page 177.

Procedure: Removing system data from a node canister

You can safely remove system information from a node canister, if you follow the proper guidelines and procedure. The information that is removed includes configuration data, cache data, and location data.

About this task

Attention: Do not remove the system data from a node canister unless instructed to do so by a service procedure. Do not use this procedure to remove the system data from the only online node canister in a system. If the system data is removed or lost from all node canisters in the system, the system is effectively deleted. Attempting a system recovery procedure to restore a deleted system is not guaranteed to recover all of your volumes.

Procedure

1. Log in to the service assistant of the node canister.
2. Use the service assistant node action to hold the node canister in service state.
3. Click **Manage System**, then click **Remove system data** to remove the system data from the node canister.

Results

The node canister restarts in service state.

What to do next

When you want the node canister to be active again, use the service assistant to leave service state. The node canister moves to candidate state, and can be added to the system. If the partner node canister is already active, the candidate node canister is added automatically.

Procedure: Deleting a system completely

You might need to completely remove all system information. When the procedure is finished, the system operates like a new installation. No data is retained.

About this task

Attention: This procedure makes all the volume data that you have on your system inaccessible. You cannot recover the data. This procedure affects all volumes that are managed by your system.

Do not continue unless you are certain that you want to remove all the volume data and configuration data from your system. This procedure is not used as part of any recovery action.

There are two stages to this procedure. First, the node canisters are reset. Second, the enclosure data is reset.

Procedure

1. Start the service assistant on one of the node canisters.
2. Use the service assistant node action to hold the node in service state.
3. Use the **Manage System** option to remove the system data from the node.
4. Repeat steps 1 through 3 on the second node canister in the enclosure.
5. On one node, open the service assistant **Configure Enclosure** and select the **Reset System ID** option. This action causes the system to reset.

Procedure: Fixing node errors

To fix node errors that are detected by node canisters in your system, use this procedure.

About this task

Node errors are reported in the service assistant when a node detects erroneous conditions in a node canister.

Procedure

1. Use the service assistant to obtain (and better understand) node canister and system information about the state of each node.
2. If possible, log into the management GUI and use the monitoring page to run the recommended fix procedure.
   a. Follow the fix procedure instructions to completion.
   b. Repeat this step for each subsequent recommended fix procedure.
3. If it is not possible to access the management GUI, or no recommended actions are listed, follow the identified user response for each reported node error.

Procedure: Changing the service IP address of a node canister

This procedure identifies many methods that you can use to change the service IP address of a node canister.

About this task

When you change an IPv4 address, you change the IP address, the subnet, mask, and gateway. When you change an IPv6 address, you change the IP address, prefix, and gateway.

Which method to use depends on the status of the system and the other node canisters in the system. Follow the methods in the order that is shown until you are successful in setting the IP address to the required value.

You can set an IPv4 address, an IPv6 address, or both, as the service address of a node. Enter the required address correctly. If you set the address to 0.0.0.0 or 0000:0000:0000:0000:0000:0000:0000, you disable the access to the port on that protocol.

Procedure

Change the service IP address.

- Use the control enclosure management GUI when the system is operating and the system is able to connect to the node with the service IP address that you want to change.
  1. Select **Settings** > **Network** from the navigation.
  2. Select **Service IP Addresses**.
  3. Complete the panel. Be sure to select the correct node to configure.
- Use the service assistant when you can connect to the service assistant on either the node canister that you want to configure or on a node canister that can connect to the node canister that you want to configure:
  1. Make the node canister that you want to configure the current node.
  2. Select **Change Service IP** from the menu.
  3. Complete the panel.
- Use one of the following procedures if you cannot connect to the node canister from another node:
  - Use the initialization tool to write the correct command file to the USB flash drive.

    For Storwize V7000 Gen2 use "Procedure: Initializing the Storwize V7000 Gen2 system using the technician port" on page 124.Go to "Using the initialization tool on a Storwize V7000 Gen1 control enclosure" on page 81.
  - Use a text editor to create the command file on the USB flash drive. Go to "Using a USB flash drive" on page 82.

Procedure: Initializing a clustered system with a USB flash drive without using the initialization tool

Use this procedure to initialize a clustered system by using a USB flash drive when you do not have a workstation to run the initialization tool or you do not have a copy of the tool.

About this task

In these situations, you must manually create an `satask.txt` file on a USB flash drive to initialize your clustered system. Use the USB flash drive that was supplied with your system or any USB flash drive that is formatted with a FAT32 file system on its first partition. (For a complete list of commands you can use in a `satask.txt` file, see "satask.txt commands" on page 75.)

Note: The USB flash drive must include only the `satask.txt` file on the formatted first partition of the drive to initialize the system. If other files are included on the formatted first partition of the drive, it might cause the initialization to fail.

Procedure

1. Open a file editor that can create ASCII text files.
2. Create a file called `satask.txt`.
3. Add a single line of command text to the file.

   If you are creating a clustered system with an IPv4 address, the command line is like the following string:

   ```
   satask mkcluster -clusterip aaa.aaa.aaa.aaa
       -gw ggg.ggg.ggg.ggg -mask mmm.mmm.mmm.mmm
   ```

where you must replace *aaa.aaa.aaa.aaa* with the management IP address, *ggg.ggg.ggg.ggg* with the network gateway address, and *mmm.mmm.mmm.mmm* with the subnet mask address.

If you are creating a clustered system with an IPv6 address, the command line is like the following string:

```
satask mkcluster -clusterip_6 aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa
    -gw_6 gggg:gggg:gggg:gggg:gggg:gggg:gggg:gggg -prefix_6 pp
```

where you must replace *aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa* with the management IPv6 address, *gggg:gggg:gggg:gggg:gggg:gggg:gggg:gggg* with the network gateway IPv6 address, and *pp* with the prefix value.

For other command options, see "Create system command" on page 78.

4. Save the file to a USB flash drive.
5. Plug the USB flash drive into a USB port on a control canister.
6. The system detects the USB flash drive, reads the `satask.txt` file, runs the command, and writes the results to the USB flash drive. The `satask.txt` file is deleted after the command is run.
7. Wait for the fault LED on the node canister to stop flashing before removing the USB flash drive.
8. Remove the USB flash drive and insert it into your workstation to view the results.
9. Use a web browser to view the results file, `satask_result.html`.

   Check that there were no errors returned by the command. If there is insufficient battery charge to protect the system, the clustered system creates successfully, but it does not start immediately. In the results, look for the `time_to_charge` field for the battery. The results provide an estimate of the time, in minutes, before the system can start. If the time is not 0, wait for the required time. Check that the node canister that you inserted the USB flash drive into has its clustered-state LED on permanently. For additional information, see "Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs" on page 112.
10. If the initialization was successful and the batteries are sufficiently charged, point a supported browser to the management IP address that you specified to start the management GUI. You see the management GUI logon panel.
11. Log on as superuser. Use `passw0rd` for the password.
12. Follow the on-screen instructions.

Results

For more information about using the USB flash drive, see "USB flash drive and the Storwize V7000 Gen1 Initialization tool interface" on page 81.

Procedure: Initializing a clustered system by using the service assistant

Use this procedure to initialize a clustered system by using the service assistant. This procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 82. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 83. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Procedure: Initializing the Storwize V7000 Gen2 system using the technician port

To initialize a new Storwize V7000 Gen2 system, you must connect a personal computer to the technician port on the rear of a node canister and run the initialization tool.

Before you begin

You must have the following items:
- A personal computer with an Ethernet port that supports Dynamic Host Configuration Protocol (DHCP)
- A supported browser that is installed on the personal computer
- An Ethernet cable to connect the personal computer to the technician port

Procedure

To initialize the system, complete the following steps.
1. Ensure that the system is powered on.

2. Configure an Ethernet port on the personal computer to enable Dynamic Host Configuration Protocol (DHCP) configuration of its IP address and DNS settings.

3. Locate the Ethernet port that is labeled **T** on the rear of a node canister. This is the Technician port. Connect an Ethernet cable between the port of the personal computer that is configured in step 2 and the technician port. A few moments after the connection is made, the node uses DHCP to configure IP and DNS settings of the personal computer.

4. After the Ethernet port of the personal computer is connected, open a supported browser and browse to address `http://install`. The browser automatically opens the initialization tool.

5. Follow the instructions that are presented by the initialization tool to configure the system with a management IP address.

6. After you complete the initialization process, disconnect the cable between the personal computer and the technician port.

What to do next

The system can now be reached by opening a supported web browser and pointing it to `http://management_IP_address`.

Procedure: Initializing a Storwize V7000 Gen1 system using the service assistant

To initialize a Storwize V7000 Gen1 system using the service assistant rather than the USB flash drive, use this procedure.

About this task

Note: The service assistant gives you the option to create a clustered system only if the node state is candidate.

Procedure

To initialize a clustered system using the service assistant, complete the following steps.

1. Point your web browser to the service assistant address of a node canister. It is best to use the node canister in slot 1; when viewed from the rear of the control enclosure, the left node canister. The default service address for this canister is *192.168.70.121/service*.

2. Log on with the superuser password. The default password is `passw0rd`. If you cannot connect, see "Problem: Cannot connect to the service assistant" on page 94.

3. Select **Manage System**.

4. Enter the system name and the management IP address.

5. Click **Create System**.

6. Point a supported browser to the management IP address that you specified to start the management GUI. The management GUI logon panel is displayed.

7. Log on as superuser. Use `passw0rd` for the password.

8. Follow the on-screen instructions.

Results

Attention: Without a USB flash drive to service the system, it is not possible to reset the superuser password or to change the system IP addresses in the event of a fault that prevents access to the management interface. It is essential that you take steps to record this information for use in the event of a failure.

Procedure: Accessing the Storwize V7000 Gen2 service assistant from the technician port

If a Storwize V7000 Gen2 node canister is inaccessible through your administrative network, use this procedure to connect a personal computer directly to the node canister to access the service assistant.

About this task

This procedure starts the initialization tool if the enclosure is not part of a system because the following conditions are true:
- The node canister is in candidate state.
- No system details are configured.
- The partner node is not in active state.

Otherwise, this procedure starts the service assistant.

Procedure

To connect a personal computer directly to the node canister, complete the following steps:

1. Configure DHCP on the Ethernet port of the personal computer to connect to the node canister.

   If the personal computer cannot support DHCP, configure static IPv4 address `192.168.0.2` on the port.

2. Connect an Ethernet cable between the port on the personal computer and the technician port.

   The technician port is labeled **T** on the rear of the node canister.

3. Open a supported web browser on the personal computer and browse to this URL:

   `http://192.168.0.1`

4. Complete the appropriate procedure.
   - If the initialization tool opens, complete the initialization as described in the installation procedure.
   - If the service assistant dialog opens, use it to service the node canister.

5. Log out of the service assistant.
6. Disconnect the Ethernet cable from the technician port.

Procedure: Accessing aStorwize V7000 Gen1 canister by using a directly attached Ethernet cable

If you need to use a direct Ethernet connection to attach a personal computer to a Storwize V7000 Gen1 node canister to run the service assistant or to use the service CLI, use this procedure.

About this task

Follow this procedure if you are not authorized to use a USB flash drive in your data center and when the service address of your nodes cannot be accessed over your Ethernet network. This situation might occur for a new installation where the default service IP addresses cannot be accessed on your network.

The default service addresses are listed in "Problem: Cannot connect to the service assistant" on page 94.

Note: Do not attempt to use a directly attached Ethernet cable to a canister that is active in a clustered system. You might disrupt access from host applications or the management GUI. If the node is active, go to **Settings** > **Network** in the management GUI to set the service IP address to one that is accessible on the network.

Procedure

Complete the following steps to access a canister by using a directly attached Ethernet cable.

1. Connect one end of an Ethernet cable to Ethernet port 1 of the node canister.

   Note: A cross-over Ethernet cable is not required.
2. Connect the other end of the Ethernet cable directly to the Ethernet port on a personal computer that has a web browser installed.
3. Get the service IP address of the node canister attached at step 1. If the service IP address is unknown, refer to "Problem: node canister service IP address unknown" on page 93.
4. Use the operating system tools on the computer to set the IP address and subnet mask of the Ethernet port that is used in step 2. Set them to the same subnet of the node canister service IP address.
5. Point the web browser to the service IP address for the node canister.
6. Log on with the superuser password. The default password is `passw0rd`.
7. Set the service address of the node canister to one that can be accessed on the network as soon as possible.
8. Wait for the action to complete.
9. Disconnect your personal computer.
10. Reconnect the node canister to the Ethernet network.

Procedure: Reseating a node canister

Use this procedure to reseat a node canister. The procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 84. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives

Table 84. Storwize V7000 Gen1 model numbers (continued)

Machine type/model	Description
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 85. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Procedure: Reseating a Storwize V7000 Gen2 node canister

Use this procedure to reseat a Storwize V7000 Gen2 node canister that is in service state or because of a service action that requires that the node canister be reseated.

About this task

Verify that you are reseating the correct node canister and that you use the correct canister handle for the node that you are reseating. A handle for each node canister is located above the canister.

Procedure

1. Verify the clustered-system status LED on the node canister. If it is permanently on, the node is active. If the node is active, no reseating is required.
2. Verify that you selected the correct node canister and verify why you are reseating it. Go to "Procedure: Identifying which Storwize V7000 Gen1 enclosure or canister to service" on page 103.
3. Rotate the handle release trigger.
4. Pull out the handle to its full extension.
5. Grasp the canister to pull it out 2 or 3 inches.
6. Push the canister to return it into the slot until the handle starts to move.
7. Finish inserting the canister by closing the handle until the locking catch clicks into place.
8. Verify that the cables were not displaced.

9. Verify that the LEDs are on.

Procedure: Reseating a Storwize V7000 Gen1 node canister

Use this procedure to reseat a canister that is in service state or because a service action has directed you.

About this task

Verify that you are reseating the correct node canister and that you use the correct canister handle for the node that you are reseating. Handles for the node canisters are located next to each other. The handle on the right operates the upper canister. The handle on the left operates the lower canister.

Procedure

1. Verify the clustered-system status LED on the node canister. If it is permanently on, the node is active. If the node is active, no reseating is required.
2. Verify that you have selected the correct node canister and verify why you are reseating it. Go to "Procedure: Identifying which Storwize V7000 Gen1 enclosure or canister to service" on page 103.

   If you reseat a node that is active, it cannot store its state data and cannot restart without other service actions.

   If the other node canister in the enclosure is not active, reseating the node canister while it is active results in loss of the data on your volumes and the system is unavailable to hosts.
3. Grasp the handle between the thumb and forefinger.
4. Squeeze them together to release the handle.
5. Pull out the handle to its full extension.
6. Grasp the canister and pull it out 2 or 3 inches.
7. Push the canister back into the slot until the handle starts to move.
8. Finish inserting the canister by closing the handle until the locking catch clicks into place.
9. Verify that the cables were not displaced.
10. Verify that the LEDs are on.

Results

Procedure: Removing a Storwize V7000 Gen2 node canister

Follow this procedure to remove a node canister.

About this task

Attention: Before a node canister can be removed it must be powered off or in service state; otherwise, loss of data or loss of access to data can result.

If a node canister was recently removed from the system and then readded, ensure that the canister is online for at least 25 minutes before you remove its partner canister. This delay allows multipath drivers to fail over to the online canister when the partner canister is removed.

Procedure

1. Read the safety information referred to in "Preparing to remove and replace parts" on page 177.

2. Follow the steps in "Procedure: Powering off a Storwize V7000 Gen2 node canister" on page 136

3. Use the LEDs on the canister to confirm that it is safe to remove the canister from the enclosure, as described in "Procedure: Understanding the Storwize V7000 Gen2 system status from the LEDs" on page 107.

4. Record which data cables are plugged into the specific ports on the rear of the node canister. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.

5. Disconnect the data cables that are connected to the node canister.

6. On the canister, unlatch the release lever and pull it open (see Figure 39). The canister moves out of the slot approximately 0.6 cm (0.25 inch). Be careful that you do not inadvertently disturb or remove any cables that are connected to other components of the system.
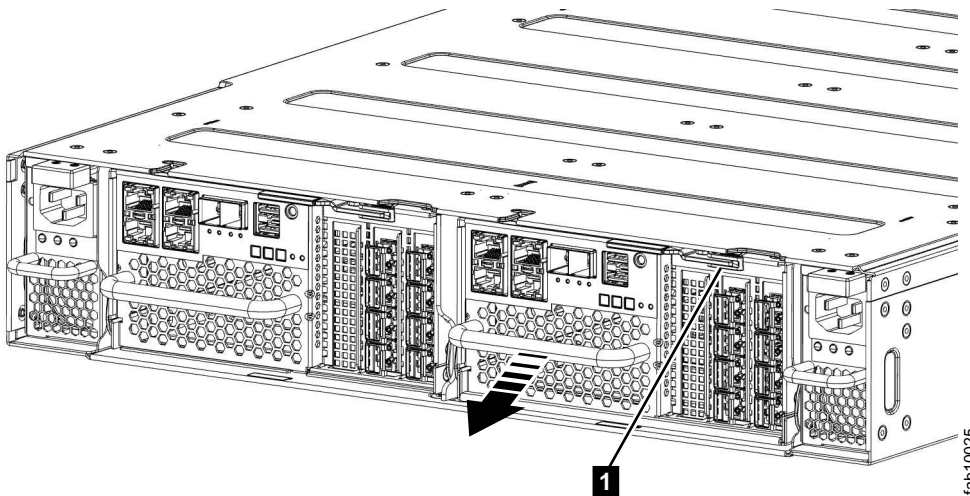
Figure 39. Removing a node canister

> **Note:** The number scale that is etched along the top and sides of the canister indicates how much of the canister is being supported by the enclosure. When you remove the canister, ensure that you are supporting the full weight of the canister before you reach "1" on the scale.

7. As you pay attention to the number scale, slide the canister out of the slot.

Procedure: Powering off your system

You must power off your Storwize V7000 system in order to service it, or to permit other maintenance actions in your data center. This procedure differs depending on the generation of your control enclosure model.

About this task

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 86. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 87. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Procedure: Powering off your Storwize V7000 Gen1 system

You must power off your Storwize V7000 system in order to service it, or to permit other maintenance actions in your data center.

About this task

To power off your Storwize V7000 system, complete the following steps:

1. Stop hosts.
2. Shut down the system by using the management GUI. Select **Monitoring** > **System**. From the **Actions** menu, select **Power Off**.
3. Wait for the power LED on both node canisters in all control enclosures to start flashing, which indicates that the shutdown operation has completed.

   The following figure shows the LEDs on the node canisters. The power LED is the LED on the left when the canister is top-side up.
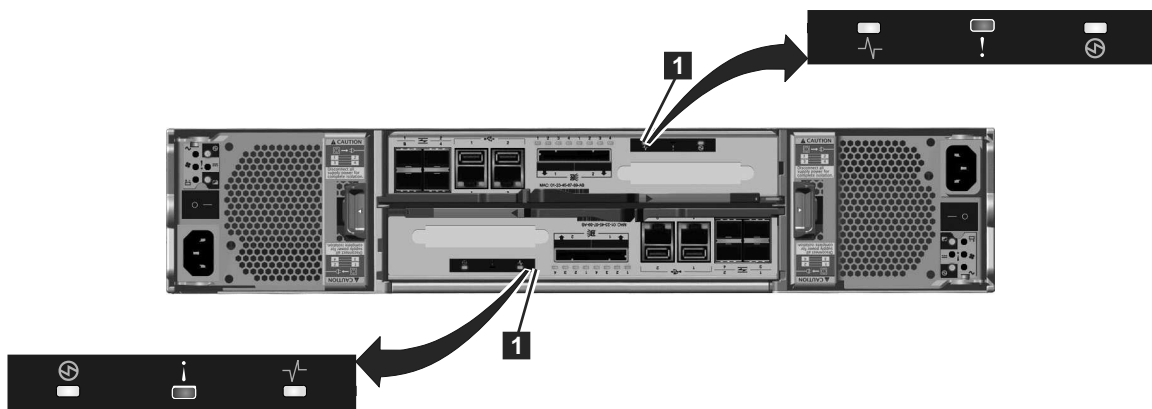
Figure 40. LEDs on the node canisters

4. Using the power switches, power off the control enclosures.
5. Using the power switches, power off the expansion enclosures.
6. (Optional) Shut down external storage systems.
7. (Optional) Shut down Fibre Channel switches.

Procedure: Powering off your Storwize V7000 Gen2 system

You must power off your Storwize V7000 Gen2 system to service it, or to allow for other maintenance actions in your data center.

Procedure

To power off your Storwize V7000 system, complete the following steps:

1. Stop all host I/O to volumes on the system.
2. Shut down the system by using the management GUI. Click **Monitoring** > **System**. From the **Actions** menu, select **Power off**.
3. Wait for the power LEDs on all node canisters in all control enclosures to blink at 1 Hz, indicating that the shutdown operation completed, as shown in Figure 41.
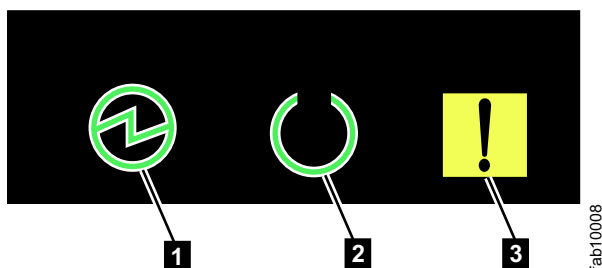
Figure 41. Power LEDs on a node canister

1 Power
2 Status
3 Fault

4. Disconnect the power cords from both power supplies in each control enclosure.
5. Disconnect the power cords from both power supplies in each expansion enclosure.

Procedure: Powering on the Storwize V7000 Gen2 system

After installing all hardware components, you must power on the Storwize V7000 Gen2 system and check its status.

About this task

Attention: Do not power on the system with any open bays or slots.
- Every unused drive bay must be occupied by a filler panel.
- Filler plates must be installed in all empty host interface adapter slots.

Open bays or slots disrupt the internal air flow, causing the drives to receive insufficient cooling.

Procedure

To power on the system, complete the following steps.
1. Power on all expansion enclosures by connecting both power supply units of the enclosure to their power sources, using the supplied power cables. If the power sources have circuit breakers or switches, ensure that they are turned on. The enclosure does not have power switches. Repeat this step for each expansion enclosure in the system.

   Note: Each enclosure has two power supply units. To provide power failure redundancy, connect the two power cords to separate power circuits.
2. Check the LEDs on each expansion canister, as displayed in Figure 42.
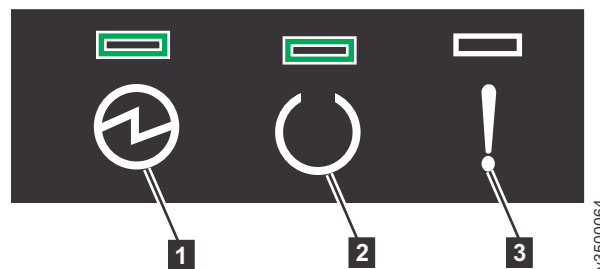
Figure 42. Expansion canister LEDs

   1 Power
   2 Status
   3 Fault

   The canister is ready with no critical errors when **Power** is illuminated, **Status** is illuminated, and **Fault** is off. If a canister is *not* ready, see "Procedure: Understanding the Storwize V7000 Gen2 system status from the LEDs" on page 107.
3. Wait for all expansion canisters to finish powering on.
4. Power on the control enclosure by connecting both power supply units of the enclosure to their power sources, using the supplied power cables. If the power sources have circuit breakers or switches, ensure that they are turned on. The enclosure does not have power switches.

   Note: Each enclosure has two power supply units. To provide power failure redundancy, connect the two power cords to separate power circuits.

5. Check the LEDs on each node canister in the control enclosure, as displayed in Figure 43.
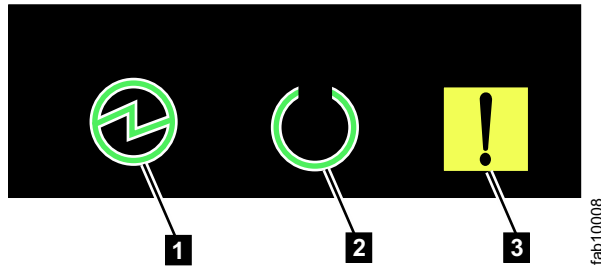
Figure 43. Node canister LEDs

1 Power
2 Status
3 Fault

The canister is ready with no critical errors when **Power** is illuminated, **Status** is blinking, and **Fault** is off. If a canister is *not* ready, refer to the "Procedure: Understanding the system status using the LEDs" topic in the troubleshooting section of the Storwize V7000 information center.

Procedure: Powering off a Storwize V7000 Gen2 control enclosure

To service a Storwize V7000 Gen2 control enclosure, you must safely power off both of the node canisters in the enclosure.

Before you begin

Host connectivity for hosts that are connected to the control enclosure is lost when the control enclosure is shut down. You must quiesce I/O activity of these hosts before completing this procedure.

About this task

If your system has a single control enclosure, complete the steps in "Procedure: Powering off your system" on page 130 instead of following this procedure. Doing so provides a more coordinated shutdown of the whole system.

Attention:

- If your system is powered on and doing I/O operations, you must power off the control enclosures correctly to ensure that no data is lost. If possible, always use the fix procedures that are presented by the management GUI to manage and maintain your system. The fix procedures ensure that node canisters are powered off safely.
- If the system includes two control enclosures, some volumes might become inaccessible when a control enclosure shuts down. Refer to "Procedure: Understanding Storwize V7000 Gen2 volume dependencies" on page 140 to determine whether it is appropriate to continue this procedure.
- If you need to power off the node canister that is operating as the configuration node, power down the other node canister first, then power down the configuration node second. This sequence prevents two failovers from happening, reducing delays in powering off the control enclosure.

Procedure

To power off a control enclosure, complete the following steps:

1. Use the management GUI to determine which two node canisters are in the control enclosure that is to be powered off. Note whether one of the two nodes is the configuration node so that you can power it off second.
2. Go to the service assistant for the first node to be powered off.
3. On the home page, select the node canister to be powered off.
4. Use the **Power off** action to power off the canister.
5. Wait for the node to appear offline.
6. Repeat steps 2 through 5 on the second node canister in the enclosure that is to be powered off.
7. If another control enclosure in the system is online, the management GUI can be used to confirm that the node status of both nodes is offline. The status LEDs on both canisters indicate whether the node is powered off, as described in "Procedure: Understanding the Storwize V7000 Gen2 system status from the LEDs" on page 107.
8. Turn off the power to the enclosure and disconnect both power cables from the enclosure.

What to do next

After you power off a control enclosure by using this procedure, you must reconnect the power cables and turn on the power. The node canisters start.

Procedure: Restarting a Storwize V7000 Gen2 control enclosure

To restart a Storwize V7000 Gen2 control enclosure, you must reseat the node canisters.

About this task

If a Storwize V7000 Gen2 control enclosure has been shut down with the GUI or CLI, you must restart it by reseating the node canisters.

Note: If enclosure input power was removed from both power supply units, the input power LED of both PSUs is off. To restart the enclosure in this case, refer to "Procedure: Powering on the Storwize V7000 Gen2 system".

Procedure

1. Refer to "Procedure: Understanding the Storwize V7000 Gen2 system status from the LEDs".
2. Locate the ac power (input) and dc power (output) LEDs on each of the two power supply units in the enclosure. If these LEDs are not lit green, go to "Procedure: Powering on the Storwize V7000 Gen2 system".
3. Locate the Power LED on the rear of node canister 1. If it is lit green and slowly flashing, reseat the node canister. The power LED of the reseated canister flashes fast for a few seconds while it starts.
4. Repeat step 3 on node canister 2.
5. Wait a few minutes until the green status LED of at least one reseated node canister is lit. The system is now online.
6. Log in to the management GUI to obtain the status of the system.

Procedure: Powering off a Storwize V7000 Gen2 node canister

You can safely power off a node canister to service the node canister.

About this task

Attention: After powering off a node canister using this procedure, a physical reseat of the canister will be required to power it back on. The reseat procedure requires physical access to the enclosure and is described in "Procedure: Reseating a Storwize V7000 Gen2 node canister" on page 128.

While a node canister is powered off, some volumes can become inaccessible. Refer to "Procedure: Understanding Storwize V7000 Gen2 volume dependencies" on page 140 to determine whether it is appropriate to continue this procedure.

If your system is powered on and doing I/O operations, it is important that the system is powered off correctly to ensure that no data is lost. If possible, always use the fix procedures that are presented by the management GUI to manage and maintain your system. The fix procedures ensure that the canister is powered off safely.

Procedure

To power off a node canister, complete the following steps:

1. Go to the service assistant for the node with the canister to shut down.
2. On the home page, select the node canister to shut down.
3. If you intend to do maintenance of the node canister, click **Identify** to light the Identify LED on the canister. Confirm that you know the location of the node canister.
4. Use the **Power off** action to power off the canister.
5. After the node is powered off, the service assistant shows that the node status is offline. The status LEDs on the canister indicate that the node is powered off.

Procedure: Collecting information for support

IBM support might ask you to collect trace files and dump files from your system to help them resolve a problem.

About this task

The management GUI and the service assistant have features to assist you in collecting the required information. The management GUI collects information from all the components in the system. The service assistant collects information from a single node canister. When the information that is collected is packaged together in a single file, the file is called a *snap*.

Special tools that are only available to the support teams are required to interpret the contents of the support package. The files are not designed for customer use.

Procedure

Always follow the instructions that are given by the support team to determine whether to collect the package by using the management GUI or the service assistant. Instruction is also given for which package content option is required.

- If you are collecting the package by using the management GUI, select **Settings** > **Support**. Click **Download Support Package**. Follow the instructions to download the appropriate log files.
- If you are collecting the package by using the service assistant, ensure that the node that you want to collect logs from is the current node. Select the **Collect Logs** option from the navigation. You can collect a support package or copy an individual file from the node canister. Follow the instructions to collect the information.

Procedure: Rescuing node canister software from another node (node rescue)

Use this procedure to complete a node rescue. This procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 88. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 89. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives

Table 89. Storwize V7000 Gen2 model numbers (continued)

Machine type/model	Description
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Procedure: Rescuing Storwize V7000 Gen2 node canister software from another node (node rescue)

Use this procedure to rescue a node.

About this task

A failure indicates that the node software is damaged and must be reinstalled.

Procedure

1. Ensure that the node you want to reinstall the code on is the current node. Go to "Accessing the service assistant" on page 74.
2. Select **Reinstall Machine Code** from the navigation.
3. Select **Rescue from another node**.

Procedure: Rescuing node canister software from another node (node rescue)

You can use the service assistant to re-install software from another node to rescue a node canister.

About this task

A failure has indicated that the node software is damaged and must be reinstalled.

Procedure

1. Log on to the service assistant. See "Accessing the service assistant" on page 74 for details.
2. Ensure that the node you want to reinstall the code on is the current node.
3. Select **Rescue from another node**.

Procedure: FCoE host-linking

About this task

If you are having problems attaching to the FCoE hosts, your problem might be related to the network, the Storwize V7000 system, or the host.

Procedure

1. If you are seeing error code 705 on the node, this means that the Fibre Channel I/O port is inactive. Note that FCoE uses Fibre Channel as a protocol and an Ethernet as an interconnect. If you are dealing with an FCoE enabled port, this means that either the Fibre Channel Forwarder (FCF) is not seen or the FCoE feature is not configured on the switch.

   a. Check that the FCoE feature is enabled on the FCF.

   b. Check the remote port (switch port) properties on the FCF.

2. If you are connecting the host through a Converged Enhanced Ethernet (CEE) switch, for network problems, you can attempt any of the following actions:

   a. Test your connectivity between the host and CEE switch.

   b. Ask the Ethernet network administrator to check the firewall and router settings.

3. Run **svcinfo lsfabric** and check that the host is seen as a remote port in the output. If not, then do the following tasks in order:

   a. Verify that Storwize V7000 and host get an fcid on FCF. If not, check the VLAN configuration.

   b. Verify that Storwize V7000 and host port are part of a zone and that zone is currently in force.

   c. Verify the volumes are mapped to the host and that they are online. See **lshostvdiskmap** and **lsvdisk** in the CLI configuration guide for more information.

4. If you still have FCoE problems, you can attempt the following action:

   a. Verify that the host adapter is in good state. You can unload and load the device driver and see the operating system utilities to verify that the device driver is installed, loaded, and operating correctly.

Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister

To remove or replace the lid of a Storwize V7000 Gen2 node canister, use this procedure.

About this task

It might be necessary to service a node canister or to replace a part that is either a customer replaceable unit (CRU) or a field replaceable unit (FRU) that is contained within the canister. To remove the lid of a node canister, follow these steps.

Attention: The lid of a node canister can be removed only after the canister was removed from its enclosure. Unless you are otherwise instructed, follow the procedure for removing a node canister to remove a node canister from its enclosure.

To remove a canister lid:

1. Place the node canister upside down on a work surface, with the release levers facing toward you.

2. Open the cover of the canister by depressing the recessed, blue touch points on the lid and sliding the lid away from you, as shown in Figure 44 on page 140.

To replace a canister lid, slide the canister lid onto the canister until the catch clicks and the lid edges are flush with the canister.
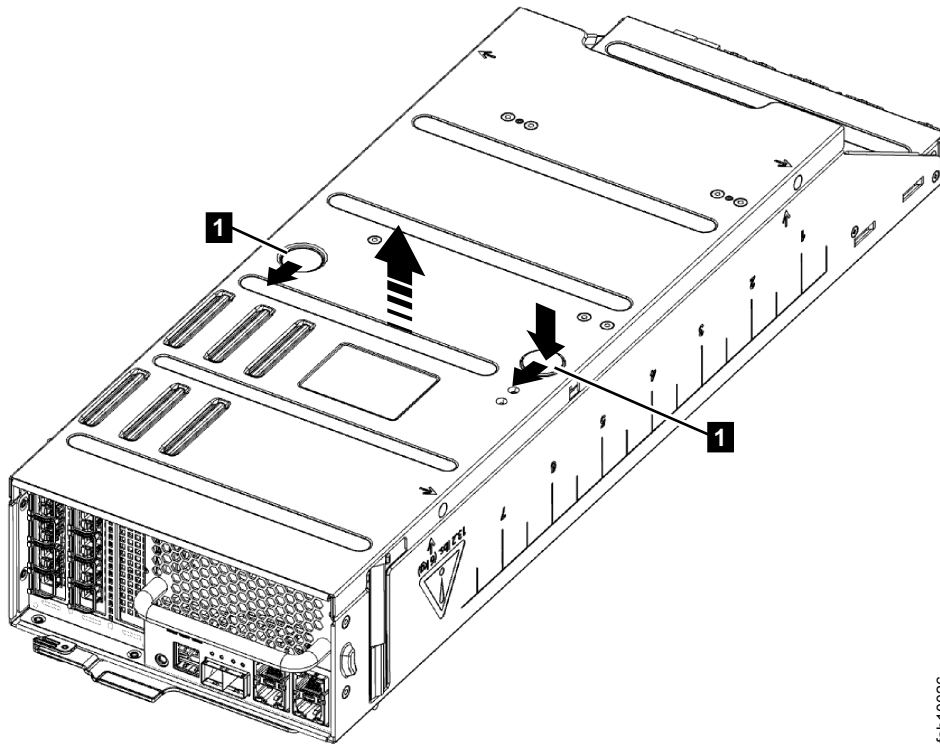
Figure 44. Replacing the canister cover

Procedure: Understanding Storwize V7000 Gen2 volume dependencies

If one component in a redundant pair is offline or powered off, host access to volumes depends on a Storwize V7000 Gen2 enclosure or canister in the system.

- If a control enclosure only has one node canister online, access to a volume depends on the online node canister if the volume is stored partially or wholly on an array that uses drives in the control enclosure or its expansion enclosures.

- If one expansion canister in an expansion enclosure is powered off, any expansion canisters further down that side of the chain become isolated from the control canister on that side of the chain. In this case, host access to volumes depends on the online canister if the volume uses drives in an isolated enclosure or the enclosure with the offline canister.

- If an entire expansion enclosure is powered off, both the left and the right side of the SAS chain are broken. In this case, host access to some volumes can be considered to depend on the entire expansion enclosure.

The impact that a service procedure might have on host access to data can be understood by using the management GUI.

1. Log on to the management GUI. Go to **Monitoring** > **System** .
2. From the dynamic graphic, right click the canister and select **Show Dependent Volumes** to see which volumes would be inaccessible if the component was taken offline or powered off.

If during a maintenance procedure, the **Show Dependent Volumes** action indicates that there are dependent volumes, you might choose to stop the procedure to

investigate whether it is possible to reinstate the redundancy in the system so that a procedure can be carried out without loss of access to data. An example would be to do procedures to ensure that both canisters in the enclosure are online before doing another procedure that powers off the only online canister in the enclosure.

Procedure: SAN problem determination

About this task

SAN failures might cause Storwize V7000 volumes to be inaccessible to host systems. Failures can be caused by SAN configuration changes or by hardware failures in SAN components.

The following list identifies some of the hardware that might cause failures:
- Power, fan, or cooling
- Application-specific integrated circuits
- Installed small form-factor pluggable (SFP) transceiver
- Fiber-optic cables

If error codes sent you here, complete the following steps:

Procedure

1. Verify that the power is turned on to all switches and storage controllers that the Storwize V7000 system uses, and that they are not reporting any hardware failures. If problems are found, resolve those problems before you proceed further.
2. Verify that the Fibre Channel cables that connect the systems to the switches are securely connected.
3. If you have a SAN management tool, use that tool to view the SAN topology and isolate the failing component.

iSCSI performance analysis and tuning

This procedure provides a solution for Internet Small Computer Systems Interface (iSCSI) host performance problems while connected to a Storwize V7000 system and its connectivity to the network switch.

About this task

Some of the attributes and host parameters that might affect iSCSI performance:
- Transmission Control Protocol (TCP) Delayed ACK
- Ethernet jumbo frame
- Network bottleneck or oversubscription
- iSCSI session login balance
- Priority flow control (PFC) setting and bandwidth allocation for iSCSI on the network

Procedure

1. Disable the TCP delayed acknowledgment feature.

   To disable this feature, refer to OS/platform documentation.
   - VMWare: http://kb.vmware.com/selfservice/microsites/microsite.do

- Windows: http://support.microsoft.com/kb/823764

The primary signature of this issue: read performance is significantly lower than write performance. Transmission Control Protocol (TCP) delayed acknowledgment is a technique that is used by some implementations of the TCP in an effort to improve network performance. However, in this scenario where the number of outstanding I/O is 1, the technique can significantly reduce I/O performance.

In essence, several ACK responses can be combined together into a single response, reducing protocol overhead. As described in RFC 1122, a host can delay sending an ACK response by up to 500 ms. Additionally, with a stream of full-sized incoming segments, ACK responses must be sent for every second segment.

Important: The host must be rebooted for these settings to take effect. A few platforms (for example, standard Linux distributions) do not provide a way to disable this feature. However, the issue was resolved with the version 7.1 release, and no host configuration changes are required to manage `TcpDelayedAck` behavior.

2. Enable jumbo frame for iSCSI.

   Jumbo frames are Ethernet frames with a size in excess of 1500 bytes. The maximum transmission unit (MTU) parameter is used to measure the size of jumbo frames.

   The Storwize V7000 supports 9000-bytes MTU. Refer to the CLI command `cfgportip` to enable jumbo frame. This command is disruptive as the link flips and the I/O operation through that port pauses.

   The network must support jumbo frames end-to-end for this to be effective; verify this by sending a ping packet to be delivered without fragmentation. For example:

   - Windows:

     ```
     ping -t <iscsi target ip> -S <iscsi initiator ip> -f -l <new mtu size - packet overhead (usually 36, might differ)>
     ```

     The following command is an example of a command that is used to check whether a 9000-bytes MTU is set correctly on a Windows 7 system:

     ```
     ping -t -S 192.168.1.117 192.168.1.217 -f -l 8964
     ```

     The following output is an example of a successful reply:

     ```
     192.168.1.217: bytes=8964 time=1ms TTL=52
     ```

   - Linux:

     ```
     ping -l <source iscsi initatior ip> -s <new mtu size> -M do <iscsi target ip>
     ```

   - ESXi:

     ```
     ping <iscsi target ip> -I <source iscsi initiator ip> -s <new mtu size - 28> -d
     ```

3. Verify the switch's port statistic where initiator/target ports are connected to make sure that packet drops are not high.

   Review network architecture to avoid any bottlenecks and oversubscription. The network needs to be balanced to avoid any packet drop; packet drop significantly reduces storage performance. Involve networking support to fix any such issues.

4. Optimize and utilize all iSCSI ports.

   To optimize Storwize V7000 resource utilization, all iSCSI ports must be used.

   - Each port is assigned to one CPU, and by balancing the login, one can maximize CPU utilization and achieve better performance. Ideally, configure subnets equal to the number of iSCSI ports on the Storwize V7000 node.

Configure each port of a node with an IP on a different subnet and keep it the same for other nodes. The following example displays an ideal configuration:

```
Node 1
Port 1: 192.168.1.11
Port 2: 192.168.2.21
Port 3: 192.168.3.31

Node 2:
Port 1: 192.168.1.12
Port 2: 192.168.2.22
Port 3: 192.168.3.33
```

- Avoid situations where 50 hosts are logged in to port 1 and only five hosts are logged in to port 2.
- Use proper subnetting to achieve a balance between the number of sessions and redundancy.

5. Troubleshoot problems with PFC settings.

   You do not need to enable PFC on the Storwize V7000 system. Storwize V7000 reads the data center bridging exchange (DCBx) packet and enables PFC for iSCSI automatically if it is enabled on the switch. In the **lsportip** command output, the fields `lossless_iscsi` and `lossless_iscsi6` show [on/off] depending on whether PFC is enabled or not for iSCSI on the system.

   If the fields `lossless_iscsi` and `lossless_iscsi6` are showing `off`, it might be due to one of the following reasons:

   a. VLAN is not set for that IP. Verify the following checks:
      - For IP address type IPv4, check the `vlan` field in the **lsportip** output. It should not be blank.
      - For IP address type IPv6, check the `vlan_6` field in the **lsportip** output. It should not be blank.
      - If the `vlan` and `vlan_6` fields are blank, set the VLAN for the IP type using Configuring VLAN for iSCSI.

   b. Host flag is not set for that IP. Verify the following checks:
      - For IP address type IPv4, check the `host` field in the **lsportip** output. It should be `yes`.
      - For IP address type IPv6, check the `host_6` field in the **lsportip** output. It should be `yes`.
      - If the `host` and `host_6` fields are not `yes`, set the host flag for the IP type using the **cfgportip** CLI command.

   c. PFC is not properly set on the switch.

      If the VLAN is properly set, and the host flag is also set, but the `lossless_iscsi` or `lossless_iscsi6` field is still showing `off`, some switch settings might be missing or incorrect.

      Verify the following settings in the switch:
      - Priority tag is set for iSCSI traffic.
      - PFC is enabled for priority tag that is assigned to iSCSI CoS.
      - DCBx is enabled on the switch.

      Also check the appropriate documentation:
      - Consult the documentation for enabling PFC on your specific switch.

- Consult the documentation for enabling PFC on Red Hat Enterprise Linux (RHEL) and Windows hosts specific to your configuration.
6. Ensure that proper bandwidth is given to iSCSI on the network.

   You can divide the bandwidth among the various types of traffic. It is important to assign proper bandwidth for good performance. To assign bandwidth for iSCSI traffic, you need to first enable the priority flow control for iSCSI.

Fibre Channel link failures

When a failure occurs on a single Fibre Channel link, the small form-factor pluggable (SFP) transceiver might need to be replaced.

Before you begin

The following items can indicate that a single Fibre Channel link failed:
- The Fibre Channel status LEDs at the rear of the node canister
- An error that indicates a single port failed

Procedure

Attempt each of these actions, in the following order, until the failure is fixed.
1. Ensure that the Fibre Channel cable is securely connected at each end.
2. Replace the Fibre Channel cable.
3. Replace the SFP transceiver for the failing port on the node.

   Note: Storwize V7000 nodes are supported by both longwave SFP transceivers and shortwave SFP transceivers. You must replace an SFP transceiver with the same type of SFP transceiver. If the SFP transceiver to replace is a longwave SFP transceiver, for example, you must provide a suitable replacement. Removing the wrong SFP transceiver might result in loss of data access.
4. Replace the Fibre Channel adapter on the node.
5. Contact the IBM Support Center for assistance in replacing the node canister.

Servicing storage systems

Storage systems that are supported for attachment to the Storwize V7000 system are designed with redundant components and access paths to enable concurrent maintenance. Hosts have continuous access to their data during component failure and replacement.

Procedure: SAN problem determination

About this task

SAN failures might cause Storwize V7000 volumes to be inaccessible to host systems. Failures can be caused by SAN configuration changes or by hardware failures in SAN components.

The following list identifies some of the hardware that might cause failures:
- Power, fan, or cooling
- Application-specific integrated circuits

- Installed small form-factor pluggable (SFP) transceiver
- Fiber-optic cables

If error codes sent you here, complete the following steps:

Procedure

1. Verify that the power is turned on to all switches and storage controllers that the Storwize V7000 system uses, and that they are not reporting any hardware failures. If problems are found, resolve those problems before you proceed further.
2. Verify that the Fibre Channel cables that connect the systems to the switches are securely connected.
3. If you have a SAN management tool, use that tool to view the SAN topology and isolate the failing component.

iSCSI performance analysis and tuning

This procedure provides a solution for Internet Small Computer Systems Interface (iSCSI) host performance problems while connected to a Storwize V7000 system and its connectivity to the network switch.

About this task

Some of the attributes and host parameters that might affect iSCSI performance:
- Transmission Control Protocol (TCP) Delayed ACK
- Ethernet jumbo frame
- Network bottleneck or oversubscription
- iSCSI session login balance
- Priority flow control (PFC) setting and bandwidth allocation for iSCSI on the network

Procedure

1. Disable the TCP delayed acknowledgment feature.

   To disable this feature, refer to OS/platform documentation.
   - VMWare: http://kb.vmware.com/selfservice/microsites/microsite.do
   - Windows: http://support.microsoft.com/kb/823764

   The primary signature of this issue: read performance is significantly lower than write performance. Transmission Control Protocol (TCP) delayed acknowledgment is a technique that is used by some implementations of the TCP in an effort to improve network performance. However, in this scenario where the number of outstanding I/O is 1, the technique can significantly reduce I/O performance.

   In essence, several ACK responses can be combined together into a single response, reducing protocol overhead. As described in RFC 1122, a host can delay sending an ACK response by up to 500 ms. Additionally, with a stream of full-sized incoming segments, ACK responses must be sent for every second segment.

   Important: The host must be rebooted for these settings to take effect. A few platforms (for example, standard Linux distributions) do not provide a way to

disable this feature. However, the issue was resolved with the version 7.1 release, and no host configuration changes are required to manage **TcpDelayedAck** behavior.

2. Enable jumbo frame for iSCSI.

   Jumbo frames are Ethernet frames with a size in excess of 1500 bytes. The maximum transmission unit (MTU) parameter is used to measure the size of jumbo frames.

   The Storwize V7000 supports 9000-bytes MTU. Refer to the CLI command **cfgportip** to enable jumbo frame. This command is disruptive as the link flips and the I/O operation through that port pauses.

   The network must support jumbo frames end-to-end for this to be effective; verify this by sending a ping packet to be delivered without fragmentation. For example:

   - Windows:

     ```
     ping -t <iscsi target ip> -S <iscsi initiator ip> -f -l <new mtu size - packet overhead (usually 36, might differ)>
     ```

     The following command is an example of a command that is used to check whether a 9000-bytes MTU is set correctly on a Windows 7 system:

     ```
     ping -t -S 192.168.1.117 192.168.1.217 -f -l 8964
     ```

     The following output is an example of a successful reply:

     ```
     192.168.1.217: bytes=8964 time=1ms TTL=52
     ```

   - Linux:

     ```
     ping -l <source iscsi initatior ip> -s <new mtu size> -M do <iscsi target ip>
     ```

   - ESXi:

     ```
     ping <iscsi target ip> -I <source iscsi initiator ip> -s <new mtu size - 28> -d
     ```

3. Verify the switch's port statistic where initiator/target ports are connected to make sure that packet drops are not high.

   Review network architecture to avoid any bottlenecks and oversubscription. The network needs to be balanced to avoid any packet drop; packet drop significantly reduces storage performance. Involve networking support to fix any such issues.

4. Optimize and utilize all iSCSI ports.

   To optimize Storwize V7000 resource utilization, all iSCSI ports must be used.

   - Each port is assigned to one CPU, and by balancing the login, one can maximize CPU utilization and achieve better performance. Ideally, configure subnets equal to the number of iSCSI ports on the Storwize V7000 node. Configure each port of a node with an IP on a different subnet and keep it the same for other nodes. The following example displays an ideal configuration:

     ```
     Node 1
     Port 1: 192.168.1.11
     Port 2: 192.168.2.21
     Port 3: 192.168.3.31

     Node 2:
     Port 1: 192.168.1.12
     Port 2: 192.168.2.22
     Port 3: 192.168.3.33
     ```

   - Avoid situations where 50 hosts are logged in to port 1 and only five hosts are logged in to port 2.

- Use proper subnetting to achieve a balance between the number of sessions and redundancy.

5. Troubleshoot problems with PFC settings.

   You do not need to enable PFC on the Storwize V7000 system. Storwize V7000 reads the data center bridging exchange (DCBx) packet and enables PFC for iSCSI automatically if it is enabled on the switch. In the **lsportip** command output, the fields `lossless_iscsi` and `lossless_iscsi6` show [on/off] depending on whether PFC is enabled or not for iSCSI on the system.

   If the fields `lossless_iscsi` and `lossless_iscsi6` are showing `off`, it might be due to one of the following reasons:

   a. VLAN is not set for that IP. Verify the following checks:
      - For IP address type IPv4, check the `vlan` field in the **lsportip** output. It should not be blank.
      - For IP address type IPv6, check the `vlan_6` field in the **lsportip** output. It should not be blank.
      - If the `vlan` and `vlan_6` fields are blank, set the VLAN for the IP type using Configuring VLAN for iSCSI.

   b. Host flag is not set for that IP. Verify the following checks:
      - For IP address type IPv4, check the `host` field in the **lsportip** output. It should be `yes`.
      - For IP address type IPv6, check the `host_6` field in the **lsportip** output. It should be `yes`.
      - If the `host` and `host_6` fields are not `yes`, set the host flag for the IP type using the **cfgportip** CLI command.

   c. PFC is not properly set on the switch.

      If the VLAN is properly set, and the host flag is also set, but the `lossless_iscsi` or `lossless_iscsi6` field is still showing `off`, some switch settings might be missing or incorrect.

      Verify the following settings in the switch:
      - Priority tag is set for iSCSI traffic.
      - PFC is enabled for priority tag that is assigned to iSCSI CoS.
      - DCBx is enabled on the switch.

      Also check the appropriate documentation:
      - Consult the documentation for enabling PFC on your specific switch.
      - Consult the documentation for enabling PFC on Red Hat Enterprise Linux (RHEL) and Windows hosts specific to your configuration.

6. Ensure that proper bandwidth is given to iSCSI on the network.

   You can divide the bandwidth among the various types of traffic. It is important to assign proper bandwidth for good performance. To assign bandwidth for iSCSI traffic, you need to first enable the priority flow control for iSCSI.

Fibre Channel link failures

When a failure occurs on a single Fibre Channel link, the small form-factor pluggable (SFP) transceiver might need to be replaced.

Before you begin

The following items can indicate that a single Fibre Channel link failed:
- The Fibre Channel status LEDs at the rear of the node canister
- An error that indicates a single port failed

Procedure

Attempt each of these actions, in the following order, until the failure is fixed.
1. Ensure that the Fibre Channel cable is securely connected at each end.
2. Replace the Fibre Channel cable.
3. Replace the SFP transceiver for the failing port on the node.

   Note: Storwize V7000 nodes are supported by both longwave SFP transceivers and shortwave SFP transceivers. You must replace an SFP transceiver with the same type of SFP transceiver. If the SFP transceiver to replace is a longwave SFP transceiver, for example, you must provide a suitable replacement. Removing the wrong SFP transceiver might result in loss of data access.
4. Replace the Fibre Channel adapter on the node.
5. Contact the IBM Support Center for assistance in replacing the node canister.

Ethernet iSCSI host-link problems

If you are having problems attaching to the Ethernet hosts, your problem might be related to the network, the Storwize V7000 system, or the host.

Note: Storwize V7000 and Host IP should be on the same VLAN. Host and Storwize V7000 nodes should not have same subnet on different VLANs.

For network problems, you can attempt any of the following actions:
- Test your connectivity between the host and Storwize V7000 ports.
- Try to ping the Storwize V7000 system from the host.
- Ask the Ethernet network administrator to check the firewall and router settings.
- Check that the subnet mask and gateway are correct for the Storwize V7000 host configuration.

Using the management GUI for Storwize V7000 problems, you can attempt any of the following actions:
- View the configured node port IP addresses.
- View the list of volumes that are mapped to a host to ensure that the volume host mappings are correct.
- Verify that the volume is online.

For host problems, you can attempt any of the following actions:
- Verify that the host iSCSI qualified name (IQN) is correctly configured.
- Use operating system utilities (such as Windows device manager) to verify that the device driver is installed, loaded, and operating correctly.
- If you configured the VLAN, check that its settings are correct. Ensure that Host Ethernet port, Storwize V7000 Ethernet ports IP address, and Switch port are on the same VLAN ID. Ensure that on each VLAN, a different subnet is used. Configuring the same subnet on different VLAN IDs can cause network connectivity problems.

Chapter 7. Recover system procedure

The recover system procedure recovers the entire storage system if the system state is lost from all control enclosure node canisters. The procedure re-creates the storage system by using saved configuration data. The saved configuration data is in the active quorum disk and the latest XML configuration backup file. The recovery might not be able to restore all volume data. This procedure is also known as Tier 3 (T3) recovery.

CAUTION:
If the system encounters a state where:
- **No nodes are active, and**
- **One or more nodes have node errors that require a node rescue, node canister replacement, or node firmware reinstallation**

STOP and contact IBM Remote Technical Support. Initiating this T3 recover system procedure while in this specific state can result in loss of the XML backup of the block volume storage configuration.

Attention:
- Run service actions only when directed by the fix procedures. If used inappropriately, service actions can cause loss of access to data or even data loss. Before you attempt to recover a storage system, investigate the cause of the failure and attempt to resolve those issues by using other fix procedures. Read and understand all of the instructions before you complete any action.
- The recovery procedure can take several hours if the system uses large-capacity devices as quorum devices.

Do not attempt the recover system procedure unless the following conditions are met:
- All of the conditions have been met in "When to run the recover system procedure" on page 150.
- All hardware errors are fixed. See "Fix hardware errors" on page 150
- All node canisters have candidate status. Otherwise, see step 1.
- All node canisters must be at the same level of code that the storage system had before the system failure. If any node canisters were modified or replaced, use the service assistant to verify the levels of code, and where necessary, to reinstall the level of code so that it matches the level that is running on the other node canisters in the system.

The system recovery procedure is one of several tasks that must be completed. The following list is an overview of the tasks and the order in which they must be completed:
1. Preparing for system recovery
   a. Review the information regarding when to run the recover system procedure.
   b. Fix your hardware errors and make sure that all nodes in the system are shown in service assistant or in the output from `sainfo lsservicenodes`.
   c. Remove the system information for node canisters with error code 550 or error code 578 by using the service assistant, but only if the recommended

user response for these node errors has already been followed. See "Removing system information for node canisters with error code 550 or error code 578 using the service assistant" on page 152.

d. For Virtual Volumes (VVols), shut down the services for any instances of Spectrum Control Base that are connecting to the system. Use the Spectrum Control Base command `service ibm_spectrum_control stop`.

2. Running the system recovery. After you prepared the system for recovery and met all the pre-conditions, run the system recovery.

   Note: Run the procedure on one system in a fabric at a time. Do not run the procedure on different node canisters in the same system. This restriction also applies to remote systems.

3. Completing actions to get your environment operational.
   - Recovering from offline volumes by using the CLI.
   - Checking your system, for example, to ensure that all mapped volumes can access the host.

When to run the recover system procedure

Attempt a recover procedure only after a complete and thorough investigation of the cause of the system failure. Attempt to resolve those issues by using other service procedures.

Attention: If you experience failures at any time while running the recover system procedure, call the IBM Support Center. Do not attempt to do further recovery actions, because these actions might prevent support from restoring the system to an operational status.

Certain conditions must be met before you run the recovery procedure. Use the following items to help you determine when to run the recovery procedure:

Note: It is important to know the number of control enclosures in the system. When the instructions indicate that every node is checked, you must check the status of both nodes in every control enclosure. For some system problems or Fibre Channel network problems, you must run the service assistant directly on the node to get its status.

1. Check that no node in the system is active and that the management IP is not accessible. If any node has active status, it is not necessary to recover the system.

2. Resolve all hardware errors in nodes so that only node errors 578 or 550 are present. If this is not the case, go to "Fix hardware errors."

3. Ensure all backend storage that is administered by the system is present before you run the recover system procedure.

4. If any nodes have been replaced, ensure that the WWNN of the replacement node matches that of the replaced node, and that no prior system data remains on this node. (See "Procedure: Removing system data from a node canister" on page 120.)

Fix hardware errors

Before running a system recovery procedure, it is important to identify and fix the root cause of the hardware issues.

Identifying and fixing the root cause can help recover a system, if these are the faults that are causing the system to fail. The following are common issues which can be easily resolved:

- The node has been powered off or the power cords were unplugged.
- Check the node status of every node canister that is part of this system. Resolve all hardware errors except node error 578 or node error 550.
  - All nodes must be reporting either a node error 578 or a node error 550. These error codes indicate that the system has lost its configuration data. If any nodes report anything other than these error codes, do not perform a recovery. You can encounter situations where non-configuration nodes report other node errors, such as a 550 node error. The 550 error can also indicate that a node is not able to join a system.
  - If any nodes show a node error 550, record the error data that is associated with the 550 error from the service assistant.
    - In addition to the node error 550, the report can show data that is separated by spaces in one of the following forms:
      - Node identifiers in the format: *<enclosure_serial>-<canister slot ID>*(7 characters, hyphen, 1 number), for example, 01234A6-2
      - Quorum drive identifiers in the format: *<enclosure_serial>:<drive slot ID>*[<drive 11S serial number>] (7 characters, colon, 1 or 2 numbers, open square bracket, 22 characters, close square bracket), for example, 01234A9:21[11S1234567890123456789]
      - Quorum MDisk identifier in the format: *WWPN/LUN* (16 hexadecimal digits followed by a forward slash and a decimal number), for example, 1234567890123456/12
    - If the error data contains a node identifier, ensure that the node that is referred to by the ID is showing node error 578. If the node is showing a node error 550, ensure that the two nodes can communicate with each other. Verify the SAN connectivity, and if the 550 error is still present, restart one of the two nodes from the service assistant by clicking **Restart Node**.
    - If the error data contains a quorum drive identifier, locate the enclosure with the reported serial number. Verify that the enclosure is powered on and that the drive in the reported slot is powered on and functioning. If the node canister that is reporting the fault is in the I/O group of the listed enclosure, ensure that it has SAS connectivity to the listed enclosure. If the node canister that is reporting the fault is in a different I/O group from the listed enclosure, ensure that the listed enclosure has SAS connectivity to both node canisters in the control enclosure in its I/O group. After verification, restart the node by clicking **Restart Node** from the service assistant.
    - If the error data contains a quorum MDisk identifier, verify the SAN connectivity between this node and that WWPN. Check the storage controller to ensure that the LUN referred to is online. After verification, if the 550 error is still present, restart the node from the service assistant by clicking **Restart Node**.
    - If there is no error data, the error is because there are insufficient connections between nodes over the Fibre Channel network. Each node must have at least two independent Fibre Channel logical connections, or logins, to every node that is not in the same enclosure. An independent connection is one where both physical ports are different. In this case, there

is a connection between the nodes, but there is not a redundant connection. If there is no error data, wait 3 minutes for the SAN to initialize. Next, verify:

- There are at least two Fibre Channel ports that are operational and connected on every node.
- The SAN zoning allows every port to connect to every port on every other node
- All redundant SANs (if used) are operational.

After verification, if the 550 error is still present, restart the node from the service assistant by clicking **Restart Node**.

Note: If after resolving all these scenarios, half or greater than half of the nodes are reporting node error 578, it is appropriate to run the recovery procedure. Call the IBM Support Center for further assistance.

- For any nodes that are reporting a node error 550, ensure that all the missing hardware that is identified by these errors is powered on and connected without faults. If you cannot contact the service assistant from any node, isolate the problems by using the LED indicators.
- If you have not been able to restart the system, and if any node other than the current node is reporting node error 550 or 578, you must remove system data from those nodes. This action acknowledges the data loss and puts the nodes into the required candidate state.

Removing system information for node canisters with error code 550 or error code 578 using the service assistant

The system recovery procedure works only when all node canisters are in candidate status. Ensure that the service assistant displays all of the node canisters with the 550 error code. The 550 error code is the expected node error when more than half of the nodes in the system are missing or when the active quorum disk cannot be found. If the service assistant displays any node canisters with error codes 550 or 578 and all the recommended actions have been completed on these nodes, you must remove their system data.

About this task

Before performing this task, ensure that you have read the introductory information in the overall recover system procedure.

To remove system information from a node canister with an error 550 or 578, follow this procedure using the service assistant:

Procedure

1. Point your browser to the service IP address of one of the nodes, for example, https://node_service_ip_address/service/.

   For Storwize V7000 Gen1 , if you do not know the IP address or if it has not been configured, you must assign an IP address with the USB flash drive method. For Storwize V7000 Gen2, if you do not know the IP address or if it has not been configured, you must assign an IP address with the technician port method. See the initialization tool interface topic for more information.

2. Log on to the service assistant.

3. Select **Manage System**.

4. Click **Remove System Data**.

5. Confirm that you want to remove the system data when prompted.
6. Remove the system data for the other nodes that display a 550 or a 578 error.

   All nodes previously in this system must have a node status of `Candidate` and have no errors listed against them.
7. Resolve any hardware errors until the error condition for all nodes in the system is **None**.
8. Ensure that all nodes in the system display a status of candidate.

Results

When all nodes display a status of candidate and all error conditions are **None**, you can run the recovery procedure.

Related reference:

"Initialization tool interface" on page 79
Use the Initialization tool interface to initialize a system and to service the node canisters in a control enclosure. Although the Initialization tool wizard interface is similar, accessing the wizard differs, depending on the generation of your control enclosure model.

Running system recovery using the service assistant

You can use the service assistant to start recovery when all node canisters that were members of the system are online and have candidate status. For any nodes that display error code 550 or 578, ensure that all nodes in the system are visible and all the recommended actions are completed before placing them into candidate status. To place a node into candidate status, remove system information for that node canister. Do not run the recovery procedure on different node canisters in the same system.

Before you begin

Note: Ensure that the web browser is not blocking pop-up windows. If it does, progress windows cannot open.

Before you begin this procedure, read the recover system procedure introductory information; see Chapter 7, "Recover system procedure," on page 149.

About this task

Attention: This service action has serious implications if not completed properly. If at any time an error is encountered not covered by this procedure, stop and call the support center.

Run the recovery from any node canisters in the system; the node canisters must not have participated in any other system.

If the system has USB encryption, run the recovery from any node canister in the system that has a USB flash drive inserted which contains the encryption key.

If a cluster contains an encrypted cloud account that uses USB encryption, a USB flash drive with the cluster master key must be present in the configuration node before the cloud account can move to the online state. This requirement is necessary when the cluster is powered down, and then restarted.

If the system has key server encryption, note the following items before you proceed with the T3 recovery.

- Run the recovery on a node that is attached to the key server. The keys are fetched remotely from the key server.
- Run the recovery procedure on a node that is not hardware replaced or node rescued. All of the information that is required for a node to successfully fetch the key from the key server resides on the node's file system. If the contents of the node's original file system are damaged or no longer exist (rescue node, hardware replacement, file system that is corrupted, and so on), then the recovery fails from this node.

Note: Each individual stage of the recovery procedure can take significant time to complete, depending on the specific configuration.

Procedure

1. Point your browser to the service IP address of one of the node canisters.

   For Storwize V7000 Gen1 systems, if you do not know the IP address or if it is not configured, you must assign an IP address with the USB flash drive method. See "Procedure: Changing the service IP address of a node canister" on page 121 for details. For Storwize V7000 Gen2 systems, if you do not know the IP address or if it is not configured, you must assign an IP address with the technician port method. See "Procedure: Initializing the Storwize V7000 Gen2 system using the technician port" on page 124 for details.

2. Log on to the service assistant.

3. Check that all node canisters that were members of the system are online and have candidate status.

   If any nodes display error code 550 or 578, remove their system data to place them into candidate status; see "Procedure: Removing system data from a node canister" on page 120.

4. Select **Recover System** from the navigation.

5. Follow the online instructions to complete the recovery procedure.

   a. Verify the date and time of the last quorum time. The time stamp must be less than 30 minutes before the failure. The time stamp format is *YYYYMMDD hh:mm*, where *YYYY* is the year, *MM* is the month, *DD* is the day, *hh* is the hour, and *mm* is the minute.

      Attention: If the time stamp is not less than 30 minutes before the failure, call the support center.

   b. Verify the date and time of the last backup date. The time stamp must be less than 24 hours before the failure. The time stamp format is *YYYYMMDD hh:mm*, where *YYYY* is the year, *MM* is the month, *DD* is the day, *hh* is the hour, and *mm* is the minute.

      Attention: If the time stamp is not less than 24 hours before the failure, call the support center.

      Changes that are made after the time of this backup date might not be restored.

Results

Any one of the following categories of messages might be displayed:

- `T3 successful`

The volumes are back online. Use the final checks to get your environment operational again.

- `T3 recovery completed with errors`

  T3 recovery completed with errors: One or more of the volumes are offline because there was fast write data in the cache. To bring the volumes online, see "Recovering from offline volumes using the CLI" for details.

- `T3 failed`

  Call the support center. Do not attempt any further action.

Verify that the environment is operational by completing the checks that are provided in "What to check after running the system recovery" on page 156.

If any errors are logged in the error log after the system recovery procedure completes, use the fix procedures to resolve these errors, especially the errors that are related to offline arrays.

If the recovery completes with offline volumes, go to "Recovering from offline volumes using the CLI."

Recovering from offline volumes using the CLI

If a Tier 3 recovery procedure completes with offline volumes, then it is likely that the data which was in the write-cache of the node canisters was lost during the failure that caused all of the node canisters to lose the block storage system cluster state. You can use the command-line interface (CLI) to acknowledge that there was data lost from the write-cache and bring the volume back online to attempt to deal with the data loss.

About this task

If you run the recovery procedure but there are offline volumes, you can complete the following steps to bring the volumes back online. Some volumes might be offline because of write-cache data loss or metadata loss during the event that led all node canisters to lose cluster state. Any data that is lost from the write-cache cannot be recovered. These volumes might need extra recovery steps after the volume is brought back online.

Note: If you encounter errors in the event log after you run the recovery procedure that is related to offline arrays, use the fix procedures to resolve the offline array errors before fixing the offline volume errors.

Example

Complete the following steps to recover an offline volume after the recovery procedure is completed:

1. Delete all IBM FlashCopy® function mappings and Metro Mirror or Global Mirror relationships that use the offline volumes.
2. If the volume is a space efficient (SE) volume, run the **repairsevdiskcopy** command. This brings the volume back online so that you can attempt to deal with the data loss.
3. If the volume is not a SE volume, run the **recovervdiskbysystem** command. This brings the volume back online so that you can attempt to deal with the data loss.

4. Refer to "What to check after running the system recovery" for what to do with volumes that are corrupted by the loss of data from the write-cache.
5. Re-create all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the volumes.

What to check after running the system recovery

Several tasks must be completed before you use the system.

The recovery procedure re-creates the old system from the quorum data. However, some things cannot be restored, such as cached data or system data managing in-flight I/O. This latter loss of state affects RAID arrays that manage internal storage. The detailed map about where data is out of synchronization has been lost, meaning that all parity information must be restored, and mirrored pairs must be brought back into synchronization. Normally this results in either old or stale data being used, so only writes in flight are affected. However, if the array lost redundancy (such as syncing, degraded, or critical RAID status) before the error requiring system recovery, then the situation is more severe. Under this situation you need to check the internal storage:

- Parity arrays will likely be syncing to restore parity; they do not have redundancy when this operation proceeds.
- Because there is no redundancy in this process, bad blocks might have been created where data is not accessible.
- Parity arrays might be marked as corrupted. This indicates that the extent of lost data is wider than in-flight I/O; to bring the array online, the data loss must be acknowledged.
- RAID6 arrays that were degraded before the system recovery might require a full restore from backup. For this reason, it is important to have at least a capacity match spare available.

Be aware of these differences regarding the recovered configuration:

- FlashCopy mappings are restored as "idle_or_copied" with 0% progress. Both volumes must have been restored to their original I/O groups.
- The management ID is different. Any scripts or associated programs that refer to the system-management ID of the clustered system (system) must be changed.
- Any FlashCopy mappings that were not in the "idle_or_copied" state with 100% progress at the point of disaster have inconsistent data on their target disks. These mappings must be restarted.
- Intersystem partnerships and relationships are not restored and must be re-created manually.
- Consistency groups are not restored and must be re-created manually.
- Intrasystem Metro Mirror relationships are restored if all dependencies were successfully restored to their original I/O groups.
- Volumes with cloud snapshots that were enabled before the recovery need to have the cloud snapshots manually reenabled.
- If hardware was replaced before the recovery, the SSL certificate might not be restored. If it is not restored, then a new self-signed certificate is generated with a validity of 30 days. Follow the associated Directed Maintenance Procedures (DMP) for a permanent resolution.
- The system time zone might not have been restored.

- Any Global Mirror secondary volumes on the recovered system might have inconsistent data if there was replication I/O from the primary volume cached on the secondary system at the point of the disaster. A full synchronization is required when re-creating and restarting these relationships.
- Immediately after the T3 recovery process runs, compressed disks do not know the correct value of their used capacity. The disks initially set the capacity as the entire real capacity. When I/O resumes, the capacity is shrunk down to the correct value.

  Similar behavior occurs when you use the `-autoexpand` option on volumes. The real capacity of a disk might increase slightly, caused by the same kind of behavior that affects compressed volumes. Again, the capacity shrinks down as I/O to the disk is resumed.

Before you use the volumes, complete the following tasks:
- Start the host systems.
- Manual actions might be necessary on the hosts to trigger them to rescan for devices. You can complete this task by disconnecting and reconnecting the Fibre Channel cables to each host bus adapter (HBA) port.
- Verify that all mapped volumes can be accessed by the hosts.
- Run file system consistency checks.
- Run the application consistency checks.

For Virtual Volumes (VVols), complete the following tasks.
- After you confirm that the T3 completed successfully, restart Spectrum Control Base (SCB) services. Use the Spectrum Control Base command **service ibm_spectrum_control start**.
- Refresh the storage system information on the SCB GUI to ensure that the systems are in sync after the recovery.
  - To complete this task, login to the SCB GUI.
  - Hover over the affected storage system, select the menu launcher, and then select **Refresh**. This step repopulates the system.
  - Repeat this step for all Spectrum Control Base instances.
- Rescan the storage providers from within the vSphere Web Client.
  - Select **vCSA > Manage > Storage Providers > select Active VP > Re-scan icon**.

For Virtual Volumes (VVols), also be aware of the following information.

FlashCopy mappings are not restored for VVols. The implications are as follows.
- The mappings that describe the VM's snapshot relationships are lost. However, the Virtual Volumes that are associated with these snapshots still exist, and the snapshots might still appear on the vSphere Web Client. This outcome might have implications on your VMware back up solution.
  - Do not attempt to revert to snapshots.
  - Use the vSphere Web Client to delete any snapshots for VMs on a VVol data store to free up disk space that is being used unnecessarily.
- The targets of any outstanding 'clone' FlashCopy relationships might not function as expected (even if the vSphere Web Client recently reported clone operations as complete). For any VMs, which are targets of recent clone operations, complete the following tasks.
  - Perform data integrity checks as is recommended for conventional volumes.

– If clones do not function as expected or show signs of corrupted data, take a fresh clone of the source VM to ensure that data integrity is maintained.

Backing up and restoring the system configuration

You can back up and restore the configuration data for the system after preliminary tasks are completed.

Configuration data for the system provides information about your system and the objects that are defined in it. The backup and restore functions of the `svcconfig` command can back up and restore only your configuration data for the Storwize V7000 system. You must regularly back up your application data by using the appropriate backup methods.

You can maintain your configuration data for the system by completing the following tasks:
* Backing up the configuration data
* Restoring the configuration data
* Deleting unwanted backup configuration data files

Before you back up your configuration data, the following prerequisites must be met:
* No independent operations that change the configuration for the system can be running while the backup command is running.
* No object name can begin with an underscore character (_).

Note:
* The default object names for controllers, I/O groups, and managed disks (MDisks) do not restore correctly if the ID of the object is different from what is recorded in the current configuration data file.
* All other objects with default names are renamed during the restore process. The new names appear in the format *name*_r where *name* is the name of the object in your system.
* Connections to iSCSI Mdisks for migration purposes are not restored.

Before you restore your configuration data, the following prerequisites must be met:
* The Security Administrator role is associated with your user name and password.
* You have a copy of your backup configuration files on a server that is accessible to the system.
* You have a backup copy of your application data that is ready to load on your system after the restore configuration operation is complete.
* You know the current license settings for your system.
* You did not remove any hardware since the last backup of your configuration.
* No zoning changes were made on the Fibre Channel fabric that would prevent communication between the Storwize V7000 and any storage controllers that are present in the configuration.
* For configurations with more than one I/O group, if a new system is created on which the configuration data is to be restored, the I/O groups for the other control enclosures must be added.

- You have at least 3 USB flash drives if encryption was enabled on the system when its configuration was backed up. The USB flash drives are used for generation of new keys as part of the restore process or for manually restoring encryption if the system has less than 3 USB ports.

Use the following steps to determine how to achieve an ideal T4 recovery:
- Open the appropriate svc.config.backup.xml (or svc.config.cron.xml) file with a suitable text editor or browser and navigate to the **node section** of the file.
- For each node entry, make a note of the value of the following properties: IO_group_id, canister_id, enclosure_serial_number .
- Use the CLI **sainfo lsservicenodes** command and the data to determine which node canisters previously belonged in each I/O group.

Restoring the system configuration should be performed by one of the nodes previously in IO group zero. For example, **property name="IO_group_id" value="0"** . The remaining enclosures should be added, as required, in the appropriate order based on the previous **IO_group_id** of its node canisters.

Note: It is not currently possible to determine which canister within the identified enclosure was previously used for cluster creation. Typically the restoration should be performed by canister 1.

The Storwize V7000 analyzes the backup configuration data file and the system to verify that the required disk controller system nodes are available.

Before you begin, hardware recovery must be complete. The following hardware must be operational: hosts, Storwize V7000 enclosures, internal flash drives and expansion enclosures (if applicable), the Ethernet network, the SAN fabric, and any external storage systems (if applicable).

Backing up the system configuration using the CLI

You can back up your configuration data using the command-line interface (CLI).

Before you begin

Before you back up your configuration data, the following prerequisites must be met:
- No independent operations that change the configuration can be running while the backup command is running.
- No object name can begin with an underscore character (_).

About this task

The backup feature of the **svcconfig** CLI command is designed to back up information about your system configuration, such as volumes, local Metro Mirror information, local Global Mirror information, storage pools, and nodes. All other data that you wrote to the volumes is *not* backed up. Any application that uses the volumes on the system as storage, must use the appropriate backup methods to back up its application data.

You must regularly back up your configuration data and your application data to avoid data loss, such as after any significant changes to the system configuration.

Note: The system automatically creates a backup of the configuration data each day at 1 AM. This backup is known as a **cron** backup and is written to /dumps/svc.config.cron.xml_*serial#* on the configuration node.

Use these instructions to generate a manual backup at any time. If a severe failure occurs, both the configuration of the system and application data might be lost. The backup of the configuration data can be used to restore the system configuration to the exact state it was in before the failure. In some cases, it might be possible to automatically recover the application data. This backup can be attempted with the Recover System Procedure, also known as a Tier 3 (T3) procedure. To restore the system configuration without attempting to recover the application data, use the Restoring the System Configuration procedure, also known as a Tier 4 (T4) recovery. Both of these procedures require a recent backup of the configuration data.

Complete the following steps to back up your configuration data:

Procedure

1. Use your preferred backup method to back up all of the application data that you stored on your volumes.
2. Issue the following CLI command to back up your configuration:

   ```
   svcconfig backup
   ```

   The following output is an example of the messages that might be displayed during the backup process:

   ```
   CMMVC6112W io_grp io_grp1 has a default name
   CMMVC6112W io_grp io_grp2 has a default name
   CMMVC6112W mdisk mdisk14 ...
   CMMVC6112W node node1 ...
   CMMVC6112W node node2 ...
   ................................................
   ```

   The **svcconfig backup** CLI command creates three files that provide information about the backup process and the configuration. These files are created in the /dumps directory of the configuration node canister.

   Table 90 describes the three files that are created by the backup process:

   Table 90. Files created by the backup process

File name	Description
svc.config.backup.xml_<serial#>	Contains your configuration data.
svc.config.backup.sh_<serial#>	Contains the names of the commands that were issued to create the backup of the system.
svc.config.backup.log_<serial#>	Contains details about the backup, including any reported errors or warnings.

3. Check that the **svcconfig backup** command completes successfully, and examine the command output for any warnings or errors. The following output is an example of the message that is displayed when the backup process is successful:

   ```
   CMMVC6155I SVCCONFIG processing completed successfully
   ```

   If the process fails, resolve the errors, and run the command again.

4. Keep backup copies of the files outside the system to protect them against a system hardware failure. Copy the backup files off the system to a secure location; use either the management GUI or scp command line. For example:

```
pscp  -unsafe superuser@cluster_ip:/dumps/svc.config.backup.*
 /offclusterstorage/
```

The cluster_ip is the IP address or DNS name of the system and **offclusterstorage** is the location where you want to store the backup files.

Tip: To maintain controlled access to your configuration data, copy the backup files to a location that is password-protected.

Restoring the system configuration

Use this procedure to restore the system configuration in the following situations:: only if the recover procedure failed or if the data that is stored on the volumes is not required. For directions on the recover procedure, see Chapter 7, "Recover system procedure," on page 149.

Before you begin

This configuration restore procedure is designed to restore information about your configuration, such as volumes, local Metro Mirror information, local Global Mirror information, storage pools, and nodes. The data that you wrote to the volumes is not restored. To restore the data on the volumes, you must restore application data from any application that uses the volumes on the clustered system as storage separately. Therefore, you must have a backup of this data before you follow the configuration recovery process.

If USB encryption was enabled on the system when its configuration was backed up, then at least 3 USB flash drives need to be present in the node canister USB ports for the configuration restore to work. The 3 USB flash drives must be inserted into the single node from which the configuration restore commands are run. Any USB flash drives in other nodes (that might become part of the cluster) are ignored. If you are not recovering a cloud backup configuration, the USB flash drives do not need to contain any keys. They are for generation of new keys as part of the restore process. If you are recovering a cloud backup configuration, the USB flash drives must contain the previous set of keys to allow the current encrypted data to be unlocked and re-encrypted with the new keys.

During T4 recovery, a new system is created with a new certificate. If the system has key server encryption, the new certificate must be exported by using the **chsystemcert -export** command, and then installed on the key server in the correct device group before you run the T4 recovery. The device group that is used is the one in which the previous system was defined. It might also be necessary to get the new cluster's certificate signed. In a T4 recovery, inform the key server administrator that the active keys are considered compromised.

About this task

You must regularly back up your configuration data and your application data to avoid data loss. If a system is lost after a severe failure occurs, both configuration for the system and application data is lost. You must restore the system to the exact state it was in before the failure, and then recover the application data.

During the restore process, the nodes and the storage enclosure are restored to the system, and then the MDisks and the array are re-created and configured. If multiple storage enclosures are involved, the arrays and MDisks are restored on the proper enclosures based on the enclosure IDs.

Important:
- There are two phases during the restore process: prepare and execute. You must not change the fabric or system between these two phases.
- For Storwize V7000 systems that contain nodes that are attached to external controllers virtualized by iSCSI, all nodes must be added into the system before you restore your data. Additionally, the system **cfgportip** settings and iSCSI storage ports must be manually reapplied before you restore your data. See step 13 on page 164.
- For VMware vSphere Virtual Volumes (sometimes referred to as VVols) environments, after a T4 restoration, some of the Virtual Volumes configuration steps are already completed: metadatavdisk created, usergroup and user created, adminlun hosts created. However, the user must then complete the last two configuration steps manually (creating a storage container on IBM Spectrum Control Base Edition and creating virtual machines on VMware vCenter). See Configuring Virtual Volumes.
- Restoring the system configuration should be performed via one of the nodes previously in IO group zero. For example, **property name="IO_group_id" value="0"**. The remaining enclosures should be added, as required, in the appropriate order based on the previous **IO_group_id** of its node canisters.
- For systems with an encrypted cloud backup configuration, during a T4 recovery the USB key that contained the system master key from the original system must be inserted into the configuration node of the new system. Alternatively, if a key server is used, the key server must contain the system master key from the original system. If the original system master key is not available, and the system data is encrypted in the cloud provider, then the data in the cloud is not accessible.
- If a cluster contains an encrypted cloud account that uses USB encryption, a USB flash drive with the cluster master key must be present in the configuration node before the cloud account can move to the online state. This requirement is necessary when the cluster is powered down, and then restarted.
- After a T4 recovery, cloud accounts are in an offline state. It is necessary to re-enter the authentication information to bring the accounts back online. See Managing cloud accounts.
- After a T4 recovery, volumes with cloud snapshots that were enabled before the recovery need to have the cloud snapshots manually reenabled. See Managing cloud volumes.

If you do not understand the instructions to run the CLI commands, see the command-line interface reference information.

To restore your configuration data, follow these steps:

Procedure
1. Verify that all nodes are available as candidate nodes before you run this recovery procedure. You must remove errors 550 or 578 to put the node in candidate state. For all nodes that display these errors, follow these steps:
    a. Point your browser to the service IP address of one of the nodes, for example, https://node_service_ip_address/service/.

b. Log on to the service assistant.

c. From the **Home** page, put the node into service state if it is not already in that state.

d. Select **Manage System**.

e. Click **Remove System Data**.

f. Confirm that you want to remove the system data when prompted.

g. Exit service state from the **Home** page. The 550 or 578 errors are removed, and the node appears as a candidate node.

h. Remove the system data for the other nodes that display a 550 or a 578 error.

   All nodes previously in this system must have a node status of `Candidate` and have no errors listed against them.

   Note: A node that is powered off might not show up in this list of nodes for the system. Diagnose hardware problems directly on the node using the service assistant IP address and by physically verifying the LEDs for the hardware components.

   Note: If you use the management GUI for the initial setup to restore the system configuration, check if a default call home email user was created. If it was created, delete the default call home email user in order for the T4 system recovery to proceed successfully.

2. Verify that all nodes are available as candidate nodes with blank system fields. Perform the following steps on one node in each control enclosure:

   a. Connect to the service assistant on either of the nodes in the control enclosure.

   b. Select **Configure Enclosure**.

   c. Select the **Reset the system ID** option. Do not make any other changes on the panel.

   d. Click **Modify**.

3. Create a system.

   - If your system is a Storwize V7000 Gen2 system, use the technician port.

   - If your system is a Storwize V7000 Gen1 system, use the initialization tool that is available on the USB flash drive. Select the **Initialize a new** Storwize V7000 (block system only) option from the `Welcome` panel of the initialization tool.

   - If your system is a hybrid Storwize V7000 system (a mix of Storwize V7000 Gen1 and Storwize V7000 Gen2 enclosures), you must begin the system creation and restore process from a Storwize V7000 Gen1 node.

4. In a supported browser, enter the IP address that you used to initialize the system and the default superuser password (`passw0rd`).

5. The setup wizard is shown. Be aware of the following items:

   a. Accept the license agreements.

   b. Set the values for the system name, date and time settings, and the system licensing. The original settings are restored during the configuration restore process.

   c. Verify the hardware. Only the control enclosure on which the clustered system was created and directly attached expansion enclosures are displayed. Any other control enclosures and expansion enclosures in other I/O groups are added to the system later.

Once the setup wizard finishes, make no other configuration changes.

6. If you set up email notification in the setup wizard, you must now remove that email user and server so that the original configuration can be restored.

   Issue the following CLI command to remove the new email user:

   `rmemailuser 0`

   Issue the following CLI command to remove the new email server:

   `rmemailserver 0`

7. Optional: From the management GUI, click **Access** > **Users** and configure an SSH key for the superuser.

8. By default, the newly initialized system is created in the storage layer. The layer of the system is not restored automatically from the configuration backup XML file. If the system you are restoring was previously configured in the replication layer, you must change the layer manually now. Refer to the System layers topic that is located under Product overview in the IBM Storwize V7000 Knowledge Center for more information.

9. If the clustered system was previously configured as replication layer, then use the **chsystem** command to change the layer setting.

10. For configurations with more than one I/O group, add the rest of the control enclosures into the clustered system by using the **addcontrolenclosure** CLI command. The remaining enclosures are added in the appropriate order based on the previous **IO_group_id** of its node canisters. The following example shows the command to add a control enclosure to I/O group 2.

    `svctask addcontrolenclosure -sernum SVT5M48 -iogrp 2`

11. Identify the configuration backup file from which you want to restore.

    The file can be either a local copy of the configuration backup XML file that you saved when you backed-up the configuration or an up-to-date file on one of the nodes.

    Configuration data is automatically backed up daily at 01:00 system time on the configuration node.

    Download and check the configuration backup files on all nodes that were previously in the system to identify the one containing the most recent complete backup

    a. From the management GUI, click **Settings** > **Support**.

    b. Click **Show full log listing**.

    c. For each node (canister) in the system, complete the following steps:

       1) Select the node to operate on from the selection box at the top of the table.

       2) Find all the files with names that match the pattern `svc.config.*.xml*`.

       3) Double-click the files to download them to your computer.

    The XML files contain a date and time that can be used to identify the most recent backup. After you identify the backup XML file that is to be used when you restore the system, rename the file to `svc.config.backup.xml`.

12. Copy onto the system the XML backup file from which you want to restore.

    `pscp full_path_to_identified_svc.config.file`
    `superuser@cluster_ip:/tmp/svc.config.backup.xml`

13. If the system contains any iSCSI storage controllers, these controllers must be detected manually now. The nodes that are connected to these controllers, the iSCSI port IP addresses, and the iSCSI storage ports must be added to the system before you restore your data.

a. To add these nodes, determine the panel name, node name, and I/O groups of any such nodes from the configuration backup file. To add the nodes to the system, run the following command:

`svctask addcontrolenclosure -iogrp` *iogrp_name_or_id* `-sernum` *enclosure_serial_number* `-site`

Where *enclosure_serial_number* is the serial number of the control enclosure, *iogrp_name_or_id* is the name or ID of the I/O group to which you want to add this node, and *site_id* is the numeric site value (1 or 2) of the control enclosure.

b. Run the following command to change the replication layer.

`chsystem -layer` *replication*

c. To restore iSCSI port IP addresses, use the **cfgportip** command.

1) To restore IPv4 address, determine id (port_id), node_id, node_name, IP_address, mask, gateway, host (0/1 stands for no/yes), remote_copy (0/1 stands for no/yes), and storage (0/1 stands for no/yes) from the configuration backup file, run the following command:

   `svctask cfgportip -node` *node_name_or_id* `-ip` *ipv4_address* `-gw` *ipv4_gw*
   `-host` *yes* | *no* `-remotecopy` *yes* | *no* `-storage` *yes* | *no* **port_id**

   Where *node_name_or_id* is the name or id of the node, *ipv4_address* is the IP v4 version protocol address of the port, and *ipv4_gw* is the IPv4 gateway address for the port.

2) To restore IPv6 address, determine id (port_id), node_id, node_name, IP_address_6, mask, gateway_6, prefix_6, host_6 (0/1 stands for no/yes), remote_copy_6 (0/1 stands for no/yes), and storage_6 (0/1 stands for no/yes) from the configuration backup file, run the following command:

   `svctask cfgportip -node` *node_name_or_id* `-ip_6` *ipv6_address* `-gw_6` *ipv6_gw*
   `-prefix_6` *prefix* `-host_6` *yes* | *no* `-remotecopy_6` *yes* | *no* `-storage_6` *yes* | *no* **port_id**

   Where *node_name_or_id* is the name or id of the node, *ipv6_address* is the IP v6 version protocol address of the port, *ipv6_gw* is the IPv6 gateway address for the port, and *prefix* is the IPv6 prefix.

Complete steps b.i and b.ii for all (earlier configured) IP ports in the *node_ethernet_portip_ip* sections from the backup configuration file.

d. Next, detect and add the iSCSI storage port candidates by using the **detectiscsistorageportcandidate** and **addiscsistorageport** commands. Make sure that you detect the iSCSI storage ports and add these ports in the same order as you see them in the configuration backup file. If you do not follow the correct order, it might result in a T4 failure. Step c.i must be followed by steps c.ii and c.iii. You must repeat these steps for all the iSCSI sessions that are listed in the backup configuration file exactly in the same order.

1) To detect iSCSI storage ports, determine src_port_id, IO_group_id (optional, not required if the value is 255), target_ipv4/target_ipv6 (the target ip that is not blank is required), iscsi_user_name (not required if blank), iscsi_chap_secret (not required if blank), and site (not required if blank) from the configuration backup file, run the following command:

   `svctask detectiscsistorageportcandidate -srcportid` *src_port_id* `-iogrp` *IO_group_id*
   `-targetip/targetip6` *target_ipv4/target_ipv6* `-username` *iscsi_user_name* `-chapsecret` *iscsi_chap_secret* `-site` *site_id_or_na*

   Where *src_port_id* is the source Ethernet port ID of the configured port, *IO_group_id* is the I/O group ID or name being detected, *target_ipv4/target_ipv6* is the IPv4/IPv6 target iSCSI controller IPv4/IPv6 address, *iscsi_user_name* is the target controller user name

being detected, *iscsi_chap_secret* is the target controller chap secret being detected, and *site_id_or_name* is the specified id or name of the site being detected.

2) Match the discovered *target_iscsiname* with the *target_iscsiname* for this particular session in the backup configuration file by running the **lsiscsistorageportcandidate** command, and use the matching index to add iSCSI storage ports in step c.iii.

   Run the **svcinfo lsiscsistorageportcandidate** command and determine the id field of the row whose *target_iscsiname* matches with the *target_iscsiname* from the configuration backup file. This is your **candidate_id** to be used in step c.iii.

3) To add the iSCSI storage port, determine IO_group_id (optional, not required if the value is 255), site (not required if blank), iscsi_user_name (not required if blank in backup file), and iscsi_chap_secret (not required if blank) from the configuration backup file, provide the *target_iscsiname_index* matched in step c.ii, and then run the following command:

   addiscsistorageport -iogrp *iogrp_id* **-username** *iscsi_user_name* **-chapsecret** *iscsi_chap_sec*

   Where *iogrp_id* is the I/O group ID or name that is added, *iscsi_user_name* is the target controller user name being added, *iscsi_chap_secret* is the target controller chap secret being added, and *site_id_or_name* specified the id or name of the site being added.

4) If the configuration is a HyperSwap® or stretched cluster, the controller name and site needs to be restored. To restore the controller name and site, determine controller_name and controller site_id from the backup xml file by matching the inter_WWPN field with the newly added iSCSI controller, and then run the following command:

   chcontroller -name *controller_name* **-site** *site_id/name controller_id/name*

   Where *controller_name* is the name of the controller from the backup xml file, *site_id/name* is the id/name of the site of iSCSI controller from the backup xml file, and *controller_id/name* is the id or current name of the controller.

14. Issue the following CLI command to compare the current configuration with the backup configuration data file:

    svcconfig restore -prepare

    This CLI command creates a log file in the /tmp directory of the configuration node. The name of the log file is svc.config.restore.prepare.log.

    Note: It can take up to a minute for each 256-MDisk batch to be discovered. If you receive error message CMMVC6200W for an MDisk after you enter this command, all the managed disks (MDisks) might not be discovered yet. Allow a suitable time to elapse and try the **svcconfig restore -prepare** command again.

15. Issue the following command to copy the log file to another server that is accessible to the system:

    pscp superuser@*cluster_ip*:/tmp/svc.config.restore.prepare.log *full_path_for_where_to_copy_log_files*

16. Open the log file from the server where the copy is now stored.

17. Check the log file for errors.

    • If you find errors, correct the condition that caused the errors and reissue the command. You must correct all errors before you can proceed to step 18 on page 167.

- If an error indicates that the system layer will not be restored, return to 8 on page 164, configure the layer setting correctly, and then continue the restore process from 12 on page 164.
- If you need assistance, contact the IBM Support Center.
18. Issue the following CLI command to restore the configuration:

    svcconfig restore -execute

    This CLI command creates a log file in the /tmp directory of the configuration node. The name of the log file is svc.config.restore.execute.log.
19. Issue the following command to copy the log file to another server that is accessible to the system:

    pscp superuser@*cluster_ip*:/tmp/svc.config.restore.execute.log
    full_path_for_where_to_copy_log_files
20. Open the log file from the server where the copy is now stored.
21. Check the log file to ensure that no errors or warnings occurred.

    Note: You might receive a warning that states that a licensed feature is not enabled. This message means that after the recovery process, the current license settings do not match the previous license settings. The recovery process continues normally and you can enter the correct license settings in the management GUI later.

    When you log in to the CLI again over SSH, you see this output:

    IBM_2076:*your_cluster_name*:superuser>

22. After the configuration is restored, verify that the quorum disks are restored to the MDisks that you want by using the **lsquorum** command. To restore the quorum disks to the correct MDisks, issue the appropriate **chquorum** CLI commands.

    Note: If IP Quorum was enabled on the system, it is not recovered automatically as the cluster certificate is regenerated. It is necessary to manually re-enable IP Quorum by downloading a java application from the **Settings**>**System**>**IP Quorum** tab in the GUI, and then installing the application on the host server.

What to do next

You can remove any unwanted configuration backup and restore files from the /tmp directory on your configuration by issuing the following CLI command:

svcconfig clear -all

Deleting backup configuration files using the CLI

You can use the command-line interface (CLI) to delete backup configuration files.

About this task

Complete the following steps to delete backup configuration files:

Procedure

1. Issue the following command to log on to the system:

   plink -i *ssh_private_key_file* superuser@*cluster_ip*

where *ssh_private_key_file* is the name of the SSH private key file for the superuser and *cluster_ip* is the IP address or DNS name of the clustered system from which you want to delete the configuration.

2. Issue the following CLI command to erase all of the files that are stored in the /tmp directory:

   `svcconfig clear -all`

Chapter 8. Storwize V7000 replaceable units

Each Storwize V7000 model consists of several replaceable units. Generic replaceable units are cables, SFP transceivers, canisters, power supply units, battery assemblies, and enclosure chassis. The parts list varies, depending on the generation of your control enclosure model.

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 91. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 92. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Storwize V7000 2076-524 Gen2 replaceable units

You might have to replace a Storwize V7000 2076-524 part. Some replaceable parts are customer-replaceable units (CRUs). Other parts are field-replaceable units (FRUs), which are replaced by IBM trained service technicians.

Table 93. Control enclosure replaceable units

Part number	Part name	CRU or FRU	Notes
31P1854	Control enclosure midplane assembly	FRU	Excludes drives, drive blanks, canisters, fan modules, bezel covers, PSUs.

Table 93. Control enclosure replaceable units (continued)

Part number	Part name	CRU or FRU	Notes
31P1845	Node canister	CRU	Includes 2x 16 GB DIMMs, system drive, and CMOS battery. Excludes node battery, interface adapters, compression adapter / passthrough.
31P1849	Control enclosure power supply unit	CRU	
31P1847	Fan module	CRU	
45W8680	Drive blank, 2.5-inch form factor	CRU	
31P1851	Control enclosure left bezel	CRU	No MTM/Serial number label on the FRU.
00Y2512	2.5-inch enclosure right bezel	CRU	
31P1856	Control enclosure rail kit	CRU	
31P1807	Node battery	CRU	
64P8453	Node canister memory (16 GB DIMM)	CRU	
33F8354	CMOS coin battery	CRU	For real-time clock
31P1861	Compression pass-through adapter	CRU	
31P1863	Compression accelerator	CRU	
64P8473	4-port 8 Gbps Fibre Channel host interface adapter	CRU	No SFPs
00AR316	4-port 10 Gbps Ethernet host interface adapter	CRU	No SFPs.
00WY984	4-port 16 Gbps Fibre Channel host interface adapter	CRU	No SFPs. Before you add this adapter, ensure that the system is running software version 7.6 or later.
00RY007	2-port 16 Gbps Fibre Channel host interface adapter	CRU	No SFPs.
31P1630	8 Gbps SW SFP	CRU	For 8 Gbps Fibre Channel host interface adapter.
00AR096	8 Gbps LW SFP	CRU	For 8 Gbps Fibre Channel host interface adapter
00RY190	16 Gbps SW SFP	CRU	For 16 Gbps Fibre Channel host interface adapter.
00RY191	16 Gbps LW SFP	CRU	For 16 Gbps Fibre Channel host interface adapter
31P1549	10 Gbps SFP	CRU	For 10 Gbps Ethernet host interface adapter.

Table 94. Expansion enclosure replaceable units

Part number	Part name	CRU or FRU	Notes
45W8680	Drive blank, 2.5-inch form factor	CRU	For models 2076-524, 2076-24F only.
64P8446	Expansion enclosure midplane assembly, 12-slot, 3.5-inch	FRU	For model 2076-12F only. Excludes drives, drive blanks, canisters, bezel covers, and PSUs.

Table 94. Expansion enclosure replaceable units (continued)

Part number	Part name	CRU or FRU	Notes
64P8447	Expansion enclosure midplane assembly, 24-slot, 2.5-inch	FRU	For model 2076-24F only. Excludes drives, drive blanks, canisters, bezel covers, and PSUs.
64P8448	Expansion Canister	CRU	
98Y2218	Expansion enclosure power supply unit	CRU	
45W8680	Drive blank, 2.5-inch form factor	CRU	For models 2076-524, 2076-24F only.
42R7992	Drive blank, 3.5-inch form factor	CRU	For model 2076-12F only.
64P8450	Expansion enclosure left bezel	CRU	No MTM/Serial number label on the FRU.
00Y2512	Enclosure right bezel, 2.5-inch form factor	CRU	
00Y2436	Enclosure right bezel, 3.5-inch form factor	CRU	
64P8449	Expansion enclosure rail kit	CRU	

Table 95. Drive replaceable units

Part number	Part name	CRU or FRU	Notes
2.5-inch form factor			
00AR323	SFF HDD - 600 GB 15 K RPM	CRU	
00AR324	SFF HDD - 300 GB 15 K RPM	CRU	
00AR325	SFF HDD - 600 GB 10K RPM	CRU	
00AR326	SFF HDD - 900 GB 10K RPM	CRU	
00AR327	SFF HDD - 1.2 TB 10K RPM	CRU	
00RX908	SFF HDD - 1.8 TB 10K RPM 12 Gbps	CRU	Requires system software version 7.4 or later.
00AR328	SFF HDD - 1 TB NL SAS 7.2 K RPM	CRU	
00AR329	SFF 200 GB SSD	CRU	
00AR330	SFF 400 GB SSD	CRU	
00AR331	SFF 800 GB SSD	CRU	
01EJ594	SFF 1.92 TB read intensive SSD	CRU	
01EJ596	SFF 3.84 TB read intensive SSD	CRU	
3.5-inch form factor			
00AR320	LFF HDD - 2 TB NL SAS	CRU	
00AR321	LFF HDD - 3 TB NL SAS	CRU	
00AR322	LFF HDD - 4 TB NL SAS	CRU	
00RX911	LFF HDD - 6 TB NL 12 Gbps SAS	CRU	Requires system software version 7.4 or later.

Table 96. Cable replaceable units

Part number	Part name	CRU or FRU	Notes
Optical			
39M5699	1 m FC cable	CRU	
39M5700	5 m FC cable	CRU	
39M5701	25 m FC cable	CRU	
41V2120	10 m OM3 FC cable	CRU	
SAS			
00AR272	0.6 m 12 Gbps SAS Cable (mini SAS HD to mini SAS HD)	CRU	For connecting expansion enclosures.
00AR311	1.5 m 12 Gbps SAS Cable (mini SAS HD to mini SAS HD)	CRU	For connecting expansion enclosures.
00AR317	3.0 m 12 Gbps SAS Cable (mini SAS HD to mini SAS HD)	CRU	For connecting expansion enclosures.
00AR439	6.0 m 12 Gbps SAS Cable (mini SAS HD to mini SAS HD)	CRU	For connecting expansion enclosures.
Power			
39M5068	Argentina 2.8 m	CRU	
39M5199	Japan 2.8 m	CRU	
39M5123	Europe 2.8 m	CRU	
39M5165	Italy 2.8 m	CRU	
39M5102	Aus/NZ 2.8 m	CRU	
39M5130	Denmark 2.8 m	CRU	
39M5144	S. Africa 2.8 m	CRU	
39M5151	UK 2.8 m	CRU	
39M5158	Switzerland 2.8 m	CRU	
39M5172	Israel 2.8 m	CRU	
39M5206	China 2.8 m	CRU	
39M5219	Korea 2.8 m	CRU	
39M5226	India 2.8 m	CRU	
39M5240	Brazil 2.8 m	CRU	
39M5247	Taiwan 2.8 m	CRU	
39M5081	US/Canada 2.8 m	CRU	
39M5377	Power jumper cord - 2.8 m	CRU	

Storwize V7000 2076-1xx and 2076-3xx Gen1 replaceable units

The Storwize V7000 system consists of several replaceable units. Generic replaceable units are cables, SFP transceivers, canisters, power supply units, battery assemblies, and enclosure chassis.

Table 97 on page 173 provides a brief description of each replaceable unit.

Table 97. Replaceable units

Part	Part number	Applicable models	FRU or customer replaced
2.5 inch SSDs	00RX920	112, 124, 212, 224 312, 324, 700, 720	FRU
2U24 enclosure chassis (empty chassis)	85Y5897	124, 224, 324	FRU
2U12 enclosure chassis (empty chassis)	85Y5896	112, 212, 312	FRU
Type 100 node canister	85Y5899	112, 124	Customer replaced
Type 300 node canister with 10 Gbps Ethernet ports	85Y6116	312, 324	Customer replaced
Expansion canister	85Y5850	212, 224	Customer replaced
764 W power supply unit	85Y5847	112, 124, 312, 324	Customer replaced
580 W power supply unit	85Y5846	212, 224	Customer replaced
Battery backup unit	85Y5898	112, 124, 312, 324	Customer replaced
1 m SAS cable	44V4041	212, 224	Customer replaced
3 m SAS cable	44V4163	212, 224	Customer replaced
6 m SAS cable	44V4164	212, 224	Customer replaced
1 m Fibre Channel cable	39M5699	112, 124, 312, 324	Customer replaced
5 m Fibre Channel cable	39M5700	112, 124, 312, 324	Customer replaced
25 m Fibre Channel cable	39M5701	112, 124, 312, 324	Customer replaced
1.8 m power cord (Chicago)	39M5080	All	Customer replaced
2.8 m power cord (EMEA)	39M5151	All	Customer replaced
2.8 m power cord (Australia)	39M5102	All	Customer replaced
2.8 m power cord (Africa)	39M5123	All	Customer replaced
2.8 m power cord (Denmark)	39M5130	All	Customer replaced
2.8 m power cord (South Africa)	39M5144	All	Customer replaced
2.8 m power cord (Switzerland)	39M5158	All	Customer replaced
2.8 m power cord (Chile)	39M5165	All	Customer replaced

Table 97. Replaceable units (continued)

Part	Part number	Applicable models	FRU or customer replaced
2.8 m power cord (Israel)	39M5172	All	Customer replaced
2.8 m power cord (Group 1 including the United States)	39M5081	All	Customer replaced
2.8 m power cord (Argentina)	39M5068	All	Customer replaced
2.8 m power cord (China)	39M5206	All	Customer replaced
2.8 m power cord (Taiwan)	39M5247	All	Customer replaced
2.8 m power cord (Brazil)	39M5233	All	Customer replaced
2.0 m jumper cable	39M5376	All	Customer replaced
2.8 m power cord (India)	39M5226	All	Customer replaced
4.3 m power cord (Japan)	39M5200	All	Customer replaced
2.8 m power cord (Korea)	39M5219	All	Customer replaced
2.5" flash drive, 300 GB, in carrier assembly	85Y5861	124, 224, 324	Customer replaced
2.5" 10 K, 300 GB, in carrier assembly	85Y5862	124, 224, 324	Customer replaced
2.5" 10 K, 450 GB, in carrier assembly	85Y5863	124, 224, 324	Customer replaced
2.5" 10 K, 600 GB drive, in carrier assembly	85Y5864	124, 224, 324	Customer replaced
2.5" 15 K, 146 GB drive, in carrier assembly	85Y6088	124, 224, 324	Customer replaced
2.5" 15 K, 300 GB drive, in carrier assembly	85Y6185	124, 224, 324	Customer replaced
2.5" 10 K, 900 GB drive, in carrier assembly	00L4680	124, 224, 324	Customer replaced
2.5" 10 K, 1.2 TB SAS drive, in carrier assembly	85Y6156	124, 224, 324	Customer replaced
2.5" 10 K, 300 GB drive, in carrier assembly	85Y6256	124, 224, 324	Customer replaced
2.5" 10 K, 600 GB drive, in carrier assembly	85Y6268	124, 224, 324	Customer replaced
2.5" 10 K, 900 GB drive, in carrier assembly	85Y6274	124, 224, 324	Customer replaced
2.5" 10 K 1.8 TB drive, in carrier assembly	00RX915	124, 224, 324	Customer replaceable. Requires system software version 7.4 or later.

Table 97. Replaceable units (continued)

Part	Part number	Applicable models	FRU or customer replaced
2.5" 7.2 K, Nearline SAS, 1 TB drive, in carrier assembly	85Y6186	124, 224, 324	Customer replaced
2.5" flash drive, 200 GB drive, in carrier assembly	85Y6188	124, 224, 324	Customer replaced
2.5" flash drive, 400 GB drive, in carrier assembly	85Y6189	124, 224, 324	Customer replaced
2.5" flash drive, 800 GB drive, in carrier assembly	00AR252	124, 224, 324	Customer replaced
3.5" 7.2 K Nearline SAS - 2 TB in carrier assembly	85Y5869	112, 212, 312	Customer replaced
3.5" 7.2 K Nearline SAS - 3 TB in carrier assembly	85Y6187	112, 212, 312	Customer replaced
3.5" 7.2 K Nearline SAS - 6 TB in carrier assembly	00RX918	112, 212, 312	Customer replaceable. Requires system software version 7.4 or later.
Blank 2.5" carrier	85Y5893	124, 224, 324	Customer replaced
Blank 3.5" carrier	85Y5894	112, 212, 312	Customer replaced
Fibre Channel shortwave small form-factor pluggable (SFP)	85Y5958	112, 124, 312, 324	Customer replaced
Fibre Channel longwave small form-factor pluggable (SFP)	85Y5957	112, 124, 312, 324	Customer replaced
Ethernet small form-factor pluggable (SFP)	31P1549	312, 324	Customer replaced
Rail kit	85Y5852	All	Customer replaced
Left enclosure cap including RID tag but no black MTM label	85Y5901	All	Customer replaced
Right enclosure cap (2U12)	85Y5903	112, 212, 312	Customer replaced
Right enclosure cap (2U24)	85Y5904	124, 224, 324	Customer replaced

Chapter 9. Replacing parts

You can remove and replace customer-replaceable units (CRUs) in control enclosures or expansion enclosures.

Attention: Even though many of these components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Each replaceable unit has its own removal procedure. Sometimes you can find that a step within a procedure might refer you to a different remove and replace procedure. You might want to complete the new procedure before you continue with the first procedure that you started.

Remove or replace parts only when you are directed to do so.

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

Preparing to remove and replace parts

Before you remove and replace parts, you must be aware of all safety issues.

Before you begin

First, read the safety precautions in the *IBM Systems Safety Notices*. These guidelines help you safely work with the Storwize V7000.

Replacing a node canister

Remove and replace a node canister.

Replacing a Storwize V7000 Gen2 node canister

To replace a faulty node canister with a new one received from CRU or FRU stock, use this procedure. When replacing a node canister, aim to maximize drive and system availability by maintaining one online node in the control enclosure with the faulty node canister. If you cannot maintain at least one node canister online in the system, then you might need to follow the system recovery procedure after replacing the faulty node canister.

Procedure

1. Follow "Procedure: Removing a Storwize V7000 Gen2 node canister" on page 129 to remove the faulty node canister.
2. Remove the lid of the faulty canister, as described in "Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister" on page 139. Do the same to the replacement canister.
3. Some components inside the faulty node canister must be transferred to the replacement canister. Transfer each of the following components as necessary:

- The battery, as described in "Replacing the battery in a Storwize V7000 Gen2 node canister" on page 198.
- The host interface adapters in one control enclosure, as described in "Replacing a Storwize V7000 Gen2 host interface adapter" on page 254.
- Memory modules, as described in "Replacing a Storwize V7000 Gen2 node canister memory module (16 GB DIMM)" on page 253.
- The compression pass-through adapter or compression accelerator, as described in the installation description of upgrading the hardware to install the compression accelerator.

4. Replace the lid of the faulty canister and the lid of the replacement canister, as described in "Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister" on page 139.
5. Open the release lever of the replacement canister.
6. Push the replacement canister into the slot of the enclosure that the faulty canister was removed from, until it stops.
7. Finish inserting the replacement canister by closing its release lever so that the orange latch engages the enclosure.
8. If the enclosure is powered and the canister is correctly installed, the canister starts automatically. Repeat from step 5, if the canister is not correctly installed.
9. Reinsert the data cables into the ports that they were originally connected.
10. If no node canisters are online, the system is not online. To recover the system in the case when no node canisters are online, see Chapter 7, "Recover system procedure," on page 149.
11. If only the replacement node is in service state with node error 503, apply "Procedure: Rescuing Storwize V7000 Gen2 node canister software from another node (node rescue)" on page 138 to rescue the replacement node canister.
12. When the node canister is powered up, it is automatically added to the system and the system automatically ensures that the machine code version on the new canister matches that of the other node canister in the control enclosure. This is reflected in the system event log.
13. When the canister is back online, check the event log for any new events that might indicate a problem with the reassembly.

Replacing a Storwize V7000 Gen1 node canister

This topic describes how to replace a node canister.

About this task

Attention: Although many components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

Attention: Do not replace one type of node canister with another type. For example, do not replace a model 2076-112 node canister with a model 2076-312 node canister.

Be aware of the following canister LED states:

- If both the power LED and system status LED are on, do not remove a node canister unless directed to do so by a service procedure.
- If the system status is off, it is acceptable to remove a node canister. However, do not remove a node canister unless directed to do so by a service procedure.
- If the power LED is flashing or off, it is safe to remove a node canister. However, do not remove a node canister unless directed to do so by a service procedure.

Attention: Even if a node canister is powered off, it is still possible to lose data. Do not remove a node canister unless directed to do so by a service procedure.

To replace the node canister, perform the following steps:

Procedure

1. Read the safety information to which "Preparing to remove and replace parts" on page 177 refers.
2. Confirm that you know which canister to replace. Go to "Procedure: Identifying which Storwize V7000 Gen1 enclosure or canister to service" on page 103.
3. Record which data cables are plugged into the specific ports of the node canister. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
4. Disconnect the data cables for each canister.
5. Grasp the handle between the thumb and forefinger.

   Note: Ensure that you are opening the correct handle. The handle locations for the node canisters and expansion canisters are slightly different.

   Handles for the node canisters are located in close proximity to each other. The handle with the finger grip on the right removes the upper canister (**1**). The handle with the finger grip on the left removes the lower canister (**2**).
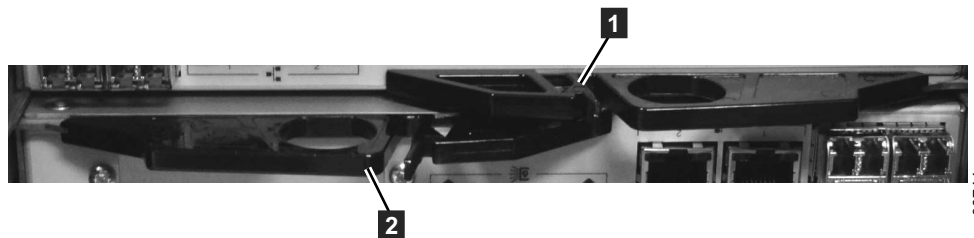
Figure 45. Rear of node canisters that shows the handles.

6. Squeeze them together to release the handle.
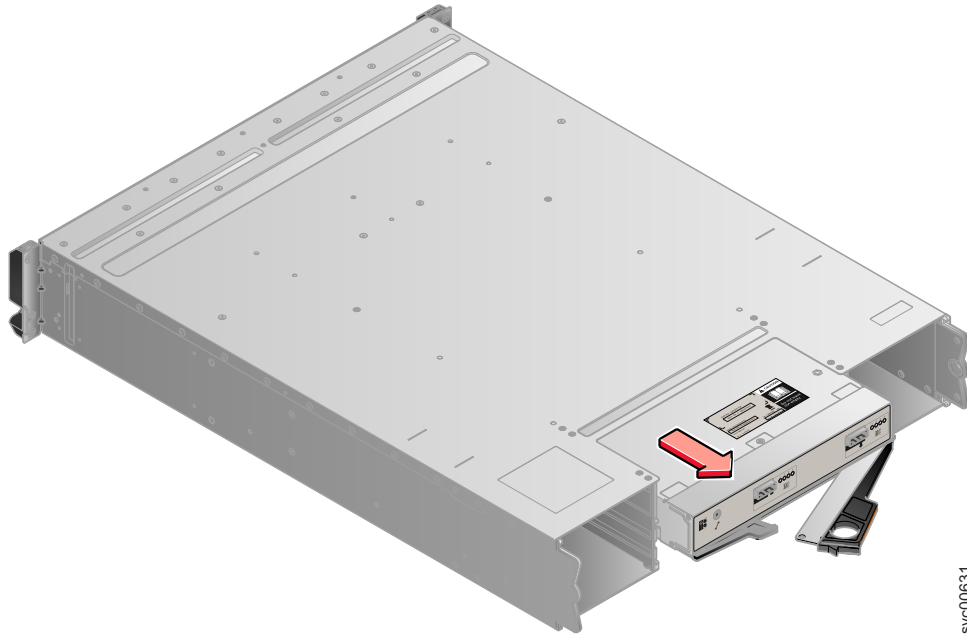
Figure 46. Removing the canister from the enclosure

7. Pull out the handle to its full extension.
8. Grasp canister and pull it out.
9. Insert the new canister into the slot with the handle pointing towards the center of the slot. Insert the unit in the same orientation as the one that you removed.
10. Push the canister back into the slot until the handle starts to move.
11. Finish inserting the canister by closing the handle until the locking catch clicks into place. Make sure that the canister is firmly and correctly seated, otherwise this can lead to problems.

    If the enclosure is powered on, the canister starts automatically.
12. Reattach the data cables.

Replacing a fan module

Remove and replace a fan module.

Replacing a Storwize V7000 Gen2 fan module

Use this procedure to replace a faulty fan module with a new one received from CRU or FRU stock.

About this task

A fan module is located behind each node canister, and is accessed using the node canister slot after the node canister is removed.

Do not remove the node canister and faulty fan module before the replacement fan is on hand. The replacement procedure described must be completed within 5 minutes of the faulty fan module being removed to ensure that components do not shut down due to excessive temperatures.

When removing a node canister, aim to maximize drive and system availability by maintaining one online node in the control enclosure. If you cannot maintain at least one node canister online in the system, then you might need to follow the system recovery procedure after the node canister is replaced into the enclosure.

Procedure

1. Remove the replacement fan module from its packaging. Familiarize yourself with the part by reading through this procedure.
2. Remove the node canister that is on the same side of the enclosure as the faulty fan module. See "Procedure: Removing a Storwize V7000 Gen2 node canister" on page 129.
3. Locate the two orange locking rings of the fan module inside the left and right edges of the node canister slot. Note their position relative to the inside of the canister slot.
4. Simultaneously rotate both rings upwards through 90 degrees, releasing the fan module from the slot. Pull the locking rings to slide the faulty fan module out from the canister slot.
5. Ensure that the orange locking rings on the replacement fan module are rotated open so that they extend out from the fan module.
6. Slide the replacement fan module into the canister slot until it stops.
7. Simultaneously rotate both locking rings downwards through 90 degrees while applying gentle pressure to push the fan module into the slot. The fan module is installed correctly when the back edges of the locking rings are flush with the relief detail inside the canister slot.
8. Replace the node canister into the canister slot until it stops.
9. Finish inserting the node canister by closing its release lever so that the orange catch engages the enclosure.
10. If the enclosure is powered and the canister is correctly installed, the canister starts automatically. Remove the canister and repeat the procedure from step 5, if the canister is not correctly installed.
11. Reinsert the data cables into the ports that they were originally connected.
12. When the canister is back online, check the event log for any new events that might indicate a problem with the reassembly.

Replacing an expansion canister

Remove and replace an expansion canister.

Replacing a Storwize V7000 Gen2 expansion canister

To replace a faulty expansion canister with a new one received from CRU / FRU stock, use this procedure.

About this task

Attention: Although many components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

Do not remove an expansion canister unless directed to do so by a service procedure.

To replace an expansion canister, do the following steps:

Procedure

1. Read the safety information in "Preparing to remove and replace parts" on page 177.
2. Refer to "Procedure: Understanding Storwize V7000 Gen2 volume dependencies" on page 140 to determine whether to do this procedure.
3. Carefully identify the expansion canister that you are replacing. If possible, go to **Monitoring** > **System** in the management GUI. Select the expansion enclosure that you are replacing and select **Actions** > **Identify** to set the canister fault LED blinking.
4. Record which SAS cables are plugged into the specific ports of the expansion canister. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
5. Disconnect the SAS cables from the canister.
6. Open the two release levers as shown in Figure 47 on page 183. The canister moves out of the slot approximately 0.6 cm (0.25 inch).
7. Slide the canister out of the slot.
8. Open the release levers of the replacement canister.
9. Push the replacement canister into the slot until it stops.
10. Finish inserting the canister by closing both release levers so that both orange latches click into place.
11. The canister is correctly installed when the rear face of the canister is flush with the rear edge of the enclosure.

    If the enclosure is powered on and the canister is correctly installed, the canister starts automatically.
12. Reattach each SAS cable into the port from which it was removed in step 5.

    a. Ensuring the SAS cable connectors are inserted with the pull tab to the bottom of the connector, gently push the connector in until a slight click is felt or heard.

    b. Verify that the connector is fully inserted by gently pulling on it (not on the tab).

    You should not be able to remove it.

    If the enclosure is powered on and the SAS connector is correctly inserted into the port, the green SAS link LED above the port lights up.
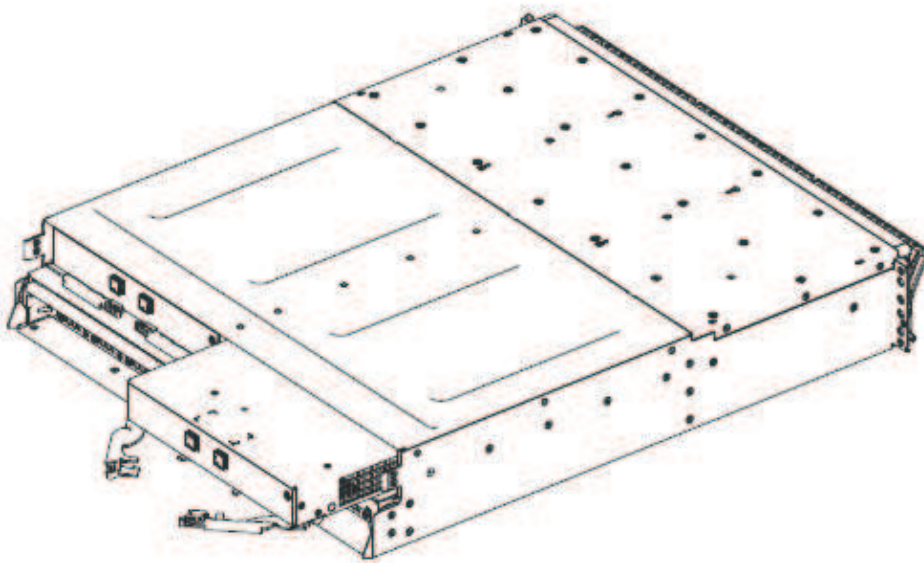
Figure 47. Removing and replacing the Storwize V7000 Gen2 expansion canister

Replacing an expansion canister

This topic describes how to replace an expansion canister.

About this task

Attention: Although many components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

Be aware of the following canister LED states:
- If the power LED is on, do not remove an expansion canister unless directed to do so by a service procedure.
- If the power LED is flashing or off, it is safe to remove an expansion canister. However, do not remove an expansion canister unless directed to do so by a service procedure.

Attention: Even if an expansion canister is powered off, it is still possible to lose data. Do not remove an expansion canister unless directed to do so by a service procedure.

To replace an expansion canister, perform the following steps:

Procedure

1. Read the safety information to which "Preparing to remove and replace parts" on page 177 refers.
2. Record which SAS cables are plugged into the specific ports of the expansion canister. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
3. Disconnect the SAS cables for each canister.
4. Grasp the handle between the thumb and forefinger.

    Note: Ensure that you are opening the correct handle. The handle locations for the node canisters and expansion canisters are slightly different.

    Handles for the upper and lower expansion canisters overlap each other. The handle with the finger grip on the left removes the upper canister (**1**). The handle with the finger grip on the right removes the lower canister (**2**).
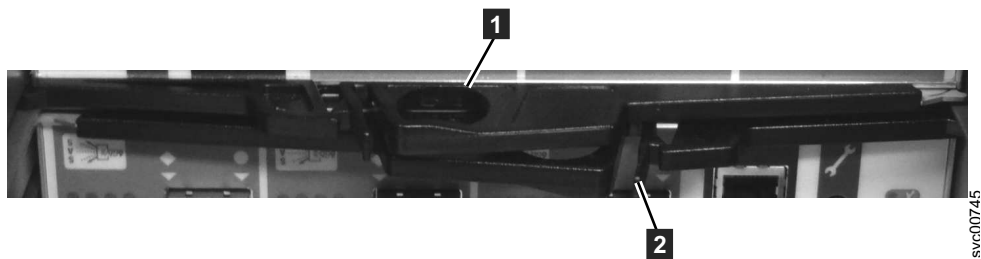
Figure 48. Rear of expansion canisters that shows the handles.

5. Squeeze them together to release the handle.
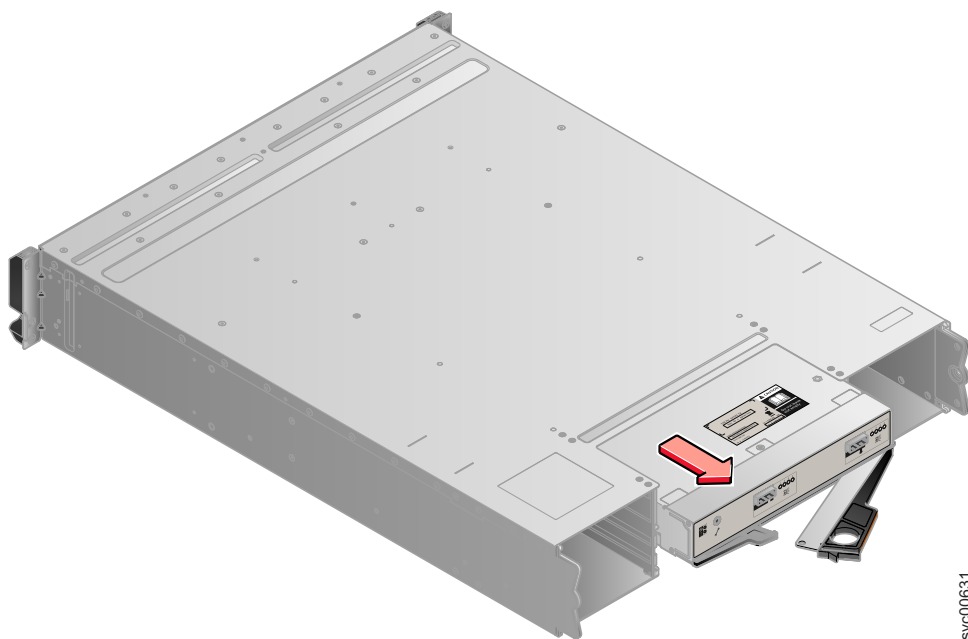
Figure 49. Removing the canister from the enclosure

6. Pull out the handle to its full extension.
7. Grasp canister and pull it out.

8. Insert the new canister into the slot with the handle pointing towards the center of the slot. Insert the unit in the same orientation as the one that you removed.

9. Push the canister back into the slot until the handle starts to move.

10. Finish inserting the canister by closing the handle until the locking catch clicks into place. Make sure that the canister is firmly and correctly seated, otherwise this can lead to problems.

11. Reattach the SAS cables.

Replacing an SFP transceiver

Remove and replace an SFP transceiver.

Replacing an SFP transceiver in a Storwize V7000 2076-524 control enclosure

When a failure occurs on an optical link, the SFP transceiver in the port that provides the link might need to be replaced. To replace a faulty SFP transceiver with a new one received from CRU or FRU stock, use this procedure.

Before you begin

Although many components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

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)

Attention: When you replace this part, you must follow recommended procedures for handling electrostatic discharge (ESD) sensitive devices.

Important: For correct operation, use the correct SFP transceivers with each adapter card. The topic "Storwize V7000 2076-524 Gen2 replaceable units" identifies the suitable IBM parts.

- Use only 8G bps SFP transceivers in the 8 Gbps Fibre Channel adapter cards.
- Use only 16 Gbps SFP transceivers in the 16 Gbps Fibre Channel adapter cards.
- Use only 10 Gbps SFP transceivers in the 10 Gbps Ethernet (FCoE/iSCSI) adapter card.

Procedure

Complete the following steps to remove and then replace an SFP transceiver.

1. Carefully determine the failing physical port connection.

Important: Removing the wrong SFP transceiver might result in loss of data access.

2. Remove the cable from the SFP.

   Figure 50 illustrates an SFP transceiver.

Figure 50. SFP transceiver

3. Remove the faulty SFP transceiver from its aperture.
   a. Unclip the handle of the SFP transceiver.
   b. Pull on the handle of the SFP transceiver.
   c. The SFP transceiver slides out of its slot.
4. Install the replacement SFP transceiver into the aperture that is vacated in step 3.
   a. Open the lock on the replacement SFP transceiver.
   b. Push the new SFP transceiver into the aperture until it stops.
   c. Close the release handle.
   d. Gently pull the SFP transceiver. If it is installed correctly, it does not move from its aperture.
5. Reconnect the optical cable.
6. Confirm that the error is now fixed. Either mark the error as fixed or restart the node, depending on the failure indication originally noted.

Replacing an SFP transceiver in a control enclosure

When a failure occurs on a single link, the SFP transceiver might need to be replaced.

Before you begin

Although many components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

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)

About this task

Complete the following steps to remove and then replace an SFP transceiver:

Procedure

1. Carefully determine the failing physical port connection.

   Important: The Fibre Channel links in the enclosures are supported by both longwave SFP transceivers and shortwave SFP transceivers. A longwave SFP transceiver has some blue components that are visible even when the SFP transceiver is plugged in. You must replace an SFP transceiver with the same type of SFP transceiver that you are replacing. If the SFP transceiver to replace is a longwave SFP transceiver, for example, you must replace with another longwave SFP transceiver. Removing the wrong SFP transceiver might result in loss of data access.

2. Remove the optical cable by pressing the release tab and pulling out the cable. Be careful to exert pressure only on the connector and do not pull on the optical cables.

3. Remove the SFP transceiver. There are a number of different handling or locking mechanisms that are used on the SFP transceivers. Some SFP transceivers might have a plastic tag. If so, pull the tag to remove the SFP transceiver.

   Important: Always check that the SFP transceiver that you replace matches the SFP transceiver that you remove.

4. Push the new SFP transceiver into the aperture and ensure that it is securely pushed home. The SFP transceiver usually locks into place without having to swing the release handle until it locks flush with the SFP transceiver. Figure 51 illustrates an SFP transceiver and its release handle.

Figure 51. SFP transceiver

5. Reconnect the optical cable.
6. Confirm that the error is now fixed. Either mark the error as fixed or restart the node, depending on the failure indication that you originally noted.

Replacing a power supply unit for a control enclosure

You might need to remove and replace the power supply units in the control enclosure.

Replacing a Storwize V7000 Gen2 power supply unit for a control enclosure

You can replace either of the two hot-swap redundant power supplies in an enclosure. These redundant power supplies operate in parallel, one continuing to power the enclosure if the other fails.

Before you begin

Attention:
- Although many components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.
- Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.
- Ensure that you are aware of the procedures for handling static-sensitive devices before you replace the power supply.

About this task

To replace the power supply, do the following steps:

Procedure

1. Read the safety information in "Preparing to remove and replace parts" on page 177.
2. Confirm that you know which power supply must be replaced. Go to "Procedure: Identifying which Storwize V7000 Gen2 enclosure or canister to service" on page 102.
3. Disconnect the power cord from the electrical outlet. Release the cable retention clip and disconnect the power cord from the power supply that you are replacing.
4. Locate the orange release tab at the top edge of the power supply unit. Press the release tab gently until it stops.
5. Using the handle, firmly pull the power supply out of the enclosure that is shown in Figure 52 on page 189.
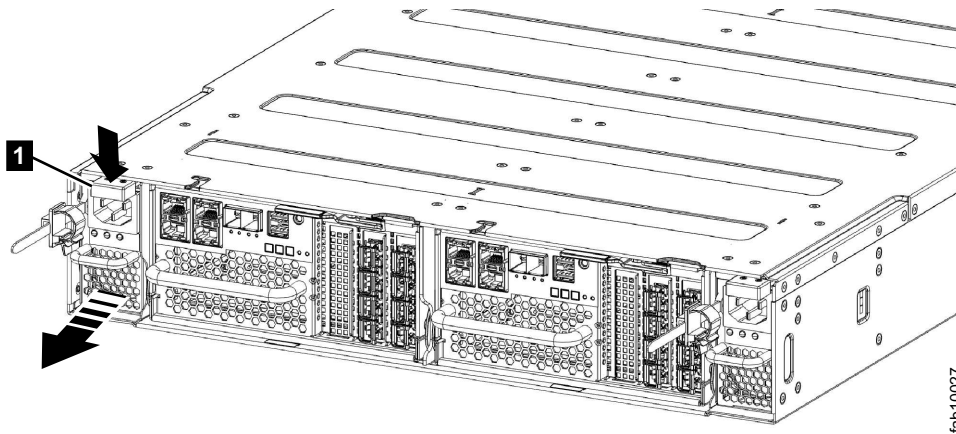
Figure 52. Removing the power supply unit (left side of enclosure)

6. Hold the new power supply so that the handle is fully extended.
7. Slide the power supply into the enclosure until it stops. Push it firmly into position until it clicks.
8. Connect the power cord to the power supply and to a properly grounded electrical outlet. Secure the cable with the cable retention clip on the rear of the power supply unit.

    Note: After the power cord is connected to the electrical outlet, make sure that the ac and dc power (green) LEDs are lit and the fault (amber) LED is off.

Replacing a Storwize V7000 Gen1 power supply unit for a control enclosure

You can replace either of the two 764-watt hot-swap redundant power supplies in the control enclosure. These redundant power supplies operate in parallel, one continuing to power the canister if the other fails.

Before you begin

DANGER

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

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

* If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
* Do not open or service any power supply assembly.
* Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
* The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
* 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.

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

Attention: Although many components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Attention: A powered-on enclosure must not have a power supply removed for more than five minutes because the cooling does not function correctly with an empty slot. Ensure that you have read and understood all these instructions and have the replacement available, and unpacked, before you remove the existing power supply.

Although many components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

Attention: In some instances, it might not be advisable to remove a power supply unit when a system is performing I/O. For example, the charge in the backup battery might not be sufficient enough within the partner power-supply unit to continue operations without causing a loss of access to the data. Wait until the partner battery is 100% charged before you replace the power supply unit.

Ensure that you are aware of the procedures for handling static-sensitive devices before you replace the power supply.

About this task

A replacement power supply unit is not shipped with a battery; therefore, transfer the battery from the existing power supply unit to the replacement unit. To transfer a battery, go to "Replacing a battery in a Storwize V7000 Gen1 power supply unit" on page 200.

To replace the power supply, perform the following steps:

Procedure

1. Read the safety information to which "Preparing to remove and replace parts" on page 177 refers.
2. Examine the Identify LED that is lit on the front of the enclosure to identify the correct enclosure.
3. Turn off the power to the power supply unit by using the switch at the back.
4. Disconnect the cable retention bracket and the power cord from the power supply that you are replacing.
5. Remove the power supply unit. Record the orientation of the power supply unit. Power supply unit 1 is top side up, and power supply unit 2 is inverted.
   a. Depress the black locking catch from the side with the colored sticker as shown in Figure 53 on page 192.
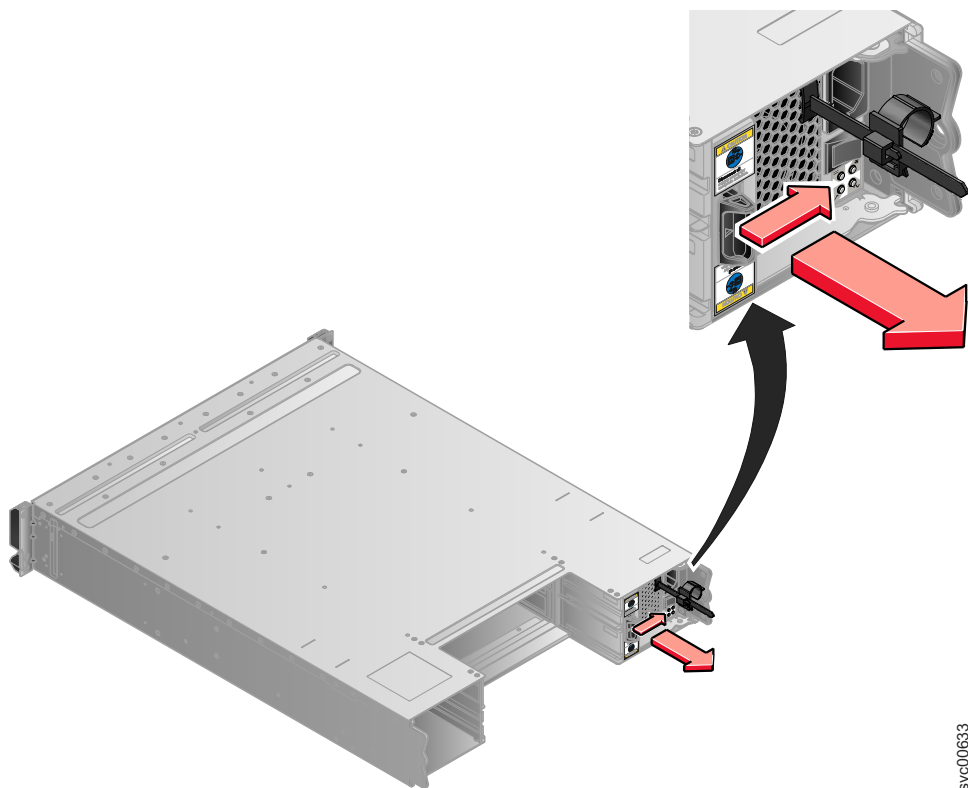
Figure 53. Directions for lifting the handle on the power supply unit

     b. Grip the handle to pull the power supply out of the enclosure as shown in Figure 54.
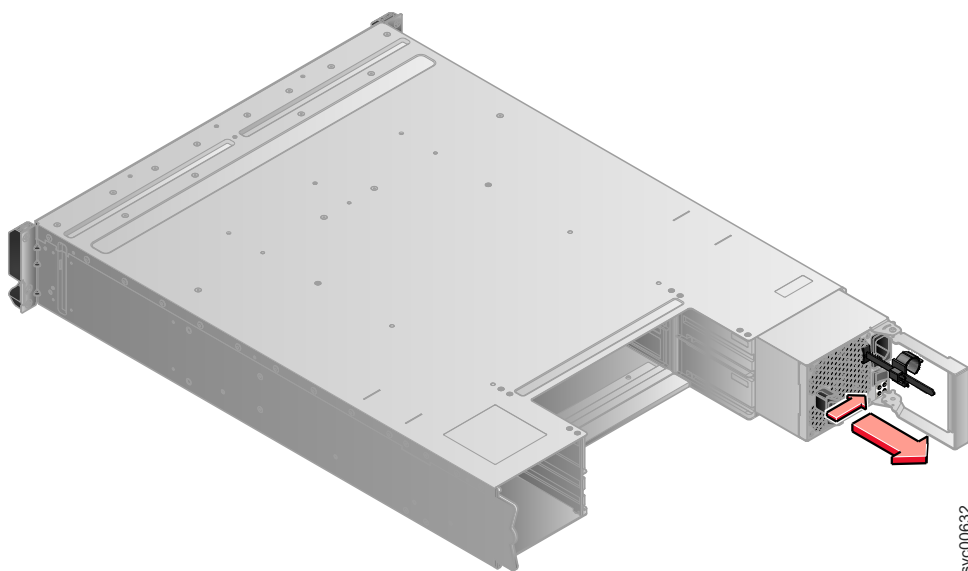
Figure 54. Using the handle to remove a power supply unit

    6. Insert the replacement power supply unit into the enclosure with the handle pointing towards the center of the enclosure. Insert the unit in the same orientation as the one that you removed.

7. Push the power supply unit back into the enclosure until the handle starts to move.
8. Finish inserting the power supply unit into the enclosure by closing the handle until the locking catch clicks into place.
9. Reattach the power cable and cable retention bracket.
10. Turn on the power switch to the power supply unit.

What to do next

If required, return the power supply. Follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Replacing a power supply unit for an expansion enclosure

Remove and replace the hot-swap redundant power supplies in the expansion enclosure.

Replacing a power supply unit for a Storwize V7000 Gen2 expansion enclosure

You can replace either of the two hot-swap redundant power supplies in an enclosure. These redundant power supplies operate in parallel, one continuing to power the canister if the other fails.

Before you begin

Attention:
- Although many components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.
- Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.
- Ensure that you are aware of the procedures for handling static-sensitive devices before you replace the power supply.

About this task

To replace the power supply, do the following steps:

Procedure
1. Read the safety information in "Preparing to remove and replace parts" on page 177.
2. Confirm that you know which power supply must be replaced. Go to "Procedure: Identifying which Storwize V7000 Gen2 enclosure or canister to service" on page 102.
3. Disconnect the power cord from the electrical outlet. Release the cable retention clip and disconnect the power cord from the power supply that you are replacing.

4. On the left side of the power supply, press the orange release tab to the right slightly to release the handle (no more than 6 mm [0.25 in.]) as you rotate the handle downward.

5. Using the handle, gently slide the power supply out of the enclosure, as shown in Figure 55.
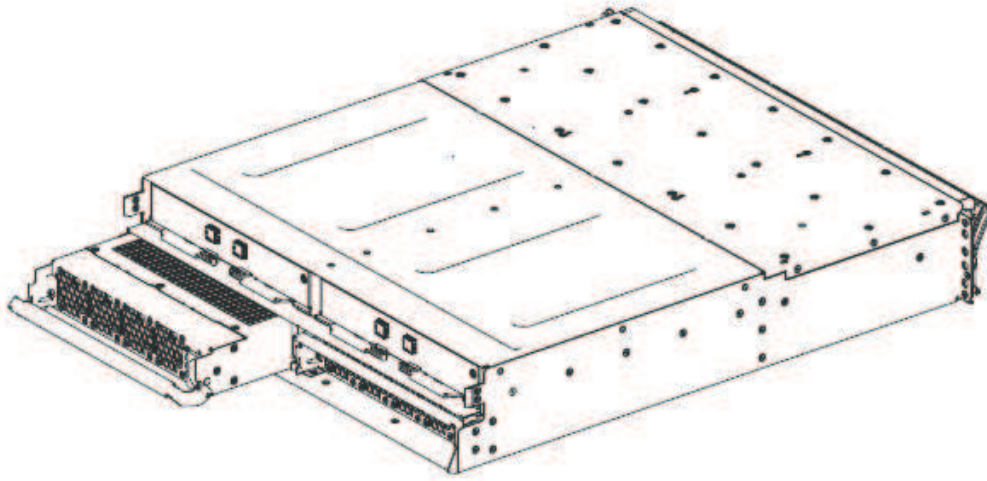
Figure 55. Removing the power supply unit from the left side of the expansion enclosure

6. Hold the new power supply so that the handle is fully extended.

7. Slide the power supply into the enclosure until it stops. Rotate the handle upward into the closed position until it clicks.

8. Hold the new power supply so that the handle is fully extended.

9. Connect the power cord to the power supply and to a properly grounded electrical outlet.

   Note: After the power cord is connected to the electrical outlet, make sure that the ac and dc power (green) LEDs are lit and the fault (amber) LED is off.

Replacing a power supply unit for an expansion enclosure

You can replace either of the two 580-watt hot-swap redundant power supplies in the expansion enclosure. These redundant power supplies operate in parallel, one continuing to power the canister if the other fails.

Before you begin

DANGER

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

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

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- 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.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

Attention: Although many components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Attention: A powered-on enclosure must not have a power supply removed for more than five minutes because the cooling does not function correctly with an empty slot. Ensure that you have read and understood all these instructions and have the replacement available, and unpacked, before you remove the existing power supply.

Although many components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

Ensure that you are aware of the procedures for handling static-sensitive devices before you replace the power supply.

About this task

To replace the power supply unit in an expansion enclosure, complete the following steps:

Procedure

1. Read the safety information to which "Preparing to remove and replace parts" on page 177 refers.
2. Examine the Identify LED that is lit on the front of the enclosure to identify the correct enclosure.
3. Turn off the power to the power supply unit by using the switch at the back of the unit.
4. Disconnect the cable retention bracket and the power cord from the power supply that you are replacing.
5. Remove the power supply unit. Record the orientation of the power supply unit. Power supply unit 1 is top side up, and power supply unit 2 is inverted.
   a. Depress the black locking catch from the side with the colored sticker as shown in Figure 56 on page 197.
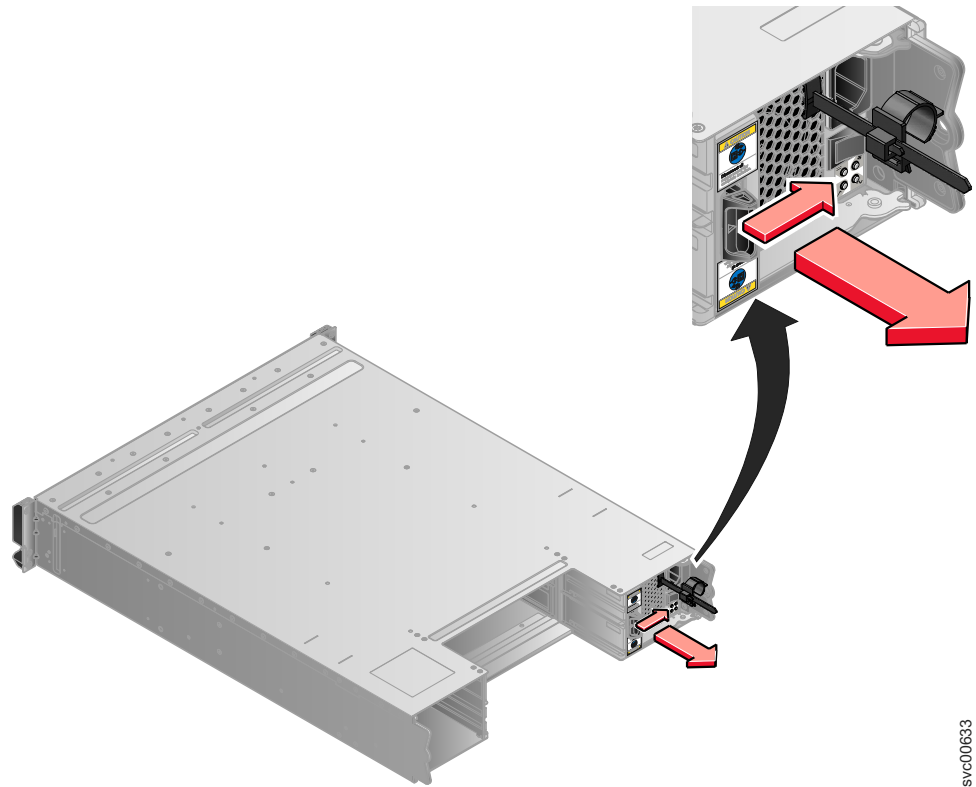
Figure 56. Directions for lifting the handle on the power supply unit

     b. Grip the handle to pull the power supply out of the enclosure as shown in Figure 57.
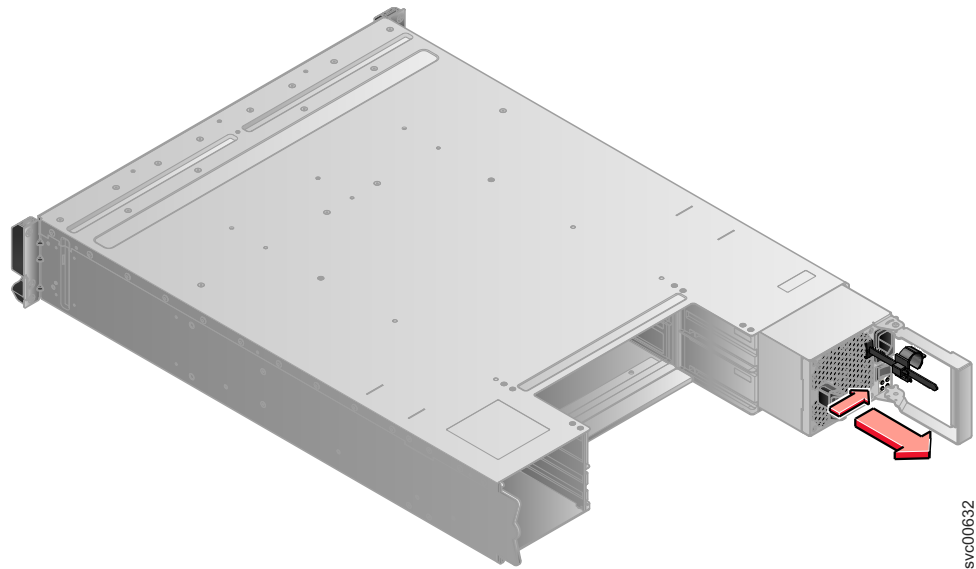
Figure 57. Using the handle to remove a power supply unit

6. Insert the replacement power supply unit into the enclosure with the handle pointing towards the center of the enclosure. Insert the unit in the same orientation as the one that you removed.

7. Push the power supply unit back into the enclosure until the handle starts to move.

8. Finish inserting the power supply unit in the enclosure by closing the handle until the locking catch clicks into place.

9. Reattach the power cable and cable retention bracket.

10. Turn on the power switch to the power supply unit.

What to do next

If required, return the power supply. Follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Replacing the battery in a node canister

Remove and replace the battery in a node canister.

Replacing the battery in a Storwize V7000 Gen2 node canister

To replace a faulty battery with a new one received from customer replaceable unit (CRU) or field replaceable unit (FRU) stock, use this procedure.

About this task

CAUTION:
The battery is a lithium ion battery. To avoid possible explosions, do not burn. Exchange only with the approved part. Recycle or discard the battery as instructed by local regulations. (C007a)

To replace a battery:

Procedure

1. Identify the node canister with the faulty battery by following the procedure "Procedure: Understanding the Storwize V7000 Gen2 system status from the LEDs" on page 107.

2. Follow the procedure "Procedure: Removing a Storwize V7000 Gen2 node canister" on page 129 to remove the node canister with the faulty battery.

3. Open the lid of the canister, as described in "Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister" on page 139.

4. Spread the two blue battery latches outward as shown in Figure 58 on page 199. Raise open both latching arms of the battery simultaneously to disconnect the battery.
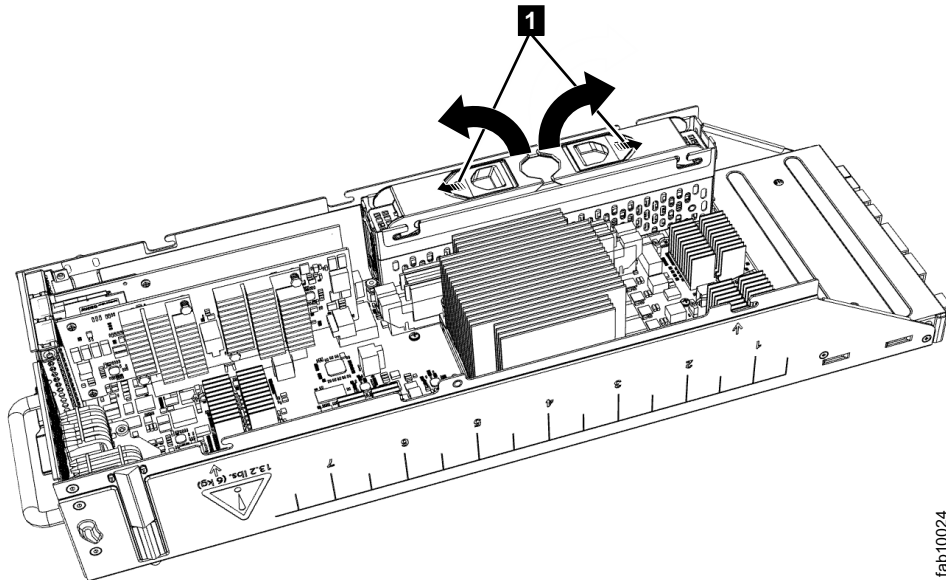
Figure 58. Opening latching arms to disconnect a Storwize V7000 Gen2 node canister battery

5. Holding the battery by its latching arms, lift it from its cradle. Place the battery in a safe place.

6. Remove the replacement battery from its package.

7. Open the latching arms of the replacement battery, then place the replacement battery into the battery cradle of the node canister so that the connectors align.

8. Apply gentle downward pressure to both battery latches so that the battery is drawn into the battery cradle. Ensure that both latches are fully engaged by spreading the two blue latches outward while you apply gentle downwards pressure.

9. Replace the canister lid, as described in "Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister" on page 139.

10. Reinstall the canister into the enclosure from which it was removed in step 2 on page 198.

    Notes:
    - If the replacement battery is sufficiently charged, the node canister starts.
    - If the replacement battery is not sufficiently charged, the node canister does not come online. The battery continues to charge.

11. Refer to "Procedure: Understanding the Storwize V7000 Gen2 system status from the LEDs" on page 107 to understand the charge level of the replacement battery. If the canister did not restart, use the management GUI to monitor the canister and battery status.

12. Reconnect the cables to the canister, ensuring that each cable goes into the same port from which it was removed in step 2 on page 198.

13. When the canister is back online, check the event log for any new events that might indicate a problem with the reassembly.

Replacing a battery in a power supply unit

Remove and replace the battery in a control enclosure power-supply unit.

Replacing a battery in a Storwize V7000 Gen1 power supply unit

This topic describes how to replace the battery in the control enclosure power-supply unit.

Before you begin

DANGER

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

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

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- 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.

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

CAUTION:
The battery is a lithium ion battery. To avoid possible explosion, do not burn. (C007)

Attention: Although many components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Although many components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

About this task

Each power supply unit in a control enclosure contains an integrated battery that is used during temporary short-term power outages. You must replace the battery with the exact same model.

To replace the battery in the power supply unit of the control enclosure, perform the following steps:

Procedure
1. Read the safety information to which "Preparing to remove and replace parts" on page 177 refers.
2. Follow the removing steps of the replacing a power-supply unit procedure. Go to "Replacing a Storwize V7000 Gen1 power supply unit for a control enclosure" on page 189.
3. Remove the battery, as shown in Figure 59 on page 202.
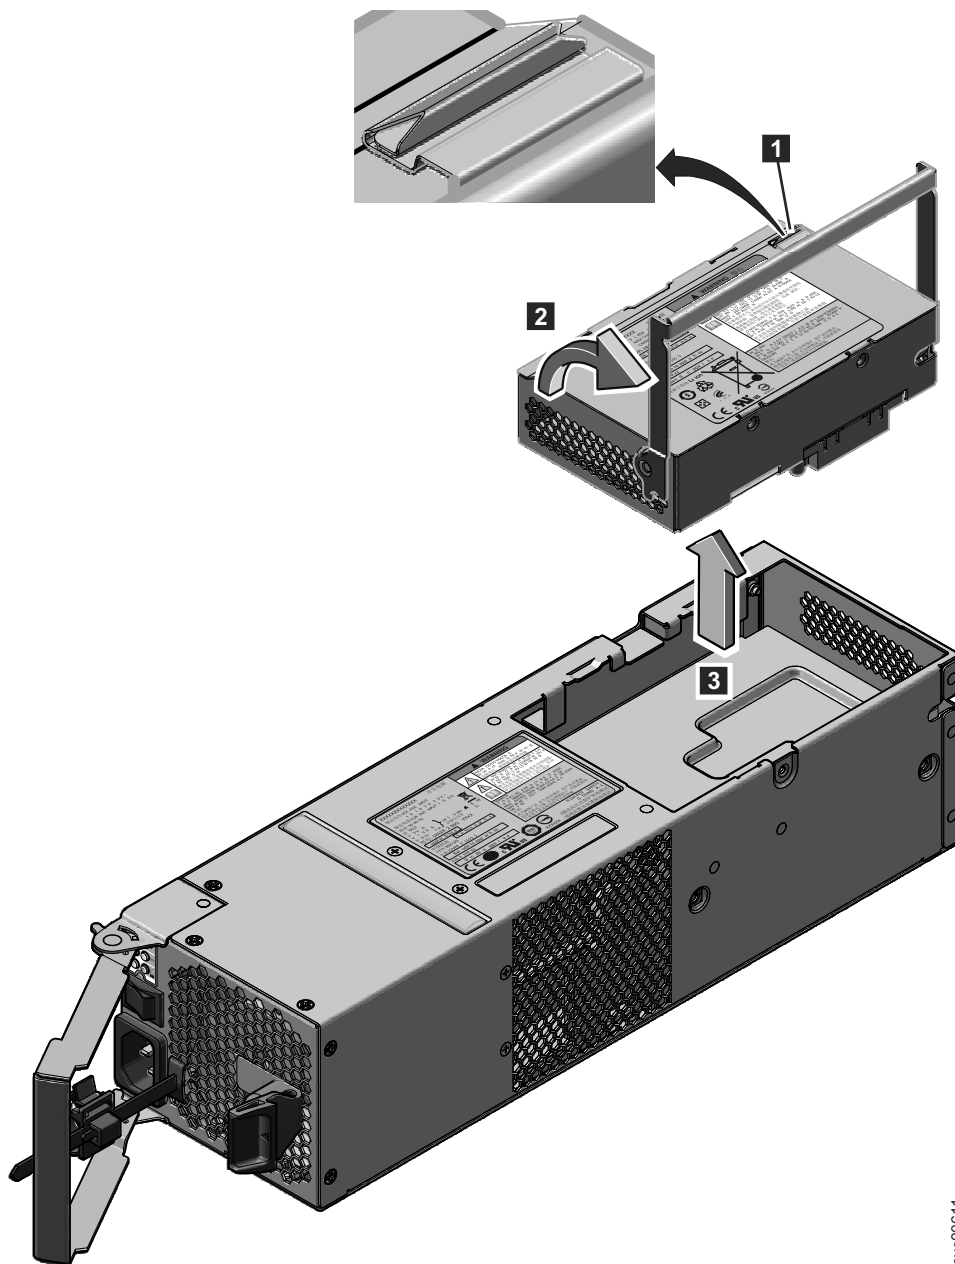
Figure 59. Removing the battery from the control enclosure power-supply unit

    a. Press the catch to release the handle **1** .

    b. Lift the handle on the battery **2** .

    c. Lift the battery out of the power supply unit **3** .

4. Install the replacement battery.

    Attention: The replacement battery has protective end caps that must be removed prior to use.

    a. Remove the battery from the packaging.

    b. Remove the end caps.

    c. Attach the end caps to both ends of the battery that you removed and place the battery in the original packaging.

d. Place the replacement battery in the opening on top of the power supply in its proper orientation.

e. Press the battery to seat the connector.

f. Place the handle in its downward location

5. Push the power supply unit back into the enclosure until the handle starts to move.

6. Finish inserting the power supply unit into the enclosure by closing the handle until the locking catch clicks into place.

7. Reattach the power cable and cable retention bracket.

8. Turn on the power switch to the power supply unit.

What to do next

If required, return the battery. Follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Releasing the cable retention bracket

Release the cable retention bracket when removing the power cords from the power supply unit.

Releasing the cable retention bracket

This topic provides instructions for releasing the cable retention bracket when removing the power cords from the power supply unit.

About this task

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

Each cable retention bracket comes attached to the back of the power supply unit by the power cord plug-in.

To release a cable retention bracket, perform these steps:

Procedure

1. Unlock the cable retention bracket that is around the end of the power cord.

2. Pull the lever next to the black plastic loop slightly towards the center of the canister.

3. Continue to pull the lever towards you as you slide the cable retention bracket away from the end of the cable.

Replacing a 3.5 inch drive assembly or blank carrier

Remove and replace a 3.5 inch drive assembly or a blank carrier.

Removing and replacing a drive assembly: Storwize V7000 Gen2

You can replace a faulty 3.5-inch drive assembly on a Storwize V7000 Gen2 expansion controller with a new one received from CRU / FRU stock.

Before you begin

Ensure that the drive is not a spare or a member of an array. The drive status is shown in **Pools** > **Internal Storage** in the management GUI.

Attention:
- Do not replace a drive unless the drive fault LED is on or you are instructed to do so by a fix procedure.
- If the drive is a member of an array, go to the management GUI and follow the fix procedures. The fix procedures mitigate loss of data and loss of access to data and manage the system's use of the drive.
- Do not leave a drive slot empty for extended periods. Do not remove a drive assembly or a blank filler without having a replacement drive or a blank filler with which to replace it.

Procedure

To prepare to replace a drive assembly, complete the following steps.

1. Read the safety information in "Preparing to remove and replace parts" on page 177.
2. Locate the slot that contains the drive assembly that you want to replace.
   a. Refer to "Procedure: Identifying which Storwize V7000 Gen2 enclosure or canister to service" on page 102 to ensure correct identification of the correct system or enclosure.
   b. The drive slots on the front are numbered 1 - 12. For example, the numbering is from left to right and top to bottom:

      ```
      1  2  3  4
      5  6  7  8
      9 10 11 12
      ```

   c. If the drive in the slot is faulty, the lit, amber fault LED on the drive helps identify it.
3. To further help identify the drive assembly, go to the management GUI to **Pools** > **Internal Storage**, select the drive to replace, and click **Actions** > **Identify**. Verify that the correct drive fault LED begins to flash.

   Attention: Never hot-swap a hard disk drive when its green activity LED is flashing. Hot-swap a drive only when its amber fault LED is lit (not flashing) or when the drive activity LED is off.

Remove a drive assembly

4. Press the latch on the right end of the tray handle to release it.
5. Pull out the tray handle to the open position (see Figure 60 on page 205).
6. Grasp the handle and pull the drive partially out of the bay.
7. Wait at least 20 seconds before you remove the drive assembly from the enclosure to enable the drive to spin down. This avoids possible damage to the drive.
8. Make sure that there is proper identification (such as a label) on the hard disk drive.
9. Gently slide it completely out of the enclosure.
10. If the drive failed, record that information on its label.

Replace a drive assembly

11. Touch the static-protective package that contains the drive assembly to any unpainted surface on the outside of the enclosure.

12. Remove the drive assembly from its package.
13. Make sure that its drive-tray handle is in the open (unlocked) position.
14. Align the drive assembly with the guide rails in the bay (see Figure 61).
15. Gently push the drive assembly into the bay until the drive stops.
16. Rotate its handle to the closed (locked) position.
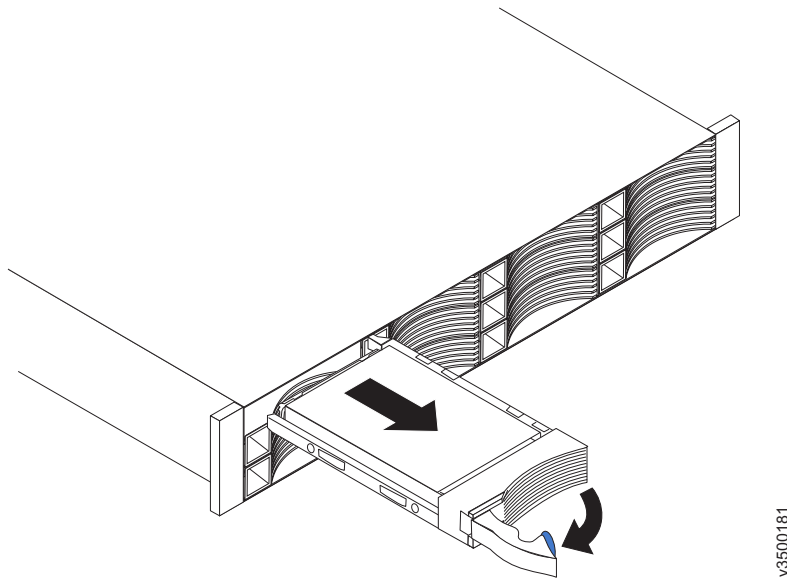
Figure 60. Unlocking and removing a 3.5-inch drive from its slot
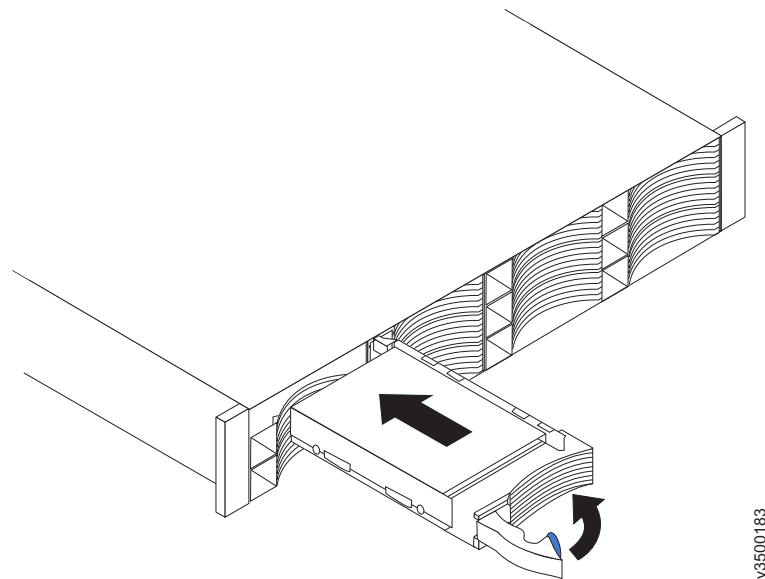
Figure 61. Installing and locking a 3.5-inch drive into its slot

Results

If the replaced drive was a failed drive, the system automatically reconfigures the replacement drive as a spare and the replaced drive is removed from the configuration. The process can take a few minutes.

Replacing a 3.5-inch drive assembly or blank carrier

This topic describes how to replace a 3.5-inch drive assembly or blank carrier.

About this task

Attention: If your drive is configured for use, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures results in loss of data or loss of access to data.

Attention: Do not leave a drive slot empty. Do not remove a drive or drive assembly before you have a replacement available.

The drives can be distinguished from the blank carriers by the color-coded striping on the drive. The drives are marked with an orange striping. The blank carriers are marked with a blue striping.

To replace the drive assembly or blank carrier, perform the following steps:

Procedure

1. Read the safety information to which "Preparing to remove and replace parts" on page 177 refers.
2. Unlock the assembly by squeezing together the tabs on the side.
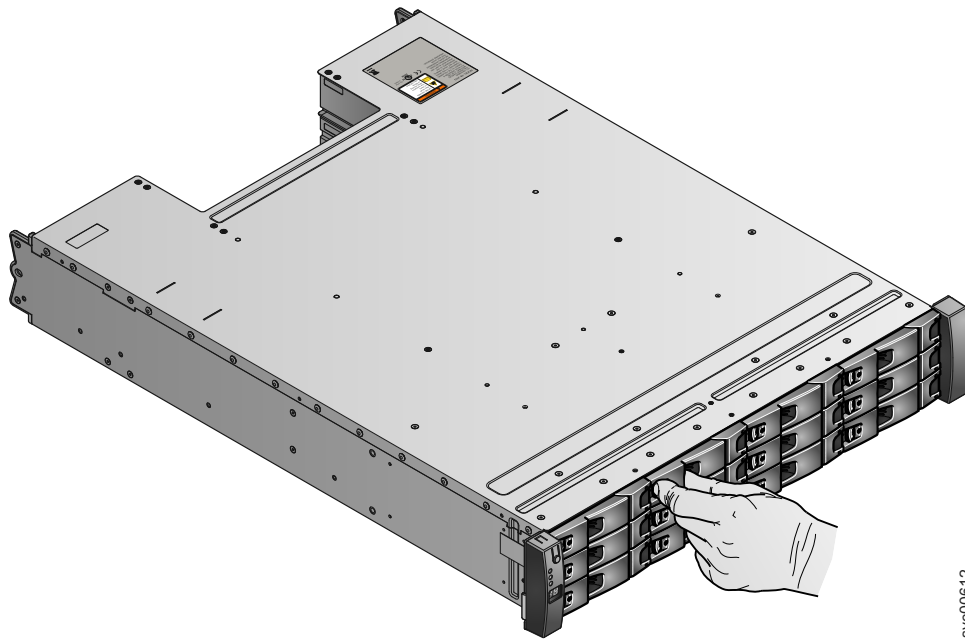
Figure 62. Unlocking the 3.5 inch drive

3. Open the handle to the full extension.
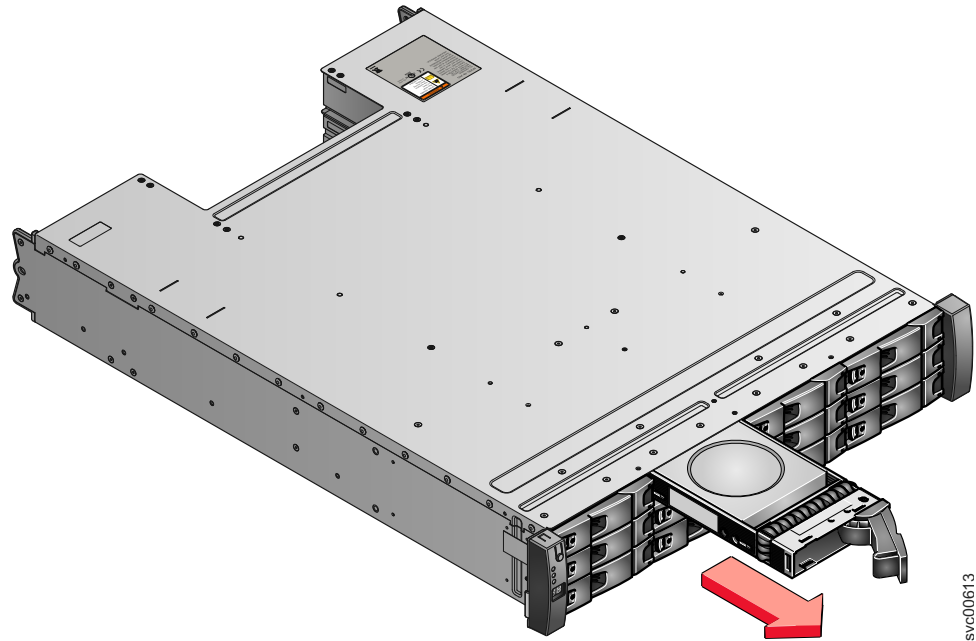
Figure 63. Removing the 3.5 inch drive

4. Pull out the drive.
5. Push the new drive back into the slot until the handle starts to move.
6. Finish inserting the drive by closing the handle until the locking catch clicks into place.

Replacing a 2.5 inch drive assembly or blank carrier

Remove and replace a 2.5 inch drive assembly or a blank carrier.

Removing and replacing a drive assembly: Storwize V7000 Gen2 or Storwize V7000 Gen2+

You can replace a faulty 2.5-inch drive assembly in the expansion controller with a new one received from CRU / FRU stock. Storwize V7000 Gen2 and Storwize V7000 Gen2+ control enclosures use the same 2.5-inch drive assembly.

Before you begin

Ensure that the drive is not a spare or a member of an array. The drive status is shown in **Pools** > **Internal Storage** in the management GUI.

Attention:

- Do not replace a drive unless the drive fault LED is on or you are instructed to do so by a fix procedure.
- If the drive is a member of an array, go to the management GUI and follow the fix procedures. The fix procedures mitigate loss of data and loss of access to data and manage use of the drive by the system.
- Do not leave a drive slot empty for extended periods. Do not remove a drive assembly or a blank filler without having a replacement drive or a blank filler with which to replace it.

Procedure

To prepare to replace a drive assembly, complete the following steps.

1. Read the safety information in "Preparing to remove and replace parts" on page 177.
2. Locate the slot that contains the drive assembly that you want to replace.
   a. Refer to "Procedure: Identifying which Storwize V7000 Gen2 enclosure or canister to service" on page 102 to ensure correct identification of the correct system or enclosure.
   b. The drive slots on the front are numbered 1 - 24, starting from the far left slot of the enclosure.
   c. If the drive in the slot is faulty, the lit, amber fault LED on the drive helps to identify it.
3. To further help identify the drive assembly, go to the management GUI to **Pools** > **Internal Storage**, select the drive to replace, and click **Actions** > **Identify**. Verify that the correct drive fault LED flashes.

   Attention: Never hot-swap a disk drive when its green activity LED is flashing. Hot-swap a drive only when its amber fault LED is lit (not flashing) or when the drive activity LED is off.

Remove a drive assembly

4. Gently slide the orange release latch up to unlock the handle.
5. Pull out the tray handle to the open position (see Figure 64 on page 209).
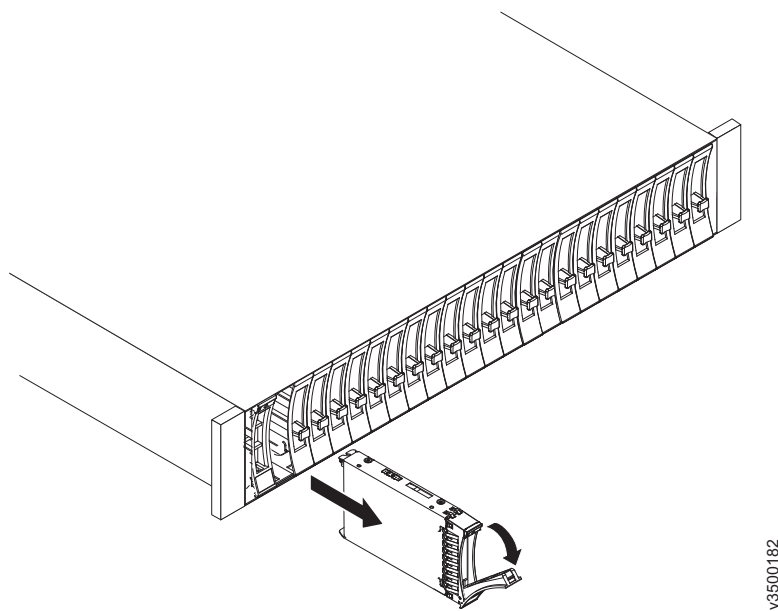
Figure 64. Unlocking and removing a 2.5-inch drive from its slot

6. Grasp the handle and pull the drive partially out of the bay.

7. To avoid possible damage to the drive, wait at least 20 seconds before you remove the drive assembly from the enclosure.

8. Gently slide the drive assembly out of the enclosure.

9. Make sure the drive assembly has proper identification (such as a label). If the drive failed, record that information on the label.

Replace a drive assembly

10. Touch the static-protective package that contains the drive assembly to any unpainted surface on the outside of the enclosure.

11. Remove the drive assembly from its package.

12. Make sure that its drive-tray handle is in the open (unlocked) position.

13. Align the drive assembly with the guide rails in the bay (see Figure 65 on page 210).
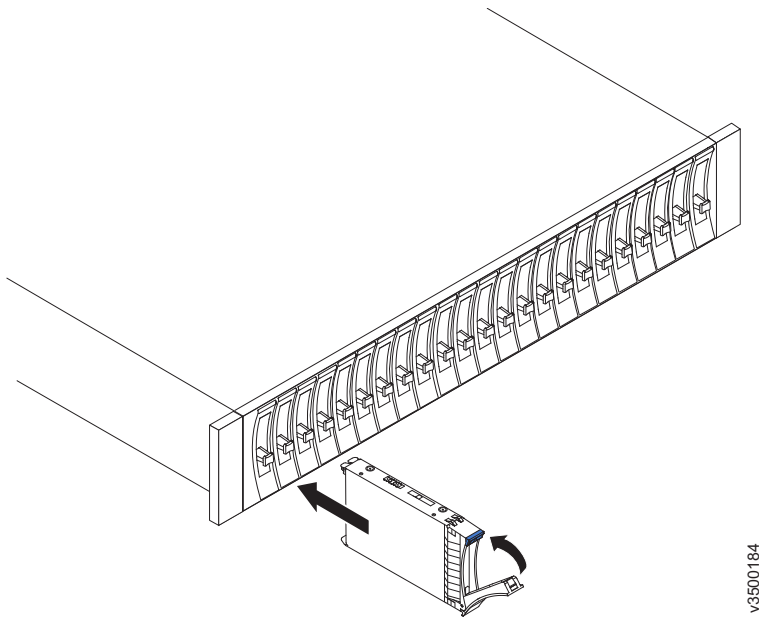
Figure 65. Installing and locking a 2.5-inch drive into its slot

14. Gently push the drive assembly into the bay until the drive stops.
15. Rotate the drive handle to the closed (locked) position.

Results

If the replaced drive was a failed drive, the system automatically reconfigures the replacement drive as a spare and the replaced drive is removed from the configuration. The process can take a few minutes.

Replacing a 2.5-inch drive assembly or blank carrier

This topic describes how to remove a 2.5-inch drive assembly or blank carrier.

About this task

Attention: If your drive is configured for use, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures results in loss of data or loss of access to data.

Attention: Do not leave a drive slot empty. Do not remove a drive or drive assembly before you have a replacement available.

To replace the drive assembly or blank carrier, perform the following steps:

Procedure

1. Read the safety information to which "Preparing to remove and replace parts" on page 177 refers.
2. Unlock the module by squeezing together the tabs at the top.
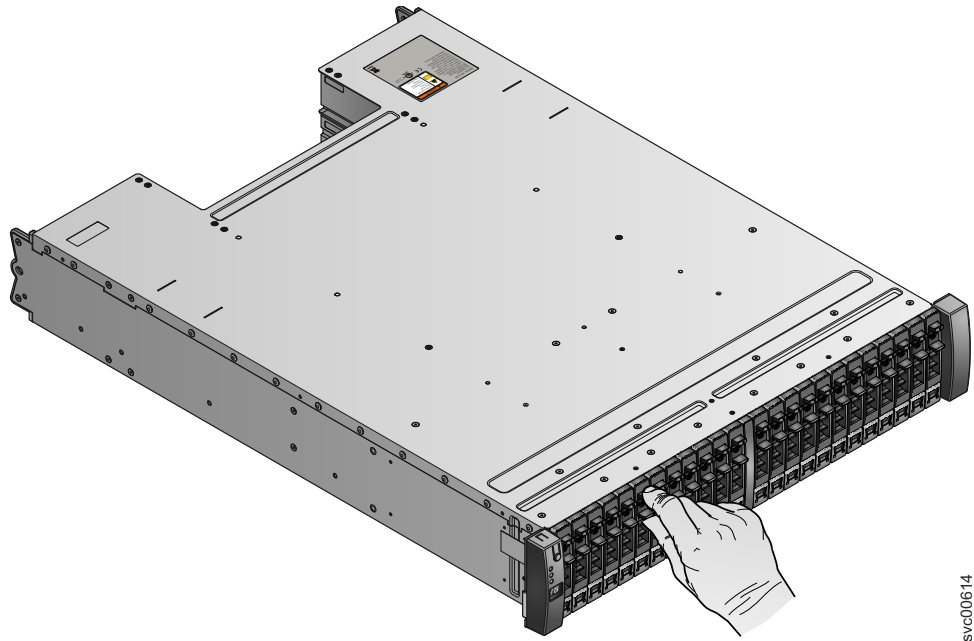
Figure 66. Unlocking the 2.5 inch drive

3. Open the handle to the full extension.
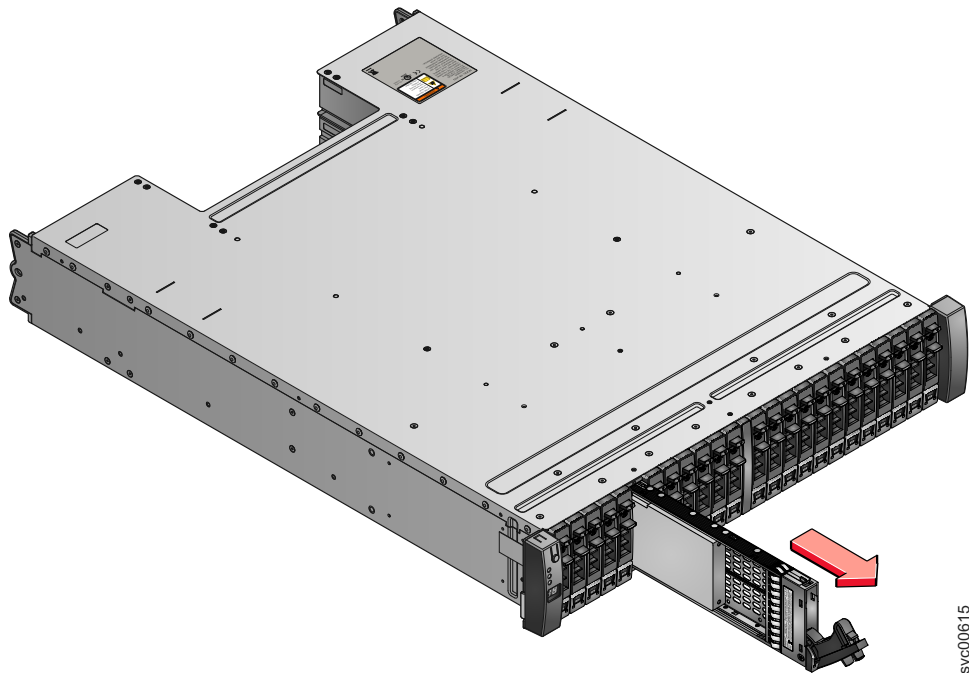
Figure 67. Removing the 2.5 inch drive

4. Pull out the drive.
5. Push the new drive back into the slot until the handle starts to move.
6. Finish inserting the drive by closing the handle until the locking catch clicks into place.

Replacing enclosure end caps

Remove and replace enclosure end caps.

Replacing Storwize V7000 Gen2 enclosure end caps

You can remove and replace enclosure end caps.

About this task

Attention: The left end cap is printed with information that helps identify the enclosure.

- Machine type and model
- Enclosure serial number

The information on the end cap should always match the information that is printed on the rear of the enclosure. It should also match the information that is stored on the enclosure midplane.

Procedure

To remove and replace either the left or right end cap, complete the following steps.

1. If the enclosure is on a table or other flat surface, elevate the enclosure front slightly or carefully extend the front over the table edge.
2. Grasp the end cap by the blue touch point and pull it until the bottom edge of the end cap is clear of the bottom tab on the chassis flange.
3. Lift the end cap off the chassis flange.
4. Fit the slot on the top of the new end cap over the tab on the top of the chassis flange.
5. Rotate the end cap down until it snaps into place. Ensure that the inside surface of the end cap is flush with the chassis.

Replacing enclosure end caps

You can remove and replace enclosure end caps.

About this task

Attention: The left end cap is printed with information that helps identify the enclosure.

- Machine type and model
- Enclosure serial number
- Machine part number

The information on the end cap should always match the information that is printed on the rear of the enclosure, and also match the information that is stored on the enclosure midplane.

Procedure

To remove and replace either the left or right end cap, complete the following steps.

1. If the enclosure is on a table or other flat surface, elevate the enclosure front slightly or carefully extend the front over the table edge.

2. Grasp the end cap by the blue touch point and pull it until the bottom edge of the end cap is clear of the bottom tab on the chassis flange.

3. Lift the end cap off the chassis flange.

4. Fit the slot on the top of the new end cap over the tab on the top of the chassis flange.

5. Rotate the end cap down until it snaps into place. Ensure that the inside surface of the end cap is flush with the chassis.

Replacing a SAS cable to an expansion enclosure

Remove and replace a SAS cable to an expansion enclosure.

Replacing a Storwize V7000 Gen2 expansion enclosure attachment SAS cable

To replace a faulty Storwize V7000 Gen2 expansion enclosure attachment SAS cable with a new one received from CRU / FRU stock, use this procedure.

About this task

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

Attention:

If you need to replace more than one cable, record which two ports, canisters, and enclosures each cable connects, so you can match the connections with the replacement cables. The system cannot operate if the expansion enclosure attachment SAS cabling is incorrect.

Expansion enclosure attachment SAS cables are connected only between SAS port 3 or 4 of a node canister and SAS port 1 of an expansion canister, or between SAS ports 1 and 2 of different expansion canisters.

More information about correct expansion enclosure attachment SAS cabling can be found in the troubleshooting description of a problem with Storwize V7000 Gen2 SAS cabling.

Procedure

To replace a SAS cable, complete the following steps.

1. Locate the connector at one end of the SAS cable that is to be removed.

2. Grasp the connector by its blue tag. Pull the tag. The connector is released and slides out of the port.

3. Repeat steps 2 and 3 on the other end of the SAS cable.

4. To connect the replacement expansion enclosure attachment SAS cable, connect each end to the vacated ports.

Attention: When you insert a SAS connector into a SAS port, ensure that the orientation of the connector matches the orientation of the port before you push the connector into the port.

- The cable connector and socket are keyed and it is important that you have proper alignment of the keys when the cable is inserted.
- Before you insert the connector into the port, ensure that the connector is rotated such that the blue tag is the lowest part.
- Figure 68 shows the correct orientation. The blue tab is always below the port for expansion enclosure attachment SAS cables.
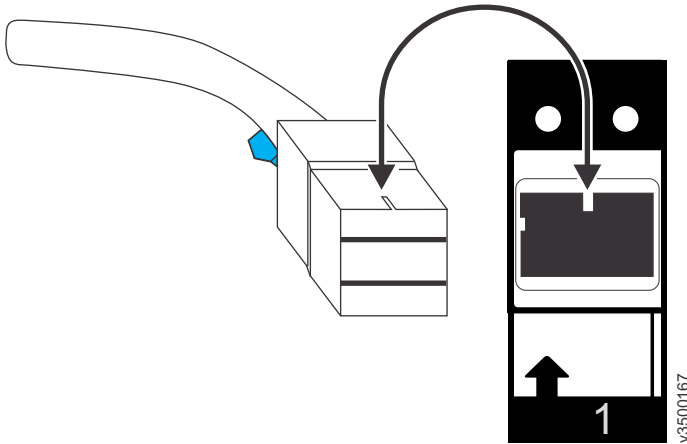
Figure 68. Proper orientation for SAS cable connector

- A click is heard or felt when the cable is successfully inserted and you should not be able to disconnect the cable without pulling on the blue tag.
- When both ends of a SAS cable are correctly connected, the green link LED next to the connected SAS ports are lit.

  For more information, see the troubleshooting procedure for finding the status of SAS connections.

Replacing a SAS cable

This topic describes how to replace a SAS cable.

About this task

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

To replace a SAS cable, perform the following steps:

Procedure

1. Record which SAS cable is plugged into the specific port of the expansion canister. The cable must be inserted back into the same port after the replacement is complete; otherwise, the system cannot function properly.

   Note: If you are replacing a single cable, this step is not necessary.
2. Pull the tab with the arrow away from the connector.

Figure 69. SAS cable

3. Plug the replacement cable into the specific port.
4. Ensure that the SAS cable is fully inserted. A click is heard when the cable is successfully inserted.

Replacing a control enclosure chassis

Remove and replace a control enclosure chassis. This procedure only applies to Storwize V7000 Gen1 control enclosure models.

About this task

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 98. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 99. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Replacing a Storwize V7000 Gen2 control enclosure

You can replace a control enclosure.

Before you begin

Note: Ensure that you know the type of enclosure that you are replacing. The procedures for replacing a control enclosure are different from those procedures for replacing an expansion enclosure chassis. For information about replacing an expansion enclosure, see "Replacing an expansion enclosure chassis" on page 226.

DANGER

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

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

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- 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.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

Attention: Perform this procedure only if instructed to do so by a service action or the IBM support center. If you have a single control enclosure, this procedure requires that you shut down your system to replace the control enclosure. If you have more than one control enclosure, you can keep part of the system running. However, you will lose access to the volumes that are on the affected I/O group and any volumes that are in other I/O groups that depend on the drives that are in the affected I/O group. If the system is still doing I/O requests in all the I/O groups, schedule the replacement during a maintenance period or other time when the I/O can be stopped.

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

Ensure that you are aware of the procedures for handling static-sensitive devices before you remove the enclosure.

Procedure

To replace a control enclosure, complete the following steps:

1. If you are able to access either of the node canisters with the service assistant, record the machine type and model of the enclosure, the serial number of the enclosure, and the two worldwide node names (WWNNs) for the enclosure.

   • From the service assistant home page, open the location data for the node. Record the machine type and model (MTM), the serial number, WWNN 1 and WWNN 2 from the enclosure column.

   • If you are replacing the enclosure because neither node canister can start, retrieve this information after you finished the replacement.

      a. Start the service assistant on one of the canisters.

      b. Go to the node location data on the home page.

      c. Record the machine type and model, the serial number, WWNN 1 and WWNN 2 from the node copy column.

         The machine type and model and the serial number are also shown on the labels at the front and back of the enclosure.

2. If the enclosure is still active, shut down the host I/O and the Metro Mirror and Global Mirror activity to all the volumes that depend on the affected enclosure.

   This statement applies to all volumes in the I/O group that are managed by this enclosure plus any volumes in other I/O groups that depend on the drives in the affected I/O group.

3. If your system contains a single I/O group and if the clustered system is still online, shut down the system by using the management GUI.

   a. From the management GUI, go to **Monitoring** > **System Details**.

   b. Select the target enclosure from the left panel and then select **Shut Down System** from the **Actions** menu.

   c. Wait for the shutdown to complete.

4. If your system contains more than one I/O group and if this I/O group is still online, shut down the I/O group by using the CLI.

   a. Identify the two nodes in the I/O group.

   b. To shut down each node, issue the following CLI command once for each of the two node canisters:

      `stopsystem -force -node <node ID>`

    c. Wait for the shutdown to complete.

5. Verify that it is safe to remove the power from the enclosure.

   For each of the canisters, verify the status of the system status LED. If the LED is lit on either of the canisters, do not continue because the system is still online. Determine why the node canisters did not shut down in step 3 on page 218 or step 4 on page 218.

   Note: If you continue while the system is still active, you risk losing the clustered system configuration and volume cache data that is stored in the canister.

6. Turn off the power to the enclosure by using the switches.

7. Record which data cables are plugged into the specific ports. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.

8. Disconnect the cable retention brackets and the power cords from the power supply units.

9. Disconnect the data cables for each canister.

10. Remove the power supply units from the enclosure.

11. Remove the canisters from the enclosure. Record the location of each canister. They must be inserted back into the same location in the new enclosure.

12. Remove the fan modules from the enclosure. Record the location of each fan module. They must be inserted back into the same location in the new enclosure.

13. Remove all the drives and blank drive assemblies from the enclosure. Record the location for each drive. They must be inserted back into the same location in the new enclosure.

14. Remove both enclosure end caps from the enclosure. Keep the left end cap because it is used again.

15. Remove the clamping screws that attached the enclosure to the rack cabinet.

16. Remove the enclosure chassis from the front of the rack cabinet and take the chassis to a work area.

17. Install the new enclosure chassis in the rack cabinet.

18. Remove the end caps from the new enclosure and install the clamping screws that attach the enclosure to the rack cabinet.

19. Replace the end caps. Use the new right end cap and use the left end cap that you removed in step 14.

    Using the left end cap that you removed preserves the model and serial number identification.

20. Reinstall the drives in the new enclosure. The drives must be inserted back into the same location from which they were removed on the old enclosure.

21. Reinstall the canisters in the enclosure. The canisters must be inserted back into the same location from which they were removed on the old enclosure.

22. Reinstall the fan modules in the enclosure. The fan modules must be inserted back into the same location from which they were removed on the old enclosure.

23. Install the power supply units.

24. Reattach the data cables to each canister by using the information that you recorded previously.

Note: The cables must be inserted back into the same ports from which they were removed on the old enclosure; otherwise, the system cannot function properly.

25. Attach the power cords and the cable retention brackets to the power supply units.

26. Write the old enclosure machine type and model (MTM) and serial number on the repair identification (RID) tag that is supplied. Attach the tag to the left flange at the back of the enclosure.

27. Turn on the power to the enclosure by using the switches.

    The node canisters boot up. The fault LEDs are on because the new enclosure was not set with the identity of the old enclosure. The node canisters report that they are in the wrong location.

    a. Connect to the service assistant on one of the node canisters to configure the machine type and model, serial number, and WWNNs that are stored in the enclosure. If you replaced a node canister, connect to the canister that was not replaced.

       You can connect by using the previous service address. However, it is not always possible to maintain this address. If you cannot connect through the original service address, attempt to connect by using the default service address. If you still cannot access the system, see the troubleshooting description of a problem when connecting to the service assistant. Problem: Cannot connect to the service assistant.

    b. Use the **Configure enclosure** panel.

    c. Select the options to **Update WWNN 1**, **Update WWNN 2**, **Update the machine type and model**, and **Update the serial number**. Do not update the system ID. Use the node copy data for each of the values. Check that these values match the values that you recorded in step 1 on page 218.

       If you were not able to record the values, use the node copy values only if none of them have all zeros as their value. If any of the node copy values are all zeros, connect the service assistant to the other node canister and configure the enclosure there. If you still do not have a full set of values, contact IBM support.

       After you modify the configuration, the node attempts to restart.

       Note: There are situations where the canisters restart and report critical node error 508. If the node canisters fail to become active after they restart when the enclosure is updated, check their status by using the service assistant. If both node canisters show critical node error 508, use the service assistant to restart the nodes. For any other node error, see "Procedure: Fixing node errors" on page 121.

       To restart a node from the service assistant, do the following steps:

       1) Log on to the service assistant.

       2) From the home page, select the node that you want to restart from the **Changed Node List**.

       3) Select **Actions** > **Restart**.

    d. The system starts and can handle I/O requests from the host systems.

       Note: The configuration changes that are described in the following steps must be done to ensure that the system is operating correctly. If you do not do these steps, the system is unable to report certain errors.

28. Start the management GUI and select **Monitoring** > **System Details**. You see an extra enclosure in the system list because the system detected the replacement control enclosure. The original control enclosure is still listed in its configuration. The original enclosure is listed with its original enclosure ID. It is offline and managed. The new enclosure has a new enclosure ID. It is online and unmanaged.

29. Select the original enclosure in the tree view.

    Verify that it is offline and managed and that the serial number is correct.

30. From the **Actions** menu, select **Remove enclosure** and confirm the action. The physical hardware was already removed. You can ignore the messages about removing the hardware. Verify that the original enclosure is no longer listed in the tree view.

31. Add the new enclosure to the system.

    a. Select the enclosure from the tree view.

    b. From the **Actions** menu, select **Add Control and Expansion Enclosures**.

    c. Because you already added the hardware, select **Next** on the first panel that asks you to install the hardware. The next panel shows the unmanaged new enclosure.

    d. Follow the steps in the wizard. The wizard changes the control enclosure to `Managed`.

    e. Select the enclosure and add it to the system.

32. Select the new enclosure in the tree view and verify that it is now online and managed.

33. Change the enclosure ID of the replaced enclosure to that of the original enclosure. From the **Enclosure ID** field, select the ID value of the original enclosure.

34. Check the status of all volumes and physical storage to ensure that everything is online.

35. Restart the host application and any FlashCopy activities, Global Mirror activities, or Metro Mirror activities that were stopped.

Replacing a Storwize V7000 Gen1 control enclosure chassis

You can replace a control enclosure chassis.

Before you begin

Note: Ensure that you know the type of enclosure chassis that you are replacing. The procedures for replacing a control enclosure chassis are different from those procedures for replacing an expansion enclosure chassis. To replace an expansion enclosure chassis, see "Replacing an expansion enclosure chassis" on page 226.

DANGER

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

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

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- 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.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

Attention: Perform this procedure only if instructed to do so by a service action or the IBM support center. If you have a single control enclosure, this procedure requires that you shut down your system to replace the control enclosure. If you have more than one control enclosure, you can keep part of the system running. However, you will lose access to the volumes that are on the affected I/O group and any volumes that are in other I/O groups that depend on the drives that are in the affected I/O group. If the system is still performing I/O requests in all the I/O groups, schedule the replacement during a maintenance period or other time when the I/O can be stopped.

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

Ensure that you are aware of the procedures for handling static-sensitive devices before you remove the enclosure.

Procedure

To replace a control enclosure chassis, complete the following steps.

1. If you are able to access either of the node canisters with the service assistant, record the machine type and model of the enclosure, the serial number of the enclosure, and the two WWNNs for the enclosure.

   a. From the service assistant home page, open the location data for the node. Record the machine type and model (MTM), the serial number, WWNN 1, and WWNN 2 from the enclosure column.

      Note: If you are replacing the enclosure because neither node canister can start, retrieve this information after you complete the replacement.

   b. Start the service assistant on one of the canisters.

   c. Go to the node location data on the home page.

   d. Record the machine type and model, the serial number, WWNN 1, and WWNN 2 from the node copy column. The machine type and model and the serial number are also shown on the labels at the front and back of the enclosure.

2. If the enclosure is still active, shut down the host I/O and the Metro Mirror and Global Mirror activity to all the volumes that depend on the affected enclosure. This statement applies to all volumes in the I/O group that are managed by this enclosure plus any volumes in other I/O groups that depend on the drives in the affected I/O group.

3. If your system contains a single I/O group and if the clustered system is still online, shut down the system by using the management GUI.

   a. From the management GUI, go to **Monitoring** > **System**.

   b. Select **Power off** from the **Actions** menu.

   c. Wait for the shutdown to complete.

4. If your system contains more than one I/O group and if this I/O group is still online, shut down the I/O group by using the CLI.

   a. Identify the two nodes in the I/O group.

   b. To shut down each node, issue the following CLI command once for each of the two node canisters:

      ```
      stopsystem -force -node <node ID>
      ```

   c. Wait for the shutdown to complete.

5. Verify that it is safe to remove the power from the enclosure. For each of the canisters, verify the status of the system status LED. If the LED is lit on either of the canisters, do not continue because the system is still online. Determine why the node canisters did not shut down in step 3 on page 223 or step 4 on page 223.

   Note: If you continue while the system is still active, you risk losing the clustered system configuration and volume cache data that is stored in the canister.

6. Turn off the power to the enclosure by using the switches on the power supply units.

7. Record which data cables are plugged into the specific ports. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.

8. Disconnect the cable retention brackets and the power cords from the power supply units.

9. Disconnect the data cables for each canister.

10. Remove the power supply units from the enclosure.

11. Remove the canisters from the enclosure. Record the location of each canister. They must be inserted back into the same location in the new enclosure.

12. Remove all the drives and blank drive assemblies from the enclosure. Record the location for each drive. They must be inserted back into the same location in the new enclosure.

13. Remove both enclosure end caps from the enclosure.

    Important: Keep the left end cap because it is used again.

14. Remove the clamping screws that attached the enclosure to the rack cabinet.

15. Remove the enclosure chassis from the front of the rack cabinet and take the chassis to a work area.

16. Install the new enclosure chassis in the rack cabinet.

17. Remove the end caps from the new enclosure and install the clamping screws that attach the enclosure to the rack cabinet.

18. Replace the end caps. Use the new right end cap and use the left end cap that you removed in step 13. Using the left end cap that you removed preserves the model and serial number identification.

19. Reinstall the drives in the new enclosure. The drives must be inserted back into the same location from which they were removed on the old enclosure.

20. Reinstall the canisters in the enclosure. The canisters must be inserted back into the same location from which they were removed on the old enclosure.

21. Install the power supply units.

22. Reattach the data cables to each canister by using the information that you recorded previously.

    Note: The cables must be inserted back into the same ports from which they were removed on the old enclosure; otherwise, the system cannot function properly.

23. Attach the power cords and the cable retention brackets to the power supply units.

24. Write the old enclosure machine type and model (MTM) and serial number on the repair identification (RID) tag that is supplied. Attach the tag to the left flange at the back of the enclosure.

25. Turn on the power to the enclosure by using the switches on the power supply units. The node canisters boot up. The fault LEDs are on because the new enclosure is not set with the identity of the old enclosure. The node canisters log node error 504, reporting that they are in the wrong location. In the system event log, these errors appear with an error code of 1192.

26. Connect to the service assistant on one of the node canisters to configure the machine type and model, serial number, and WWNNs that are stored in the enclosure. If you replace a node canister, connect to the canister that has not been replaced. You can connect by using the previous service address. However, it is not always possible to maintain this address. If you cannot connect through the original service address, attempt to connect by using the default service address. If you still cannot access the system, see "Problem: Cannot connect to the service assistant" on page 94.

27. Use the **Configure enclosure** panel.

28. Select the options to **Update WWNN 1**, **Update WWNN 2**, **Update the machine type and model**, and **Update the serial number**. Do not update the system ID. Use the node copy data for each of the values. Check that these values match the values that you recorded in step 1 on page 223. If you were not able to record the values, use the node copy values only if none of them have all zeros as their value. If any of the node copy values are all zeros, connect the service assistant to the other node canister and configure the enclosure there. If you still do not have a full set of values, contact IBM support.

    Important: Step 29 writes the enclosure identity into the replacement midplane. The replacement midplane cannot be used as a replacement part for a different enclosure after step 29 is completed.

29. Click the **Modify** button. The node writes the data and attempts to restart.

    Note: There are situations where the canisters restart and report critical node error 508. If the node canisters fail to become active after they restart when the enclosure is updated, check their status by using the service assistant. If both node canisters show critical node error 508, use the service assistant to restart the nodes. For any other node error, see "Procedure: Fixing node errors" on page 121. To restart a node from the service assistant, perform the following steps:

    a. Log on to the service assistant.

    b. From the home page, select the node that you want to restart from the **Changed Node List**.

    c. Select **Actions** > **Restart**.

    The system starts and can handle I/O requests from the host systems.

    Note: The configuration changes that are described in the following steps must be performed to ensure that the system is operating correctly. If you do not perform these steps, the system is unable to report certain errors.

30. Start the management GUI and select **Monitoring** > **System**. You see an additional enclosure in the system list because the system has detected the replacement control enclosure. The original control enclosure is still listed in its configuration. The original enclosure is listed with its original enclosure ID. It is offline and managed. The new enclosure has a new enclosure ID. It is online and unmanaged.

31. Select the original enclosure. Verify that it is offline and managed and that the serial number is correct.

32. Right-click the enclosure and select **Remove**. The physical hardware has already been removed. You can ignore the messages about removing the hardware. Verify that the original enclosure is no longer listed in the tree view.

33. Add the new enclosure to the system.

    a. From the **Actions** menu, select **Add Enclosures**.

    b. Because you already added the hardware, select **Next** on the first panel that asks you to install the hardware. The next panel shows the unmanaged new enclosure.

    c. Follow the steps in the wizard. The wizard changes the control enclosure to `Managed`.

    d. Select the enclosure and add it to the system.

34. Select the new enclosure and verify that it is now online and managed.

35. Change the enclosure ID of the replaced enclosure to that of the original enclosure. Right-click the control enclosure and select **Modify ID** and change the ID to that of the original enclosure.

36. Check the status of all volumes and physical storage to ensure that everything is online.

37. Restart the host application and any FlashCopy activities, Global Mirror activities, or Metro Mirror activities that were stopped.

Results

Replacing an expansion enclosure chassis

Remove and replace an expansion enclosure chassis. This procedure only applies to Storwize V7000 Gen1 enclosure models.

About this task

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 100. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 101. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives

Table 101. Storwize V7000 Gen2 model numbers (continued)

Machine type/model	Description
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Replacing a Storwize V7000 Gen2 expansion enclosure

You can replace an expansion enclosure.

Before you begin

Note: Ensure that you know the type of enclosure chassis that you are replacing. The procedures for replacing an expansion enclosure chassis are different from those procedures for replacing a control enclosure chassis. To replace a control enclosure chassis, see "Replacing a control enclosure chassis" on page 215.

DANGER

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

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

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- 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.

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

Attention: If your system is powered on and performing I/O operations, go to the management GUI and follow the fix procedures. Performing the replacement actions without the assistance of the fix procedures can result in loss of data or access to data.

Even though many parts are hot-swappable, these procedures are intended to be used only when your system is not up and running and performing I/O operations. Unless your system is offline, go to the management GUI and follow the fix procedures.

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

Ensure that you are aware of the procedures for handling static-sensitive devices before you remove the enclosure.

About this task

Note: If your system is online, replacing an expansion enclosure can cause one or more of your volumes to go offline or your quorum disks to be inaccessible. Before you proceed with these procedures, verify which volumes might go offline. From the management GUI, go to **Home** > **Manage Devices**. Select the enclosure that you want to replace. Then, select **Show Dependent Volumes** in the **Actions** menu.

Procedure

To replace an expansion enclosure chassis, perform the following steps:
1. Shut down the I/O activity to the enclosure, which includes host access, FlashCopy, Metro Mirror, and Global Mirror access.
2. Turn off the power to the enclosure by disconnecting the power cable.
3. Record which data cable is plugged into each specific port. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
4. Disconnect the data cables for each canister.
5. Remove the power supply units from the enclosure.
6. Remove the canisters from the enclosure.
7. Remove all the drives and blank drive assemblies from the enclosure. Record the location for each drive. They must be inserted back into the same location in the new enclosure.
8. Remove both enclosure end caps from the enclosure. Keep the left end cap because it is used again.
9. Remove the clamping screws that attached the enclosure to the rack cabinet.
10. Remove the enclosure chassis from the front of the rack cabinet and take the chassis to a work area.
11. Install the new enclosure chassis in the rack cabinet.
12. Remove the end caps from the new enclosure and install the clamping screws that attach the enclosure to the rack cabinet.
13. Replace the end caps. Use the new right end cap and use the left end cap that you removed in step 8.

    Using the left end cap that you removed preserves the model and serial number identification.

14. Reinstall drives in the new enclosure. You must insert the drives back into the same location from which they were removed on the old enclosure.
15. Reinstall the canisters (and drives) in the enclosure.
16. Install the power supply units.
17. Use the information that you recorded previously to reattach the data cables to each canister.

    Note: The cables must be inserted back into the same ports from which they were removed on the old enclosure; otherwise, the system cannot function properly.
18. Attach the power cords and the cable retention brackets to the power supply units.
19. Write the old enclosure machine type and model (MTM) and serial number on the repair identification (RID) tag that is supplied. Attach the tag to the left flange at the back of the enclosure.

Results

The system records an error that indicates that an enclosure FRU replacement was detected. Go to the management GUI to use the fix procedure to change the machine type and model and serial number in the expansion enclosure.

Replacing a Storwize V7000 Gen1 expansion enclosure chassis

You can replace an expansion enclosure chassis.

Before you begin

Note: Ensure that you know the type of enclosure chassis that you are replacing. The procedures for replacing an expansion enclosure chassis are different from those procedures for replacing a control enclosure chassis. To replace a control enclosure chassis, see "Replacing a control enclosure chassis" on page 215.

DANGER

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

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

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- 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.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

Attention: If your system is powered on and performing I/O operations, go to the management GUI and follow the fix procedures. Performing the replacement actions without the assistance of the fix procedures can result in loss of data or access to data.

Even though many of the parts are hot-swappable, these procedures are intended to be used only when your system is not up and running and performing I/O operations. Unless your system is offline, go to the management GUI and follow the fix procedures.

Be careful when you are replacing the hardware components that are located in the back of the system. Do not inadvertently disturb or remove any cables that you are not instructed to remove.

Ensure that you are aware of the procedures for handling static-sensitive devices before you remove the enclosure.

About this task

Note: If your system is online, replacing an expansion enclosure can cause one or more of your volumes to go offline or your quorum disks to be inaccessible. Before you proceed with these procedures, verify which volumes might go offline. From the management GUI, go to **Home** > **Manage Devices**. Select the enclosure that you want to replace. Then select **Show Dependent Volumes** in the **Actions** menu.

To replace an expansion enclosure chassis, perform the following steps:

Procedure

1. Shut down the I/O activity to the enclosure, which includes host access, FlashCopy, Metro Mirror, and Global Mirror access.
2. Turn off the power to the enclosure by using the switches on the power supply units.
3. Record which data cables are plugged into the specific ports. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
4. Disconnect the cable retention brackets and the power cords from the power supply units.
5. Disconnect the data cables for each canister.
6. Remove the power supply units from the enclosure.
7. Remove the canisters from the enclosure.
8. Remove all the drives and blank drive assemblies from the enclosure. Record the location for each drive. They must be inserted back into the same location in the new enclosure.
9. Remove both enclosure end caps from the enclosure. Keep the left end cap because it is used again.
10. Remove the clamping screws that attached the enclosure to the rack cabinet.
11. Remove the enclosure chassis from the front of the rack cabinet and take the chassis to a work area.
12. Install the new enclosure chassis in the rack cabinet.
13. Remove the end caps from the new enclosure and install the clamping screws that attach the enclosure to the rack cabinet.

14. Replace the end caps. Use the new right end cap and use the left end cap that you removed in step 9 on page 232. Using the left end cap that you removed preserves the model and serial number identification.

15. Reinstall the drives in the new enclosure. The drives must be inserted back into the same location from which they were removed on the old enclosure.

16. Reinstall the canisters in the enclosure.

17. Install the power supply units.

18. Reattach the data cables to each canister by using the information that you recorded previously.

    Note: The cables must be inserted back into the same ports from which they were removed on the old enclosure; otherwise, the system cannot function properly.

19. Attach the power cords and the cable retention brackets to the power supply units.

20. Write the old enclosure machine type and model (MTM) and serial number on the repair identification (RID) tag that is supplied. Attach the tag to the left flange at the back of the enclosure.

21. Turn on the power to the enclosure by using the switches on the power supply units. The system records an error that indicates that an enclosure FRU replacement was detected.

    Important: Step 22 writes the enclosure identity into the replacement midplane. The replacement midplane cannot be used as a replacement part for a different enclosure after step 22 is completed.

22. Go to the management GUI to use the fix procedure to change the machine type and model and serial number in the expansion enclosure.

Replacing a Storwize V7000 Gen2 enclosure midplane

A trained service provider must replace the midplane assembly of a Storwize V7000 Gen2 enclosure.

About this task

DANGER

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

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

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- 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.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

Attention:

- The enclosure midplane must be replaced only by a trained service provider. Perform this procedure only if instructed to do so by a service action or the IBM support center.
- Be careful when you are replacing the hardware components that are in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.
- Ensure that you are aware of the procedures for handling static-sensitive devices before you remove the enclosure.

Replacing a Storwize V7000 Gen2 control enclosure midplane assembly

A trained service provider can use this procedure to replace a faulty Storwize V7000 Gen2 control enclosure midplane with a new one received from CRU / FRU stock. Ensure that your control enclosure midplane assembly is replaced only by a trained service provider.

Before you begin

Three persons are required at step 14 on page 238.

About this task

Follow all safety precautions when you complete this procedure.

DANGER

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

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

- **If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.**
- **Do not open or service any power supply assembly.**
- **Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.**
- **The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.**
- **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.**

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

Attention:

The control enclosure must be replaced only by a trained service provider. Complete this procedure only if instructed to do so by a service action or the IBM support center.

If you have a single control enclosure, this procedure requires that you shut down your system to replace the control enclosure midplane assembly. If you have more than one control enclosure, you can keep part of the system running. However, you lose access to the volumes that are on the affected I/O group and any volumes that are in other I/O groups that depend on the drives that are in the affected I/O group. If the system is still doing I/O requests in all the I/O groups, schedule the replacement during a maintenance period or other time when the I/O can be stopped.

When you replace hardware components in the back of the enclosure, ensure that you do not inadvertently disturb or remove cables that you are not instructed to remove.

Ensure that you are aware of procedures for handling static-sensitive devices before you remove the enclosure.

Procedure

To replace the control enclosure midplane, complete the following steps:

1. Log in to the service assistant on one of the node canisters in the control enclosure.
2. Navigate to the **Enclosure Information** panel.

    Important: Do NOT select the **Reset the system ID** check box.

    Record the following information for use in subsequent steps:
    - WWNN 1
    - WWNN 2
    - Machine type and model
    - Serial number
3. Read the safety information in "Preparing to remove and replace parts" on page 177.
4. If the control enclosure is still active, stop host I/O and Metro Mirror and Global Mirror activity on all the volumes that depend on the enclosure. This step applies to all I/O group volumes that are managed by this enclosure plus any volumes in other I/O groups that depend on the drives in the affected I/O group.
5. Complete "Procedure: Powering off a Storwize V7000 Gen2 control enclosure" on page 134 for the control enclosure that requires the midplane assembly replacement.
6. Disconnect both power cables from the rear of the enclosure.
7. Write down which port connects to which cable before you disconnect all cables from the rear of the enclosure.
8. Carefully remove each drive and label it with the drive slot from which it was removed.

    You can use the drive-slot information to insert the drives into the correct drive slots at the end of this procedure.

9. Remove the two power supplies from the enclosure. Refer to "Replacing a Storwize V7000 Gen2 power supply unit for a control enclosure" on page 188 for guidance.

10. Remove the node canisters from the enclosure. Label them to indicate what canister came from each canister slot.

11. Remove the fan modules from the enclosure, as described in "Replacing a Storwize V7000 Gen2 fan module" on page 180.

12. Remove the end caps from the enclosure, as described in "Replacing Storwize V7000 Gen2 enclosure end caps" on page 212.

13. Remove the two M5 screws from the front of the enclosure to free the enclosure from the rack.

14. Slide the enclosure from the rack, and then place the enclosure on a work surface, so that the underside of the enclosure faces upward, and the enclosure front is facing toward you.

    CAUTION:
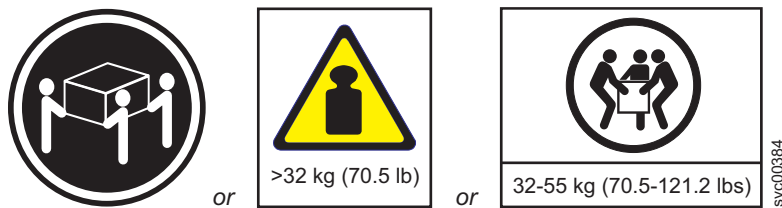
    or   >32 kg (70.5 lb)   or   32-55 kg (70.5-121.2 lbs)   svc00384

    The weight of this part or unit is between 32 and 55 kg (70.5 and 121.2 lb). It takes three persons to safely lift this part or unit. (C010)

15. Remove the four screws from the bottom of the enclosure. Three screws are near the front and one is near the middle. Label these screws to indicate the location from which they are removed and place them aside. Figure 70 illustrates the location of the screws on the bottom of the enclosure.

Figure 70. Bottom enclosure screws

> **Note:** A PH1 screw driver is used for the screws in this step and the following steps. A pair of pliers is needed for the screw-pins in the following steps.

16. Turn the enclosure over again so that the top of the enclosure is facing upward and the front is facing towards you.

17. Remove the three screws and one screw-pin on the right side that secure the midplane assembly to the enclosure. Label each screw to indicate the removal location and place the screws aside. Figure 71 illustrates the location of the screws and screw-pin on the right-side of the enclosure.

Figure 71. Right-side enclosure screws

18. Remove the three screws and one screw-pin on the left side that secure the midplane assembly to the enclosure. Label each screw to indicate the removal location and place the screws aside. Figure 72 illustrates the location of the screws and screw-pin on the left-side of the enclosure.

Figure 72. Left-side enclosure screws

19. Remove the midplane assembly from the chassis by rotating up the midplane assembly to about 45°, then withdraw the midplane assembly from the front of the enclosure. Figure 73 on page 240 shows the midplane assembly at a 45-degree angle.

Figure 73. Angled midplane assembly

20. Unpack the replacement midplane assembly. Grasp the midplane assembly with two hands to hold the assembly at a 45° angle.

21. Insert the tabs on the midplane assembly into the tab holes in the enclosure and rotate down the front of the assembly.

22. Secure the midplane assembly to the enclosure chassis on both the right and left sides by using six screws and two screw-pins that you removed in steps 16 on page 239 and 17 on page 239.

23. Turn over the bottom of the enclosure to face upward, then insert the four screws on the bottom of the enclosure that were removed in step 15 on page 238.

24. Reinstall the enclosure in the rack cabinet, securing it with two screws that were removed at step 13 on page 238.

25. Reinstall the end caps at the front of the enclosure, as described in "Replacing Storwize V7000 Gen2 enclosure end caps" on page 212.

26. Reinstall the hard disk drives at the front of the enclosure, making sure that each drive is inserted into the same slot from which it was removed.

27. Replace the fan modules, as described in "Replacing a Storwize V7000 Gen2 fan module" on page 180.

28. Reinstall the canisters into the same canister slots from which you removed them.

29. Reinstall the two power supplies.

30. Reconnect the data cables at the rear of the enclosure into the same connectors from which you removed them.

31. Reconnect power to the control enclosure. The node canisters restart. The fault LEDs are on because the new enclosure has not been set with the identity of the old enclosure. The node canisters log node error 504, reporting that they are in the wrong location. In the system event log, the error code is 1192.

32. Connect to the service assistant on one of the node canisters to configure the machine type and model, serial number, and WWNNs that are stored in the enclosure. If you have replaced a node canister, connect to the canister that has not been replaced. The service assistant retains a copy of the same information that was on the faulty enclosure midplane assembly. You can connect by using the previous service address. However, it is not always possible to maintain this address. If you cannot connect through the original service address, attempt to connect by using the default service address. If you still cannot access the system, see "Problem: Cannot connect to the service assistant" on page 94.

33. Use the **Configure enclosure** panel.

34. Use the node copy data that you recorded in step 2 on page 237 to update each of these values: **Update WWNN 1**, **Update WWNN 2**, **Update the machine type and model**, and **Update the serial number**.

    Attention: Do **not** update the system ID.
    If you were not able to record the values, use the node copy values only if none of them have all zeros as their value. If any of the node copy values are all zeros, connect the service assistant to the other node canister and configure the enclosure there. If you still do not have a full set of values, contact IBM support.

    Important: Step 35 writes the enclosure identity into the replacement midplane. The replacement midplane cannot be used as a replacement part for a different enclosure after step 35 is completed.

35. In the **Enclosure Information** panel, click **Modify**. The node canisters restart. When the restart finishes, the system comes online with both node canisters online.

    Note: In some situations, the canisters restart and report critical node error 508. If the node canisters fail to become active after they restart when the enclosure is updated, check their status by using the service assistant. If both node canisters show critical node error 508, use the service assistant to restart the nodes. For any other node error, see "Procedure: Fixing node errors" on page 121. To restart a node from the service assistant, complete the following steps:

    a. Log on to the service assistant.

    b. From the home page, select the node that you want to restart from the **Changed Node List**.

    c. Select **Actions** > **Restart**.

    The system starts and can handle I/O requests from any host systems.

36. Use the management GUI to check the status of all volumes and physical storage to ensure that everything is online.

37. Go to **Monitoring** > **Events** to check the event log for other events or errors.

38. Restart the host application and any FlashCopy activities, Global Mirror activities, or Metro Mirror activities that were stopped.

Replacing a Storwize V7000 Gen2 expansion enclosure midplane assembly

A trained service provider can use this procedure to replace a faulty Storwize V7000 Gen2 expansion enclosure midplane assembly with a new one received from CRU / FRU stock.

Before you begin

Three persons are required at step 11 on page 243.

About this task

Attention: To prevent data loss, you must shut down the system before you begin the procedure to replace an expansion enclosure midplane assembly.

The expansion enclosure midplane assembly must be replaced only by a trained service provider.

There are two models of expansion enclosure. Before you replace an expansion enclosure midplane assembly, ensure the FRU part number of the replacement part matches that of the enclosure that is being repaired.

Procedure

To replace the expansion enclosure midplane, complete the following steps.

1. Read the safety information in "Preparing to remove and replace parts" on page 177.
2. Read "Procedure: Understanding Storwize V7000 Gen2 volume dependencies" on page 140 to determine whether to continue this procedure.
3. Disconnect each power supply unit in the expansion enclosure from its power outlet so that the expansion enclosure is powered off.
4. Confirm that all the LEDs on the rear of the enclosure are off.
5. Disconnect all cables, labeling each cable to record exactly which port it was attached to (so that the cables can be inserted back into the same ports).
6. Carefully remove each hard disk drive and label it with the drive slot from which it was removed (so that the drives can be inserted back into the same slots). Refer to Figure 74 or Figure 75 on page 243.
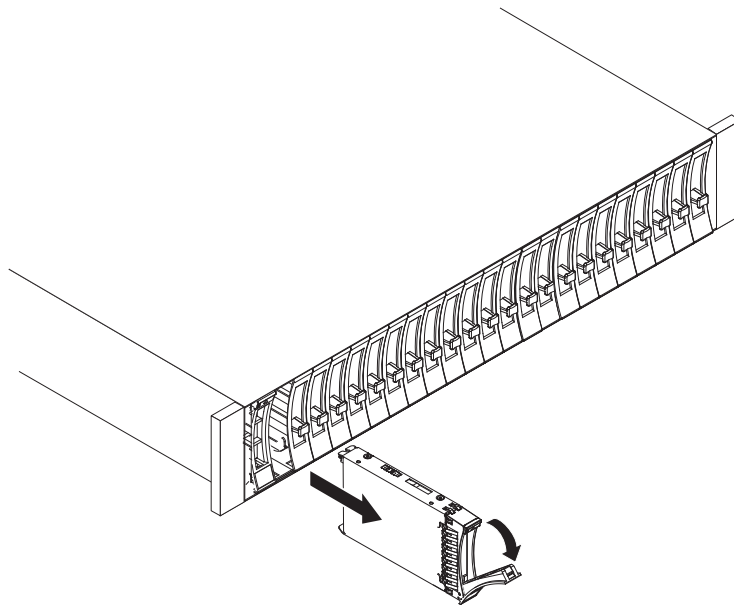
Figure 74. Removing a vertical style hard disk drive
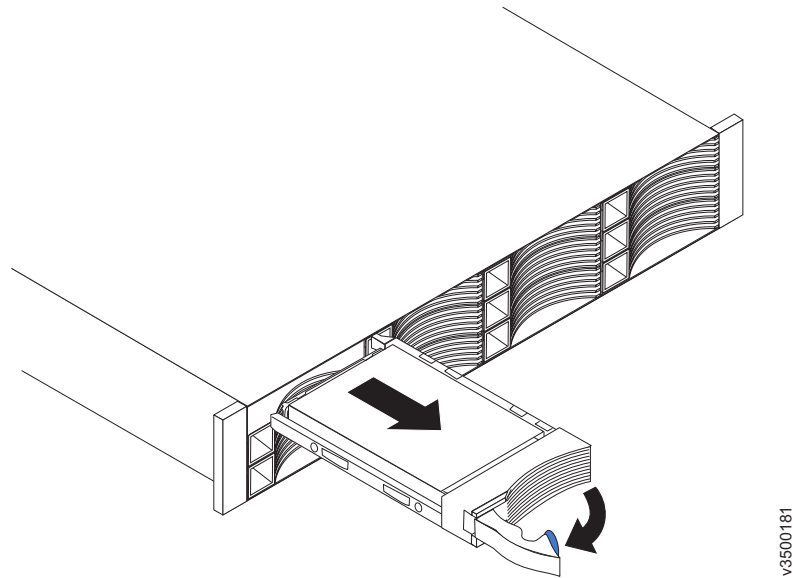
Figure 75. Removing a horizontal style hard disk drive

7. Remove the two power supplies from the enclosure. Refer to "Replacing a power supply unit for a Storwize V7000 Gen2 expansion enclosure" on page 193 for guidance.

8. Remove the expansion canisters from the enclosure. Label them to indicate which canister came from which slot.

9. Remove the end caps from the enclosure, as described in "Replacing Storwize V7000 Gen2 enclosure end caps" on page 212.

10. Remove the two screws that secure the front of the enclosure into the rack. Label these screws to indicate the location from which they are removed and place them aside.

11. Slide the enclosure from the rack cabinet, turn it onto its back so that the bottom of the enclosure is facing upwards, and place the enclosure on a flat surface.

    CAUTION:
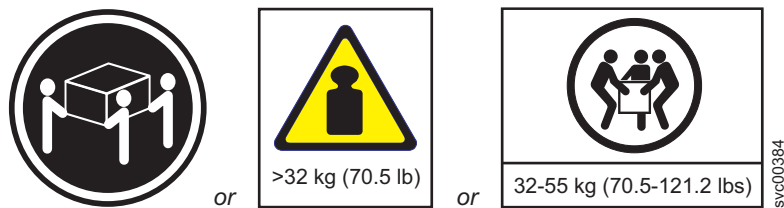
or >32 kg (70.5 lb) *or* 32-55 kg (70.5-121.2 lbs)

The weight of this part or unit is between 32 and 55 kg (70.5 and 121.2 lb). It takes three persons to safely lift this part or unit. (C010)

12. Remove the four screws from the bottom of the enclosure (see Figure 76 on page 244). Remove the three screws that are near the front and the screw that is near the middle. Label these screws to indicate the location from which they are removed and place them aside.
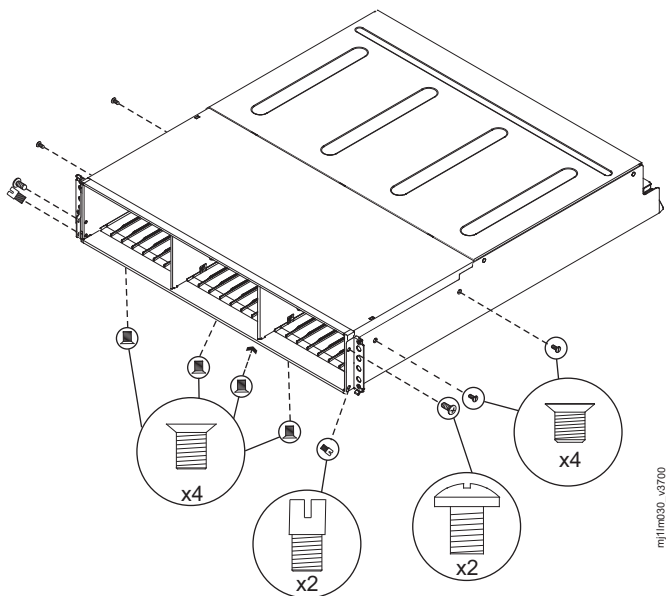
Figure 76. Removing the screws of an expansion enclosure assembly

13. Turn the enclosure top side up and place it on a flat surface.

14. Remove the three screws and one screw-pin on the right side that secure the midplane assembly to the enclosure (see Figure 76). Label the screws to indicate the location from which they are removed and place them aside.

15. Remove the three screws and one screw-pin on the left side that secure the midplane assembly to the enclosure (see Figure 76). Label the screws to indicate the location from which they are removed and place them aside. See Figure 4.

16. Remove the midplane assembly from the chassis by rotating the midplane assembly up about 45° and then lifting it out. Set the midplane assembly on a flat surface.

17. Unpack the replacement midplane assembly. Grasp the midplane assembly with two hands and hold it at a 45° angle.

18. Insert the tabs on the midplane assembly into the tab holes in the enclosure and rotate the front of the assembly down.

19. Secure the midplane assembly to the chassis on both the right and left sides of the enclosure by using the six screws and two screw-pins that you removed in steps 14 and 15.

20. Turn the enclosure over so the bottom faces upwards and insert the four screws on the bottom of the enclosure that you removed in step 12 on page 243.

21. Reinstall the enclosure in the rack cabinet, securing it with the two screws that are removed at step 10 on page 243.

22. Reinstall the end caps at the front of the enclosure, as described in "Replacing Storwize V7000 Gen2 enclosure end caps" on page 212.

23. Reinstall the hard disk drives at the front of the enclosure. Ensure that each drive is inserted back in the same slot from which it was removed.

24. Reinstall the canisters into the same slots they were removed from.

25. Reinstall the two power supplies.

26. Reconnect the data cables at the rear of the enclosure.

27. Reconnect the power to the expansion enclosure. The expansion canisters restart and the system logs an error in the event log alerting you to the unrecognized enclosure.

   Important: Step 28 writes the enclosure identity into the replacement midplane. The replacement midplane cannot be used as a replacement part for a different enclosure after step 28 is completed.

28. Go to **Monitoring** > **Events** in the management GUI. Find the error that relates to the enclosure ID of the replaced enclosure and run the fix procedure for the error.

Replacing the support rails

Remove and replace the support rails. The procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 *Gen1* refers to the enclosure models in the following table:

Table 102. Storwize V7000 Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 expansion enclosure for 2.5-inch drives

Storwize V7000 *Gen2* refers to the newer generation of enclosures in the following table:

Table 103. Storwize V7000 Gen2 model numbers

Machine type/model	Description
2076-AF6	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) flash drives
2076-624	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-524	Storwize V7000 control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-AFF	Storwize V7000 expansion enclosure, for up to 24 2.5-inch (6.35 cm) flash drives
2076-12F	Storwize V7000 expansion enclosure, for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 expansion enclosure, for up to 2.5-inch (6.35 cm) drives
2076-92F	Storwize V7000 expansion enclosure for up to 92 3.5-inch (8.89 cm) or 2.5-inch (6.35 cm) drives
2076-A9F	Storwize V7000 expansion enclosure for up to 92 flash drives

Replacing the Storwize V7000 Gen2 control enclosure support rails

You can replace faulty support rails with new ones that are received from CRU / FRU stock.

Before you begin

Three persons are required at step 7

About this task

Follow all safety precautions when you complete this procedure.

Procedure

To replace the support rails, complete the following steps.

1. Identify the enclosure that is mounted on the rails that are being replaced.

   Follow the steps in "Procedure: Identifying which Storwize V7000 Gen2 enclosure or canister to service" on page 102 to ensure that you identify the correct enclosure.

2. Shut down the system by following the steps in "Procedure: Powering off your Storwize V7000 Gen2 system" on page 132.

3. Remove power from the enclosure by unplugging both power cables from the electrical outlets.

4. Ensuring you identify which port each cable connects to, remove all cables from the back of the enclosure that has faulty support rails.

5. Remove the end caps from the front flanges of the enclosure by following the removal instructions in topic "Replacing Storwize V7000 Gen2 enclosure end caps" on page 212.

6. Unscrew the M5 screw from the left flange.

   Repeat with the M5 screw in the right flange.

7. Slide the enclosure from the rack.

   CAUTION:

   or >32 kg (70.5 lb) *or* 32-55 kg (70.5-121.2 lbs) svc00384

   The weight of this part or unit is between 32 and 55 kg (70.5 and 121.2 lb). It takes three persons to safely lift this part or unit. (C010)

8. Locate the left support rail.

   Record the shelf number of the support rail so that the replacement rails can be installed into the same position.

9. At the rear of the rack, remove the securing M5 screw from the bottom hole of the rear bracket of the rail, then open the rear hinge bracket (Figure 77 on page 247).
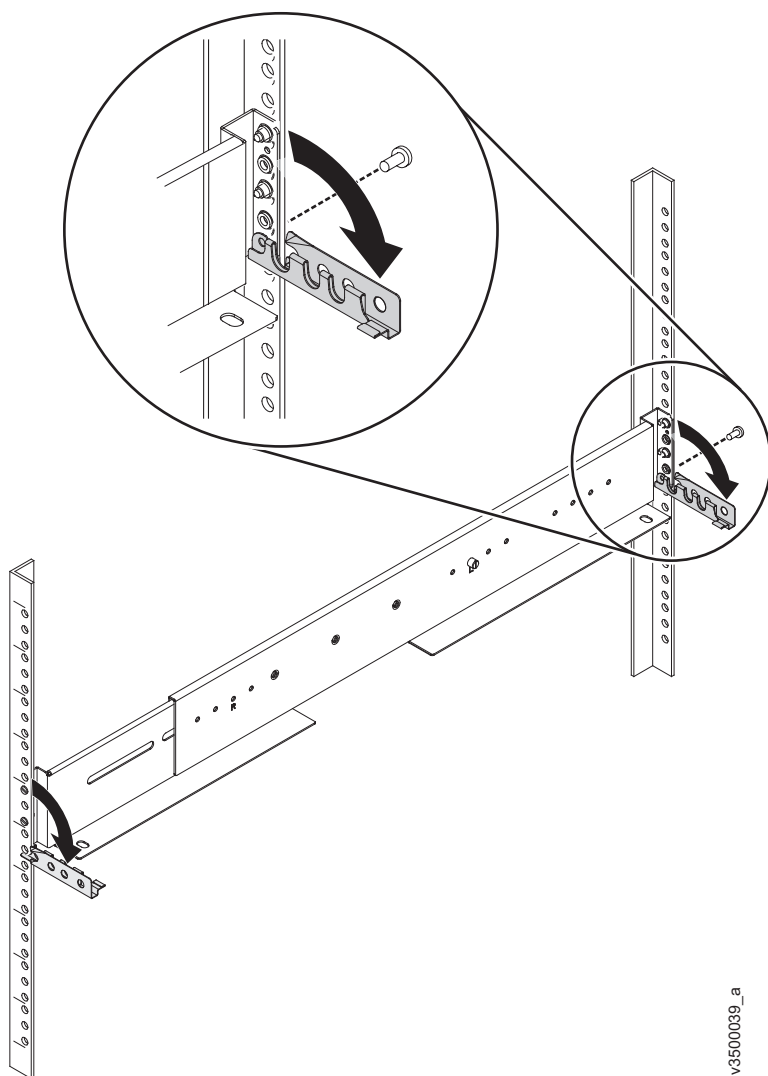
Figure 77. Opening rear hinge bracket of mounting rail

10. At the front of the rack, hold onto the rail and open the front hinge bracket.
11. Compress the rail against its spring to shorten it, then remove it from inside the rack (Figure 78 on page 248).

v3500116_a
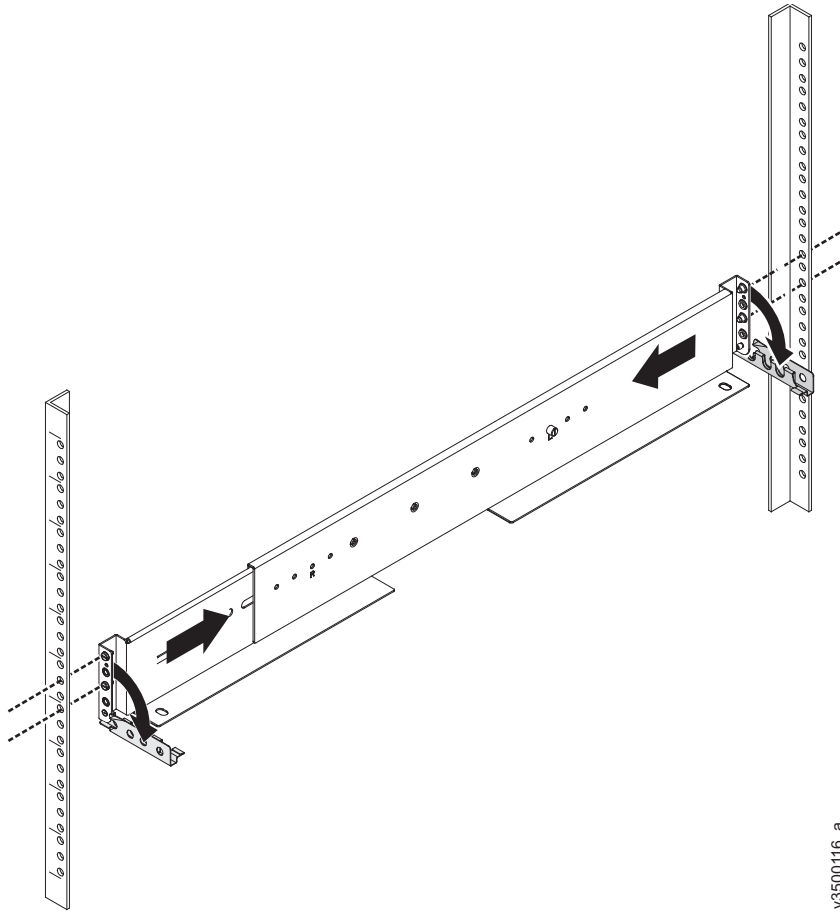
Figure 78. Compressing rail for removal from rack

12. Repeat steps 9 on page 246 to 11 on page 247 on the right support rail.

13. Install the new support rails at the rack position that is recorded at step 8 on page 246 by following the instructions in Installing support rails for the Storwize V7000 Gen2 and Storwize V7000 Gen2+ control enclosure.

14. Reinstall the enclosure (removed at step 7 on page 246) and the end caps (removed at step 5 on page 246) by following the instructions in Installing the enclosures.

15. If components were removed from the enclosure at step 7 on page 246, return each canister, drive assembly, and power supply unit to its labeled slot.

16. Reconnect the cables, ensuring that they are connected to their original ports.

17. Reconnect the power supply cables to their original power supply and electrical outlet.

    The system starts.

18. After the system is online, use the management GUI to verify that the system is correct.

Replacing the Storwize V7000 Gen2 expansion enclosure support rails

You can replace faulty support rails with new ones that are received from CRU / FRU stock.

Before you begin

Two persons are required at step 7

Procedure

To replace the support rails, complete the following steps.

1. Identify the enclosure that is mounted on the rails that are being replaced.

   Follow the steps in "Procedure: Identifying which Storwize V7000 Gen2 enclosure or canister to service" on page 102 to ensure that you identify the correct enclosure.

2. Shut down the system by following the steps in "Procedure: Powering off your Storwize V7000 Gen2 system" on page 132.

3. Remove power from the enclosure by unplugging both power cables from the electrical outlets.

4. Ensuring you identify which port each cable connects to, remove all cables from the back of the enclosure that has faulty support rails.

5. Remove the end caps from the front flanges of the enclosure by following the removal instructions in topic "Replacing Storwize V7000 Gen2 enclosure end caps" on page 212.

6. Unscrew the M5 screw from the left flange.

   Repeat with the M5 screw in the right flange.

7. Slide the enclosure from the rack.

   CAUTION:
   The weight of this part or unit is between 18 and 32 kg (39.7 and 70.5 lb). It takes two persons to safely lift this part or unit. (C009)

8. Locate the left support rail.

   Record the shelf number of the support rail so that the replacement rails can be installed into the same position.

9. At the rear of the rack, remove the securing M5 screw from the bottom hole of the rear bracket of the rail, then open the rear hinge bracket (Figure 79 on page 250).
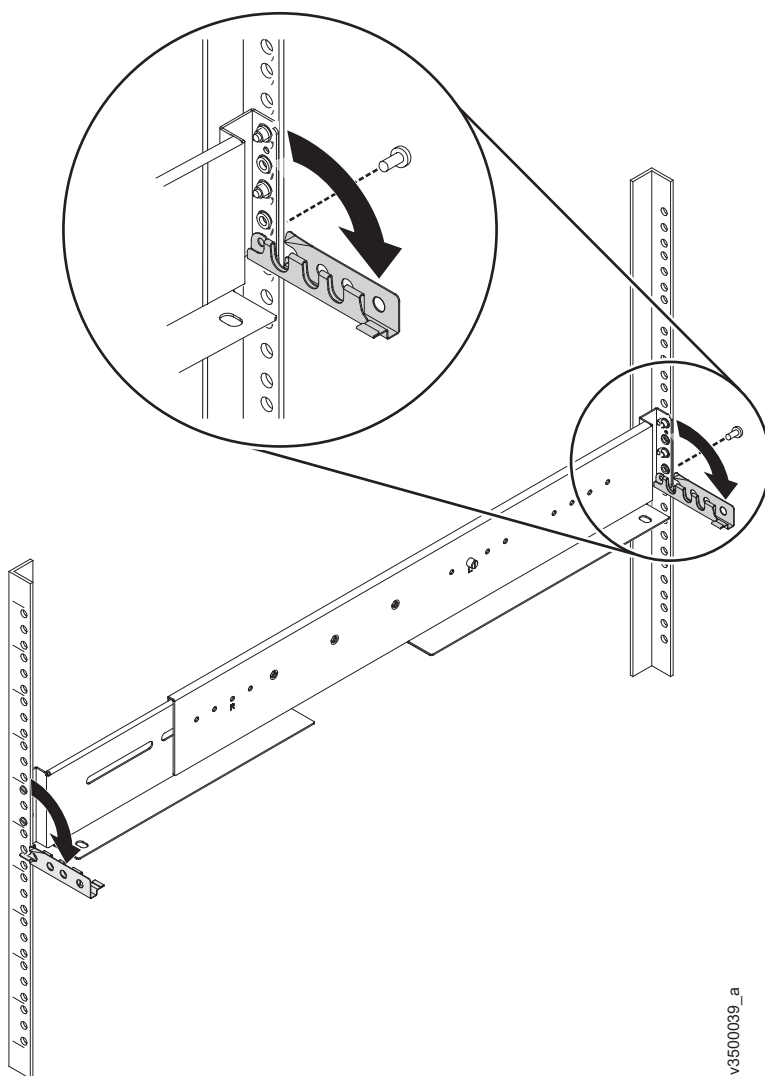
Figure 79. Opening rear hinge bracket of mounting rail

10. At the front of the rack, hold onto the rail and open the front hinge bracket.
11. Compress the rail against its spring to shorten it, then remove it from inside the rack (Figure 80 on page 251).
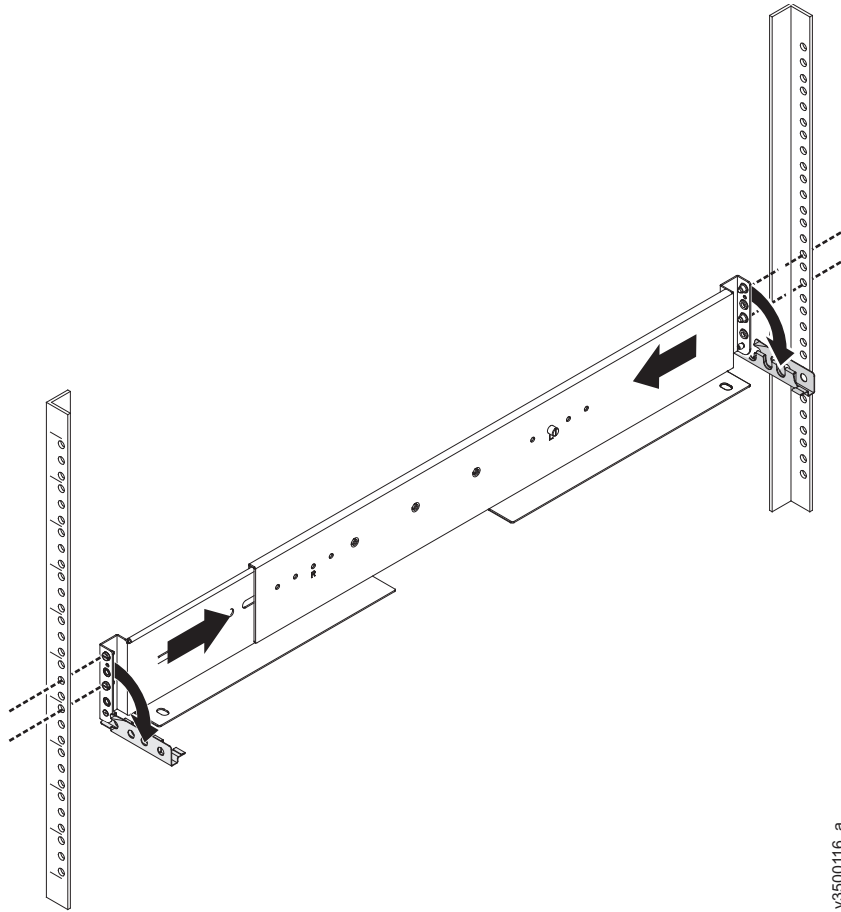
Figure 80. Compressing rail for removal from rack

12. Repeat steps 9 on page 249 to Figure 80 on the right support rail.

13. Install the new support rails at the rack position that is recorded at step 8 on page 249 by following the instructions in Installing support rails for Storwize V7000 Gen2 and Storwize V7000 Gen2+ expansion enclosures .

14. Reinstall the enclosure (removed at step 7 on page 249) and the end caps (removed at step 5 on page 249) by following the instructions in Installing the enclosures.

15. If components were removed from the enclosure at step 7 on page 249, return each canister, drive assembly, and power supply unit to its labeled slot.

16. Reconnect the cables, ensuring that they are connected to their original ports.

17. Reconnect the power supply cables to their original power supply and electrical outlet.

    The system starts.

18. After the system is online, use the management GUI to verify that the system is correct.

Replacing the Storwize V7000 Gen1 support rails

You can replace the support rails.

Procedure

To replace the support rails, complete the following steps:

1. Remove the enclosure.
2. Record the location of the rail assembly in the rack cabinet.
3. Working from the back of the rack cabinet, remove the clamping screw **1** from the rail assembly on both sides of the rack cabinet.
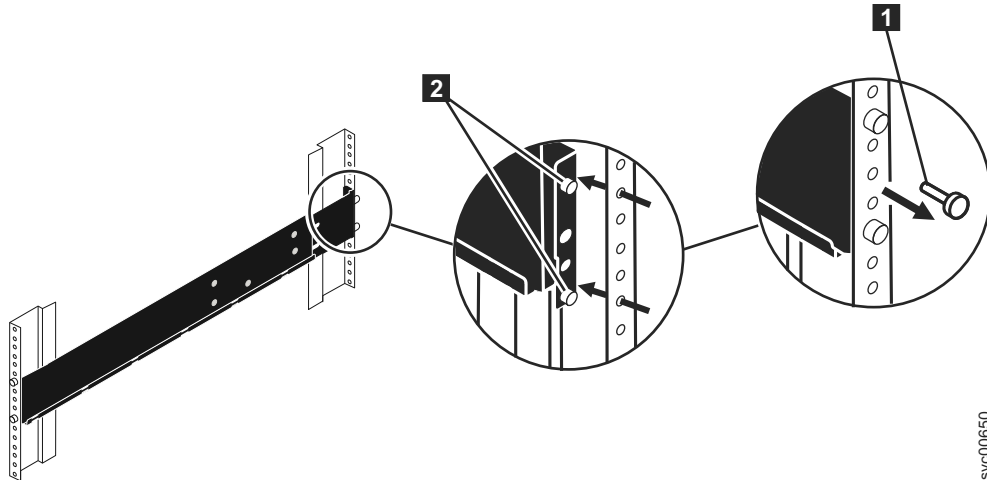
Figure 81. Removing a rail assembly from a rack cabinet

4. Working from the front of the rack cabinet, remove the clamping screw from the rail assembly on both sides of the rack cabinet.
5. From one side of the rack cabinet, grip the rail and slide the rail pieces together to shorten the rail.
6. Disengage the rail location pins **2** .
7. From the other side the rack cabinet, grip the rail and slide the rail pieces together to shorten the rail.
8. Disengage the rail location pins **2** .
9. Starting from the location of the previous rail assembly, align the bottom of the rail with the bottom of the two rack units. Insert the rail location pins through the holes in the rack cabinet.
10. Insert a clamping screw into the upper mounting hole between the rail location pins.
11. Tighten the screw to secure the rail to the rack.
12. Working from the rear of the rack cabinet, extend the rail that you secured to the front to align the bottom of the rail with the bottom of the two rack units.

    Note: Ensure that the rail is level between the front and the back.

13. Insert the rail location pins through the holes in the rack cabinet.
14. Insert a clamping screw into the upper mounting hole between the rail location pins.
15. Tighten the screw to secure the rail to the rack from the back side.
16. Repeat the steps to secure the opposite rail to the rack cabinet.

Replacing node canister memory modules

Remove and replace node canister memory modules.

Replacing a Storwize V7000 Gen2 node canister memory module (16 GB DIMM)

You can replace a faulty node canister memory module (16 GB DIMM) with a new one received from CRU / FRU stock.

Before you begin

Note: Storwize V7000 2076-524 and Storwize V7000 2076-624 models use different DIMMs. Be certain to use the correct replacement DIMM for your model. The memory type for each model is as follows.

- Storwize V7000 2076-524 DIMMs are 16 GB, 240 p socket, DDR3.
- Storwize V7000 2076-624 DIMMs are 16 GB, 288 p socket, DDR4.

Procedure

1. Follow "Procedure: Removing a Storwize V7000 Gen2 node canister" on page 129 to disconnect and remove the node canister with the faulty memory.
2. Remove the lid of the canister, as described in "Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister" on page 139.
3. Locate the DIMM slot with the faulty DIMM. Slot 1 is next to the battery area. Slot 2 is next to the processor. The slots are marked **1** , **2** , **3** , **4** as shown in Figure 82 .
4. Remove the faulty DIMM by applying gentle, outwards pressure simultaneously to the retaining clips at each end of the DIMM slot until the DIMM is levered out of the slot.
5. Touch the replacement DIMM packaging onto a metal area of the case, then remove the replacement DIMM from its package.
6. Ensure that the retaining clips of the DIMM slot are open.
7. Gently place the DIMM in the slot, ensuring that the notches in the DIMM align with the shape of the slot, as shown in Figure 82.
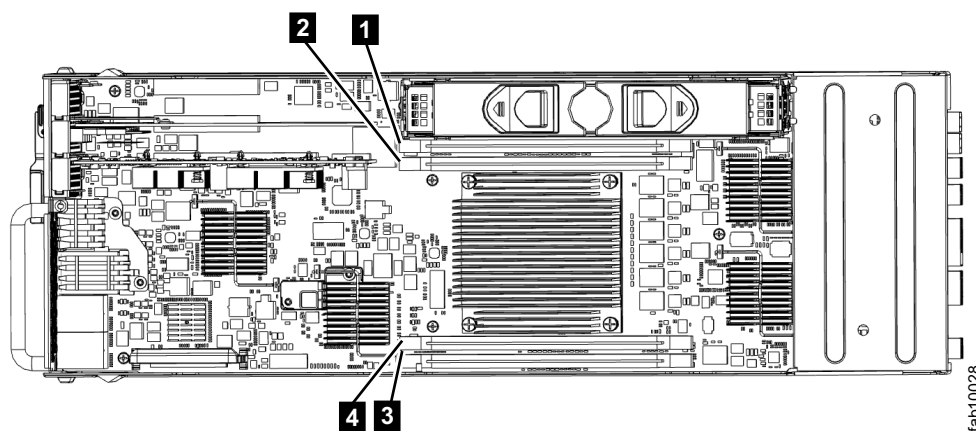
Figure 82. Installing a Storwize V7000 2076-524 node canister memory module

8. Apply even, firm, downwards pressure on the DIMM in its slot until the retaining clips move inwards and engage the edges of the DIMM.

9. Ensure that the retaining clips are fully engaged with the edges of the DIMM. Gently pull the DIMM upwards and ensure that it does not become dislodged.

10. Replace the canister lid, as described in "Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister" on page 139.

11. Reinstall the canister, as described in "Replacing a Storwize V7000 Gen2 node canister" on page 177, into the enclosure from which it was removed in step 1 on page 253. The node canister starts.

12. Reconnect the cables to the canister, ensuring cables go into the same ports from which they were removed in step 1 on page 253.

13. When the canister is back online, check the event log for new events, particularly events that relate to hardware changes.

Replacing a host interface adapter

Remove and replace a host interface adapter.

Replacing a Storwize V7000 Gen2 host interface adapter

To replace a faulty host interface adapter in a Storwize V7000 2076-524 with a new one received from customer replaceable unit (CRU) or field replaceable unit (FRU) stock, use this procedure.

About this task

For lists of supported host interface adapters, refer to "Storwize V7000 2076-524 Gen2 replaceable units" on page 169.

Important: For correct operation, use the correct SFP transceivers with each adapter card. The topic "Storwize V7000 2076-524 Gen2 replaceable units" identifies the suitable IBM parts.

- Use only 8G bps SFP transceivers in the 8 Gbps Fibre Channel adapter cards.
- Use only 16 Gbps SFP transceivers in the 16 Gbps Fibre Channel adapter cards.
- Use only 10 Gbps SFP transceivers in the 10 Gbps Ethernet (FCoE/iSCSI) adapter card.

Procedure

Complete the following steps to replace a host interface adapter.

1. Complete "Procedure: Removing a Storwize V7000 Gen2 node canister" on page 129 to remove the Storwize V7000 2076-524 node canister with the faulty host interface adapter.

2. Identify which host interface adapter is to be removed. The interface adapters are in slots numbered 2 and 3

3. Remove any small form-factor pluggable SFP transceiver from each rear-facing port of the host interface adapter and put safely to one side.

4. Complete "Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister" on page 139 to remove and replace the lid of a Storwize V7000 2076-524 node canister.

5. Gently pull the host interface adapter upward to disconnect it **2** , and then carefully remove it from the canister **1** . Figure 83 on page 255 displays removing the host interface adapter.
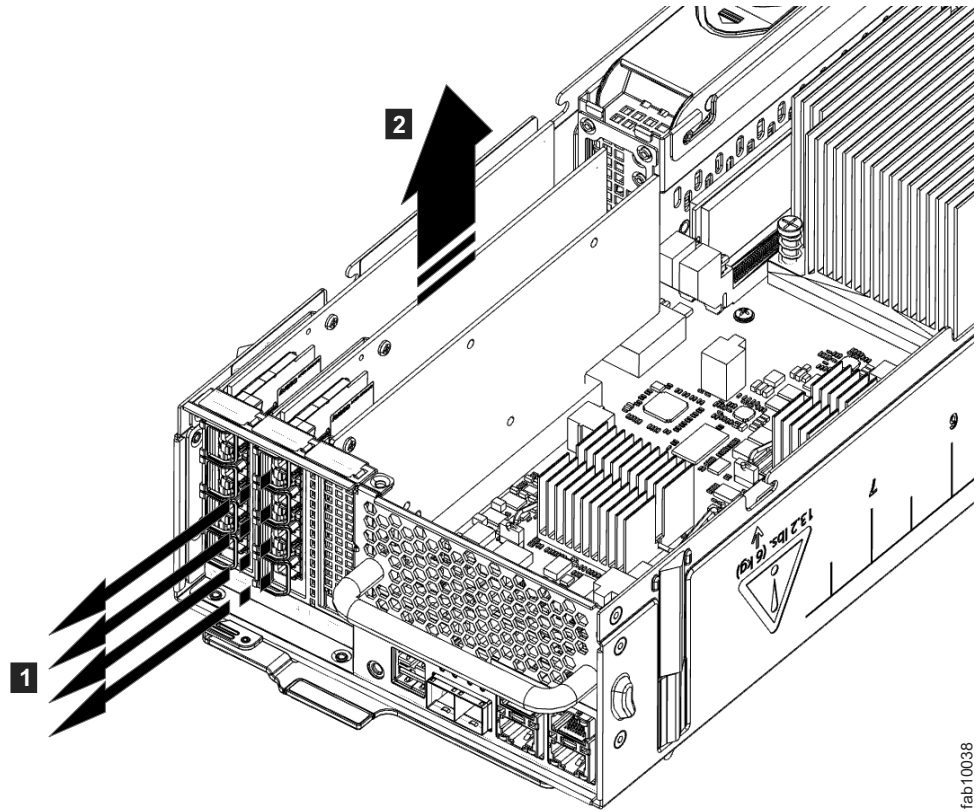
Figure 83. Removing the host interface adapter

6. Remove the replacement host interface adapter **1** from its package. Figure 84 on page 256 displays installing the host interface adapter.

7. Set the connecting edge of the replacement host interface adapter **3** on the host interface adapter connector so that the connectors are aligned.

8. Ensure that the adapter is perpendicular to the canister main board so that the small tab on the top of the bracket **2** is aligned with the alignment hole in the top edge of the slot.

9. Maintain alignment while applying pressure to the top edge of the host interface adapter opposite the connecting edge to push the host interface adapter into the connector **4** and **5** .
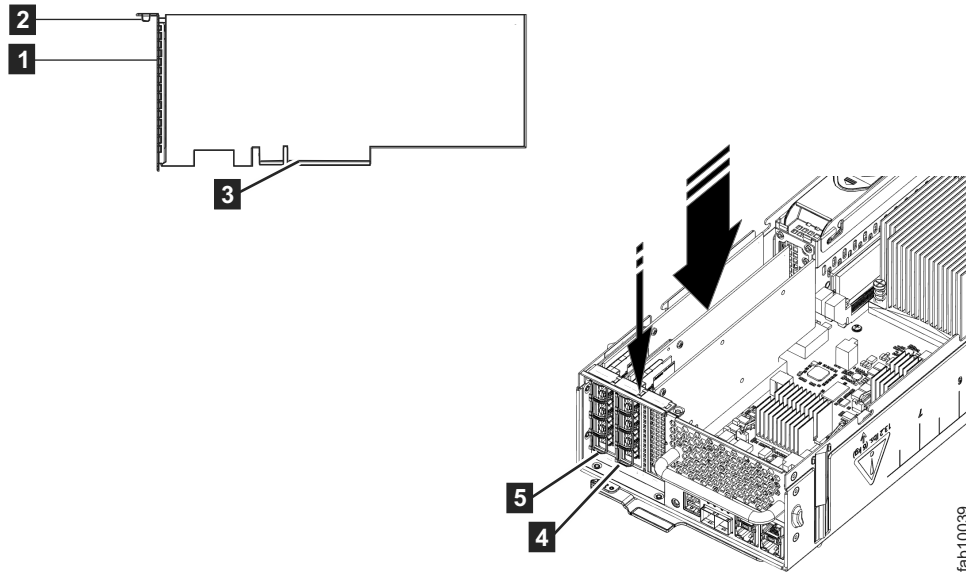
Figure 84. Installing the host interface adapter

10. Check that the host interface adapter is installed squarely in its slot. If the small tab of the mounting bracket is not positioned correctly, repeat steps 5 on page 254 onward to install the adapter correctly.

11. Replace the canister lid, as described in "Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister" on page 139.

12. If any SFP transceivers were removed from the rear-facing ports of the host interface adapter at step 2 on page 254, ensure each one is reinstalled by following the installation steps of "Replacing an SFP transceiver in a Storwize V7000 2076-524 control enclosure" on page 185.

13. Reinstall the canister into the enclosure from which it was removed in step 1 on page 254 following "Replacing a Storwize V7000 Gen2 node canister" on page 177. The node canister starts.

14. Reconnect the cables to the canister, ensuring cables go into the same ports from which they were removed in step 1 on page 254.

15. When the canister is back online, check the event log for any new events relating to hardware changes.

Replacing Storwize V7000 Gen2 host interface adapters in two control enclosures concurrently

It is possible to reconfigure one node canister of each control enclosure at the same time. During the procedure, both I/O groups (control enclosures) are online with no redundancy, but the total maintenance period is reduced.

To replace host interface adapters in both control enclosures concurrently, use the procedure for replacing the host interface adapter in a single enclosure, but complete each step in both enclosures before you continue to the next step. Table 104 on page 257 shows how to sequence the step in each node. Work your way down the table, completing each row before you start the next row.

For the procedure for replacing the host interface adapter in a single enclosure, refer to the "Replacing a Storwize V7000 Gen2 host interface adapter" on page 254.

Table 104. Replacing host interface adapters in two control enclosures concurrently

Control enclosure 1		Control enclosure 2	
Node canister 1	Node canister 2	Node canister 1	Node canister 2
Step 1		Step 1	
Step 2		Step 2	
...		...	
Final step		Final step	
	Step 1		Step 1
	Step 2		Step 2
	...		...
	Final step		Final step

Replacing a CMOS battery

Remove and replace the complementary metal-oxide semiconductor (CMOS) battery.

Replacing a Storwize V7000 Gen2 CMOS battery

The complementary metal-oxide semiconductor (CMOS) battery is a coin-shaped power cell that is mounted inside a node canister. It is used to keep the system time when there is no power to the node canister. The lithium battery must be handled correctly to avoid possible danger. If you replace the battery, you must adhere to all safety instructions.

About this task

Use this procedure to replace a CMOS battery. Dispose of the faulty battery properly.

CAUTION: If your system has a module containing a lithium battery, replace it only with the same module type made by the same manufacturer. The battery contains lithium and can explode if not properly used, handled, or disposed of.

Do not:
- Throw or immerse into water
- Heat to more than 100°C (212°F)
- Repair or disassemble

Dispose of the battery as required by local ordinances or regulations. (C045)

Procedure

1. Complete "Procedure: Removing a Storwize V7000 Gen2 node canister" on page 129
2. Open the canister and remove the lid as described in "Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister" on page 139.
3. Locate the CMOS battery inside the node canister.. See Figure 85 on page 258
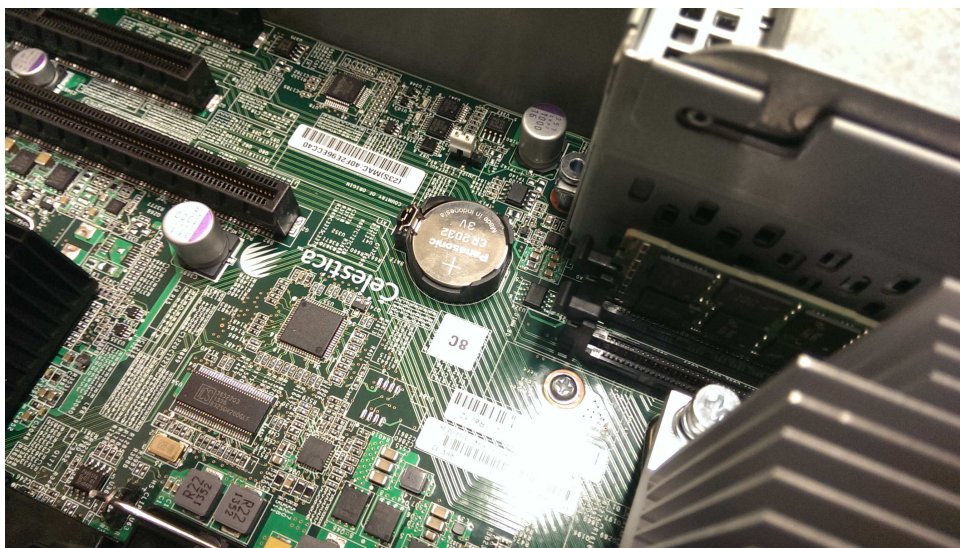
Figure 85. Replacing a CMOS Gen2 battery

4. Push the coin cell latch to the side to release the coin cell from its holder, then remove the expired coin cell.

5. Orient the replacement coin cell with the flat side upwards and place it down onto the coin cell holder.

6. Gently push the coin cell down into the holder so that it clicks under the latch and sits parallel with the canister main board.

7. Replace the canister lid as described in "Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister" on page 139.

8. Reinsert the canister into the slot from which it came.

9. Reconnect all cables.

10. Open the management GUI.

11. Use the management GUI to check that the time and date settings of the system are correct.

12. In the event log view, if a CMOS battery error is present, run the fix procedure.

Replacing a Storwize V7000 Gen2 compression-accelerator

Use this procedure to replace a faulty compression accelerator- with a new one received from customer replaceable unit (CRU) or field replaceable unit (FRU) stock.

Procedure

1. Complete "Procedure: Removing a Storwize V7000 Gen2 node canister" on page 129 to remove the faulty compression-accelerator.

2. Follow the "Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister" on page 139 to remove the lid.

3. Locate the faulty compression-accelerator in slot 1. Slot 1 is next to the two host interface adapters in slot 2 and slot 3.

4. Gently pull the faulty compression-accelerator upwards to disconnect it and then carefully remove it from the canister.

5. Remove the replacement compression-accelerator from its package.

6. Set the connecting edge of the replacement compression-accelerator on the connector so that the connectors are aligned.

7. Ensure that the compression-accelerator is perpendicular to the canister main board so that the small tab on the top of the bracket is aligned with the alignment hole in the top edge of the slot.

8. Maintain alignment while applying pressure to the top edge of the compression-accelerator opposite the connecting edge to push the compression-accelerator into the connector.

9. Check that the compression-accelerator is installed squarely in its slot. If the small tab of the mounting bracket is not positioned correctly, repeat steps 6 onwards to install the compression-accelerator correctly.

10. Replace the canister lid as described in "Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister" on page 139.

11. Reinstall the canister in to the enclosure from which it was removed in step 1 on page 258 following the "Replacing a Storwize V7000 Gen2 node canister" on page 177 procedure.

12. Reconnect the cables to the canister and ensure that the cables go into the same ports from which they were removed in step 1 on page 258.

13. When the canister is back online, check the event log for any new events relating to hardware changes.

Replacing a Storwize V7000 Gen2 compression pass-through adapter

Use this procedure to replace a faulty compression pass-through adapter with a new one received from customer replaceable unit (CRU) or field replaceable unit (FRU) stock.

Procedure

1. Complete "Procedure: Removing a Storwize V7000 Gen2 node canister" on page 129 to remove the faulty compression pass-through adapter.

2. Follow the "Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister" on page 139 to remove the lid.

3. Locate the faulty compression pass-through adapter in slot 1. Slot 1 is next to the two host interface adapters in slot 2 and slot 3.

4. Gently pull the faulty compression pass-through adapter upwards to disconnect it and then carefully remove it from the canister.

5. Remove the replacement compression pass through adapter from its package.

6. Set the connecting edge of the replacement compression pass-through adapter on the connector so that the connectors are aligned.

7. Ensure that the adapter is perpendicular to the canister main board so that the small tab on the top of the bracket is aligned with the alignment hole in the top edge of the slot.

8. Maintain alignment while applying pressure to the top edge of the compression pass-through adapter opposite the connecting edge to push the compression pass-through adapter into the connector.

9. Check that the compression pass through adapter is installed squarely in its slot. If the small tab of the mounting bracket is not positioned correctly, repeat steps 6 onwards to install the adapter correctly.

10. Replace the canister lid as described in "Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister" on page 139.

11. Reinstall the canister in to the enclosure from which it was removed in step 1 on page 259 following the "Replacing a Storwize V7000 Gen2 node canister" on page 177 procedure.

12. Reconnect the cables to the canister and ensure that the cables go into the same ports from which they were removed in step 1 on page 259.

13. When the canister is back online, check the event log for any new events relating to hardware changes.

Chapter 10. Event reporting

Events that are detected are saved in an event log. As soon as an entry is made in this event log, the condition is analyzed. If any service activity is required, a notification is sent, if you have set up notifications.

Event reporting process

The following methods are used to notify you and the IBM Support Center of a new event:

- If you enabled Simple Network Management Protocol (SNMP), an SNMP trap is sent to an SNMP manager that is configured by the customer.
- If enabled, log messages can be forwarded on an IP network by using the syslog protocol.
- If enabled, event notifications can be forwarded by email by using Simple Mail Transfer Protocol (SMTP).
- Call Home can be enabled so that critical faults generate a problem management record (PMR) that is then sent directly to the appropriate IBM Support Center by using email.

Understanding events

When a significant change in status is detected, an event is logged in the event log.

Error data

Events are classified as either alerts or messages:

- An alert is logged when the event requires some action. Some alerts have an associated error code that defines the service action that is required. The service actions are automated through the fix procedures. If the alert does not have an error code, the alert represents an unexpected change in state. This situation must be investigated to see if it is expected or represents a failure. Investigate an alert and resolve it as soon as it is reported.
- A message is logged when a change that is expected is reported, for instance, an IBM FlashCopy operation completes.

Viewing the event log

You can view the event log by using the management GUI or the command-line interface (CLI).

About this task

You can view the event log by using the **Monitoring** > **Events** options in the management GUI. The event log contains many entries. You can, however, select only the type of information that you need.

You can also view the event log by using the command-line interface (`lseventlog`). See the "Command-line interface" topic for the command details.

Managing the event log

The event log has a limited size. After it is full, newer entries replace entries that are no longer required.

To avoid having a repeated event that fills the event log, some records in the event log refer to multiple occurrences of the same event. When event log entries are coalesced in this way, the time stamp of the first occurrence and the last occurrence of the problem is saved in the log entry. A count of the number of times that the error condition has occurred is also saved in the log entry. Other data refers to the last occurrence of the event.

Describing the fields in the event log

The event log includes fields with information that you can use to diagnose problems.

Table 105 describes some of the fields that are available to assist you in diagnosing problems.

Table 105. Description of data fields for the event log

Data field	Description
Event ID	This number precisely identifies why the event was logged.
Description	A short description of the event.
Status	Indicates whether the event requires some attention. Alert: if a red icon with a cross is shown, follow the fix procedure or service action to resolve the event and turn the status green. Monitoring: the event is not yet of concern. Expired: the event no longer represents a concern. Message: provide useful information about system activity.
Error code	Indicates that the event represents an error in the system that can be fixed by following the fix procedure or service action that is identified by the error code. Not all events have an error code. Different events have the same error code if the same service action is required for each.
Sequence number	Identifies the event within the system.
Event count	The number of events that are coalesced into this event log record.
Object type	The object type to which the event relates.
Object ID	Uniquely identifies the object within the system to which the event relates.
Object name	The name of the object in the system to which the event relates.
Copy ID	If the object is a volume and the event refers to a specific copy of the volume, this field is the number of the copy to which the event relates.
Reporting node ID	Typically identifies the node responsible for the object to which the event relates. For events that relate to nodes, it identifies the node that logged the event, which can be different from the node that is identified by the object ID.
Reporting node name	Typically identifies the node that contains the object to which the event relates. For events that relate to nodes, it identifies the node that logged the event, which can be different from the node that is identified by the object name.

Table 105. Description of data fields for the event log (continued)

Data field	Description
Fixed	Where an alert is shown for an error or warning condition, it indicates that the user marked the event as fixed, completed the fix procedure, or that the condition was resolved automatically. For a message event, this field can be used to acknowledge the message.
First time stamp	The time when this error event was reported. If events of a similar type are being coalesced together, so that one event log record represents more than one event, this field is the time the first error event was logged.
Last time stamp	The time when the last instance of this error event was recorded into this event log record.
Root sequence number	If set, it is the sequence number of an event that represents an error that probably caused this event to be reported. Resolve the root event first.
Sense data	Additional data that gives the details of the condition that caused the event to be logged.

Event notifications

Storwize V7000 can use Simple Network Management Protocol (SNMP) traps, syslog messages, emails and Call Homes to notify you and IBM(r) Remote Technical Support when significant events are detected. Any combination of these notification methods can be used simultaneously. Notifications are normally sent immediately after an event is raised. However, there are some events that might occur because of service actions that are being performed. If a recommended service action is active, these events are notified only if they are still unfixed when the service action completes.

Only events recorded in the event log can be notified. Most CLI messages in response to some CLI commands are not recorded in the event log so do not cause an event notification.

Table 106 describes the levels of event notifications.

Table 106. Notification levels

Notification level	Description
Error	Error notification is sent to indicate a problem that must be corrected as soon as possible. This notification indicates a serious problem with the system. For example, the event that is being reported could indicate a loss of redundancy in the system, and it is possible that another failure could result in loss of access to data. The most typical reason that this type of notification is sent is because of a hardware failure, but some configuration errors or fabric errors also are included in this notification level. Error notifications can be configured to be sent as a call home message to your support center.

Table 106. Notification levels (continued)

Notification level	Description
Warning	A warning notification is sent to indicate a problem or unexpected condition with the system. Always immediately investigate this type of notification to determine the effect that it might have on your operation, and make any necessary corrections. A warning notification does not require any replacement parts and therefore should not require involvement from your support center. The allocation of notification type Warning does not imply that the event is less serious than one that has notification level Error.
Information	An informational notification is sent to indicate that an expected event has occurred: for example, a FlashCopy operation has completed. No remedial action is required when these notifications are sent.

Power-on self-test

When you turn on the system, the node canisters complete self-tests.

A series of tests is completed to check the operation of components and some of the options that have been installed when the units are first turned on. This series of tests is called the power-on self-test (POST).

If a critical failure is detected during the POST, the software is not loaded and the fault LED is illuminated. To determine if there is a POST error on a canister, go to "Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs" on page 112.

When the code is loaded, additional testing takes place, which ensures that all of the required hardware and code components are installed and functioning correctly.

Understanding the error codes

Error codes are generated by the event-log analysis and system configuration code.

Error codes help you to identify the cause of a problem, a failing component, and the service actions that might be needed to solve the problem.

Event IDs

The Storwize V7000 software generates events, such as informational events and error events. An event ID or number is associated with the event and indicates the reason for the event.

Informational events provide information about the status of an operation. Informational events are recorded in the event log, and, depending on the configuration, informational event notifications can be sent through email, SNMP, or syslog.

Error events are generated when a service action is required. An error event maps to an alert with an associated error code. Depending on the configuration, error event notifications can be sent through email, SNMP, or syslog.

Informational events

The informational events provide information about the status of an operation.

Informational events are recorded in the event log and, based on notification type, can generate notifications through email, SNMP, or syslog. Informational events are distinguished from error events, which are associated with error codes and might require service procedures. For a list of error events, see "Error event IDs and error codes" on page 269.

Informational events can be either notification type I (information) or notification type W (warning). An informational event report of type (W) might require user attention. Table 107 provides a list of informational events, the notification type, and the reason for the event.

Table 107. Informational events

Event ID	Notification type	Description
062004	I	Type conversion completed and the original copy has been deleted.
980221	I	The error log is cleared.
980230	I	The SSH key was discarded for the service login user.
980231	I	User name has changed.
980301	I	Degraded or offline managed disk is now online.
980310	I	A degraded or offline storage pool is now online.
980320	I	Offline volume is now online.
980321	W	Volume is offline because of degraded or offline storage pool.
980330	I	All nodes can see the port.
980349	I	A node has been successfully added to the cluster (system).
980350	I	The node is now a functional member of the cluster (system).
980351	I	A noncritical hardware error occurred.
980352	I	Attempt to automatically recover offline node starting.
980370	I	Both nodes in the I/O group are available.
980371	I	One node in the I/O group is unavailable.
980372	W	Both nodes in the I/O group are unavailable.
980380	I	Maintenance mode was started.
980381	I	Maintenance mode has ended.
980392	I	Cluster (system) recovery completed.
980435	W	Failed to obtain directory listing from remote node.
980440	W	Failed to transfer file from remote node.
980445	I	The migration is complete.
980446	I	The secure delete is complete.
980501	W	The virtualization amount is close to the limit that is licensed.

Table 107. Informational events (continued)

Event ID	Notification type	Description
980502	W	The FlashCopy feature is close to the limit that is licensed.
980503	W	The Metro Mirror or Global Mirror feature is close to the limit that is licensed.
980504	I	The limit was reached for the external virtualization feature.
980505	I	The limit was reached for the compression feature license.
981002	I	Fibre Channel discovery occurred; configuration changes are pending.
981003	I	Fibre Channel discovery occurred; configuration changes are complete.
981004	I	Fibre Channel discovery occurred; no configuration changes were detected.
981007	W	The managed disk is not on the preferred path.
981009	W	The initialization for the managed disk failed.
981014	W	The LUN discovery has failed. The cluster (system) has a connection to a device through this node but this node cannot discover the unmanaged or managed disk that is associated with this LUN.
981015	W	The LUN capacity equals or exceeds the maximum. Only part of the disk can be accessed.
981020	W	The managed disk error count warning threshold has been met.
981022	I	Managed disk offline imminent, offline prevention started
981025	I	Drive firmware download completed successfully
981026	I	Drive FPGA download completed successfully
981027	I	Drive firmware download started
981028	I	Drive FPGA download started
981029	I	Drive firmware download cancelled by user
981101	I	SAS discovery occurred; no configuration changes were detected.
981102	I	SAS discovery occurred; configuration changes are pending.
981103	I	SAS discovery occurred; configuration changes are complete.
981104	W	The LUN capacity equals or exceeds the maximum capacity. Only the first 1 PB of disk will be accessed.
981105	I	The drive format has started.
981106	I	The drive recovery was started.
981110	I	iSCSI discovery occurred, configuration changes pending.
981111	I	iSCSI discovery occurred, configuration changes complete.

Table 107. Informational events (continued)

Event ID	Notification type	Description
981112	I	iSCSI discovery occurred, no configuration changes were detected.
982003	W	Insufficient virtual extents.
982004	W	The migration suspended because of insufficient virtual extents or too many media errors on the source managed disk.
982007	W	Migration has stopped.
982009	I	Migration is complete.
982010	W	Copied disk I/O medium error.
983001	I	The FlashCopy operation is prepared.
983002	I	The FlashCopy operation is complete.
983003	W	The FlashCopy operation has stopped.
984001	W	First customer data being pinned in a volume working set.
984002	I	All customer data in a volume working set is now unpinned.
984003	W	The volume working set cache mode is in the process of changing to synchronous destage because the volume working set has too much pinned data.
984004	I	Volume working set cache mode updated to allow asynchronous destage because enough customer data has been unpinned for the volume working set.
984005	W	Volumes have been taken offline because the storage pool's cache is full of data that cannot be destaged.
984006	W	Volume copies have been taken offline because the storage pool's cache is full of data that cannot be destaged.
984007	W	Volumes have been taken offline because a node's cache is full of data that cannot be destaged.
984008	W	Volume copies have been taken offline because a node's cache is full of data that cannot be destaged.
984501	I	The firmware level of an enclosure component is being updated.
984502	I	The firmware level updated has completed.
984503	I	The battery conditioning completed.
984504	I	The battery conditioning started.
984505	I	The statesave information for the enclosure was collected.
984506	I	The debug from an IERR was extracted to disk.
984507	I	An attempt was made to power on the slots.
984508	I	All the expanders on the strand were reset.
984509	I	The component firmware update paused to allow the battery charging to finish.
984511	I	The update for the component firmware paused because the system was put into maintenance mode.

Table 107. Informational events (continued)

Event ID	Notification type	Description
984512	I	A component firmware update is needed but is prevented from running.
985001	I	The Metro Mirror or Global Mirror background copy is complete.
985002	I	The Metro Mirror or Global Mirror is ready to restart.
985003	W	Unable to find path to disk in the remote cluster (system) within the timeout period.
986001	W	The thin-provisioned volume copy data in a node is pinned.
986002	I	All thin-provisioned volume copy data in a node is unpinned.
986010	I	The thin-provisioned volume copy import has failed and the new volume is offline; either update the Storwize V7000 software to the required version or delete the volume.
986011	I	The thin-provisioned volume copy import is successful.
986020	W	A thin-provisioned volume copy space warning has occurred.
986030	I	A thin-provisioned volume copy repair has started.
986031	I	A thin-provisioned volume copy repair is successful.
986032	I	A thin-provisioned volume copy validation is started.
986033	I	A thin-provisioned volume copy validation is successful.
986034	I	The import of the compressed-virtual volume copy was successful.
986035	W	A compressed-virtual volume copy space warning has occurred.
986036	I	A compressed-virtual volume copy repair has started.
986037	I	A compressed-virtual volume copy repair is successful.
986038	I	A compressed-virtual volume copy has too many bad blocks.
986201	I	A medium error has been repaired for the mirrored copy.
986203	W	A mirror copy repair, using the validate option cannot complete.
986204	I	A mirror disk repair is complete and no differences are found.
986205	I	A mirror disk repair is complete and the differences are resolved.
986206	W	A mirror disk repair is complete and the differences are marked as medium errors.
986207	I	The mirror disk repair has been started.
986208	W	A mirror copy repair, using the set medium error option, cannot complete.

Table 107. Informational events (continued)

Event ID	Notification type	Description
986209	W	A mirror copy repair, using the resync option, cannot complete.
987102	W	Node coldstarted.
987103	W	A node power-off has been requested from the power switch.
987104	I	Additional Fibre Channel ports were connected.
987106	I	Additional ethernet ports connected
987107	I	Additional fibre channel IO ports connected
987301	W	The connection to a configured remote cluster (system) has been lost.
987400	W	The node unexpectedly lost power but has now been restored to the cluster (system).
988022	I	The rebuild for an array MDisk was started. Performance may be affected, wait for rebuild to complete.
988023	I	The rebuild for an array MDisk has finished.
988028	I	Array validation started.
988029	I	Array validation complete.
988100	W	An overnight maintenance procedure has failed to complete. Resolve any hardware and configuration problems that you are experiencing on the cluster (system). If the problem persists, contact your IBM service representative for assistance.
988300	W	An array MDisk is offline because it has too many missing members.
988304	I	A RAID array has started exchanging an array member.
988305	I	A RAID array has completed exchanging an array member.
988306	I	A RAID array needs resynchronization.
988307	I	A failed drive has been re-seated or replaced. The system has automatically configured the device.
988308	I	Distributed array MDisk rebuild started.
988309	I	Distributed array MDisk rebuild completed.
988310	I	Distributed array MDisk copyback started.
988311	I	Distributed array MDisk copyback completed.
988312	I	Distributed array MDisk initialization started.
988313	I	Distributed array MDisk initialization completed.
988314	I	Distributed array MDisk needs resynchronization.
989001	W	A storage pool space warning has occurred.

Error event IDs and error codes

Error codes describe a service procedure that must be followed. Each event ID that requires service has an associated error code.

Note: Service procedures that involve field-replaceable units (FRUs) do not apply to the IBM Spectrum Virtualize™ product, which is software based. For information about possible user actions that relate to FRU replacements, refer to your hardware manufacturer's documentation.

Error codes can be either notification type E (error) or notification type W (warning). Table 108 lists the event IDs that have corresponding error codes, and shows the error code, the notification type, and the condition for each event. For a list of informational events, which do not have associated error codes, see "Informational events" on page 265.

The 07*nnnn* event ID range refers to node errors that were logged by the system. The last 3 digits represent the error that was reported by the node. You can find these codes in the list of error codes at the end of this topic.

Table 108. Error event IDs and error codes

Event ID	Notification type	Condition	Error code
009020	E	A system recovery has run. All configuration commands are blocked.	1001
009040	E	The error event log is full.	1002
009052	W	The following causes are possible: • The node is missing. • The node is no longer a functional member of the system.	1196
009053	E	A node has been missing for 30 minutes.	1195
009054	W	Node has been shut down.	1707
009100	W	The software install process has failed.	2010
009101	W	Software install package cannot be delivered to all nodes.	2010
009110		Software install process stalled due to lack of redundancy	2010
009115		Software downgrade process stalled due to lack of redundancy	2008
009150	W	Unable to connect to the SMTP (email) server.	2600
009151	W	Unable to send mail through the SMTP (email) server.	2601
009170	W	Remote Copy feature capacity is not set.	3030
009171	W	The FlashCopy feature capacity is not set.	3031
009172	W	The Virtualization feature has exceeded the amount that is licensed.	3032
009173	W	The FlashCopy feature has exceeded the amount that is licensed.	3032
009174	W	Remote Copy feature license limit exceeded.	3032
009175	W	Thin-provisioned volume usage not licensed.	3033
009176	W	The value set for the virtualization feature capacity is not valid.	3029
009177	E	A physical disk FlashCopy feature license is required.	3035

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
009178	E	A physical disk Metro Mirror and Global Mirror feature license is required.	3036
009179	E	A virtualization feature license is required.	3025
009180	E	Automatic recovery of offline node failed.	1194
009181	W	Unable to send email to any of the configured email servers.	3081
009182	W	The external virtualization feature license limit was exceeded.	3032
009183	W	Unable to connect to LDAP server.	2251
009184	W	The LDAP configuration is not valid.	2250
009185	E	The limit for the compression feature license was exceeded.	3032
009186	E	The limit for the compression feature license was exceeded.	3032
009187	E	Unable to connect to LDAP server that has been automatically configured.	2256
009188	E	Invalid LDAP configuration for automatically configured server.	2255
009189	W	A licensable feature's trial-timer has reached 0. The feature has now been deactivated.	3082
009190	W	A trial of a licensable feature will expire in 5 days.	3083
009191	W	A trial of a licensable feature will expire in 10 days.	3084
009192	W	A trial of a licensable feature will expire in 15 days.	3085
009193	W	A trial of a licensable feature will expire in 45 days.	3086
009194	W	Easy Tier feature license limit exceeded.	3032
009195	W	FlashCopy feature license limit exceeded.	3032
009196	W	External virtualization feature license limit exceeded.	3032
009197	W	Remote copy feature license limit exceeded.	3032
009198	W	System update completion is required.	2050
009199	W	System update completion has stalled.	2012
009200	W	Encryption feature license limit exceeded	3032
009201	W	The quorum application is out of date and needs to be redeployed.	3123
009202	W	System SSL certificate will expire within the next 30 days.	3130
009203	W	System SSL certificate has expired.	2258
009205	W	No active quorum device found on this cluster.	3124

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
010002	E	The node ran out of base event sources. As a result, the node has stopped and exited the system.	2030
010003	W	The number of device logins has reduced.	1630
010004	W	Device excluded due to excessive errors on all Managed Disks	1640
010006	E	Access beyond end of disk, or Managed Disk missing.	2030
010008	E	The block size is invalid, the capacity or LUN identity has changed during the managed disk initialization.	1660
010010	E	The managed disk is excluded because of excessive errors.	1310
010011	E	The remote port is excluded for a managed disk and node.	1220
010012	E	The local port is excluded.	1210
010013	E	The login is excluded.	1230
010014	E	The local port is excluded.	1211
010015	E	Timeout due to non-responsive device	1340
010016	E	Timeout due to lost command	1340
010017	E	A timeout has occurred as a result of excessive processing time.	1340
010018	E	An error recovery procedure has occurred.	1370
010019	E	A managed disk is reporting excessive errors.	1310
010020	E	The managed disk error count threshold has exceeded.	1310
010021	W	There are too many devices presented to the system.	1200
010022	W	There are too many managed disks presented to the system.	1200
010023	W	There are too many LUNs presented to a node.	1200
010024	W	There are too many drives presented to a system.	1200
010025	W	A disk I/O medium error has occurred.	1320
010026	W	A suitable MDisk or drive for use as a quorum disk was not found.	1330
010027	W	The quorum disk is not available.	1335
010028	W	A controller configuration is not supported.	1625
010029	E	A login transport fault has occurred.	1360

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
010030	E	A managed disk error recovery procedure (ERP) has occurred. The node or controller reported the following: • Sense • Key • Code • Qualifier	1370
010031	E	One or more MDisks on a controller are degraded.	1623
010032	W	The controller configuration limits failover.	1625
010033	E	The controller configuration uses the RDAC mode; this is not supported.	1624
010034	W	Persistent unsupported controller configuration.	1695
010035	W	Controller has quorum disabled, but quorum disk is configured	1570
010040	E	The controller system device is only connected to the node through a single initiator port.	1627
010041	E	The controller system device is only connected to the node through a single target port.	1627
010042	E	The controller system device is only connected to the nodes through a single target port.	1627
010043	E	The controller system device is only connected to the nodes through half of the expected target ports.	1627
010044	E	The controller system device has disconnected all target ports to the nodes.	1627
010045	W	Number of Device paths from the controller site allowed accessible nodes has reduced	1630
010050	W	Solid state drive failed and rebuild is required.	1201
010051		A Solid state drive is missing from the configuration	1202
010052	E	Solid state drive offline as a result of a drive hardware error.	1205
010053	E	Solid state drive reporting PFA errors.	1215
010054	E	Solid state drive reporting too many errors.	1215
010055	W	An unrecognized SAS device.	1665
010056	E	SAS error counts exceeded the warning thresholds.	1216
010057	E	SAS errors exceeded critical thresholds.	1216

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
010058	E	The drive initialization failed because of an unknown block size or a block size that is not valid; an unknown capacity or a capacity that is not valid; or was not able to set the required mode pages.	1661
010059	E	Solid state drive offline due to excessive errors.	1311
010060	E	Solid state drive exceeded warning temperature threshold.	1217
010061	E	Solid state drive exceeded offline temperature threshold.	1218
010062	E	A drive exceeded the warning temperature threshold.	1217
010063	W	Drive medium error.	1321
010066	W	Controller indicates that it does not support descriptor sense for LUNs that are greater than 2 TBs.	1625
010067	W	Too many enclosures were presented to a system.	1200
010068	E	Format of the solid state drive is corrupted.	1204
010069	E	Solid state drive has incorrect block size.	1204
010070	W	Too many controller target ports were presented to the system.	1200
010071	W	Too many target ports were presented to the system from a single controller.	1200
010072	E	The drive is offline as a result of a drive hardware error.	1680
010073	E	The drive is reporting predictive failure analysis (PFA) errors.	1680
010080	E	The drive is reporting too many errors.	1680
010081	E	The drive format is corrupted.	1206
010082	E	The block size for the drive was incorrect.	1206
010083	E	The drive is offline due to excessive errors.	1680
010084	E	The error counts for the SAS drive exceeded the warning thresholds.	1285
010085	W	The SAS device was not recognized.	1666
010086	W	The SAS enclosure was not recognized.	1666
010087	W	The SAS device was not able to be identified.	1666
010088	E	There were excessive medium errors on the drive.	1680
010089	E	There were excessive overall timeout errors on the drive.	1680
010090	E	There were excessive times when the drive stopped.	1680

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
010091	E	A drive failed validation testing.	1680
010092	E	Excessive medium errors on the solid state drive.	1215
010093	E	Excessive overall timeout errors on the solid state drive.	1204
010094	E	Login excluded.	1231
010095	E	Drive failed.	1687
010096	E	The drive initialization failed because of an unknown block size or a block size that is not valid; an unknown capacity or a capacity that is not valid; or was not able to set the required mode pages.	1680
010097	E	A drive is reporting excessive errors.	1685
010098	W	There are too many drives presented to a system.	1200
010100	W	Incorrect connection detected to a port.	1669
010101	E	Too many long IOs to drive.	1215
010102	E	A drive is reported as continuously slow with contributory factors.	1215
010103	E	Too many long IOs to drive (Mercury drives).	1215
010104	E	A drive is reported as continuously slow with contributory factors (Mercury drives).	1215
010105	W	Storage system connected to unsupported port	2080
010106	E	Drive reporting too many t10dif errors.	1680
010107	W	Encrypting MDisk is no longer encrypted	2580
010110	W	Drive firmware download canceled because of system changes.	3090
010111	W	Drive firmware download canceled because of a drive download problem.	3090
010117	W	A disk controller is not accessible from a node allowed to access the device by site policy	1627
010118	W	Too many drives attached to the system.	1179
010119	W	Drive data integrity error.	1322
010120	W	A member drive has been forced to turn off protection information support.	2035
010121	E	Drive exchange required.	1693
010123	W	Performance of external MDisk has changed.	2115
010124	W	iSCSI session excluded.	1230
010125	W	A Flash drive is expected to fail within six months due to limited write endurance.	"1215" on page 356
010126	W	A Flash drive with high write endurance usage rate.	"2560" on page 384

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
020001	E	There are too many medium errors on the MDisk.	1610
020002	E	A storage pool is offline.	1620
020003	W	There are insufficient virtual extents.	2030
020008	E	Storage optimization services disabled.	3023
029001	E	The MDisk has bad blocks.	1840
029002	W	The system failed to create a bad block because MDisk already has the maximum number of allowed bad blocks.	1226
029003	W	The system failed to create a bad block because the system already has the maximum number of allowed bad blocks.	1225
030000	W	FlashCopy prepare failed due to cache flush failure.	1900
030010	W	FlashCopy has been stopped due to the error indicated in the data.	1910
030020	W	Unrecovered FlashCopy mappings.	1895
045001	E	One or more power supply unit fans have failed.	1124
045002	E	A fan is operating outside the expected range.	1126
045003	E	There was a fan status communications failure.	1126
045004	W	The power supply unit is not installed.	1128
045005	W	The power supply unit has indicated an input power failure.	1138
045006	E	The power supply unit has indicated an output failure.	1126
045007	E	The power supply unit has failed.	1124
045008	E	Cannot communicate with the power supply unit.	1148
045009	E	The PSU type is not valid for this enclosure type.	1124
045010	E	The power supply unit type is unknown to this product.	1124
045011	E	PSU 11S serial number not valid.	1124
045012	W	The canister temperature is at the warning level.	1098
045013	W	The canister temperature is at the critical level.	1095
045014	E	A SAS cable was excluded due to internal errors.	1260
045015	E	A SAS cable was excluded because too many change events were caused.	1260
045016	E	A SAS cable was excluded.	1255

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
045017	E	A SAS cable is operating at a reduced speed.	1260
045018	E	A SAS cable was excluded because frames were dropped.	1260
045019	E	A SAS cable was excluded because the enclosure discovery timed out.	1260
045020	W	A SAS cable is not present.	1265
045021	W	A canister was removed from the system.	1036
045022	E	A canister has been in a degraded state for too long and cannot be recovered.	1034
045023	E	A canister is encountering communication problems.	1038
045024	E	The canister VPD is not valid.	1032
045025	E	The canister has experienced too many resets.	1032
045026	E	The drive slot is causing the network to be unstable.	1686
045027	E	The drive slot is not running at 6 Gbps.	1686
045028	E	The drive slot is dropping frames.	1686
045029	E	The drive is visible through only one SAS port.	1686
045031	E	The drive power control is not functional.	1008
045032	E	A drive is present in the slot but does not respond to any SCSI or SAS commands.	1686
045033	E	The drive slot contains a device that is not responding to queries.	1685
045034	E	The managed enclosure is not visible from any node canisters.	1042
045035	E	The electronics in the enclosure has failed.	1694
045036	E	The electronics in the enclosure has experienced a critical failure.	1008
045037	E	The SAS network has too many errors.	1048
045038	E	The SAS network has too many errors.	1048
045040	W	The firmware update for the enclosure component has failed.	3015
045041	W	More than one initiator port was detected on the same strand.	1005
045042	W	The order of the enclosures is different on each strand.	1005
045044	W	Multiple canisters are connected to a single canister port.	1005
045045	W	Canister one is connected to canister two.	1005
045046	W	An enclosure is connected to more than one I/O group.	1005

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
045047	W	A managed enclosure is connected to the wrong I/O group.	1005
045048	W	An enclosure is connected to more than one chain.	1005
045049	W	Too many canisters are connected to a strand.	1005
045050	W	The canister is connected to the wrong port.	1005
045051	E	A SAS cable is excluded because of single port active drives.	1260
045052	W	More than one canister was detected at the same hop count.	1005
045053	E	The node location is not able to be detected.	1031
045054	E	An enclosure display cannot be updated.	1694
045055	E	There is an enclosure battery fault.	1118
045056	E	An enclosure battery is missing.	1112
045057	E	An enclosure battery is nearing end of life.	1114
045058	E	An enclosure battery is at end of life.	1113
045060	E	Drive slot cannot be excluded.	1048
045062	W	An enclosure battery conditioning is required but not possible.	1131
045063	E	There was an enclosure battery communications error.	1116
045064	W	A SAS port is active, but no enclosures can be detected.	1005
045065	E	There is a connectivity problem between a canister and an enclosure.	1036
045066	E	The FRU identity of the enclosure is not valid.	1008
045067	W	A new enclosure FRU was detected and needs to be configured.	1041
045068	E	The internal device on a node canister was excluded because of too many change events.	1034
045069	E	The internal connector on the node canister was excluded as the cause of single ported drives.	1034
045070	W	The canister temperature sensor cannot be read.	1034
045071	W	The enclosure contains both a node canister and an expansion canister.	1037
045072	E	The discovery failed to complete.	1048
045073	E	The VPD for the enclosure cannot be read.	1048
045074	W	The PSU temperature is at the warning level.	1098
045075	W	The PSU temperature is at the critical level.	1095

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
045076	E	Error reading PSU temperature sensor.	1124
045077	E	Attempts to exclude connector have failed.	1260
045080	E	There are too many self-initiated resets on the enclosure.	1048
045081	E	The PSU has attempted too many internal recovery actions.	1124
045082	E	The slots are powered off.	1048
045083	W	The enclosure configuration is not supported.	1005
045084	W	There are too many SDE enclosures on strand.	1005
045085	W	The SDE is detected in a chassis that is not valid.	1005
045086	E	Node canister internal connector failed because of dropped frames.	1034
045102	W	SAS cable is not working at full capacity	1260
045103	E	An attempt to automatically configure a reseated or replaced drive has failed.	1686
045104	W	Drives are single ported due to a spare node	3200
045105	E	Enclosure secondary expander module has failed	"1267" on page 357
045106	E	Enclosure secondary expander module FRU identity is not valid	"1266" on page 357
045107	E	Enclosure secondary expander module temperature sensor cannot be read	"1267" on page 357
045108	E	Enclosure secondary expander module temperature has passed warning threshold	"1098" on page 343
045109	E	Enclosure secondary expander module temperature has passed critical threshold	"1095" on page 341
045110	E	Enclosure display panel is not installed	"1268" on page 358
045111	E	Enclosure display panel temperature sensor cannot be read	"1268" on page 358
045112	E	Enclosure display panel temperature has passed warning threshold	"1098" on page 343
045113	E	Enclosure display panel temperature has passed critical threshold	"1095" on page 341
045114	E	Enclosure secondary expander module connector excluded due to too many change events	"1267" on page 357
045119	E	Enclosure display panel VPD cannot be read	"1268" on page 358
045120	E	Enclosure secondary expander module is missing	"1267" on page 357
045121	E	Enclosure secondary expander module connector excluded due to dropped frames	"1267" on page 357

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
045122	E	Enclosure secondary expander module connector is excluded and cannot be unexcluded	"1267" on page 357
045123	E	Enclosure secondary expander module connectors excluded as the cause of single ported drives	"1267" on page 357
045124	E	Enclosure secondary expander module leaf expander connector excluded as the cause of single ported drives	"1267" on page 357
050001	W	The Metro Mirror or Global Mirror relationship cannot be recovered.	1700
050002	W	A Metro Mirror or Global Mirror relationship or consistency group exists within a system, but its partnership has been deleted.	3080
050010	W	A Global Mirror relationship has stopped because of a persistent I/O error.	1920
050011	W	A remote copy has stopped because of a persistent I/O error.	1915
050020	W	Remote Copy relationship or consistency groups lost synchronization.	1720
050030	W	There are too many system partnerships. The number of partnerships has been reduced.	1710
050031	W	There are too many system partnerships. The system has been excluded.	1710
050040	W	Background copy process for the remote copy was blocked.	1960
050041	W	Partner cluster IP address unreachable	2021
050042	W	Cannot authenticate with partner cluster.	2022
050043	W	Unexpected cluster ID for partner cluster	2023
050050	E	The Global Mirror secondary volume is offline. The relationship has pinned hardened write data for this volume.	1925
050060	E	The Global Mirror secondary volume is offline due to missing I/O group partner node. The relationship has pinned hardened write data for this volume but the node containing the required data is currently offline.	1730
050070	E	Global Mirror performance is likely to be impacted. A large amount of pinned data for the offline volumes has reduced the resource available to the global mirror secondary disks.	1925
050080	W	HyperSwap volume has lost synchronization between sites.	1940
050081	W	HyperSwap consistency group has lost synchronization between sites.	1940

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
050082	E	Compression has stopped unexpectedly	3131
060001	W	A thin-provisioned volume copy is offline because of insufficient space.	1865
060002	E	A thin-provisioned volume copy is offline because of corrupt metadata.	1862
060003	E	A thin-provisioned volume copy is offline because of a failed repair.	1860
060004	W	A compressed volume copy is offline because of insufficient space.	1865
060005	E	A compressed volume copy is offline because of corrupt metadata.	1862
060006	E	A compressed volume copy is offline because of a failed repair.	1860
060007	E	A compressed volume copy has bad blocks.	1850
062001	W	System is unable to mirror medium error.	1950
062002	E	Mirrored volume is offline because it cannot synchronize data.	1870
062003	W	Repair of a mirrored volume stopped because of difference.	1600
064001	W	A host port has more than four logins to a node	2016
070000	E	Unrecognized node error.	1083
070510	E	Detected memory size does not match the expected memory size.	1022
070511	E	DIMMs are incorrectly installed.	1009
070517	E	The WWNN that is stored on the service controller and the WWNN that is stored on the drive do not match.	1192
070521	E	Unable to detect any Fibre Channel adapter.	1016
070522	E	The system board processor has failed.	1020
070523	E	The internal disk file system of the node is damaged.	1187
070524	E	Unable to update BIOS settings.	1027
070525	E	Unable to update the service processor firmware for the system board.	1020
070528	E	The ambient temperature is too high while the system is starting.	1182
070534	E	System board fault	1026
070536	E	A system board device breached critical temperature threshold.	1084
070538	E	A PCI Riser breached critical temperature threshold.	1085
070541	E	Multiple hardware failures	1184
070542	E	A processor has failed.	1024

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
070543	E	No usable persistent data could be found on the boot drives.	1035
070544	E	The boot drives do not belong in this node.	1035
070545	E	Boot drive and system board mismatch.	1035
070547	E	Pluggable TPM is missing or broken	1051
070550	W	Cannot form system due to lack of resources.	1192
070551	W	Cannot form cluster due to lack of cluster resources, overridequorum possible	1192
070556	E	Duplicate WWNN detected on the SAN.	1192
070558	E	A node is unable to communicate with other nodes.	1192
070562	E	The node hardware does not meet minimum requirements.	1183
070564	E	Too many software failures.	1188
070572	E	Battery protection temporarily unavailable; both batteries are expected to be available soon.	1473
070573	E	Node software inconsistent	1192
070574	E	The node software is damaged.	1187
070576	E	The system data cannot be read.	1030
070578	E	The system data was not saved when power was lost.	1194
070580	E	Unable to read the service controller ID.	1044
070690	W	Node held in service state.	1189
070700	W	Fibre Channel adapter missing	1045
070701	E	Fibre Channel adapter failed	1046
070702	E	Fibre Channel adapter PCI error	1046
070703	E	Fibre Channel adapter degraded	1046
070704	W	Fewer Fibre Channel ports operational.	1060
070705	W	Fewer Fibre Channel I/O ports operational.	1450
070706	W	Fibre Channel clustered system path failure.	1550
070711	E	SAS adapter failed	1046
070712	E	SAS adapter PCI error	1046
070713	E	SAS adapter degraded	1046
070715	W	Fewer SAS ports operational	1046
070717	W	SAS ports degraded	1046
070718	W	SASA port has unsupported SAS device	1046
070720	W	Ethernet adapter missing	1045
070721	E	Ethernet adapter failed	1046
070722	E	Ethernet adapter PCI error	1046
070723	E	Ethernet adapter degraded	1046

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
070724	W	Fewer Ethernet ports	1046
070730		Bus adapter missing	1192
070731		Bus adapter failed	1192
070732		Bus adapter PCI error	1192
070733		Bus adapter degraded	1192
070734		Fewer bus ports operational	1006
070747	W	Technician connected.	747
070760	E	Voltage fault	1110
070761	E	Voltage high	1100
070762	E	Voltage low	1105
070765	E	Fan error	1089
070768	W	Ambient temperature warning	1094
070769	W	CPU temperature warning	1093
070770	W	Shutdown temperature reached	1092
070830	W	Encryption key required	1328
070831	W	Encryption key invalid	2555
070832	W	Encryption key not found	2555
070833	W	USB device (such as hub) unsupported	2555
070836	W	Encryption key required	1328
070842	W	Fibre Channel IO port mapping failed	1059
070860	W	Fibre-channel network fabric is too big.	1800
071500	W	Incorrect enclosure	1021
071501	E	Incorrect slot	1192
071502	E	No enclosure id and cannot get status from partner	1192
071503	E	Incorrect enclosure type	1192
071504	E	No enclosure id & partner matches	1192
071505	E	No enclosure id and partner has cluster data does not match	1192
071506	E	No enclosure id and no cluster state on partner	1192
071507	E	No enclosure id and no cluster state	1192
071508	W	Cluster id different between enclosure and node	1023
071509	E	Cannot read enclosure identity	1036
071510	E	The detected memory size does not match the expected memory size	1032
071522	E	The system board processor has failed.	1034
071523	E	Internal disk file system is damaged	1187
071524	E	Unable to update BIOS settings	1034

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
071525	E	Unable to update system board service processor firmware	1034
071528	W	Ambient temperature too high while system starting	1092
071535	E	Canister internal PCIe switch failed	1034
071541	E	Multiple hardware failures	1184
071547	E	Pluggable TPM is missing or broken	1051
071550	W	Cannot form cluster due to lack of cluster resources	1192
071551	W	Cannot form cluster due to lack of cluster resources, overridequorum possible	1192
071556	W	Duplicate WWNN detected on SAN	1133
071562	E	The node's hardware configuration does not meet minimum requirements	1034
071564	W	Too many software failures	1188
071565	E	The node's internal drive is failing.	1032
071569	E	CPU over temp.	1032
071573	E	Node software inconsistent	1187
071574	E	Node software is damaged	1187
071576	E	Cluster state and configuration data cannot be read	1032
071578	E	State data was not saved on power loss	1194
071671	W	The available battery charge is not enough to allow the node canister to start . Two batteries are charging.	1176
071672	W	The available battery charge is not enough to allow the node canister to start . One battery is charging.	1176
071673	E	The available battery charge is not enough to allow the node canister to start . No batteries are charging.	1004
071690	W	Node held in service state	1189
071700	W	Fibre Channel adapter missing	1032
071701	E	Fibre Channel adapter failed	1032
071702	E	Fibre Channel adapter PCI error	1034
071703	E	Fibre Channel adapter degraded	1034
071704	W	Fewer Fibre Channel ports operational.	1061
071705	W	Fewer Fibre Channel I/O ports operational.	1450
071706	W	Fibre Channel clustered system path failure.	1550
071710	W	SAS adapter missing	1032
071711	E	SAS adapter failed	1032
071712	E	SAS adapter PCI error	1034
071713	E	SAS adapter degraded	1034

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
071715	W	Fewer SAS ports operational	1034
071717	W	SAS ports degraded	1034
071718	W	SASA port has unsupported SAS device	1034
071720	W	Ethernet adapter missing	1032
071721	E	Ethernet adapter failed	1032
071722	E	Ethernet adapter PCI error	1034
071723	E	Ethernet adapter degraded	1034
071724	W	Fewer Ethernet ports	1401
071730	W	Bus adapter missing	1032
071731	E	Bus adapter failed	1032
071732	E	Bus adapter PCI error	1034
071733	E	Bus adapter degraded	1034
071734	W	Fewer bus ports operational	1006
071746	W	Technician port connection is not valid.	3024
071747	W	Technician connected.	747
071768	W	Ambient temperature warning	1094
071769	W	CPU temperature warning	1093
071781	W	Battery cold	1156
071782	W	Battery hot	1157
071786	E	Battery VPD checksum	1154
071820	W	Node canister has the incorrect model for the enclosure.	3020
071830	W	Encryption key required	1328
071831	W	Encryption key invalid	2555
071832	W	Encryption key not found	2555
071833	W	USB device (such as hub) unsupported	2555
071836	W	Encryption key required	1328
071840	W	Detected hardware is not a valid configuration.	1198
071841	W	Detected hardware needs activation.	1199
071850	W	Canister battery is nearing end of life	1159
072001	E	System board has more or less processors detected.	1020
072500	W	Incorrect enclosure	1021
072501	E	Incorrect slot	1192
072502	E	No enclosure id and cannot get status from partner	1192
072503	E	Incorrect enclosure type	1192
072504	E	No enclosure id & partner matches	1192
072505	E	No enclosure id and partner has cluster data that does not match	1192

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
072506	E	No enclosure id and no cluster state on partner	1192
072507	E	No enclosure id and no cluster state	1192
072508	W	Cluster id different between enclosure and node	1023
072509	E	Cannot read enclosure identity	1036
072510	E	The detected memory size does not match the expected memory size	1032
072522	E	The system board processor has failed	1033
072523	E	Internal disk file system is damaged	1187
072525	E	Unable to update system board service processor firmware	1034
072535	E	Canister internal PCIe switch failed	1192
072541	E	Multiple hardware failures	1184
072550	W	Cannot form cluster due to lack of cluster resources	1192
072551	W	Cannot form cluster due to lack of cluster resources, overridequorum possible	1192
072556	E	Duplicate WWNN detected on SAN	1133
072562	E	The node's hardware configuration does not meet minimum requirements	1034
072564	E	Too many software failures	1188
072565	E	The node's internal drive is failing.	1032
072569	E	CPU over temp.	1032
072573	E	Node software inconsistent	1187
072574	E	Node software is damaged	1187
072576	E	Cluster state and configuration data cannot be read	1032
072578	E	State data was not saved on power loss	1194
072650	W	The canister battery is not supported.	1149
072651	W	The canister battery is missing.	1153
072652	E	The canister battery has failed.	1154
072655	E	Canister battery communications error	1158
072656	W	Canister battery has insufficient charge to support a firehose dump	1197
072690	W	Node held in service state	1189
072700	W	Fibre Channel adapter missing	1045
072701	E	Fibre Channel adapter failed	1046
072702	E	Fibre Channel adapter PCI error	1046
072703	E	Fibre Channel adapter degraded	1046
072704	W	Fewer Fibre Channel ports operational.	1062
072705	W	Fewer Fibre Channel I/O ports operational.	1450

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
072706	W	Fibre Channel clustered system path failure.	1550
072710	W	SAS adapter missing	1045
072711	E	SAS adapter failed	1046
072712	E	SAS adapter PCI error	1046
072713	E	SAS adapter degraded	1046
072715	W	Fewer SAS ports operational	1046
072717	W	SAS ports degraded	1046
072718	W	SASA port has unsupported SAS device	1046
072720	W	Ethernet adapter missing	1045
072721	E	Ethernet adapter failed	1046
072722	E	Ethernet adapter PCI error	1046
072723	E	Ethernet adapter degraded	1046
072724	W	Fewer Ethernet ports	1402
072730	W	Bus adapter missing	1032
072731	E	Bus adapter failed	1032
072732	E	Bus adapter PCI error	1032
072733	E	Bus adapter degraded	1032
072734	W	Fewer bus ports operational	1006
072766	E	CMOS error	1670
072840	W	A hardware change was made that is not supported by software. User action is required to repair the hardware or update the software.	1198
072841	W	A supported hardware change was made to this node. User action is required to activate the new hardware.	1199
072850	W	Canister battery is nearing end of life	1159
072860	W	Fibre-channel network fabric is too big.	1800
072900	E	There was a PCIe link failure between canisters.	1006
072901	E	The PCIe link is degraded between canisters.	1052
072902	E	Inter-canister PCIe link recovered.	1006
072911	E	The PCIe link for the CPU is degraded.	1034
073001	E	More or lesss Fibre Channel adapters detected.	1010
073002	E	Fibre Channel adapter is faulty.	1050
073003	W	The Fibre Channel ports are not operational.	1060
073004	E	Fibre Channel adapter detected PCI bus error.	1012
073005	E	System path has a failure.	1550

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
073006	W	The SAN is not correctly zoned. As a result, more than 512 ports on the SAN have logged into one Storwize V7000 port.	1800
073007	W	There are fewer Fibre Channel ports operational than are configured.	1061
073305	W	Fibre Channel speed has changed.	1065
073310	E	Duplicate Fibre Channel frame is detected.	1203
073402	E	The Fibre Channel adapter has a failure.	1032
073404	E	Fibre Channel adapter has detected PCI bus error.	1032
073500	W	Incorrect enclosure	1021
073501	E	The canister position is not correct.	1192
073502	E	No enclosure id and cannot get status from partner	1192
073503	E	The enclosure type is not correct.	1192
073504	E	The enclosure identity and partner does not match.	1192
073505	E	No enclosure id and partner has cluster data that does not match	1192
073506	E	The state and enclosure identity cannot be detected on the partner.	1192
073507	E	The enclosure identity and node state cannot be detected.	1192
073508	W	The system identity is different on the enclosure and the node.	1023
073509	E	The enclosure identity cannot be read.	1036
073510	E	Detected memory size does not match the expected memory size.	1032
073512	E	Enclosure VPD is inconsistent.	1008
073522	E	The system board service processor has failed.	1034
073523	W	The internal disk file system of the node is damaged.	1187
073524	E	Unable to update BIOS settings	1034
073525	E	Unable to update the service processor firmware of the system board.	1034
073528	E	Ambient temperature is too high during system startup.	1098
073535	E	The internal PCIe switch of the node canister failed.	1034
073541	E	Multiple hardware failures	1184
073550	W	System cannot be created because of lack of resources.	1192
073551	W	Cannot form cluster due to lack of cluster resources, overridequorum possible	1192

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
073556	W	Duplicate WWNN are detected on the SAN.	1133
073562	E	The node hardware does not meet the minimum requirements.	1034
073564	W	Too many software failures	1188
073565	E	The internal drive of the node is failing.	1032
073569	E	CPU over temp.	1032
073573	E	Node software inconsistent (not raised on this platform).	1187
073574	E	The system data cannot be read.	1187
073576	E	Cluster state and configuration data cannot be read	1032
073578	E	The system data was not saved when power was lost.	1194
073650	W	The canister battery is not supported.	1149
073651	E	The canister battery is missing.	1153
073652	E	The canister battery has failed.	1154
073653	E	The canister battery's temperature is too low.	1156
073654	E	The canister battery's temperature is too high.	1157
073655	E	The canister has a battery communications fault.	1158
073656	E	The canister battery has insufficient charge.	1184
073690	W	Node held in service state	1189
073700	E	The Fibre Channel adapter is missing.	1045
073701	E	The Fibre Channel adapter failed.	1046
073702	E	The Fibre Channel adapter has a PCI error.	1046
073703	E	The Fibre Channel adapter is degraded.	1045
073704	W	Fewer Fibre Channel ports are operational.	1061
073705	W	Fewer Fibre Channel I/O ports are operational.	1450
073710	E	The SAS adapter is missing.	1045
073711	E	The SAS adapter has failed.	1046
073712	E	The SAS adapter has a PCI error.	1046
073713	E	The SAS adapter is degraded.	1046
073715	W	Fewer SAS ports operational	1046
073717	W	SAS ports degraded	1046
073718	W	SASA port has unsupported SAS device	1669
073720	E	The Ethernet adapter is missing.	1045
073721	E	The Ethernet adapter has failed.	1046
073722	E	Ethernet adapter PCI error.	1046
073723	E	Ethernet adapter degraded.	1046

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
073724	W	Fewer Ethernet ports are operational.	1401
073730	E	The bus adapter is missing.	1032
073731	E	The bus adapter has failed.	1032
073732	E	The bus adapter has a PCI error.	1032
073733	E	The bus adapter is degraded.	1032
073734	W	The inter-canister PCIe has a link failure.	1006
073748	W	Technician port remains enabled.	748
073766	E	CMOS error	1670
073768	W	Ambient temperature warning.	1094
073769	W	CPU temperature warning.	1093
073820	W	The node canister has detected that it has a hardware type that is not compatible with the control enclosure MTM	3020
073830	W	Encryption key required	1328
073831	W	Encryption key invalid	2555
073832	W	Encryption key not found	2555
073833	W	USB device (such as hub) unsupported	2555
073836	W	Encryption key required	1328
073840	E	Detected hardware is not a valid configuration.	1198
073841	E	Detected hardware needs activation.	1199
073850	W	Canister battery is nearing end of life	1159
073860	W	The fabric is too large.	1800
074001	W	System is unable to determine VPD for a FRU.	2040
074002	E	The node warm started after a software error.	2030
074003	W	A connection to a configured remote system has been lost because of a connectivity problem.	1715
074004	W	A connection to a configured remote system has been lost because of too many minor errors.	1716
074500	W	Incorrect enclosure	1021
074501	E	Incorrect slot	1192
074502	E	No enclosure id and cannot get status from partner	1192
074503	E	Incorrect enclosure type	1192
074504	E	No enclosure id & partner matches	1192
074505	E	No enclosure id and partner has cluster data that does not match	1192
074506	E	No enclosure id and no cluster state on partner	1192

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
074507	E	No enclosure id and no cluster state	1192
074508	W	Cluster id different between enclosure and node	1023
074509	E	Cannot read enclosure identity	1043
074510	E	The detected memory size does not match the expected memory size	1039
074512	E	Enclosure VPD is inconsistent	1029
074521	E	Unable to detect any fibre-channel adapter	1192
074522	E	The system board processor has failed	1088
074523	E	Internal disk file system is damaged	1187
074524	E	Unable to update BIOS settings	1034
074525	E	Unable to update system board service processor firmware	1192
074528	W	Ambient temperature too high while system starting	1087
074534	E	System board fault	1039
074535	E	Canister internal PCIe switch failed	1034
074536	E	A device on the system board is too hot	1192
074538	E	PCI Riser too hot	1192
074541	E	Multiple hardware failures	1184
074550	W	Cannot form cluster due to lack of cluster resources	1192
074551	W	Cannot form cluster due to lack of cluster resources, overridequorum possible	1192
074556	W	Duplicate WWNN detected on SAN	1133
074562	E	The node's hardware configuration does not meet minimum requirements	1034
074564	E	Too many software failures	1188
074565	E	The node's internal drive is failing.	1039
074569	E	CPU over temp.	1192
074573	E	Node software inconsistent	1192
074574	E	Node software is damaged	1187
074576	E	Cluster state and configuration data cannot be read	1039
074578	E	State data was not saved on power loss	1194
074650	W	The canister battery is not supported.	1192
074651	W	The canister battery is missing.	1192
074652	E	The canister battery has failed.	1192
074653	W	The canister battery is below minimum operating temperature.	1192
074654	W	The canister battery is above maximum operating temperature.	1192

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
074655	E	Canister battery communications error	1192
074656	W	Canister battery has insufficient charge to support a firehose dump	1192
074657	E	Not enough battery to support graceful shutdown.	1111
074690	W	Node held in service state	1189
074710	W	SAS adapter missing	1192
074711	E	SAS adapter failed	1192
074712	E	SAS adapter PCI error	1192
074713	E	SAS adapter degraded	1192
074715	W	Fewer SAS ports operational	1192
074717	W	SAS ports degraded	1192
074718	W	SASA port has unsupported SAS device	1192
074720	W	Ethernet adapter missing	1039
074721	E	Ethernet adapter failed	1039
074722	E	Ethernet adapter PCI error	1034
074723	E	Ethernet adapter degraded	1034
074724	W	Fewer Ethernet ports	1401
074730	W	Bus adapter missing	1039
074731	E	Bus adapter failed	1039
074732	E	Bus adapter PCI error	1034
074733	E	Bus adapter degraded	1034
074734	W	Fewer bus ports operational	1007
074768	W	Ambient temperature warning	1099
074830	W	Encryption key required	1328
074831	W	Encryption key invalid	2555
074832	W	Encryption key not found	2555
074833	W	USB device (such as hub) unsupported	2555
074840	W	A hardware change was made that is not supported by software. User action is required to repair the hardware or update the software.	1198
074841	W	A supported hardware change was made to this node. User action is required to activate the new hardware.	1199
076001	E	The internal disk for a node has failed.	1030
076002	E	The hard disk is full and cannot capture any more output.	2030
076401	E	One of the two power supply units in the node has failed.	1096
076402	E	One of the two power supply units in the node cannot be detected.	1096

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
076403	E	One of the two power supply units in the node is without power.	1097
076502	E	The PCIe lanes on a high-speed SAS adapter are degraded.	1121
076503	E	A PCI bus error occurred on a high-speed SAS adapter.	1121
076504	E	A high-speed SAS adapter requires a PCI bus reset.	1122
076505	E	The SAS adapter has an internal fault.	1121
076511	E	A high-speed SAS controller is missing.	1032
076512	E	The PCIe lanes on a high-speed SAS adapter are degraded.	1032
076513	E	A PCI bus error occurred on a high-speed SAS adapter.	1032
076514	E	A high-speed SAS adapter requires a PCI bus reset.	1034
077001	E	The node service processor indicated Fan 1 failure.	1070
077002	E	The node service processor indicated Fan 2 failure.	1070
077003	E	The node service processor indicated Fan 3 failure.	1070
077004	E	The node service processor indicated Fan 4 failure.	1070
077005	E	The node service processor indicated Fan 5 failure.	1071
077011	E	The node service processor indicated that the ambient temperature threshold was exceeded.	1075
077012	E	The node service processor indicated temperature warning threshold was exceeded.	1076
077013	E	The node service processor indicated soft or hard shutdown temperature threshold was exceeded.	1077
077021	E	The node service processor indicated 12-volt supply has exceeded thresholds.	1080
077022	E	The node service processor indicated 5-volt supply has exceeded thresholds.	1080
077023	E	The node service processor indicated 3.3-volt supply has exceeded thresholds.	1080
077024	E	The node service processor indicated 2.5-volt supply has exceeded thresholds.	1081
077025	E	The node service processor indicated 1.5-volt supply has exceeded thresholds.	1081
077026	E	The node service processor indicated 1.25-volt supply has exceeded thresholds.	1081

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
077027	E	The node service processor indicated that the CPU supply has exceeded thresholds.	1081
079500	W	The limit on the number of system secure shell (SSH) sessions has been reached.	2500
079501	W	Unable to access the Network Time Protocol (NTP) network time server.	2700
079503	W	Unable to connect to NTP server that has been automatically configured.	2702
079504	W	Hardware configurations of nodes differ in an I/O group.	1470
079505	W	Stretch cluster reconfiguration is required to restore a dual site configuration	1178
081002	E	An Ethernet port failure has occurred.	1401
082001	E	A server error has occurred.	2100
082002	W	Service failure has occured.	2100
083001	E	System failed to communicate with UPS.	1145
083002	E	UPS output loading was unexpectedly high.	1165
083003	E	Battery has reached end of life.	1190
083004	E	UPS battery has a fault.	1180
083005	E	UPS electronics has a fault.	1170
083006	E	UPS frame has a fault.	1175
083007	E	UPS is overcurrent.	1160
083008	E	UPS has a fault but no specific FRU is identified.	1185
083009	E	The UPS has detected an input power fault.	1140
083010	E	UPS has a cabling error.	1150
083011	E	UPS ambient temperature threshold has exceeded.	1135
083012	E	UPS ambient temperature is high.	3000
083013	E	UPS crossed-cable test is bypassed because of an internal UPS software error.	3010
084000	W	An array MDisk has deconfigured members and has lost redundancy.	1689
084050	W	An array MDisk is expected to fail within six months due to limited write endurance of member drives	"3060" on page 388
084100	E	An array MDisk is corrupt because of lost metadata.	1240
084200	W	Array MDisk has taken a spare member that does not match array goals.	1692
084201	W	An array has members that are located in a different I/O group.	1688
084300	W	An array MDisk is no longer protected by an appropriate number of suitable spares.	1690

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
084301	W	No spare protection exists for one or more array MDisks.	1690
084302	W	Distributed array MDisk has fewer rebuild areas available than threshold.	1690
084400	W	A background scrub process has found an inconsistency between data and parity on the array.	1691
084420	W	Array MDisk has been forced to disable hardware data integrity checking on member drives.	2035
084500	E	An array MDisk is offline. The metadata for the inflight writes is on a missing node.	1243
084600	E	An array MDisk is offline. Metadata on the missing node contains needed state information.	1243
084700	W	Array response time too high.	1750
084701	W	Distributed array MDisk member slow write count threshold exceeded.	1750
084800	E	Distributed array MDisk offline due to I/O timeout.	1340
085047	W	Battery reconditioning required but not possible	1131
085052	E	Interface card has degraded PCI link	1039
085055	W	External FC data link degraded	1064
085056	W	External IB data link degraded	1064
085063	E	Canister is missing an interface card	1045
085091	W	External iSCSI port not operational	1403
085092	W	Too many ISCSI host logins	1803
085118	W	System update halted	2010
085160	W	Check the air filter	1820
085161	E	Array data compromised	1048
085198	W	Too many enclosures visible on fabric	1807
085199	W	Enclosure visible on fabric managed by another system	1706
085200	W	Cabling error. Internal cabling connectivity has changed.	1440
085201	W	Enclosure connectivity undetermined. Connectivity to an enclosure can no longer be determined	1440
085202	W	Minimal enclosure connectivity not met.	1705
085203	W	Config node cannot communicate with canister.	1034
085204	W	Managed enclosure is not visible from config node.	1042
085205	W	Canister internal error.	1705

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
085221	I	Successful write to USB Flash Drive	n/a
085222	W	Write failed to USB Flash Drive	1790
086001	E	Encryption key unavailable	1739
086002	W	Encryption key on USB flash drive removed	2550
086003	W	Write to USB Flash Drive failure	1790
086004	I	Write to USB Flash Drive successful	n/a
086005	W	Encryption not committed	1780
086006	E	Key Server reported KMIP error	"1785" on page 373
086007	E	Key Server reported vendor information error	"1785" on page 373
086008	E	Failed to connect to Key Server	"1785" on page 373
087001	E	Cloud gateway service restarted	"2031" on page 380
087002	E	Cloud gateway service restarted too often	"1404" on page 362
087003	W	Cloud account SSL certificate will expire within the next 30 days	"3140" on page 391
087004	W	Cloud account not available, cannot resolve hostname	"1580" on page 364
087005	W	Cloud account not available, cannot contact cloud provider	"2310" on page 383
087006	W	Cloud account not available, cannot communicate with cloud provider	"2320" on page 383
087007	W	Cloud account not available, no matching CA certificate	"2300" on page 382
087008	W	Cloud account not available, no matching CA certificate	"2300" on page 382
087009	W	Cloud account not available, cannot establish secure connection with cloud provider	"3100" on page 389
087010	W	Cloud account not available, cannot authenticate with cloud provider	"2330" on page 383
087011	W	Cloud account not available, cannot obtain permission to use cloud storage	"2330" on page 383 "2305" on page 382
087012	W	Cloud account not available, cannot complete cloud storage operation	"3100" on page 389
087013	W	Cloud account not available, cannot access cloud object storage	"2105" on page 381
087014	W	Cloud account not available, incompatible object data format	"3135" on page 391
087016	W	Cloud account not available, cloud object storage encrypted	"1656 " on page 368

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
087017	W	Cloud account not available, cloud object storage not encrypted	"1656 " on page 368
087018	W	Cloud account not available, cloud object storage encrypted with the wrong key	"1657" on page 369
087019	W	No permission to use cloud storage snapshot operation	"2305" on page 382
087020	W	Cloud account out of space during cloud storage snapshot operation	"2125" on page 381
087021	W	Cannot create container object to cloud object storage during cloud snapshot operation	"2305" on page 382
087022	W	A cloud object could not be found during cloud snapshot operation.	"3108" on page 389
087023	W	A cloud object was found to be corrupt during cloud snapshot operation.	"3108" on page 389
087024	W	A cloud object was found to be corrupt during cloud snapshot decompression operation.	"3108" on page 389
087025	W	Etag integrity error during cloud snapshot operation	"3108" on page 389
087026	W	Internal Read error during cloud snapshot operation	"2120" on page 381
087027	W	Unexpected error occurred, cannot complete cloud snapshot operation	"3108" on page 389
087028	W	No permission to use cloud snapshot restore operation	"2305" on page 382
087029	W	A cloud object could not be found during a cloud snapshot restore operation	"3108" on page 389
087030	W	A cloud object was found to be corrupt during a cloud snapshot restore operation	"3108" on page 389
087031	W	A cloud object was found to be corrupt during a cloud snapshot restore decompression operation	"3108" on page 389
087032	W	Etag integrity error during cloud snapshot restore operation	"3108" on page 389
087033	W	Internal write error during cloud snapshot operation	"2120" on page 381
087034	W	Cannot create bad blocks on a managed disk during cloud snapshot restore operation.	"3108" on page 389
087035	W	Unexpected error occurred, cannot complete cloud snapshot restore operation	"3108" on page 389
087036	W	No permission to use cloud snapshot delete operation	"2305" on page 382
087037	W	A cloud object could not be found during a cloud snapshot delete operation	"3108" on page 389
087038	W	A cloud object was found to be corrupt during cloud snapshot delete operation	"3108" on page 389

Table 108. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
087039	W	A cloud object was found to be corrupt during cloud snapshot delete decompression operation	"3108" on page 389
087040	W	Unexpected error occurred, cannot complete cloud snapshot delete operation	"3108" on page 389
087044	W	Cloud account out of space during cloud snapshot restore commit operation	"2125" on page 381
087045	W	Cloud account out of space during cloud snapshot delete operation	"2125" on page 381
087046	W	Transparent Cloud Tiering feature license limit exceeded.	3032
087048	W	Too many node restarts have occurred, cloud backup operations paused	3104
087049	W	Internal FlashCopy error on volume enabled for cloud snapshots.	2118
088000	E	An IO port cannot be started	1300
088001	E	A fibrechannel target port mode transition was not successful	1300
088002	W	Equivalent fibrechannel ports are reporting that they are connected to different fabrics	3220
088003	W	A spare node in this cluster is not providing additional redundancy	1380
088004	W	A spare node could not be automatically removed from the cluster	3180

Node error code overview

Node error codes describe failures that relate to a specific node canister.

Use the service assistant GUI by the Technician port to view node errors on a node that does not have a front panel display such as a 2145-DH8 node.

Because node errors are specific to a node, for example, memory failures, the errors might be reported only on that node. However, some of the conditions that the node detects relate to the shared components of the enclosure. In these cases, both node canisters in the enclosure report the error.

There are two types of node errors: critical node errors and noncritical node errors.

Critical errors

A critical error means that the node is not able to participate in a clustered system until the issue that is preventing it from joining a clustered system is resolved. This error occurs because part of the hardware failed or the system detects that the code is corrupted. If it is possible to communicate with the canister with a node error, an alert that describes the error is logged in the event log. If the system cannot communicate with the node canister, a Node missing alert is reported. If a node has a critical node error, it is in service state, and the fault LED on the node is on. The exception is when the node cannot connect to enough resources to form a clustered

system. It shows a critical node error but is in the starting state. The range of errors that are reserved for critical errors are 500 - 699.

Some critical errors might be accompanied by error codes 1021, 1036, 1188, and 1189.

Noncritical errors

A noncritical error code is logged when a hardware or code failure that is related to just one specific node. These errors do not stop the node from entering active state and joining a clustered system. If the node is part of a clustered system, an alert describes the error condition. The node error is shown to make it clear which of the node canisters the alert refers to. The range of errors that are reserved for noncritical errors are 800 - 899.

Clustered-system code overview

Recovery codes for clustered systems indicate that a critical software error has occurred that might corrupt your system. Each error-code topic includes an error code number, a description, action, and possible field-replaceable units (FRUs).

Error codes for recovering a clustered system

You must perform software problem analysis before you can perform further operations to avoid the possibility of corrupting your configuration.

Error code range

This topic shows the number range for each message classification.

Table 109 lists the number range for each message classification.

Table 109. Message classification number range

Message classification	Range	
Node errors	Critical node errors	500-699 Some critical errors might be accompanied by error codes 1021, 1036, 1188, and 1189.
	Noncritical node errors	700-899
Error codes when recovering a clustered system	920, 990	

100	Boot is running

Explanation: The node has started. It is running diagnostics and loading the runtime code.

User response: Go to the hardware boot MAP to resolve the problem.

Possible Cause-FRUs or other:

2145-CG8 or 2145-CF8

• Service controller (47%)

• Service controller cable (47%)

• System board assembly (6%)

120	Disk drive hardware error

Explanation: The internal disk drive of the node has reported an error. The node is unable to start.

User response: Ensure that the boot disk drive and all related cabling is properly connected, then exchange the FRU for a new FRU. (See "Possible Cause-FRUs or other.")

Possible Cause-FRUs or other:

2145-CF8 or 2145-CG8
- Disk drive (50%)
- Disk controller (30%)
- Disk backplane (10%)
- Disk signal cable (8%)
- Disk power cable (1%)
- System board (1%)

---

130   Checking the internal disk file system

Explanation: The file system on the internal disk drive of the node is being checked for inconsistencies.

User response: If the progress bar has been stopped for at least five minutes, power off the node and then power on the node. If the boot process stops again at this point, run the node rescue procedure.

Possible Cause-FRUs or other:
- None.

---

132   Updating BIOS settings of the node

Explanation: The system has found that changes are required to the BIOS settings of the node. These changes are being made. The node will restart once the changes are complete.

User response: None for Storwize V7000 systems.

---

135   Verifying the software

Explanation: The software packages of the node are being checked for integrity.

User response: Allow the verification process to complete.

---

137   Updating system board service processor firmware

Explanation: The service processor firmware of the node is being updated to a new level. This process can take 90 minutes. Do not restart the node while this is in progress.

User response: Allow the updating process to complete.

---

150   Loading cluster code

Explanation: The SAN Volume Controller code is being loaded.

User response: If the progress bar has been stopped for at least 90 seconds, power off the node and then power on the node. If the boot process stops again at this point, run the node rescue procedure.

Possible Cause-FRUs or other:
- None.

---

155   Loading cluster data

Explanation: The saved cluster state and cache data is being loaded.

User response: If the progress bar has been stopped for at least 5 minutes, power off the node and then power on the node. If the boot process stops again at this point, run the node rescue procedure.

Possible Cause-FRUs or other:
- None.

---

160   Updating the service controller

Explanation: The firmware on the service controller is being updated. This can take 30 minutes.

User response: When a node rescue is occurring, if the progress bar has been stopped for at least 30 minutes, exchange the FRU for a new FRU. When a node rescue is not occurring, if the progress bar has been stopped for at least 15 minutes, exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

2145-CG8 or 2145-CF8
- Service controller (95%)
- Service controller cable (5%)

All previous 2145 models
- Service Controller (100%)

---

168   The command cannot be initiated because authentication credentials for the current SSH session have expired.

Explanation: Authentication credentials for the current SSH session have expired, and all authorization for the current session has been revoked. A system administrator may have cleared the authentication cache.

User response: Begin a new SSH session and re-issue the command.

---

170   A flash module hardware error has occurred.

Explanation: A flash module hardware error has occurred.

User response: Exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

2145-CG8 or 2145-CF8
- Service controller (95%)
- Service controller cable (5%)

All previous 2145 models
- Service controller (100%)

---

182   Checking uninterruptible power supply

Explanation: The node is checking whether the uninterruptible power supply is operating correctly.

User response: Allow the checking process to complete.

---

232   Checking uninterruptible power supply connections

Explanation: The node is checking whether the power and signal cable connections to the uninterruptible power supply are correct.

User response: Allow the checking process to complete.

---

300   The 2145 is running node rescue.

Explanation: The 2145 is running node rescue.

User response: If the progress bar has been stopped for at least two minutes, exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

2145-CG8 or 2145-CF8
- Service controller (95%)
- Service controller cable (5%)
- Service controller (100%)

---

310   The 2145 is running a format operation.

Explanation: The 2145 is running a format operation.

User response: If the progress bar has been stopped for two minutes, exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

2145-CG8 or 2145-CF8
- Disk drive (50%)
- Disk controller (30%)
- Disk backplane (10%)
- Disk signal cable (8%)
- Disk power cable (1%)
- System board (1%)
- Disk drive assembly (90%)
- Disk cable assembly (10%)

---

320   A 2145 format operation has failed.

Explanation: A 2145 format operation has failed.

User response: Exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

2145-CG8 or 2145-CF8
- Disk drive (50%)

- Disk controller (30%)
- Disk backplane (10%)
- Disk signal cable (8%)
- Disk power cable (1%)
- System board (1%)
- Disk drive assembly (90%)
- Disk cable assembly (10%)

---

330   The 2145 is partitioning its disk drive.

Explanation: The 2145 is partitioning its disk drive.

User response: If the progress bar has been stopped for two minutes, exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

2145-CG8 or 2145-CF8
- Disk drive (50%)
- Disk controller (30%)
- Disk backplane (10%)
- Disk signal cable (8%)
- Disk power cable (1%)
- System board (1%)
- Disk drive assembly (90%)
- Disk cable assembly (10%)

Other:
- Configuration problem
- Software error

---

340   The 2145 is searching for donor node.

Explanation: The 2145 is searching for donor node.

User response: If the progress bar has been stopped for more than two minutes, exchange the FRU for a new FRU.

Possible Cause-FRUs or other:
- Fibre Channel adapter (100%)

---

345   The 2145 is searching for a donor node from which to copy the software.

Explanation: The node is searching at 1 Gb/s for a donor node.

User response: If the progress bar has stopped for more than two minutes, exchange the FRU for a new FRU.

Possible Cause-FRUs or other:
- Fibre Channel adapter (100%)

---

350 **The 2145 cannot find a donor node.**

Explanation: The 2145 cannot find a donor node.

User response: If the progress bar has stopped for more than two minutes, perform the following steps:

1. Ensure that all of the Fibre Channel cables are connected correctly and securely to the cluster.

2. Ensure that at least one other node is operational, is connected to the same Fibre Channel network, and is a donor node candidate. A node is a donor node candidate if the version of software that is installed on that node supports the model type of the node that is being rescued.

3. Ensure that the Fibre Channel zoning allows a connection between the node that is being rescued and the donor node candidate.

4. Perform the problem determination procedures for the network.

Possible Cause-FRUs or other:

- None

Other:

- Fibre Channel network problem

---

360 **The 2145 is loading software from the donor.**

Explanation: The 2145 is loading software from the donor.

User response: If the progress bar has been stopped for at least two minutes, restart the node rescue procedure.

Possible Cause-FRUs or other:

- None

---

365 **Cannot load SW from donor**

Explanation: None.

User response: None.

---

370 **Installing software**

Explanation: The 2145 is installing software.

User response:

1. If this code is displayed and the progress bar has been stopped for at least ten minutes, the software install process has failed with an unexpected software error.

2. Power off the 2145 and wait for 60 seconds.

3. Power on the 2145. The software update operation continues.

4. Report this problem immediately to your Software Support Center.

Possible Cause-FRUs or other:

- None

---

500 **Incorrect enclosure**

Explanation: The node canister has saved cluster information, which indicates that the canister is now located in a different enclosure from where it was previously used. Using the node canister in this state might corrupt the data held on the enclosure drives.

User response: Follow troubleshooting procedures to move the nodes to the correct location.

1. Follow the "Procedure: Getting node canister and system information using the service assistant" task to review the node canister saved location information and the status of the other node canister in the enclosure (the partner canister). Determine if the enclosure is part of an active system with volumes that contain required data.

2. If you have unintentionally moved the canister into this enclosure, move the canister back to its original location, and put the original canister back in this enclosure. Follow the "Replacing a node canister" procedure.

3. If you have intentionally moved the node canister into this enclosure you should check it is safe to continue or whether you will lose data on the enclosure you removed it from. Do not continue if the system the node canister was removed from is offline, rather return the node canister to that system.

4. If you have determined that you can continue, follow the "Procedure: Removing system data from a node canister" task to remove cluster data from node canister.

5. If the partner node in this enclosure is not online, or is not present, you will have to perform a system recovery. Do not create a new system, you will lose all the volume data.

Possible Cause-FRUs or other cause:

- None

---

501 **Incorrect slot**

Explanation: The node canister has saved cluster information, which indicates that the canister is not located in the expected enclosure, but in a different slot from where it was previously used. Using the node canister in this state might mean that hosts are not able to connect correctly.

User response: Follow troubleshooting procedures to relocate the node canister to the correct location.

1. Follow the "Procedure: Getting node canister and system information using the service assistant" task to review the node canister saved location information and the status of the other node canister in the enclosure (the partner canister). If the

node canister has been inadvertently swapped, the other node canister will have the same error.

2. If the canisters have been swapped, use the "Replacing a node canister" procedure to swap the canisters. The system should start.

3. If the partner canister is in candidate state, use the hardware remove and replace canister procedure to swap the canisters. The system should start.

4. If the partner canister is in active state, it is running the cluster on this enclosure and has replaced the original use of this canister. Follow the "Procedure: Removing system data from a node canister" task to remove cluster data from this node canister. The node canister will then become active in the cluster in its current slot.

5. If the partner canister is in service state, review its node error to determine the correct action. Generally, you will fix the errors reported on the partner node in priority order, and review the situation again after each change. If you have to replace the partner canister with a new one, you should move this canister back to the correct location at the same time.

Possible Cause-FRUs or other:

- None

---

502   No enclosure identity exists and a status from the partner node could not be obtained.

Explanation: The enclosure has been replaced and communication with the other node canister (partner node) in the enclosure is not possible. The partner node could be missing, powered off, unable to boot, or an internode communication failure may exist.

User response: Follow troubleshooting procedures to configure the enclosure:

1. Follow the procedures to resolve a problem to get the partner node started. An error will still exist because the enclosure has no identity. If the error has changed, follow the service procedure for that error.

2. If the partner has started and is showing a location error (probably this one), then the PCI link is probably broken. Since the enclosure midplane was recently replaced, this is likely the problem. Obtain a replacement enclosure midplane, and replace it. See "Replacing a Storwize V7000 Gen1 control enclosure chassis" on page 221.

3. If this action does not resolve the issue, contact IBM Support Center. They will work with you to ensure that the system state data is not lost while resolving the problem.

Possible Cause—FRUs or other:

- Enclosure midplane (100%)

---

503   Incorrect enclosure type

Explanation: The node canister has been moved to an expansion enclosure. A node canister will not operate in this environment. This can also be reported when a replacement node canister is installed for the first time.

User response: If this node error is reported when a node canister is being replaced with a field-replaceable unit (FRU) or a customer-replaceable unit (CRU), apply "Procedure: Rescuing Storwize V7000 Gen2 node canister software from another node (node rescue)" on page 138 to rescue the replacement node canister.

Follow troubleshooting procedures to relocate the nodes to the correct location.

1. Follow the procedure Getting node canister and system information using a USB flash drive and review the saved location information of the node canister to determine which control enclosure the node canister belongs in.

2. Follow the procedure to move the node canister to the correct location, then follow the procedure to move the expansion canister that is probably in that location to the correct location. If there is a node canister that is in active state where this node canister must be, do not replace that node canister with this one.

---

504   No enclosure identity and partner node matches.

Explanation: The enclosure vital product data indicates that the enclosure midplane has been replaced. This node canister and the other node canister in the enclosure were previously operating in the same enclosure midplane.

User response: Follow troubleshooting procedures to configure the enclosure.

1. This is an expected situation during the hardware remove and replace procedure for a control enclosure midplane. Continue following the remove and replace procedure and configure the new enclosure.

Possible Cause—FRUs or other:

- None

---

505   No enclosure identity and partner has system data that does not match.

Explanation: The enclosure vital product data indicates that the enclosure midplane has been replaced. This node canister and the other node canister in the enclosure do not come from the same original enclosure.

User response: Follow troubleshooting procedures to relocate nodes to the correct location.

1. Follow the "Procedure: Getting node canister and system information using the service assistant" task to review the node canister saved location information and the status of the other node canister in the enclosure (the partner canister). Determine if the enclosure is part of an active system with volumes that contain required data.

2. Decide what to do with the node canister that did not come from the enclosure that is being replaced.

   a. If the other node canister from the enclosure being replaced is available, use the hardware remove and replace canister procedures to remove the incorrect canister and replace it with the second node canister from the enclosure being replaced. Restart both canisters. The two node canister should show node error 504 and the actions for that error should be followed.

   b. If the other node canister from the enclosure being replaced is not available, check the enclosure of the node canister that did not come from the replaced enclosure. Do not use this canister in this enclosure if you require the volume data on the system from which the node canister was removed, and that system is not running with two online nodes. You should return the canister to its original enclosure and use a different canister in this enclosure.

   c. When you have checked that it is not required elsewhere, follow the "Procedure: Removing system data from a node canister" task to remove cluster data from the node canister that did not come from the enclosure that is being replaced.

   d. Restart both nodes. Expect node error 506 to be reported now, then follow the service procedures for that error.

Possible Cause—FRUs or other:

• None

---

506   No enclosure identity and no node state on partner

Explanation: The enclosure vital product data indicates that the enclosure midplane has been replaced. There is no cluster state information on the other node canister in the enclosure (the partner canister), so both node canisters from the original enclosure have not been moved to this one.

User response: Follow troubleshooting procedures to relocate nodes to the correct location:

1. Follow the procedure: Getting node canister and system information and review the saved location information of the node canister and determine why the second node canister from the original enclosure was not moved into this enclosure.

2. If you are sure that this node canister came from the enclosure that is being replaced, and the original

partner canister is available, use the "Replacing a node canister" procedure to install the second node canister in this enclosure. Restart the node canister. The two node canisters should show node error 504, and the actions for that error should be followed.

3. If you are sure this node canister came from the enclosure that is being replaced, and that the original partner canister has failed, continue following the remove and replace procedure for an enclosure midplane and configure the new enclosure.

Possible Cause—FRUs or other:

• None

---

507   No enclosure identity and no node state

Explanation: The node canister has been placed in a replacement enclosure midplane. The node canister is also a replacement or has had all cluster state removed from it.

User response: Follow troubleshooting procedures to relocate the nodes to the correct location.

1. Check the status of the other node in the enclosure. Unless it also shows error 507, check the errors on the other node and follow the corresponding procedures to resolve the errors. It typically shows node error 506.

2. If the other node in the enclosure is also reporting 507, the enclosure and both node canisters have no state information. Contact IBM support. They will assist you in setting the enclosure vital product data and running cluster recovery.

Possible Cause-FRUs or other:

• None

---

508   Cluster identifier is different between enclosure and node

Explanation: The node canister location information shows it is in the correct enclosure, however the enclosure has had a new clustered system created on it since the node was last shut down. Therefore, the clustered system state data stored on the node is not valid.

User response: Follow troubleshooting procedures to correctly relocate the nodes.

1. Check whether a new clustered system has been created on this enclosure while this canister was not operating or whether the node canister was recently installed in the enclosure.

2. Follow the "Procedure: Getting node canister and system information using the service assistant" task, and check the partner node canister to see if it is also reporting node error 508. If it is, check that the saved system information on this and the partner node match.

If the system information on both nodes matches, follow the "Replacing a control enclosure midplane" procedure to change the enclosure midplane.

3. If this node canister is the one to be used in this enclosure, follow the "Procedure: Removing system data from a node canister" task to remove clustered system data from the node canister. It will then join the clustered system.

4. If this is not the node canister that you intended to use, follow the "Replacing a node canister" procedure to replace the node canister with the one intended for use.

Possible Cause—FRUs or other:
- Service procedure error (90%)
- Enclosure midplane (10%)

---

509   The enclosure identity cannot be read.

Explanation: The canister was unable to read vital product data (VPD) from the enclosure. The canister requires this data to be able to initialize correctly.

User response: Follow troubleshooting procedures to fix the hardware:

1. Check errors reported on the other node canister in this enclosure (the partner canister).

2. If it is reporting the same error, follow the hardware remove and replace procedure to replace the enclosure midplane.

3. If the partner canister is not reporting this error, follow the hardware remove and replace procedure to replace this canister.

Note: If a newly installed system has this error on both node canisters, the data that needs to be written to the enclosure will not be available on the canisters; contact IBM support for the WWNNs to use.

Remember: Review the `lsservicenodes` output for what the node is reporting.
Possible Cause—FRUs or other:
- Node canister (50%)
- Enclosure midplane (50%)

---

510   The detected memory size does not match the expected memory size.

Explanation: The amount of memory detected in the node canister differs from the amount required for the canister to operate as an active member of a system. The error code data shows the detected memory (in MB) followed by the minimum required memory (in MB). A series of values indicates the amount of memory (in GB) detected in each memory slot.

Data:
- Detected memory in MB
- Minimum required memory in MB

- Memory in slot 1 in GB
- Memory in slot 2 in GB
- ...
- Memory in slot *n* in GB

User response: Follow troubleshooting procedures to fix the hardware:

1. Use the hardware remove and replace node canister procedure to install a new node canister.

Possible Cause-FRUs or other:
- Node canister (100%)

---

511   Memory bank 1 of the 2145 is failing. For the 2145-DH8 only, the DIMMS are incorrectly installed.

Explanation: Memory bank 1 of the 2145 is failing.

For the 2145-DH8 only, the DIMMS are incorrectly installed. This will degrade performance.

User response: For the 2145-DH8 only, shut down the node and adjust the DIMM placement as per the install directions.

Possible Cause-FRUs or other:
- Memory module (100%)

---

512   Enclosure VPD is inconsistent

Explanation: The enclosure midplane VPD is not consistent. The machine part number is not compatible with the machine type and model. This indicates that the enclosure VPD is corrupted.

User response:

1. Check the support site for a code update.

2. Use the remove and replace procedures to replace the enclosure midplane.

Possible Cause—FRUs or other:
- Enclosure midplane (100%)

---

513   Memory bank 2 of the 2145 is failing.

Explanation: Memory bank 2 of the 2145 is failing.

User response: For the 2145-8G4 and 2145-8A4, go to the light path MAP to resolve this problem.

Possible Cause-FRUs or other:
- Memory module (100%)

---

514   Memory bank 3 of the 2145 is failing.

Explanation: Memory bank 3 of the 2145 is failing.

User response: For the 2145-8G4 and 2145-8A4, go to the light path MAP to resolve this problem.

Possible Cause-FRUs or other:

• Memory module (100%)

---

515   **Memory bank 4 of the 2145 is failing.**

Explanation: Memory bank 4 of the 2145 is failing.

User response: For the 2145-8G4 and 2145-8A4, go to the light path MAP to resolve this problem.

Possible Cause-FRUs or other:

• Memory module (100%)

---

517   **The WWNNs of the service controller and the disk do not match.**

Explanation: The node is unable to determine the WWNN that it must use. This is because of the service controller or the nodes internal drive being replaced.

User response: Follow troubleshooting procedures to configure the WWNN of the node.

1. Continue to follow the hardware remove and replace procedure for the service controller or disk.

2. If you have not followed the hardware remove and replace procedures, determine the correct WWNN. If you do not have this information recorded, examine your Fibre Channel switch configuration to see whether it is listed there. Follow the procedures to change the WWNN of a node.

Possible Cause-FRUs or other:

• None

---

521   **Unable to detect a Fibre Channel adapter**

Explanation: The system cannot detect any Fibre Channel adapters.

User response: Ensure that a Fibre Channel adapter has been installed. Ensure that the Fibre Channel adapter is seated correctly in the riser card. Ensure that the riser card is seated correctly on the system board. If the problem persists, exchange FRUs for new FRUs in the order shown.

Possible Cause-FRUs or other:

2145-CG8 or 2145-CF8

• 4-port Fibre Channel host bus adapter assembly (95%)

• System board assembly (5%)

---

522   **The system board service processor has failed.**

Explanation: The service processor (PSOC) in the canister has failed or is not communicating.

User response:

1. Reseat the node canister.

2. If the error persists, use the remove and replace procedures to replace the node canister.

Possible Cause-FRUs or other:

• Node canister

---

523   **The internal disk file system is damaged.**

Explanation: The node startup procedures have found problems with the file system on the internal disk of the node.

User response: Follow troubleshooting procedures to reload the software.

1. Follow the procedures to rescue the software of a node from another node.

2. If the rescue node does not succeed, use the hardware remove and replace procedures for the node canister.

Possible Cause—FRUs or other:

• Node canister (100%)

---

524   **Unable to update BIOS settings.**

Explanation: Unable to update BIOS settings.

User response: Power off node, wait 30 seconds, and then power on again. If the error code is still reported, replace the system board.

Possible Cause-FRUs or other:

• System board (100%)

---

525   **Unable to update system board service processor firmware.**

Explanation: The node startup procedures have been unable to update the firmware configuration of the node canister.

User response: Follow troubleshooting procedures to fix the hardware:

1. Follow the hardware remove and replace procedures for the node canister.

Possible Cause—FRUs or other:

• Node canister (100%)

---

528   **Ambient temperature is too high during system startup.**

Explanation: The ambient temperature in the enclosure, read during the node canister startup procedures, is too high for the node canister to continue. The startup procedure will continue when the temperature is within range.

User response: Reduce the temperature around the system.

1. Resolve the issue with the ambient temperature, by checking and correcting:
   a. Room temperature and air conditioning
   b. Ventilation around the rack
   c. Airflow within the rack

Possible Cause—FRUs or other:
- Environment issue (100%)

---

530   **A problem with one of the node's power supplies has been detected.**

Explanation: The 530 error code is followed by two numbers. The first number is either 1 or 2 to indicate which power supply has the problem.

The second number is either 1, 2 or 3 to indicate the reason.

1   The power supply is not detected.

2   The power supply failed.

3   No input power is available to the power supply.

If the node is a member of a cluster, the cluster reports error code 1096 or 1097, depending on the error reason.

The error will automatically clear when the problem is fixed.

User response:
1. Ensure that the power supply is seated correctly and that the power cable is attached correctly to both the node and to a power source.
2. If the error has not been automatically marked fixed after two minutes, note the status of the three LEDs on the back of the power supply. For the 2145-CG8 or 2145-CF8, the AC LED is the top green LED, the DC LED is the middle green LED and the error LED is the bottom amber LED.
3. If the power supply error LED is off and the AC and DC power LEDs are both on, this is the normal condition. If the error has not been automatically fixed after two minutes, replace the system board.
4. Follow the action specified for the LED states noted in the list below.
5. If the error has not been automatically fixed after two minutes, contact support.

Error, AC, DC: Action

ON,ON or OFF,ON or OFF:The power supply has a fault. Replace the power supply.

OFF,OFF,OFF:There is no power detected. Ensure that the power cable is connected at the node and to a power source. If the AC LED does not light, check your power source. If you are connected to a 2145 UPS-1U

that is showing an error, follow MAP 5150 2145 UPS-1U. Otherwise, replace the power cable. If the AC LED still does not light, replace the power supply.

OFF,OFF,ON:The power supply has a fault. Replace the power supply.

OFF,ON,OFF:Ensure that the power supply is installed correctly. If the DC LED does not light, replace the power supply.

Possible Cause-FRUs or other:

Reason 1: A power supply is not detected.
- Power supply (19%)
- System board (1%)
- Other: Power supply is not installed correctly (80%)

Reason 2: The power supply has failed.
- Power supply (90%)
- Power cable assembly (5%)
- System board (5%)

Reason 3: There is no input power to the power supply.
- Power cable assembly (25%)
- UPS-1U assembly (4%)
- System board (1%)
- Other: Power supply is not installed correctly (70%)

---

534   **System board fault**

Explanation: There is a unrecoverable error condition in a device on the system board.

User response: For a storage enclosure, replace the canister and reuse the interface adapters and fans.

For a control enclosure, refer to the additional details supplied with the error to determine the proper parts replacement sequence.
- Pwr rail A: Replace CPU 1.

  Replace the power supply if the OVER SPEC LED on the light path diagnostics panel is still lit.
- Pwr rail B: Replace CPU 2.

  Replace the power supply if the OVER SPEC LED on the light path diagnostics panel is still lit.
- Pwr rail C: Replace the following components until "Pwr rail C" is no longer reported:
  - DIMMs 1 - 6
  - PCI riser-card assembly 1
  - Fan 1
  - Optional adapters that are installed in PCI riser-card assembly 1
  - Replace the power supply if the OVER SPEC LED on the light path diagnostics panel is still lit.

- Pwr rail D: Replace the following components until "Pwr rail D" is no longer reported:
  - DIMMs 7 - 12
  - Fan 2
  - Optional PCI adapter power cable
  - Replace the power supply if the OVER SPEC LED on the light path diagnostics panel is still lit.
- Pwr rail E: Replace the following components until "Pwr rail E" is no longer reported:
  - DIMMs 13 - 18
  - Hard disk drives
  - Replace the power supply if the OVER SPEC LED on the light path diagnostics panel is still lit.
- Pwr rail F: Replace the following components until "Pwr rail F" is no longer reported:
  - DIMMs 19 - 24
  - Fan 4
  - Optional adapters that are installed in PCI riser-card assembly 2
  - PCI riser-card assembly 2
  - Replace the power supply if the OVER SPEC LED on the light path diagnostics panel is still lit.
- Pwr rail G: Replace the following components until "Pwr rail G" is no longer reported:
  - Hard disk drive backplane assembly
  - Hard disk drives
  - Fan 3
  - Optional PCI adapter power cable
- Pwr rail H: Replace the following components until "Pwr rail H" is no longer reported:
  - Optional adapters that are installed in PCI riser-card assembly 2
  - Optional PCI adapter power cable

Possible Cause—FRUs or other:

- Hardware (100%)

---

535   **Canister internal PCIe switch failed**

Explanation: The PCI Express switch has failed or cannot be detected. In this situation, the only connectivity to the node canister is through the Ethernet ports.

User response: Follow troubleshooting procedures to fix the hardware.

1. Follow the procedure for reseating a node canister. See "Procedure: Reseating a Storwize V7000 Gen1 node canister" on page 129.
2. If reseating the canister does not resolve the situation, follow the "Replacing a Storwize V7000 Gen1 node canister" on page 178 procedure to replace the canister.

Possible Cause—FRUs or other:

- Node canister (100%)

---

536   **The temperature of a device on the system board is greater than or equal to the critical threshold.**

Explanation: The temperature of a device on the system board is greater than or equal to the critical threshold.

User response: Check for external and internal air flow blockages or damage.

1. Remove the top of the machine case and check for missing baffles, damaged heat sinks, or internal blockages.
2. If the error persists, replace system board.

Possible Cause-FRUs or other:

- None

---

538   **The temperature of a PCI riser card is greater than or equal to the critical threshold.**

Explanation: The temperature of a PCI riser card is greater than or equal to the critical threshold.

User response: Improve cooling.

1. If the problem persists, replace the PCI riser

Possible Cause-FRUs or other:

- None

---

541   **Multiple, undetermined, hardware errors**

Explanation: Multiple hardware failures were reported on the data paths within the node, and the threshold of the number of acceptable errors within a given time frame was reached. It was not possible to isolate the errors to a single component.

After this node error is raised, all ports on the node are deactivated. The node is considered unstable, and has the potential to corrupt data.

User response:

1. Follow the procedure for collecting information for support, and contact your support organization.
2. A software update may resolve the issue.
3. Replace the node.

---

542   **An installed CPU has failed or been removed.**

Explanation: An installed CPU has failed or been removed.

User response: Replace the CPU.

Possible Cause-FRUs or other:

- CPU (100%)

---

543   **None of the node serial numbers that are stored in the three locations match.**

Explanation: When the Storwize V7000 software starts, it reads the node serial number from the system board and compares this serial number to the node serial numbers stored on the two boot drives. There must be at least two matching node serial numbers for the Storwize V7000 software to assume that node serial number is good.

User response: Look at a boot drive view for the node to work out what to do.

1. Replace missing or failed drives.
2. Put any drive that belongs to a different node back where it belongs.
3. If you intend to use a drive from a different node in this node from now on, the node error changes to a different node error when the other drive is replaced.
4. If you replaced the system board, then the panel name is now 0000000, and if you replaced one of the drives, then the slot status of that drive is uninitialized. If the node serial number of the other boot drive matches the MT-M S/N label on the front of the node, then run **satask rescuenode** to initialize the uninitialized drive. Initializing the drive should lead to the 545 node error.

Possible Cause-FRUs or other:

- None

---

544   **Boot drives are from other nodes.**

Explanation: Boot drives are from other nodes.

User response: Look at a boot drive view for the node to determine what to do.

1. Put any drive that belongs to a different node back where it belongs.
2. If you intend to use a drive from a different node in this node from now on, the node error changes to a different node error when the other drive is replaced.
3. See error code 1035 for additional information regarding boot drive problems.

Possible Cause-FRUs or other:

- None

---

545   **The node serial number on the boot drives match each other, but they do not match the product serial number on the system board.**

Explanation: The node serial number on the boot

drives match each other, but they do not match the product serial number on the system board.

User response: Check the S/N value on the MT-M S/N label on the front of the node. Look at a boot drive view to see the node serial number of the system board and the node serial number of each drive.

1. Replace the boot drives with the correct boot drives if needed.
2. Set the system board serial number using the following command:

   satask chvpd -type *<value>* **-serial** *<S/N value from the MT-M S/N label>*

Possible Cause-FRUs or other:

- None

---

547   **Pluggable TPM is missing or broken.**

Explanation: The Trusted Platform Module (TPM) for the system is not functioning.

User response:

Important: Confirm that the system is running on at least one other node before you commence this repair. Each node uses its TPM to securely store encryption keys on its boot drive. When the TPM or boot drive of a node is replaced, the node loses its encryption key, and must be able to join an existing system to obtain the keys. If this error occurred on the last node in a system, do not replace the TPM, boot drive, or node hardware until the system contains at least one online node with valid keys.

1. Shut down the node and remove the node hardware.
2. Locate the TPM in the node hardware and ensure that it is correctly seated.
3. Reinsert the node hardware and apply power to the node.
4. If the error persists, replace the TPM with one from FRU stock.
5. If the error persists, replace the system board or the node hardware with one from FRU stock.

You do not need to return the faulty TPM to IBM.

Note: It is unlikely that the failure of a TPM can cause the loss of the System Master Key (SMK):

- The SMK is sealed by the TPM, using its unique encryption key, and the result is stored on the system boot drive.
- The working copy of the SMK is on the RAM disk, and so is unaffected by a sudden TPM failure.
- If the failure happens at boot time, the node is held in an unrecoverable error state because the TPM is a FRU.

- The SMK is also mirrored by the other nodes in the system. When the node with replacement TPM joins the system, it determines that it does not have the SMK, requests it, gets it, and then seals with the new TPM.

---

550   **A clustered system cannot be formed because of a lack of clustered system resources.**

Explanation: The node canister cannot become active in a clustered system because it is unable to connect to enough clustered system resources. The clustered system resources are the node canisters in the system and the active quorum disk or drive. The node must be able to connect to most of the resources before that group forms an online clustered system. This connection prevents the clustered system from splitting into two or more active parts, with both parts independently performing I/O.

The error data lists the missing resources. This information includes a list of node canisters and optionally a drive that is operating as the quorum drive or a LUN on an external storage system that is operating as the quorum disk.

If a drive in one of the system enclosures is the missing quorum disk, it is listed as enclosure:slot[part identification] where enclosure:slot is the location of the drive when the node shutdown, enclosure is the seven-digit product serial number of the enclosure, slot is a number 1 - 24. The part identification is the 22 character string that starts with "11S" found on a label on a drive. The part identification cannot be seen until the drive is removed from the enclosure.

If a LUN on an external storage system is the missing quorum disk, it is listed as WWWWWWWWWWWWWWWW/LL, where WWWWWWWWWWWWWWWW is a worldwide port name (WWPN) on the storage system that contains the missing quorum disk and LL is the Logical Unit Number (LUN).

User response: Follow troubleshooting procedures to correct connectivity issues between the system canisters and the quorum devices.

1. Check the status of other node canisters in the system and resolve any faults.

2. Check that all enclosures in the system are powered on and that the SAS cabling between the enclosures has not been disturbed. If any wiring changes have been made, check that all cables are securely connected and that the cabling rules have been followed.

   Check that all nodes in the system are shown in the service assistant or by using the service command: **sainfo lsservicenodes**. Investigate any missing nodes.

3. Check all nodes and quorum disks shown in the error data and check the communication links from this node to those nodes and quorum disks.

   a. If a quorum drive in a system enclosure is shown as missing, find the drive and check that it is working. The drive may have been moved from the location shown. In that case, find the drive and ensure it is installed and working. If the drive is not located in the control enclosure, try moving it to the control enclosure. A problem in SAS connectivity might be the issue.

      Note: If you are able to reestablish the system's operation, you will be able to use the extra diagnostics the system provides to diagnose problem on SAS cables and expansion enclosures.

   b. If a quorum disk on an external storage system is shown as missing, find the storage controller and confirm that the LUN is available. Check that the Fibre Channel connections between the storage controller and the 2076 are working and that any changes made to the SAN configuration and zoning have not effected the connectivity. Check the status of the Fibre Channel ports on the node and resolve any issues.

4. If all canisters have either node error 578 or 550, attempt to reestablish a clustered system by following the service procedures for the nodes showing node error 578. If this is not successful, follow the system recovery procedures.

---

551   **A cluster cannot be formed because of a lack of cluster resources.**

Explanation: The node does not have sufficient connectivity to other nodes or the quorum device to form a cluster.

Attempt to repair the fabric or quorum device to establish connectivity. If a disaster occurred and the nodes at the other site cannot be recovered, then it is possible to allow the nodes at the surviving site to form a system by using local storage.

User response: Follow troubleshooting procedures to correct connectivity issues between the cluster nodes and the quorum devices.

1. Check for any node errors that indicate issues with Fibre Channel connectivity. Resolve any issues.

2. Ensure that the other nodes in the cluster are powered on and operational.

3. Using the SAT GUI or CLI (sainfo lsservicestatus), display the Fibre Channel port status. If any port is not active, perform the Fibre Channel port problem determination procedures.

4. Ensure that Fibre Channel network zoning changes have not restricted communication between nodes or between the nodes and the quorum disk.

5. Perform the problem determination procedures for the network.

6. The quorum disk failed or cannot be accessed. Perform the problem determination procedures for the disk controller.

7. As a last resort when the nodes at the other site cannot be recovered, then it is possible to allow the nodes at the surviving site to form a system by using local site storage:

   To avoid data corruption ensure that all host servers that were previously accessing the system have had all volumes unmounted or have been rebooted. Ensure that the nodes at the other site are not operational and are unable to form a system in the future.

   After starting this command, a full resynchronization of all mirrored volumes is completed when the other site is recovered. This is likely to take many hours or days to complete.

   Contact IBM support personnel if you are unsure.

   Note: Before continuing, confirm that you have taken the following actions - failure to perform these actions can lead to data corruption that is undetected by the system but affects host applications.

   a. All host servers that were previously accessing the system have had all volumes unmounted or have been rebooted.

   b. Ensure that the nodes at the other site are not operating as a system and actions have been taken to prevent them from forming a system in the future.

   After these actions have been taken, the `satask overridequorum` can be used to allow the nodes at the surviving site to form a system that uses local storage.

---

555   **Power Domain error**

Explanation: Both 2145s in an I/O group that are being powered by the same uninterruptible power supply. The ID of the other 2145 is displayed with the node error code on the front panel.

User response: Ensure that the configuration is correct and that each 2145 is in an I/O group is connected from a separate uninterruptible power supply.

---

556   **A duplicate WWNN has been detected.**

Explanation: The node has detected another device that has the same World Wide Node Name (WWNN) on the Fibre Channel network. A WWNN is 16 hexadecimal digits long. For the Storwize V7000 system, the first 11 digits are always 50050768020. The last 5 digits of the WWNN are given in the additional data of the error. The Fibre Channel ports of the node

are disabled to prevent disruption of the Fibre Channel network. One or both nodes with the same WWNN can show the error. Because of the way WWNNs are allocated, a device with a duplicate WWNN is normally another Storwize V7000 node.

User response:

1. Find the Storwize V7000 node with the same WWNN as the node reporting the error. The WWNN for a Storwize V7000 node can be found from the node Vital Product Data (VPD) or from the node details shown by the service assistant. The node with the duplicate WWNN need not be part of the same cluster as the node reporting the error; it could be remote from the node reporting the error on a part of the fabric connected through an inter-switch link. The two nodes within a control enclosure must have different WWNNs. The WWNN of the node is stored within the enclosure chassis, so the duplication is most likely caused by the replacement of a control enclosure chassis.

2. If a Storwize V7000 node with a duplicate WWNN is found, determine whether it, or the node reporting the error, has the incorrect WWNN. Generally, it is the node that has had its enclosure chassis recently replaced or had its WWNN changed incorrectly. Also, consider how the SAN is zoned when making your decision.

3. Determine the correct WWNN for the node with the incorrect WWNN. If the enclosure chassis has been replaced as part of a service action, the WWNN for the node should have been written down. If the correct WWNN cannot be determined, contact your support center for assistance.

4. Use the service assistant to modify the incorrect WWNN. If it is the node showing the error that should be modified, this can safely be done immediately. If it is an active node that should be modified, use caution because the node will restart when the WWNN is changed. If this node is the only operational node in an enclosure, access to the volumes that it is managing will be lost. You should ensure that the host systems are in the correct state before you change the WWNN.

5. If the node showing the error had the correct WWNN, it can be restarted, using the service assistant, after the node with the duplicate WWNN is updated.

6. If you are unable to find a Storwize V7000 node with the same WWNN as the node showing the error, use the SAN monitoring tools to determine whether there is another device on the SAN with the same WWNN. This device should not be using a WWNN assigned to a Storwize V7000, so you should follow the service procedures for the device to change its WWNN. Once the duplicate has been removed, restart the node.

558	The node is unable to communicate with other nodes.

Explanation: The system cannot see the Fibre Channel fabric or the Fibre Channel adapter port speed might be set to a different speed than that of the Fibre Channel fabric.

User response: Ensure that:

1. The Fibre Channel network fabric switch is powered-on.
2. At least one Fibre Channel cable connects the system to the Fibre Channel network fabric.
3. The Fibre Channel adapter port speed is equal to that of the Fibre Channel fabric.
4. At least one Fibre Channel adapter is installed in the system.
5. Go to the Fibre Channel MAP.

Possible Cause-FRUs or other:

- None

560	Battery cabling fault

Explanation: A fault exists in one of the cables connecting the battery backplane to the rest of the system.

User response: Follow troubleshooting procedures to fix the hardware:

1. Reseat the cable.
2. If reseating the cable does not fix the problem, replace the cable.
3. If replacing the cable does not fix the problem, replace the battery backplane.

561	Battery backplane or cabling fault

Explanation: Either the battery backplane has failed, or the power or LPC cables connecting the battery backplane to the rest of the system are not connected properly.

User response: Follow troubleshooting procedures to fix the hardware:

1. Check the cables connecting the battery backplane.
2. Reseat the power and LPC cables.
3. If reseating the cables does not fix the problem, replace the cables.
4. Once the cables are well connected, but the problem persists, replace the battery backplane.
5. Conduct the corrective service procedure described in "1108" on page 344.

562	The nodes hardware configuration does not meet the minimum requirements

Explanation: The node hardware is not at the minimum specification for the node to become active in a cluster. This may be because of hardware failure, but is also possible after a service action has used an incorrect replacement part.

User response: Follow troubleshooting procedures to fix the hardware:

1. It is not possible to service parts within the node canister. Reseat the existing node canister to see whether the problem fixes. If it does not, use the hardware node canister remove and replace procedures to change the node canister.

564	Too many machine code crashes have occurred.

Explanation: The node has been determined to be unstable because of multiple resets. The cause of the resets can be that the system encountered an unexpected state or has executed instructions that were not valid. The node has entered the service state so that diagnostic data can be recovered.

The node error does not persist across restarts of the machine code on the node.

User response: Follow troubleshooting procedures to reload the machine code:

1. Get a support package (snap), including dumps, from the node, using the management GUI or the service assistant.
2. If more than one node is reporting this error, contact IBM technical support for assistance. The support package from each node will be required.
3. Check the support site to see whether the issue is known and whether a machine code update exists to resolve the issue. Update the cluster machine code if a resolution is available. Use the manual update process on the node that reported the error first.
4. If the problem remains unresolved, contact IBM technical support and send them the support package.

Possible Cause—FRUs or other:

- None

565	The internal drive of the node is failing.

Explanation: The internal drive within the node is reporting too many errors. It is no longer safe to rely on the integrity of the drive. Replacement is recommended.

User response: Follow troubleshooting procedures to fix the hardware:

1. The drive of the node canister cannot be replaced individually. Follow the hardware remove and replace instructions to change the node canister.

Possible Cause—FRUs or other:

- Node canister (100%)

---

569   At boot time: the CPU reached a temperature that is greater than or equal to the warning threshold. During normal running: the CPU reached a temperature that is greater than or equal to the critical threshold.

Explanation: At boot time: the CPU reached a temperature that is greater than or equal to the warning threshold. During normal running: the CPU reached a temperature that is greater than or equal to the critical threshold.

User response: Check for external and internal air flow blockages or damage.

1. Remove the top of the machine case and check for missing baffles, damaged heat sinks, or internal blockages.
2. If problem persists, replace the CPU/heat sink.

Possible Cause-FRUs or other:

- CPU
- Heat sink

---

570   Battery protection unavailable

Explanation: The node cannot start because battery protection is not available. Both batteries require user intervention before they can become available.

User response: Follow troubleshooting procedures to fix hardware.

The appropriate service action will be indicated by an accompanying non-fatal node error. Examine the event log to determine the accompanying node error.

---

571   Battery protection temporarily unavailable; one battery is expected to be available soon

Explanation: The node cannot start because battery protection is not available. One battery is expected to become available shortly with no user intervention required, but the other battery will not become available.

User response: Follow troubleshooting procedures to fix hardware.

The appropriate service action will be indicated by an accompanying non-fatal node error. Examine the event log to determine the accompanying node error.

---

572   Battery protection temporarily unavailable; both batteries are expected to be available soon

Explanation: The node cannot start because battery protection is not available. Both batteries are expected to become available shortly with no user intervention required.

User response: Wait for sufficient battery charge for enclosure to start.

---

573   The node machine code is inconsistent.

Explanation: Parts of the node machine code package are receiving unexpected results; there may be an inconsistent set of subpackages installed, or one subpackage may be damaged.

User response: Follow troubleshooting procedures to reload the machine code.

1. Follow the procedure to run a node rescue.
2. If the error occurs again, contact IBM technical support.

Possible Cause—FRUs or other:

- None

---

574   The node machine code is damaged.

Explanation: A checksum failure has indicated that the node machine code is damaged and needs to be reinstalled.

User response:

1. If the other nodes are operational, run node rescue; otherwise, install new machine code using the service assistant. Node rescue failures, as well as the repeated return of this node error after reinstallation, are symptomatic of a hardware fault with the node.

Possible Cause—FRUs or other:

- None

---

576   The cluster state and configuration data cannot be read.

Explanation: The node was unable to read the saved cluster state and configuration data from its internal drive because of a read or medium error.

User response: Follow troubleshooting procedures to fix the hardware:

1. The drive of the node canister cannot be replaced individually. Follow the hardware remove and replace instructions to change the node canister.

Possible Cause—FRUs or other:

- None

578	**The state data was not saved following a power loss.**

Explanation: On startup, the node was unable to read its state data. When this happens, it expects to be automatically added back into a clustered system. However, if it is not joined to a clustered system in 60 sec, it raises this node error. This error is a critical node error, and user action is required before the node can become a candidate to join a clustered system.

User response: Follow troubleshooting procedures to correct connectivity issues between the clustered system nodes and the quorum devices.

1. Manual intervention is required once the node reports this error.

2. Attempt to reestablish the clustered system by using other nodes. This step might involve fixing hardware issues on other nodes or fixing connectivity issues between nodes.

3. If you are able to reestablish the clustered system, remove the system data from the node that shows error 578 so it goes to a candidate state. It is then automatically added back to the clustered system.

   a. To remove the system data from the node, go to the service assistant, select the radio button for the node with a 578, click **Manage System**, and then choose **Remove System Data**.

   b. Or use the CLI command `satask leavecluster -force`.

   If the node does not automatically add back to the clustered system, note the name and I/O group of the node, and then delete the node from the clustered system configuration (if this has not already happened). Add the node back to the clustered system using the same name and I/O group.

4. If all nodes have either node error 578 or 550, follow the recommended user response for node error 550.

5. Attempt to determine what caused the nodes to shut down.

Possible Cause—FRUs or other:

• None

579	**Battery subsystem has insufficient charge to save system data**

Explanation: Not enough capacity is available from the battery subsystem to save system data in response to a series of battery and boot-drive faults.

User response: Follow troubleshooting procedures to fix hardware.

The appropriate service actions are indicated by the series of battery and boot-drive faults. Examine the event log to determine the accompanying faults. Service the other faults.

580	**The service controller ID could not be read.**

Explanation: The system cannot read the unique ID from the service controller, so the Fibre Channel adapters cannot be started.

User response: In the sequence shown, exchange the following FRUs for new FRUs.

Possible Cause-FRUs or other:

2145-CG8 or 2145-CF8

• Service controller (70%)

• Service controller cable (30%)

Service controller (100%)

Other:

• None

581	**A serial link error in the 2145 UPS-1U has occurred.**

Explanation: There is a fault in the communications cable, the serial interface in the uninterruptible power supply 2145 UPS-1U, or 2145.

User response: Check that the communications cable is correctly plugged into the 2145 and the 2145 UPS-1U. If the cable is plugged in correctly, replace the FRUs in the order shown.

Possible Cause-FRUs or other:

2145-CF8, or 2145-CG8

• 2145 power cable assembly (40%)

• 2145 UPS-1U assembly (30%)

• 2145 system board (30%)

582	**A battery error in the 2145 UPS-1U has occurred.**

Explanation: A problem has occurred with the uninterruptible power supply 2145 UPS-1U battery.

User response: Exchange the FRU for a new FRU. After replacing the battery assembly, if the 2145 UPS-1U service indicator is on, press and hold the 2145 UPS-1U Test button for three seconds to start the self-test and verify the repair. During the self-test, the rightmost four LEDs on the 2145 UPS-1U front-panel assembly flash in sequence.

Possible Cause-FRUs or other:

• UPS-1U battery assembly (50%)

• UPS-1U assembly (50%)

---

583   **An electronics error in the 2145 UPS-1U has occurred.**

Explanation: A problem has occurred with the 2145 UPS-1U electronics.

User response: Exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

- 2145 UPS-1U assembly

---

584   **The 2145 UPS-1U is overloaded.**

Explanation: A problem with output overload has been reported by the uninterruptible power supply 2145 UPS-1U. The Overload Indicator on the 2145 UPS-1U front panel is illuminated red.

User response:

1. Ensure that only one 2145 is receiving power from the 2145 UPS-1U. Also ensure that no other devices are connected to the 2145 UPS-1U.
2. Disconnect the 2145 from the 2145 UPS-1U. If the Overload Indicator is still illuminated, on the disconnected 2145 replace the 2145 UPS-1U.
3. If the Overload Indicator is now off, and the node is a 2145-CG8 or 2145-CF8, on the disconnected 2145, with all outputs disconnected, determine whether it is one of the two power supplies or the power cable assembly that must be replaced. Plug just one power cable into the left hand power supply and start the node and see whether the error is reported. Then shut down the node and connect the other power cable into the left hand power supply and start the node and see whether the error is repeated. Then repeat the two tests for the right hand power supply. If the error is repeated for both cables on one power supply but not the other, replace the power supply that showed the error; otherwise, replace the power cable assembly.

Possible Cause-FRUs or other:

- Power cable assembly (45%)
- Power supply assembly (45%)
- UPS-1U assembly (10%)

---

586   **The power supply to the 2145 UPS-1U does not meet requirements.**

Explanation: None.

User response: Follow troubleshooting procedures to fix the hardware.

---

587   **An incorrect type of uninterruptible power supply has been detected.**

Explanation: An incorrect type of 2145 UPS-1U was installed.

User response: Exchange the 2145 UPS-1U for one of the correct type.

Possible Cause-FRUs or other:

- 2145 UPS-1U (100%)

---

588   **The 2145 UPS-1U is not cabled correctly.**

Explanation: The signal cable or the 2145 power cables are probably not connected correctly. The power cable and signal cable might be connected to different 2145 UPS-1U assemblies.

User response:

1. Connect the cables correctly.
2. Restart the node.

Possible Cause-FRUs or other:

- None.

Other:

- Cabling error (100%)

---

589   **The 2145 UPS-1U ambient temperature limit has been exceeded.**

Explanation: The ambient temperature threshold for the 2145 UPS-1U has been exceeded.

User response: Reduce the temperature around the system:

1. Turn off the 2145 UPS-1U and unplug it from the power source.
2. Clear the vents and remove any heat sources.
3. Ensure that the air flow around the 2145 UPS-1U is not restricted.
4. Wait at least five minutes, and then restart the 2145 UPS-1U. If the problem remains, exchange 2145 UPS-1U assembly.

---

590   **Repetitive node transitions into standby mode from normal mode because of power subsystem-related node errors.**

Explanation: Multiple node restarts occurred because of 2145 UPS-1U errors, which can be reported on any node type

This error means that the node made the transition into standby from normal mode because of power subsystem-related node errors too many times within a short period. Too many times are defined as three, and a short period is defined as 1 hour. This error alerts the user that something might be wrong with the power subsystem as it is clearly not normal for the node to repeatedly go in and out of standby.

If the actions of the tester or engineer are expected to cause many frequent transitions from normal to standby and back, then this error does not imply that

there is any actual fault with the system.

User response: Follow troubleshooting procedures to fix the hardware:

1. Verify that the room temperature is within specified limits and that the input power is stable.
2. If a 2145 UPS-1U is connected, verify that the 2145 UPS-1U signal cable is fastened securely at both ends.
3. Look in the system event log for the node error that is repeating.

Note: The condition is reset by powering off the node from the node front panel.

---

650   The canister battery is not supported

Explanation: The canister battery shows product data that indicates it cannot be used with the code version of the canister.

User response: This is resolved by either obtaining a battery which is supported by the system's code level, or the canister's code level is updated to a level which supports the battery.

1. Remove the canister and its lid and check the FRU part number of the new battery matches that of the replaced battery. Obtain the correct FRU part if it does not.
2. If the canister has just been replaced, check the code level of the partner node canister and use the service assistant to update this canister's code level to the same level.

Possible cause—FRUs or other cause

• canister battery

---

651   The canister battery is missing

Explanation: The canister battery cannot be detected.

User response:

1. Use the remove and replace procedures to remove the node canister and its lid.
2. Use the remove and replace procedures to install a battery.
3. If a battery is present, ensure that it is fully inserted. Replace the canister.
4. If this error persists, use the remove and replace procedures to replace the battery.

Possible cause—FRUs or other cause

• Canister battery

---

652   The canister battery has failed

Explanation: The canister battery has failed. The battery may be showing an error state, it may have reached the end of life, or it may have failed to charge.

Data

Number indicators with failure reasons

• 1—battery reports a failure
• 2—end of life
• 3—failure to charge

User response:

1. Use the remove and replace procedures to replace the battery.

Possible cause—FRUs or other cause

• canister battery

---

653   The canister battery's temperature is too low

Explanation: The canister battery's temperature is below its minimum operating temperature.

User response:

• Wait for the battery to warm up, the error will clear when its minimum working temperature is reached.
• If the error persists for more than an hour when the ambient temperature is normal, use the remove and replace procedures to replace the battery.

Possible cause—FRUs or other cause

• canister battery

---

654   The canister battery's temperature is too high

Explanation: The canister battery's temperature is above its safe operating temperature.

User response:

• If necessary, reduce the ambient temperature.
• Wait for the battery to cool down, the error will clear when normal working temperature is reached. Keep checking the reported error as the system may determine the battery has failed.
• If the node error persists for more than two hours after the ambient temperature returns to the normal operating range, use the remove and replace procedures to replace the battery.

Possible cause—FRUs or other cause

• canister battery

---

655   **Canister battery communications fault.**

Explanation: The canister cannot communicate with the battery.

User response:

- Use the remove and replace procedures to replace the battery.
- If the node error persists, use the remove and replace procedures to replace the node canister.

Possible Cause-FRUs or other cause:

- Canister battery
- Node canister

656   **The canister battery has insufficient charge**

Explanation: The canister battery has insufficient charge to save the canister's state and cache data to the internal drive if power were to fail.

User response:

- Wait for the battery to charge, the battery does not need to be fully charged for the error to automatically clear.

Possible cause—FRUs or other cause

- none

657   **Not enough battery charge to support graceful shutdown of the storage enclosure.**

Explanation: Insufficient power available for the enclosure.

User response: If a battery is missing, failed or having a communication error, replace the battery.

If a battery is failed, replace the battery.

If a battery is charging, this error should go away when the battery is charged.

If a battery is too hot, the system can be started after it has cooled.

If running on a single power supply with low input power (110 V AC), "low voltage" will be seen in the extra data. If this is the case, the failed or missing power supply should be replaced. This will only happen if a single power supply is running with input power that is too low.

668   **The remote setting is not available for users for the current system.**

Explanation: On the current systems, users cannot be set to remote.

User response: Any user defined on the system must be a local user. To create a remote user the user must not be defined on the local system.

670   **The UPS battery charge is not enough to allow the node to start.**

Explanation: The uninterruptible power supply connected to the node does not have sufficient battery charge for the node to safely become active in a cluster. The node will not start until a sufficient charge exists to store the state and configuration data held in the node memory if power were to fail. The front panel of the node will show "charging".

User response: Wait for sufficient battery charge for enclosure to start:

1. Wait for the node to automatically fix the error when there is sufficient charge.
2. Ensure that no error conditions are indicated on the uninterruptible power supply.

671   **The available battery charge is not enough to allow the node canister to start. Two batteries are charging.**

Explanation: The battery charge within the enclosure is not sufficient for the node to safely become active in a cluster. The node will not start until sufficient charge exists to store the state and configuration data held in the node canister memory if power were to fail. Two batteries are within the enclosure, one in each of the power supplies. Neither of the batteries indicate an error—both are charging.

The node will start automatically when sufficient charge is available. The batteries do not have to be fully charged before the nodes can become active.

Both nodes within the enclosure share the battery charge, so both node canisters report this error. The service assistant shows the estimated start time in the node canister hardware details.

User response: Wait for the node to automatically fix the error when sufficient charge becomes available.

672   **The available battery charge is not enough to allow the node canister to start. One battery is charging.**

Explanation: The battery charge within the enclosure is not sufficient for the node to safely become active in a cluster. The node will not start until sufficient charge exists to store the state and configuration data held in the node canister memory if power were to fail. Two batteries are within the enclosure, one in each of the power supplies. Only one of the batteries is charging, so the time to reach sufficient charge will be extended.

The node will start automatically when sufficient charge is available. The batteries do not have to be fully charged before the nodes can become active.

Both nodes within the enclosure share the battery charge, so both node canisters report this error.

The service assistant shows the estimated start time,

and the battery status, in the node canister hardware details.

Possible Cause-FRUs or other:

- None

User response:

1. Wait for the node to automatically fix the error when sufficient charge becomes available.
2. If possible, determine why one battery is not charging. Use the battery status shown in the node canister hardware details and the indicator LEDs on the PSUs in the enclosure to diagnose the problem. If the issue cannot be resolved, wait until the cluster is operational and use the troubleshooting options in the management GUI to assist in resolving the issue.

Possible Cause-FRUs or other:

- Battery (33%)
- Control power supply (33%)
- Power cord (33%)

---

673   **The available battery charge is not enough to allow the node canister to start. No batteries are charging.**

Explanation: A node cannot be in active state if it does not have sufficient battery power to store configuration and cache data from memory to internal disk after a power failure. The system has determined that both batteries have failed or are missing. The problem with the batteries must be resolved to allow the system to start.

User response: Follow troubleshooting procedures to fix hardware:

1. Resolve problems in both batteries by following the procedure to determine status using the LEDs.
2. If the LEDs do not show a fault on the power supplies or batteries, power off both power supplies in the enclosure and remove the power cords. Wait 20 seconds, then replace the power cords and restore power to both power supplies. If both node canisters continue to report this error replace the enclosure chassis.

Possible Cause-FRUs or other:

- Battery (33%)
- Power supply (33%)
- Power cord (33%)
- Enclosure chassis (1%)

---

674   **The cycling mode of a Metro Mirror object cannot be changed.**

Explanation: The cycling mode may only be set for Global Mirror objects. Metro Mirror objects cannot have a cycling mode defined.

User response: The object's type must be set to 'global' before or when setting the cycling mode.

---

690   **The node is held in the service state.**

Explanation: The node is in service state and has been instructed to remain in service state. While in service state, the node will not run as part of a cluster. A node must not be in service state for longer than necessary while the cluster is online because a loss of redundancy will result. A node can be set to remain in service state either because of a service assistant user action or because the node was deleted from the cluster.

User response: When it is no longer necessary to hold the node in the service state, exit the service state to allow the node to run:

1. Use the service assistant action to release the service state.

Possible Cause—FRUs or other:

- none

---

700   **The Fibre Channel adapter that was previously present has not been detected.**

Explanation: A Fibre Channel adapter that was previously present has not been detected. For Storwize V7000, the adapter is located on the node canister system board.

This node error does not, in itself, stop the node canister from becoming active in the system; however, the Fibre Channel network might be being used to communicate between the node canisters in a clustered system. It is possible that this node error indicates why the critical node error 550 A cluster cannot be formed because of a lack of cluster resources is reported on the node canister.

Data:

- Location—A number indicating the adapter location. Location 0 indicates the adapter integrated into the system board is being reported.

User response:

1. If possible, this noncritical node error should be serviced using the management GUI and running the recommended actions for the service error code.
2. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

   There are a number of possibilities.

   a. If you have deliberately removed the adapter (possibly replacing it with a different adapter type), you will need to follow the management GUI recommended actions to mark the hardware change as intentional.

b. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause—FRUs or other cause:

- Node canister 100%

---

701   A Fibre Channel adapter has failed.

Explanation: A Fibre Channel adapter has failed. The adapter is located on the node canister system board.

This node error does not, in itself, stop the node canister becoming active in the system. However, the Fibre Channel network might be being used to communicate between the node canisters in a clustered system. Therefore, it is possible that this node error indicates the reason why the critical node error 550 A `cluster cannot be formed because of a lack of cluster resources` is reported on the node canister.

Data:

- A number indicating the adapter location. Location 0 indicates the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

- Node canister

---

702   A Fibre Channel adapter has a PCI error.

Explanation: A Fibre Channel adapter has a PCI error. The adapter is located on the node canister system board.

This node error does not, in itself, stop the node canister becoming active in the system. However, the Fibre Channel network might be being used to communicate between the node canisters in a clustered system. Therefore, it is possible that this node error indicates the reason why the critical node error 550 A `cluster cannot be formed because of a lack of cluster resources` is reported on the node canister.

Data:

- A number indicating the adapter location. Location 0 indicates the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.

2. Use the procedures to restart (physically remove and reseat) a node canister.
3. As the adapter is located on the system board, replace the node canister by using the remove and replace procedures.

Possible Cause-FRUs or other cause:

- Node canister

---

703   A Fibre Channel adapter is degraded.

Explanation: A Fibre Channel adapter is degraded. The adapter is located on the node canister system board.

This node error does not, in itself, stop the node canister becoming active in the system. However, the Fibre Channel network might be being used to communicate between the node canisters in a clustered system. Therefore, it is possible that this node error indicates the reason why the critical node error 550 A `cluster cannot be formed because of a lack of cluster resources` is reported on the node canister.

Data:

- A number indicating the adapter location. Location 0 indicates the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Use the procedures to restart (physically remove and reseat) a node canister .
3. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause FRUs or other cause:

- Node canister

---

704   Fewer Fibre Channel ports operational.

Explanation: A Fibre Channel port that was previously operational is no longer operational. The physical link is down.

This node error does not, in itself, stop the node canister becoming active in the system. However, the Fibre Channel network might be being used to communicate between the node canisters in a clustered system. Therefore, it is possible that this node error indicates the reason why the critical node error 550 A `cluster cannot be formed because of a lack of cluster resources` is reported on the node canister.

Data:

Three numeric values are listed:

- The ID of the first unexpected inactive port. This ID is a decimal number.
- The ports that are expected to be active, which is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is expected to be active.
- The ports that are actually active, which is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is active.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Possibilities:
   - If the port has been intentionally disconnected, use the management GUI recommended action for the service error code and acknowledge the intended change.
   - Check that the Fibre Channel cable is connected at both ends and is not damaged. If necessary, replace the cable.
   - Check the switch port or other device that the cable is connected to is powered and enabled in a compatible mode. Rectify any issue. The device service interface might indicate the issue.
   - Use the remove and replace procedures to replace the SFP transceiver in the Storwize V7000 and the SFP transceiver in the connected switch or device.
   - As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:
- Fibre Channel cable
- SFP transceiver
- Node canister

---

705   **Fewer Fibre Channel I/O ports operational.**

Explanation: One or more Fibre Channel I/O ports that have previously been active are now inactive. This situation has continued for one minute.

A Fibre Channel I/O port might be established on either a Fibre Channel platform port or an Ethernet platform port using FCoE. This error is expected if the associated Fibre Channel or Ethernet port is not operational.

Data:

Three numeric values are listed:
- The ID of the first unexpected inactive port. This ID is a decimal number.
- The ports that are expected to be active, which is a hexadecimal number. Each bit position represents a

port, with the least significant bit representing port 1. The bit is 1 if the port is expected to be active.
- The ports that are actually active, which is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is active.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Follow the procedure for mapping I/O ports to platform ports to determine which platform port is providing this I/O port.
3. Check for any 704 (`Fibre channel platform port not operational`) or 724 (`Ethernet platform port not operational`) node errors reported for the platform port.
4. Possibilities:
   - If the port has been intentionally disconnected, use the management GUI recommended action for the service error code and acknowledge the intended change.
   - Resolve the 704 or 724 error.
   - If this is an FCoE connection, use the information the view gives about the Fibre Channel forwarder (FCF) to troubleshoot the connection between the port and the FCF.

Possible Cause-FRUs or other cause:
- None

---

706   **Fibre Channel clustered system path failure.**

Explanation: One or more Fibre Channel (FC) input/output (I/O) ports that have previously been able to see all required online node canisters can no longer see them. This situation has continued for 5 minutes. This error is not reported unless a node is active in a clustered system.

A Fibre Channel I/O port might be established on either a FC platform port or an Ethernet platform port using Fiber Channel over Ethernet (FCoE).

Data:

Three numeric values are listed:
- The ID of the first FC I/O port that does not have connectivity. This is a decimal number.
- The ports that are expected to have connections. This is a hexadecimal number, and each bit position represents a port - with the least significant bit representing port 1. The bit is 1 if the port is expected to have a connection to all online node canisters.
- The ports that actually have connections. This is a hexadecimal number, each bit position represents a

port, with the least significant bit representing port 1. The bit is 1 if the port has a connection to all online nodes.

User response:

1. If possible, this noncritical node error should be serviced using the management GUI and running the recommended actions for the service error code.

2. Follow the procedure: Mapping I/O ports to platform ports to determine which platform port does not have connectivity.

3. There are a number of possibilities.
   * If the port's connectivity has been intentionally reconfigured, use the management GUI recommended action for the service error code and acknowledge the intended change. You must have at least two I/O ports with connections to all other node canisters, except the node canisters in the same enclosure.
   * Resolve other node errors relating to this platform port or I/O port.
   * Check that the SAN zoning is correct.

Possible Cause: FRUs or other cause:

* None.

---

710   The SAS adapter that was previously present has not been detected.

Explanation: A SAS adapter that was previously present has not been detected. The adapter is located on the node canister system board.

Data:

* A number indicating the adapter location. Location 0 indicates the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.

2. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

* Node canister

---

711   A SAS adapter has failed.

Explanation: A SAS adapter has failed. The adapter is located on the node canister system board.

Data:

* A number indicating the adapter location. Location 0 indicates that the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.

2. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

* Node canister

---

712   A SAS adapter has a PCI error.

Explanation: A SAS adapter has a PCI error. The adapter is located on the node canister system board.

Data:

* A number indicating the adapter location. Location 0 indicates the adapter that is integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.

2. Use the procedures to restart (physically remove and reseat) a node canister.

3. Locate the adapter on the system board and replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

* Node canister

---

713   A SAS adapter is degraded.

Explanation: A SAS adapter is degraded. The adapter is located on the node canister system board.

Data:

* A number indicating the adapter location. Location 0 indicates that the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.

2. Use the procedures to restart (physically remove and reseat) a node canister.

3. Locate the adapter on the system board and replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

* Node canister

---

715   Fewer SAS host ports operational

Explanation: A SAS port that was previously operational is no longer operational. The physical link is down.

Data:

Three numeric values are listed:

- The ID of the first unexpected inactive port. This ID is a decimal number.
- The ports that are expected to be active, which is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is expected to be active.
- The ports that are actually active, which is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is active.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Possibilities:
   - If the port has been intentionally disconnected, use the management GUI recommended action for the service error code and acknowledge the intended change.
   - Check that the SAS cable is connected at both ends and is not damaged. If necessary, replace the cable.
   - Check the switch port or other device that the cable is connected to is powered and enabled in a compatible mode. Rectify any issue. The device service interface might indicate the issue.
   - Use the remove and replace procedures to replace the adapter.

Possible Cause-FRUs or other cause:

- SAS cable
- SAS adapter

720   Ethernet adapter that was previously present has not been detected.

Explanation: An Ethernet adapter that was previously present has not been detected. The adapters form a part of the canister assembly.

Data:

- A number indicating the adapter location. The location indicates an adapter slot. See the node canister description for the definition of the adapter slot locations. If the location is 0, the adapter is integrated into the system board or directly connected to it, that is, not in a PCI express expansion slot.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Because the Ethernet adapters are integrated into the node canisters, replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

- Node canister

721   An Ethernet adapter has failed.

Explanation: An Ethernet adapter failed. The adapters form part of the canister assembly.

Data:

-
   A number indicating the adapter location. The location indicates an adapter slot. See the node canister description for the definition of the adapter slot locations. If the location is 0, the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. As the Ethernet adapters are integrated into the node canisters, use the remove and replace procedures to replace the node canister.

Possible Cause—FRUs or other cause:

- Node canister

722   An Ethernet adapter has a PCI error.

Explanation: An Ethernet adapter has a PCI error. The adapters form part of the canister assembly.

Data:

-
   A number indicating the adapter location. The location indicates an adapter slot. See the node canister description for the definition of the adapter slot locations. If the location is 0, the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. As the Ethernet adapters are integrated into the node canisters, use the remove and replace procedures to replace the node canister.

Possible Cause—FRUs or other cause:

- Node canister

723   An Ethernet adapter is degraded.

Explanation: An Ethernet adapter is degraded. The adapters form part of the canister assembly.

Data:

*

A number indicating the adapter location. The location indicates an adapter slot. See the node canister description for the definition of the adapter slot locations. If the location is 0, the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.

2. As the Ethernet adapters are integrated into the node canisters, use the remove and replace procedures to replace the node canister.

Possible Cause—FRUs or other cause:

* Node canister

---

724   Fewer Ethernet ports active.

Explanation: An Ethernet port that was previously operational is no longer operational. The physical link is down.

Data:

Three numeric values are listed:

* The ID of the first unexpected inactive port. This is a decimal number.

* The ports that are expected to be active. This is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is expected to be active.

* The ports that are actually active. This is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is active.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.

2. Possibilities:

    a. If the port has been intentionally disconnected, use the management GUI recommended action for the service error code and acknowledge the intended change.

    b. Make sure the Ethernet cable is connected at both ends and is undamaged. If necessary, replace the cable.

    c. Make sure the switch port or other device the cable is connected to is powered and enabled in a compatible mode. Rectify any issue. The device service interface might indicate the issue.

    d. If this is a 10 Gb/s port, use the remove and replace procedures to replace the SFP transceiver in the Storwize V7000 and the SFP transceiver in the connected switch or device.

    e. Replace the node canister using the remove and replace procedures.

Possible Cause—FRUs or other cause:

* Ethernet cable
* Ethernet SFP transceiver
* Node canister

---

730   The bus adapter has not been detected.

Explanation: The bus adapter that connects the canister to the enclosure midplane has not been detected.

This node error does not, in itself, stop the node canister becoming active in the system. However, the bus might be being used to communicate between the node canisters in a clustered system. Therefore, it is possible that this node error indicates the reason why the critical node error 550 A cluster cannot be formed because of a lack of cluster resources is reported on the node canister.

Data:

* A number indicating the adapter location. Location 0 indicates that the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.

2. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

* Node canister

---

731   The bus adapter has failed.

Explanation: The bus adapter that connects the canister to the enclosure midplane has failed.

This node error does not, in itself, stop the node canister becoming active in the system. However, the bus might be being used to communicate between the node canisters in a clustered system. Therefore, it is possible that this node error indicates the reason why the critical node error 550 A cluster cannot be formed because of a lack of cluster resources is reported on the node canister.

Data:

* A number indicating the adapter location. Location 0 indicates that the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.

2. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

- Node canister

---

732   **The bus adapter has a PCI error.**

Explanation: The bus adapter that connects the canister to the enclosure midplane has a PCI error.

This node error does not, in itself, stop the node canister becoming active in the system. However, the bus might be being used to communicate between the node canisters in a clustered system; therefore it is possible that this node error indicates the reason why the critical node error 550 A cluster cannot be formed because of a lack of cluster resources is reported on the node canister.

Data:

- A number indicating the adapter location. Location 0 indicates that the adapter integrated into the system board is being reported.

User response:

1. If possible, this noncritical node error should be serviced using the management GUI and running the recommended actions for the service error code.

2. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

- Node canister

---

733   **The bus adapter degraded.**

Explanation: The bus adapter that connects the canister to the enclosure midplane is degraded.

This node error does not, in itself, stop the node canister from becoming active in the system. However, the bus might be being used to communicate between the node canisters in a clustered system. Therefore, it is possible that this node error indicates the reason why the critical node error 550 A cluster cannot be formed because of a lack of cluster resources is reported on the node canister.

Data:

- A number indicating the adapter location. Location 0 indicates that the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.

2. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

- Node canister

---

734   **Fewer bus ports.**

Explanation: One or more PCI bus ports that have previously been active are now inactive. This condition has existed for over one minute. That is, the internode link has been down at the protocol level.

This could be a link issue but is more likely caused by the partner node unexpectedly failing to respond.

Data:

Three numeric values are listed:

- The ID of the first unexpected inactive port. This is a decimal number.

- The ports that are expected to be active. This is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is expected to be active.

- The ports that are actually active. This is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is active.

User response:

1. If possible, this noncritical node error should be serviced using the management GUI and running the recommended actions for the service error code.

2. Follow the procedure for getting node canister and clustered-system information and determine the state of the partner node canister in the enclosure. Fix any errors reported on the partner node canister.

3. Use the remove and replace procedures to replace the enclosure.

Possible Cause-FRUs or other cause:

- Node canister
- Enclosure midplane

---

736   **The temperature of a device on the system board is greater than or equal to the warning threshold.**

Explanation: The temperature of a device on the system board is greater than or equal to the warning threshold.

User response: Check for external and internal air flow blockages or damage.

1. Remove the top of the machine case and check for missing baffles, damaged heat sinks, or internal blockages.
2. If problem persists, replace the system board.

Possible Cause-FRUs or other:

- System board

---

737   **The temperature of a power supply is greater than or equal to the warning or critical threshold.**

Explanation: The temperature of a power supply is greater than or equal to the warning or critical threshold.

User response: Check for external and internal air flow blockages or damage.

1. Remove the top of the machine case and check for missing baffles, damaged heat sinks, or internal blockages.
2. If the problem persists, replace the power supply.

Possible Cause-FRUs or other:

- Power supply

---

738   **The temperature of a PCI riser card is greater than or equal to the warning threshold.**

Explanation: The temperature of a PCI riser card is greater than or equal to the warning threshold.

User response: Check for external and internal air flow blockages or damage.

1. Remove the top of the machine case and check for missing PCI riser card 2, missing baffles, or internal blockages.
2. Check all of the PCI cards plugged into the riser that is identified by the extra data to find if any are faulty, and replace as necessary.
3. If the problem persists, replace the PCI riser.

Possible Cause-FRUs or other:

- PCI riser

---

740   **The command failed because of a wiring error described in the event log.**

Explanation: It is dangerous to exclude a sas port while the topology is invalid, so we forbid the user from attempting it to avoid any potential loss of data access.

User response: Correct the topology, then retry the command.

---

741   **CPU missing**

Explanation: A CPU that was previously present has not been detected. The CPU might not be correctly installed or it might have failed.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Select one of the following actions:
   - If removing the CPU was deliberate, follow the management GUI recommended actions to mark the hardware change as intentional.
   - If it is not possible to isolate the problem, use the remove and replace procedures to replace the CPU.
   - Replace the system board.

---

743   **A boot drive is offline, missing, out of sync, or the persistent data is not usable.**

Explanation: A boot drive is offline, missing, out of sync, or the persistent data is not usable.

User response: Look at a boot drive view to determine the problem.

1. If slot status is `out of sync`, then re-sync the boot drives by running the command **satask chbootdrive**.
2. If slot status is `missing`, then put the original drive back in this slot or install a FRU drive.
3. If slot status is `failed`, then replace the drive.

Possible Cause-FRUs or other:

- Boot drive

---

744   **A boot drive is in the wrong location.**

Explanation: A boot drive is in the wrong slot or comes from another node.

User response: Look at a boot drive view to determine the problem.

1. Replace the boot drive with the correct drive and put this drive back in the node that it came from.
2. Sync the boot drive if you choose to use it in this node.

Possible Cause-FRUs or other:

- None

---

745   **A boot drive is in an unsupported slot.**

Explanation: A boot drive is in an unsupported slot. This means that at least one of the first two drives are online and at least one invalid slot (3-8) is occupied.

User response: Look at a boot drive view to

determine which invalid slot(s) are occupied and remove the drive(s).

Possible Cause-FRUs or other:

- None

---

746   **Technician port connection invalid.**

Explanation: The code has detected more than one MAC address though the connection, or the DHCP has given out more than one address. The code thus believes there is a switch attached.

User response:

1. Plug a cable from the technician port to a switch, and plug 2 or more machines into that switch. They must have IP addresses in the range 192.168.0.1 - 192.168.0.30

2. Request a DHCP lease to trigger the detection.

---

747   **The Technician port is being used.**

Explanation: The Technician port is active and being used

User response: No service action is required. Use the workstation to configure the node.

---

748   **The technician port is enabled.**

Explanation: The technician port is enabled initially for easy configuration, and then disabled, so that the port can be used for iSCSI connection. When all connectivity to the node fails, the technician port can be reenabled for emergency use but must not remain enabled. This event is to remind you to disable the technician port. While the technician port is enabled, do not connect it to the LAN/SAN.

User response: Complete the following step to resolve this problem.

1. Turn off technician port by using the following CLI command:

   satask chserviceip -techport *disable*

Possible Cause-FRUs or other:

- N/A

---

750   **Compression accelerator missing**

Explanation: A compression adapter that was previously present was not detected.

User response:

1. Use the **svcinfo lsnodehw** command to review the hardware on the node indicated by this event.

2. If all missing and changed hardware is as expected, use the **chnodehw** command to accept the current node hardware configuration.

3. Otherwise, complete each of the following steps in turn until the event automatically marks as fixed:

   a. Shut down the node. Ensure the correct hardware is installed in its correct location. Reseat any hardware that are indicated as missing. Bring the node back online. Go back to step 1.

   b. Shut down the node. Replace any hardware that is indicated as missing. Bring the node back online. Go back to step 1.

   c. Shut down the node. Replace the system board or canister. Bring the node back online. Go back to step 1.

---

751   **Compression accelerator failed**

Explanation: A compression adapter has failed.

User response:

1. Shut down the node.

2. Replace the adapter in the slot indicated by the event log with a new adapter of the same type.

   Note: For the Storwize V7000 Gen2, the two compression cards share the same location.

3. Bring the node back online.

4. If the error does not auto-fix, shut down the node and replace the system board or canister, then bring the node back online.

---

766   **CMOS battery failure.**

Explanation: CMOS battery failure.

User response: Replace the CMOS battery.

Possible Cause-FRUs or other:

- CMOS battery

---

768   **Ambient temperature warning.**

Explanation: Data:

- A text string identifying the thermal sensor reporting the warning level and the current temperature in degrees (Celsius).

User response: Possible Cause-FRUs or other cause:

- None

---

769   **CPU temperature warning.**

Explanation: Data:

- A text string identifying the thermal sensor reporting the warning level and the current temperature in degrees (Celsius).

User response: Possible Cause—FRUs or other cause:

- CPU

---

770 Shutdown temperature reached

Explanation: The node temperature has reached the point at which it is must shut down to protect electronics and data. This is most likely an ambient temperature problem, but it could be a hardware issue.

Data:

- A text string identifying the thermal sensor reporting the warning level and the current temperature in degrees (Celsius).

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Check the temperature of the room and correct any air conditioning or ventilation problems.
3. Check the airflow around the system and make sure no vents are blocked.

Possible Cause-FRUs or other cause:

- CPU

775 Power supply problem.

Explanation: A power supply has a fault condition.

User response: Replace the power supply.

Possible Cause-FRUs or other:

- Power supply

776 Power supply mains cable unplugged.

Explanation: A power supply mains cable is not plugged in.

User response: Plug in power supply mains cable.

Possible Cause-FRUs or other:

- None

777 Power supply missing.

Explanation: A power supply is missing.

User response: Install power supply.

Possible Cause-FRUs or other:

- Power supply

779 Battery is missing

Explanation: The battery is not installed in the system.

User response: Install the battery.

You can power up the system without the battery installed.

Possible Cause-FRUs or other:

- Battery (100%)

780 Battery has failed

Explanation:

1. The battery has failed.
2. The battery is past the end of its useful life.
3. The battery failed to provide power on a previous occasion and is therefore, regarded as unfit for its purpose.

User response: Replace the battery.

Possible Cause-FRUs or other:

- Battery (100%)

781 Battery is below the minimum operating temperature

Explanation: The battery cannot perform the required function because it is below the minimum operating temperature.

This error is reported only if the battery subsystem cannot provide full protection.

An inability to charge is not reported if the combined charge available from all installed batteries can provide full protection at the current charge levels.

User response: No service action required, use the console to manage the node.

Wait for the battery to warm up.

782 Battery is above the maximum operating temperature

Explanation: The battery cannot perform the required function because it is above the maximum operating temperature.

This error is reported only if the battery subsystem cannot provide full protection.

An inability to charge is not reported if the combined charge available from all installed batteries can provide full protection at the current charge levels.

User response: No service action required, use the console to manage the node.

Wait for the battery to cool down.

783 Battery communications error

Explanation: A battery is installed, but communications via I2C are not functioning.

This might be either a fault in the battery unit or a fault in the battery backplane.

User response: No service action required, use the console to manage the node.

Replace the battery. If the problem persists, conduct the corrective service procedure described in "1109" on page 344.

---

784   **Battery is nearing end of life**

Explanation: The battery is near the end of its useful life. You should replace it at the earliest convenient opportunity.

This might be either a fault in the battery unit or a fault in the battery backplane.

User response: No service action required, use the console to manage the node.

Replace the battery.

---

785   **Battery capacity is reduced because of cell imbalance**

Explanation: The charge levels of the cells within the battery pack are out of balance.

Some cells become fully charged before others, which causes charging to terminate early, before the entire battery pack is fully charged.

Ending recharging prematurely effectively reduces the available capacity of the pack.

Circuitry within the battery pack corrects such errors normally, but can take tens of hours to complete.

If this error is not fixed after 24 hours, or if the error reoccurs after it fixes itself, the error is likely indicative of a problem in the battery cells. In such a case, replace the battery pack.

User response: No service action required, use the console to manage the node.

Wait for the cells to balance.

---

786   **Battery VPD checksum error**

Explanation: The checksum on the vital product data (VPD) stored in the battery EEPROM is incorrect.

User response: No service action required, use the console to manage the node.

Replace the battery.

---

787   **Battery is at a hardware revision level not supported by the current code level**

Explanation: The battery currently installed is at a hardware revision level that is not supported by the current code level.

User response: No service action required, use the console to manage the node.

Either update the code level to one that supports the currently installed battery or replace the battery with one that is supported by the current code level.

---

803   **Fibre Channel adapter not working**

Explanation: A problem has been detected on the node's Fibre Channel (FC) adapter.

User response: None.

User response: Follow troubleshooting procedures to fix the hardware.

1. If possible, use the management GUI to run the recommended actions for the associated service error code.

Possible Cause-FRUs or other cause:
- None

---

818   **Unable to recover the service controller flash disk.**

Explanation: A nonrecoverable error occurred when accessing the service controller persistent memory.

User response:
1. Restart the node and see if it recovers.
2. Replace the field replaceable units (FRUs) in the order listed.

Possible Cause-FRUs or other cause:
- Service controller
- Service controller cable

---

820   **Canister type is incompatible with enclosure model**

Explanation: The node canister has detected that it has a hardware type that is not compatible with the control enclosure MTM, such as node canister type 300 in an enclosure with MTM 2076-112.

This is an expected condition when a control enclosure is being updated to a different type of node canister.

User response:
1. Check that all the update instructions have been followed completely.
2. Use the management GUI to run the recommended actions for the associated service error code.

Possible Cause-FRUs or other cause:
- None

---

830   **Encryption key required.**

Explanation: It is necessary to provide an encryption key before the system can become fully operational. This node error occurs when a system with encryption enabled is restarted without an encryption key available.

User response: Insert a USB flash drive containing a valid key into one of the node canisters.

831   **Encryption key is not valid.**

Explanation: It is necessary to provide an encryption key before the system can become fully operational. This node error occurs when the encryption key identified is invalid. A file with the correct name was found but the key in the file is corrupt.

This node error will clear once the USB flash drive containing the invalid key is removed.

User response: Remove the USB flash drive from the port.

832   **Encryption key file not found.**

Explanation: A USB flash drive containing an encryption key is present but he expected file cannot be located. This can occur if a key for a different system or an old key for this system has been provided .

Additionally, other user-created files that match the key file name format can cause this error if the USB flash drive does not contain the expected key.

This node error will clear when the USB flash drive identified has been removed.

User response: Remove the USB flash drive from the port.

833   **Unsupported USB device.**

Explanation: An unsupported device has been connected to a USB port.

Only USB flash drives are supported and this node error will be raised if another type of device is connected to a USB port.

User response: Remove the unsupported device.

836   **Encryption key required**

Explanation: It is necessary to provide an encryption key before the system can become fully operational. This error occurs when a system with encryption enabled is restarted without an encryption key available.

User response: Connect a key server that contains the current key for this system to one or more of the nodes.

840   **Unsupported hardware change detected.**

Explanation:

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.

2. If the hardware configuration is unexpectedly reduced, make sure the component has not been unseated. Hardware replacement might be necessary.

If the hardware detected does not match the expected configuration, replace the hardware component that is reported incorrectly.

Possible Cause-FRUs or other cause:

• One of the optional hardware components might require replacement

841   **Supported hardware change detected.**

Explanation:

User response:

1. Use the management GUI to run the recommended actions for the associated service error code. Use the directed maintenance to accept or reject the new configuration.

842   **Fibre Channel IO port mapping failed**

Explanation: A Fibre Channel or Fibre Channel over Ethernet port is installed but is not included in the Fibre Channel I/O port mapping, and so the port cannot be used for Fibre Channel I/O. This error is raised in one of the following situations:

• A node hardware installation
• A change of I/O adapters
• The application of an incorrect Fibre Channel port map
• A node canister upgrade that combined old and new I/O adapters that could not be mapped automatically

These tasks are normally performed by service representatives.

User response: Your service representative can use the Service Assistant to modify the Fibre Channel I/O port mappings to include all the installed ports capable of Fibre Channel I/O. The following command is used:

satask chvpd -fcportmap

850   **The canister battery is reaching the end of its useful life.**

Explanation: The canister battery is reaching the end of its useful life. It should be replaced within a week of the node error first being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.

2. Replace the node canister battery by using the remove and replace procedures.

Possible Cause-FRUs or other cause:

- Canister battery

---

860   **Fibre Channel network fabric is too big.**

Explanation: The number of Fibre Channel (FC) logins made to the node canister exceeds the allowed limit. The node canister continues to operate, but only communicates with the logins made before the limit was reached. The order in which other devices log into the node canister cannot be determined, so the node canister's FC connectivity might vary after each restart. The connection might be with host systems, other storage systems, or with other node canisters.

This error might be the reason the node canister is unable to participate in a system.

The number of allowed logins per node is 1024.

Data:

- None

User response: This error indicates a problem with the Fibre Channel fabric configuration. It is resolved by reconfiguring the FC switch:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Rezone the FC network so only the ports the node canister needs to connect to are visible to it.

Possible Cause-FRUs or other cause:

- None

---

870   **Too many cluster creations made on node**

Explanation: Too many Storwize V7000 clustered systems have been created on this node. The count of the number of clustered systems created on the node is stored within the node service controller.

Data:

- None

User response:

1. Try to create the clustered system on a different node.
2. Replace the service controller using the remove and replace procedures.

Possible cause-FRUs or other cause:

- Service controller

---

871   **Failed to increment cluster ID**

Explanation: The clustered system create option failed because the clustered system, which is stored in the service controller, could not be updated.

Data:

- None

User response:

1. Try to create the clustered system on a different node.
2. Replace the service controller using the remove and replace procedures.

Possible cause-FRUs or other cause:

- Service controller

---

875   **Request to cluster rejected.**

Explanation: A candidate node canister could not be added to the clustered system. The node canister contains hardware or firmware that is not supported in the clustered system.

Data:

This node error and extra data is viewable through **sainfo lsservicestatus** on the candidate node only. The extra data lists a full set of feature codes that are required by the node canister to run in the clustered system.

User response:

- Choose a different candidate that is compatible with the clustered system.
- Update the clustered system to code that is supported by all components.
- Do not add a candidate to the clustered system.
- Where applicable, remove and replace the hardware that is preventing the candidate from joining the clustered system.

Possible Cause—FRUs or other cause.

For information on feature codes available, see the SAN Volume Controller and Storwize family Characteristic Interoperability Matrix on the support website: www.ibm.com/support.

---

878   **Attempting recovery after loss of state data.**

Explanation: During startup, the node canister cannot read its state data. It reports this error while waiting to be added back into a clustered system. If the node canister is not added back into a clustered system within a set time, node error 578 is reported.

User response:

1. Allow time for recovery. No further action is required.
2. Keep monitoring in case the error changes to error code 578.

---

888   Too many Fibre Channel logins between nodes.

Explanation: The system has determined that the user has zoned the fabric such that this node has received more than 16 unmasked logins originating from another node or node canister - this can be any non service mode node or canister in the local cluster or in a remote cluster with a partnership. An unmasked login is from a port whose corresponding bit in the FC port mask is '1'. If the error is raised against a node in the local cluster, then it is the local FC port mask that is applied. If the error is raised against a node in a remote cluster, then it is the partner FC port masks from both clusters that apply.

More than 16 logins is not a supported configuration as it increases internode communication and can affect bandwidth and performance. For example, if node A has 8 ports and node B has 8 ports where the nodes are in different clusters, if node A has a partner FC port mask of 00000011 and node B has a partner FC port mask of 11000000 there are 4 unmasked logins possible (1,7 1,8 2,7 2,8). Fabric zoning may be used to reduce this amount further, i.e. if node B port 8 is removed from the zone there are only 2 (1,7 and 2,7). The combination of masks and zoning must leave 16 or fewer possible logins.

Note: This count includes both FC and Fibre Channel over Ethernet (FCoE) logins. The log-in count will not include masked ports.
When this event is logged. the cluster id and node id of the first node whose logins exceed this limit on the local node will be reported, as well as the WWNN of said node. If logins change, the error is automatically fixed and another error is logged if appropriate (this may or may not choose the same node to report in the sense data if the same node is still over the maximum allowed).

Data

Text string showing
- WWNN of the other node
- Cluster ID of other node
- Arbitrary node ID of one other node that is logged into this node. (node ID as it appears in `lsnode`)

User response: The error is resolved by either re-configuring the system to change which type of connection is allowed on a port, or by changing the SAN fabric configuration so ports are not in the same zone. A combination of both options may be used.

The system reconfiguration is to change the Fibre Channel ports mask to reduce which ports can be used for internode communication.

The local Fibre Channel port mask should be modified if the cluster id reported matches the cluster id of the node logging the error.

The partner Fibre Channel port mask should be modified if the cluster id reported does not match the cluster id of the node logging the error. The partner Fibre Channel port mask may need to be changed for one or both clusters.

SAN fabric configuration is set using the switch configuration utilities.

Use the `lsfabric` command to view the current number of logins between nodes.

Possible Cause-FRUs or other cause:
- None

Service error code

1801

---

889   Failed to create remote IP connection.

Explanation: Despite a request to create a remote IP partnership port connection, the action has failed or timed out.

User response: Fix the remote IP link so that traffic can flow correctly. Once the connection is made, the error will auto-correct.

---

920   Unable to perform cluster recovery because of a lack of cluster resources.

Explanation: The node is looking for a quorum of resources which also require cluster recovery.

User response: Contact IBM technical support.

---

921   Unable to perform cluster recovery because of a lack of cluster resources.

Explanation: The node does not have sufficient connectivity to other nodes or quorum device to form a cluster. If a disaster has occurred and the nodes at the other site cannot be recovered, then it is possible to allow the nodes at the surviving site to form a system using local storage.

User response: Repair the fabric or quorum device to establish connectivity. As a last resort when the nodes at the other site cannot be recovered, then it is possible to allow the nodes at the surviving site to form a system using local site storage as described below:

To avoid data corruption ensure that all host servers that were previously accessing the system have had all volumes un-mounted or have been rebooted. Ensure that the nodes at the other site are not operational and are unable to form a system in the future.

After invoking this command a full re-synchronization of all mirrored volumes will be performed when the other site is recovered. This is likely to take many hours or days to complete.

Contact IBM support personnel if you are unsure.

Note: Before continuing confirm that you have taken the following actions - failure to perform these actions can lead to data corruption that will be undetected by the system but will affect host applications.

1. All host servers that were previously accessing the system have had all volumes un-mounted or have been rebooted.

2. Ensure that the nodes at the other site are not operating as a system and actions have been taken to prevent them from forming a system in the future.

After these actions have been taken the **satask overridequorum** can be used to allow the nodes at the surviving site to form a system using local storage.

---

950   **Special update mode.**

Explanation: Special update mode.

User response: None.

---

990   **Cluster recovery has failed.**

Explanation: Cluster recovery has failed.

User response: Contact IBM technical support.

1001   **Automatic cluster recovery has run.**

Explanation: All cluster configuration commands are blocked.

User response: Call your software support center.

Caution: You can unblock the configuration commands through the cluster GUI, but you must first consult with your software support to avoid corrupting your cluster configuration.

Possible Cause-FRUs or other:

*   None

1002   **Event log full.**

Explanation: Event log full.

User response: To fix the errors in the event log, go to the start MAP.

Possible Cause-FRUs or other:

*   Unfixed errors in the log.

1007   **Canister to canister communication error.**

Explanation: A canister to canister communication error can appear when one canister cannot communicate with the other.

User response: Reseat the passive canister, and then

try reseating the active canister. If neither resolve the alert, try replacing the passive canister, and then the other canister.

A canister can be safely reseated or replaced while the system is in production. Make sure that the other canister is the active node before removing this canister. It is preferable that this canister shuts down completely before removing it, but it is not required.

1. Reseat the passive canister (a failover is not required).

2. Reseat the second canister (a failover is required).

3. If necessary, replace the passive canister (a failover is not required).

4. If necessary, replace the active canister (a failover is required).

    If a second new canister is not available, the previously removed canister can be used, as it apparently is not at fault.

5. An enclosure replacement might be necessary. Contact IBM support.

Possible Cause-FRUs or other:

Canister (95%)

Enclosure (5%)

1009   **DIMMs are incorrectly installed.**

Explanation: DIMMs are incorrectly installed.

User response: Ensure that memory DIMMs are spread evenly across all memory channels.

1. Shut down the node.

2. Ensure that memory DIMMs are spread evenly across all memory channels.

3. Restart the node.

4. If the error persists, replace system board.

Possible Cause-FRUs or other:

*   None

1011   **Fibre Channel adapter (4 port) in slot 1 is missing.**

Explanation: Fibre Channel adapter (4 port) in slot 1 is missing.

User response:

1. In the sequence shown, exchange the FRUs for new FRUs.

2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired as "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.

3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-CG8 or 2145-CF8

- 4-port Fibre Channel host bus adapter (98%)
- System board (2%)

1013   **Fibre Channel adapter (4-port) in slot 1 PCI fault.**

Explanation: Fibre Channel adapter (4-port) in slot 1 PCI fault.

User response:

1. In the sequence shown, exchange the FRUs for new FRUs.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired as "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-CG8 or 2145-CF8

- 4-port Fibre Channel host bus adapter (98%)
- System board (2%)

1014   **Fibre Channel adapter in slot 1 is missing.**

Explanation: The Fibre Channel adapter in slot 1 is missing.

User response:

1. In the sequence that is shown in the log, replace any failing FRUs with new FRUs.
2. Check node status:
   - If all nodes show a status of **online**, mark the error as **fixed**.
   - If any nodes do not show a status of **online**, go to the start MAP.
   - If you return to this step, contact your support center to resolve the problem with the node.
3. Go to the repair verification MAP.

Possible Cause, FRUs, or other:

- N/A

1015   **Fibre Channel adapter in slot 2 is missing.**

Explanation: Fibre Channel adapter in slot 2 is missing.

User response:

1. In the sequence that is shown in the log, replace any failing FRUs for new FRUs.

2. Check the node status:
   - If all nodes show a status of **online**, mark the error as **fixed**.
   - If any node does not show a status of **online**, go to the start MAP.
   - If you return to this step, contact your support center to resolve the problem with the node.
3. Go to the repair verification MAP.

Possible Cause, FRUs, or other:
- N/A

1016   **Fibre Channel adapter (4 port) in slot 2 is missing.**

Explanation: The four-port Fibre Channel adapter in PCI slot 2 is missing.

User response:

1. In the sequence that is shown in the log, replace any failing FRUs with new FRUs.
2. Check node status:
   - If all nodes show a status of **online**, mark the error as **fixed**.
   - If any nodes do not show a status of **online**, go to the start MAP.
   - If you return to this step, contact your support center to resolve the problem with the node.
3. Go to the repair verification MAP.

Possible Cause, FRUs, or other:
- Fibre Channel host bus adapter (90%)
- PCI riser card (5%)
- Other (5%)

1017   **Fibre Channel adapter in slot 1 PCI bus error.**

Explanation: The Fibre Channel adapter in PCI slot 1 is failing with a PCI bus error.

User response:

1. In the sequence that is shown in the log, replace any failing FRUs with new FRUs.
2. Check node status:
   - If all nodes show a status of **online**, mark the error as **fixed**.
   - If any nodes do not show a status of **online**, go to the start MAP.
   - If you return to this step, contact your support center to resolve the problem with the node.
3. Go to the repair verification MAP.

Possible Cause, FRUs, or other:
- Fibre Channel host bus adapter (80%)
- PCI riser card (10%)

• Other (10%)

1018   Fibre Channel adapter in slot 2 PCI fault.

Explanation: The Fibre Channel adapter in slot 2 is failing with a PCI fault.

User response:
1. In the sequence that is shown in the log, replace any failing FRUs with new FRUs.
2. Check node status:
   • If all nodes show a status of **online**, mark the error as **fixed**.
   • If any nodes do not show a status of **online**, go to the start MAP.
   • If you return to this step, contact your support center to resolve the problem with the node.
3. Go to the repair verification MAP.

Possible Cause, FRUs, or other:
• Dual port Fibre Channel host bus adapter - full height (80%)
• PCI riser card (10%)
• Other (10%)

1019   Fibre Channel adapter (four-port) in slot 2 PCI fault.

Explanation: The four-port Fibre Channel adapter in slot 2 is failing with a PCI fault.

User response:
1. In the sequence that is shown in the log, replace any failing FRUs with new FRUs.
2. Check node status:
   • If all nodes show a status of **online**, mark the error as **fixed**.
   • If any nodes do not show a status of **online**, go to the start MAP.
   • If you return to this step, contact your support center to resolve the problem with the node.
3. Go to the repair verification MAP.

Possible Cause, FRUs, or other:
• Four-port Fibre Channel host bus adapter (80%)
• PCI Express riser card (10%)
• Other (10%)

1020   The system board service processor has failed.

Explanation: The cluster is reporting that a node is not operational because of critical node error 522. See the details of node error 522 for more information.

User response: See node error 522.

1021   Incorrect enclosure

Explanation: The cluster is reporting that a node is not operational because of critical node error 500. See the details of node error 500 for more information.

User response: See node error 500.

1022   The detected memory size does not match the expected memory size.

Explanation: The cluster is reporting that a node is not operational because of critical node error 510. See the details of node error 510 for more information.

User response: See node error 510.

1024   CPU is broken or missing.

Explanation: CPU is broken or missing.

User response: Review the node hardware using the **svcinfo lsnodehw** command on the node indicated by this event.
1. Shutdown the node. Replace the CPU that is broken as indicated by the light path and event data.
2. If error persist, replace system board.

Note: Intentional removal is not permitted on a clustered node. To use the node with only one processor, you must **rmnode**, and then **readd**. Otherwise, shutdown the node and replace the processor that was removed.
Possible Cause-FRUs or other:
• CPU (80%)
• System board (20%)

1025   Processor missing

Explanation: The system assembly is failing.

User response:
1. Go to the light path diagnostic MAP and complete the light path diagnostic procedures.
2. If the light path diagnostic procedure isolates the FRU, mark this error as **fixed**. Then, go to the repair verification MAP.
3. If you replace a FRU, but it does not correct the problem, ensure that the FRU is installed correctly. Then, go to the next step.
4. Replace the system board as indicated in the Possible Cause list.
5. Check the node status:
   • If all nodes show a status of **online**, mark the error as **fixed**.
   • If any nodes do not show a status of **online**, go to the start MAP.
   • If you return to this step, contact your support center to resolve the problem with the node.

6. Go to the repair verification MAP.

Possible Cause, FRUs, or other:

2145-CF8 or 2145-CG8

- The FRUs that are indicated by the Light path diagnostics (98%)
- System board (2%)

1026   **System board device problem.**

Explanation: System board device problem.

User response: The action depends on the extra data that is provided with the node error and the light path diagnostics.

Possible Cause-FRUs or other:

- Variable

1027   **Unable to update BIOS settings.**

Explanation: The cluster is reporting that a node is not operational because of critical node error 524. See the details of node error 524 for more information.

User response: See node error 524.

1028   **System board service processor failed.**

Explanation: System board service processor failed.

User response: Complete the following steps:
1. Shut down the node.
2. Remove the main power cable.
3. Wait for the lights to stop flashing.
4. Plug in the power cable.
5. Wait for node to boot.
6. If the node still reports the error, replace system board.

Possible Cause-FRUs or other:

- System board

1029   **Enclosure VPD is unavailable or invalid.**

Explanation: Enclosure VPD is unavailable or invalid.

User response: Overwrite the enclosure VPD or replace the power interposer board.

Possible Cause-FRUs or other:

PIB card (10%)

Other:

No FRU (90%)

1030   **The internal disk of a node has failed.**

Explanation: An error has occurred while attempting to read or write data to the internal disk of one of the nodes in the cluster. The disk has failed.

User response: Determine which node's internal disk has failed using the node information in the error. Replace the FRUs in the order shown. Mark the error as fixed.

Possible Cause-FRUs or other:

2072 - Node Canister (100%)
- disk drive (50%)
- Disk controller (30%)
- Disk backplane (10%)
- Disk signal cable (8%)
- Disk power cable (1%)
- System board (1%)

1031   **Node canister location unknown.**

Explanation: Node canister location unknown.

User response: Complete the following steps to resolve this problem.
1. List all enclosure canisters for all control enclosures. Look for an online canister that does not have a node ID associated with it. This canister is the one with the problem.
2. Unplug the SAS cable from port 2 of the canister that is identified in step 1.
3. Run the command `lsenclosurecanister`, and see whether there is a node ID present. If step 2 fixes the error (a node ID is present), then something failed in one of the attached devices.
4. Reconnect the expansion enclosures and see whether the system is able to isolate the fault.
5. Reseat all the canisters on that strand and replace the canister that is identified in step 1 if step 4 does not fix the error.

Possible Cause-FRUs or other:

- Nothing (80%)
- Canister (20%)

1032   **Fibre Channel adapter not working**

Explanation: A problem has been detected on the node's Fibre Channel (FC) adapter. This node error is reported only on SAN Volume Controller 2145-CG8 or older nodes.

User response: Follow troubleshooting procedures to fix the hardware.
1. If possible, use the management GUI to run the recommended actions for the associated service error code.

Possible Cause-FRUs or other cause:

• None

1034   Canister fault type 2

Explanation: There is a canister internal error.

User response: Reseat the canister, and then replace the canister if the error continues.

Possible Cause-FRUs or other:

Canister (80%)

Other:

No FRU (20%)

1035   Boot drive problems

Explanation: Boot drive problems

User response: Complete the following steps:

1. Look at a boot drive view to determine the problems.

2. Run the commands **lsnodebootdrive** / **lsbootdrive** to display a status for each slot for users and DMPs to diagnose and repair problems.

3. If you plan to move any drives, shut down the node if `booted yes` is shown for that drive in the boot drive view (**lsbootdrive**). After you move the drives, a different node error will probably be displayed for you to work on.

4. If you plan to set the serial number of the system board, see **satask chvpd**.

5. If there is still no usable persistent data on the boot drives, then contact IBM Remote Technical Support.

Possible Cause-FRUs or other:

• System drive

1036   The enclosure identity cannot be read.

Explanation: The cluster is reporting that a node is not operational because of critical node error 509. See the details of node error 509 for more information.

User response: See node error 509.

1039   Canister failure, canister replacement required.

Explanation: Canister failure, canister replacement required.

User response: Replace the canister.

A canister can be safely replaced while the system is in production. Make sure that the other canister is the active node before removing this canister. It is preferable that this canister shuts down completely before removing it, but it is not required.

Possible cause-FRUs or other:

Interface adapter (50%)

SFP (20%)

Canister (20%)

Internal interface adapter cable (10%)

1040   Node flash disk fault

Explanation: A flash module error has occurred after a successful system start. Note: The node containing the flash module has not been rejected by the cluster.

User response:

1. Replace the FRUs below in the order listed

2. Check node status. If all nodes show a status of Online, mark the error that you have just repaired "fixed". If any nodes do not show a status of Online, go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145.

3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-CG8 or 2145-CF8

• Service controller (50%)

• Service controller cable (50%)

1044   A service controller read failure occurred

Explanation: A service controller read failure occurred.

User response:

1. Replace the FRUs below in the order listed.

2. Check node status. If all nodes show a status of Online, mark the error that you have just repaired "fixed". If any nodes do not show a status of Online, go to start MAP. If you return to this step, contact your support center to resolve the problem.

3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-CG8 or 2145-CF8

• Service controller (50%)

• Service controller cable (50%)

1048   Unexpected enclosure fault.

Explanation: Unexpected enclosure fault.

User response: Use the bottom snap option in the management GUI. This performs the following functions:

• Generates new enclosure dumps for all enclosures.

- Generates `livedump` from all nodes in the cluster.
- Runs an **`svc_snap dumpall`**.
1. Contact IBM support for further analysis.

Possible Cause-FRUs or other:
- None

1051   Pluggable TPM failed or missing

Explanation: The Trusted Platform Module (TPM) for the system is not functioning.

User response:

Important: Confirm that the system is running on at least one other node before you commence this repair. Each node uses its TPM to securely store encryption keys on its boot drive. When the TPM or boot drive of a node is replaced, the node loses its encryption key, and must be able to join an existing system to obtain the keys. If this error occurred on the last node in a system, do not replace the TPM, boot drive, or node hardware until the system contains at least one online node with valid keys.

1. Shut down the node and remove the node hardware.
2. Locate the TPM in the node hardware and ensure that it is correctly seated.
3. Reinsert the node hardware and apply power to the node.
4. If the error persists, replace the TPM with one from FRU stock.
5. If the error persists, replace the system board or the node hardware with one from FRU stock.

You do not need to return the faulty TPM to IBM.

Note: It is unlikely that the failure of a TPM can cause the loss of the System Master Key (SMK):

- The SMK is sealed by the TPM, using its unique encryption key, and the result is stored on the system boot drive.
- The working copy of the SMK is on the RAM disk, and so is unaffected by a sudden TPM failure.
- If the failure happens at boot time, the node is held in an unrecoverable error state because the TPM is a FRU.
- The SMK is also mirrored by the other nodes in the system. When the node with replacement TPM joins the system, it determines that it does not have the SMK, requests it, gets it, and then seals with the new TPM.

1052   Inter-canister PCIe link degraded

Explanation: The inter-canister PCI Express link is operating with reduced capacity. This might be a fault in either canister or the midplane. However, it is most likely to be a fault in the node canister that reported the problem.

User response:

- Unless both nodes in the I/O group are online, fix the problem that is causing the node to be offline first.
- Find which node reported the 1052 error by using the `object_id` or `object_name` value for the event in the event log. Check which node canister has that `node_id`.
- Use the CLI command **`svctask rmnode`** to delete the node canister that reported the problem.
- Replace the canister that reported the problem with the new canister.
- Wait for the new canister to be automatically added to the system.
- If the error is auto-fixed, then finish; otherwise, continue.
- Use the CLI command **`svctask rmnode`** to delete the other node canister.
- Replace the other node canister with the old node canister.
- Wait for that node canister to be automatically added to the system.
- If the error is auto-fixed, then finish; otherwise, continue.
- Replace the enclosure.
- If the error is auto-fixed, then finish; otherwise, call IBM for further support.

Possible Cause-FRUs or other:
- Canister that reported the problem (50%)
- Other canister (30%)
- Enclosure (20%)

1053   Internal SAS connector failure, service action required.

Explanation: An error occurred involving an internal SAS connector. Any of the following alerts might be associated with this error code.

- 045116 SAS connector to an enclosure secondary expander module is not working at full capacity
- 045117 SAS connector to an enclosure secondary expander module is offline
- 045118 The state of an enclosure secondary expander module connector cannot be determined

User response: Complete the following steps:

1. Enable maintenance mode for the I/O group.

2. Slide the enclosure out of the rack sufficiently to open the access lid.

3. Reseat the affected secondary expander module (SEM).

4. If the error does not clear, reseat the canister on the side of the affected SEM.

5. If the error does not clear, replace the affected SEM.

6. If the error does not clear, replace the canister on the side of the affected SEM.

7. If the error does not clear, contact your service support representative. You might need to replace the enclosure.

1054   **Fibre Channel adapter in slot 1 adapter present but failed.**

Explanation: The Fibre Channel adapter in PCI slot 1 is present but is failing.

User response:

1. In the sequence that is shown in the log, replace any failing FRUs with new FRUs.

2. Check node status:
   - If all nodes show a status of **online**, mark the error as **fixed**.
   - If any nodes do not show a status of **online**, go to the start MAP.
   - If you return to this step, contact your support center to resolve the problem with the node.

3. Go to the repair verification MAP.

Possible Cause, FRUs, or other:
- Fibre Channel host bus adapter (100%)

1055   **Fibre Channel adapter (4 port) in slot 1 adapter present but failed.**

Explanation: Fibre Channel adapter (4 port) in slot 1 adapter present but failed.

User response:

1. Exchange the FRU for new FRU.

2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.

3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-CF8, or 2145-CG8
- 4-port Fibre Channel host bus adapter (100%)

1056   **The Fibre Channel adapter in slot 2 is present but is failing.**

Explanation: The Fibre Channel adapter in slot 2 is present but is failing.

User response:

1. In the sequence that is shown in the log, replace any failing FRUs with new FRUs.

2. Check node status:
   - If all nodes show a status of **online**, mark the error as **fixed**.
   - If any nodes do not show a status of **online**, go to the start MAP.
   - If you return to this step, contact your support center to resolve the problem with the node.

3. Go to the repair verification MAP.

Possible Cause, FRUs, or other:
- N/A

1057   **Fibre Channel adapter (four-port) in slot 2 adapter is present but failing.**

Explanation: The four-port Fibre Channel adapter in slot 2 is present but failing.

User response:

1. In the sequence that is shown in the log, replace any failing FRUs with new FRUs.

2. Check node status:
   - If all nodes show a status of **online**, mark the error as **fixed**.
   - If any nodes do not show a status of **online**, go to the start MAP.
   - If you return to this step, contact your support center to resolve the problem with the node.

3. Go to the repair verification MAP.

Possible Cause, FRUs, or other:
- N/A

1059   **Fibre Channel IO port mapping failed**

Explanation: A Fibre Channel or Fibre Channel over Ethernet port is installed but is not included in the Fibre Channel I/O port mapping, and so the port cannot be used for Fibre Channel I/O. This error is raised in one of the following situations:

- A node hardware installation
- A change of I/O adapters
- The application of an incorrect Fibre Channel port map
- A node canister upgrade that combined old and new I/O adapters that could not be mapped automatically

These tasks are normally performed by service representatives.

User response: Your service representative can use the Service Assistant to modify the Fibre Channel I/O port mappings to include all of the installed ports that are capable of Fibre Channel I/O. The following command is used:

`satask chvpd -fcportmap`

1060   One or more Fibre Channel ports on the 2072 are not operational.

Explanation: One or more Fibre Channel ports on the 2072 are not operational.

User response:

1. Go to MAP 5600: Fibre Channel to isolate and repair the problem.
2. Go to the repair verification MAP.

Possible Cause-FRUs or other:

- Fibre Channel cable (80%)
- Small Form-factor Pluggable (SFP) connector (5%)
- 4-port Fibre Channel host bus adapter (5%)

Other:

- Fibre Channel network fabric (10%)

1061   Fibre Channel ports are not operational.

Explanation: Fibre Channel ports are not operational.

User response: An offline port can have many causes and so it is necessary to check them all. Start with the easiest and least intrusive possibility such as resetting the Fibre Channel or FCoE port via CLI command.

Possible Cause-FRUs or other:

External (cable, HBA/CNA, switch, and so on) (75%)

SFP (10%)

Interface (10%)

Node (5%)

1065   One or more Fibre Channel ports are running at lower than the previously saved speed.

Explanation: The Fibre Channel ports will normally operate at the highest speed permitted by the Fibre Channel switch, but this speed might be reduced if the signal quality on the Fibre Channel connection is poor. The Fibre Channel switch could have been set to operate at a lower speed by the user, or the quality of the Fibre Channel signal has deteriorated.

User response:

- Go to MAP 5600: Fibre Channel to resolve the problem.

Possible Cause-FRUs or other:

2072 - Node Canister (100%)

- Fibre Channel cable (50%)
- Small Form-factor Pluggable (SFP) connector (20%)
- 4-port Fibre Channel host bus adapter (5%)

Other:

- Fibre Channel switch, SFP connector, or GBIC (25%)

1067   Fan fault type 1

Explanation: The fan has failed.

User response: Replace the fan.

Possible Cause-FRUs or other:

Fan (100%)

1068   Fan fault type 2

Explanation: The fan is missing.

User response: Reseat the fan, and then replace the fan if reseating the fan does not correct the error.

Note: If replacing the fan does not correct the error, then the canister will need to be replaced.
Possible Cause-FRUs or other:

Fan (80%)

Other:

No FRU (20%)

1083   Unrecognized node error

Explanation: The cluster is reporting that a node is not operational because of critical node error 562. See the details of node error 562 for more information.

User response: See node error 562.

1084   System board device exceeded temperature threshold.

Explanation: System board device exceeded temperature threshold.

User response: Complete the following steps:

1. Check for external air flow blockages.
2. Remove the top of the machine case and check for missing baffles, damaged heat sinks, or internal blockages.
3. If problem persists, follow the service instructions for replacing the system board FRU in question.

Possible Cause-FRUs or other:

• Variable

1085   **PCI Riser card exceeded temperature threshold.**

Explanation: PCI Riser card exceeded temperature threshold.

User response: Complete the following steps:

1. Check airflow.
2. Remove the top of the machine case and check for missing baffles or internal blockages.
3. Check for faulty PCI cards and replace as necessary.
4. If problem persists, replace PCI Riser FRU.

Possible Cause-FRUs or other:

• None

1087   **Shutdown temperature threshold exceeded**

Explanation: Shutdown temperature threshold exceeded.

User response: Inspect the enclosure and the enclosure environment.

1. Check environmental temperature.
2. Ensure that all of the components are installed or that there are fillers in each bay.
3. Check that all of the fans are installed and operating properly.
4. Check for any obstructions to airflow, proper clearance for fresh inlet air, and exhaust air.
5. Handle any specific obstructed airflow errors that are related to the drive, the battery, and the power supply unit.
6. Bring the system back online. If the system performed a hard shutdown, the power must be removed and reapplied.

Possible Cause-FRUs or other:

Node (2%)

Battery (1%)

Power supply unit (1%)

Drive (1%)

Other:

Environment (95%)

1089   **One or more fans are failing.**

Explanation: One or more fans are failing.

User response:

1. Determine the failing fan(s) from the fan indicator on the system board or from the text of the error data in the log. The reported fan for the 2145-CF8, or 2145-CG8 matches the fan module position. Each fan module contains two fans.
2. Exchange the FRU for a new FRU.
3. Go to repair verification MAP.

• Fan number: Fan module position

• 1 or 2 :1

• 3 or 4 :2

• 5 or 6 :3

• 7 or 8 :4

• 9 or 10:5

• 11 or 12:6

Possible Cause-FRUs or other:

• Fan module (100%)

1090   **One or more fans (40x40x28) are failing.**

Explanation: One or more fans (40x40x28) are failing.

User response:

1. Determine the failing fans from the fan indicator on the system board or from the text of the error data in the log.
2. Verify that the cable between the fan backplane and the system board is connected:
    • If all fans on the fan backplane are failing
    • If no fan fault lights are illuminated
3. Exchange the FRU for a new FRU.
4. Go to repair verification MAP.

Possible Cause, FRUs, or other:

• N/A

1091   **One or more fans (40x40x56) are failing.**

Explanation: One or more fans (40x40x56) are failing.

User response:

1. Determine the failing fans from the fan indicator on the system board or from the text of the error data in the log.
2. Verify that the cable between the fan backplane and the system board is connected:
    • If all fans on the fan backplane are failing
    • If no fan fault lights are illuminated
3. Exchange the FRU for a new FRU.
4. Go to repair verification MAP.

Possible Cause, FRUs, or other:

- N/A

**1092   The temperature soft or hard shutdown
           threshold of the 2072 has been exceeded.
           The 2072 has automatically powered off.**

Explanation: The temperature soft or hard shutdown
threshold of the 2072 has been exceeded. The 2072 has
automatically powered off.

User response:

1. Ensure that the operating environment meets
   specifications.
2. Ensure that the airflow is not obstructed.
3. Ensure that the fans are operational.
4. Go to the light path diagnostic MAP and perform
   the light path diagnostic procedures.
5. Check node status. If all nodes show a status of
   "online", mark the error that you have just repaired
   as "fixed". If any nodes do not show a status of
   "online", go to the start MAP. If you return to this
   step, contact your support center to resolve the
   problem.
6. Go to the repair verification MAP.

Possible Cause-FRUs or other:

2072 - Node Canister (100%)

- The FRU that is indicated by the Light path
  diagnostics (25%)
- System board (5%)

Other:

System environment or airflow blockage (70%)

**1093   Temperature warning threshold
           exceeded**

Explanation: The system internal temperature sensor
has reported that the temperature warning threshold
has been exceeded.

User response:

1. Ensure that the internal airflow of the node has not
   been obstructed.
2. Check node status. If all nodes show a status of
   "online", mark the error that you have just repaired
   "fixed". If any nodes do not show a status of
   "online", go to the start MAP. If you return to this
   step, contact your support center to resolve the
   problem.
3. Go to repair verification MAP.

For the 2145-DH8 and 2145-CF8 only:

1. Check for external air flow blockages.

2. Remove the top of the machine case and check for
   missing baffles, damaged heatsinks, or internal
   blockages.
3. If the problem persists after taking these measures,
   replace the CPU assembly FRU if 2145-DH8 or
   2145-CF8.

Possible Cause-FRUs or other:

2145-CF8 or 2145-CG8

- Fan module (20%)
- System board (5%)
- Canister (5%)

Possible Cause-FRUs or other:

2145-DH8 or 2145-CF8

- CPU assembly (30%)

Other:

Airflow blockage (70%)

**1094   The ambient temperature threshold has
           been exceeded.**

Explanation: The ambient temperature threshold has
been exceeded.

User response:

1. Check that the room temperature is within the
   limits allowed.
2. Check for obstructions in the air flow.
3. Mark the errors as fixed.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

None

Other:

System environment (100%)

**1095   Enclosure temperature has passed
           critical threshold.**

Explanation: Enclosure temperature has passed critical
threshold.

User response: Check for external and internal air
flow blockages or damage.

1. Check environmental temperature.
2. Check for any impedance to airflow.
3. If the enclosure has shut down, then turn off both
   power switches on the enclosure and power both
   back on.

Possible Cause-FRUs or other:

- None

1096   **A Power Supply Unit is missing or has failed.**

Explanation: One of the two power supply units in the node is either missing or has failed.

NOTE: This error is reported when a hot-swap power supply is removed from an active node, so it might be reported when a faulty power supply is removed for replacement. Both the missing and faulty conditions report this error code.

User response: Error code 1096 is reported when the power supply either cannot be detected or reports an error.

1. Ensure that the power supply is seated correctly and that the power cable is attached correctly to both the node and to the 2145 UPS-1U.
2. If the error has not been automatically marked fixed after two minutes, note the status of the three LEDs on the back of the power supply. For the 2145-CG8 or 2145-CF8, the AC LED is the top green LED, the DC LED is the middle green LED and the error LED is the bottom amber LED.
3. If the power supply error LED is off and the AC and DC power LEDs are both on, this is the normal condition. If the error has not been automatically fixed after two minutes, replace the system board.
4. Follow the action specified for the LED states noted in the table below.
5. If the error has not been automatically fixed after two minutes, contact support.
6. Go to repair verification MAP.

Error,AC,DC:Action

ON,ON or OFF,ON or OFF:The power supply has a fault. Replace the power supply.

OFF,OFF,OFF:There is no power detected. Ensure that the power cable is connected at the node and 2145 UPS-1U. If the AC LED does not light, check the status of the 2145 UPS-1U to which the power supply is connected. Follow MAP 5150 2145 UPS-1U if the UPS-1U is showing no power or an error; otherwise, replace the power cable. If the AC LED still does not light, replace the power supply.

OFF,OFF,ON:The power supply has a fault. Replace the power supply.

OFF,ON,OFF:Ensure that the power supply is installed correctly. If the DC LED does not light, replace the power supply.

Possible Cause-FRUs or other:

Failed PSU:

- Power supply (90%)
- Power cable assembly (5%)
- System board (5%)

Missing PSU:

- Power supply (19%)
- System board (1%)
- Other: Power supply not correctly installed (80%)

1097   **PSU problem**

Explanation: One of the power supply units in the node is reporting that no main power is detected.

User response:

1. Ensure that the power supply is attached correctly to both the node and to the UPS.
2. If the error is not automatically marked fixed after 2 minutes, note the status of the three LEDs on the back of the power supply. For the 2145-CG8 or 2145-CF8, the AC LED is the top green LED. The DC LED is the middle green LED, and the error LED is the bottom amber LED.
3. If the power supply error LED is off and the AC and DC power LEDs are both on, this state is the normal condition. If the error is not automatically fixed after 2 minutes, replace the system board.
4. Follow the action that is specified for the LED states noted in the following list.
5. If the error is not automatically fixed after 2 minutes, contact support.
6. Go to repair verification MAP.

Error,AC,DC:Action

ON,ON or OFF,ON or OFF:The power supply has a fault. Replace the power supply.

OFF,OFF,OFF:There is no power detected. Ensure that the power cable is connected at the node and UPS. If the AC LED does not light, check whether the UPS is showing any errors. Follow MAP 5150 2145 UPS-1U if the UPS is showing an error; otherwise, replace the power cable. If the AC LED still does not light, replace the power supply.

OFF,OFF,ON:The power supply has a fault. Replace the power supply.

OFF,ON,OFF:Ensure that the power supply is installed correctly. If the DC LED does not light, replace the power supply.

Possible Cause-FRUs or other:

- Power cable assembly (85%)
- UPS-1U assembly (10%)

- System board (5%)

1098 **Enclosure temperature has passed warning threshold.**

Explanation: Enclosure temperature has passed warning threshold.

User response: Check for external and internal air flow blockages or damage.

1. Check environmental temperature.
2. Check for any impedance to airflow.

Possible Cause-FRUs or other:

- None

1099 **Temperature exceeded warning threshold**

Explanation: Temperature exceeded warning threshold.

User response: Inspect the enclosure and the enclosure environment.

1. Check environmental temperature.
2. Ensure that all of the components are installed or that there are fillers in each bay.
3. Check that all of the fans are installed and operating properly.
4. Check for any obstructions to airflow, proper clearance for fresh inlet air, and exhaust air.
5. Wait for the component to cool.

Possible Cause-FRUs or other:

Hardware component (5%)

Other:

Environment (95%)

1100 **One of the voltages that is monitored on the system board is over the set threshold.**

Explanation: One of the voltages that is monitored on the system board is over the set threshold.

User response:

1. See the light path diagnostic MAP.
2. If the light path diagnostic MAP does not resolve the issue, exchange the frame assembly.
3. Check node status. If all nodes show a status of "online", mark the error that you have just repaired as "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.
4. Go to repair verification MAP.

1101 **One of the voltages that is monitored on the system board is over the set threshold.**

Explanation: One of the voltages that is monitored on the system board is over the set threshold.

User response:

1. See the light path diagnostic MAP.
2. If the light path diagnostic MAP does not resolve the issue, exchange the system board assembly.
3. Check node status. If all nodes show a status of "online", mark the error that you have just repaired as "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

- Light path diagnostic MAP FRUs (98%)
- System board (2%)

1105 **One of the voltages that is monitored on the system board is under the set threshold.**

Explanation: One of the voltages that is monitored on the system board is under the set threshold.

User response:

1. Check the cable connections.
2. See the light path diagnostic MAP.
3. If the light path diagnostic MAP does not resolve the issue, exchange the frame assembly.
4. Check node status. If all nodes show a status of "online", mark the error that you have just repaired as "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.
5. Go to repair verification MAP.

1106 **One of the voltages that is monitored on the system board is under the set threshold.**

Explanation: One of the voltages that is monitored on the system board is under the set threshold.

User response:

1. Check the cable connections.
2. See the light path diagnostic MAP.
3. If the light path diagnostic MAP does not resolve the issue, exchange the system board assembly.
4. Check node status. If all nodes show a status of "online", mark the error that you have just repaired as "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.

5. Go to repair verification MAP.

Possible Cause-FRUs or other:
- Light path diagnostic MAP FRUs (98%)
- System board (2%)

| 1107 | **The battery subsystem has insufficient capacity to save system data due to multiple faults.** |

Explanation: This message is an indication of other problems to solve before the system can successfully recharge the batteries.

User response: No service action is required for this error, but other errors must be fixed. Look at other indications to see if the batteries can recharge without being put into use.

| 1108 | **Battery backplane cabling faulty or possible battery backplane requires replacing.** |

Explanation: Faulty cabling or a faulty backplane are preventing the system from full communication with and control of the batteries.

User response: Check the cabling to the battery backplane, making sure that all the connectors are properly mated.

Four signal cables (EPOW, LPC, PWR_SENSE & LED) and one power cable (which uses 12 red and 12 black heavy gauge wires) are involved:

- The EPOW cable runs to a 20-pin connector at the front of the system planar, which is the edge nearest the drive bays, near the left side.

  To check that this connector is mated properly, it is necessary to remove the plastic airflow baffle, which lifts up.

  A number of wires run from the same connector to the disk backplane located to the left of the battery backplane.

- The LPC cable runs to a small adapter that is plugged into the back of the system planar between two PCI Express adapter cages. It is helpful to remove the left adapter cage when checking that these connectors are mated properly.

- The PWR_SENSE cable runs to a 24-pin connector at the back of the system planar between the PSUs and the left adapter cage. Check the connections of both a female connector (to the system planar) and a male connector (to the connector from the top PSU). Again, it can be helpful to remove the left adapter cage to check the proper mating of the connectors.

- The power cable runs to the system planar between the PSUs and the left adapter cage. It is located just in front of the PWR_SENSE connector. This cable has both a female connector that connects to the system planar, and a male connector that mates with the

connector from the top PSU. Due to the bulk of this cable, care must be taken to not disturb PWR_SENSE connections when dressing it away in the space between the PSUs and the left adapter cage.

- The LED cable runs to a small PCB on the front bezel. The only consequence of this cable not being mated correctly is that the LEDs do not work.

If no problems exist, replace the battery backplane as described in the service action for "1109."

You do not replace either battery at this time.

To verify that the battery backplane works after replacing it, check that the node error is fixed.

Possible Cause-FRUs or other:
- Battery backplane (50%)

| 1109 | **Battery or possibly battery backplane requires replacing.** |

Explanation: Battery or possibly battery backplane requires replacing.

User response: Complete the following steps:
1. Replace the drive bay battery.
2. Check to see whether the node error is fixed. If not, replace the battery backplane.
3. To verify that the new battery backplane is working correctly, check that the node error is fixed.

Possible Cause-FRUs or other:
- Drive bay battery (95%)
- Battery backplane (5%)

| 1110 | **The power management board detected a voltage that is outside of the set thresholds.** |

Explanation: The power management board detected a voltage that is outside of the set thresholds.

User response:
1. In the sequence that is shown in the log, replace any failing FRUs with new FRUs.
2. Check node status:
   - If all nodes show a status of **online**, mark the error as **fixed**.
   - If any nodes do not show a status of **online**, go to the start MAP.
   - If you return to this step, contact your support center to resolve the problem with the node.
3. Go to repair verification MAP.

Possible Cause, FRUs, or other:

2145-CG8 or 2145-CF8

- Power supply unit (50%)
- System board (50%)

1111   **Batteries have insufficient charge.**

Explanation: The insufficient charge message can appear for various reasons such as the battery is charging; the battery is missing or has failed; there is a communication error, or there has been an over temperature event.

User response: This node error can be corrected by correcting each of the underlying battery problems.

1. If a battery is missing, replace the battery.
2. If a battery is failed, replace the battery.
3. If a battery is charging, this error should go away when the battery is charged.
4. If a battery is having a communication error (comm errror), try to reseat the battery as described in the replacement procedure. If reseating the battery does not correct the problem, replace the battery.
5. If a battery is too hot, the system can be started after it has cooled.

    Inspect the battery for damage after an over-temperature event.

Possible Cause - FRUs or other:

If both batteries have errors, battery charging might be underway. (No FRU)

If both batteries have errors that do not resolve after a sufficient time to charge, battery charging might be impaired, such as by a faulty battery backplane FRU.

Communication errors are often correctable by reseating the battery or by allowing the temperature of the battery to cool without the need to replace the battery. (No FRU)

If a battery is missing or failed, the solution is to replace the battery FRU.

Battery (50%)

Other:

No FRU (50%)

1112   **Enclosure battery is missing.**

Explanation: Enclosure battery is missing.

User response: Install a battery in the missing slot. If the battery is present in the slot, reseat the battery.

Attention: Do not reseat a battery unless the other battery has enough charge, or data loss might occur.

Possible Cause-FRUs or other:

Battery (95%)

Other:

No FRU (5%)

1114   **Enclosure battery fault type 1**

Explanation: Enclosure battery fault type 1.

User response: Replace the battery.

Possible Cause-FRUs or other:

Battery (100%)

1115   **Enclosure Battery fault type 4**

Explanation: Enclosure Battery fault type 4.

User response: Reseat the battery. Replace the battery if the error continues.

Note: Do not reseat a battery unless the other battery has enough charge, or data loss might occur.
Possible Cause-FRUs or other:

Battery (95%)

Other:

Bad connection (5%)

1120   **A high speed SAS adapter is missing**

Explanation: This node has detected that a high speed SAS adapter that was previously installed is no longer present.

User response: If the high speed SAS adapter was deliberately removed, mark the error "fixed."

Otherwise, the high speed SAS adapter has failed and must be replaced. In the sequence shown, exchange the FRUs for new FRUs.

Go to the repair verification MAP.

Possible Cause-FRUs or other:

1. High speed SAS adapter (90%)
2. System board (10%)

1121   **A high speed SAS adapter has failed.**

Explanation: A fault has been detected on a high speed SAS adapter.

User response: In the sequence shown, exchange the FRUs for new FRUs.

Go to the repair verification MAP.

Possible Cause-FRUs or other:

1. High speed SAS adapter (90%)
2. System board (10%)

1122　　　**A high speed SAS adapter error has occurred.**

Explanation:　The high speed SAS adapter has detected a PCI bus error and requires service before it can be restarted. The high speed SAS adapter failure has caused all of the flash drives that were being accessed through this adapter to go Offline.

User response:　If this is the first time that this error has occurred on this node, complete the following steps:

1. Power off the node.
2. Reseat the high speed SAS adapter.
3. Power on the node.
4. Submit the **lsmdisk** task and ensure that all of the flash drive managed disks that are located in this node have a status of Online.

If the sequence of actions above has not resolved the problem or the error occurs again on the same node, complete the following steps:

1. In the sequence shown, exchange the FRUs for new FRUs.
2. Submit the **lsmdisk** task and ensure that all of the flash drive managed disks that are located in this node have a status of Online.
3. Go to the repair verification MAP.

Possible Cause-FRUs or other:

1. High speed SAS adapter (90%)
2. System board (10%)

1124　　　**Power Supply Unit fault type 1**

Explanation:　A fault has been detected on a power supply unit (PSU).

User response:　Replace the PSU.

Attention:　To avoid losing state and data from the node, use the **satask startservice** command to put the node into service state so that it no longer processes I/O. Then you can remove and replace the top power supply unit (PSU 2). This precaution is due to a limitation in the power-supply configuration. Once the service action is complete, run the **satask stopservice** command to let the node rejoin the system.

Possible Cause-FRUs or other:

PSU (100%)

1125　　　**Power Supply Unit fault type 1**

Explanation:　The power supply unit (PSU) is not supported.

User response:　Replace the PSU with a supported version.

Attention:　To avoid losing state and data from the node, use the **satask startservice** command to put the node into service state so that it no longer processes I/O. Then you can remove and replace the top power supply unit (PSU 2). This precaution is due to a limitation in the power-supply configuration. Once the service action is complete, run the **satask stopservice** command to let the node rejoin the system.

Possible Cause-FRUs or other:

PSU (100%)

1126　　　**Power Supply Unit fault type 2**

Explanation:　A fault exists on the power supply unit (PSU).

User response:

1. Reseat the PSU in the enclosure.

   Attention:　To avoid losing state and data from the node, use the **satask startservice** command to put the node into service state so that it no longer processes I/O. Then you can remove and replace the top power supply unit (PSU 2). This precaution is due to a limitation in the power-supply configuration. Once the service action is complete, run the **satask stopservice** command to let the node rejoin the system.
2. If the fault is not resolved, replace the PSU.

Possible Cause-FRUs or other:

1. No Part (30%)
2. PSU (70 %)

1128　　　**Power Supply Unit missing**

Explanation:　The power supply unit (PSU) is not seated in the enclosure, or no PSU is installed.

User response:

1. If no PSU is installed, install a PSU.
2. If a PSU is installed, reseat the PSU in the enclosure.

Attention: To avoid losing state and data from the node, use the `satask startservice` command to put the node into service state so that it no longer processes I/O. Then you can remove and replace the top power supply unit (PSU 2). This precaution is due to a limitation in the power-supply configuration. Once the service action is complete, run the `satask stopservice` command to let the node rejoin the system.

Possible Cause-FRUs or other:

1. No Part (5%)
2. PSU (95%)

Reseat the power supply unit in the enclosure.

Possible Cause-FRUs or other:

Power supply (100%)

1129   The node battery is missing.

Explanation: Install new batteries to enable the node to join a clustered system.

User response: Install a battery in battery slot 1 (on the left from the front) and in battery slot 2 (on the right). Leave the node running as you add the batteries.

Align each battery so that the guide rails in the enclosure engage the **guide rail slots** on the battery. Push the battery firmly into the battery bay until it stops. The cam on the front of the battery remains closed during this installation.

To verify that the new battery works correctly, check that the node error is fixed. After the node joins a clustered system, use the `lsnodebattery` command to view information about the battery.

Possible Cause-FRUs or other:

- Battery (100%)

1130   The node battery requires replacing.

Explanation: When a battery must be replaced, you get this message. The proper response is to install new batteries.

User response: Battery 1 is on the left (from the front), and battery 2 is on the right. Remove the old battery by disengaging and pulling down the cam handle to lever out the battery enough to pull the battery from the enclosure.

This service procedure is intended for a failed or offline battery. To prevent losing data from a battery that is online, run the `svctask chnodebattery -remove -battery` *battery_ID node_ID*. Running the command verifies when it is safe to remove the battery.

Install new batteries in battery slot 1 and in battery slot 2. Leave the node running as you add the batteries.

Align each battery so that the guide rails in the enclosure engage the **guide rail slots** on the battery. Push the battery firmly into the battery bay until it stops. The cam on the front of the battery remains closed during this installation.

To verify that the new battery works correctly, check that the node error is fixed. After the node joins a clustered system, use the `lsnodebattery` command to view information about the battery.

1131   Battery conditioning is required but not possible.

Explanation: Battery conditioning is required but not possible.

One possible cause of this error is that battery reconditioning cannot occur in a clustered node if the partner node is not online.

User response: This error can be corrected on its own. For example, if the partner node comes online, the reconditioning begins.

Possible Cause-FRUs or other:

Other:

Wait, or address other errors.

1133   A duplicate WWNN has been detected.

Explanation: The cluster is reporting that a node is not operational because of critical node error 556. See the details of node error 556 for more information.

User response: See node error 556.

1135   UPS ambient temperature threshold exceeded

Explanation: The system UPS reported an ambient over temperature. The UPS switches to Bypass mode to cool down.

User response:

1. Power off the nodes that are attached to the UPS.
2. Turn off the UPS, and then unplug the UPS from the main power source.
3. Ensure that the air vents of the UPS are not obstructed.
4. Ensure that the air flow around the UPS is not restricted.
5. Wait for at least five minutes, and then restart the UPS. If the problem remains, check the ambient temperature. Correct the problem. Otherwise, exchange the FRU for a new FRU.
6. Check node status. If all nodes show a status of "online", mark the error that you have just repaired as "fixed". If any nodes do not show a status of

"online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the UPS.

7. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145 UPS electronics unit (50%)

Other:

The system ambient temperature is outside the specification (50%)

1136   UPS ambient temperature threshold exceeded

Explanation: The system UPS has reported an ambient over temperature.

User response:
1. Power off the node attached to the UPS.
2. Turn off the UPS, and then unplug the UPS from the main power source.
3. Ensure that the UPS air vents are not obstructed.
4. Ensure that the air flow around the UPS is not restricted.
5. Wait for at least five minutes, and then restart the UPS. If the problem remains, check the ambient temperature. Correct the problem. Otherwise, exchange the FRU for a new FRU.
6. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.
7. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145 UPS-1U assembly (50%)

Other:

The system ambient temperature is outside the specification (50%)

1138   Power supply unit input power failed.

Explanation: Power Supply Unit input power failed.

User response: Check the power cord.
1. Check that the power cord is plugged in.
2. Check that the wall power is good.
3. Replace the power cable.
4. Replace the power supply unit.

Possible Cause-FRUs or other:

Power cord (20%)

PSU (5%)

Other:

No FRU (75%)

1140   UPS AC input power fault

Explanation: The UPS has reported that it has a problem with the input AC power.

User response:
1. Check the input AC power, whether it is missing or out of specification. Correct if necessary. Otherwise, exchange the FRU for a new FRU.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:
- UPS input power cable (10%)
- Electronics assembly (10%)

Other:
- The input AC power is missing (40%)
- The input AC power is not in specification (40%)

1141   UPS AC input power fault

Explanation: The UPS has reported that it has a problem with the input AC power.

User response:
1. Check the input AC power, whether it is missing or out of specification. Correct if necessary. Otherwise, exchange the FRU for a new FRU.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:
- UPS input power cable (10%)
- UPS assembly (10%)

Other:
- The input AC power is missing (40%)

- The input AC power is not in specification (40%)

1145   UPS communications fault

Explanation: The signal connection between the system and its UPS is failing.

User response:
1. If other nodes that are using this UPS are reporting this error, exchange the UPS for a new one.
2. If only this node is reporting the problem, check the signal cable and exchange the FRUs for new FRUs in the sequence shown.
3. Check the node status:
   - If all nodes show a status of **online**, mark the error as **fixed**.
   - If any nodes do not show a status of **online**, go to the start MAP.
   - If you return to this step, contact your support center to resolve the problem.
4. Go to the repair verification MAP.

Possible Cause, FRUs, or other:

2145-CF8 or 2145-CG8

N/A

1146   UPS communications fault

Explanation: The signal connection between a node and its UPS is failing.

User response:
1. In the sequence that is shown in the log, replace any failing FRUs with new FRUs.
2. Check node status:
   - If all nodes show a status of **online**, mark the error as **fixed**.
   - If any nodes do not show a status of **online**, go to the start MAP.
   - If you return to this step, contact your support center to resolve the problem with the node.
3. Go to the repair verification MAP.

Possible Cause, FRUs, or other:

2145-CF8 or 2145-CG8
- Power cable assembly (40%)
- 2145 UPS-1U assembly (30%)
- System board (30%)

1150   UPS configuration error

Explanation: Data that the system received from the UPS suggests that the UPS power cable, the signal cable, or both, are not connected correctly.

User response:
1. Connect the cables correctly. See your product installation guide.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:
- None

Other:
- Configuration error

1151   UPS configuration error

Explanation: Data that the system received from the UPS suggests that the UPS power cable, the signal cable, or both, are not connected correctly.

User response:
1. Connect the cables correctly. See your product's installation guide.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:
- None

Other:
- Configuration error

1152   UPS configuration error

Explanation: The cluster is reporting that a node is not operational because of critical node error 587. See the details of node error 587 for more information.

User response: See node error 587.

1155   A power domain error has occurred.

Explanation: Both 2145s of a pair are powered by the same uninterruptible power supply.

User response:

1. List the 2145s of the cluster and check that 2145s in the same I/O group are connected to a different uninterruptible power supply.

2. Connect one of the 2145s as identified in step 1 to a different uninterruptible power supply.

3. Mark the error that you have just repaired, "fixed".

4. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Configuration error

1160   UPS output overcurrent

Explanation: The UPS reports that too much power is being drawn from it. The power overload warning LED, which is above the load level indicators on the UPS, will be lit.

User response:

1. Determine the UPS that is reporting the error from the error event data. Perform the following steps on just this UPS.

2. Check that the UPS is still reporting the error. If the power overload warning LED is no longer on, go to step 6.

3. Ensure that only appropriate systems are receiving power from the UPS. Ensure that there are no switches or disk controllers that are connected to the UPS.

4. Remove each connected input power in turn until the output overload is removed.

5. Exchange the FRUs for new FRUs in the sequence shown, on the overcurrent system.

6. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.

7. Go to repair verification MAP.

Possible Cause-FRUs or other:

- Power cable assembly (50%)
- Power supply assembly (40%)
- UPS electronics assembly (10%)

1161   UPS output overcurrent

Explanation: The output load on the UPS exceeds the specifications, as reported by UPS alarm bits.

User response:

1. Ensure that only appropriate systems are receiving power from the UPS. Also, ensure that no other devices are connected to the UPS.

2. Exchange, in the sequence shown, the FRUs for new FRUs. If the Overload Indicator is still illuminated with all outputs disconnected, replace the UPS.

3. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.

4. Go to repair verification MAP.

Possible Cause-FRUs or other:

- Power cable assembly (50%)
- Power supply assembly (40%)
- UPS assembly (10%)

1165   UPS output load high

Explanation: The UPS output load is unexpectedly high. The UPS output is possibly connected to an extra mismatched device.

User response:

1. Ensure that only appropriate systems are receiving power from the UPS. Ensure that there are no switches or disk controllers that are connected to the UPS.

2. Check node status. If all nodes show a status of "online", the problem no longer exists. Mark the error that you have just repaired "fixed" and go to the repair verification MAP.

3. Go to repair verification MAP.

Possible Cause-FRUs or other:

None

Other:

- Configuration error

1166   UPS output load high

Explanation: The uninterruptible power supply output is possibly connected to a mismatched device.

User response:

1. Ensure that there are no other devices that are connected to the UPS.

2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired

---

"fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145 UPS-1U.

3. Go to repair verification MAP.

Possible Cause-FRUs or other:

• UPS assembly (5%)

Other:

• Configuration error (95%)

1170   UPS electronics fault

Explanation: UPS electronics fault, reported by the UPS alarm bits.

User response:

1. Replace the UPS electronics assembly.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

UPS electronics assembly (100%)

1171   UPS electronics fault

Explanation: UPS electronics fault, reported by the UPS alarm bits.

User response:

1. Replace the UPS assembly.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

UPS assembly (100%)

1175   A problem has occurred with the uninterruptible power supply frame fault (reported by uninterruptible power supply alarm bits).

Explanation: A problem has occurred with the uninterruptible power supply frame fault (reported by the uninterruptible power supply alarm bits).

User response:

1. Replace the uninterruptible power supply assembly.

2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.

3. Go to repair verification MAP.

Possible Cause-FRUs or other:

Uninterruptible power supply assembly (100%)

1179   Too many drives attached to the system.

Explanation: The cluster only supports a fixed number of drives. A drive has been added that makes the number of drives larger than the total number of supported drives per cluster.

User response:

1. Disconnect any excessive unmanaged enclosures from the system.
2. Unmanage any offline drives that are not present in the system.
3. Identify unused drives and remove them from the enclosures.
4. Identify arrays of drives that are no longer required.
5. Remove the arrays and remove the drives from the enclosures if they are present.
6. Once there are fewer than 4096 drives in the system, consider re-engineering system capacity by migrating data from small arrays onto large arrays, then removing the small arrays and the drives that formed them. Consider the need for an additional Storwize system in your SAN solution.

1180   UPS battery fault

Explanation: UPS battery fault, reported by UPS alarm bits.

User response:

1. Replace the UPS battery assembly.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

UPS battery assembly (100%)

1181   **UPS battery fault**

Explanation: UPS battery fault, reported by UPS alarm bits.

User response:

1. Replace the UPS battery assembly.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

UPS battery assembly (100%)

1182   **Ambient temperature is too high during system startup.**

Explanation: The cluster is reporting that a node is not operational because of critical node error 528. See the details of node error 528 for more information.

User response: See node error 528.

1183   **The nodes hardware configuration does not meet the minimum requirements.**

Explanation: The cluster is reporting that a node is not operational because of critical node error 562. See the details of node error 562 for more information.

User response: See node error 562.

1185   **UPS fault**

Explanation: UPS fault, with no specific unit identified, reported by UPS alarm bits.

User response:

1. In the sequence shown, exchange the FRU for a new FRU.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:
- 2145 UPS electronics assembly (60%)
- 2145 UPS battery assembly (20%)
- 2145 UPS assembly (20%)

1186   **UPS fault**

Explanation: A problem has occurred in the UPS, with no specific FRU identified, as reported by UPS alarm bits.

User response:

1. In the sequence shown, exchange the FRU for a new FRU.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

UPS assembly (100%)

1187   **Node software is inconsistent or damaged**

Explanation: The cluster is reporting that a node is not operational because of critical node errors 523, 573, 574. See the details of node errors 523, 573, 574 for more information.

User response: See node errors 523, 573, 574.

1188   **Too many software crashes have occurred.**

Explanation: The cluster is reporting that a node is not operational because of critical node error 564. See the details of node error 564 for more information.

User response: See node error 564.

1189   **The node is held in the service state.**

Explanation: The cluster is reporting that a node is not operational because of critical node error 690. See the details of node error 690 for more information.

User response: See node error 690.

1190   **UPS battery end of life**

Explanation: The UPS battery has reached its end of life.

User response:

1. Replace the UPS battery assembly.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

UPS battery assembly (100%)

1191 UPS battery end of life

Explanation: The UPS battery has reached its end of life.

User response:

1. Replace the UPS battery assembly.

2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.

3. Go to repair verification MAP.

Possible Cause-FRUs or other:

UPS battery assembly (100%)

1192 Unexpected node error

Explanation: A node is missing from the cluster. The error that it is reporting is not recognized by the system.

User response: Find the node that is in service state and use the service assistant to determine why it is not active.

1193 Insufficient uninterruptible power supply charge

Explanation: The cluster is reporting that a node is not operational because of critical node error 587, indicating that an incorrect type of UPS was installed.

User response: Exchange the UPS for one of the correct type.

1194 Automatic recovery of offline node has failed.

Explanation: The cluster has an offline node and has determined that one of the candidate nodes matches the characteristics of the offline node. The cluster has attempted but failed to add the node back into the cluster. The cluster has stopped attempting to automatically add the node back into the cluster.

If a node has incomplete state data, it remains offline after it starts. This occurs if the node has had a loss of power or a hardware failure that prevented it from completing the writing of all of the state data to disk. The node reports a node error 578 when it is in this state.

If three attempts to automatically add a matching candidate node to a cluster have been made, but the

node has not returned online for 24 hours, the cluster stops automatic attempts to add the node and logs error code 1194 "Automatic recovery of offline node failed".

Two possible scenarios when this error event is logged are:

1. The node has failed without saving all of its state data. The node has restarted, possibly after a repair, and shows node error 578 and is a candidate node for joining the cluster. The cluster attempts to add the node into the cluster but does not succeed. After 15 minutes, the cluster makes a second attempt to add the node into the cluster and again does not succeed. After another 15 minutes, the cluster makes a third attempt to add the node into the cluster and again does not succeed. After another 15 minutes, the cluster logs error code 1194. The node never came online during the attempts to add it to the cluster.

2. The node has failed without saving all of its state data. The node has restarted, possibly after a repair, and shows node error 578 and is a candidate node for joining the cluster. The cluster attempts to add the node into the cluster and succeeds and the node becomes online. Within 24 hours the node fails again without saving its state data. The node restarts and shows node error 578 and is a candidate node for joining the cluster. The cluster again attempts to add the node into the cluster, succeeds, and the node becomes online; however, the node again fails within 24 hours. The cluster attempts a third time to add the node into the cluster, succeeds, and the node becomes online; however, the node again fails within 24 hours. After another 15 minutes, the cluster logs error code 1194.

A combination of these scenarios is also possible.

Note: If the node is manually removed from the cluster, the count of automatic recovery attempts is reset to zero.

User response:

1. If the node has been continuously online in the cluster for more than 24 hours, mark the error as fixed and go to the Repair Verification MAP.

2. Determine the history of events for this node by locating events for this node name in the event log. Note that the node ID will change, so match on the WWNN and node name. Also, check the service records. Specifically, note entries indicating one of three events: 1) the node is missing from the cluster (cluster error 1195 event 009052), 2) an attempt to automatically recover the offline node is starting (event 980352), 3) the node has been added to the cluster (event 980349).

3. If the node has not been added to the cluster since the recovery process started, there is probably a hardware problem. The node's internal disk might

be failing in a manner that it is unable to modify its software level to match the software level of the cluster. If you have not yet determined the root cause of the problem, you can attempt to manually remove the node from the cluster and add the node back into the cluster. Continuously monitor the status of the nodes in the cluster while the cluster is attempting to add the node. Note: If the node type is not supported by the software version of the cluster, the node will not appear as a candidate node. Therefore, incompatible hardware is not a potential root cause of this error.

4. If the node was added to the cluster but failed again before it has been online for 24 hours, investigate the root cause of the failure. If no events in the event log indicate the reason for the node failure, collect dumps and contact IBM technical support for assistance.

5. When you have fixed the problem with the node, you must use either the cluster console or the command line interface to manually remove the node from the cluster and add the node into the cluster.

6. Mark the error as fixed and go to the verification MAP.

Possible Cause-FRUs or other:

None, although investigation might indicate a hardware failure.

1195   Node missing.

Explanation: You can resolve this problem by repairing the failure on the missing 3700.

User response:

1. If it is not obvious which node in the cluster has failed, check the status of the nodes and find the 3700 with a status of offline.

2. Go to the Start MAP and perform the repair on the failing node.

3. When the repair has been completed, this error is automatically marked as fixed.

4. Check node status. If all nodes show a status of "online", but the error in the log has not been marked as fixed, manually mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the 3700.

5. Go to repair verification MAP.

Possible Cause-FRUs or other:
• None

1198   Detected hardware is not a valid configuration.

Explanation: A hardware change was made to this node that is not supported by its software. Either a hardware component failed, or the node was incorrectly upgraded.

User response: Complete the following steps:

1. If required, power the node off for servicing.

2. If new hardware is correctly installed, but it is listed as an invalid configuration, then update the software to a level that supports the new hardware. Use the management GUI to install this level if necessary.

3. If you upgraded the software to make the hardware work, there is a new event after the upgrade requesting that you enable the new hardware.

Possible Cause-FRUs or other:
• None

1200   The configuration is not valid. Too many devices, MDisks, or targets have been presented to the system.

Explanation: The configuration is not valid. Too many devices, MDisks, or targets have been presented to the system.

User response:

1. Remove unwanted devices from the Fibre Channel network fabric.

2. Start a cluster discovery operation to find devices/disks by rescanning the Fibre Channel network.

3. List all connected managed disks. Check with the customer that the configuration is as expected. Mark the error that you have just repaired fixed.

4. Go to repair verification MAP.

Possible Cause-FRUs or other:
• None

Other:

Fibre Channel network fabric fault (100%)

1201   A flash drive requires a recovery.

Explanation: The flash drive that is identified by this error needs to be recovered.

User response: To recover this flash drive, submit the following command: `chdrive -task recover` *drive_id* where *drive_id* is the identity of the drive that needs to be recovered.

1202   **A flash drive is missing from the configuration.**

Explanation: The offline flash drive identified by this error must be repaired.

User response: In the management GUI, click **Troubleshooting** > **Recommended Actions** to run the recommended action for this error. Otherwise, use MAP 6000 to replace the drive.

1203   **A duplicate Fibre Channel frame has been received.**

Explanation: A duplicate Fibre Channel frame should never be detected. Receiving a duplicate Fibre Channel frame indicates that there is a problem with the Fibre Channel fabric. Other errors related to the Fibre Channel fabric might be generated.

User response:

1. Use the transmitting and receiving WWPNs indicated in the error data to determine the section of the Fibre Channel fabric that has generated the duplicate frame. Search for the cause of the problem by using fabric monitoring tools. The duplicate frame might be caused by a design error in the topology of the fabric, by a configuration error, or by a software or hardware fault in one of the components of the Fibre Channel fabric, including inter-switch links.

2. When you are satisfied that the problem has been corrected, mark the error that you have just repaired "fixed".

3. Go to MAP 5700: Repair verification.

Possible Cause-FRUs or other:

- Fibre Channel cable assembly (1%)
- Fibre Channel adapter (1%)

Other:

- Fibre Channel network fabric fault (98%)

1210   **A local Fibre Channel port has been excluded.**

Explanation: A local Fibre Channel port has been excluded.

User response:

1. Repair faults in the order shown.

2. Check the status of the disk controllers. If all disk controllers show a "good" status, mark the error that you just repaired as "fixed".

3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- Fibre Channel cable assembly (75%)
- Small Form-factor Pluggable (SFP) connector (10%)

- Fibre Channel adapter (5%)

Other:

- Fibre Channel network fabric fault (10%)

1212   **Power supply exceeded temperature threshold.**

Explanation: Power supply exceeded temperature threshold.

User response: Complete the following steps:

1. Check airflow. Remove the top of the machine case and check for missing baffles or internal blockages.

2. If problem persists, replace the power supply.

Possible Cause-FRUs or other:

- Power supply

1213   **Boot drive missing, out of sync, or failed.**

Explanation: Boot drive missing, out of sync, or failed.

User response: Complete the following steps:

1. Look at a boot drive view to determine the missing, failed or out of sync drive.

2. Insert a missing drive.

3. Replace a failed drive.

4. Synchronize an out of sync drive by running the commands **svctask chnodebootdrive -sync** and/or **satask chbootdrive -sync**.

Possible Cause-FRUs or other:

- System drive

1214   **Boot drive is in the wrong slot.**

Explanation: Boot drive is in the wrong slot.

User response: Complete the following steps:

1. Look at a boot drive view to determine which drive is in the wrong slot, which node and slot it belongs in, and which drive must be in this slot.

2. Swap the drive for the correct one but shut down the node first if `booted yes` is shown for that drive in boot drive view.

3. If you want to use the drive in this node, synchronize the boot drives by running the commands **svctask chnodebootdrive -sync** and/or **satask chbootdrive -sync**.

4. The node error clears, or a new node error is displayed for you to work on.

Possible Cause-FRUs or other:

- None

1215   A flash drive is failing.

Explanation: The flash drive has detected faults that indicate that the drive is likely to fail soon. The drive should be replaced. The cluster event log will identify a drive ID for the flash drive that caused the error.

User response: In the management GUI, click **Troubleshooting** > **Recommended Actions** to run the recommended action for this error. If this does not resolve the issue, contact your next level of support.

1216   SAS errors have exceeded thresholds.

Explanation: The cluster has experienced a large number of SAS communication errors, which indicates a faulty SAS component that must be replaced.

User response: In the sequence shown, exchange the FRUs for new FRUs.

Go to the repair verification MAP.

Possible Cause-FRUs or other:

1. SAS Cable (70%)
2. High speed SAS adapter (20%)
3. SAS drive backplane (5%)
4. flash drive (5%)

1217   A flash drive has exceeded the temperature warning threshold.

Explanation: The flash drive identified by this error has reported that its temperature is higher than the warning threshold.

User response: Take steps to reduce the temperature of the drive.

1. Determine the temperature of the room, and reduce the room temperature if this action is appropriate.
2. Replace any failed fans.
3. Ensure that there are no obstructions to air flow for the node.
4. Mark the error as fixed. If the error recurs, contact hardware support for further investigation.

Possible Cause-FRUs or other:

- Flash drive (10%)

Other:

- System environment or airflow blockage (90%)

1220   A remote Fibre Channel port has been excluded.

Explanation: A remote Fibre Channel port has been excluded.

User response:

1. View the event log. Note the MDisk ID associated with the error code.
2. From the MDisk, determine the failing disk controller ID.
3. Refer to the service documentation for the disk controller and the Fibre Channel network to resolve the reported problem.
4. After the disk drive is repaired, start a cluster discovery operation to recover the excluded Fibre Channel port by rescanning the Fibre Channel network.
5. To restore MDisk online status, include the managed disk that you noted in step 1.
6. Check the status of the disk controller. If all disk controllers show a "good" status, mark the error that you have just repaired, "fixed".
7. If all disk controllers do not show a good status, contact your support center to resolve the problem with the disk controller.
8. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Enclosure/controller fault (50%)
- Fibre Channel network fabric (50%)

1230   A login has been excluded.

Explanation: A port to port fabric connection, or login, between the cluster node and either a controller or another cluster has had excessive errors. The login has therefore been excluded, and will not be used for I/O operations.

User response: Determine the remote system, which might be either a controller or a cluster. Check the event log for other 1230 errors. Ensure that all higher priority errors are fixed.

This error event is usually caused by a fabric problem. If possible, use the fabric switch or other fabric diagnostic tools to determine which link or port is reporting the errors. If there are error events for links from this node to a number of different controllers or clusters, then it is probably the node to switch link that is causing the errors. Unless there are other contrary indications, first replace the cable between the switch and the remote system.

1. From the fabric analysis, determine the FRU that is most likely causing the error. If this FRU has recently been replaced while resolving a 1230 error, choose the next most likely FRU that has not been replaced recently. Exchange the FRU for a new FRU.
2. Mark the error as fixed. If the FRU replacement has not fixed the problem, the error will be logged

again; however, depending on the severity of the problem, the error might not be logged again immediately.

3. Start a cluster discovery operation to recover the login by re-scanning the Fibre Channel network.

4. Check the status of the disk controller or remote cluster. If the status is not "good", go to the Start MAP.

5. Go to repair verification MAP.

Possible Cause-FRUs or other:

- Fibre Channel cable, switch to remote port, (30%)
- Switch or remote device SFP connector or adapter, (30%)
- Fibre Channel cable, local port to switch, (30%)
- Cluster SFP connector, (9%)
- Cluster Fibre Channel adapter, (1%)

Note: The first two FRUs are not cluster FRUs.

1260   SAS cable fault type 2.

Explanation: The associated alert event contains more information about the error:

045014 SAS cable excluded due to internal errors
 The cable was excluded because one or more phys (lanes of communication) are missing.

045015 SAS cable excluded due to causing too many change events
 The connector port caused too many change events.

045017 SAS cable is operating at reduced speed
 If the cable is not the last path to data, reduced speed causes it to be excluded.

045018 SAS cable excluded due to dropped frames
 Frame errors occurred.

045019 SAS cable excluded due to enclosure discovery timing out
 Enclosure discovery timed out before the cable could be identified.

045051 SAS cable excluded due to Single Port Active drives The connector or the attached canister might be the cause of multiple single-ported drives.

045077 Attempts to exclude connector have failed
 Multiple attempts to exclude the failing connector did not change the connector state.

045102 SAS cable is not working at full capacity
 Some of the physical data paths in the cable are not working properly. This error is logged only if no other events are logged on the cable.

In all cases, the user response is the same.

User response: Complete the following steps:

Note: After each action, check to see whether the canister ports at both ends of the cable are excluded. If the ports are excluded, then enable them by issuing the following command:

`chenclosurecanister -excludesasport no -port X`

1. Reset this canister and the upstream canister.

 The upstream canister is identified in sense data as enclosureid2, faultobjectlocation2...

2. Reseat the cable between the two ports that are identified in the sense data.

3. Replace the cable between the two ports that are identified in the sense data.

4. Replace this canister.

5. Replace the other canister (enclosureid2).

Possible Cause-FRUs or other:

- SAS cable
- Canister

1266   SEM Fault Type 1

Explanation: An unrecoverable error occurred involving a secondary expander module (SEM). The SEM must be replaced.

User response: Complete the following steps:

1. Enable maintenance mode for the I/O group.

2. Slide the enclosure out of the rack sufficiently to open the access lid.

3. Remove the failed SEM.

4. Insert the replacement SEM.

5. Close the access lid.

6. Slide the enclosure back into the rack.

7. Maintenance mode will disable automatically after 30 minutes, or you can disable it manually

8. If the error does not autofix, contact your service support representative.

1267   Enclosure secondary expander module is missing

Explanation: An error occurred involving a secondary expander module (SEM). You might be able to resolve the problem by reseating the SEM. The alert event gives more information about the error.

045105 Enclosure secondary expander module has failed A SEM is offline and might have failed.

045107 Enclosure secondary expander module temperature sensor cannot be read
 A SEM temperature sensor could not be read.

045114 Enclosure secondary expander module connector excluded due to too many change events
 A SEM is in degraded state due to too many transient errors.

045120 Enclosure secondary expander module is missing

    A SEM was removed from the disk drawer for an enclosure.

045121 Enclosure secondary expander module connector excluded due to dropped frames

    An internal SAS connector in the enclosure is in a degraded state due to too many Virtual LUN Manager login errors.

045122 Enclosure secondary expander module connector is excluded and cannot be unexcluded

    An internal SAS connector in the enclosure was excluded and cannot be included.

045123 Enclosure secondary expander module connectors excluded as the cause of single ported drives SEM connectors were excluded because slot ports under them were unreachable.

045124 Enclosure secondary expander module leaf expander connector excluded as the cause of single ported drives

    An SEM leaf expander connector was excluded because slot ports under it were unreachable.

User response: Complete the following steps:

1. Reseat the SEM:
   a. Enable maintenance mode for the I/O group.
   b. Slide the enclosure out of the rack sufficiently to open the access lid.
   c. Remove the designated SEM.
   d. Reinsert the designated SEM.
   e. Maintenance mode will disable automatically after 30 minutes, or you can disable it manually.
2. If the error autofixes, close up the enclosure:
   a. Close the access lid.
   b. Slide the enclosure back into the rack.
3. If the error does not autofix, replace the SEM:
   a. Enable maintenance mode for the I/O group.
   b. Slide the enclosure out of the rack sufficiently to open the access lid.
   c. Remove the failed SEM.
   d. Insert the replacement SEM.
   e. Close the access lid.
   f. Slide the enclosure back into the rack.
   g. Maintenance mode will disable automatically after 30 minutes, or you can disable it manually.

1268   Enclosure Display Panel Fault Type 2

Explanation: A problem was found with the display panel for the enclosure. The alert event gives more information about the error.

045110 Enclosure display panel is not installed

    The display panel is offline and might be missing.

045111 Enclosure display panel temperature sensor cannot be read

    The temperature sensor for the display panel could not be read.

045119 Enclosure display panel VPD cannot be read

    The Vital Product Data (VPD) for the display panel could not be read.

User response: Complete the following steps:

1. Reseat the display panel:
   a. Put the system into maintenance mode.
   b. Slide the enclosure out of the rack sufficiently to remove the top cover and remove the top cover.
   c. Locate the display panel access handle.
   d. Pinch the sides of the display panel handle and remove the display panel module
   e. Reinsert the display panel module.
   f. Replace the cover and slide the enclosure back into the rack.
   g. Turn off maintenance mode.
2. If the error does not clear, replace the display panel:
   a. Turn on maintenance mode.
   b. Slide the enclosure out of the rack sufficiently to remove the top cover and remove the top cover.
   c. Locate the display panel access handle.
   d. Pinch the sides of the display panel handle and remove the display panel module.
   e. Insert the replacement display panel module.
   f. Replace the cover and slide the enclosure back into the rack.
   g. Turn off maintenance mode
3. If the error does not clear, the enclosure might need to be replaced. Contact your service support representative.

1298   A node has encountered an error updating.

Explanation: One or more nodes has failed the update.

User response: Check `lsupdate` for the node that failed and continue troubleshooting with the error code it provides.

1300   IO port configuration issue

Explanation: A port that was configured for N_Port ID virtualization (NPIV) is off line.

User response: Complete both of the following procedures:

1. Check the switch configuration to ensure that NPIV is enabled and that resource limits are sufficient.
2. Run the **detectmdisks** command and wait 30 seconds after the discovery completes to see if the event fixes itself.
3. If the event does not fix itself, contact IBM Support.

1310   **A managed disk is reporting excessive errors.**

Explanation: A managed disk is reporting excessive errors.

User response:

1. Repair the enclosure/controller fault.
2. Check the managed disk status. If all managed disks show a status of "online", mark the error that you have just repaired as "fixed". If any managed disks show a status of "excluded", include the excluded managed disks and then mark the error as "fixed".
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

• None

Other:

Enclosure/controller fault (100%)

1311   **A flash drive is offline due to excessive errors.**

Explanation: The drive that is reporting excessive errors has been taken offline.

User response: In the management GUI, click **Troubleshooting** > **Recommended Actions** to run the recommended action for this error. If this does not resolve the issue, contact your next level of support.

1320   **A disk I/O medium error has occurred.**

Explanation: A disk I/O medium error has occurred.

User response:

1. Check whether the volume the error is reported against is mirrored. If it is, check if there is a "1870 Mirrored volume offline because a hardware read error has occurred" error relating to this volume in the event log. Also check if one of the mirror copies is synchronizing. If all these tests are true then you must delete the volume copy that is not synchronized from the volume. Check that the volume is online before continuing with the following actions. Wait until the medium error is corrected before trying to re-create the volume mirror.
2. If the medium error was detected by a read from a host, ask the customer to rewrite the incorrect data

to the block logical block address (LBA) that is reported in the host systems SCSI sense data. If an individual block cannot be recovered it will be necessary to restore the volume from backup. (If this error has occurred during a migration, the host system does not notice the error until the target device is accessed.)

3. If the medium error was detected during a mirrored volume synchronization, the block might not be being used for host data. The medium error must still be corrected before the mirror can be established. It may be possible to fix the block that is in error using the disk controller or host tools. Otherwise, it will be necessary to use the host tools to copy the volume content that is being used to a new volume. Depending on the circumstances, this new volume can be kept and mirrored, or the original volume can be repaired and the data copied back again.
4. Check managed disk status. If all managed disks show a status of "online", mark the error that you have just repaired as "fixed". If any managed disks do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the disk controller.
5. Go to repair verification MAP.

Possible Cause-FRUs or other:

• None

Other:

Enclosure/controller fault (100%)

1322   **Data protection information mismatch.**

Explanation: This error occurs when something has broken the protection information in read or write commands.

User response:

1. Determine if there is a single or multiple drives logging the error. Because the SAS transport layer can cause multiple drive errors, it is necessary to fix other hardware errors first.
2. Check related higher priority hardware errors. Fix higher priority errors before continuing.
3. Use **lseventlog** to determine if more than one drive with this error has been logged in the last 24 hours. If so, contact IBM support.
4. If only a single drive with this error has been logged, the system is monitoring the drive for health and will fail if RAID is used to correct too many errors of this kind.

1328   **Encryption key required.**

Explanation: It is necessary to provide an encryption key before the system can become fully operational. This error occurs when a system with encryption enabled is restarted without an encryption key available.

User response: Connect a USB flash drive or a key server that contains the current key for this system to one or more of the nodes.

1330   **A suitable managed disk (MDisk) or drive for use as a quorum disk was not found.**

Explanation: A quorum disk is needed to enable a tie-break when some cluster members are missing. Three quorum disks are usually defined. By default, the cluster automatically allocates quorum disks when managed disks are created; however, the option exists to manually assign quorum disks. This error is reported when there are managed disks or image mode disks but no quorum disks.

To become a quorum disk:

- The MDisk must be accessible by all nodes in the cluster.
- The MDisk must be managed; that is, it must be a member of a storage pool.
- The MDisk must have free extents.
- The MDisk must be associated with a controller that is enabled for quorum support. If the controller has multiple WWNNs, all of the controller components must be enabled for quorum support.

A quorum disk might not be available because of a Fibre Channel network failure or because of a Fibre Channel switch zoning problem.

User response:

1. Resolve any known Fibre Channel network problems.
2. Ask the customer to confirm that MDisks have been added to storage pools and that those MDisks have free extents and are on a controller that is enabled for use as a provider of quorum disks. Ensure that any controller with multiple WWNNs has all of its components enabled to provide quorum disks. Either create a suitable MDisk or if possible enable quorum support on controllers with which existing MDisks are associated. If at least one managed disk shows a mode of managed and has a non-zero quorum index, mark the error that you have just repaired as "fixed".
3. If the customer is unable to make the appropriate changes, ask your software support center for assistance.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

Configuration error (100%)

1335   **Quorum disk not available.**

Explanation: Quorum disk not available.

User response:

1. View the event log entry to identify the managed disk (MDisk) being used as a quorum disk, that is no longer available.
2. Perform the disk controller problem determination and repair procedures for the MDisk identified in step 1.
3. Include the MDisks into the cluster.
4. Check the managed disk status. If the managed disk identified in step 1 shows a status of "online", mark the error that you have just repaired as "fixed". If the managed disk does not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the disk controller.
5. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

Enclosure/controller fault (100%)

1340   **A managed disk has timed out.**

Explanation: This error was reported because a large number of disk timeout conditions have been detected. The problem is probably caused by a failure of some other component on the SAN.

User response:

1. Repair problems on all enclosures or controllers and switches on the same SAN as this 2145 cluster.
2. If problems are found, mark this error as "fixed".
3. If no switch or disk controller failures can be found, take an event log dump and call your hardware support center.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Enclosure/controller fault
- Fibre Channel (FC) switch

1350   IB ports are not operational.

Explanation: IB ports are not operational.

User response: An offline port can have many causes and so it is necessary to check them all. Start with the easiest and least intrusive possibility.

1. Reset the IB port with CLI command.

2. If the IB port is connected to a switch, double-check the switch configuration for issues.

3. Reseat the IB cable on both the IB side and the HBA/switch side.

4. Run a temporary second IB cable to replace the current one to check for a cable fault.

5. If the system is in production, schedule a maintenance downtime before continuing to the next step. Other ports will be affected.

6. Reset the IB interface adapter; reset the node; reboot the system.

Possible Cause-FRUs or other:

External (cable, HCA, switch, and so on) (85%)

Interface (10%)

Node (5%)

1360   A SAN transport error occurred.

Explanation: This error has been reported because the 2145 performed error recovery procedures in response to SAN component associated transport errors. The problem is probably caused by a failure of some other component on the SAN.

User response:

1. View the event log entry to determine the node that logged the problem. Determine the 2145 node or controller that the problem was logged against.

2. Perform Fibre Channel (FC) switch problem determination and repair procedures for the switches connected to the 2145 node or controller.

3. Perform FC cabling problem determination and repair procedures for the cables connected to the 2145 node or controller.

4. If any problems are found and resolved in step 2 and 3, mark this error as "fixed".

5. If no switch or cable failures were found in steps 2 and 3, take an event log dump. Call your hardware support center.

6. Go to repair verification MAP.

Possible Cause-FRUs or other:
- None

Other:
- FC switch

- FC cabling

1370   A managed disk error recovery procedure (ERP) has occurred.

Explanation: This error was reported because a large number of disk error recovery procedures have been performed by the disk controller. The problem is probably caused by a failure of some other component on the SAN.

User response:

1. View the event log entry and determine the managed disk that was being accessed when the problem was detected.

2. Perform the disk controller problem determination and repair procedures for the MDisk determined in step 1.

3. Perform problem determination and repair procedures for the Fibre Channel (FC) switches connected to the 2145 and any other FC network components.

4. If any problems are found and resolved in steps 2 and 3, mark this error as "fixed".

5. If no switch or disk controller failures were found in steps 2 and 3, take an event log dump. Call your hardware support center.

6. Go to repair verification MAP.

Possible Cause-FRUs or other:
- None

Other:
- Enclosure/controller fault
- Fibre Channel (FC) switch

1400   Ethernet port failure

Explanation: The system cannot detect an Ethernet connection.

User response:

1. Go to the Ethernet MAP.

2. Go to the repair verification MAP.

Possible Cause-FRUs or other:

2145-CF8, or 2145-CG8
- Ethernet cable (25%)
- System board (25%)

Other:
- Ethernet cable is disconnected or damaged (25%)
- Ethernet hub fault (25%)

1403   **External port not operational.**

Explanation: If this error occurs when a port was initially online and subsequently went offline, it indicates:

- the server, HBA, CNA or switch has been turned off.
- there is a physical issue.

If this error occurs during an initial setup or during a setup change, it is most likely a configuration issue rather than a physical issue.

User response:
1. Reset the port via the CLI command `Maintenance`. If the port is now online, the DMP is complete.
2. If the port is connected to a switch, check the switch to make sure the port is not disabled. Check the switch vendor troubleshooting documentation for other possibilities. If the port is now online, the DMP is complete.
3. Reseat the cable. This includes plugging in the cable and SFP if not already done. If the port is now online, the DMP is complete.
4. Reseat the hot swap SFPs (optics modules). If the port is now online, the DMP is complete.
5. Try using a new cable.
6. Try using a new SFP.
7. Try using a new port on the switch.

   Note: Continuing from here will affect other ports connected on the adapter.
8. Reset the adapter.
9. Reset the node.

1404   **Cloud gateway service restarted too often**

Explanation: The system reported a persistent error with the cloud gateway service. Cloud storage functions are not available.

User response: Try the following actions:
1. Check the IP network. For example, ensure that all network switches report good status.
2. Update the system to the latest code.
3. If the problem persists, contact your service support representative.

1450   **Fewer Fibre Channel I/O ports operational.**

Explanation: One or more Fibre Channel I/O ports that have previously been active are now inactive. This situation has continued for one minute.

A Fibre Channel I/O port might be established on either a Fibre Channel platform port or an Ethernet platform port using FCoE. This error is expected if the associated Fibre Channel or Ethernet port is not operational.

Data:

Three numeric values are listed:

- The ID of the first unexpected inactive port. This ID is a decimal number.
- The ports that are expected to be active, which is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is expected to be active.
- The ports that are actually active, which is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is active.

User response:
1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Follow the procedure for mapping I/O ports to platform ports to determine which platform port is providing this I/O port.
3. Check for any 704 (`Fibre channel platform port not operational`) or 724 (`Ethernet platform port not operational`) node errors reported for the platform port.
4. Possibilities:
   - If the port has been intentionally disconnected, use the management GUI recommended action for the service error code and acknowledge the intended change.
   - Resolve the 704 or 724 error.
   - If this is an FCoE connection, use the information the view gives about the Fibre Channel forwarder (FCF) to troubleshoot the connection between the port and the FCF.

Possible Cause-FRUs or other cause:

- None

1471   **Interface card is unsupported.**

Explanation: Interface adapter is unsupported.

User response: Replace the wrong interface adapter with the correct type.

Possible Cause-FRUs or other:

Interface adapter (100%)

1472   **Boot drive is in an unsupported slot.**

Explanation: Boot drive is in an unsupported slot.

User response: Complete the following steps:
1. Look at a boot drive view to determine which drive is in an unsupported slot.

2. Move the drive back to its correct node and slot, but shut down the node first if `booted yes` is shown for that drive in boot drive view.

3. The node error clears, or a new node error is displayed for you to work on.

Possible Cause-FRUs or other:

• None

1473   **The installed battery is at a hardware revision level that is not supported by the current code level.**

Explanation: The installed battery is at a hardware revision level that is not supported by the current code level.

User response: To replace the battery with one that is supported by the current code level, follow the service action for "1130" on page 347. To update the code level to one that supports the currently installed battery, perform a service mode code update. Always install the latest level of the system software to avoid problems with upgrades and component compatibility.

Possible Cause-FRUs or other:

• Battery (50%)

1474   **Battery is nearing end of life.**

Explanation: When a battery nears the end of its life, you must replace it if you intend to preserve the capacity to failover power to batteries.

User response: Replace the battery by following this procedure as soon as you can.

If the node is in a clustered system, ensure that the battery is not being relied upon to provide data protection before you remove it. Issue the **chnodebattery -remove -battery** *battery_ID node_ID* command to establish the lack of reliance on the battery.

If the command returns with a "The command has failed because the specified battery is offline"(BATTERY_OFFLINE) error, replace the battery immediately.

If the command returns with a "The command has failed because the specified battery is not redundant" (BATTERY_NOT_REDUNDANT) error, do not remove the relied-on battery. Removing the battery compromises data protection.

In this case, without other battery-related errors, use the **chnodebattery -remove -battery** *battery_ID node_ID* command periodically to force the system to remove reliance on the battery. The system often removes reliance within one hour (TBC).

Alternatively, remove the node from the clustered system. Once the node is independent, you can replace its battery immediately. If the node is not part of a

cluster, or the battery is offline, or the **chnodebattery** command returns without error, conduct the service action for "1130" on page 347.

Possible Cause-FRUs or other:

• Battery (100%)

1475   **Battery is too hot.**

Explanation: Battery is too hot.

User response: The battery might be slow to cool if the ambient temperature is high. You must wait for the battery to cool down before it can resume its normal operation.

If node error 768 is reported, service that as well.

1476   **Battery is too cold.**

Explanation: You must wait for the battery to warm before it can resume its normal operation.

User response: The battery might be slow to warm if the ambient temperature is low. If node error 768 is reported, service that as well.

Otherwise, wait for the battery to warm.

1550   **A cluster path has failed.**

Explanation: One of the Fibre Channel ports is unable to communicate with all of the other ports in the cluster.

User response:

1. Check for incorrect switch zoning.

2. Repair the fault in the Fibre Channel network fabric.

3. Check the status of the node ports that are not excluded via the system's local port mask. If the status of the node ports shows as active, mark the error that you have repaired as "fixed". If any node ports do not show a status of active, go to start MAP. If you return to this step contact your support center to resolve the problem.

4. Go to repair verification MAP.

Possible Cause-FRUs or other:

• None

Other:

Fibre Channel network fabric fault (100%)

1570   **Quorum disk configured on controller that has quorum disabled**

Explanation: This error can occur with a storage controller that can be accessed through multiple WWNNs and have a default setting of not allowing quorum disks. When these controllers are detected by a

cluster, although multiple component controller definitions are created, the cluster recognizes that all of the component controllers belong to the same storage system. To enable the creation of a quorum disk on this storage system, all of the controller components must be configured to allow quorum.

A configuration change to the SAN, or to a storage system with multiple WWNNs, might result in the cluster discovering new component controllers for the storage system. These components will take the default setting for allowing quorum. This error is reported if there is a quorum disk associated with the controller and the default setting is not to allow quorum.

User response:
- Determine if there should be a quorum disk on this storage system. Ensure that the controller supports quorum before you allow quorum disks on any disk controller. You can check www.ibm.com/support for more information.
- If a quorum disk is required on this storage system, allow quorum on the controller component that is reported in the error. If the quorum disk should not be on this storage system, move it elsewhere.
- Mark the error as "fixed".

Possible Cause-FRUs or other:
- None

Other:

Fibre Channel network fabric fault (100%)

1580   Hostname cannot be resolved

Explanation: The system cannot determine the IP address to connect to.

User response: Try the following actions to determine the source of the problem:
1. Verify that the configured DNS server settings are correct.
   a. Check the output from the `lsdnserver` command and verify that the configured IP addresses are correct.
   b. Try to ping the configured DNS servers by entering `svctask ping dns_server`.
   c. If the ping command fails, enter `sainfo traceroute dns_server` and save the output. Contact your service support representative.
2. Verify that DNS is working by entering `sainfo host www.example.com`.
3. Verify the host name by entering `sainfo host host_name` where *host_name* is the name of the host for which the error was raised. If the system is able to resolve this host name, the issue is now resolved. Manually mark the alert as fixed.

4. If the system cannot resolve the host name, contact your service support representative.

1585   Could not connect to DNS server

Explanation: An invalid DNS server IP was provided, or the DNS server was unresponsive.

User response: Try the following actions:
1. Check the output from the `lsdnserver` command and verify that the configured IP addresses are correct.
2. Try to ping the configured DNS servers by entering `svctask ping dns_server`.
3. If the ping command fails, enter `sainfo traceroute dns_server` and save the output. Contact your service support representative.

1590   Invalid hostname specified

Explanation: An invalid host name was specified, or the DNS server was not able to resolve the host name in its database.

User response: Try the following actions:
1. Check that the host name looks correct.
2. Try to ping the host by entering `svctask ping host_name`.
3. Verify that DNS is working by entering `sainfo host www.example.com`.
4. Verify the host name by entering `sainfo host host_name`. If the system is able to resolve this host name, the issue is now resolved. Manually mark the alert as fixed.
5. If the system cannot resolve the host name, contact your service support representative.

1600   Mirrored disk repair halted because of difference.

Explanation: During the repair of a mirrored volume two copy disks were found to contain different data for the same logical block address (LBA). The validate option was used, so the repair process has halted.

Read operations to the LBAs that differ might return the data of either volume copy. Therefore it is important not to use the volume unless you are sure that the host applications will not read the LBAs that differ or can manage the different data that potentially can be returned.

User response: Perform one of the following actions:
- Continue the repair starting with the next LBA after the difference to see how many differences there are for the whole mirrored volume. This can help you decide which of the following actions to take.
- Choose a primary disk and run repair resynchronizing differences.

- Run a repair and create medium errors for differences.
- Restore all or part of the volume from a backup.
- Decide which disk has correct data, then delete the copy that is different and re-create it allowing it to be synchronized.

Then mark the error as "fixed".

Possible Cause-FRUs or other:

- None

1610   There are too many copied media errors on a managed disk.

Explanation: The cluster maintains a virtual medium error table for each MDisk. This table is a list of logical block addresses on the managed disk that contain data that is not valid and cannot be read. The virtual medium error table has a fixed length. This error event indicates that the system has attempted to add an entry to the table, but the attempt has failed because the table is already full.

There are two circumstances that will cause an entry to be added to the virtual medium error table:

1. FlashCopy, data migration and mirrored volume synchronization operations copy data from one managed disk extent to another. If the source extent contains either a virtual medium error or the RAID controller reports a real medium error, the system creates a matching virtual medium error on the target extent.

2. The mirrored volume validate and repair process has the option to create virtual medium errors on sectors that do not match on all volume copies. Normally zero, or very few, differences are expected; however, if the copies have been marked as synchronized inappropriately, then a large number of virtual medium errors could be created.

User response: Ensure that all higher priority errors are fixed before you attempt to resolve this error.

Determine whether the excessive number of virtual medium errors occurred because of a mirrored disk validate and repair operation that created errors for differences, or whether the errors were created because of a copy operation. Follow the corresponding option shown below.

1. If the virtual medium errors occurred because of a mirrored disk validate and repair operation that created medium errors for differences, then also ensure that the volume copies had been fully synchronized prior to starting the operation. If the copies had been synchronized, there should be only a few virtual medium errors created by the validate and repair operation. In this case, it might be possible to rewrite only the data that was not consistent on the copies using the local data recovery process. If the copies had not been

synchronized, it is likely that there are now a large number of medium errors on all of the volume copies. Even if the virtual medium errors are expected to be only for blocks that have never been written, it is important to clear the virtual medium errors to avoid inhibition of other operations. To recover the data for all of these virtual medium errors it is likely that the volume will have to be recovered from a backup using a process that rewrites all sectors of the volume.

2. If the virtual medium errors have been created by a copy operation, it is best practice to correct any medium errors on the source volume and to not propagate the medium errors to copies of the volume. Fixing higher priority errors in the event log would have corrected the medium error on the source volume. Once the medium errors have been fixed, you must run the copy operation again to clear the virtual medium errors from the target volume. It might be necessary to repeat a sequence of copy operations if copies have been made of already copied medium errors.

An alternative that does not address the root cause is to delete volumes on the target managed disk that have the virtual medium errors. This volume deletion reduces the number of virtual medium error entries in the MDisk table. Migrating the volume to a different managed disk will also delete entries in the MDisk table, but will create more entries on the MDisk table of the MDisk to which the volume is migrated.

Possible Cause-FRUs or other:

- None

1620   A storage pool is offline.

Explanation: A storage pool is offline.

User response:

1. Repair the faults in the order shown.

2. Start a cluster discovery operation by rescanning the Fibre Channel network.

3. Check managed disk (MDisk) status. If all MDisks show a status of "online", mark the error that you have just repaired as "fixed". If any MDisks do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the disk controller.

4. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Fibre Channel network fabric fault (50%)
- Enclosure/controller fault (50%)

1623   One or more MDisks on a controller are degraded.

Explanation: At least one MDisk on a controller is degraded because the MDisk is not available through one or more nodes. The MDisk is available through at least one node. Access to data might be lost if another failure occurs.

In a correctly configured system, each node accesses all of the MDisks on a controller through all of the controller's ports.

This error is only logged once per controller. There might be more than one MDisk on this controller that has been configured incorrectly, but the error is only logged for one MDisk.

To prevent this error from being logged because of short-term fabric maintenance activities, this error condition must have existed for one hour before the error is logged.

User response:
1. Determine which MDisks are degraded. Look for MDisks with a path count lower than the number of nodes. Do not use only the MDisk status, since other errors can also cause degraded MDisks.
2. Ensure that the controller is zoned correctly with all of the nodes.
3. Ensure that the logical unit is mapped to all of the nodes.
4. Ensure that the logical unit is mapped to all of the nodes using the same LUN.
5. Run the console or CLI command to discover MDisks and ensure that the command completes.
6. Mark the error that you have just repaired as "fixed". When you mark the error as "fixed", the controller's MDisk availability is tested and the error will be logged again immediately if the error persists for any MDisks. It is possible that the new error will report a different MDisk.
7. Go to repair verification MAP.

Possible Cause-FRUs or other:
• None

Other:
• Fibre Channel network fabric fault (50%)
• Enclosure/controller fault (50%)

1624   Controller configuration has unsupported RDAC mode.

Explanation: The cluster has detected that an IBM DS series disk controller's configuration is not supported by the cluster. The disk controller is operating in RDAC mode. The disk controller might appear to be operating with the cluster; however, the configuration is

unsupported because it is known to not work with the cluster.

User response:
1. Using the IBM DS series console, ensure that the host type is set to 'IBM TS SAN VCE' and that the AVT option is enabled. (The AVT and RDAC options are mutually exclusive).
2. Mark the error that you have just repaired as "fixed". If the problem has not been fixed it will be logged again; this could take a few minutes.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:
• None

Other:
• Enclosure/controller fault

1625   Incorrect disk controller configuration.

Explanation: While running an MDisk discovery, the cluster has detected that a disk controller's configuration is not supported by the cluster. The disk controller might appear to be operating with the cluster; however, the configuration detected can potentially cause issues and should not be used. The unsupported configuration is shown in the event data.

User response:
1. Use the event data to determine changes required on the disk controller and reconfigure the disk controller to use a supported configuration.
2. Mark the error that you have just repaired as "fixed". If the problem has not been fixed it will be logged again by the managed disk discovery that automatically runs at this time; this could take a few minutes.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:
• None

Other:
• Enclosure/controller fault

1627   The cluster has insufficient redundancy in its controller connectivity.

Explanation: The cluster has detected that it does not have sufficient redundancy in its connections to the disk controllers. This means that another failure in the SAN could result in loss of access to the application data. The cluster SAN environment should have redundant connections to every disk controller. This redundancy allows for continued operation when there is a failure in one of the SAN components.

To provide recommended redundancy, a cluster should be configured so that:

- each node can access each disk controller through two or more different initiator ports on the node.
- each node can access each disk controller through two or more different controller target ports. **Note:** Some disk controllers only provide a single target port.
- each node can access each disk controller target port through at least one initiator port on the node.

If there are no higher-priority errors being reported, this error usually indicates a problem with the SAN design, a problem with the SAN zoning or a problem with the disk controller.

If there are unfixed higher-priority errors that relate to the SAN or to disk controllers, those errors should be fixed before resolving this error because they might indicate the reason for the lack of redundancy. Error codes that must be fixed first are:

- 1210 Local FC port excluded
- 1230 Login has been excluded

Note: This error can be reported if the required action, to rescan the Fibre Channel network for new MDisks, has not been performed after a deliberate reconfiguration of a disk controller or after SAN rezoning.

The 1627 error code is reported for a number of different error IDs. The error ID indicates the area where there is a lack of redundancy. The data reported in an event log entry indicates where the condition was found.

The meaning of the error IDs is shown below. For each error ID the most likely reason for the condition is given. If the problem is not found in the suggested areas, check the configuration and state of all of the SAN components (switches, controllers, disks, cables and cluster) to determine where there is a single point of failure.

010040 A disk controller is only accessible from a single node port.

- A node has detected that it only has a connection to the disk controller through exactly one initiator port, and more than one initiator port is operational.
- The error data indicates the device WWNN and the WWPN of the connected port.
- A zoning issue or a Fibre Channel connection hardware fault might cause this condition.

010041 A disk controller is only accessible from a single port on the controller.

- A node has detected that it is only connected to exactly one target port on a disk controller, and more than one target port connection is expected.
- The error data indicates the WWPN of the disk controller port that is connected.
- A zoning issue or a Fibre Channel connection hardware fault might cause this condition.

010042 Only a single port on a disk controller is accessible from every node in the cluster.

- Only a single port on a disk controller is accessible to every node when there are multiple ports on the controller that could be connected.
- The error data indicates the WWPN of the disk controller port that is connected.
- A zoning issue or a Fibre Channel connection hardware fault might cause this condition.

010043 A disk controller is accessible through only half, or less, of the previously configured controller ports.

- Although there might still be multiple ports that are accessible on the disk controller, a hardware component of the controller might have failed or one of the SAN fabrics has failed such that the operational system configuration has been reduced to a single point of failure.
- The error data indicates a port on the disk controller that is still connected, and also lists controller ports that are expected but that are not connected.
- A disk controller issue, switch hardware issue, zoning issue or cable fault might cause this condition.

010044 A disk controller is not accessible from a node.

- A node has detected that it has no access to a disk controller. The controller is still accessible from the partner node in the I/O group, so its data is still accessible to the host applications.
- The error data indicates the WWPN of the missing disk controller.
- A zoning issue or a cabling error might cause this condition.

010117 A disk controller is not accessible from a node allowed to access the device by site policy

- A disk controller is not accessible from a node that is allowed to access the device by site policy. If a disk controller has multiple WWNNs, the disk controller may still be accessible to the node through one of the other WWNNs.
- The error data indicates the WWNN of the inaccessible disk controller.
- A zoning issue or a fibre channel connection hardware fault might cause this condition.

User response:

1. Check the error ID and data for a more detailed description of the error.

2. Determine if there has been an intentional change to the SAN zoning or to a disk controller configuration that reduces the cluster's access to the indicated disk controller. If either action has occurred, continue with step 8.

3. Use the GUI or the CLI command **lsfabric** to ensure that all disk controller WWPNs are reported as expected.

4. Ensure that all disk controller WWPNs are zoned appropriately for use by the cluster.

5. Check for any unfixed errors on the disk controllers.

6. Ensure that all of the Fibre Channel cables are connected to the correct ports at each end.

7. Check for failures in the Fibre Channel cables and connectors.

8. When you have resolved the issues, use the GUI or the CLI command **detectmdisk** to rescan the Fibre Channel network for changes to the MDisks. **Note:** Do not attempt to detect MDisks unless you are sure that all problems have been fixed. Detecting MDisks prematurely might mask an issue.

9. Mark the error that you have just repaired as fixed. The cluster will revalidate the redundancy and will report another error if there is still not sufficient redundancy.

10. Go to MAP 5700: Repair verification.

Possible Cause-FRUs or other:

• None

1630   The number of device logins was reduced.

Explanation: The number of port to port fabric connections, or logins, between the node and a storage controller has decreased. This situation might be caused by a problem on the SAN or by a deliberate reconfiguration of the SAN.

The 1630 error code is reported for a number of different error IDs. The error ID indicates more specifics about the problem. The data reported in an event log entry indicates where the condition was found.

010045 Number of Device paths from the controller site allowed accessible nodes has reduced

• The controller now has fewer logins from the controller site that allocated accessible nodes to the storage controller.

• The error data indicates the WWNN or IP address of the disk controller, and the current path count from each node.

• A controller fault or a Fibre Channel network fabric fault might cause this condition.

User response:

1. Check the error in the cluster event log to identify the object ID associated with the error.

2. Check the availability of the failing device using the following command line: **lscontroller object_ID**. If the command fails with the message "CMMVC6014E The command failed because the requested object is either unavailable or does not exist," ask the customer if this device was removed from the system.

    • If "yes", mark the error as fixed in the cluster event log and continue with the repair verification MAP.

    • If "no" or if the command lists details of the failing controller, continue with the next step.

3. Check whether the device has regained connectivity. If it has not, check the cable connection to the remote-device port.

4. If all attempts to log in to a remote-device port have failed and you cannot solve the problem by changing cables, check the condition of the remote-device port and the condition of the remote device.

5. Start a cluster discovery operation by rescanning the Fibre Channel network.

6. Check the status of the disk controller. If all disk controllers show a "good" status, mark the error that you have just repaired as "fixed". If any disk controllers do not show "good" status, go to the start MAP. If you return to this step, contact the support center to resolve the problem with the disk controller.

7. Go to repair verification MAP.

Possible Cause-FRUs or other:

• None

Other:

• Fibre Channel network fabric fault (50%)

• Enclosure/controller fault (50%)

1656   Cloud account not available, encryption setting mismatch

Explanation: The system encountered a mismatch between cloud object storage and cluster encryption state. Cloud backup services remain unavailable until this alert is fixed. The associated alert code gives more information.

087016 Cloud account not available, cloud object storage encrypted

The cloud object data is encrypted and the cluster cloud account is not configured with encryption enabled.

087017 Cloud account not available, cloud object storage not encrypted

The cloud data is not encrypted and the cluster cloud account is configured with encryption enabled.

User response: Ensure that you specified the correct cloud account. If not, retry the command with the correct account.

You cannot change the encryption setting for the cloud account. If the specified cloud account is correct, you must delete the account by using the `rmcloudaccount` command and re-create the account by using the `mkcloudaccount` command, this time with an encryption setting that matches the setting for the cloud data.

1657	Cloud account not available, cloud object storage encrypted with the wrong key

Explanation: The master key that is associated with the cloud data does not match the cluster master key that was used when the cluster cloud account was created. Cloud backup services remain unavailable until this alert is fixed.

The error code is associated with the following alert event:

087018 Cloud account not available, cloud object storage encrypted with the wrong key

User response: Compete the following steps:

1. Make the correct master key available in one of the following ways:
   - Insert a USB drive that contains the key
   - Ensure that the system is attached to a Network Key Server that contains the key.
2. Run the `testcloudaccount` command. If the command completes with good status, mark the error as fixed.
3. If the command does not complete with good status, contact your service support representative.

1660	The initialization of the managed disk has failed.

Explanation: The initialization of the managed disk has failed.

User response:

1. View the event log entry to identify the managed disk (MDisk) that was being accessed when the problem was detected.
2. Perform the disk controller problem determination and repair procedures for the MDisk identified in step 1.
3. Include the MDisk into the cluster.
4. Check the managed disk status. If all managed disks show a status of "online", mark the error that you have just repaired as "fixed". If any managed

disks do not show a status of "online", go to the start MAP. If you return to this step, contact your support center to resolve the problem with the disk controller.

5. Go to repair verification MAP.

Possible Cause-FRUs or other:
- None

Other:

Enclosure/controller fault (100%)

1670	The CMOS battery on the system board failed.

Explanation: The CMOS battery on the system board failed.

User response: Replace the node until the FRU is available.

Possible Cause-FRUs or other:

CMOS battery (100%)

1680	Drive fault type 1

Explanation: Drive fault type 1

User response: Replace the drive.

Possible Cause-FRUs or other:

Drive (95%)

Canister (3%)

Midplane (2%)

1684	Drive is missing.

Explanation: Drive is missing.

User response: Install the missing drive. The drive is typically a data drive that was previously part of the array.

Possible Cause-FRUs or other:

Drive (100%)

1686	Drive fault type 3.

Explanation: Drive fault type 3.

User response: Complete the following steps to resolve this problem.

1. Reseat the drive.
2. Replace the drive.
3. Replace the canister as identified in the sense data.
4. Replace the enclosure.

Note: The removal of the exclusion on the drive slot will happen automatically, but only after this error has been marked as fixed.

Possible Cause-FRUs or other:

- Drive (46%)
- Canister (46%)
- Enclosure (8%)

1689   Array MDisk has lost redundancy.

Explanation: Array MDisk has lost redundancy. The RAID 5 system is missing a data drive.

User response: Replace the missing or failed drive.

Possible Cause-FRUs or other:

Drives removed or failed (100%)

1690   No spare protection exists for one or more array MDisks.

Explanation: The system spare pool cannot immediately provide a spare of any suitability to one or more arrays.

User response:

1. Configure an array but no spares.
2. Configure many arrays and a single spare. Cause that spare to be consumed or change its use.

For a distributed array, unused or candidate drives are converted into array members.

1. Decode/explain the number of rebuild areas available and the threshold set.
2. Check for unfixed higher priority errors.
3. Check for unused and candidate drives that are suitable for the distributed array. Run the `lsarraymembergoals` command to determine drive suitability by using tech_type, capacity, and rpm information.
   - Offer to add the drives into the array. Allow up to the number of missing array members to be added.
   - Recheck after array members are added.
4. If no drives are available, explain that drives need to be added to restore the wanted number of rebuild areas.
   - If the threshold is greater than the number of rebuild areas available, and the threshold is greater than 1, offer to reduce the threshold to the number of drives that are available.

1691   A background scrub process has found an inconsistency between data and parity on the array.

Explanation: The array has at least one stride where the data and parity do not match. RAID has found an inconsistency between the data stored on the drives and the parity information. This could either mean that the data has been corrupted, or that the parity information has been corrupted.

User response: Follow the directed maintenance procedure for inconsistent arrays.

1692   Array MDisk has taken a spare member that does not match array goals.

Explanation:

1. A member of the array MDisk either has technology or capability that does not match exactly with the established goals of the array.
2. The array is configured to want location matches, and the drive location does not match all the location goals.

User response: The error will fix itself automatically as soon as the rebuild or exchange is queued up. It does not wait until the array is showing `balanced = exact` (which indicates that all populated members have exact capability match and exact location match).

1693   Drive exchange required.

Explanation: Drive exchange required.

User response: Complete the following steps to resolve this problem.

1. Exchange the failed drive.

Possible Cause-FRUs or other:

- Drive (100%)

1695   Persistent unsupported disk controller configuration.

Explanation: A disk controller configuration that might prevent failover for the cluster has persisted for more than four hours. The problem was originally logged through a 010032 event, service error code 1625.

User response:

1. Fix any higher priority error. In particular, follow the service actions to fix the 1625 error indicated by this error's root event. This error will be marked as "fixed" when the root event is marked as "fixed".
2. If the root event cannot be found, or is marked as "fixed", perform an MDisk discovery and mark this error as "fixed".
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Enclosure/controller fault

1700   **Unrecovered remote copy relationship**

Explanation: This error might be reported after the recovery action for a clustered system failure or a complete I/O group failure. The error is reported because some remote copy relationships, whose control data is stored by the I/O group, could not be recovered.

User response: To fix this error it is necessary to delete all of the relationships that might not be recovered, and then re-create the relationships.

1. Note the I/O group index against which the error is logged.

2. List all of the relationships that have either a master or an auxiliary volume in this I/O group. Use the volume view to determine which volumes in the I/O group you noted have a relationship that is defined.

3. Note the details of the relationships that are listed so that they can be re-created.

   If the affected I/O group has active-active relationships that are in a consistency group, run the command `chrcrelationship -noconsistgrp rc_rel_name` for each active-active relationship that was not recovered. Then, use the command `lsrcrelatioship` in case volume labels are changed and to see the value of the primary attributes.

4. Delete all of the relationships that are listed in step 2, except any active-active relationship that has host applications that use the auxiliary volume via the master volume unique ID. (that is, the primary attribute value is auxiliary in the output from `lsrcrelationship`).

   For the active-active relationships that have the primary attribute value of auxiliary, use the `rmvolumecopy` CLI command (which also deletes the relationship). For example, `rmvolumecopy` *master_volume_id/name*.

   Note: The error is automatically marked as "fixed" once the last relationship on the I/O group is deleted. New relationships must not be created until the error is fixed.

5. Re-create all the relationships that you deleted by using the details noted in step 3.

   Note: For Metro Mirror and Global Mirror relationships, you are able to delete a relationship from either the master or auxiliary system; however, you must re-create the relationship on the master system. Therefore, it might be necessary to go to another system to complete this service action.

Possible Cause-FRUs or other:
- None

1710   **There are too many cluster partnerships. The number of cluster partnerships has been reduced.**

Explanation: A cluster can have a Metro Mirror and Global Mirror cluster partnership with one or more other clusters. Partnership sets consist of clusters that are either in direct partnership with each other or are in indirect partnership by having a partnership with the same intermediate cluster. The topology of the partnership set is not fixed; the topology might be a star, a loop, a chain or a mesh. The maximum supported number of clusters in a partnership set is four. A cluster is a member of a partnership set if it has a partnership with another cluster in the set, regardless of whether that partnership has any defined consistency groups or relationships.

These are examples of valid partnership sets for five unique clusters labelled A, B, C, D, and E where a partnership is indicated by a dash between two cluster names:

- A-B, A-C, A-D. E has no partnerships defined and therefore is not a member of the set.
- A-B, A-D, B-C, C-D. E has no partnerships defined and therefore is not a member of the set.
- A-B, B-C, C-D. E has no partnerships defined and therefore is not a member of the set.
- A-B, A-C, A-D, B-C, B-D, C-D. E has no partnerships defined and therefore is not a member of the set.
- A-B, A-C, B-C. D-E. There are two partnership sets. One contains clusters A, B, and C. The other contains clusters D and E.

These are examples of unsupported configurations because the number of clusters in the set is five, which exceeds the supported maximum of four clusters:

- A-B, A-C, A-D, A-E.
- A-B, A-D, B-C, C-D, C-E.
- A-B, B-C, C-D, D-E.

The cluster prevents you from creating a new Metro Mirror and Global Mirror cluster partnership if a resulting partnership set would exceed the maximum of four clusters. However, if you restore a broken link between two clusters that have a partnership, the number of clusters in the set might exceed four. If this occurs, Metro Mirror and Global Mirror cluster partnerships are excluded from the set until only four clusters remain in the set. A cluster partnership that is excluded from a set has all of its Metro Mirror and Global Mirror cluster partnerships excluded.

Event ID 0x050030 is reported if the cluster is retained in the partnership set. Event ID 0x050031 is reported if the cluster is excluded from the partnership set. All clusters that were in the partnership set report error 1710.

All inter-cluster Metro Mirror or Global Mirror relationships that involve an excluded cluster will lose connectivity. If any of these relationships are in the consistent_synchronized state and they receive a write I/O, they will stop with error code 1720.

User response: To fix this error it is necessary to delete all of the relationships that could not be recovered and then re-create the relationships.

1. Determine which clusters are still connected and members of the partnership set, and which clusters have been excluded.
2. Determine the Metro Mirror and Global Mirror relationships that exist on those clusters.
3. Determine which of the Metro Mirror and Global Mirror relationships you want to maintain, which determines which cluster partnerships you want to maintain. Ensure that the partnership set or sets that would result from configuring the cluster partnerships that you want contain no more than four clusters in each set. NOTE: The reduced partnership set created by the cluster might not contain the clusters that you want in the set.
4. Remove all of the Metro Mirror and Global Mirror relationships that you do not want to retain.
5. Remove all of the Metro Mirror and Global Mirror cluster partnerships that you do not want to retain.
6. Restart all relationships and consistency groups that were stopped.
7. Go to repair verification MAP.

Possible Cause-FRUs or other:

• None

**1720   Metro Mirror (remote copy) -
               Relationship has stopped and lost
               synchronization, for reason other than a
               persistent I/O error (LSYNC)**

Explanation: A remote copy relationship or consistency group needs to be restarted. In a Metro Mirror (remote copy) or Global Mirror operation, the relationship has stopped and lost synchronization, for a reason other than a persistent I/O error.

User response: The administrator must examine the state of the system to validate that everything is online to allow a restart to work. Examining the state of the system also requires checking the partner Fibre Channel (FC) port masks on both clusters.

1. If the partner FC port mask was changed recently, check that the correct mask was selected.
2. Perform whatever steps are needed to maintain a consistent secondary volume, if desired.
3. The administrator must issue a start command.

Possible Cause-FRUs or other:

• None

1740   Recovery encryption key not available.

Explanation: Recovery encryption key is not available.

User response: Make the recovery encryption key available.

1. If the key is not available:
   • Install a USB drive with the encryption key.
   • Ensure correct file is on the USB drive.
2. If the key is not valid:
   • Get a USB drive with a valid key for this MTMS. The key does not have a valid CRC.

Possible Cause-FRUs or other:

No FRU

1741   Flash module is predicted to fail.

Explanation: The Flash module is predicted to fail due to low health (event ID 085023) or due to an encryption issue (event ID 085158). In either case, the drive should be replaced.

User response: A replacement drive of the same size is needed to correct this error.

If any higher array events exist, correct those first.

If no other array events exist, replace the drive. If the array is RAID5, replace and format the drive.

If the array is RAID0, correcting this issue will result in a loss of all data. If the data is needed, do the following:

1. Backup all array data.
2. Replace the drives using the **recoverarray** format.
3. Restore array data.

If the array data is not needed, replace the drive(s) using the **recoverarray** format.

1750   Array response time too high.

Explanation: A number of causes can lead to higher-than-usual array response time.

User response:

1. Fix higher priority errors first.
2. Fix any other known errors.
3. Change the array into redundancy mode by using the charray interface.

Possible Cause-FRUs or other:

Environment or configuration issues:

Volume config 30%

Slow drive 30%

Enclosure 20%

SAS port 20%

1780   Encryption key changes are not committed.

Explanation: Changes were made to the encryption key, but the pending changes were not committed. A directed maintenance procedure (DMP) was launched to cancel the changes.

User response: Press **Next** to cancel the pending key changes. Launch the GUI to restart the operation.

1785   A problem occurred with the Key Server

Explanation: The meaning of the error code depends on the associated event code. All of these errors involve the key server validation process, which can be triggered by the **mkkeyserver**, **chkeyserver**, or **testkeyserver** commands, or by the regular validation timer.

086006 Key Server reported KMIP error

While key server validation was running, the server reported a nonzero KMIP error code. Because the key server can report a wide range of KMIP error codes, the sense data includes the following additional information about the error:

- KMIP Error Code
- KMIP Result Status
- KMIP Result Reason
- An error string that contains the KMIP Result Message

086007 Key Server reported vendor information error

While key server validation was running, the server reported one of the following conditions:

- Unsupported type of key server
- Unsupported code level on the key server

086008 Failed to connect to Key Server

While key server validation was running, the node was unable to connect to the key server.

User response: For event code 086006:

1. The key server reported a server-side problem. The sense data of this event includes more details to help pinpoint the problem on the key server. Run the **testkeyserver** command to determine whether the problem is fixed. The **testkeyserver** command either automatically fixes the error, or raises the event again.

2. Check that the cluster certificate was accepted on the key server. For more information, search your product documentation for "Certificates that are used for key servers".

3. Ensure that ISKLM has been configured to use TLS v1.2. Failure to do so can cause an SSL connection error.

For event code 086007:

1. The key server reported that it is running an unsupported software version. Verify that you are using the correct key server and that the IP address, port address, and other characteristics are all correct. If not, use the **chkeyserver** command to change this information. The **chkeyserver** command automatically starts the validation process to confirm that the error is fixed, and either auto-fixes this event or raises it again.

2. Verify that you are using a supported key server type and version. A list of supported key servers is provided in the documentation. The sense data of this event includes the version information reported by the key server.

   - The minimum supported version of Key Management Interoperability Protocol (KMIP) is 1.3.
   - The supported key server type is ISKLM only.
   - The supported versions of ISKLM are 2.6.0.0 and later.

For event code 086008:

1. Check that a service IP address is configured for all nodes in the cluster (IPv4 if you use IPv4 key servers, IPv6 if you use IPv6 key servers). If not, configure these IP addresses and run the **testkeyserver** command. If the **testkeyserver** command is successful, the event is automatically fixed.

2. Confirm that all nodes in the cluster have their Ethernet cable plugged in correctly. If not, plug them in and run the **testkeyserver** command. If the **testkeyserver** command is successful, the event is automatically fixed.

3. Confirm that the IP address and IP port of the key server object is correct. If not, change the key server details by using the **chkeyserver** command. The **chkeyserver** command automatically starts the validation process to confirm that the error is fixed, and either auto-fixes this event or raises it again.

4. Confirm that any SSL certificates for the key server are valid. Certificates must have correct start and end dates and must be in the PEM format.

For event code 086009:

1. Run the **lskeyserver** command to show the current status of the key servers. One of these servers has the **primary** field incorrectly set to **yes**.

2. Determine which server should correctly be designated as primary. Do this on the server side by identifying the IP address and port that points to the real primary server. The primary server has the role of "MASTER" in the replication relationship in SKLM. For more information about this process, refer to your SKLM documentation. If the primary server in the **lskeyserver** command appears to be correct, contact your service support representative.

3. Otherwise, run the following command:

   chkeyserver -primary *server_id*

   where *server_id* is the ID of the correct primary server.

4. The **chkeyserver** command automatically validates the new primary key server. To fix the event, complete one of the following actions:
   - Manually mark the event as fixed by using the **cheventlog -fix** command
   - Wait for the periodic validation of the old primary key server
   - Manually validate the old server by using the **testkeyserver** command

   If the problem persists, contact your service support representative.

1800   The SAN has been zoned incorrectly.

Explanation: This has resulted in more than 512 other ports on the SAN logging into one port of a 2145 node.

User response:
1. Ask the user to reconfigure the SAN.
2. Mark the error as "fixed".
3. Go to repair verification MAP.

Possible Cause-FRUs or other:
- None

Other:
- Fibre Channel (FC) switch configuration error
- FC switch

1801   A node has received too many Fibre Channel logins from another node.

Explanation: This event was logged because the node has received more than sixteen Fibre Channel logins originating from another node. This indicates that the Fibre Channel storage area network that connects the two nodes is not correctly configured.

Data:
- None

User response: Change the zoning and/or Fibre Channel port masking so that no more than 16 logins are possible between a pair of nodes.

See Non-critical node error "888" on page 331 for details.

Use the **lsfabric** command to view the current number of logins between nodes.

Possible Cause-FRUs or other cause:
- None

1802   Fibre Channel network settings

Explanation: Fibre Channel network settings

User response: Follow these troubleshooting steps to reduce the number of hosts that are logged in to the port:
1. Increase the granularity of the switch zoning to reduce unnecessary host port logins.
2. Change switch zoning to spread out host ports across other available ports.
3. Use interfaces with more ports, if not already at the maximum.
4. Scale out by using another FlashSystem enclosure.

Possible Cause-FRUs or other:

No FRU

1804   IB network settings

Explanation: IB network settings

User response: Follow these troubleshooting steps to reduce the number of hosts that are logged in to the port:
1. Increase the granularity of the switch zoning to reduce unnecessary host port logins.
2. Change switch zoning to spread out host ports across other available ports.
3. Use interfaces with more ports, if not already at the maximum.
4. Scale out by using another FlashSystem enclosure.

Possible Cause-FRUs or other:

No FRU

1840   The managed disk has bad blocks.

Explanation: These are "virtual" medium errors which are created when copying a volume where the source has medium errors. During data moves or duplication, such as during a flash copy, an attempt is made to move medium errors; to achieve this, virtual medium errors called "bad blocks" are created. Once a bad block has been created, no attempt will be made to read the underlying data, as there is no guarantee that the old data still exists once the "bad block" is created. Therefore, it is possible to have "bad blocks", and thus medium errors, reported on a target volume, without

medium errors actually existing on the underlying storage. The "bad block" records are removed when the data is overwritten by a host.

Note: On an external controller, this error can only result from a copied medium error.

User response:

1. The support center will direct the user to restore the data on the affected volumes.

2. When the volume data has been restored, or the user has chosen not to restore the data, mark the error as "fixed".

3. Go to repair verification MAP.

Possible Cause-FRUs or other:

• None

1850 Compressed volume copy has bad blocks

Explanation: A system recovery operation was performed, but data on one or more volumes was not recovered; this is normally caused by a combination of hardware faults. If data containing a medium error is copied or migrated to another volume, bad blocks will be recorded. If a host attempts to read the data in any of the bad block regions, the read will fail with a medium error.

User response:

1. The support center will direct the user to restore the data on the affected volumes.

2. When the volume data has been restored, or the user has chosen not to restore the data, mark the error as "fixed".

3. Go to repair verification MAP.

Possible Cause-FRUs or other:

• None

1860 Thin-provisioned volume copy offline because of failed repair.

Explanation: The attempt to repair the metadata of a thin-provisioned volume that describes the disk contents has failed because of problems with the automatically maintained backup copy of this data. The error event data describes the problem.

User response: Delete the thin-provisioned volume and reconstruct a new one from a backup or mirror copy. Mark the error as "fixed". Also mark the original 1862 error as "fixed".

Possible Cause-FRUs or other:

• None

1862 Thin-provisioned volume copy offline because of corrupt metadata.

Explanation: A thin-provisioned volume has been taken offline because there is an inconsistency in the cluster metadata that describes the disk contents. This might occur because of corruption of data on the physical disk (e.g., medium error or data miscompare), the loss of cached metadata (because of a cluster recovery) or because of a software error. The event data gives information on the reason.

The cluster maintains backup copies of the metadata and it might be possible to repair the thin-provisioned volume using this data.

User response: The cluster is able to repair the inconsistency in some circumstances. Run the repair volume option to start the repair process. This repair process, however, can take some time. In some situations it might be more appropriate to delete the thin-provisioned volume and reconstruct a new one from a backup or mirror copy.

If you run the repair procedure and it completes, this error is automatically marked as "fixed"; otherwise, another error event (error code 1860) is logged to indicate that the repair action has failed.

Possible Cause-FRUs or other:

• None

1864 Compressed volume size limitation breached, diagnosis required

Explanation: The system indicates that the virtual or real capacity of at least one compressed volume exceeds the system limits.

User response: For information about how to deal with this issue, see www.ibm.com/support/docview.wss?uid=ssg1S1005731.

1865 Thin-provisioned volume copy offline because of insufficient space.

Explanation: A thin-provisioned volume has been taken offline because there is insufficient allocated real capacity available on the volume for the used space to increase further. If the thin-provisioned volume is auto-expand enabled, then the storage pool it is in also has no free space.

User response: The service action differs depending on whether the thin-provisioned volume copy is auto-expand enabled or not. Whether the disk is auto-expand enabled or not is indicated in the error event data.

If the volume copy is auto-expand enabled, perform one or more of the following actions. When you have performed all of the actions that you intend to perform, mark the error as "fixed"; the volume copy will then return online.

- Determine why the storage pool free space has been depleted. Any of the thin-provisioned volume copies, with auto-expand enabled, in this storage pool might have expanded at an unexpected rate; this could indicate an application error. New volume copies might have been created in, or migrated to, the storage pool.
- Increase the capacity of the storage pool that is associated with the thin-provisioned volume copy by adding more MDisks to the storage pool.
- Provide some free capacity in the storage pool by reducing the used space. Volume copies that are no longer required can be deleted, the size of volume copies can be reduced or volume copies can be migrated to a different storage pool.
- Migrate the thin-provisioned volume copy to a storage pool that has sufficient unused capacity.
- Consider reducing the value of the storage pool warning threshold to give more time to allocate extra space.

If the volume copy is not auto-expand enabled, perform one or more of the following actions. In this case the error will automatically be marked as "fixed", and the volume copy will return online when space is available.

- Determine why the thin-provisioned volume copy used space has grown at the rate that it has. There might be an application error.
- Increase the real capacity of the volume copy.
- Enable auto-expand for the thin-provisioned volume copy.
- Consider reducing the value of the thin-provisioned volume copy warning threshold to give more time to allocate more real space.

Possible Cause-FRUs or other:
- None

1870   Mirrored volume offline because a hardware read error has occurred.

Explanation: While attempting to maintain the volume mirror, a hardware read error occurred on all of the synchronized volume copies.

The volume copies might be inconsistent, so the volume is now offline.

User response:
- Fix all higher priority errors. In particular, fix any read errors that are listed in the sense data. This error event will automatically be fixed when the root event is marked as "fixed".
- If you cannot fix the root error, but the read errors on some of the volume copies have been fixed, mark this error as "fixed" to run without the mirror. You can then delete the volume copy that cannot read data and re-create it on different MDisks.

Possible Cause-FRUs or other:
- None

1895   Unrecovered FlashCopy mappings

Explanation: This error might be reported after the recovery action for a cluster failure or a complete I/O group failure. The error is reported because some FlashCopies, whose control data is stored by the I/O group, were active at the time of the failure and the current state of the mapping could not be recovered.

User response: To fix this error it is necessary to delete all of the FlashCopy mappings on the I/O group that failed.

1. Note the I/O group index against which the error is logged.
2. List all of the FlashCopy mappings that are using this I/O group for their bitmaps. You should get the detailed view of every possible FlashCopy ID. Note the IDs of the mappings whose IO_group_id matches the ID of the I/O group against which this error is logged.
3. Note the details of the FlashCopy mappings that are listed so that they can be re-created.
4. Delete all of the FlashCopy mappings that are listed. Note: The error will automatically be marked as "fixed" once the last mapping on the I/O group is deleted. New mappings cannot be created until the error is fixed.
5. Using the details noted in step 3, re-create all of the FlashCopy mappings that you just deleted.

Possible Cause-FRUs or other:
- None

1900   A FlashCopy, Trigger Prepare command has failed because a cache flush has failed.

Explanation: A FlashCopy, Trigger Prepare command has failed because a cache flush has failed.

User response:
1. Correct higher priority errors, and then try the Trigger Prepare command again.
2. Mark the error that you have just repaired as "fixed".
3. Go to repair verification MAP.

Possible Cause-FRUs or other:
- None

Other:

Cache flush error (100%)

1910   **A FlashCopy mapping task was stopped because of the error that is indicated in the sense data.**

Explanation: A stopped FlashCopy might affect the status of other volumes in the same I/O group. Preparing the stopped FlashCopy operations as soon as possible is advised.

User response:

1. Correct higher priority errors, and then prepare and start the FlashCopy task again.

2. Mark the error that you have just repaired as "fixed".

3. Go to repair verification MAP.

Possible Cause-FRUs or other:

• None

1920   **Global and Metro Mirror persistent error.**

Explanation: This error might be caused by a problem on the primary system, a problem on the secondary system, or a problem on the intersystem link. The problem might be a failure of a component, a component becoming unavailable or having reduced performance because of a service action or it might be that the performance of a component dropped to a level where the Metro Mirror or Global Mirror relationship cannot be maintained. Alternatively the error might be caused by a change in the performance requirements of the applications using Metro Mirror or Global Mirror.

This error is reported on the primary system when the copy relationship has not progressed sufficiently over a period of time. Therefore, if the relationship is restarted before all of the problems are fixed, the error might be reported again when the time period next expires (the default period is five minutes).

This error might also be reported because the primary system has encountered read errors.

You might need to refer to the Copy Services features information in the software installation and configuration documentation while you diagnose this error.

User response:

1. If the 1920 error has occurred previously on Metro Mirror or Global Mirror between the same systems and all the following actions have been attempted, contact your product support center to resolve the problem.

2. On both systems, check the partner fc port mask to ensure that there is sufficient connectivity. If the partner fc port mask was changed recently, ensure that the mask is correct.

3. On the primary system that is reporting the error, correct any higher priority errors.

4. On the secondary system, review the maintenance logs to determine whether the system was operating with reduced capability at the time the error was reported. The reduced capability might be because of a software upgrade, hardware maintenance to a node, maintenance to a backend disk system or maintenance to the SAN.

5. On the secondary system, correct any errors that are not fixed.

6. On the intersystem link, review the logs of each link component for any incidents that would cause reduced capability at the time of the error. Ensure that the problems are fixed.

7. If a reason for the error has been found and corrected, go to Action 11.

8. On the primary system that is reporting the error, examine the statistics using a SAN productivity monitoring tool and confirm that all the Metro Mirror and Global Mirror requirements that are described in the planning documentation are met. Ensure that any changes to the applications that use Metro Mirror or Global Mirror have been taken into account. Resolve any issues.

9. On the secondary system, examine the statistics by using a SAN productivity monitoring tool and confirm that all the Metro Mirror and Global Mirror requirements described in the software installation and configuration documentation are met. Resolve any issues.

10. On the intersystem link, examine the performance of each component by using an appropriate SAN productivity monitoring tool to ensure that they are operating as expected. Resolve any issues.

11. Mark the error as "fixed" and restart the Metro Mirror or Global Mirror relationship.

When you restart the Metro Mirror or Global Mirror relationship there is an initial period during which Metro Mirror or Global Mirror performs a background copy to resynchronize the volume data on the primary and secondary systems. During this period the data on the Metro Mirror or Global Mirror auxiliary volumes on the secondary system is inconsistent and the volumes cannot be used as backup disks by your applications.

Note: To ensure the system has the capacity to handle the background copy load, you might want to delay restarting the Metro Mirror or Global Mirror relationship until there is a quiet period when the secondary system and the SAN fabric (including the intersystem link) have the required capacity. If the required capacity is not available you might experience another 1920 error and the Metro Mirror or Global Mirror relationship will stop in an inconsistent state.

Note: If the Metro Mirror or Global Mirror relationship has stopped in a consistent state ("consistent-stopped") it is possible to use the data on the Metro Mirror or Global Mirror auxiliary volumes on the secondary system as backup disks by your applications. You might therefore want to start a Flash Copy of your Metro Mirror or Global Mirror auxiliary disks on the secondary system before restarting the Metro Mirror or Global Mirror relationship. This means you maintain the current, consistent, image until the time when the Metro Mirror or Global Mirror relationship is again synchronized and in a consistent state.

Possible Cause-FRUs or other:

• None

Other:

• Primary system or SAN fabric problem (10%)
• Primary system or SAN fabric configuration (10%)
• Secondary system or SAN fabric problem (15%)
• Secondary system or SAN fabric configuration (25%)
• Intersystem link problem (15%)
• Intersystem link configuration (25%)

1925   Cached data cannot be destaged.

Explanation: Problem diagnosis is required.

User response:

1. Run the directed maintenance procedure to fix all errors of a higher priority. This will allow the cached data to be destaged and the originating event to be marked fixed.

Possible Cause-FRUs or other:

• None

1930   Migration suspended.

Explanation: Migration suspended.

User response:

1. Ensure that all error codes of a higher priority have already been fixed.
2. Ask the customer to ensure that all storage pools that are the destination of suspended migrate operations have available free extents.
3. Mark this error as "fixed". This causes the migrate operation to be restarted. If the restart fails, a new error is logged.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

• None

1940   HyperSwap volume or consistency group has lost synchronization between sites.

Explanation: HyperSwap volume or consistency group has lost synchronization between sites.

User response: Complete the following steps to resolve this problem.

1. Check the event log for any higher priority unfixed errors.
2. HyperSwap volumes will automatically resynchronize when the underlying problem has been resolved.

Possible Cause-FRUs or other:

• N/A

1950   Unable to mirror medium error.

Explanation: During the synchronization of a mirrored volume copy it was necessary to duplicate the record of a medium error onto the volume copy, creating a virtual medium error. Each managed disk has a table of virtual medium errors. The virtual medium error could not be created because the table is full. The volume copy is in an inconsistent state and has been taken offline.

User response: Three different approaches can be taken to resolving this problem: 1) the source volume copy can be fixed so that it does not contain medium errors, 2) the number of virtual medium errors on the target managed disk can be reduced or 3) the target volume copy can be moved to a managed disk with more free virtual medium error entries.

The managed disk with a full medium error table can be determined from the data of the root event.

Approach 1) - This is the preferred procedure because it restores the source volume copy to a state where all of the data can be read. Use the normal service procedures for fixing a medium error (rewrite block or volume from backup or regenerate the data using local procedures).

Approach 2) - This method can be used if the majority of the virtual medium errors on the target managed disk do not relate to the volume copy. Determine where the virtual medium errors are using the event log events and re-write the block or volume from backup.

Approach 3) - Delete the offline volume copy and create a new one either forcing the use of different MDisks in the storage pool or using a completely different storage pool.

Follow your selection option(s) and then mark the error as "fixed".

Possible Cause-FRUs or other:

• None

2008 **A software downgrade has failed.**

Explanation: Cluster configuration changes are restricted until the downgrade is completed. The cluster downgrade process waits for user intervention when this error is logged.

User response: The action required to recover from a stalled downgrade depends on the current state of the cluster being downgraded. Call IBM Support for an action plan to resolve this problem.

Possible Cause-FRUs or other:

- None

Other:

System software (100%)

2010 **A software update has failed.**

Explanation: Cluster configuration changes are restricted until the update is completed or rolled back. The cluster update process waits for user intervention when this error is logged.

User response: The action required to recover from a stalled update depends on the current state of the cluster being updated. Call IBM technical support for an action plan to resolve this problem.

Possible Cause-FRUs or other:

- None

Other:

System software (100%)

2016 **A host port has more than four logins to a node**

Explanation: More than 4 logins have been made to at least one host port or WWPN on at least one node. The network might not be zoned correctly.

User response: Complete the following steps. If at any point you need additional assistance, contact your service support representative.

1. Create a list of the problem hosts , WWPNs, and nodes:
   a. Run the `svcinfo lsfabric -host` command and parse the output into a human readable format.
   b. Sort by WWPN, then by node.
   c. For any WWPN and node combination that shows more than 4 logins:
      1) Get the host port mask from the mask field of the `lshost` detailed view.
      2) Ignore any row where the local_port field does not match the appropriate bit in the host port mask.

3) Make a note of any hosts that still show more than 4 logins after the host port mask is applied.

2. Fix the issue either by changing the zoning or by changing the host port mask.

3. The event will auto-fix when all of the host ports have login counts of 4 or less on every node.

2020 **IP Remote Copy link unavailable.**

Explanation: IP Remote Copy link is unavailable.

User response: Fix the remote IP link so that traffic can flow correctly. Once the connection is made, the error will auto-correct.

2021 **Partner cluster IP address unreachable.**

Explanation: Partner cluster IP address unreachable.

User response:

1. Verify the system IP address of the remote system forming the partnership.
2. Check if remote cluster IP address is reachable from local cluster. The following can be done to verify accessibility:
   a. Use `svctask` to ping the remote cluster IP address. If the ping works, there may be a block on the specific port traffic that needs to be opened in the network. If the ping does not work, there may be no route between the system. Check the IP gateway configuration on the SAN Volume Controller nodes and the IP network configuration.
   b. Check the configuration of the routers and firewall to ensure that TCP/IP port 3620 used for IP partnership is not blocked.
   c. Use the `ssh` command from another system to attempt to establish a session with the problematic remote cluster IP address to confirm that the remote cluster is operational.

2022 **Cannot authenticate with partner cluster.**

Explanation: Cannot authenticate with partner cluster.

User response: Verify the CHAP secret set of partnership using `mkippartnership` or `chpartnership` CLIs match remote system CHAP secret set using `chsystem` CLI. If they don't match, use appropriate commands to set the right CHAP secrets.

2023 **Unexpected cluster ID for partner cluster.**

Explanation: Unexpected cluster ID for partner cluster.

User response: After deleting all relationships and consistency group, remove the partnership.

This is an unrecoverable error when one of the sites has

undergone a T3 recovery and lost all partnership information. Contact IBM support.

2030   **Software error.**

Explanation: The software has restarted because of a problem in the cluster, on a disk system or on the Fibre Channel fabric.

User response:

1. Collect the software dump file(s) generated at the time the error was logged on the cluster.
2. Contact your product support center to investigate and resolve the problem.
3. Ensure that the software is at the latest level on the cluster and on the disk systems.
4. Use the available SAN monitoring tools to check for any problems on the fabric.
5. Mark the error that you have just repaired as "fixed".
6. Go to repair verification Map.

Possible Cause-FRUs or other:

- Your support center might indicate a FRU based on their problem analysis (2%)

Other:

- Software (48%)
- Enclosure/controller software (25%)
- Fibre Channel switch or switch configuration (25%)

2031   **Cloud gateway service restarted**

Explanation: The system detected that an error occurred with the cloud gateway service and the service was restarted.

User response: Try the following actions:

1. Check the IP network. For example, ensure that all network switches report good status.
2. Update the system to the latest code.
3. If the problem persists, contact your service support representative.

2035   **Drive has disabled protection information support.**

Explanation: An array has been interrupted in the process of establishing data integrity protection information on or more of its members by initial writes or rebuild writes.

In order to ensure the array is usable, the system has turned off hardware data protection for the member drive.

User response: If many or all the member drives in an array have logged this error, and sufficient storage exists in the pool to migrate the allocated extents, then

the simplest strategy is to delete the array and recreate it once the drive service action has been accomplished.

If a small number of drives are affected then it is simplest to remove these drives from the array and service them individually. This option is not possible if the array is currently syncing post recovery.

2040   **A software update is required.**

Explanation: The software cannot determine the VPD for a FRU. Probably, a new FRU has been installed and the software does not recognize that FRU.

User response:

1. If a FRU has been replaced, ensure that the correct replacement part was used. The node VPD indicates which part is not recognized.
2. Ensure that the cluster software is at the latest level.
3. Save dump data with configuration dump and logged data dump.
4. Contact your product support center to resolve the problem.
5. Mark the error that you have just repaired as "fixed".
6. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

System software (100%)

2055   **System reboot required.**

Explanation: A system restart is required.

User response: The software update is not complete. Restart the system.

The system is not available for I/O or systems management during the system reset.

2060   **Manual discharge of batteries required.**

Explanation: Manual discharge of batteries required.

User response: Use `chenclosureslot -battery -slot 1 -recondition on` to cause battery calibration.

2070   **A drive has been detected in an enclosure that does not support that drive.**

Explanation: A drive has been detected in an enclosure that does not support that drive.

User response: Remove the drive. If the result is an invalid number of drives, replace the drive with a valid drive.

Possible Cause-FRUs or other:

Drive (100%)

2100   A software error has occurred.

Explanation: One of the V3700 server software components (sshd, crond, or httpd) has failed and reported an error.

User response:

1. Ensure that the software is at the latest level on the cluster.
2. Save dump data with configuration dump and logged data dump.
3. Contact your product support center to resolve the problem.
4. Mark the error that you have just repaired as "fixed".
5. Go to repair verification MAP.

Possible Cause-FRUs or other:

• None

Other:

V3700 software (100%)

2105   Cloud account not available, cannot access cloud object storage

Explanation: The system encountered a problem in trying to read, write, or search for data in the cloud object storage.

User response: Try the following actions:

1. Mark the error as fixed to retry the operation.
2. Check the cloud provider console for errors, if available.
3. Report the problem to the cloud provider. Include the following information:

   • Check the sense data to determine whether the system was attempting to read, write, or search.
   • Reconstruct the container name from the container prefix in the cloud account object, and the container suffix in the sense data.
   • Check the sense data to learn the BLOB name that the system was working with.

2115   Performance of external MDisk has changed

Explanation: The system identified a change in the performance category of an external MDisk. A storage device in the external system might have been replaced with a device that has different performance characteristics to the original. The ID of the MDisk is logged in the event (Bytes 5-8 of the sense data). It might be necessary to re-configure the tier of the

MDisk so that EasyTier makes best use of the storage.

User response: Run the fix procedure for this event, assisting you with the following tasks:

1. Run the **Detect MDisks** task, so that the system determines the current performance category of each Mdisk. When the detection task is complete, if performance has reverted, the event is automatically marked as `fixed`.
2. If the event is not automatically fixed, you can change the tier of the MDisk to the recommended tier shown in the event properties. The recommended tier is logged in the event (Bytes 9-13 of the sense data. A value of 10 hex indicates flash tier, a value of 20 hex indicates enterprise tier).
3. If you choose not to change the tier configuration, mark the event as `fixed`.

2120   Internal IO error occurred while doing cloud operation.

Explanation: An internal error occurred while the system was trying to create a cloud snapshot or complete a restore operation. More information is provided by the associated alert event:

• 087026 Internal Read error during cloud snapshot operation
• 087033 Internal write error during cloud snapshot operation

User response: Complete the following steps:

1. Fix for any unfixed errors on the volume where the error was reported or on the volume that was being restored. To determine the name of the volume that was being restored, use the `lsvolumerestoreprogress` command.
2. Mark the error as fixed to have the system retry the operation.
3. If the error persists, contact your service support representative.

2125   Cloud account out of space

Explanation: The operation during which the cloud account ran out of space is indicated by the associated event code:

• 087020 Cloud account out of space during cloud storage snapshot operation
• 087044 Cloud account out of space during cloud snapshot restore commit operation
• 087045 Cloud account out of space during cloud snapshot delete operation

The user response is the same in all cases.

User response: Contact your cloud service provider to add more cloud storage space.

2258   System SSL certificate has expired.

Explanation: System SSL certificate has expired.

Connections to the GUI, service assistant, and CIMOM are likely to generate security exceptions.

User response: Complete the following steps to resolve this problem.

1. Access the CLI by using ssh.
2. Check that the system time and date is correct. If it is incorrect, it can cause the certificate to be incorrectly marked as expired.
3. Create a new self-signed system certificate, or create a certificate request. Get it signed by your certificate authority and install the signed request.

   Note: If it takes some time to get a certificate signed, you can also create a self-signed certificate to use while you wait for your request to be signed.

Possible Cause-FRUs or other:

- N/A

2259   Storwize V7000 Gen1 compatibility mode can now be disabled on this system.

Explanation: No more Storwize V7000 Gen1 canisters are attached to the system.

User response: Complete one of the following actions:

- If you want to disable Storwize V7000 Gen1 compatibility mode, enter the following command:

  `chsystem -gen1compatibilitymode no`

- If you want to maintain Storwize V7000 Gen1 compatibility mode, you can reattach Storwize V7000 Gen1 canisters to the cluster.

2300   Cloud account not available, SSL certificate problem

Explanation: The cloud account is using SSL (https:// URL or Amazon) and a problem was found with the certificate. The most likely outcome is that a new certificate must be installed. The exact meaning of the error code depends on the associated event code.

087007 Cloud account not available, no matching CA certificate

The cloud account provider that is associated with the account presented an SSL certificate. The system cannot access a matching root CA (certificate authority) certificate.

087008 Cloud account not available, expired SSL certificate

The SSL certificate that is installed on the system that is associated with the cloud account is expired or is not

yet active. Cloud backup services remain paused until the alert is fixed.

User response: For event code 087007:

- For a private cloud, contact the administrator of the cloud. Request the CA certificate and install it.
- For a public cloud, it is likely that you need to upgrade your Storwize V7000 software.

For event code 087008:

1. Check the valid_not_before and valid_not_after dates from alert sense data.
2. Verify that the system time is correct.
3. Complete one of the following actions:
   - For a private cloud, contact the administrator of the cloud. Request a new certificate and install it.
   - For a public cloud, you might need to update your Storwize V7000 license. If your license is correct, contact the administrator of the cloud, request a new certificate, and install it.

2305   No authorization to perform cloud operation

Explanation: The cloud account was configured with credentials (for Amazon, AWS access key; for Swift, user/tenant/password) that are not sufficient to use the cloud storage. The system can log in, but the specified user does not have permission to complete one or more of the following operations:

- Upload data. Required to create a cloud snapshot.
- Create a container in cloud storage. Required to create a cloud snapshot.
- Download data. Required to complete a restore operation.
- Delete data. Required to delete a cloud snapshot.

The error code is associated with the following alert event:

087011 Cloud account not available, cannot obtain permission to use cloud storage

User response: Complete the following steps:

1. Use the `lscloudaccount` command to display cloud account information and verify that everything is correct.
2. Verify that the system time is correct. Some cloud providers are sensitive to time differences.
3. Check the cloud service provider console or contact the cloud administrator to confirm that the correct permissions are in place for the user.
4. Fix the alert to retry the cloud operation.

2310	Cloud account not available, cannot contact cloud provider

Explanation: The system cannot make an IP connection over the management network from the config node to the cloud.

User response: Try the following actions:

1. Check for higher-priority unfixed errors. The system might be reporting network errors. Fix these errors first, and this alert might then auto-fix.

2. For a SWIFT cloud account, check the endpoint URL. If this URL is changed to one that is working, the event auto-fixes.

3. Use **ping** or **traceroute** with the cloud endpoint IP address to try to locate where the connection is being lost. For Amazon Web Services, use s3.amazonaws.com as the endpoint address.

2320	Cloud account not available, cannot communicate with cloud provider

Explanation: The local system can make an IP connection to the server, but the server is not replying properly to cloud storage protocol commands. The most likely problem is a configuration error on the local system, such as an IP address that needs updating after the server changed its IP address. The remaining problems are on the server side. This error is most likely to occur with private cloud installations.

User response: Try the following actions:

1. Check your configuration settings. If you change a setting that results in a valid configuration, the event auto-fixes.

2. Contact the cloud service provider administrator.

2330	Cloud account not available, cloud provider login error

Explanation: A problem was reported with the credentials that were submitted to the cloud account object. For Amazon, the credential is an AWS access key. For SWIFT, the credentials consist of a user name, tenant, and password. The meaning of the error code depends on the associated event code.

087010 Cloud account not available, cannot authenticate with cloud provider

The cloud service provider rejected the credentials that are associated with the cloud account. Cloud backup services remain paused until the alert is fixed. For some public cloud providers, including AWS S3, this alert can occur if the system time deviates more than 15 minutes from standard time. This alert can also occur after a full system (T4) recovery if your credentials are lost.

087011 Cloud account not available, cannot obtain permission to use cloud storage

The cloud service provider accepted the credentials that

are associated with the cloud account, but the system is not allowed to run cloud storage operations. Cloud backup services remain paused until the alert is fixed.

User response: For event code 087010:

1. Verify that you are using the correct credentials.

2. Verify that the system time is correct.

3. Contact the cloud service provider to see whether your password was changed on the cloud side.

4. Fix the alert to retry the login.

For event code 087011:

1. Verify that you are using the correct credentials.

2. Contact the cloud service provider to provide sufficient permission for your account.

3. Fix the alert to retry the login.

2500	A secure shell (SSH) session limit for the cluster has been reached.

Explanation: Secure Shell (SSH) sessions are used by applications that manage the cluster. An example of such an application is the command-line interface (CLI). An application must initially log in to the cluster to create an SSH session. The cluster imposes a limit on the number of SSH sessions that can be open at one time. This error indicates that the limit on the number of SSH sessions has been reached and that no more logins can be accepted until a current session logs out.

The limit on the number of SSH sessions is usually reached because multiple users have opened an SSH session but have forgotten to close the SSH session when they are no longer using the application.

User response:

• Because this error indicates a problem with the number of sessions that are attempting external access to the cluster, determine the reason that so many SSH sessions have been opened.

• Run the Fix Procedure for this error on the panel at **Management GUI Troubleshooting** > **Recommended Actions** to view and manage the open SSH sessions.

2550	Encryption key on USB flash drive removed

Explanation: The USB flash drive in a particular node or port has been removed. This USB flash drive contained a valid encryption key for the system. Unauthorized removal can compromise data security.

User response: If your data has been compromised, perform a rekey operation immediately.

2555   **Encryption key error on USB flash drive.**

Explanation: It is necessary to provide an encryption key before the system can become fully operational. This error occurs when the encryption key identified is invalid. A file with the correct name was found but the key in the file is corrupt.

User response: Remove the USB flash drive from the port.

2560   **drive write endurance usage rate high**

Explanation: Flash drives have a limited write endurance. A high usage rate is leading a drive to failure earlier than expected.

User response: Complete the following steps:

1. Check the event log for the ID of the drive with the high usage rate.
2. Run the **lsdrive** command and note the date in the Predicted Failure Date field.
3. If the predicted failure date is approaching, consider replacing the drive.
4. Mark the event as fixed.

2600   **The cluster was unable to send an email.**

Explanation: The cluster has attempted to send an email in response to an event, but there was no acknowledgement that it was successfully received by the SMTP mail server. It might have failed because the cluster was unable to connect to the configured SMTP server, the email might have been rejected by the server, or a timeout might have occurred. The SMTP server might not be running or might not be correctly configured, or the cluster might not be correctly configured. This error is not logged by the test email function because it responds immediately with a result code.

User response:

- Ensure that the SMTP email server is active.
- Ensure that the SMTP server TCP/IP address and port are correctly configured in the cluster email configuration.
- Send a test email and validate that the change has corrected the issue.
- Mark the error that you have just repaired as fixed.
- Go to MAP 5700: Repair verification.

Possible Cause-FRUs or other:

- None

2601   **Error detected while sending an email.**

Explanation: An error has occured while the cluster was attempting to send an email in response to an event. The cluster is unable to determine if the email has been sent and will attempt to resend it. The problem might be with the SMTP server or with the cluster email configuration. The problem might also be caused by a failover of the configuration node. This error is not logged by the test email function because it responds immediately with a result code.

User response:

- If there are higher-priority unfixed errors in the log, fix those errors first.
- Ensure that the SMTP email server is active.
- Ensure that the SMTP server TCP/IP address and port are correctly configured in the cluster email configuration.
- Send a test email and validate that the change has corrected the issue.
- Mark the error that you have just repaired as fixed.
- Go to MAP 5700: Repair verification.

Possible Cause-FRUs or other:

- None

2700   **Unable to access NTP network time server**

Explanation: Cluster time cannot be synchronized with the NTP network time server that is configured.

User response: There are three main causes to examine:

- The cluster NTP network time server configuration is incorrect. Ensure that the configured IP address matches that of the NTP network time server.
- The NTP network time server is not operational. Check the status of the NTP network time server.
- The TCP/IP network is not configured correctly. Check the configuration of the routers, gateways and firewalls. Ensure that the cluster can access the NTP network time server and that the NTP protocol is permitted.

The error will automatically fix when the cluster is able to synchronize its time with the NTP network time server.

Possible Cause-FRUs or other:

- None

2702   **Check configuration settings of the NTP server on the CMM**

Explanation: The node is configured to automatically set the time using an NTP server within the CMM. It is not possible to connect to the NTP server during

authentication. The NTP server configuration cannot be changed within S-ITE. Within the CMM, there are changeable NTP settings. However, these settings configure how the CMM gets the time and date - the internal CMM NTP server that is used by the S-ITE cannot be changed or configured. This event is only raised when an attempt is made to use the server - once every half hour.

Note: The NTP configuration settings are re-read from the CMM before each connection.
The reason for a connection error can be due to the following:

- all suitable Ethernet ports are offline
- the CMM hardware is not operational
- the CMM is active but the CMM NTP server is **offline**.

The reason for an authentication issue can be due to the following:

- the authentication values provided were invalid
- the NTP server rejected the authentication key provided to the node by the CMM.

If the NTP port is an unsupported value, a port error can display. Currently, only port 123 is supported. Only the current configuration node attempts to resync with the server.

User response:

1. Make sure that CMM is operational by logging in and confirming its time.
2. Check that the IP address in the event log can be pinged from the node.
3. If there is an error, try rebooting the CMM.

3000   UPS Ambient temperature high

Explanation: The temperature sensor in the UPS is reporting a temperature that is close to the operational limit of the unit. If the temperature continues to rise, the UPS will power off for safety reasons. The sensor is probably reporting an excessively high temperature because the environment in which the UPS is operating is too hot.

User response:

1. Ensure that the room ambient temperature is within the permitted limits.
2. Ensure that the air vents at the front and back of the UPS are not obstructed.
3. Ensure that other devices in the same rack are not overheating.
4. When you are satisfied that the cause of the overheating has been resolved, mark the error "fixed".

3001   UPS Ambient temperature high

Explanation: The temperature sensor in the UPS is reporting a temperature that is close to the operational limit of the unit. If the temperature continues to rise, the UPS will power off for safety reasons. The sensor is probably reporting an excessively high temperature because the environment in which the UPS is operating is too hot.

User response:

1. Ensure that the room ambient temperature is within the permitted limits.
2. Ensure that the air vents at the front and back of the UPS are not obstructed.
3. Ensure that other devices in the same rack are not overheating.
4. When you are satisfied that the cause of the overheating has been resolved, mark the error "fixed".

**3010   Internal uninterruptible power supply
               software error detected.**

Explanation: Some of the tests that are performed during node startup did not complete because some of the data reported by the uninterruptible power supply during node startup is inconsistent because of a software error in the uninterruptible power supply. The node has determined that the uninterruptible power supply is functioning sufficiently for the node to continue operations. The operation of the cluster is not affected by this error. This error is usually resolved by power cycling the uninterruptible power supply.

User response:

1. Power cycle the uninterruptible power supply at a convenient time. The one or two nodes attached to the uninterruptible power supply should be powered off before powering off the uninterruptible power supply. Once the nodes have powered off, wait 5 minutes for the uninterruptible power supply to go into standby mode (flashing green AC LED). If this does not happen automatically then check the cabling to confirm that all nodes powered by this uninterruptible power supply have been powered off. Remove the power input cable from the uninterruptible power supply and wait at least 2 minutes for the uninterruptible power supply to clear its internal state. Reconnect the uninterruptible power supply power input cable. Press the uninterruptible power supply ON button. Power on the nodes connected to this uninterruptible power supply.
2. If the error is reported again after the nodes are restarted replace the 2145 UPS electronics assembly.

Possible Cause-FRUs or other:

- 2145 UPS electronics assembly (5%)

Other:

* Transient 2145 UPS error (95%)

3024   Technician port connection invalid

Explanation: The code has detected more than one MAC address through the connection, or the DHCP has given out more than one address. The code thus believes there is a switch attached.

User response:

1. Remove the cable from the technician port.
2. (Optional) Disable additional network adapters on the laptop to which it is to connected.
3. Ensure DHCP is enabled on the network adapter.
4. If this was not possible, manually set the IP to 192.168.0.2
5. Connect a standard Ethernet cable between the network adapter and the technician port.
6. If this still does not work, reboot the node and repeat the above steps.
7. This event will auto-fix once either no connection or a valid connection has been detected.

3025   A virtualization feature license is required.

Explanation: The cluster has no virtualization feature license registered. You should have either an Entry Edition Physical Disk virtualization feature license or a Capacity virtualization feature license that covers the cluster.

The cluster will continue to operate, but it might be violating the license conditions.

User response:

* If you do not have a virtualization feature license that is valid and sufficient for this cluster, contact your IBM sales representative, arrange a license and change the license settings for the cluster to register the license.
* The error will automatically fix when the situation is resolved.

Possible Cause-FRUs or other:

* None

3029   Virtualization feature capacity is not valid.

Explanation: The setting for the amount of space that can be virtualized is not valid. The value must be an integer number of terabytes.

This error event is created when a cluster is upgraded from a version prior to 4.3.0 to version 4.3.0 or later. Prior to version 4.3.0 the virtualization feature capacity value was in gigabytes and therefore could be set to a fraction of a terabyte. With version 4.3.0 and later the

licensed capacity for the virtualization feature must be an integer number of terabytes.

User response:

* Review the license conditions for the virtualization feature. If you have one cluster, change the license settings for the cluster to match the capacity that is licensed. If your license covers more than one cluster, apportion an integer number of terabytes to each cluster. You might have to change the virtualization capacity that is set on the other clusters to ensure that the sum of the capacities for all of the clusters does not exceed the licensed capacity.
* You can view the event data or the feature log to ensure that the licensed capacity is sufficient for the space that is actually being used. Contact your IBM sales representative if you want to change the capacity of the license.
* This error will automatically be fixed when a valid configuration is entered.

Possible Cause-FRUs or other:

* None

3030   Global and Metro Mirror feature capacity not set.

Explanation: The Global and Metro Mirror feature is set to On for the system, but the capacity has not been set.

User response: Perform one of the following actions:

* Change the Global and Metro Mirror license settings for the system either to the licensed Global and Metro Mirror capacity, or if the license applies to more than one system, to the portion of the license allocated to this system. Set the licensed Global and Metro Mirror capacity to zero if it is no longer being used.
* View the event data or the feature log to ensure that the licensed Global and Metro Mirror capacity is sufficient for the space actually being used. Contact your IBM sales representative if you want to change the licensed Global and Metro Mirror capacity.
* The error will automatically be fixed when a valid configuration is entered.

Possible Cause-FRUs or other:

* None

3031   FlashCopy feature capacity not set.

Explanation: The FlashCopy feature is set to On for the system, but the capacity has not been set.

User response: Perform one of the following actions:

* Change the FlashCopy license settings for the system either to the licensed FlashCopy capacity, or if the license applies to more than one system, to the

portion of the license allocated to this system. Set the licensed FlashCopy capacity to zero if it is no longer being used.

- View the event data or the feature log to ensure that the licensed FlashCopy capacity is sufficient for the space actually being used. Contact your IBM sales representative if you want to change the licensed FlashCopy capacity.
- The error will automatically be fixed when a valid configuration is entered.

Possible Cause-FRUs or other:

- None

3032   Feature license limit exceeded.

Explanation: The amount of space that is licensed for a cluster feature is being exceeded.

The feature that is being exceeded might be:
- Virtualization (event identifier 009172)
- FlashCopy (event identifier 009173)
- Global and Metro Mirror (event identifier 009174)
- Transparent cloud tiering (event identifier 087046)

The cluster will continue to operate, but it might be violating the license conditions.

User response:

- Determine which feature license limit has been exceeded. This might be:
  – Virtualization (event identifier 009172)
  – FlashCopy (event identifier 009173)
  – Global and Metro Mirror (event identifier 009174)
  – Transparent cloud tiering (event identifier 087046)
- Use the `lslicense` command to view the current license settings.
- Ensure that the feature capacity that is reported by the cluster has been set to match either the licensed size, or if the license applies to more than one cluster, to the portion of the license that is allocated to this cluster.
- Decide whether to increase the feature capacity or to reduce the space that is being used by this feature.
- To increase the feature capacity, contact your IBM sales representative and arrange an increased license capacity. Change the license settings for the cluster to set the new licensed capacity. Alternatively, if the license applies to more than one cluster modify how the licensed capacity is apportioned between the clusters. Update every cluster so that the sum of the license capacity for all of the clusters does not exceed the licensed capacity for the location.
- To reduce the amount of disk space that is virtualized, delete some of the managed disks or

image mode volumes. The used virtualization size is the sum of the capacities of all of the managed disks and image mode disks.

- To reduce the FlashCopy capacity delete some FlashCopy mappings. The used FlashCopy size is the sum of all of the volumes that are the source volume of a FlashCopy mapping.
- To reduce Global and Metro Mirror capacity delete some Global Mirror or Metro Mirror relationships. The used Global and Metro Mirror size is the sum of the capacities of all of the volumes that are in a Metro Mirror or Global Mirror relationship; both master and auxiliary volumes are counted.
- To reduce the number of I/O groups that use Transparent cloud tiering, disable cloud snapshots for all cloud snapshot-enabled volumes from individual I/O groups until the total number of I/O groups using transparent cloud tiering is below the license limit.
- The error will automatically be fixed when the licensed capacity is greater than the capacity that is being used.

Possible Cause-FRUs or other:

- None

3035   Physical Disk FlashCopy feature license required

Explanation: The Entry Edition cluster has some FlashCopy mappings defined. There is, however, no Physical Disk FlashCopy license registered on the cluster. The cluster will continue to operate, but it might be violating the license conditions.

User response:

- Check whether you have an Entry Edition Physical Disk FlashCopy license for this cluster that you have not registered on the cluster. Update the cluster license configuration if you have a license.
- Decide whether you want to continue to use the FlashCopy feature or not.
- If you want to use the FlashCopy feature contact your IBM sales representative, arrange a license and change the license settings for the cluster to register the license.
- If you do not want to use the FlashCopy feature, you must delete all of the FlashCopy mappings.
- The error will automatically fix when the situation is resolved.

Possible Cause-FRUs or other:

- None

3036   Physical Disk Global and Metro Mirror feature license required

Explanation: The Entry Edition cluster has some Global Mirror or Metro Mirror relationships defined. There is, however, no Physical Disk Global and Metro Mirror license registered on the cluster. The cluster will continue to operate, but it might be violating the license conditions.

User response:
- Check if you have an Entry Edition Physical Disk Global and Metro Mirror license for this cluster that you have not registered on the cluster. Update the cluster license configuration if you have a license.
- Decide whether you want to continue to use the Global Mirror or Metro Mirror features or not.
- If you want to use either the Global Mirror or Metro Mirror feature contact your IBM sales representative, arrange a license and change the license settings for the cluster to register the license.
- If you do not want to use both the Global Mirror and Metro Mirror features, you must delete all of the Global Mirror and Metro Mirror relationships.
- The error will automatically fix when the situation is resolved.

Possible Cause-FRUs or other:
- None

3060   Array write endurance limited

Explanation: A RAID MDisk is affected by member flash drives that have a limited remaining write endurance.

User response: Complete the following steps:
1. Check the event log for the ID of the MDisk with limited remaining write endurance.
2. Run the `lsmdisk` and `lsdrive` commands to display information about the array and the individual drives. Note the date in the Replacement Date field for each drive in the `lsdrive` results.
3. If the replacement date or dates are approaching, consider replacing individual drives, or replacing the entire array.
4. Mark the event as fixed.

3080   Global or Metro Mirror relationship or consistency group with deleted partnership

Explanation: A Global Mirror or Metro Mirror relationship or consistency group exists with a cluster whose partnership is deleted.

This configuration is not supported and the problem should be resolved.

User response: The issue can be resolved either by

deleting all of the Global Mirror or Metro Mirror relationships or consistency groups that exist with a cluster whose partnership is deleted, or by recreating all of the partnerships that they were using.

The error will automatically fix when the situation is resolved.
1. List all of the Global Mirror and Metro Mirror relationships and note those where the master cluster name or the auxiliary cluster name is blank. For each of these relationships, also note the cluster ID of the remote cluster.
2. List all of the Global Mirror and Metro Mirror consistency groups and note those where the master cluster name or the auxiliary cluster name is blank. For each of these consistency groups, also note the cluster ID of the remote cluster.
3. Determine how many unique remote cluster IDs there are among all of the Global Mirror and Metro Mirror relationships and consistency groups that you have identified in the first two steps. For each of these remote clusters, decide if you want to re-establish the partnership with that cluster. Ensure that the total number of partnerships that you want to have with remote clusters does not exceed the cluster limit. If you re-establish a partnership, you will not have to delete the Global Mirror and Metro Mirror relationships and consistency groups that use the partnership.
4. Re-establish any selected partnerships.
5. Delete all of the Global Mirror and Metro Mirror relationships and consistency groups that you listed in either of the first two steps whose remote cluster partnership has not been re-established.
6. Check that the error has been marked as fixed by the system. If it has not, return to the first step and determine which Global Mirror or Metro Mirror relationships or consistency groups are still causing the issue.

Possible Cause-FRUs or other:
- None

3081   Unable to send email to any of the configured email servers.

Explanation: Either the system was not able to connect to any of the SMTP email servers, or the email transmission has failed. A maximum of six email servers can be configured. Error event 2600 or 2601 is raised when an individual email server is found to be not working. This error indicates that all of the email servers were found to be not working.

User response:
- Check the event log for all unresolved 2600 and 2601 errors and fix those problems.
- If this error has not already been automatically marked fixed, mark this error as fixed.

- Perform the check email function to test that an email server is operating properly.

Possible Cause-FRUs or other:

- None

3090   **Drive firmware download is cancelled by user or system, problem diagnosis required.**

Explanation: The drive firmware download has been cancelled by the user or the system and problem diagnosis required.

User response: If you cancelled the download using `applydrivesoftware -cancel` then this error is to be expected.

If you changed the state of any drive while the download was ongoing, this error is to be expected, however you will have to rerun the `applydrivesoftware` to ensure all your drive firmware has been updated.

Otherwise:

1. Check the drive states using `lsdrive`, in particular look at drives which are status=degraded, offline or use=failed.
2. Check node states using `lsnode` or `lsnodecanister`, and confirm all nodes are online.
3. Use `lsdependentvdisks -drive <drive_id>` to check for vdisks that are dependent on specific drives.
4. If the drive is a member of a RAID0 array, consider whether to introduce additional redundancy to protect the data on that drive.
5. If the drive is not a member of a RAID0 array, fix any errors in the event log that relate to the array.
6. Consider using the `-force` option. With any drive software upgrade there is a risk that the drive might become unusable. Only use the `-force` option if you accept this risk.
7. Reissue the `applydrivesoftware` again.

Note: The `lsdriveupgradeprogress` command can be used to check the progress of the `applydrivesoftware` command as it updates each drive.

3100   **Cloud account not available, unexpected error**

Explanation: The meaning of the error code depends on the associated event code.

087009 Cloud account not available, cannot establish secure connection with cloud provider

The network connection between the system and the cloud service provider is configured to use SSL. The SSL connection cannot be established. Cloud backup services remain paused until the alert is fixed.

The issue is *not* that the system cannot locate the CA certificate for the cloud service provider, or that the CA certificate is expired.

087012 Cloud account not available, cannot complete cloud storage operation

An unexpected error occurred when the system attempted to complete a cloud storage operation.

User response: Try the following actions for either event code:

1. Mark the error as fixed so that the system retries the operation.
2. If the errors repeat, check the cloud provider console or contact the cloud service provider. Look for errors and for changes since the last successful connection. The SSL connection worked at the time that the cloud account object was created.
3. Contact your service support representative. If possible, provide your representative with debug data from livedump and snap.

3108   **Unexpected error occurred while doing cloud operation**

Explanation: The associated event codes provide more information about a specific error:

087022 A cloud object could not be found during cloud snapshot operation.
> The system encountered a problem when it tried to read a particular object from cloud storage. The object is missing in the cloud.

087023 A cloud object was found to be corrupt during cloud snapshot operation.
> The system encountered a problem when it tried to read a particular object from cloud storage. The object format is wrong or the object longitudinal redundancy check (LRC) failed.

087024 A cloud object was found to be corrupt during cloud snapshot decompression operation.
> The system encountered a checksum failure while it was decompressing a particular object from cloud storage.

087025 Etag integrity error during cloud snapshot operation
> While the system was creating a snapshot in cloud storage, it encountered an HTML entity tag integrity error.

087027 Unexpected error occurred, cannot complete cloud snapshot operation
> An unanticipated error occurred during a snapshot operation.

087029 A cloud object could not be found during a cloud snapshot restore operation
> The system encountered a problem when it

tried to read a particular object from cloud storage during a restore operation. The object is missing in the cloud.

087030 A cloud object was found to be corrupt during a cloud snapshot restore operation
The system encountered a problem when it tried to read a particular object from cloud storage during a restore operation. The object format is wrong or the object longitudinal redundancy check (LRC) failed.

087031 A cloud object was found to be corrupt during a cloud snapshot restore decompression operation
The system encountered a checksum failure while it was decompressing a particular object from cloud storage during a restore operation.

087032 Etag integrity error during cloud snapshot restore operation
During a restore operation, the system encountered an HTML entity tag integrity error.

087034 Cannot create bad blocks on a managed disk during cloud snapshot restore operation.
The system cannot work around medium errors on the cloud volume during a restore operation.

087035 Unexpected error occurred, cannot complete cloud snapshot restore operation
An unanticipated error occurred during a restore operation.

087037 A cloud object could not be found during a cloud snapshot delete operation
The system encountered a problem when it tried to read a particular object from cloud storage during a delete operation. The object is missing in the cloud.

087038 A cloud object was found to be corrupt during cloud snapshot delete operation
The system encountered a problem when it tried to read a particular object from cloud storage during a delete operation. The object format is wrong or the object longitudinal redundancy check (LRC) failed.

087039 A cloud object was found to be corrupt during cloud snapshot delete decompression operation
The system encountered a checksum failure while it was decompressing a particular object from cloud storage during a delete operation.

087040 Unexpected error occurred, cannot complete cloud snapshot delete operation
An unanticipated error occurred during a delete operation.

In all cases, the job remains paused until the alert is fixed.

User response: Contact your support service representative.

3123   **The quorum application needs to be redeployed.**

Explanation: A setting specific to the quorum application changed, which means that the quorum application might not be able to function as the active quorum device. Any of the following problems might be involved:

- A service IP was changed.
- A change in the IP network prevented the quorum application from reaching all the nodes.
- One or more nodes were permanently added to or removed from the cluster.
- The certificate was changed.

User response: Complete the following steps:

1. Make sure that all Ethernet cables are connected correctly.
2. Make sure that the service IP addresses are set for all nodes.
3. Make sure that you can ping all nodes from the quorum application host.
4. Regenerate the JAR file that contains the new configuration by using the management GUI or the command line.
5. Transfer the new application to the deployment locations or the host or hosts.
6. Stop the old application.
7. Start the new application.
8. Verify that the cluster is using the quorum application as the active quorum device by using the `lsquorum` command.

3124   **No active quorum device found.**

Explanation: A quorum device must be active to avoid an I/O outage if the node fails.

User response: Use the `lsquorum` command to verify that a quorum device is active. The **active** field should have a value of **yes**. If no quorum devices are active, complete one of the following actions:

- On HyperSwap or stretched systems, deploy a new IP quorum application or create a third Fibre Channel quorum site.
- On regular systems, create some managed storage or deploy a new IP quorum application.

3130   **System SSL certificate expires within the next 30 days.**

Explanation: System SSL certificate expires within the next 30 days.

The system SSL certificate that is used to authenticate connections to the GUI, service assistant, and the CIMOM is about to expire.

User response: Complete the following steps to resolve this problem.

1. If you are using a self-signed certificate, then generate a new self-signed certificate.

2. If you are using a certificate that is signed by a certificate authority, generate a new certificate request and get this certificate signed by your certificate authority. The existing certificate can continue to be used until the expiry date to provide time to get the new certificate request signed and installed.

Possible Cause-FRUs or other:

- N/A

3135   Cloud account not available, incompatible object data format

Explanation: The cloud account is in import mode, accessing data from another system. The code on that system was updated to a level higher than the level on your current system. The other system made updates to the cloud storage that your current system cannot interpret.

User response: Try the following actions:

1. Contact the administrator of the other system to determine its code level and the changes that are planned. Use `lscloudaccount` to get the ID and name for the other system.

2. Update your current system to a compatible level of code.

3. Alternatively, change the cloud account back to normal mode.

3140   Cloud account SSL certificate will expire within the next 30 days

Explanation: A cloud account SSL certificate was presented that is due to expire.

User response: Try the following actions:

1. Verify certificate validity start and end times from the alert event sense data.

2. Verify that the system time is correct.

3. Contact the cloud service provider for a new certificate.

Note: The alert does not auto-fix until the certificate becomes valid or the account is switched out of SSL mode.

3220   Equivalent ports may be on different fabrics

Explanation: Mismatched fabric World Wide Names (WWNs) were detected.

User response: Complete the following steps:

1. Run the `lsportfc` command to get the fabric World Wide Name (WWN) of each port.

2. List all partnered ports (that is, all ports for which the platform port ID is the same, and the node is in the same I/O group) that have mismatched fabric WWNs.

3. Verify that the listed ports are on the same fabric.

4. Rewire if needed. For information about wiring requirements, see "Zoning considerations for N_Port ID Virtualization" in your product documentation. After all ports are on the same fabric, the event corrects itself.

5. This error might be displayed by mistake. If you confirm that all remaining ports are on the same fabric, despite apparent mismatches that remain, mark the event as fixed.

Appendix. Accessibility features for Storwize V7000

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

These are the major accessibility features for the Storwize V7000:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. HTML documents have been tested using JAWS version 15.0.
- This product uses standard Windows navigation keys.
- Interfaces are commonly used by screen readers.
- Industry-standard devices, ports, and connectors.

The Storwize V7000 online documentation and its related publications are accessibility-enabled. The accessibility features of the online documentation are described in Viewing information in the information center .

Keyboard navigation

You can use keys or key combinations to perform operations and initiate menu actions that can also be done through mouse actions. You can navigate the Storwize V7000 online documentation from the keyboard by using the shortcut keys for your browser or screen-reader software. See your browser or screen-reader software Help for a list of shortcut keys that it supports.

IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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

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

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those

websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

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

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application

programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other product and service names might be trademarks of IBM or other companies.

Electronic emission notices

This section contains the electronic emission notices or statements for the United States and other countries.

Federal Communications Commission (FCC) statement

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, might 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 might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

Industry Canada compliance statement

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

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

Australia and New Zealand Class A Statement

Attention: This is a Class A product. In a domestic environment this product might cause radio interference in which case the user might be required to take adequate measures.

European Union Electromagnetic Compatibility Directive

This product is in conformity with the protection requirements of European Union (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.

Attention: This is an EN 55022 Class A product. In a domestic environment this product might cause radio interference in which case the user might be required to take adequate measures.

European community contact:

IBM Deutschland GmbH
Technical Regulations, Department M372
IBM-Allee 1, 71139 Ehningen, Germany
Tele: +49 (0) 800 225 5423 or +49 (0) 180 331 3233
Email: halloibm@de.ibm.com

Germany Electromagnetic Compatibility Directive

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2004/108/EG zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55022 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung der IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung der IBM gesteckt/eingebaut werden.

EN 55022 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:

"Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Mabnahmen 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
Tele: +49 (0) 800 225 5423 or +49 (0) 180 331 3233
Email: halloibm@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 Klasse A.

People's Republic of China Class A Statement

中华人民共和国"A类"警告声明

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。

Taiwan Class A compliance statement

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。

Taiwan Contact Information

This topic contains the product service contact information for Taiwan.

```
IBM Taiwan Product Service Contact Information:
IBM Taiwan Corporation
3F, No 7, Song Ren Rd., Taipei Taiwan
Tel: 0800-016-888
```

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

f2c00790

Japan Voluntary Control Council for Interference Class A statement

This explains the Japan Voluntary Control Council for Interference (VCCI) statement.

この装置は、クラス A 情報技術装置です。この装置を家庭環境で使用する
と電波妨害を引き起こすことがあります。この場合には使用者が適切な対策
を講ずるよう要求されることがあります。　　　　　　　VCCI-A

Japan Electronics and Information Technology Industries Association Statement

This statement explains the Japan JIS C 61000-3-2 product wattage compliance.

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値：Knowledge Center を参照

This statement explains the Japan Electronics and Information Technology Industries Association (JEITA) statement for products less than or equal to 20 A per phase.

高周波電流規格　JIS C 61000-3-2 適合品

This statement explains the JEITA statement for products greater than 20 A, single phase.

高周波電流規格　JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」
対象機器（高調波発生機器）です。
• 回路分類　　　　：6（単相、ＰＦＣ回路付）
• 換算係数　　　　：0

This statement explains the JEITA statement for products greater than 20 A per phase, three-phase.

高周波電流規格　JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」
対象機器（高調波発生機器）です。
・回路分類　　　：5（3相、ＰＦＣ回路付）
・換算係数　　　：0

Korean Electromagnetic Interference (EMI) Statement

This explains the Korean Electromagnetic Interference (EMI) statement.

이 기기는 업무용 환경에서 사용할 목적으로 적합성평가를 받은 기기로서
가정용 환경에서 사용하는 경우 전파간섭의 우려가 있습니다.

Russia Electromagnetic Interference Class A Statement

This statement explains the Russia Electromagnetic Interference (EMI) statement.

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать
радиопомехи, для снижения которых необходимы
дополнительные меры

rusemi

IBM ®

Printed in USA