IBM System Storage SAN Volume Controller

Troubleshooting Guide

Note

Before using this information and the product it supports, read the following information:

- The general information in "Notices" on page 317
- The information in the Safety and environmental notices
- The information in the *IBM Environmental Notices and User Guide* (provided on a DVD)

This edition applies to version 8, release 1, modification 3, and to all subsequent modifications until otherwise indicated in new editions.

Contents

Figures

Tables

About this guide

This guide describes how to troubleshoot the IBM® SAN Volume Controller .

The chapters that follow introduce you to the SAN Volume Controller , expansion enclosure, the redundant AC-power switch, and the uninterruptible power supply. They describe how you can configure and check the status of one SAN Volume Controller node or a clustered system of nodes through the front panel, with the service assistant GUI, or with the management GUI.

The vital product data (VPD) chapter provides information about the VPD that uniquely defines each hardware and microcode element that is in the SAN Volume Controller . You can also learn how to diagnose problems using the SAN Volume Controller .

The maintenance analysis procedures (MAPs) can help you analyze failures that occur in a SAN Volume Controller . With the MAPs, you can isolate the field-replaceable units (FRUs) of the SAN Volume Controller that fail. Begin all problem determination and repair procedures from "MAP 5000: Start" on page 273.

Who should use this guide

This guide is intended for the system administrator or systems services representative who uses and diagnoses problems with the SAN Volume Controller , the redundant AC-power switch, and the uninterruptible power supply.

Emphasis

Different typefaces are used in this guide to show emphasis.

The following typefaces are used to show emphasis.

Emphasis	Meaning
Boldface	Text in **boldface** represents menu items.
`Bold monospace`	Text in **`bold monospace`** represents command names.
Italics	Text in *italics* is used to emphasize a word. In command syntax, it is used for variables for which you supply actual values, such as a default directory or the name of a system.
`Monospace`	Text in `monospace` identifies the data or commands that you type, samples of command output, examples of program code or messages from the system, or names of command flags, parameters, arguments, and name-value pairs.

Library and related publications

Product manuals, other publications, and websites that contain information that is related to your system are available.

IBM Knowledge Center for SAN Volume Controller

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

SAN Volume Controller library

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 SAN Volume Controller (2145)	www.ibm.com/support
Support for IBM System Storage® and IBM TotalStorage products	www.ibm.com/support

Each PDF publication in theTable 2 library is available in the IBM Knowledge Center by clicking the title in the "Link to PDF" column:

Table 2. SAN Volume Controller library

Title	Description	Link to PDF file
IBM SAN Volume Controller Model 2145-SV1 Hardware Installation Guide	The guide provides the instructions that the IBM service representative uses to install the hardware for SAN Volume Controller model 2145-SV1.	Hardware Installation Guide [PDF]
IBM SAN Volume Controller Hardware Maintenance Guide	The guide provides the instructions that the IBM service representative uses to service the SAN Volume Controller hardware, including the removal and replacement of parts.	Hardware Maintenance Guide [PDF]
IBM SAN Volume Controller Troubleshooting Guide	The guide describes the features of each SAN Volume Controller model, explains how to use the front panel or service assistant GUI, and provides maintenance analysis procedures to help you diagnose and solve problems with the SAN Volume Controller .	Troubleshooting Guide [PDF]
IBM Spectrum Virtualize for Public Cloud, 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 SAN Volume Controller command-line interface (CLI).	Command-Line Interface User's Guide [PDF]

Table 2. SAN Volume Controller library (continued)

Title	Description	Link to PDF file
IBM Spectrum Virtualize REST API	This document provides information on the REST API and related CLI commands.	

IBM documentation and related websites

Table 3 lists websites that provide publications and other information about the SAN Volume Controller 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

Related websites

The following websites provide information about the system, related products, or technologies.

Type of information	Website
SAN Volume Controller 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 ibmkc@us.ibm.com. Be sure to include the following information:

- Exact publication title and version
- 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 want more information about IBM products, you can 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.

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

Website	Address
Directory of worldwide contacts	http://www.ibm.com/planetwide
Support for SAN Volume Controller (2145)	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 you call 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 you call 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 you call for assistance. Choose the **software** option if you are uncertain if the problem involves the SAN Volume Controller software or hardware. Choose the **hardware** option only if you are certain the problem solely involves the SAN Volume Controller hardware. When you call IBM to service the product, follow these guidelines for the **software** and **hardware** options:

Software option
Identify the SAN Volume Controller 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 might 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 SAN Volume Controller , the machine type is 2145.

In the US and Canada, hardware service and support can be extended to 24 x 7 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 take steps to try to solve the problem yourself before you call.

Some suggestions for resolving the problem before you call 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 this information, 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 for your system, and in the future.

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

Chapter 1. SAN Volume Controller overview

The SAN Volume Controller system combines software and hardware into a comprehensive, modular appliance that provides symmetric virtualization.

Symmetric virtualization is achieved by creating a pool of managed disks (MDisks) from the attached storage systems and optional SAS expansion enclosures.Volumes can be created in a pool for use by attached host systems. System administrators can view and access a common pool of storage on the storage area network (SAN). or local area network (LAN). This functionality helps administrators to use storage resources more efficiently and provides a common base for advanced functions.

A *SAN* is a high-speed Fibre Channel network that connects host systems and storage devices. A *LAN* is a high-speed Ethernet network that connects host systems and storage devices. In a SAN and LAN, a host system can be connected to a storage device across the network. The connections are made through units such as routers and switches. The area of the network that contains these units is known as the *fabric* of the network.

IBM Real-time Compression™ software

IBM SAN Volume Controller system is built with IBM Spectrum Virtualize™ software, which is part of the IBM Spectrum Storage™ family.

IBM Spectrum Virtualize is a key member of the IBM Spectrum Storage portfolio. It is a highly flexible storage solution that enables rapid deployment of block storage services for new and traditional workloads, on-premises, off-premises and in a combination of both. Designed to help enable cloud environments, it is based on the proven technology. For more information about the IBM Spectrum Storage portfolio, see the following website.

`http://www.ibm.com/systems/storage/spectrum`

The software provides these functions for the host systems that attach to the system:

- Creates a single pool of storage
- Provides logical unit virtualization
- Manages logical volumes
- Mirrors logical volumes

The system also provides the following functions:

- Large scalable cache
- Copy Services:
  - IBM FlashCopy® (point-in-time copy) function, including thin-provisioned FlashCopy to make multiple targets affordable
  - IBM HyperSwap® (active-active copy) function
  - Metro Mirror (synchronous copy)
  - Global Mirror (asynchronous copy)
  - Data migration
- Space management:

- IBM Easy Tier® function to migrate the most frequently used data to higher-performance storage
- Metering of service quality when combined with IBM Spectrum Control Base Edition. For information, refer to the IBM Spectrum Control Base Edition documentation.
- Thin-provisioned logical volumes
- Compressed volumes to consolidate storage

Figure 1 shows hosts, system nodes, and RAID storage systems connected to a SAN fabric. The redundant SAN fabric comprises a fault-tolerant arrangement of two or more counterpart SANs that provide alternative paths for each SAN-attached device.

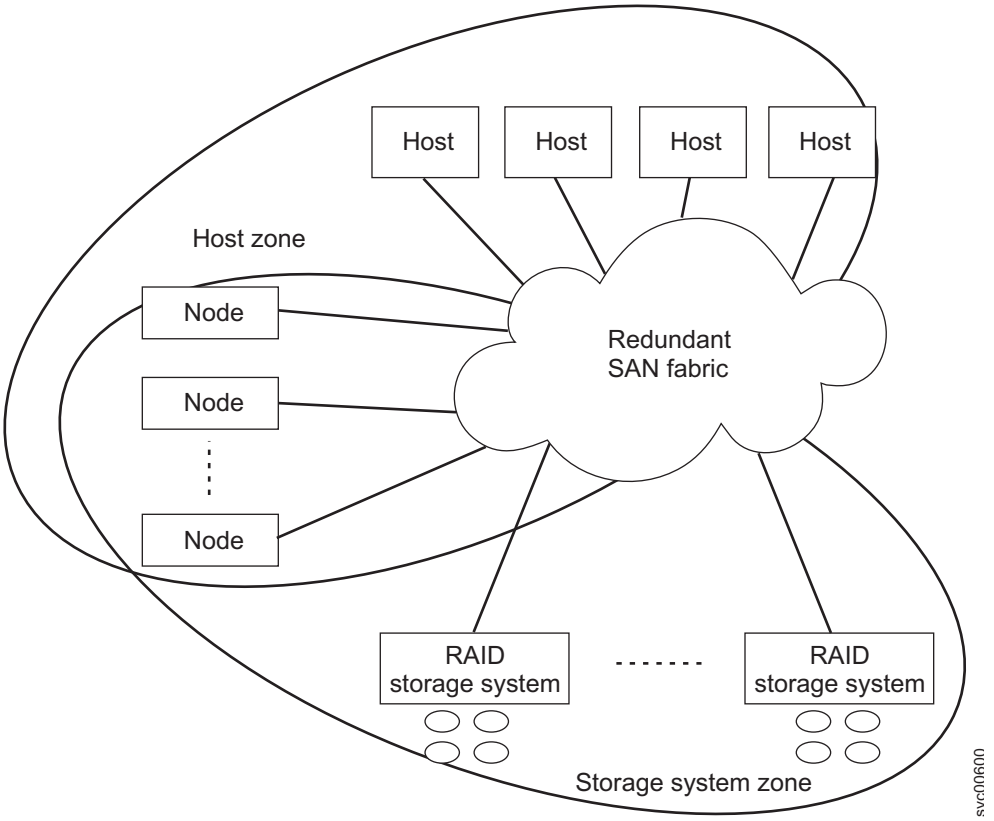

Figure 1. Example of a system in a fabric

Volumes

System nodes present volumes to the hosts. Most of the advanced system functions are defined on volumes. These volumes are created from managed disks (MDisks) that are presented by the RAID storage systems. The volumes can also be created by arrays that are provided by flash drives in an expansion enclosure. All data transfer occurs through the system node, which is described as *symmetric virtualization*.

Figure 2 on page 3 shows the data flow across the fabric.

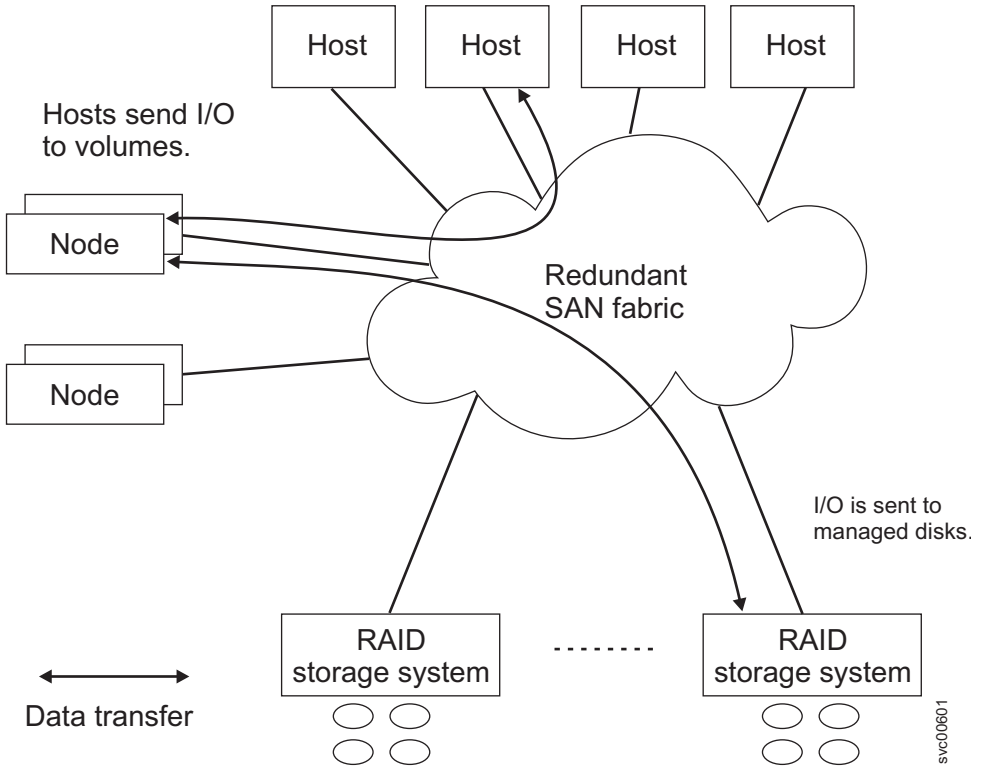

Figure 2. Data flow in a system

The nodes in a system are arranged into pairs that are known as *I/O groups*. A single pair is responsible for serving I/O on a volume. Because a volume is served by two nodes, no loss of availability occurs if one node fails or is taken offline. The Asymmetric Logical Unit Access (ALUA) features of SCSI are used to disable the I/O for a node before it is taken offline or when a volume cannot be accessed via that node.

Volume types

You can create the following types of volumes on the system:

- *Basic volumes*, where a single copy of the volume is cached in one I/O group. Basic volumes can be established in any system topology; however, Figure 3 on page 4 shows a standard system topology.

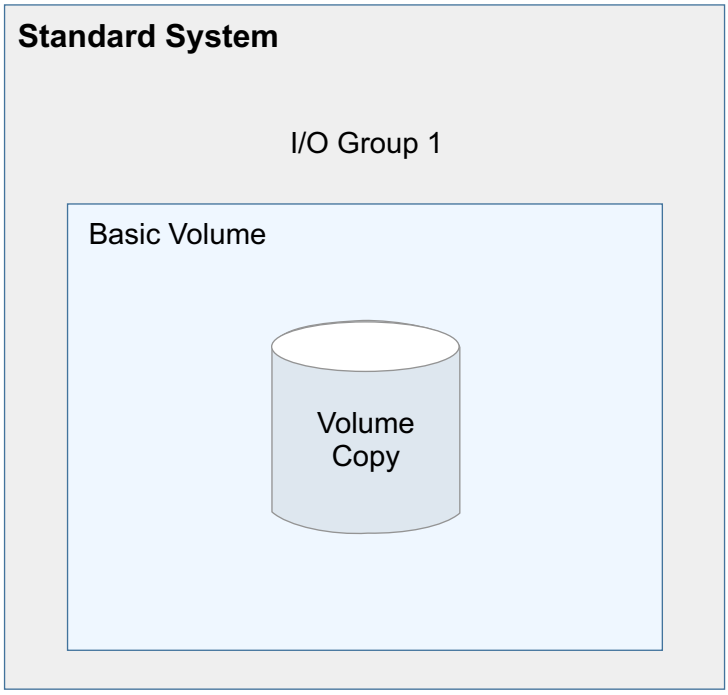

Figure 3. Example of a basic volume

- *Mirrored volumes*, where copies of the volume can either be in the same storage pool or in different storage pools. As Figure 4 shows, the volume is cached in a single I/O group. Typically, mirrored volumes are established in a standard system topology.

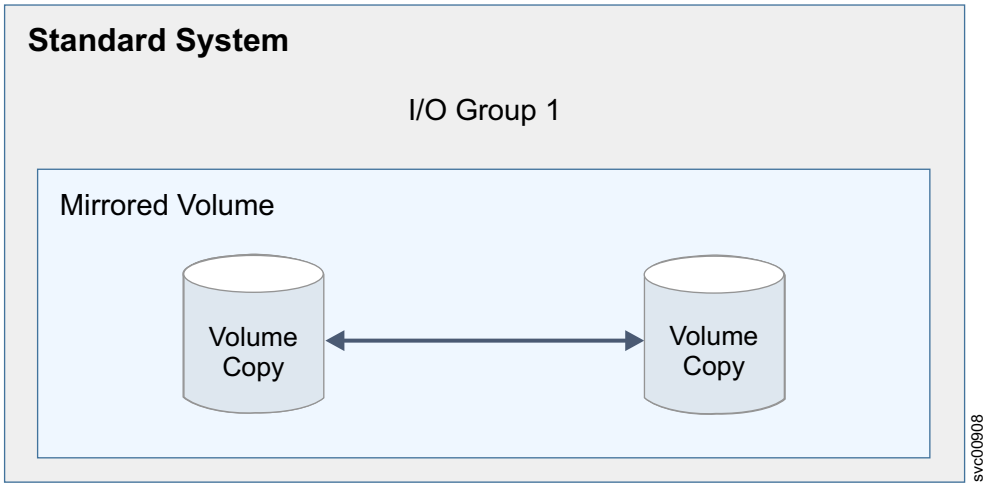

Figure 4. Example of mirrored volumes

- *Stretched volumes*, where copies of a single volume are in different storage pools at different sites. As Figure 5 on page 5 shows, the volume is cached in one I/O group. Stretched volumes are only available in stretched topology systems.

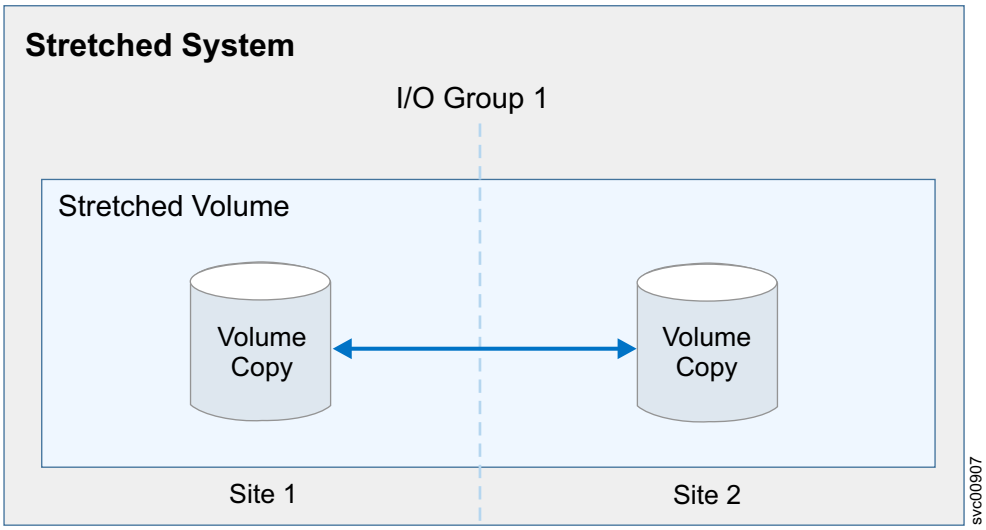

Figure 5. Example of stretched volumes

- *HyperSwap volumes*, where copies of a single volume are in different storage pools that are on different sites. The volume is cached in two I/O groups that are on different sites, as Figure 6 on page 6 shows. These volumes can be created only when the system topology is HyperSwap.

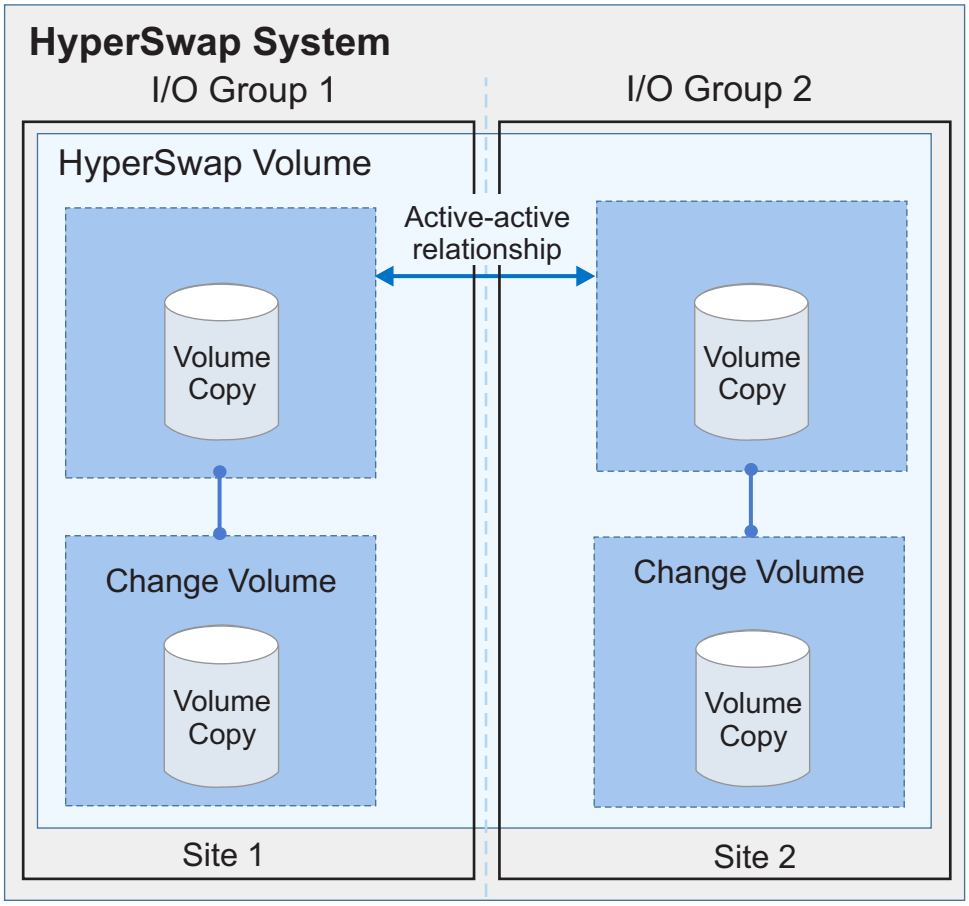

Figure 6. Example of HyperSwap volumes

System topology

The topology property of a system can be set to one of the following states:

Note: You cannot mix I/O groups of different topologies in the same system.

- *Standard* topology, where all nodes in the system are at the same site.

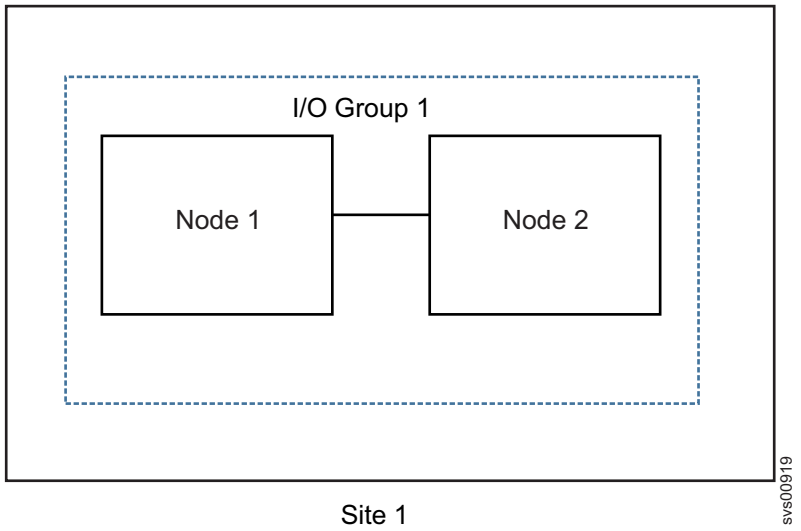

Figure 7. Example of a standard system topology

- *Stretched* topology, where each node of an I/O group is at a different site. When one site is not available, access to a volume can continue but with reduced performance.

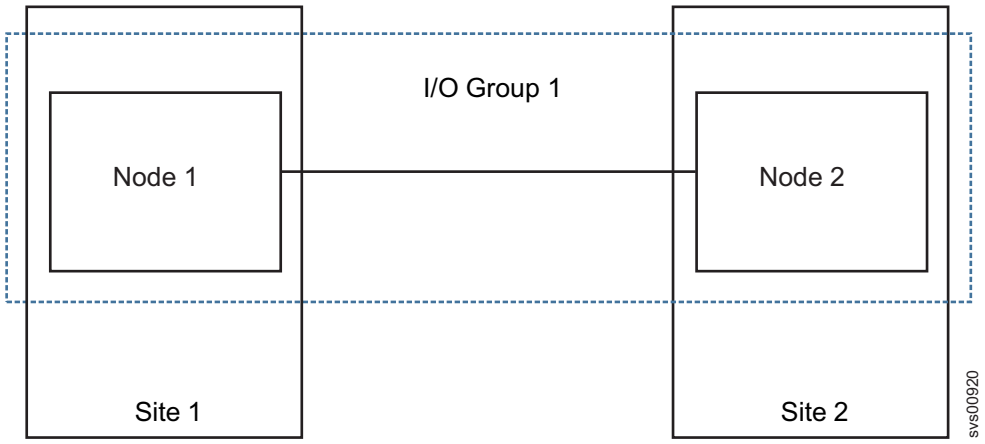

Figure 8. Example of a stretched system topology

- *HyperSwap* topology, where the system consists of at least two I/O groups. Each I/O group is at a different site. Both nodes of an I/O group are at the same site. A volume can be active on two I/O groups so that it can immediately be accessed by the other site when a site is not available.

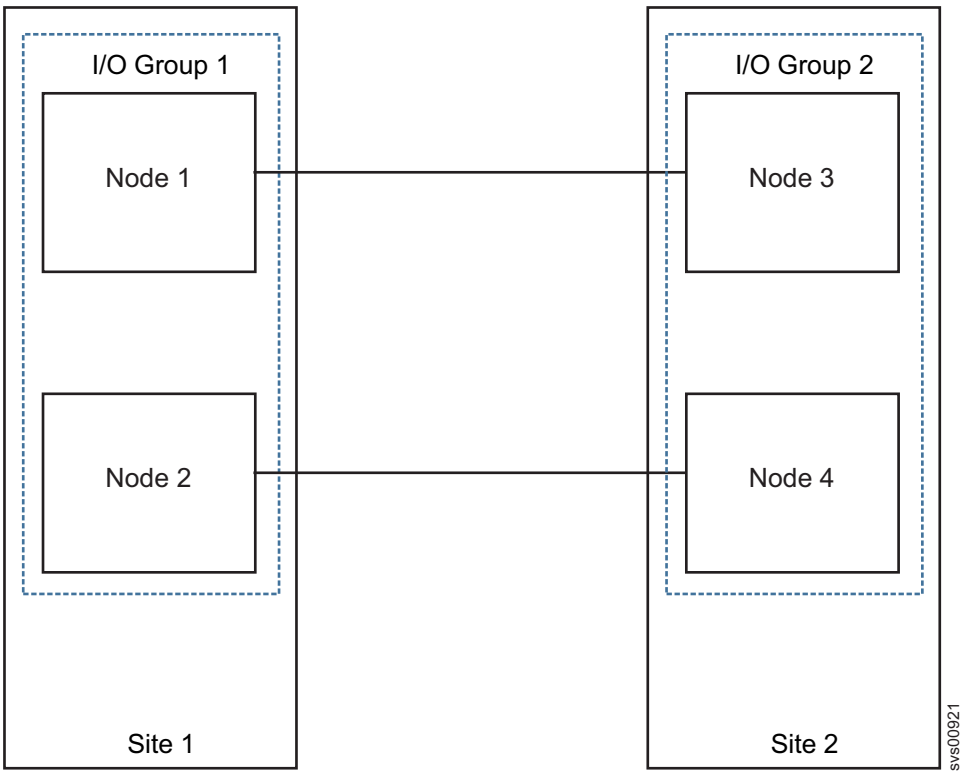

Figure 9. Example of a HyperSwap system topology

Summary of system topology and volumes

Table 5 summarizes the types of volumes that can be associated with each system topology.

Table 5. System topology and volume summary

Topology	Volume Type				
	Basic	**Mirrored**	**Stretched**	**HyperSwap**	**Custom**
Standard	X	X			X
Stretched	X		X		X
HyperSwap	X			X	X

System management

A system is composed of individual nodes that present a single point of control for system management and service. System management and error reporting are provided through an Ethernet interface to one of the nodes in the system, which is called the *configuration node*. The configuration node runs a web server and provides a command-line interface (CLI). Any node in the system can be the configuration node. If the current configuration node fails, a new configuration node is selected from the remaining nodes. Each node also provides a command-line interface and web interface for initiating hardware service actions.

Fabric types

I/O operations between hosts and system nodes and between the nodes and arrays use the SCSI standard. The nodes communicate with each other through private SCSI commands.

All nodes that run system software version 6.4 or later can support Fibre Channel over Ethernet (FCoE) connectivity.

Table 6 shows the fabric type that can be used for communicating between hosts, nodes, and RAID storage systems. These fabric types can be used at the same time.

Table 6. System communications types

Communications type	Host to system nodes	System nodes to storage system	System nodes to system nodes
Fibre Channel SAN	Yes	Yes	Yes
iSCSI • 1 Gbps Ethernet (SAN Volume Controller 2145-DH8 only) • 10 Gbps Ethernet • 25 Gbps Ethernet (SAN Volume Controller 2145-SV1 only)	Yes	Yes	No
Fibre Channel Over Ethernet SAN (10 Gbps Ethernet)	Yes	Yes	Yes

Flash drives

Some system nodes are attached to expansion enclosures that contain flash drives. These flash drives can be used to create RAID-managed disks (MDisks) that in turn can be used to create volumes. Flash drives are in an expansion enclosure that is connected to both sides of an I/O group.

Flash drives provide host servers with a pool of high-performance storage for critical applications. MDisks on flash drives can also be placed in a storage pool with MDisks from regular RAID storage systems. IBM Easy Tier performs automatic data placement within that storage pool by moving high-activity data onto better-performing storage.

SAN Volume Controller nodes

Each node is an individual server in a SAN Volume Controller clustered system on which the SAN Volume Controller software runs.

The nodes are always installed in pairs; a minimum of one pair and a maximum of four pairs of nodes constitute a *system*. Each pair of nodes is known as an *I/O group*.

I/O groups take the storage that is presented to the SAN by the storage systems as MDisks and transforms the storage into logical disks (volumes) that are used by applications on the hosts. A node is in only one I/O group and provides access to the volumes in that I/O group.

SAN Volume Controller 2145-SV1 node features

The SAN Volume Controller 2145-SV1 system has the following features.

- A 19-inch rack-mounted enclosure
- Two 8-core processors
- 64 GB base memory per processor. Optionally, by adding 64 GB memory modules, the processor can support 128 GB, 192 GB, or 256 GB of memory.
- Eight small form factor (SFF) drive bays at the front of the control enclosure
- Support for various optional host adapters, including:
  - 4-port 16 Gbps Fibre Channel adapters
  - 4-port 10 Gbps Fibre Channel over Ethernet (FCoE) adapters for host attachment
  - 4-port 12 Gbps SAS cards to attach to expansion enclosures
- Support for iSCSI host attachment (10 Gbps Ethernet or 25 Gbps Ethernet)
- Support for expansion enclosures to support more drives
  - SAN Volume Controller 2145-92F expansion enclosure houses up to 92 flash drives (SFF or LFF drives) and two secondary expander modules
  - SAN Volume Controller 2145-24F houses up to 24 SFF flash drives
  - SAN Volume Controller 2145-12F houses up to 12 large form factor (LFF) HDD or flash drives
- Support for optional Compression Accelerator cards for IBM Real-time Compression
- Dual redundant power supplies
- Dual redundant batteries
- A dedicated technician port to initialize or service the system

SAN Volume Controller 2147-SV1 node features

The SAN Volume Controller 2147-SV1 system includes all of the features of the SAN Volume Controller 2145-SV1 system plus Enterprise Class Support and a three-year warranty.

SAN Volume Controller 2145-DH8 node features

The SAN Volume Controller 2145-DH8 node has the following features:

- A 19-inch rack-mounted enclosure
- At least one Fibre Channel adapter or one 10 Gbps Ethernet adapter
- Optional second, third, and fourth Fibre Channel adapters
- 32 GB memory per processor
- One or two, eight-core processors
- Dual redundant power supplies
- Dual redundant batteries for better reliability, availability, and serviceability
- SAN Volume Controller 2145-92F expansion enclosure to house up to 92 flash drives (SFF or LFF drives) and two secondary expander modules
- Up to two SAN Volume Controller 2145-24F expansion enclosures to house up to 24 flash drives each
- SAN Volume Controller 2145-12F expansion enclosures to house up to 12 LFF HDD or flash drives
- iSCSI host attachment (1 Gbps Ethernet and optional 10 Gbps Ethernet)

- Supports optional IBM Real-time Compression
- A dedicated technician port for local access to the initialization tool or the service assistant interface.

Systems

Systems are collections of nodes. Systems can consist of between two to eight nodes.

All configuration settings are replicated across all nodes in the system. Management IP addresses are assigned to the system. Each interface accesses the system remotely through the Ethernet system-management addresses, also known as the primary, and secondary system IP addresses.

Configuration node

A *configuration node* is a single node that manages configuration activity of the system.

If the configuration node fails, the system chooses a new configuration node. This action is called configuration node failover. The new configuration node takes over the management IP addresses. Thus, you can access the system through the same IP addresses although the original configuration node has failed. During the failover, there is a short period when you cannot use the command-line tools or management GUI.

Figure 10 shows an example of a clustered system that contains four nodes. Node 1 is the configuration node. User requests (**1**) are handled by node 1.

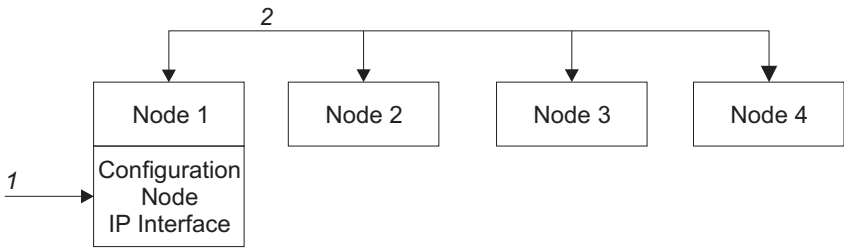

Figure 10. Configuration node

Configuration node addressing

At any given time, only one node within a SAN Volume Controller clustered system is assigned an IP addresses.

An IP address for the clustered system must be assigned to Ethernet port 1. An IP address can also be assigned to Ethernet port 2. These are the only ports that can be assigned management IP addresses.

This node then acts as the focal point for all configuration and other requests that are made from the management GUI application or the CLI. This node is known as the *configuration node*.

If the configuration node is stopped or fails, the remaining nodes in the system determine which node will take on the role of configuration node. The new

configuration node binds the management IP addresses to its Ethernet ports. It broadcasts this new mapping so that connections to the system configuration interface can be resumed.

The new configuration node broadcasts the new IP address mapping using the Address Resolution Protocol (ARP). You must configure some switches to forward the ARP packet on to other devices on the subnetwork. Ensure that all Ethernet devices are configured to pass on unsolicited ARP packets. Otherwise, if the ARP packet is not forwarded, a device loses its connection to the SAN Volume Controller system.

If a device loses its connection to the SAN Volume Controller system, it can regenerate the address quickly if the device is on the same subnetwork as the system. However, if the device is not on the same subnetwork, it might take hours for the address resolution cache of the gateway to refresh. In this case, you can restore the connection by establishing a command line connection to the system from a terminal that is on the same subnetwork, and then by starting a secure copy to the device that has lost its connection.

Management IP failover

If the configuration node fails, the IP addresses for the clustered system are transferred to a new node. The system services are used to manage the transfer of the management IP addresses from the failed configuration node to the new configuration node.

The following changes are performed by the system service:

- If software on the failed configuration node is still operational, the software shuts down the management IP interfaces. If the software cannot shut down the management IP interfaces, the hardware service forces the node to shut down.
- When the management IP interfaces shut down, all remaining nodes choose a new node to host the configuration interfaces.
- The new configuration initializes the configuration daemons, including SSHD and HTTPD, and then binds the management IP interfaces to its Ethernet ports.
- The router is configured as the default gateway for the new configuration.
- The routing tables are established on the new configuration for the management IP addresses. The new configuration sends five unsolicited address resolution protocol (ARP) packets for each IP address to the local subnet broadcast address. The ARP packets contain the management IP and the Media Access Control (MAC) address for the new configuration node. All systems that receive ARP packets are forced to update their ARP tables. After the ARP tables are updated, these systems can connect to the new configuration node.

  Note: Some Ethernet devices might not forward ARP packets. If the ARP packets are not forwarded, connectivity to the new configuration node cannot be established automatically. To avoid this problem, configure all Ethernet devices to pass unsolicited ARP packets. You can restore lost connectivity by logging in to the system and starting a secure copy to the affected system. Starting a secure copy forces an update to the ARP cache for all systems that are connected to the same switch as the affected system.

Ethernet link failures

If the Ethernet link to the system fails because of an event that is unrelated to the system, the system does not attempt to fail over the configuration node to restore

management IP access. For example, the Ethernet link can fail if a cable is disconnected or an Ethernet router fails. To protect against this type of failure, the system provides the option for two Ethernet ports that each have a management IP address. If you cannot connect through one IP address, attempt to access the system through the alternative IP address.

Note: IP addresses that are used by hosts to access the system over an Ethernet connection are different from management IP addresses.

Routing considerations for event notification and Network Time Protocol

The system supports the following protocols that make outbound connections:

- Email
- Simple Network Mail Protocol (SNMP)
- Syslog
- Network Time Protocol (NTP)

These protocols operate only on a port that is configured with a management IP address. When it is making outbound connections, the system uses the following routing decisions:

- If the destination IP address is in the same subnet as one of the management IP addresses, the system sends the packet immediately.
- If the destination IP address is not in the same subnet as either of the management IP addresses, the system sends the packet to the default gateway for Ethernet port 1.
- If the destination IP address is not in the same subnet as either of the management IP addresses and Ethernet port 1 is not connected to the Ethernet network, the system sends the packet to the default gateway for Ethernet port 2.

When you configure any of these protocols for event notifications, use these routing decisions to ensure that error notification works correctly if the network fails.

SAN fabric overview

The *SAN fabric* is an area of the network that contains routers and switches. A SAN is configured into a number of zones. A device that uses the SAN can communicate only with devices that are included in the same zones that it is in. A system requires several distinct types of zones: a system zone, host zones, and disk zones. The intersystem zone is optional.

In the host zone, the host systems can identify and address the nodes. You can have more than one host zone and more than one disk zone. Unless you are using a dual-core fabric design, the system zone contains all ports from all nodes in the system. Create one zone for each host Fibre Channel port. In a disk zone, the nodes identify the storage systems. Generally, create one zone for each external storage system. If you are using the Metro Mirror and Global Mirror feature, create a zone with at least one port from each node in each system; up to four systems are supported.

Note: Some operating systems cannot tolerate other operating systems in the same host zone, although you might have more than one host type in the SAN fabric.

For example, you can have a SAN that contains one host that runs on an IBM AIX® operating system and another host that runs on a Microsoft Windows operating system.

All communication between the system nodes is performed through the SAN. All of the system configuration and service commands are sent to the system through an Ethernet network.

Chapter 2. Introducing the SAN Volume Controller hardware components

A SAN Volume Controller system consists of SAN Volume Controller nodes and related hardware components, such as uninterruptible power supply units and the optional redundant AC-power switches. Note that nodes and uninterruptible power supply units are installed in pairs.

SAN Volume Controller nodes

The system supports several different types of models.

The following nodes are supported:

- The SAN Volume Controller 2145-DH8 node is available for purchase, with the following features:
  - At least one Fibre Channel adapter or one 10 Gbps Ethernet adapter
  - Optional second and third Fibre Channel adapters
  - Up to two SAN Volume Controller 2145-24F expansion enclosures to house optional flash drives
  - iSCSI host attachment (1 Gbps Ethernet and optional 10 Gbps Ethernet)
- The SAN Volume Controller 2145-SV1 node is available for purchase, with the following features:
  - At least one Fibre Channel adapter or one 10 Gbps Ethernet adapter
  - Optional second and third Fibre Channel adapters
  - Up to two SAN Volume Controller 2145-24F expansion enclosures to house optional flash drives
  - iSCSI host attachment (1 Gbps Ethernet and optional 10 Gbps Ethernet)

A label on the front of the node indicates the node type, hardware revision (if appropriate), and serial number.

Optional features

The SAN Volume Controller 2145-SV1 and SAN Volume Controller 2145-DH8 nodes support optional features, which can be installed concurrently.

Features or models

Table 7 lists the optional features and models that can be installed on your SAN Volume Controller 2145-SV1 or SAN Volume Controller 2145-DH8 system. Only an IBM service support representative (SSR) can remove or install adapters on the system.

Table 7. Optional features and models

Feature or model	Description	Minimum software level required	Maximum per 2145-DH8 node	Maximum per 2145-SV1 node
2145-92F	SAN Volume Controller expansion enclosure for 92 3.5-inch SAS drive slots	7.8.0	8	8

Table 7. Optional features and models (continued)

Feature or model	Description	Minimum software level required	Maximum per 2145-DH8 node	Maximum per 2145-SV1 node
2145-24F	SAN Volume Controller expansion enclosure, which is required for 2.5-inch SAS drives	7.3.0 for up to two enclosures 7.7.0 for up to two enclosures (12F, 24F, or both) 7.7.1 for up to 20 enclosures	20	20
2145-12F	SAN Volume Controller expansion enclosure, which is required for 3.5-inch SAS drives	7.3.0 for up to two enclosures 7.7.0 for up to two enclosures (12F, 24F, or both) 7.7.1 for up to 20 enclosures	20	20
AH10	4-port 8 Gbps Fibre Channel adapter with four short-wave SFP transceivers **Notes:** • The maximum fiber length is 10 km when used with feature AH1T and single-mode fiber optic cable. • If you are installing four adapters, software level 7.6.0.3 is required.	7.3.0 or 7.6.0.3	4	0
AH11	2-port 16 Gbps Fibre Channel adapter with two short-wave SFP transceivers The maximum fiber length is 10 km when used with feature ACHU and single-mode fiber optic cable.	7.4.0	4	0
AH12	4-port 10 Gbps Ethernet adapter with four SFP transceivers	7.3.0 8.1.1 (Required to support two adapters.)	2	1
AH13	4-port 12 Gbps SAS adapter. Required for Model 24F attachment	7.3.0	1	1
AH14	4-port 16 Gbps Fibre Channel adapter with four short-wave SFP transceivers **Note:** The maximum fiber length is 5 km when used with feature ACHU and single-mode fiber optic cable.	7.6.0	4	4
AH16	2-port 25 Gbps Ethernet (RoCE) adapter with two SFP28 transceivers for iSCSI	8.1.1.1	0	3
AH17	2-port 25 Gbps Ethernet (iWARP) adapter with two SFP28 transceivers for iSCSI	8.1.1.1	0	3
AH1T	Two 8 Gbps Fibre Channel long-wave SFP Transceivers for optional use with feature AH10	7.3.0		N/A

Table 7. Optional features and models (continued)

Feature or model	Description	Minimum software level required	Maximum per 2145-DH8 node	Maximum per 2145-SV1 node
ACHU	Two 16 Gbps Fibre Channel long-wave SFP Transceivers for optional use with feature AH11 or AH14	7.3.0		
AH1A	Compression accelerator. Requires feature AH1B	7.3.0	2	2
AH1B	Second microprocessor and 32 GB RAM	7.3.0	1	0 (in base)
AH20	200 GB 12 Gbps SAS 2.5-inch tier 0 flash drive	7.3.0		
AH21	400 GB 12 Gbps SAS 2.5-inch tier 0 flash drive	7.3.0		
AH22	800 GB 12 Gbps SAS 2.5-inch tier 0 flash drive	7.3.0		
AH23	1.6 TB 12 Gbps SAS 2.5-inch tier 0 flash drive	7.3.0		
AH24	3.2 TB 12 Gbps SAS 2.5-inch tier 0 flash drive	7.4.0		
AH30	4.0 TB 7.2 K RPM 3.5-inch nearline disk drive	7.7.0		
AH31	6.0 TB 7.2 K RPM 3.5-inch nearline disk drive	7.7.0		
AH32	8.0 TB 7.2 K RPM 3.5-inch nearline disk drive	7.7.0		
AH33	10.0 TB 7.2 K RPM 3.5-inch nearline disk drive	7.8.0		
AH34	12 TB 7.2 K RPM 2.5-inch near-line SAS disk drive	7.6.1.4		
AH40	300 GB 15 K RPM 2.5-inch disk drive	7.6.1.4		
AH41	600 GB 15 K RPM 2.5-inch disk drive	7.6.1.4		
AH50	900 GB 10 K RPM 2.5-inch disk drive	7.6.1.4		
AH51	1.2 TB 10 K RPM 2.5-inch disk drive	7.6.1.4		
AH52	1.8 TB 10 K RPM 2.5-inch disk drive	7.6.1.4		
AH60	2.0 TB 7.2 K RPM 2.5-inch nearline disk drive	7.6.1.4		
AH2A	1.92 TB 2.5-inch SAS tier 1 flash drive	7.6.1.4		
AH2B	3.84 TB 2.5-inch SAS tier 1 flash drive	7.7.0		
AH2C	7.68 TB 2.5-inch SAS tier 1 flash drive	7.8.0		
AH2D	15.3 TB 2.5-inch SAS tier 1 flash drive	7.8.0		
AH42	900 GB 15 K RPM 2.5-inch SAS disk drive	7.6.1.4		
AH53	2.4 TB 10 K RPM 2.5-inch SAS disk drive	7.6.1.4		
AH70	600 GB 15 K RPM SAS disk drive for 92F	7.8.0		
AH73	1.2 TB 10 K RPM SAS disk drive for 92F	7.8.0		
AH74	1.8 TB 10 K RPM SAS disk drive for 92F	7.8.0		
AH75	2.4 TB 10 K RPM SAS disk drive for 92F	7.8.0		

Table 7. Optional features and models (continued)

Feature or model	Description	Minimum software level required	Maximum per 2145-DH8 node	Maximum per 2145-SV1 node
AH77	6 TB 7.2 K RPM nearline SAS disk drive for 92F	7.8.0		
AH78	8 TB 7.2 K RPM nearline SAS disk drive for 92F	7.8.0		
AH79	10 TB 7.2 K RPM nearline SAS disk drive for 92F	7.8.0		
AH7A	12 TB 7.2 K Near-Line SAS disk drive	7.6.1.4		
AH7D	1.6 TB SAS tier 0 flash drive for 92F	7.8.0		
AH7E	3.2 TB SAS tier 0 flash drive for 92F	7.8.0		
AH7J	1.92 TB SAS tier 1 flash drive for 92F	7.8.0		
AH7K	3.84 TB SAS tier 1 flash drive for 92F	7.8.0		
AH7L	7.68 TB SAS tier 1 flash drive for 92F	7.8.0		
AH7M	15.3 TB SAS tier 1 flash drive for 92F	7.8.0		

2145-SV1 PCI express expansion slot rules

Table 8 lists the optional adapters that are supported in each PCI express expansion slot.

Table 8. PCI express expansion slot rules for 2145-SV1 nodes

PCIe slot	Options that are supported in the specified slot
1	None
2	12 Gbps SAS adapter
3	4-port 16 Gbps Fibre Channel 10 Gbps Ethernet (see Note) 2-port 25 Gbps Ethernet
4	4-port 16 Gbps Fibre Channel 10 Gbps Ethernet (see Note) 2-port 25 Gbps Ethernet
5	12 Gbps SAS adapter Compression accelerator 2-port 25 Gbps Ethernet
6	4-port 16 Gbps Fibre Channel 10 Gbps Ethernet (see Note) 2-port 25 Gbps Ethernet
7	4-port 16 Gbps Fibre Channel 10 Gbps Ethernet (see Note) 2-port 25 Gbps Ethernet
8	Compression accelerator
Note: With software level 8.1.1, each node can support one 10 Gbps Ethernet adapter (see Table 7 on page 15).	

2145-DH8 PCI express expansion slot rules

Use the rules in Table 9 to see which adapters are supported in each PCI express expansion slot.

Table 9. PCI express expansion slot rules for 2145-DH8 nodes

PCIe slot	Options that are supported in the specified 2145-DH8 slot
1	8 Gbps Fibre Channel 2-port 16 Gbps Fibre Channel 4-port 16 Gbps Fibre Channel 10 Gbps Ethernet (see note 1)
2	8 Gbps Fibre Channel 2-port 16 Gbps Fibre Channel 4-port 16 Gbps Fibre Channel 10 Gbps Ethernet (see note 1)
3	8 Gbps Fibre Channel (see note 2) 2-port 16 Gbps Fibre Channel 4-port 16 Gbps Fibre Channel 10 Gbps Ethernet (see note 1) 12 Gbps SAS
4	Compression accelerator (see note 3)
5	8 Gbps Fibre Channel 2-port 16 Gbps Fibre Channel 4-port 16 Gbps Fibre Channel 10 Gbps Ethernet (see note 1)
6	Compression accelerator (see note 3)

Notes:

1. With software level 8.1.1, each node can support two 10 Gbps Ethernet adapters (see Table 7 on page 15). **Note:** If a 10 Gbps Ethernet adapter is being used for FCoE connectivity, avoid installing the second 10 Gbps Ethernet adapter in a PCIe expansion slot that is below the first adapter. When more than one 10 Gbps Ethernet adapter is installed, only the first four Ethernet ports that are detected by the system and displayed by the **lsportip** command support FCoE. The remaining ports do not support FCoE and any existing FCoE zones break. When the node is added back to the system cluster, you must manually reconfigure the zoning to make the first 10 Gbps Ethernet adapter ports visible to the host again.
2. An 8 Gbps Fibre Channel adapter in slot 3 needs a minimum software level of 7.6.0.3.
3. If the system has only one compression accelerator, it can be installed in either slot 4 or 6.

Node controls and indicators

The controls and indicators provide information about the system status and activity. They also help to identiy the node.

SAN Volume Controller 2145-SV1 front panel controls and indicators

The controls and indicators on the front panel are used for power and to indicate information such as system activity, node failures, and node identification.

Figure 11 on page 20 shows the controls and indicators on the front panel of the SAN Volume Controller 2145-SV1.

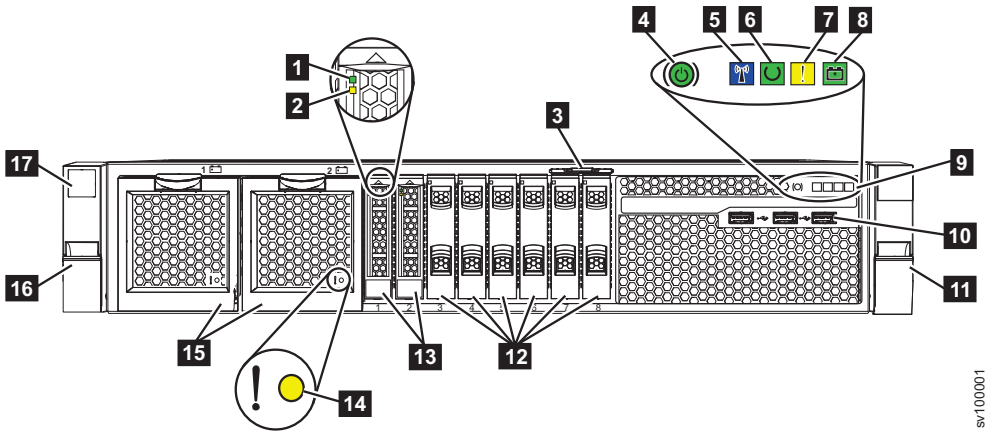

Figure 11. SAN Volume Controller 2145-SV1 front panel

1. Boot drive activity LED
2. Boot drive status LED
3. Pullout tab showing 11S serial number
4. Power-control button and power-on LED
5. Identify LED
6. Node status LED
7. Node fault LED
8. Battery status LED
9. Operator-information panel
10. Front USB ports 1-3
11. Right side latch (releases chassis to slide out on rails)
12. Drive slot fillers (no empty slots can be used)
13. Boot drives
14. Battery fault LED
15. Batteries
16. Left side latch (releases chassis to slide out on rails)
17. Machine type and model (MTM) and serial number

Boot drive activity LED

The green drive activity LED indicates one of the following conditions.

Off The drive is not ready for use.

Flashing
The drive is in use.

On The drive is ready for use, but is not in use.

Boot drive status LED

The amber drive status LED indicates one of the following conditions.

Off The drive is in a good state or has no power.

Flashing
The drive is being identified.

On The drive failed.

Battery fault LED

The amber Battery fault LED indicates one of the following conditions.

Off The battery is functioning normally.

Flashing
The battery is being identified.

On The battery failed.

SAN Volume Controller 2145-DH8 front panel controls and indicators

The controls and indicators on the front panel are used for power and to indicate information such as system activity, node failures, and node identification.

Figure 12 shows the controls and indicators on the front panel of the SAN Volume Controller 2145-DH8.

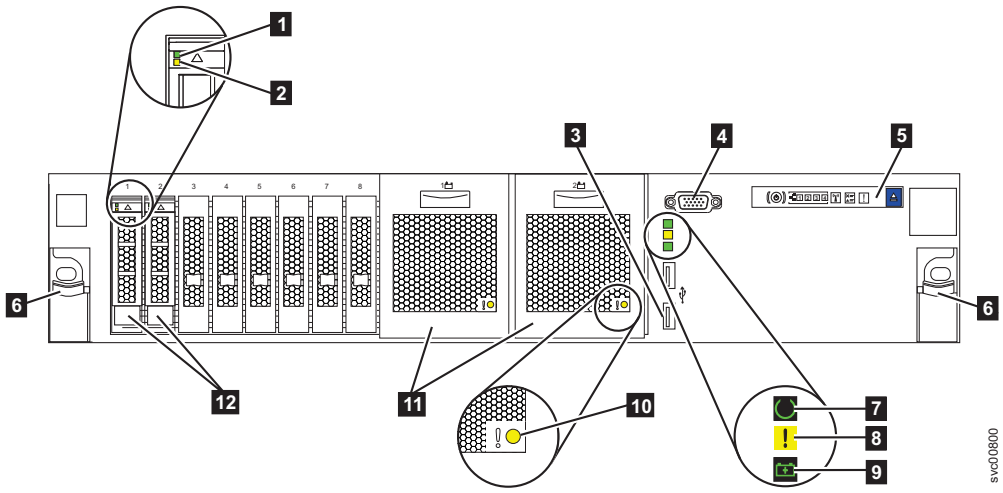

Figure 12. SAN Volume Controller 2145-DH8 front panel

1. Hard disk drive activity LED
2. Hard disk drive status LED
3. USB port
4. Video connector
5. Operator-information panel
6. Rack release latch
7. Node status LED
8. Node fault LED
9. Battery status LED
10. Battery fault LED
11. Batteries
12. Hard disk drives (boot drives)

Node status LED

The node status LED provides the following system activity indicators:

Off The node is not operating as a member of a system.

On The node is operating as a member of a system.

Slow blinking
The node is in candidate or service state.

Fast blinking
The node is dumping cache and state data to the local disk in anticipation of a system restart from a pending power-off action or other controlled restart sequence.

Node fault LED

A node fault is indicated by the amber node-fault LED.

Off The node does not have any errors that prevent it from doing I/O or the system software is not running on the node.

On The node has an unrecoverable node error and is not part of the system.

Battery status LED

The green battery status LED indicates one of the following battery conditions.

Off The system software is not running on the node or the state of the system cannot be saved if power to the node is lost.

Fast blinking
Battery charge level is too low for the state of the system to be saved if power to the node is lost. Batteries are charging.

Slow blinking
Battery charge level is sufficient for the state of the system to be saved **once** if power to the node is lost.

On Battery charge level is sufficient for the state of the system to be saved **twice** if power to the node is lost.

Battery fault LED

The amber battery fault LED indicates one of the following battery conditions.

Off The system software is not running on the node or this battery does not have a fault.

Blinking
This battery is being identified.

On This battery has a fault. It cannot be used to save the system state if power to the node is lost.

Hard disk drive activity LED

The green drive activity LED indicates one of the following conditions.

Off The drive is not ready for use.

Flashing
The drive is in use.

On The drive is ready for use, but is not in use.

Hard disk drive status LED

The amber drive status LED indicates one of the following conditions.

Off The drive is in a good state or has no power.

Blinking
The drive is being identified.

On The drive has failed.

Node status LED

System activity is indicated through the green node-status LED.

The node status LED provides the following system activity indicators:

Off The node is not operating as a member of a system.

On The node is operating as a member of a system.

Slow blinking
The node is in candidate or service state.

Fast blinking
The node is dumping cache and state data to the local disk in anticipation of a system reboot from a pending power-off action or other controlled restart sequence.

Product serial number

The node contains a product serial number that is written to the system board hardware. The product serial number is also printed on the serial number label that is on the front panel.

This number is used for warranty and service entitlement checking and is included in the data that is sent with error reports.

Remember: Do not change this number during the life of the product. If the system board is replaced, you must follow the system board replacement instructions carefully and rewrite the serial number on the system board.

Nodeoperator-information panel

The operator-information panel is located on the front panel of the node.

SAN Volume Controller 2145-SV1 operator-information panel

The operator-information panel contains buttons and indicators such as the power-control button, and LEDs that provide node information.

Figure 13 on page 24 shows the operator-information panel for the SAN Volume Controller 2145-SV1.

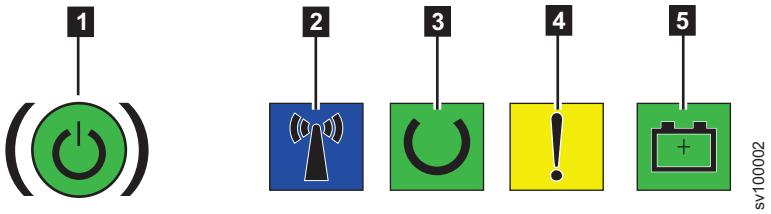

Figure 13. SAN Volume Controller 2145-SV1 operator-information panel

1. Power-control button and power-on LED
2. Identify LED
3. Node status LED
4. Node fault LED
5. Battery status LED

Power LED

The green power LED indicates one of the following power conditions.

Off One or more of the following are true:

- No power is present at the power supply input.
- The power supply has failed.
- The LED has failed.

On The node is turned on.

Blinking
The node is turned off, but is still connected to a power source.

Power button

The power button turns main power on or off for the SAN Volume Controller .

- To turn on the power, press and release the power button.
- To turn off the power, press and release the power button. For more information about what to check before you turn off the SAN Volume Controller node, see "MAP 5350: Powering off a node."

Attention: When the node is operational and you press and immediately release the power button, the SAN Volume Controller writes its control data to its internal disk and then turns off. This process can take up to 5 minutes.

Identify LED

This LED blinks if the Identify button on the back of node is pressed. The Identify LED blinks on both the front and rear panels. Use this feature to find a specific node in the data center. After the SAN Volume Controller system is initialized and initial setup is completed, you can use the Management GUI to identify a node by making the Identify LED on the node blink.

Node status LED

The green Node status LED has the following states:

Off The SAN Volume Controller software is not running or cannot communicate with this LED.

On This node is active in a SAN Volume Controller system.

Slow blink
This node is not active. It has Candidate or Service status.

Fast blink
The node is dumping cache and state data to the local disk in anticipation of a system reboot from a pending power-off action or other controlled restart sequence.

Node fault LED

The yellow Node fault LED has the following states:

Off No warning or critical error is shown in the baseboard management controller (BMC) event log, and no fatal node error is reported by the SAN Volume Controller software.

On The SAN Volume Controller software indicates a fatal node error.

Blinking
A warning or critical error is shown in the BMC event log.

Battery status LED

The green battery status LED has the following states:

Off Hardened data is not saved if there is power loss or the SAN Volume Controller software is not running.

On Battery charge level is sufficient for the hardened data to be saved twice if power to the node is lost.

Slow blink
Battery charge level is sufficient for the hardened data to be saved once if power to the node is lost.

Fast blink
Battery charge level is too low for the hardened data to be saved if power to the node is lost. Batteries are charging.

SAN Volume Controller 2145-DH8 operator information panel

The operator-information panel indicates information such as system board errors, Ethernet activity, and power status.

Figure 14 on page 26 shows the operator-information panel for the SAN Volume Controller 2145-DH8.

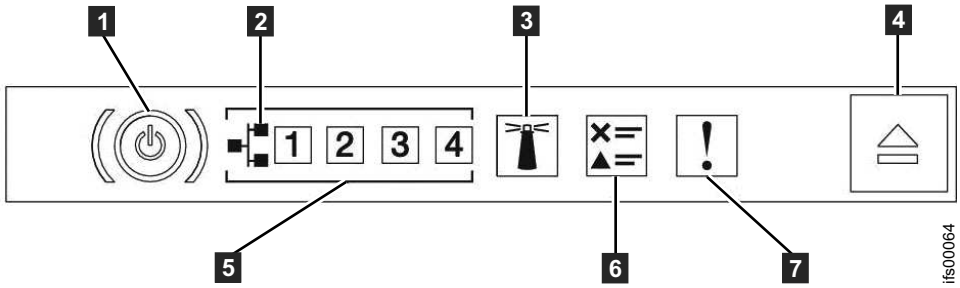

Figure 14. SAN Volume Controller 2145-DH8 operator information panel

1. **Power-control** button and power-on LED (green)
2. **Ethernet** icon
3. **System-locator** button and LED (blue)
4. Release latch for the light path diagnostics panel
5. **Ethernet activity** LEDs
6. **Check log** LED
7. **System-error** LED (yellow)

Note: If the node has more than four Ethernet ports, activity on ports five and above is not reflected on the operator-information panel Ethernet activity LEDs.

System-error LED

When it is lit, the system-error LED indicates that a system-board error has occurred.

This amber LED lights up if the hardware detects an unrecoverable error that requires a new field-replaceable unit (FRU). To help you isolate the faulty FRU, see MAP 5800: Light path to help you isolate the faulty FRU.

Disk drive activity LED

When it is lit, the green disk drive activity LED indicates that the disk drive is in use.

Reset button

If a reset button is available on your SAN Volume Controller node, do not use it.

Attention: If you use the reset button, the node restarts immediately without the SAN Volume Controller control data that is being written to the disk. Service actions are then required to make the node operational again.

Power button

The power button turns main power on or off for the SAN Volume Controller .

To turn on the power, press and release the power button. You must have a pointed device, such as a pen, to press the button.

To turn off the power, press and release the power button. For more information about how to turn off the SAN Volume Controller node, see MAP 5350: Powering off a SAN Volume Controller node.

Attention: When the node is operational and you press and immediately release the power button, the SAN Volume Controller writes its control data to its internal disk and then turns off. This process can take up to 5 minutes. If you press the power button but do not release it, the node turns off immediately without the SAN Volume Controller control data that is being written to disk. Service actions are then required to make the SAN Volume Controller operational again. Therefore, during a power-off operation, do not press and hold the power button for more than 2 seconds.

Power LED

The green power LED indicates the power status of the system.

The power LED has the following properties:

Off One or more of the following are true:

- No power is present at the power supply input.
- The power supply has failed.
- The LED has failed.

On The node is turned on.

Blinking
The node is turned off, but is still connected to a power source.

System-information LED

When the system-information LED is lit, a noncritical event occurs.

Check the light path diagnostics panel and the event log. Light path diagnostics are described in more detail in the light path maintenance analysis procedure (MAP).

Locator LED

The SAN Volume Controller does not use the locator LED.

Ethernet-activity LED

An Ethernet-activity LED beside each Ethernet port indicates that theSAN Volume Controller node is communicating on the Ethernet network that is connected to the Ethernet port.

The operator-information panel LEDs refer to the Ethernet ports that are mounted on the system board. If you install the 10 Gbps Ethernet card on a SAN Volume Controller 2145-CG8, the port activity is not reflected on the activity LEDs.

Node rear-panel indicators and connectors

The rear-panel indicators for the node are on the back-panel assembly. The external connectors are on the node and the power supply assembly.

SAN Volume Controller 2145-SV1 rear-panel indicators

The rear-panel indicators consist of LEDs that indicate the status of the Fibre Channel ports, Ethernet connection and activity, power, and electrical current.

Figure 15 on page 28 shows the rear-panel indicators on the SAN Volume Controller 2145-SV1 back-panel assembly.

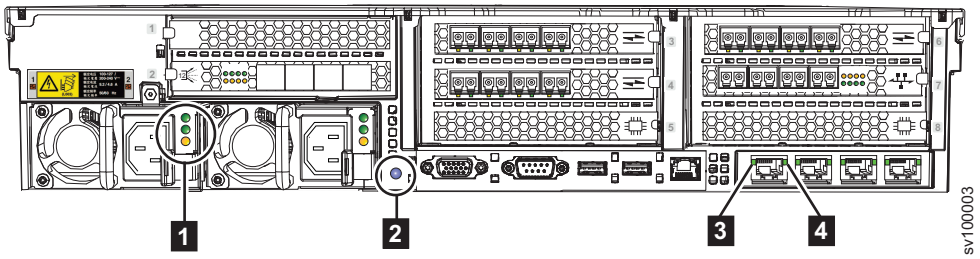

Figure 15. SAN Volume Controller 2145-SV1 rear-panel indicators

1. AC, DC, and power-supply fault LEDs
2. Identify button and LED
3. Ethernet-link LED
4. Ethernet-activity LED

SAN Volume Controller 2145-DH8 rear-panel indicators

The rear-panel indicators consist of LEDs that indicate the status of the Fibre Channel ports, Ethernet connection and activity, power, electrical current, and system-board errors.

Figure 16 shows the rear-panel indicators on the SAN Volume Controller 2145-DH8 back-panel assembly.

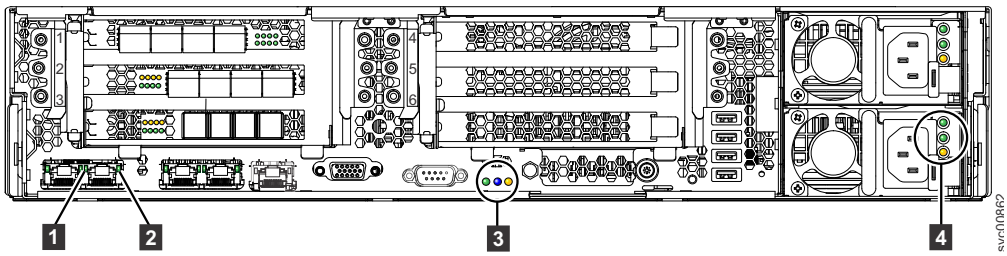

Figure 16. SAN Volume Controller 2145-DH8 rear-panel indicators

1. Ethernet-link LED
2. Ethernet-activity LED
3. Power, location, and system-error LEDs
4. AC, DC, and power-supply error LEDs

SAN Volume Controller 2145-SV1 connectors

The SAN Volume Controller 2145-SV1 includes multiple external connectors for data, video, and power.

Figure 17 on page 29 shows the external connectors on the SAN Volume Controller 2145-SV1 back panel assembly.

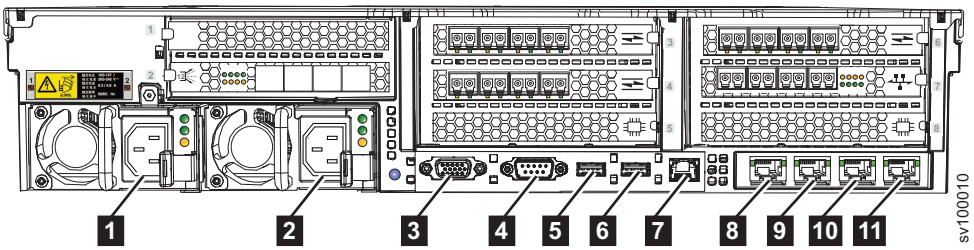

Figure 17. Connectors on the rear of the SAN Volume Controller 2145-SV1

1. Power supply 1
2. Power supply 2
3. Video port
4. Serial port (not used)
5. Rear USB port 1
6. Rear USB port 2
7. Unused Ethernet port
8. 10 Gbps Ethernet port 1
9. 10 Gbps Ethernet port 2
10. 10 Gbps Ethernet port 3
11. Technician port (Ethernet)

Figure 18 shows the type of connector that is on each power-supply assembly.

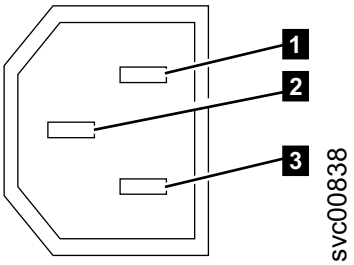

Figure 18. Power connector

1. Neutral
2. Ground
3. Live

Note: Optional host interface adapters provide extra connectors for 10 Gbps Ethernet, Fibre Channel, or SAS.

SAN Volume Controller 2145-SV1 ports used during service procedures:

The SAN Volume Controller 2145-SV1 contains a number of ports that are used during service procedures.

The following figure shows ports that are used during service procedures.

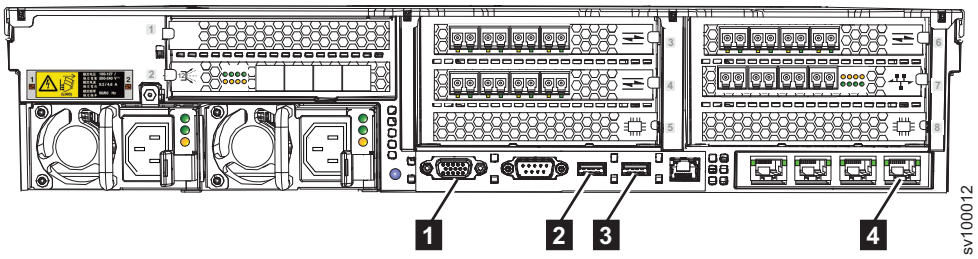

Figure 19. SAN Volume Controller 2145-SV1 service ports

1. VGA port
2. Rear USB port 1
3. Rear USB port 2
4. Technician port (Ethernet)

Any of these ports other than the Technician port can be used during normal operation. Connect a device to the Technician port only when you are directed to do so by a service procedure or by your IBM service representative.

SAN Volume Controller 2145-SV1 unused ports:

The SAN Volume Controller 2145-SV1 includes one Ethernet port and one serial port that are not used.

The following figure shows the Ethernet port that is not used during service procedures or normal operation. This port is disabled in software to make the port inactive.

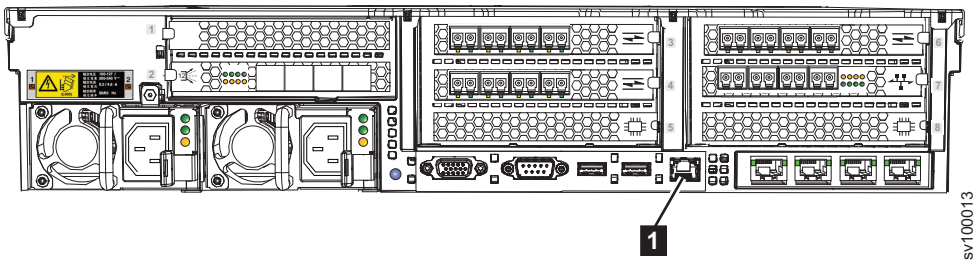

Figure 20. SAN Volume Controller 2145-SV1 unused Ethernet port

1. Unused Ethernet port

Although not disabled, the serial port is also not used in normal operation.

SAN Volume Controller 2145-SV1 Fibre Channel and Ethernet port numbers:

Fibre Channel port numbers for the SAN Volume Controller 2145-SV1 vary, depending on how many network adapters are installed, and in which slots. Port numbers also depend on the number and configuration of the Ethernet adapters.

Table 10 on page 31 lists the SAN Volume Controller 2145-SV1 expansion slots and the adapters that can be installed.

Table 10. PCIe expansion slots in which an adapter can be used

PCIe expansion slot number	Adapter
1	Not used
2	12 Gbps SAS adapter
3	16 Gbps Fibre Channel adapter, 10 Gbps Ethernet adapter, or 25 Gbps Ethernet adapter*
4	16 Gbps Fibre Channel adapter, 10 Gbps Ethernet adapter, or 25 Gbps Ethernet adapter
5	Compression Accelerator
6	16 Gbps Fibre Channel adapter, 10 Gbps Ethernet adapter, or 25 Gbps Ethernet adapter
7	16 Gbps Fibre Channel adapter, 10 Gbps Ethernet adapter, or 25 Gbps Ethernet adapter
8	Compression Accelerator

1. Slots 3, 4, 6, and 7 can contain a 10 Gbps Ethernet adapter, but only one 10 Gbps Ethernet adapter is supported.
2. Slots 3, 4, 6, and 7 can contain a 25 Gbps Ethernet adapter; however, the system only supports three 25 Gbps Ethernet adapters.

Figure 21 shows the physical Fibre Channel (FC) port numbers when the 10 Gbps Optical Ethernet adapter is configured for Fibre Channel over Ethernet (FCoE) communications.

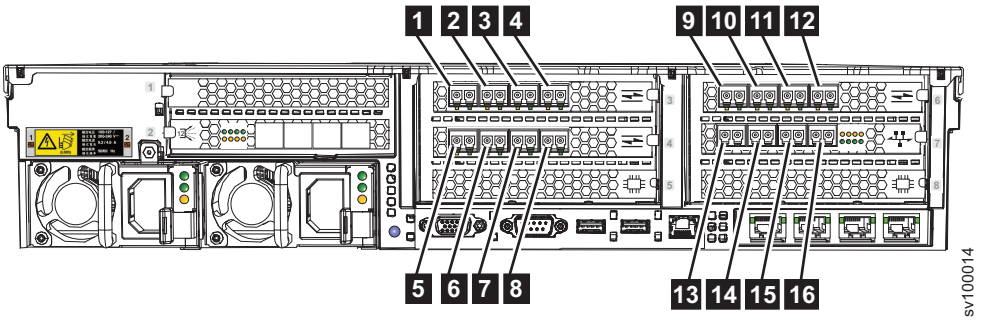

Figure 21. Fibre Channel port numbers in a typical configuration

1 - **16** Fibre Channel ports 1-16

Figure 22 on page 32 shows the Ethernet port numbers for the SAN Volume Controller 2145-SV1 when the 10 Gbps Optical Ethernet adapter is configured for iSCSI communications.

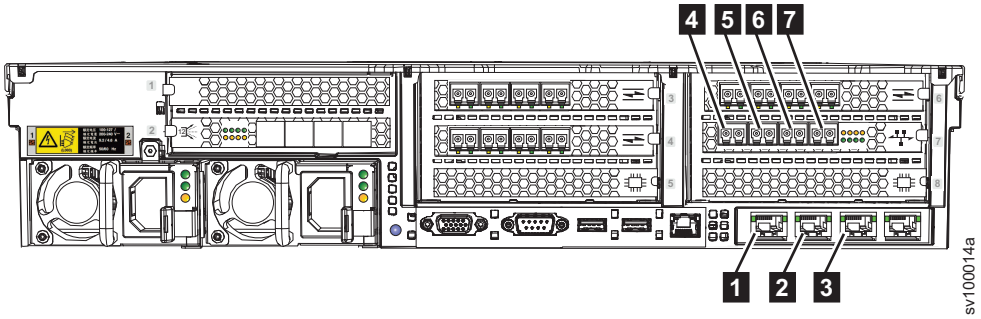

Figure 22. Ethernet port numbers for iSCSI communication (10 Gbps Ethernet adapter)

1 - **3** 10 Gbps Ethernet ports 1-3 (onboard)
4 - **7** 10 Gbps optical Ethernet ports 4-7

Figure 23 shows the Ethernet port numbers for the SAN Volume Controller 2145-SV1 when two 2-port 25 Gbps Optical Ethernet (RoCE) adapters are configured. Ethernet ports 4 and 5 are located in the Ethernet adapter that is installed in the lowest PCIe expansion slot number.

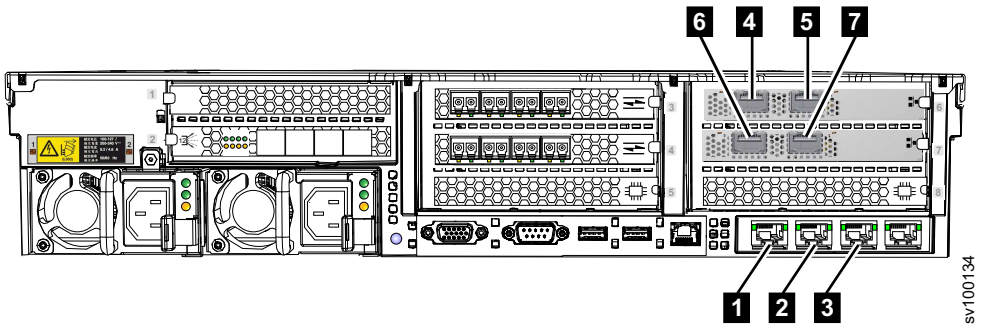

Figure 23. Ethernet port numbers for 25 Gbps adapter

1 - **3** 10 Gbps Ethernet ports 1-3 (onboard)
4 - **7** 25 Gbps optical Ethernet ports 4-7

SAN Volume Controller 2145-DH8 connectors

The SAN Volume Controller 2145-DH8 includes multiple external connectors for data, video, and power.

Figure 24 on page 33 shows the external connectors on the SAN Volume Controller 2145-DH8 back panel assembly.

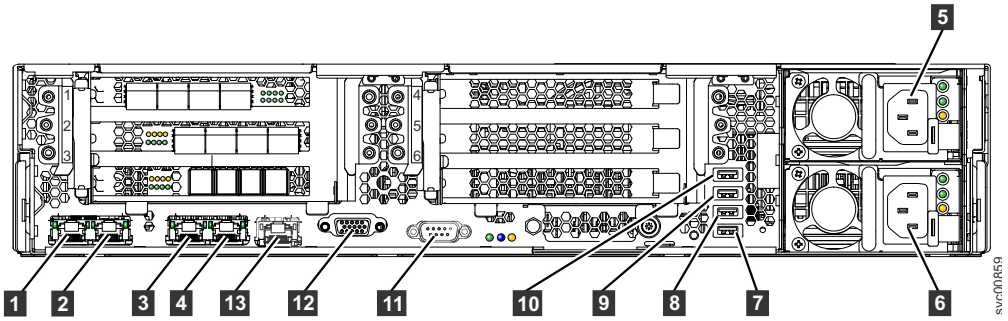

Figure 24. Connectors on the rear of the SAN Volume Controller 2145-DH8

1. 1 Gbps Ethernet port 1
2. 1 Gbps Ethernet port 2
3. 1 Gbps Ethernet port 3
4. Technician port (Ethernet)
5. Power supply 2
6. Power supply 1
7. USB 6
8. USB 5
9. USB 4
10. USB 3
11. Serial
12. Video
13. Unused Ethernet port

Figure 25 shows the type of connector that is on each power-supply assembly.

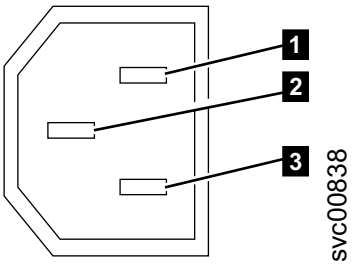

Figure 25. Power connector

1. Neutral
2. Ground
3. Live

Note: Optional host interface adapters provide extra connectors for 10 Gbps Ethernet, Fibre Channel, or SAS connections.

SAN Volume Controller 2145-DH8 ports used during service procedures:

The SAN Volume Controller 2145-DH8 contains a number of ports that are only used during service procedures.

Figure 26 shows ports that are used only during service procedures.

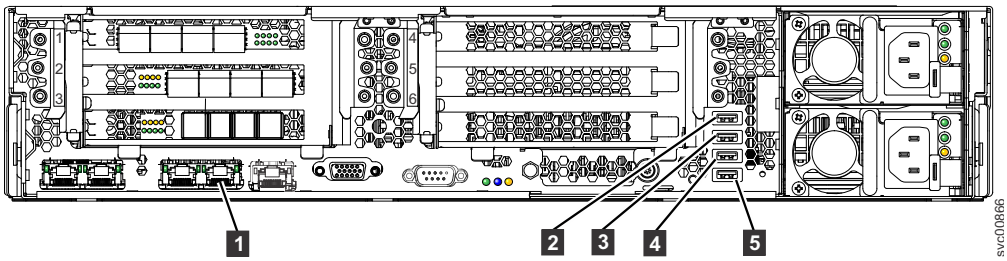

Figure 26. SAN Volume Controller 2145-DH8 service ports

1. Technician port (Ethernet)
2. USB 3
3. USB 4
4. USB 5
5. USB 6

During normal operations, none of these ports are used. Connect a device to any of these ports only when you are directed to do so by a service procedure or by an IBM service representative.

SAN Volume Controller 2145-DH8 unused ports:

The SAN Volume Controller 2145-DH8 includes one port that is not used.

Figure 27 shows the one port that is not used during service procedures or normal operation. This port is disabled in software to make the port inactive.

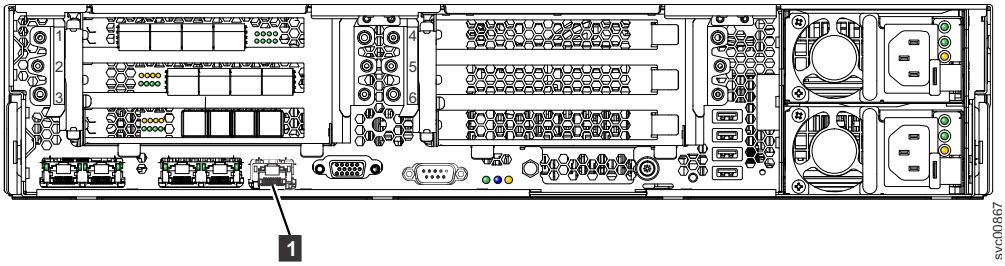

Figure 27. SAN Volume Controller 2145-DH8 unused Ethernet port

1. Unused Ethernet port

Fibre Channel LEDs

The Fibre Channel LEDs indicate the status of the Fibre Channel ports on the SAN Volume Controller 2145-DH8 node.

The SAN Volume Controller 2145-DH8 uses two light-emitting diodes (LEDs) per Fibre Channel port, which are arranged one above the other. The LEDs are arranged in the same order as the ports. Figure 28 on page 35 shows the location of the LEDs.

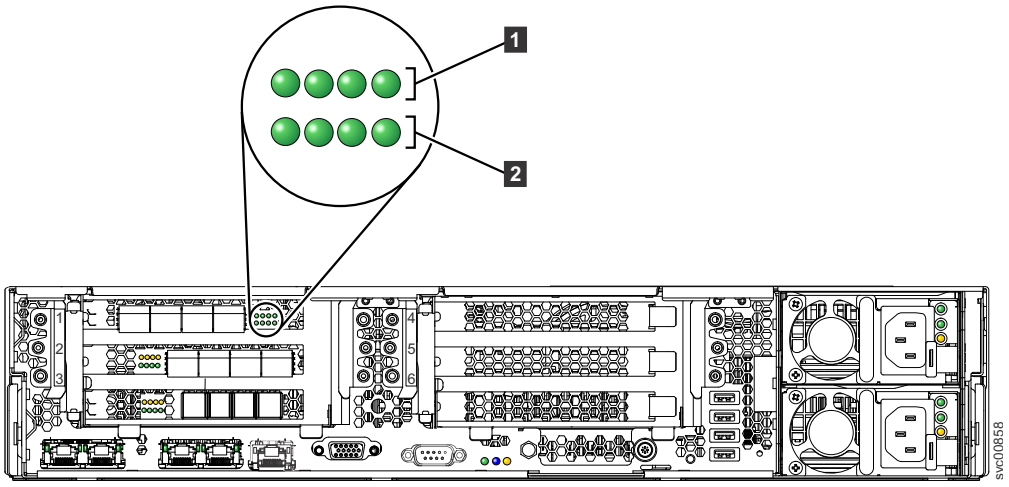

Figure 28. Fibre Channel LEDs

1. Link speed LEDs
2. Link activity LEDs

The following table lists the link status values for the Fibre Channel LEDs.

Table 11. Link status values for Fibre Channel LEDs

Top LED (link speed)	Bottom LED (link activity) Flashing indicates I/O activity.	Link status
Off	Off	Inactive
Off	On / Flashing	Active 2 Gbps
Blinking	On / Flashing	Active 4 Gbps
On	On / Flashing	Active 8 Gbps
Note: To accommodate the different Fibre Channel speed ranges, LEDs are effectively OFF=slow, FLASHING=medium, and ON=fast.		

Ethernet activity LED

The Ethernet activity LED indicates that the node is communicating with the Ethernet network that is connected to the Ethernet port.

There is a set of LEDs for each Ethernet connector. The top LED is the Ethernet link LED. When it is lit, it indicates that there is an active connection on the Ethernet port. The bottom LED is the Ethernet activity LED. When it flashes, it indicates that data is being transmitted or received between the server and a network device.

Ethernet link LED

The Ethernet link LED indicates that there is an active connection on the Ethernet port.

There is a set of LEDs for each Ethernet connector. The top LED is the Ethernet link LED. When it is lit, it indicates that there is an active connection on the Ethernet port. The bottom LED is the Ethernet activity LED. When it flashes, it indicates that data is being transmitted or received between the server and a network device.

Power, location, and system-error LEDs

The power, location, and system-error LEDs are housed on the rear of the SAN Volume Controller . These three LEDs are duplicates of the same LEDs that are shown on the front of the node.

The following terms describe the power, location, and system-error LEDs:

Power LED

This is the top of the three LEDs and indicates the following states:

Off One or more of the following are true:

- No power is present at the power supply input
- The power supply has failed
- The LED has failed

On The SAN Volume Controller is powered on.

Blinking

The SAN Volume Controller is turned off but is still connected to a power source.

Location LED

This is the middle of the three LEDs and is not used by the SAN Volume Controller .

System-error LED

This is the bottom of the three LEDs that indicates that a system board error has occurred. The light path diagnostics provide more information.

AC and DC LEDs

The AC and DC LEDs indicate whether the node is receiving electrical current.

AC LED

The upper LED indicates that AC current is present on the node.

DC LED

The lower LED indicates that DC current is present on the node.

AC, DC, and power-supply error LEDs:

The AC, DC, and power-supply error LEDs indicate whether the node is receiving electrical current.

Figure 29 on page 37 shows the location of the SAN Volume Controller 2145-DH8 AC, DC, and power-supply error LEDs.

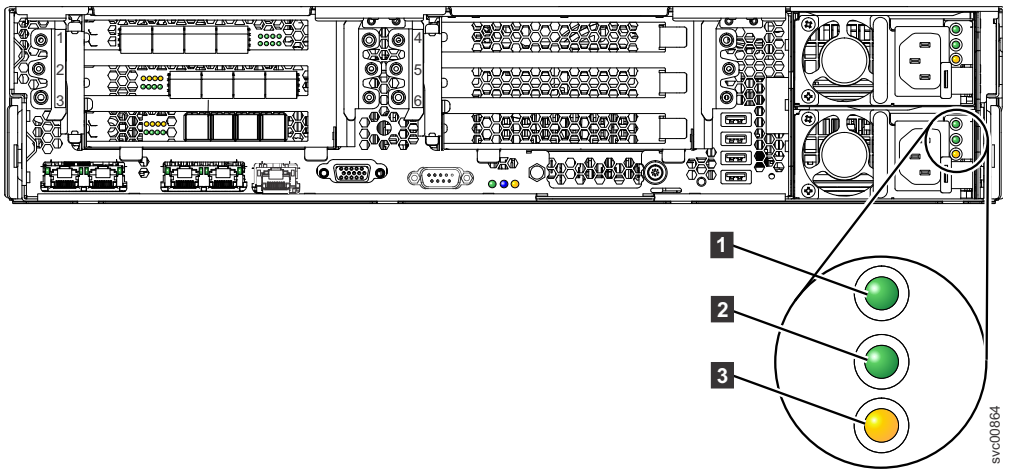

Figure 29. SAN Volume Controller 2145-DH8 AC, DC, and power-error LEDs

Each of the two power supplies has its own set of LEDs.

1 Indicates that AC current is present on the node.

2 Indicates that DC current is present on the node.

3 Indicates a problem with the power supply.

Fibre Channel port numbers and worldwide port names

Fibre Channel (FC) ports are identified by their physical port number and by a worldwide port name (WWPN).

The physical port numbers identify Fibre Channel adapters and cable connections when you run service tasks. Worldwide port names (WWPNs), which uniquely identify the devices on the SAN, are used for tasks such as Fibre Channel switch configuration. The WWPNs are derived from the worldwide node name (WWNN) of the node in which the ports are installed.

Requirements for the SAN Volume Controller environment

Certain specifications for the physical site of the SAN Volume Controller must be met before the IBM representative can set up your SAN Volume Controller environment.

SAN Volume Controller 2145-SV1 environment requirements

Before the SAN Volume Controller 2145-SV1 is installed, the physical environment must meet certain requirements. This includes verifying that adequate space is available and that requirements for power and environmental conditions are met.

Input-voltage requirements

Ensure that your environment meets the voltage requirements that are shown in Table 12.

Table 12. Input-voltage requirements

Voltage	Frequency
100-127 / 200-240Vac	50 Hz or 60 Hz

Maximum power requirements for each node

Ensure that your environment meets the power requirements as shown in Table 13.

The maximum power that is required depends on the node type and the optional features that are installed.

Table 13. Power consumption

Components	Power requirements
SAN Volume Controller 2145-SV1	~450 W typical, 700 W maximum (200 - 240V ac, 50/60 Hz)

Environment requirements without redundant AC power

Ensure that your environment falls within the following ranges if you are not using redundant AC power.

If you are not using redundant ac power, ensure that your environment falls within the ranges that are shown in Table 14.

Table 14. Physical specifications

Environment	Temperature	Altitude	Relative humidity	Maximum dew point
Operating in lower altitudes	5°C to 40°C (41°F to 104°F)	0 - 950 m (0 ft to 3,117 ft)	8% to 85%	24°C (75°F)
Operating in higher altitudes	5°C to 28°C (41°F to 82°F)	951 m to 3,050 m (3,118 ft to 10,000 ft)		
Turned off (with standby power)	5°C to 45°C (41°F to 113°F)	0 m to 3,050 m (0 ft to 10,000 ft)	8% to 85%	27°C (80.6°F)
Storing	1°C to 60°C (33.8°F to 140.0°F)	0 m to 3,050 m (0 ft to 10,000 ft)	5% to 80%	29°C (84.2°F)
Shipping	-40°C to 60°C (-40°F to 140.0°F)	0 m to 10,700 m (0 ft to 34,991 ft)	5% to 100%	29°C (84.2°F)

Note: Decrease the maximum system temperature by 1°C for every 175 m increase in altitude.

Preparing your environment

The following tables list the physical characteristics of a SAN Volume Controller 2145-SV1 node.

Dimensions and weight

Use the parameters that are shown in Table 15 on page 39 to ensure that space is available in a rack capable of supporting the node.

Table 15. Dimensions and weight

Height	Width	Depth	Maximum weight
87 mm (3.4 in.)	447 mm (17.6 in)	746 mm (30.1 in)	25 kg (55 lb) to 30 kg (65 lb) depending on configuration

Additional space requirements

Ensure that space is available in the rack for the additional space requirements around the node, as shown in Table 16.

Table 16. Additional space requirements

Location	Additional space requirements	Reason
Left side and right side	Minimum: 50 mm (2 in.)	Cooling air flow
Back	Minimum: 100 mm (4 in.) If the cable-management arm is used, allow 177 mm (7 in.)	Cable exit

Maximum heat output of each SAN Volume Controller 2145-SV1 node

The node dissipates the maximum heat output that is given in Table 17.

Table 17. Maximum heat output of each SAN Volume Controller 2145-SV1 node

Model	Heat output per node
SAN Volume Controller 2145-SV1	• Minimum configuration: 419.68 Btu per hour (AC 123 watts) • Maximum configuration: 3480.24 Btu per hour (AC 1020 watts)

SAN Volume Controller 2145-DH8 environment requirements

Before the SAN Volume Controller 2145-DH8 is installed, the physical environment must meet certain requirements. This includes verifying that adequate space is available and that requirements for power and environmental conditions are met.

Input-voltage requirements

Ensure that your environment meets the voltage requirements that are shown in Table 18.

Table 18. Input-voltage requirements

Voltage	Frequency
100-127 / 200-240Vac	50 Hz or 60 Hz

Maximum power requirements for each node

Ensure that your environment meets the power requirements as shown in Table 19 on page 40.

The maximum power that is required depends on the node type and the optional features that are installed.

Table 19. Power consumption

Components	Power requirements
SAN Volume Controller 2145-DH8	200 W typical, 750 W maximum (200 - 240V ac, 50/60 Hz)

Note: You cannot mix ac and dc power sources; the power sources must match.

Environment requirements without redundant AC power

Ensure that your environment falls within the following ranges if you are not using redundant AC power.

If you are not using redundant ac power, ensure that your environment falls within the ranges that are shown in Table 20.

Table 20. Physical specifications

Environment	Temperature	Altitude	Relative humidity	Maximum dew point
Operating in lower altitudes	5°C to 40°C (41°F to 104°F)	0 to 950 m (0 ft to 3,117 ft)	8% to 85%	24°C (75°F)
Operating in higher altitudes	5°C to 28°C (41°F to 82°F)	951 m to 3,050 m (3,118 ft to 10,000 ft)		
Turned off (with standby power)	5°C to 45°C (41°F to 113°F)	0 m to 3,050 m (0 ft to 10,000 ft)	8% to 85%	27°C (80.6°F)
Storing	1°C to 60°C (33.8°F to 140.0°F)	0 m to 3,050 m (0 ft to 10,000 ft)	5% to 80%	29°C (84.2°F)
Shipping	-40°C to 60°C (-40°F to 140.0°F)	0 m to 10,700 m (0 ft to 34,991 ft)	5% to 100%	29°C (84.2°F)

Note: Decrease the maximum system temperature by 1°C for every 175 m increase in altitude.

Preparing your environment

The following tables list the physical characteristics of the 2145-DH8 node.

Dimensions and weight

Use the parameters that are shown in Table 21 to ensure that space is available in a rack capable of supporting the node.

Table 21. Dimensions and weight

Height	Width	Depth	Maximum weight
86 mm (3.4 in.)	445 mm (17.5 in)	746 mm (29.4 in)	25 kg (55 lb) to 30 kg (65 lb) depending on configuration

Additional space requirements

Ensure that space is available in the rack for the additional space requirements around the node, as shown in Table 22.

Table 22. Additional space requirements

Location	Additional space requirements	Reason
Left side and right side	Minimum: 50 mm (2 in.)	Cooling air flow
Back	Minimum: 100 mm (4 in.)	Cable exit

Maximum heat output of each 2145-DH8 node

The node dissipates the maximum heat output that is given in Table 23.

Table 23. Maximum heat output of each 2145-DH8 node

Model	Heat output per node
2145-DH8	• Minimum configuration: 419.68 Btu per hour (AC 123 watts) • Maximum configuration: 3480.24 Btu per hour (AC 1020 watts)

Parts listing

Part numbers are available for the different parts and field-replaceable units (FRUs) of the nodes, expansion enclosures, the redundant AC-power switch, and the uninterruptible power-supply unit.

The system supports several different types of models. A label on the front of the node indicates the node type, hardware revision (if appropriate), and serial number.

SAN Volume Controller 2145-SV1 parts

The only replaceable SAN Volume Controller 2145-SV1 parts are the field-replaceable units (FRUs) which are replaced by service support representatives (SSRs). There are no customer replaceable parts (CRUs).

For more information about the terms of the warranty and getting service and assistance, see the *Warranty and Support Information* document.

SAN Volume Controller 2145-SV1 replaceable units

Table 24 provides the part numbers and brief descriptions of the SAN Volume Controller 2145-SV1 parts.

Table 24. FRUs in the SAN Volume Controller 2145-SV1 parts assembly

FRU part Number	Quantity	Description
01EJ624	2	Battery
00RY543	1	3.0-volt CMOS battery
01AF423	6	Drive slot filler

Table 24. FRUs in the SAN Volume Controller 2145-SV1 parts assembly (continued)

FRU part Number	Quantity	Description
01EJ360	2	Intel E5-2667v4 8c 3.2 GHz 135W microprocessor
01EJ361	4, 8, 12, or 16	16 GB DDR4 DIMM
01EJ260	2	240 GB SATA flash drive assembly
01EJ362	1	Battery backplane power cable
01EJ363	1	Battery backplane power sense cable
01EJ364	1	Battery backplane LPC cable
01EJ365	1 set	Slide rails
01EJ366	1	Cable management arm (CMA)
01EJ367	1	Chassis metalwork kit (the enclosure without all the other FRUs)
01EJ368	1	SV1 operator information panel
01EJ369	1	Front left ear assembly
01EJ370	1	Front right ear assembly
01EJ372	1	Operator information panel USB cable
01EJ373	1	Operator information panel LED and power button cable
01EJ374 01YM716	1	SATA drive backplane
01EJ375	1	SATA drive backplane power cable
01EJ376	2	SATA drive backplane SATA cable
01EJ377	2	AC power supply unit
01EJ378	6	Fan module
01EJ379	1	Fan cage assembly
01EJ380	1	Trusted Platform Module (TPM)
01EJ381 01YM718	1	Main board with tray
01EJ382	1	Microprocessor heat sink
01EJ383	2	3-slot PCIe riser assembly
01EJ384	1	1-slot PCIe riser assembly
01EJ385	1	4-port Ethernet edge board

Table 24. FRUs in the SAN Volume Controller 2145-SV1 parts assembly (continued)

FRU part Number	Quantity	Description
01EJ387	1	Top cover, front
01EJ389	1	Top cover, back
01LJ163	1	Battery backplane
00WY983	0 - 4	4-port 16 Gbps Fibre Channel adapter
01LJ590	0 - 3	2-port 25 Gbps Ethernet (RoCE) adapter
01LJ591	0 - 3	2-port 25 Gbps Ethernet (iWARP) adapter
00AR319	0 or 1	4-port 10 Gbps optical Ethernet adapter
01AC573	0 or 1	12 Gbps SAS adapter
00RY191	0 - 4	16 Gbps long-wave SFP
31P1549	0 - 4	10 Gbps short-wave SFP
00RY190	0 - 16	16 Gbps short-wave SFP
01FT777	0 - 3	25 Gbps short wave SFP28 (RoCE)
01NN193	0 - 3	25 Gbps short wave SFP28 (iWARP)
01EJ817	0 - 2	Compression accelerator
39M5700	0 - 16	5 m fiber cable
39M5701	0 - 16	25 m fiber cable
45D4774	0 - 3	5 m OM3 optical cable
41V2120	0 - 4	10 m OM3 fiber cable
15R8848	0 - 3	25 m OM3 optical cable
39M5068	0 or 2	Power cord, Argentina
39M5080	0 or 2	Power cord, Chicago
39M5081	0 or 2	Power cord, US / group 1
39M5102	0 or 2	Power cord, Austrailia / New Zealand
39M5123	0 or 2	Power cord, Europe / Africa
39M5130	0 or 2	Power cord, Denmark
39M5144	0 or 2	Power cord, South Africa
39M5151	0 or 2	Power cord, EMEA
39M5158	0 or 2	Power cord, Switzerland
39M5165	0 or 2	Power cord, Chile / Italy
39M5172	0 or 2	Power cord, Israel
39M5199	0 or 2	Power cord, Japan
39M5206	0 or 2	Power cord, China
39M5219	0 or 2	Power cord, Korea
39M5226	0 or 2	Power cord, India

Table 24. FRUs in the SAN Volume Controller 2145-SV1 parts assembly (continued)

FRU part Number	Quantity	Description
39M5240	0 or 2	Power cord, Brazil
39M5247	0 or 2	Power cord, Taiwan
39M5377	0 or 2	Power cord, PDU connection
41Y9292	1	Thermal grease
59P4739	1	Alcohol wipes

SAN Volume Controller 2145-DH8 parts

The only replaceable SAN Volume Controller 2145-DH8 parts are the field-replaceable units (FRUs) which are replaced by IBM Service Support Representatives (SSRs). No customer replaceable parts (CRUs) are available.

For information about the terms of the warranty and getting service and assistance, see the *Warranty and Support Information* document.

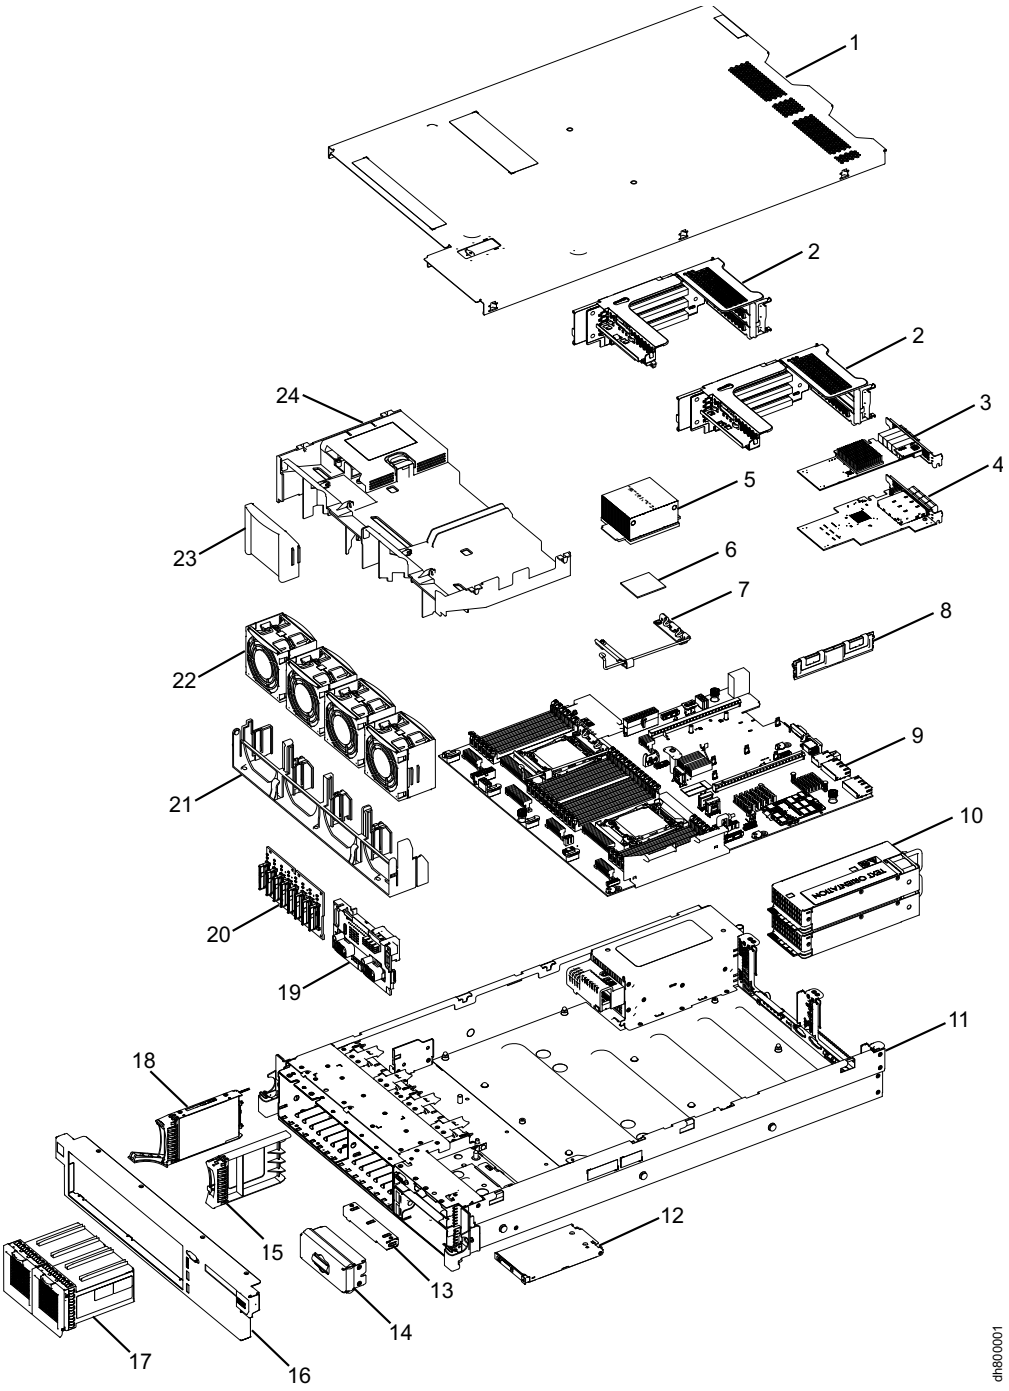

Figure 30. SAN Volume Controller 2145-DH8 replaceable parts in exploded view diagram

SAN Volume Controller 2145-DH8 replaceable units

The following tables identify part numbers and provide brief descriptions of the SAN Volume Controller 2145-DH8 parts. Use the assembly index number to locate and identify the parts that are shown in Figure 30.

- Table 25 on page 46 calls out the FRUs that are referred to in service procedures.
- Table 26 on page 48 calls out the FRUs that are not referred to by any SAN Volume Controller 2145-DH8 service procedure, but that might be replaced in some circumstances.

- Table 27 on page 49 calls out the FRU parts that are required by the long-wave small-form factor pluggable (SFP) transceiver feature.

Table 25. FRUs in the SAN Volume Controller 2145-DH8 parts assembly

Figure Index	FRU part Number	Quantity	Description
1	94Y6622	1	Top cover assembly
2	94Y6704	2	PCI Express riser card assembly. Each expansion slot might contain one of the optional adapters. There must be at least one Fibre Channel (FC) or one 10 gigabits-per-second (Gbps) Ethernet adapter in riser card assembly 1.
3	64P8485	0-1	12 Gbps SAS adapter (optional). This adapter connects the SAN Volume Controller 2145-DH8 to the SAN Volume Controller 2145-24F expansion enclosure. It is installed into PCI express expansion slot 3.
4	31P1702	0-3	A 4-port 8 Gbps FC adapter (optional). **Important:** If the system is using alternative SFPs, replace the SFPs on the FRU part with the SFPs from the FC adapter that is being replaced.
	31P1630	0-12	8 Gbps Shortwave small form-factor pluggable (SFP) transceiver. This SFP transceiver provides an auto-negotiating 2, 4, or 8 Gbps shortwave optical connection on a 8 Gbps FC adapter. **Important:** It is possible that SFPs other than those that are shipped with the product are in use on the FC host bus adapter. It is a customer responsibility to obtain replacement parts for such SFPs. The FRU part number is shown as "Non-standard - supplied by customer" in the vital product data.
	00RY004	0-4	2-port 16 Gbps FC host bus adapter (optional). **Important:** If the system is using alternative SFPs, replace the SFPs on the FRU part with the SFPs from the FC adapter that is being replaced.
	00WY983	0-4	4-port 16 Gbps FC adapter (optional). **Important:** • If the system is using alternative SFPs, replace the SFPs on the FRU part with the SFPs from the FC adapter that is being replaced. • Before you add this adapter, ensure that the system is running software version 7.6 or later.
	00RY190	0-16	16 Gbps Shortwave small form-factor pluggable (SFP) transceiver. This SFP transceiver provides an auto-negotiating 2, 4, 8 or 16 Gbps shortwave optical connection on a 16 Gbps FC adapter. **Important:** It is possible that SFPs other than those that are shipped with the product are in use on the FC adapter. It is the customer responsibility to obtain replacement parts for such SFPs. The FRU part number is shown as "Non-standard - supplied by customer" in the vital product data.
	00AR319	0-1	10 Gbps Ethernet adapter (optional). This includes a 10 Gbps Ethernet adapter that provides connectivity for up to four 10 Gbps fiber optic Ethernet cables. These cables are used for Fibre Channel over Ethernet (FCoE) and for iSCSI communications.
	31P1549	0-4	10 Gbps Shortwave SFP small form-factor pluggable (SFP) transceiver.
	00AR065	0-2	Compression accelerator (optional). This option accelerates I/O between nodes and compressed volumes. The second microprocessor and eight memory modules must be installed. The compression accelerator can be installed only in PCI expansion slots 4 and 6.

Table 25. FRUs in the SAN Volume Controller 2145-DH8 parts assembly (continued)

Figure Index	FRU part Number	Quantity	Description
5	94Y6618	1-2	Heat sink. 95 W heat sink for the microprocessor. When you replace this part, you need alcohol wipes and thermal grease.
6	00Y2783	1-2	Microprocessor. Intel Xeon E5-2650V2, 2.60 GHz, 8 core, 20 MB cache, 95 W. **Important:** This part is the microprocessor only. When replaced, you must also have alcohol wipes and thermal grease.
7	94Y7739	1-2	Heat sink retention module.
8	00D5034	4-8	Memory module. 8 GB, single-rank, 1.5 V, DDR3, 1866 MHz, RDIMM. Four memory modules are installed if there is one microprocessor. Eight memory modules are installed if two microprocessors are available.
9	00AM209	1	System board. **Important:** This part is also called the *planar*, and is the system board only. When you replace this part, you must use the microprocessor, DIMMs, and CMOS battery from the system board that you are replacing.
	33F8354	1	CMOS battery. 3.0 V. This part maintains the system BIOS settings.
10	94Y8114 or 94Y8116	2	Power supply unit. Two power units are shown in Figure 30 on page 45.
11	94Y6619	1	Safety cover. 240 V AC.
12	00AM393	1	Operator-information panel This assembly includes the information panel that contains the power-control button and diagnostic LEDs.
	90Y4768	1	Operator-information panel cable.
13	00KA089	1	DVD bay EMC shield.
14	00AR186	1	Tape bay EMC shield.
15	44T2248	6	Drive-slot blank EMC filler assembly.
16	00WY584	1	Bezel with node LEDs.
	00NV626	1	Bezel overlay This part fits over the bezel.
17	01EJ624	2	Battery. The batteries provide temporary power to save the write cache and node status to disk if main power is lost. Two batteries are shown in Figure 30 on page 45.
18	90Y8878	2	Boot disk drive. 300 GB, SAS, 2.5 inches.

Table 25. FRUs in the SAN Volume Controller 2145-DH8 parts assembly (continued)

Figure Index	FRU part Number	Quantity	Description
19	00RY001	1	Battery backplane. This part manages the batteries and switches the node to battery power if main power is lost.
	81Y6674	2	SAS signal cable. 820 mm, SAS. Connects the disk drive backplane to the system board.
	81Y6773	1	Disk drive backplane configuration cable.
20	46W9187	1	Disk drive backplane. Hot-swappable, SAS, 2.5 inches.
	00FK347	1	Disk and battery backplane power and emergency power off warning (EPOW) cable. The EPOW cable is a Y cable; one end connects to the system board and the other two ends connect to the disk drive backplane and the battery backplane.
	00AR497	1	Battery backplane power cable. Supplied with dummy DIMMs.
	00RY335	1	Battery backplane voltage sense cable.
	00AR499	1	Battery backplane low-pin count (LPC) cable.
	00AR496	1	Battery backplane LPC cable converter with clip. This connects the battery backplane LPC cable to the system board.
21	00AM212	1	Fan cage.
22	94Y6620	3-4	Fan assembly. This part is used in each of the 4 fan positions. Four assemblies are shown in Figure 30 on page 45.
23	94Y6736	0-1	Fan blank. This part is used in place of fan 4 when only one microprocessor is installed.
24	94Y6624	1	Airflow baffle.

SAN Volume Controller 2145-DH8 cable replaceable units

Table 26. FRUs to which SAN Volume Controller 2145-DH8 service procedures do not refer

Description	FRU part number
Microprocessor installation tool	94Y9955
Thermal grease	41Y9292
Alcohol wipes	59P4739
Support rails	94Y6719
Cable management arm assembly (2U)	90Y6464
VGA cable	81Y6775
USB cable	81Y6770
USB module	94Y6629
Power paddle card	69Y5787

Table 26. FRUs to which SAN Volume Controller 2145-DH8 service procedures do not refer (continued)

Description	FRU part number
Miscellaneous parts kit	94Y6746
EIA set kit	49Y5356
Bezel screws	00D3010
5 m FC cable	39M5700
25 m FC cable	39M5701
Ethernet Cat 5E cable	46X0581
2.0 m jumper cable	39M5376

SAN Volume Controller 2145-DH8 SFP replaceable units

Table 27. FRU parts for the long-wave small form-factor pluggable (SFP) transceiver feature

Description	FRU part number	Feature Code
8 Gbps Long-wave SFP transceiver. **Important:** It is possible that SFP transceivers other than those shipped with the product are in use on the FC host bus adapter. It is a customer responsibility to obtain replacement parts for the SFP transceiver. The FRU part number is shown as "Non standard - supplied by customer" in the vital product data.	31P1658	AH1T
16 Gbps Long-wave SFP transceiver (pack of 2). **Important:** It is possible that SFP transceivers other than those shipped with the product are in use on the FC host bus adapter. It is the customer responsibility to obtain replacement parts for the SFP transceiver. The FRU part number is shown as “Non standard - supplied by customer” in the vital product data.	00RY191	ACHU

SAN Volume Controller 2145-92F expansion enclosure parts

On the 2145-92F expansion enclosure, all replaceable parts are field-replaceable units (FRUs). FRUs are replaced by your IBM service support representatives (SSRs). The expansion enclosure does not have any customer replaceable parts (CRUs).

Note: All of the information that is listed in the following tables for the 2145-92F expansion enclosure is also applicable to the 2147-92F expansion enclosure.

Expansion enclosure drives

Table 28 on page 50 summarizes the types of SAS drives that are supported by the 2145-92F expansion enclosure on SAN Volume Controller 2145-DH8 and SAN Volume Controller 2145-SV1 systems.

Table 28. Supported expansion enclosure SAS drives

Description	FRU part number	Feature code
600 GB 15 K disk drive	01LJ061	AH70
900 GB 15 K disk drive	01LJ827	AH71
1.2 TB 10 K disk drive	01LJ062	AH73
1.8 TB 10 K disk drive	01LJ063	AH74
2.4 TB 10 K disk drive	01YM178	AH75
6 TB 7.2 K Near-Line SAS disk drive	01LJ064	AH77
8 TB 7.2 K Near-Line SAS disk drive	01LJ065	AH78
10 TB 7.2 K Near-Line SAS disk drive	01LJ066	AH79
12 TB 7.2 K Near-Line SAS disk drive	01YM179	AH7A
1.6 TB tier 0 flash drive	01LJ073	AH7D
3.2 TB tier 0 flash drive	01LJ074	AH7E
1.92 TB tier 1 flash drive	01LJ075	AH7J
3.84 TB tier 1 flash drive	01LJ076	AH7K
7.68 TB tier 1 flash drive	01LJ077	AH7L
15.36 TB tier 1 flash drive	01LJ078	AH7M

Other expansion enclosure parts

Table 29 summarizes the part numbers and feature codes for other parts. The values are the same for all SAN Volume Controller systems that support the 2145-92F expansion enclosure.

Table 29. Other expansion enclosure parts

Description	FRU part number	Feature code	Comments
3 m 12 Gb SAS Cable (mSAS HD)	00AR317	ACUC	
6 m 12 Gb SAS Cable (mSAS HD)	00AR439	ACUD	
16A power cord C19 / C20 2 m	39M5388	AHP5	
Enclosure	01LJ607 **Note:** Replaces enclosure FRU P/N 01LJ112.		Includes the drive board, signal interconnect board, and internal power cables, in an otherwise empty enclosure.
Rail kit	01LJ114		
Front fascia (4U front cover)	01LJ116		
Display panel assembly	01LJ118		
PSU fascia (1U cover)	01LJ120		The fascia must be removed to access the power supply units.
Power supply unit (PSU)	01LJ122		The expansion enclosure contains 2 PSUs. Each PSU requires a C19 / C20 power cord.

Table 29. Other expansion enclosure parts (continued)

Description	FRU part number	Feature code	Comments
Secondary expansion module	01LJ124 (for use with enclosure FRU P/N 01LJ112) 01LJ860 (for use with enclosure FRU P/N 01LJ607		The expansion enclosure supports 2 secondary expansion modules. **CAUTION:** **Use caution when you are removing or replacing a secondary expansion module from an enclosure with FRU part number 01LJ112. Avoid contact with the connectors on the main board.**
Fan module	01LJ126		The expansion enclosure contains 4 fan modules.
Expansion canister	01LJ128		
Cable management arms (CMA)	01LJ130		The part contains the upper and lower CMA.
Top cover	01LJ132		
Fan interface board	01LJ134		

SAN Volume Controller 2145-12F expansion enclosure parts

The only replaceable SAN Volume Controller parts are the field-replaceable units (FRUs) which are replaced by your IBM service support representatives (SSRs). There are no customer replaceable parts (CRUs).

For information about the terms of the warranty and getting service and assistance, refer to your product warranty and support information.

Table 30. Expansion enclosure field replaceable units

Part Number	Part name	Notes
01AC555	Expansion enclosure drive bay with midplane assembly, 12-slot, 3.5-inch	Excludes drives, drive blanks, canisters, bezel covers, PSUs.
01AC579	Expansion Canister	N/A
01AC404	Expansion enclosure power supply unit	N/A
42R7992	Drive blank, 3.5-inch form factor	N/A
00Y2450	Expansion enclosure left bezel	No MTM/Serial number label on the FRU.
00Y2436	Enclosure right bezel, 3.5-inch form factor	N/A
00RY309	Expansion enclosure rail kit	N/A

Table 31. Drive field replaceable units

Part Number	Part name	Notes
00AR322	4 TB Near Line SAS hard disk drive	N/A
00RX911	6 TB NearLine SAS hard disk drive	N/A
00WK782	8 TB NearLine SAS hard disk drive	N/A

Table 31. Drive field replaceable units (continued)

Part Number	Part name	Notes
01EJ990	10 TB NearLine SAS disk drive	N/A
01YM177	12 TB NearLine SAS disk drive	N/A

Table 32. Cable field replaceable SAS units

Part Number	Part name	Notes
00AR311	1.5 m 12 Gbps SAS Cable (mini SAS HD to mini SAS HD)	For connecting expansion enclosures to nodes
00AR317	3.0 m 12 Gbps SAS Cable (mini SAS HD to mini SAS HD)	For connecting expansion enclosures to nodes
00AR439	6.0 m 12 Gbps SAS Cable (mini SAS HD to mini SAS HD)	For connecting expansion enclosures to nodes

Table 33. Cable field replaceable power units

Part Number	Part name	Notes
39M5068	Argentina 2.8 m	N/A
39M5199	Japan 2.8 m	N/A
39M5123	Europe 2.8 m	N/A
39M5165	Italy 2.8m	N/A
39M5102	Australia / New Zealand 2.8 m	N/A
39M5130	Denmark 2.8 m	N/A
39M5144	South Africa 2.8 m	N/A
39M5151	United Kingdom 2.8 m	N/A
39M5158	Switzerland 2.8 m	N/A
39M5172	Israel 2.8m	N/A
39M5206	China 2.8m	N/A
39M5219	Korea 2.8m	N/A
39M5226	India 2.8m	N/A
39M5240	Brazil 2.8m	N/A
39M5247	Taiwan 2.8m	N/A
39M5081	United States / Canada 2.8m	N/A
39M5377	Power jumper cord 2.8 m	N/A

SAN Volume Controller 2145-24F expansion enclosure parts

The only replaceable SAN Volume Controller parts are the field-replaceable units (FRUs) which are replaced by your service support representatives (SSRs). There are no customer replaceable parts (CRUs).

For information about the terms of the warranty and getting service and assistance, refer to your product warranty and support information.

Table 34. Expansion enclosure field replaceable units

Part Number	Part name	Notes
64P8445	Expansion enclosure midplane assembly, 24-slot, 2.5-inch	Excludes drives, drive blanks, canisters, bezel covers, and PSUs.
01AC579	Expansion Canister	N/A
01AC381	Expansion enclosure power supply unit	N/A
45W8680	Drive blank, 2.5-inch form factor	N/A
06Y2450	Expansion enclosure left bezel	No MTM/Serial number label on the FRU.
00Y2512	Enclosure right bezel, 2.5-inch form factor	N/A
00RY309	Expansion enclosure rail kit	N/A

Table 35. Small-form factor SAS drives field replaceable units

Part Number	Part name
31P1818	200 GB tier 0 flash drive
31P1819	400 GB tier 0 flash drive
31P1820	800 GB tier 0 flash drive
00RX914	1.6 TB tier 0 flash drive
01EJ983	3.2 TB tier 0 flash drive
00AR324	15 K RPM, 300 GB disk drive
00AR323	15 K RPM, 600 GB disk drive
00AR326	10 K RPM, 900 GB disk drive
00AR327	10 K RPM, 1.2 TB disk drive
00RX908	10 K RPM, 1.8 TB disk drive
00WK780	7.2 K RPM, 2 TB near line SAS drive
01EJ601	1.92 TB tier 1 flash drive
01EJ602	3.84 TB tier 1 flash drive
01EJ991	7.68 TB tier 1 flash drive
01EJ992	15.36 TB tier 1 flash drive

Table 36. Cable field replaceable units

Part Number	Part name	Notes
SAS		
00AR311	1.5 m 12 Gbps SAS Cable (mini SAS HD to mini SAS HD)	For connecting expansion enclosures to nodes
00AR317	3.0 m 12 Gbps SAS Cable (mini SAS HD to mini SAS HD)	For connecting expansion enclosures to nodes
00AR439	6.0 m 12 Gbps SAS Cable (mini SAS HD to mini SAS HD)	For connecting expansion enclosures to nodes
Power		
39M5068	Argentina 2.8 m	N/A
39M5081	United States / Canada 2.8 m	N/A

Table 36. Cable field replaceable units (continued)

Part Number	Part name	Notes
39M5102	Australia / New Zealand 2.8 m	N/A
39M5123	Europe 2.8 m	N/A
39M5130	Denmark 2.8 m	N/A
39M5144	South Africa 2.8 m	N/A
39M5151	United Kingdom 2.8 m	N/A
39M5158	Switzerland 2.8 m	N/A
39M5165	Italy 2.8 m	N/A
39M5172	Israel 2.8 m	N/A
39M5199	Japan 2.8 m	N/A
39M5206	China 2.8 m	N/A
39M5219	Korea 2.8 m	N/A
39M5226	India 2.8 m	N/A
39M5240	Brazil 2.8 m	N/A
39M5247	Taiwan 2.8 m	N/A
39M5377	Power jumper cord 2.8 m	N/A

Chapter 3. User interfaces for servicing your system

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

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

  Note: The front panel display is replaced by a technician port on some models.

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 four 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 alerts
Displays only the alerts 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.

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.

Deleting a node from a clustered system by using the management GUI

Remove a node from a system if the node fails and is being replaced with a new node, or if a repair causes that node to be unrecognizable by the system.

Before you begin

The cache on the selected node is flushed before the node is taken offline. In some circumstances, such as when the system is already degraded (for example, when both nodes in the I/O group are online and the volumes within the I/O group are degraded), the system ensures that data loss does not occur as a result of deleting the only node with the cache data. If a failure occurs on the other node in the I/O group, the cache is flushed before the node is removed to prevent data loss.

Before you delete a node from the system, record the node serial number , worldwide node name (WWNN), all worldwide port names (WWPNs), and the I/O group that the node is part of. If the node is added to the system later, recording this node information now can avoid data corruption.

Attention:

- If you are removing a single node and the remaining node in the I/O group is online, the data on the remaining node goes into write-through mode. This data can be exposed to a single point of failure if the remaining node fails.
- If the volumes are already degraded before you remove a node, redundancy to the volumes is degraded. Removing a node might result in a loss of access to data and data loss.
- Removing the last node in the system destroys the system. Before you remove the last node in the system, ensure that you want to destroy the system.
- When you remove a node, you remove all redundancy from the I/O group. As a result, new or existing failures can cause I/O errors on the hosts. The following failures can occur:
  - Host configuration errors
  - Zoning errors
  - Multipathing-software configuration errors
- If you are deleting the last node in an I/O group and there are volumes that are assigned to the I/O group, you cannot remove the node from the system if the node is online. You must back up or migrate all data that you want to save before you remove the node. If the node is offline, you can remove the node.
- When you remove the configuration node, the configuration function moves to a different node within the system. This process can take a short time, typically less than a minute. The management GUI reattaches to the new configuration node transparently.
- If you turn on the power to the node that is removed and it is still connected to the same fabric or zone, it attempts to rejoin the system. The system tells the node to remove itself from the system and the node becomes a candidate for addition to this system or another system.
- If you are adding this node into the system, ensure that you add it to the same I/O group that it was previously a member of. Failure to do so can result in data corruption.

This task assumes that you access the management GUI.

About this task

Complete the following steps to remove a node from a system:

Procedure

1. Select **Monitoring** > **System**.
2. Right-click the node that you want to remove and select **Remove**.

   If the node that you want to remove is shown as `Offline`, then the node is not participating in the system.

   If the node that you want to remove is shown as `Online`, deleting the node can result in the dependent volumes to also go offline. Verify whether the node has any dependent volumes.
3. To check for dependent volumes before you attempt to remove the node, right-click the node and select **Show Dependent Volumes**.

   If any volumes are listed, determine why and if access to the volumes is required while the node is removed from the system. If the volumes are assigned from storage pools that contain flash drives that are located in the node, check why the volume mirror, if it is configured, is not synchronized. There can also be dependent volumes because the partner node in the I/O

group is offline. Fabric issues can also prevent the volume from communicating with the storage systems. Resolve these problems before you continue with the node removal.

4. Click **Remove**.
5. Click **Yes** to remove the node. Before a node is removed, the system checks to determine whether there are any volumes that depend on that node.

   If the node that you selected contains volumes within the following situations, the volumes go offline and become unavailable if the node is removed:

   - The node contains flash drives and also contains the only synchronized copy of a mirrored volume.
   - The other node in the I/O group is offline.

   If you select a node to remove that has these dependencies, another panel displays confirming the removal.

Adding a node to a system

You can add a node to the system by using the CLI or management GUI. A node can be added to the system if the node previously failed and is being replaced with a new node or if a repair action causes the node to be unrecognizable by the system. When you add nodes, ensure that they are added in pairs to create a full I/O group. Adding a node to the system typically increases the capacity of the entire system. Adding spare nodes to a system does not increase the capacity of the system.

You can use either the management GUI or the command-line interface to add a node to the system. Some models might require you to use the front panel to verify that the new node was added correctly.

Before you add a node to a system, you must make sure that the switch zoning is configured such that the node that is being added is in the same zone as all other nodes in the system. If you are replacing a node and the switch is zoned by worldwide port name (WWPN) rather than by switch port, make sure that the switch is configured such that the node that is being added is in the same VSAN or zone.

Note: It is recommended that you use a consistent method (either only the management GUI, or only the CLI) when you add, remove, and re-add nodes. If a node is added by using the CLI and later re-added by using the GUI, it might get a different node name than it originally had.

Rules and restrictions for adding a node to a system

If you are using hot-spare nodes, the following considerations might not all be applicable. For more information, see the topic on adding a hot-spare node and the **swapnode** command.

If you are adding a node that was used previously, either within a different I/O group within this system or within a different system, if you add a node without changing its worldwide node name (WWNN), hosts might detect the node and use it as if it were in its old location. This action might cause the hosts to access the wrong volumes.

- You must ensure that the model type of the new node is supported by the software level that is installed on the system. If the model type is not supported by the software level, update the system to a software level that supports the model type of the new node.
- Each node in an I/O group must be connected to a different uninterruptible power supply.
- If you are adding a node back to the same I/O group after a service action required it to be deleted from the system, and if the physical node did not change, then no special procedures are required to add it back to the system.
- If you are replacing a node in a system either because of a node failure or an update, you must change the WWNN of the new node to match that of the original node before you connect the node to the Fibre Channel network and add the node to the system.
- If you are adding a node to the network again, to avoid data corruption, ensure that you are adding the node to the same I/O group from which it was removed. You must use the information that was recorded when the node was originally added to the system. If you do not have access to this information, contact the support center for assistance with adding the node back into the system so that data is not corrupted.
- For each external storage system, the LUNs that are presented to the ports on the new node must be the same as the LUNs that are presented to the nodes that currently exist in the system. You must ensure that the LUNs are the same before you add the new node to the system.
- If you create an I/O group in the system and add a node, no special procedures are needed because this node was never added to a system.
- If you create an I/O group in the system and add a node that was added to a system before, the host system might still be configured to the node WWPNs and the node might still be zoned in the fabric. Because you cannot change the WWNN for the node, you must ensure that other components in your fabric are configured correctly. Verify that any host that was previously configured to use the node was correctly updated.
- If the node that you are adding was previously replaced, either for a node repair or update, you might use the WWNN of that node for the replacement node. Ensure that the WWNN of this node was updated so that you do not have two nodes with the same WWNN attached to your fabric. Also, ensure that the WWNN of the node that you are adding is not 00000. If it is 00000, contact your support representative.
- The new node must be running a software level that supports encryption.
- If you are adding the new node to a system with either a HyperSwap or stretched system topology, you must assign the node to a specific site.

Rules and restrictions for using multipathing device drivers

- Applications on the host systems direct I/O operations to file systems or logical volumes that are mapped by the operating system to virtual paths (*vpaths*), which are pseudo disk objects that are supported by the multipathing device drivers. Multipathing device drivers maintain an association between a vpath and a volume. This association uses an identifier (UID) which is unique to the volume and is never reused. The UID allows multipathing device drivers to directly associate vpaths with volumes.
- Multipathing device drivers operate within a protocol stack that contains disk and Fibre Channel device drivers that are used to communicate with the system by using the SCSI protocol over Fibre Channel as defined by the ANSI FCS standard. The addressing scheme that is provided by these SCSI and Fibre

Channel device drivers uses a combination of a SCSI logical unit number (LUN) and the worldwide node name (WWNN) for the Fibre Channel node and ports.

- If an error occurs, the error recovery procedures (ERPs) operate at various tiers in the protocol stack. Some of these ERPs cause I/O to be redriven by using the same WWNN and LUN numbers that were previously used.
- Multipathing device drivers do not check the association of the volume with the vpath on every I/O operation that it performs.

You can use either the **addnode** command or the **Add Node** wizard in the management GUI. To access the **Add Node** wizard, select **Monitoring** > **System**. On the image, click the new node to start the wizard. Complete the wizard and verify the new node. If the new node is not displayed in the image, it indicates a potential cabling issue. Check the installation information to ensure that your node was cabled correctly.

To add a node to a system by using the command-line interface, complete these steps:

1. Enter this command to verify that the node is detected on the network:

   ```
   svcinfo lsnodecandidate
   ```

   This example shows the output for this command:

   ```
   # svcinfo lsnodecandidate
   id               panel_name UPS_serial_number UPS_unique_id    hardware serial_number product_mtm machine_signature
   500507680C007B00 KD0N8AM                      500507680C007B00  DH8     KD0N8AM       2145-DH8    0123-4567-89AB-CDEF
   ```

   The **id** parameter displays the WWNN for the node. If the node is not detected, verify cabling to the node.
2. Enter this command to determine the I/O group where the node must be added:

   ```
   lsiogrp
   ```
3. Record the name or ID of the first I/O group that has a node count of zero. You need the name or ID for the next step. Note: You must do this step for the first node that is added. You do not do this step for the second node of the pair because it uses the same I/O group number.
4. Enter this command to add the node to the system:

   ```
   addnode -wwnodename WWNN -iogrp iogrp_name -name new_name_arg -site site_name
   ```

   Where **WWNN** is the WWNN of the node, **iogrp_name** is the name of the I/O group that you want to add the node to and `new_name_arg` is the name that you want to assign to the node. If you do not specify a new node name, a default name is assigned. Typically, you specify a meaningful node name. The `site_name` specifies the name of the site location of the new node. This parameter is only required if the topology is a HyperSwap or stretched system.

   Note: Adding the node might take a considerable amount of time.
5. Record this information for future reference:
   - Serial number.
   - Worldwide node name.
   - All of the worldwide port names.
   - The name or ID of the I/O group

Service assistant interface

The service assistant interface is a browser-based GUI that is used to service your nodes.

When to use the service assistant

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

Attention: Complete service actions on nodes 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 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 system to run the recommended actions
- When the recommended action directs you to use the service assistant.

The management GUI operates only when there is an online clustered system. Use the service assistant if you are unable to create a clustered system.

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 software package from the support site or rescue the software from another node.
- Update software on nodes manually versus completing a standard update procedure.
- Change the service IP address that is assigned to Ethernet port 1 for the current node.
- Install a temporary SSH key if a key is not installed and CLI access is required.
- Restart the services used by the system.

Accessing the service assistant

The service assistant is a web application that helps troubleshoot and resolve problems on a node. The service assistant can be accessed through a service IP address. On SAN Volume Controller 2145-DH8, you can connect to the service assistant 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 that you want to work on.
2. Log on to the service assistant using 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.

Results

Complete the service assistant actions on the correct node.

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 SAN Volume Controller 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 to initialize and use a CLI session.

Service command-line interface

Use the service command-line interface (CLI) to manage a node 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 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, review in the Command-line interface topic of this product information.

USB flash drive interface

Use a USB flash drive to help service a node.

When a USB flash drive is inserted into one of the USB ports on a node, 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.

When to use the USB flash drive

The USB flash drive can be used for service 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 node canister 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 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.

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

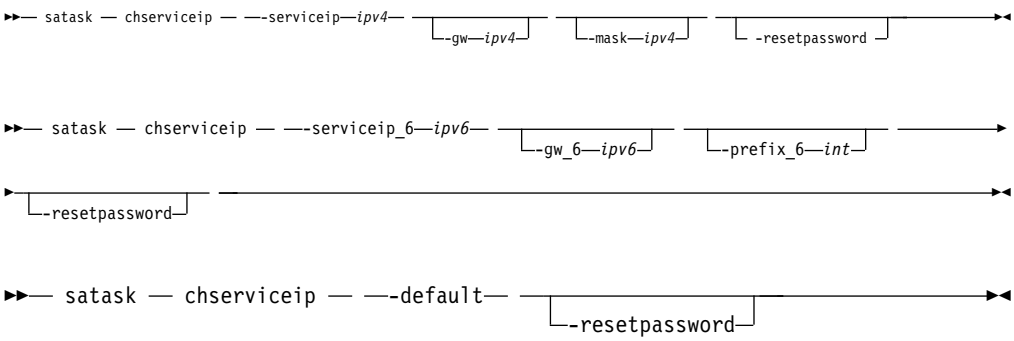

Parameters

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

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

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

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

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

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

-resetpassword
Sets the service assistant password to the default value.

Description

This command resets the service assistant IP address to the default value. 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 log on to the system because you forget the superuser password, and you wish to reset it.

Syntax

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.

satask snap:

Use the **satask snap** command to collect diagnostic information from the node and to write the output to a USB flash drive, or to upload specified support information.

Syntax

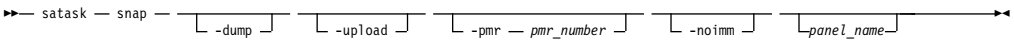

Parameters

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

-upload
(Optional) Specifies that the snap file be uploaded after it is generated.

-pmr *pmr_number*
(Optional) Specifies the PMR number to use to upload the snap file. The format for a PMR must be a 13-character alphanumeric string. If the specified PMR is invalid or unknown, it is uploaded to a generic location on the server with the prefix:

`unknown_pmr_`*`pmr_number`*`_`

If this option is not supplied, the snap file is uploaded using the machine type and serial number attributes.

-noimm
(Optional) Indicates the `/dumps/imm.ffdc` file must 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 and uploads support information.

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 5 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.

Specify the **lsdumps** command to view the file that you create.

An invocation example

```
satask snap
```

The resulting output:

```
No feedback
```

Important: The name of the output file (placed on the specified node) is `snap.single.`*`nodeid`*`.`*`date`*`.`*`time`*`.tgz`.

An invocation example

```
satask snap -noimm
```

The resulting output:

```
No feedback
```

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

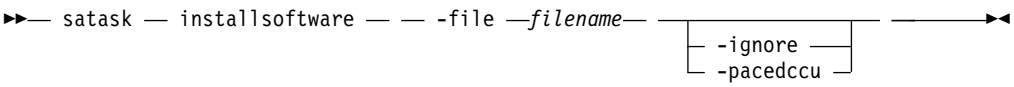

Parameters

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

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

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

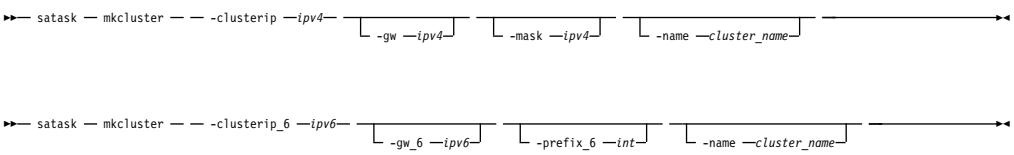

Parameters

-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.

Change system IP address:

Use this command to change the system IP address of the storage system.

It is best to use the initialization tool to create this command in `satask.txt` together with the associated `clitask.txt` file that changes the file modules management IP addresses.

Syntax

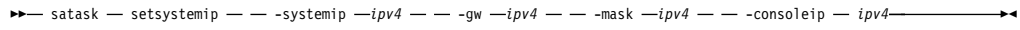

Parameters

-systemip
The IPv4 address for Ethernet port 1 on the system.

-gw
The IPv4 gateway for Ethernet port 1 on the system.

-mask
The IPv4 subnet for Ethernet port 1 on the system.

-consolip
The management IPv4 address of SAN Volume Controller system.

Description

This command is only supported in the satask.txt file on a USB flash drive.

It calls the **svctask chsystemip** command if the USB flash drive is inserted in the configuration node canister. Otherwise, it flashes the amber identify LED of the node canister that is the configuration node.

If the amber identify LED for a different node canister starts to flash, move the USB flash drive over to that node canister because it is the configuration node.

When the amber LED turns off, you can move the USB flash drive to one of the file modules so that it uses the clitask.txt file to change the file module management IP addresses.

Leave the USB flash drive in the file module for at least 2 minutes before you remove it. Use a workstation to check the `clitask_results.txt` and `satask.txt` results files on the USB flash drive.

If the IP address change was successful, then you must run the `startmgtsrv -r` command to restart the management service so that it does not continue to issue commands to the old system IP address of the volume storage system.

For example, on a Linux workstation with network access to the new management IP address:

```
satask setsystemip -systemip 123.123.123.20 -gw 123.123.123.1 -mask 255.255.255.0
 -consoleip 123.123.123.10
```

You can now access the management GUI, which you can use to change any other IP address that needs to be changed.

The following text is an example of what might be in the clitask.txt file:

```
chnwmgt --serviceip1 123.123.123.11 --serviceip2 123.123.123.12
--mgtip 123.123.123.10 --gateway 123.123.123.1 --netmask 255.255.255.0 --force
chstoragesystem --ip1 123.123.123.20
```

The following text is an example of what might be in the satask.txt file:

```
satask setsystemip -systemip 123.123.123.20 -gw 123.123.123.1 -mask 255.255.255.0
 -consoleip 123.123.123.10
```

Query status command:

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

Syntax

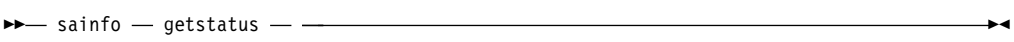

Parameters

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.

Technician port

The technician port is an Ethernet port on the back panel of 2145-SV1 and 2145-DH8 nodes that you use to configure the node.

You can use the technician port to do most of the system configuration operations that are provided by the front panel of earlier system models, which includes the following tasks:

- Defining a management IP address.
- Initializing a new system.
- Servicing the system.

To use the technician port, plug one end of an Ethernet cable into the technician port. Then, plug the other end into the Ethernet port of a personal computer with Dynamic Host Configuration Protocol (DHCP) configured and a web browser that is installed. Run the system configuration tool by going to address `http://install` with your browser. If you do not have DCHP, open a supported browser and go to the default static IP address `192.168.0.1` for the node.

Note: When your personal computer is configured with DHCP, the technician port uses DHCP to reconfigure network services on your personal computer. Software

on your personal computer that was using these services might experience network problems while it is connected to the technician port. For example, selecting a link in a web page that was loaded before you connect to the technician port might result in an error message.

2145-SV1 node

On the back of a 2145-SV1 node, the technician port is on the bottom right side of the node. Figure 31 shows the location of the technician port and other ports that are used to service the node.

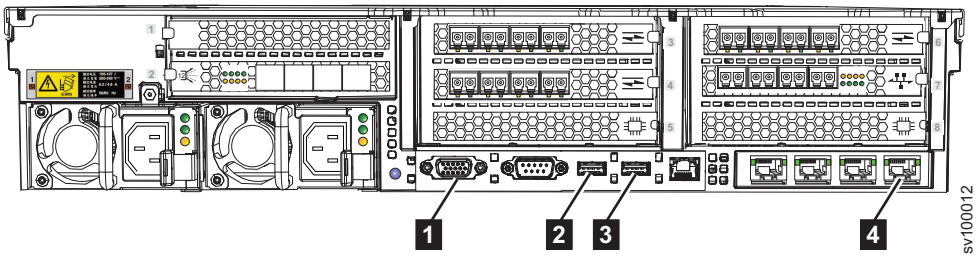

Figure 31. 2145-SV1 technician port

1. VGA port
2. Rear USB port 1
3. Rear USB port 2
4. Technician port (Ethernet)

2145-DH8 node

Starting from the left at the rear of the SAN Volume Controller 2145-DH8 node, the technician port is the fourth Ethernet port to the right. Figure 32 shows the rear of the SAN Volume Controller node, where **1** is the technician port.

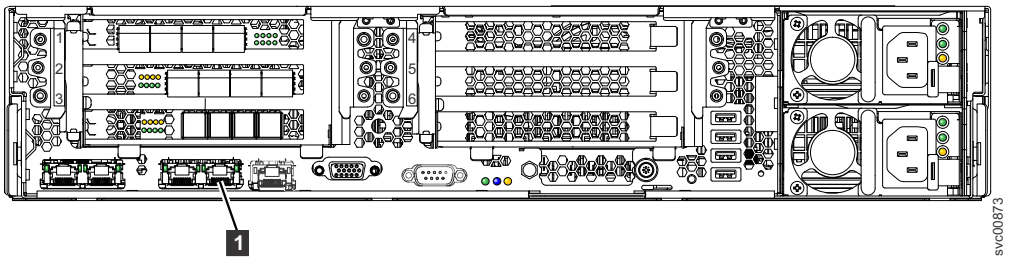

Figure 32. 2145-DH8 technician port

Chapter 4. Performing recovery actions using the SAN Volume Controller CLI

The SAN Volume Controller command-line interface (CLI) is a collection of commands that you can use to manage SAN Volume Controller clusters. See the Command-line interface documentation for the specific details about the commands provided here.

Validating and repairing mirrored volume copies by using the CLI

You can use the **repairvdiskcopy** command from the command-line interface (CLI) to validate and repair mirrored volume copies.

Attention: Run the **repairvdiskcopy** command only if all volume copies are synchronized.

When you issue the **repairvdiskcopy** command, you must use only one of the **-validate**, **-medium**, or **-resync** parameters. You must also specify the name or ID of the volume to be validated and repaired as the last entry on the command line. After you issue the command, no output is displayed.

-validate
Use this parameter only if you want to verify that the mirrored volume copies are identical. If any difference is found, the command stops and logs an error that includes the logical block address (LBA) and the length of the first difference. You can use this parameter, starting at a different LBA each time to count the number of differences on a volume.

-medium
Use this parameter to convert sectors on all volume copies that contain different contents into virtual medium errors. Upon completion, the command logs an event, which indicates the number of differences that were found, the number that were converted into medium errors, and the number that were not converted. Use this option if you are unsure what the correct data is, and you do not want an incorrect version of the data to be used.

-resync
Use this parameter to overwrite contents from the specified primary volume copy to the other volume copy. The command corrects any differing sectors by copying the sectors from the primary copy to the copies that are being compared. Upon completion, the command process logs an event, which indicates the number of differences that were corrected. Use this action if you are sure that either the primary volume copy data is correct or that your host applications can handle incorrect data.

-startlba *lba*
Optionally, use this parameter to specify the starting Logical Block Address (LBA) from which to start the validation and repair. If you previously used the **validate** parameter, an error was logged with the LBA where the first difference, if any, was found. Reissue **repairvdiskcopy** with that LBA to avoid reprocessing the initial sectors that compared identically. Continue to reissue **repairvdiskcopy** by using this parameter to list all the differences.

Issue the following command to validate and, if necessary, automatically repair mirrored copies of the specified volume:

```
repairvdiskcopy -resync -startlba 20 vdisk8
```

Notes:

1. Only one **repairvdiskcopy** command can run on a volume at a time.
2. After you start the **repairvdiskcopy** command, you cannot use the command to stop processing.
3. The primary copy of a mirrored volume cannot be changed while the **repairvdiskcopy -resync** command is running.
4. If there is only one mirrored copy, the command returns immediately with an error.
5. If a copy that is being compared goes offline, the command is halted with an error. The command is not automatically resumed when the copy is brought back online.
6. In the case where one copy is readable but the other copy has a medium error, the command process automatically attempts to fix the medium error by writing the read data from the other copy.
7. If no differing sectors are found during **repairvdiskcopy** processing, an informational error is logged at the end of the process.

Checking the progress of validation and repair of volume copies by using the CLI

Use the **lsrepairvdiskcopyprogress** command to display the progress of mirrored volume validation and repairs. You can specify a volume copy by using the **-copy** *id* parameter. To display the volume that has two or more copies with an active task, specify the command with no parameters; it is not possible to have only one volume copy with an active task.

To check the progress of validation and repair of mirrored volumes, issue the following command:

```
lsrepairvdiskcopyprogress –delim :
```

The following example shows how the command output is displayed:

```
vdisk_id:vdisk_name:copy id:task:progress:estimated_completion_time
0:vdisk0:0:medium:50:070301120000
0:vdisk0:1:medium:50:070301120000
```

Repairing a thin-provisioned volume using the CLI

You can use the **repairsevdiskcopy** command from the command-line interface to repair the metadata on a thin-provisioned volume.

The **repairsevdiskcopy** command automatically detects and repairs corrupted metadata. The command holds the volume offline during the repair, but does not prevent the disk from being moved between I/O groups.

If a repair operation completes successfully and the volume was previously offline because of corrupted metadata, the command brings the volume back online. The only limit on the number of concurrent repair operations is the number of volume copies in the configuration.

When you issue the **repairsevdiskcopy** command, you must specify the name or ID of the volume to be repaired as the last entry on the command line. Once started, a repair operation cannot be paused or canceled; the repair can be terminated only by deleting the copy.

Attention: Use this command only to repair a thin-provisioned volume that has reported corrupt metadata.

Issue the following command to repair the metadata on a thin-provisioned volume:

```
repairsevdiskcopy vdisk8
```

After you issue the command, no output is displayed.

Notes:

1. Because the volume is offline to the host, any I/O that is submitted to the volume while it is being repaired fails.
2. When the repair operation completes successfully, the corrupted metadata error is marked as fixed.
3. If the repair operation fails, the volume is held offline and an error is logged.

Checking the progress of the repair of a thin-provisioned volume by using the CLI

Issue the **lsrepairsevdiskcopyprogress** command to list the repair progress for thin-provisioned volume copies of the specified volume. If you do not specify a volume, the command lists the repair progress for all thin-provisioned copies in the system.

Note: Run this command only after you run the **repairsevdiskcopy** command, which you must run only as required by the fix procedures that are recommended by your support team.

Recovering offline volumes using the CLI

If a node or an I/O group fails, you can use the command-line interface (CLI) to recover offline volumes.

About this task

If you lose both nodes in an I/O group, you lose access to all volumes that are associated with the I/O group. To regain access to the volumes, you must perform one of the following procedures. Depending on the failure type, you might have lost data that was cached for these volumes and the volumes are now offline.

Data loss scenario 1

One node in an I/O group failed and failover started on the second node. During the failover process, the second node in the I/O group fails before the data in the write cache is flushed to the backend. The first node is successfully repaired but its hardened data is not the most recent version that is committed to the data store, therefore, it cannot be used. The second node is repaired or replaced and lost its hardened data, and the node has no way of recognizing that it is part of the system.

Complete the following steps to recover an offline volume when one node has down-level hardened data and the other node loses hardened data.

Procedure

1. Recover the node and add it back into the system.
2. Delete all IBM FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the offline volumes.
3. Run the **recovervdisk**, **recovervdiskbyiogrp** or **recovervdiskbysystem** command.
4. Re-create all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the volumes.

Example

Data loss scenario 2

Both nodes in the I/O group failed and have been repaired. Therefore, the nodes that lost their hardened data and have no way of recognizing that they are part of the system.

Complete the following steps to recover an offline volume when both nodes that have lost their hardened data and cannot be recognized by the system.

1. Delete all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the offline volumes.
2. Run the **recovervdisk**, **recovervdiskbyiogrp** or **recovervdiskbysystem** command.
3. Re-create all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the volumes.

Chapter 5. Viewing the vital product data

Vital product data (VPD) is information that uniquely records each element in the SAN Volume Controller . The data is updated automatically by the system when the configuration is changed.

The VPD lists the following types of information:

- System-related values such as the software version, space in storage pools, and space allocated to volumes.
- Node-related values that include the specific hardware that is installed in each node. Examples include the FRU part number for the system board and the level of BIOS firmware that is installed. The node VPD is held by the system that makes it possible to get most of the VPD for the nodes that are powered off.

Using different sets of commands, you can view the system VPD and the node VPD. You can also view the VPD through the management GUI.

Downloading the vital product data using the management GUI

You can download the vital product data for a node from the management GUI.

Procedure

1. In the management GUI, select **Monitoring** > **System**.
2. From the dynamic graphic of the system, select the node and click the icon to the right of the Actions menu to download VPD information.

Displaying the vital product data using the CLI

You can use the command-line interface (CLI) to display the system or node vital product data (VPD).

Issue the following CLI commands to display the VPD:

```
sainfo lsservicestatus
lsnodehw
lsnodevpd nodename
lssystem system_name
lssystemip
lsdrive
```

Displaying node properties by using the CLI

You can use the command-line interface (CLI) to display node properties.

About this task

To display the node properties:

Procedure

1. Use the **lsnode** CLI command to display a concise list of nodes in the clustered system.

   Issue this CLI command to list the system nodes:

   ```
   lsnode -delim :
   ```

2. Issue the **lsnode** CLI command and specify the node ID or name of the node that you want to receive detailed output.

   The following example is a CLI command that you can use to list detailed output for a node in the system:

   ```
   lsnode -delim : group1node1
   ```

   Where *group1node1* is the name of the node for which you want to view detailed output.

Displaying clustered system properties by using the CLI

You can use the command-line interface (CLI) to display the properties for a clustered system (system).

About this task

These actions help you display your system property information.

Procedure

Issue the **lssystem** command to display the properties for a system.
The following command is an example of the **lssystem** command you can issue:

```
lssystem -delim : build1
```

where *build1* is the name of the system.

Results

```
id:000002007A00A0FE
name:build1
location:local
partnership:
bandwidth:
total_mdisk_capacity:90.7GB
space_in_mdisk_grps:90.7GB
space_allocated_to_vdisks:14.99GB
total_free_space:75.7GB
statistics_status:on
statistics_frequency:15
required_memory:0
cluster_locale:en_US
time_zone:522 UTC
code_level:6.1.0.0 (build 47.3.1009031000)
FC_port_speed:2Gb
console_IP:9.71.46.186:443
id_alias:000002007A00A0FE
gm_link_tolerance:300
gm_inter_cluster_delay_simulation:0
gm_intra_cluster_delay_simulation:0
email_reply:
email_contact:
email_contact_primary:
email_contact_alternate:
email_contact_location:
email_state:stopped
inventory_mail_interval:0
total_vdiskcopy_capacity:15.71GB
total_used_capacity:13.78GB
total_overallocation:17
total_vdisk_capacity:11.72GB
cluster_ntp_IP_address:
cluster_isns_IP_address:
iscsi_auth_method:none
iscsi_chap_secret:
auth_service_configured:no
auth_service_enabled:no
auth_service_url:
auth_service_user_name:
auth_service_pwd_set:no
auth_service_cert_set:no
relationship_bandwidth_limit:25
gm_max_host_delay:5
tier:generic_ssd
tier_capacity:0.00MB
tier_free_capacity:0.00MB
tier:generic_hdd
tier_capacity:90.67GB
tier_free_capacity:75.34GB
email_contact2:
email_contact2_primary:
email_contact2_alternate:
total_allocated_extent_capacity:16.12GB
```

Fields for the node VPD

The node vital product data (VPD) provides information for items such as the system board, batteries, processor, fans, memory module, adapter, devices, software, front panel assembly, serial-attached SCSI (SAS) flash drive and SAS host bus adapter (HBA).

Table 37 on page 80 shows the fields that you see for the system board.

Table 37. Fields for the system board

Item	Field name
System board	Part number
	System serial number
	Number of processors
	Number of memory slots
	Number of fans
	Number of Fibre Channel adapters
	Number of SCSI, IDE, SATA, or SAS devices
	Number of compression accelerator adapters
	Number of power supplies
	Number of high-speed SAS adapters
	BIOS manufacturer
	BIOS version
	BIOS release date
	System manufacturer
	System product
	Planar manufacturer
	Power supply part number
	CMOS battery part number
	Power cable assembly part number
	Service processor firmware
	SAS controller part number

Table 38 shows the fields that you see for the batteries.

Table 38. Fields for the batteries

Item	Field name
Batteries	Battery_FRU_part
	Battery_part_identity
	Battery_fault_led
	Battery_charging_status
	Battery_cycle_count
	Battery_power_on_hours
	Battery_last_recondition
	Battery_midplane_FRU_part
	Battery_midplane_part_identity
	Battery_midplane_FW_version
	Battery_power_cable_FRU_part
	Battery_power_sense_cable_FRU_part
	Battery_comms_cable_FRU_part
	Battery_EPOW_cable_FRU_part

Table 39 shows the fields that you see for each processor that is installed.

Table 39. Fields for the processors

Item	Field name
Processor	Part number
	Processor location
	Manufacturer
	Version
	Speed
	Status
	Processor serial number

Table 40 shows the fields that you see for each fan that is installed.

Table 40. Fields for the fans

Item	Field name
Fan	Part number
	Location

Table 41 shows the fields that are repeated for each installed memory module.

Table 41. Fields that are repeated for each installed memory module

Item	Field name
Memory module	Part number
	Device location
	Bank location
	Size (MB)
	Manufacturer (if available)
	Serial number (if available)

Table 42 shows the fields that are repeated for each installed adapter.

Table 42. Fields that are repeated for each adapter that is installed

Item	Field name
Adapter	Adapter type
	Part number
	Port numbers
	Location
	Device serial number
	Manufacturer
	Device
	Adapter revision
	Chip revision

Table 43 shows the fields that are repeated for each device that is installed.

Table 43. Fields that are repeated for each SCSI, IDE, SATA, and SAS device that is installed

Item	Field name
Device	Part number
	Bus
	Device
	Model
	Revision
	Serial number
	Approximate capacity
	Hardware revision
	Manufacturer

Table 44 shows the fields that are specific to the node software.

Table 44. Fields that are specific to the node software

Item	Field name
Software	Code level
	Node name
	Worldwide node name
	ID
	Unique string that is used in dump file names for this node

Table 45 shows the fields that are provided for the front panel assembly.

Table 45. Fields that are provided for the front panel assembly

Item	Field name
Front panel	Part number
	Front panel ID
	Front panel locale

Table 46 shows the fields that are provided for the Ethernet port.

Table 46. Fields that are provided for the Ethernet port

Item	Field name
Ethernet port	Port number
	Ethernet port status
	MAC address
	Supported speeds

Table 47 on page 83 shows the fields that are provided for the power supplies in the node.

Table 47. Fields that are provided for the power supplies in the node

Item	Field name
Power supplies	Part number
	Location

Table 48 shows the fields that are provided for the SAS host bus adapter (HBA).

Table 48. Fields that are provided for the SAS host bus adapter (HBA)

Item	Field name
SAS HBA	Part number
	Port numbers
	Device serial number
	Manufacturer
	Device
	Adapter revision
	Chip revision

Table 49 shows the fields that are provided for the SAS flash drive.

Table 49. Fields that are provided for the SAS flash drive

Item	Field name
SAS SSD	Part number
	Manufacturer
	Device serial number
	Model
	Type
	UID
	Firmware
	Slot
	FPGA firmware
	Speed
	Capacity
	Expansion tray
	Connection type

Table 50 on page 84 shows the fields that are provided for the small form factor pluggable (SFP) transceiver.

Table 50. Fields that are provided for the small form factor pluggable (SFP) transceiver

Item	Field name
Small form factor pluggable (SFP) transceiver	Part number
	Manufacturer
	Device
	Serial number
	Supported speeds
	Connector type
	Transmitter type
	Wavelength
	Maximum distance by cable type
	Hardware revision
	Port number
	Worldwide port name

Fields for the system VPD

The system vital product data (VPD) provides various information about the system, including its ID, name, location, IP address, email contact, code level, and total free space.

Table 51 shows the fields that are provided for the system properties as shown by the management GUI.

Table 51. Fields that are provided for the system properties

Item	Field name
General	ID **Note:** This value is the unique identifier for the system.
	Name
	Location
	Time Zone
	Required Memory
	Licensed Code Version
	Channel Port Speed
IP Addresses[1]	Ethernet Port 1 (attributes for both IPv4 and IPv6) • IP Address • Service IP Address • Subnet Mask • Prefix • Default Gateway
	Ethernet Port 2 (attributes for both IPv4 and IPv6) • IP Address • Service IP Address • Subnet Mask • Prefix • Default Gateway

Table 51. Fields that are provided for the system properties (continued)

Item	Field name
Remote Authentication	Remote Authentication
	Web Address
	User Name
	Password
	SSL Certificate
Space	Total MDisk Capacity
	Space in Storage Pools
	Space Allocated to Volumes
	Total Free Space
	Total Used Capacity
	Total Allocation
	Total Volume Copy Capacity
	Total Volume Capacity
Statistics	Statistics Status
	Statistics Frequency
Metro and Global Mirror	Link Tolerance
	Intersystem Delay Simulation
	Intrasystem Delay Simulation
	Partnership
	Bandwidth
Email	SMTP Email Server
	Email Server Port
	Reply Email Address
	Contact Person Name
	Primary Contact Phone Number
	Alternate Contact Phone Number
	Physical Location of the System Reporting Error
	Email Status
	Inventory Email Interval
iSCSI	iSNS Server Address
	Supported Authentication Methods
	CHAP Secret
[1] You can also use the **lssystemip** CLI command to view this data.	

Chapter 6. Diagnosing problems

You can diagnose problems with the control and indicators, the command-line interface (CLI), the management GUI, or the Service Assistant GUI. The diagnostic LEDs on the SAN Volume Controller nodes and uninterruptible power supply units also help you diagnose hardware problems.

Event logs

By understanding the event log, you can do the following tasks:

- Manage the event log
- View the event log
- Describe the fields in the event log

Error codes

The following topics provide information to help you understand and process the error codes:

- Event reporting
- Understanding the events
- Understanding the error codes
- Determining a hardware boot failure

If the node is showing a boot message, failure message, or node error message, and you determined that the problem was caused by a software or firmware failure, you can restart the node to see whether that might resolve the problem. Perform the following steps to properly shut down and restart the node:

1. Follow the instructions in "MAP 5350: Powering off a node" on page 284.
2. Restart only one node at a time.
3. Do not shut down the second node in an I/O group for at least 30 minutes after you shut down and restart the first node.

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 52 describes the statistics collection for MDisks, for individual nodes.

Table 52. 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 started 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).
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 that are 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 that is taken for write commands to complete from the time they join the queue.

Table 53 on page 89 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 53. Statistic collection for volumes for individual nodes

Statistic name	
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 volumes read operations that are processed (since the node started running).
ub	Indicates the cumulative number of blocks of data unmapped (since the node is running).
ul	Indicates the cumulative unmap response time in milliseconds for each volume. The cumulative response time for volume unmaps is calculated by starting a timer when a SCSI unmap command is received and stopped when the command completes successfully. The elapsed time is added to the cumulative counter.
ulw	Indicates the worst unmap response time in milliseconds for each volume. The worst response time for volume unmaps is calculated by starting a timer when a SCSI unmap command is received and stopped when the command completes successfully.
uo	Indicates the cumulative number of volume unmap operations that were processed (since the node started running).
uou	Indicates the cumulative number of volume unmap operations that are not aligned on an 8 K boundary (according to the alignment/granularity setting in Block Limits VPD Page (0xb0).
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.
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 volumes write operations that are processed (since the node is running).
wou	Indicates the cumulative number of volumes write operations that are not aligned on a 4 K 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 SAN Volume Controller .

Note: For unmap statistics, it is where an unmap operation is a **SCSI unmap** or **Write same with unmap** command.

Table 54 describes the VDisk information that is related to Metro Mirror or Global Mirror relationships that is reported for individual nodes.

Table 54. 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 at system version 4.3.1, this statistic 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 system version earlier than 4.3.1, this value does not increment.
gws	Indicates the total number of write requests that are issued to the secondary site.

Table 55 describes the port information that is reported for individual nodes.

Table 55. Statistic collection for node ports

Statistic name	Description
bbcz	Indicates the total time in microseconds for which the buffer credit counter was at zero. That this statistic is only reported by 8 Gbps Fibre Channel ports. For other port types, this statistic is 0.
cbr	Indicates the bytes received from controllers.
cbt	Indicates the bytes transmitted to disk controllers.
cer	Indicates the commands that are received from disk controllers. **Note:** The cer metric is always 0.
cet	Indicates the commands that are initiated to disk controllers.
dtdc	Indicates the number of transfers that experienced excessive data transmission delay.
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 that are received from hosts.
het	Indicates the commands that are initiated to hosts. **Note:** The het metric is always 0.
icrc	Indicates the number of CRC that is not valid.

Table 55. Statistic collection for node ports (continued)

Statistic name	Description
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 that are received from other nodes in the same cluster.
lnet	Indicates the commands that are 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 that are received from other nodes in the other clusters.
rmet	Indicates the commands that are initiated to other nodes in the other clusters.
wwpn	Indicates the worldwide port name for the node.

Table 56 describes the node information that is reported for each node.

Table 56. 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 spends polling, waiting for work versus 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 spends polling, waiting for work versus doing work. This statistic accumulates from zero. This statistic is the same information as the information provided with the `cpu busy` statistic and eventually replaces the `cpu busy` statistic.
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.

Table 56. Statistic collection for nodes (continued)

Statistic name	Description
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 that is taken for the remote node to respond.

Table 57 describes the statistics collection for volumes.

Table 57. Cache statistics collection for volumes and volume copies

Statistic	Acronym	Cache statistics for volumes	Cache statistics for volume copies	Cache partition statistics for volumes	Cache partition statistics for volume copies	Overall node cache statistics	Cache statistics for mdisks	Units and state	Cache statistics for data reduction pools
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	
prefetches	p		Yes					sectors, cumulative	
prefetch hits (prefetch data that is read)	ph		Yes					sectors, cumulative	

Table 57. Cache statistics collection for volumes and volume copies (continued)

Statistic	Acronym	Cache statistics for volumes	Cache statistics for volume copies	Cache partition statistics for volumes	Cache partition statistics for volume copies	Overall node cache statistics	Cache statistics for mdisks	Units and state	Cache statistics for data reduction pools
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	Yes
fullness Max	fmx			Yes	Yes			%, non-cumulative	Yes
fullness Min	fmn			Yes	Yes			%, non-cumulative	Yes
Destage Target Average	dtav				Yes		Yes	IOs capped 9999, non-cumulative	Yes
Destage Target Max	dtmx				Yes			IOs, non-cumulative	Yes
Destage Target Min	dtmn				Yes			IOs, non-cumulative	Yes
Destage In Flight Average	dfav				Yes		Yes	IOs capped 9999, non-cumulative	Yes
Destage In Flight Max	dfmx				Yes			IOs, non-cumulative	Yes
Destage In Flight Min	dfmn				Yes			IOs, non-cumulative	Yes
destage latency average	dav	Yes	Yes	Yes	Yes	Yes	Yes	µs capped 9999999, non-cumulative	Yes
destage latency max	dmx			Yes	Yes	Yes		µs capped 9999999, non-cumulative	Yes

Table 57. Cache statistics collection for volumes and volume copies (continued)

Statistic	Acronym	Cache statistics for volumes	Cache statistics for volume copies	Cache partition statistics for volumes	Cache partition statistics for volume copies	Overall node cache statistics	Cache statistics for mdisks	Units and state	Cache statistics for data reduction pools
destage latency min	dmn			Yes	Yes	Yes		µs capped 9999999, non-cumulative	Yes
destage count	dcn	Yes	Yes	Yes	Yes	Yes		ios, non-cumulative	Yes
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	
Read Cache Fullness Max	rfmx					Yes		%, non-cumulative	

Table 57. Cache statistics collection for volumes and volume copies (continued)

Statistic	Acronym	Cache statistics for volumes	Cache statistics for volume copies	Cache partition statistics for volumes	Cache partition statistics for volume copies	Overall node cache statistics	Cache statistics for mdisks	Units and state	Cache statistics for data reduction pools
Read Cache Fullness Min	rfmn					Yes		%, non-cumulative	
Pinned Percent	pp	Yes	Yes	Yes	Yes	Yes		% of total cache snapshot, non-cumulative	Yes
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).

There are three types of data reduction properties per data reduction pool.

- dca - these statistics are related to the data stored within the data reduction pool.
- rca - these statistics are related to I/O to manage the background garbage collection processes of the data reduction pool.
- jca - these statistics are related to journaling operations for the metadata that manages the data reduction pool.

Table 58 describes the statistic collection for volume cache per individual nodes.

Table 58. 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, which are 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 that is 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.

Table 58. Statistic collection for volume cache per individual nodes (continued). This table describes the volume cache information that is reported for individual nodes.

Statistic name	Description
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 59 describes the garbage collection statistics for data reduction pools.

Table 59. Garbage collection statistics for data reduction pools

Statistic name	Description	State
cm	Consumed Mb (the number of MBs of host rewrites).	Cumulative
ext col	Extents collected (the number of extents that garbage collection has processed).	Cumulative
id	The internal repository identified to which the statistics reported refers.	
mdg	The mdisk group id for the data reduction pool repository.	
mm	Moved Mb (the number of MBs of data moved by garbage collection).	Cumulative

Table 59. Garbage collection statistics for data reduction pools (continued)

Statistic name	Description	State
nm	New Mb (the number of MBs of host writes to new addresses).	Cumulative
rec	Reclaimable capacity in the pool, for this node, current value, in MBs.	Cumulative
rm	Recovered Mb (the number of MBs of space recovered by garbage collection).	Cumulative

Table 60 describes the XML statistics specific to an IP Partnership port.

Table 60. 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) takes place.
iptx	Indicates the bytes transmitted to other nodes in other clusters by the IP partnership driver.

Table 61 describes the offload data transfer (ODX) Vdisk and node level I/O statistics.

Table 61. 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).

Table 61. ODX VDisk and node level statistics (continued)

Statistic name	Acronym	Description
Total transferred ODX I/O read blocks	oro	Cumulative total number of blocks that are 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 that are 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 that are 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 revocation and expiration.

Table 62 describes the statistics collection for cloud per cloud account ID.

Table 62. 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 that is taken to transfer the data to the cloud
Down Latency	dwlt	Total time that is taken to download the data from the cloud
Down Error Latency	dwerlt	Time that is taken for the GET errors

Table 62. Statistics collection for cloud per cloud account ID (continued)

Statistic name	Acronym	Description
Part Error Latency	pterlt	Total time that is taken for part errors In SAN Volume Controller it might always zero as no MPU scenario gets triggered.
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 you might have a 100 bytes file, of which you 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 that is taken to download the data from the cloud that were part of successful GET operation
Persisted Up Latency	pruplt	Total time that is 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 if multi part upload occurs. The part refers to the multi-part upload scenario. In SAN Volume Controller , it always remains zero as MPU size is 32 MiB. SAN Volume Controller blob size ranges from a few KBs to1 MiB.

Table 62. 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 63 describes the statistics collection for cloud per VDisk.

Table 63. Statistics collection for cloud per VDisk

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

Note: A block is 512 bytes.

Actions

The following actions are available to the user:

OK Click to change statistic collection.

Cancel
Click 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 it is 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, the following text is the output 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.

The following text shows the cache statistics for data reduction pools and volume copy cache statistics nodes and partitions:

```
<lca><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"
/>
<dca p="2089792" rh="305754" ph="178873" d="0"
ft="0" wt="0" fw="0" wh="0"
v="10348585" m="3334742" pm="1120" ri="10720"
wi="0" r="3923240" dav="0" dcn="0"
sav="59926" scn="6045" pav="48350" pcn="2723"
```

```
teav="0"  tsav="0"  tav="0"  pp="0"/>
<rca p="2089792" rh="305754" ph="178873" d="0"
ft="0" wt="0" fw="0" wh="0"
v="10348585" m="3334742" pm="1120" ri="10720"
wi="0" r="3923240" dav="0" dcn="0"
sav="59926" scn="6045" pav="48350" pcn="2723"
teav="0"  tsav="0"  tav="0"  pp="0"/>
<jca p="2089792" rh="305754" ph="178873" d="0"
ft="0" wt="0" fw="0" wh="0"
v="10348585" m="3334742" pm="1120" ri="10720"
wi="0" r="3923240" dav="0" dcn="0"
sav="59926" scn="6045" pav="48350" pcn="2723"
teav="0"  tsav="0"  tav="0"  pp="0"/>
</partition>
```

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 set up notifications.

Event reporting process

The following methods are used to identify 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 sent in an email to the appropriate support center.

Power-on self-test

When you turn on the system, the system board performs self-tests. During the initial tests, the hardware boot symbol is displayed.

All models perform a series of tests to check the operation of components and some of the options that are installed when the units are first turned on. This series of tests is called the power-on self-test (POST).

The node status LED is off until booting finishes and the system software is loaded If a critical failure is detected during the POST, the software is not loaded and the system error LED on the operator information panel is illuminated. If this failure occurs, use “MAP 5000: Start” on page 273 to help isolate the cause of the failure.

When the software is loaded, extra testing takes place, which ensures that all of the required hardware and software components are installed and functioning correctly.

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 whether 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, including an IBM FlashCopy operation completes.

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.

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.

Describing the fields in the event log

The event log includes fields with information that you can use to diagnose problems.

Table 64 describes some of the fields that are available to assist you in diagnosing problems.

Table 64. 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.

Table 64. Description of data fields for the event log (continued)

Data field	Description
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.
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	Extra data that gives the details of the condition that caused the event to be logged.

Event notifications

The system can use Simple Network Management Protocol (SNMP) traps, syslog messages, and call home emails to notify you and the support center 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 active service actions. If a recommended service action is active, these events are notified only if they are still unfixed when the service action completes.

Each event that the system detects is assigned a notification type of Error, Warning, Information, or Inventory. When you configure notifications, you specify where the

notifications should be sent and which notification types are sent to that recipient. The following table describes the types of event notifications.

Table 65. 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.
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. No remedial action is required when these notifications are sent.
Inventory	Inventory notifications contain summaries of system status and configuration settings.

Events with notification type “Error” or “Warning” are shown as alerts in the event log. Events with notification type “Information” are shown as messages.

SNMP traps

Simple Network Management Protocol (SNMP) is a standard protocol for managing networks and exchanging messages. The system can send SNMP messages that notify personnel about an event. You can use an SNMP manager to view the SNMP messages that the system sends. You can use the management GUI or the command-line interface to configure and modify your SNMP settings. You can specify up to a maximum of six SNMP servers.

You can use the Management Information Base (MIB) file for SNMP to configure a network management program to receive SNMP messages that are sent by the system. This file can be used with SNMP messages from all versions of the software. More information about the MIB file for SNMP is available at this website:

www.ibm.com/support

Search for the name of your storage system, and then search for “MIB file for SNMP”. Go to the downloads results to find **IBM Management Information Base (MIB) file for SNMP**. Click this link to find download options.

Syslog messages

The syslog protocol is a standard protocol for forwarding log messages from a sender to a receiver on an IP network. The system can send syslog messages that notify personnel about an event. The system can transmit syslog messages in either expanded or concise format. Servers configured with facility values of 0 - 3 receive syslog messages in concise format. Servers configured with facility values of 4 - 7 receive syslog messages in fully-expanded format. The default value is 0. The facility number used in syslog messages also identifies the origin of the message to the receiving server. You can use a syslog manager to view the syslog messages that the system sends. The system uses the User Datagram Protocol (UDP) to transmit the syslog message. You can specify up to a maximum of six syslog servers. You can use the management GUI or the command-line interface to configure and modify your syslog settings.

Table 66 shows how system notification codes map to syslog security-level codes.

Table 66. System notification types and corresponding syslog level codes

System notification type	Syslog level code	Description
ERROR	LOG_ALERT	Fault that might require hardware replacement that needs immediate attention.
WARNING	LOG_ERROR	Fault that needs immediate attention. Hardware replacement is not expected.
INFORMATIONAL	LOG_INFO	Information message used, for example, when a configuration change takes place or an operation completes.
TEST	LOG_DEBUG	Test message

Table 67 shows how system values of user-defined message origin identifiers map to syslog facility codes.

Table 67. System values of user-defined message origin identifiers and syslog facility codes

System value	Syslog value	Syslog facility code	Message format
0	16	LOG_LOCAL0	Full
1	17	LOG_LOCAL1	Full
2	18	LOG_LOCAL2	Full
3	19	LOG_LOCAL3	Full
4	20	LOG_LOCAL4	Concise
5	21	LOG_LOCAL5	Concise
6	22	LOG_LOCAL6	Concise
7	23	LOG_LOCAL7	Concise

Call home email

The call home function sends enhanced reports that include operational and event-related data and specific configuration information to the support center. When configured, this function alerts the support center about hardware failures and potentially serious configuration or environmental issues. The support center

can use configuration information to automatically generate best practices or recommendations that are based on your actual configuration.

To send email, you must configure at least one Simple Mail Transfer Protocol (SMTP) server. You can specify as many as 5 additional SMTP servers for backup purposes. The SMTP server must accept the relaying of email from the management IP address. Set the reply address to a valid email address. Send a test email to check that all connections and infrastructure are set up correctly. If you only want error and inventory information sent to the support center you can choose to hide sensitive entries, such as object names, cloud accounts, network information, certificates, host and user information from the reports.

Data that is sent with notifications

Notifications can be sent by using email, SNMP, or syslog. The data that is sent for each type of notification is the same. It includes:

- Record type
- Machine type
- Machine serial number
- Error ID
- Error code
- Software version
- FRU part number
- Cluster (system) name
- Node ID
- Error sequence number
- Time stamp
- Object type
- Object ID
- Problem data

Emails contain the following additional information that allows the Support Center to contact you:

- Contact names for first and second contacts
- Contact phone numbers for first and second contacts
- Alternate contact numbers for first and second contacts
- Offshift phone number
- Contact email address
- Machine location

Inventory information email

An inventory information email summarizes the hardware components and configuration of a system. Service personnel can use this information to contact you when relevant software updates are available or when an issue that can affect your configuration is discovered. It is a good practice to enable inventory reporting.

Because inventory information is sent using the call home email function, you must meet the call home function requirements and enable the call home email function before you can attempt to send inventory information email. You can

adjust the contact information, adjust the frequency of inventory email, or manually send an inventory email using the management GUI or the command-line interface.

The call home function sends enhanced reports that include specific configuration information to the support center. The support center can use this information to automatically generate recommendations that are based on your actual configuration.

The inventory email includes the following information about the clustered system on which the call home function is enabled. Sensitive information such as IP addresses is not included.

- Licensing information
- Details about the following objects and functions:
  - Drives
  - External storage systems
  - Hosts
  - MDisks
  - Volumes
  - Array types and levels
  - Easy Tier
  - FlashCopy
  - Metro Mirror and Global Mirror
  - HyperSwap

Example email

Figure 33 on page 110 shows an example of the header and VPD information that is included in an email. For details about other specific information that is included in the call home inventory for your system, configure your system to send an inventory email to yourself.

```
# Timestamp = Sun Mar  18 12:09:16 2018
# Timezone = +0000, UTC
# Organization =
# Machine Address =
# Machine City =
# Machine State = XX
# Machine Zip =
# Machine Country =
# Contact Name = lens
# Alternate Contact Name = N/A
# Contact Phone Number = 12357
# Alternate Contact Phone Number = N/A
# Offshift Phone Number = N/A
# Alternate Offshift Phone Number = N/A
# Contact Email = developer@system.com
# Machine Location = town
# Machine Type = 2076524
# Serial Number = 7836531
# Machine Part Number =
# System Version = 9.9.9 (build 140.12.0000000000000)
# Record Type = 6
# Frequency = 0
# Cluster Alias = 0x10036600202
# IBM Customer Number =
# IBM Component ID =
# IBM Country Code =
# Spectrum Virtualize Unique ID = 10036800202

# Cluster_VPD:

id:0000010036600202
name:mcr-fab1-cluster-29
location:local
partnership:
bandwidth:
total_mdisk_capacity:3.2TB
space_in_mdisk_grps:3.2TB
space_allocated_to_vdisks:1001.00GB
total_free_space:2.3TB
statistics_status:on
statistics_frequency:1
required_memory:32768
cluster_locale:en_US
time_zone:522 UTC
code_level:9.9.9 (build 140.12.0000000000000)
FC_port_speed:2Gb
id_alias:0000010036600202
gm_link_tolerance:300
gm_inter_cluster_delay_simulation:0
gm_intra_cluster_delay_simulation:0
email_reply:stevenfr@system.com
email_contact:lens
.
. (many lines were removed from this example)
.
```

Figure 33. Example of inventory information email

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.

Note: If more than one error occurs during an operation, the highest priority error code is displayed on the front panel. The lower the number for the error code, the higher the priority. For example, error code 1020 has a higher priority than error code 1370.

Using the error code tables

The error code tables list the various error codes and describe the actions that you can take.

About this task

Complete the following steps to use the error code tables:

Procedure

1. Locate the error code in one of the tables. If you cannot find a particular code in any table, call IBM Support Center for assistance.
2. Read about the action you must complete to correct the problem. Do not exchange field replaceable units (FRUs) unless you are instructed to do so.
3. Normally, exchange only one FRU at a time, starting from the top of the FRU list for that error code.

Event IDs

The system 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 119.

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 68 on page 112 provides a list of informational events, the notification type, and the reason for the event.

Table 68. Informational events

Event ID	Notification type	Description
060011	I	Error occurred during the recovery of the pool and some data has possibly been lost to one through to all vdisks
062004	I	Type conversion completed and the original copy has been deleted.
070570	I	Battery protection unavailable.
070571	I	Battery protection temporarily unavailable; one battery is expected to be available soon.
070572	I	Battery protection temporarily unavailable; both batteries are expected to be available soon.
070785	I	Battery capacity is reduced because of cell imbalance.
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.
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.
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.
981002	I	Fibre Channel discovery occurred; configuration changes are pending.

Table 68. Informational events (continued)

Event ID	Notification type	Description
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.
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.

Table 68. Informational events (continued)

Event ID	Notification type	Description
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.
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.
984512	I	A component firmware update is needed but is prevented from running.
984514	I	Node battery conditioning started.
984515	I	Node battery conditioning completed.
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 system 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.

Table 68. Informational events (continued)

Event ID	Notification type	Description
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.
986039	I	A data reduction pool repair process has begun.
986040	I	A data reduction pool repair process has completed successfully.
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.
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.

Table 68. Informational events (continued)

Event ID	Notification type	Description
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 support 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.

SCSI event reporting

Nodes can notify their hosts of events for SCSI commands that are issued.

SCSI status

Some events are part of the SCSI architecture and are handled by the host application or device drivers without reporting an event. Some events, such as read and write I/O events and events that are associated with the loss of nodes or loss of access to backend devices, cause application I/O to fail. To help troubleshoot these events, SCSI commands are returned with the `Check Condition` status and a 32-bit event identifier is included with the sense information. The identifier relates to a specific event in the event log.

If the host application or device driver captures and stores this information, you can relate the application failure to the event log.

Table 69 describes the SCSI status and codes that are returned by the nodes.

Table 69. SCSI status

Status	Code	Description
Good	00h	The command was successful.
Check condition	02h	The command failed and sense data is available.

Table 69. SCSI status (continued)

Status	Code	Description
Condition met	04h	N/A
Busy	08h	An Auto-Contingent Allegiance condition exists and the command specified NACA=0.
Intermediate	10h	N/A
Intermediate - condition met	14h	N/A
Reservation conflict	18h	Returned as specified in SPC2 and SAM-2 where a reserve or persistent reserve condition exists.
Task set full	28h	The initiator has at least one task queued for that LUN on this port.
ACA active	30h	This code is reported as specified in SAM-2.
Task aborted	40h	This code is returned if TAS is set in the control mode page 0Ch. The node has a default setting of TAS=0, which cannot be changed; therefore, the node does not report this status.

SCSI sense

Nodes notify the hosts of events on SCSI commands. Table 70 defines the SCSI sense keys, codes, and qualifiers that are returned by the nodes.

Table 70. SCSI sense keys, codes, and qualifiers

Key	Code	Qualifier	Definition	Description
2h	04h	01h	Not Ready. The logical unit is in the process of becoming ready.	The node lost sight of the system and cannot perform I/O operations. The additional sense does not have additional information.
2h	04h	0Ch	Not Ready. The target port is in the state of unavailable.	The following conditions are possible: • The node lost sight of the system and cannot perform I/O operations. The additional sense does not have additional information. • The node is in contact with the system but cannot perform I/O operations to the specified logical unit because of either a loss of connectivity to the backend controller or some algorithmic problem. This sense is returned for offline volumes.
3h	00h	00h	Medium event	This is only returned for read or write I/Os. The I/O suffered an event at a specific LBA within its scope. The location of the event is reported within the sense data. The additional sense also includes a reason code that relates the event to the corresponding event log entry. For example, a RAID controller event or a migrated medium event.

Table 70. SCSI sense keys, codes, and qualifiers (continued)

Key	Code	Qualifier	Definition	Description
4h	08h	00h	Hardware event. A command to logical unit communication failure has occurred.	The I/O suffered an event that is associated with an I/O event that is returned by a RAID controller. The additional sense includes a reason code that points to the sense data that is returned by the controller. This is only returned for I/O type commands. This event is also returned from FlashCopy target volumes in the prepared and preparing state.
5h	25h	00h	Illegal request. The logical unit is not supported.	The logical unit does not exist or is not mapped to the sender of the command.

Reason codes

The reason code appears in bytes 20-23 of the sense data. The reason code provides the node with a specific log entry. The field is a 32-bit unsigned number that is presented with the most significant byte first. Table 71 lists the reason codes and their definitions.

If the reason code is not listed in Table 71, the code refers to a specific event in the event log that corresponds to the sequence number of the relevant event log entry.

Table 71. Reason codes

Reason code (decimal)	Description
40	The resource is part of a stopped FlashCopy mapping.
50	The resource is part of a Metro Mirror or Global Mirror relationship and the secondary LUN in the offline.
51	The resource is part of a Metro Mirror or Global Mirror and the secondary LUN is read only.
60	The node is offline.
71	The resource is not bound to any domain.
72	The resource is bound to a domain that was recreated.
73	Running on a node that is contracted out for some reason that is not attributable to any path that is going offline.
80	Wait for the repair to complete, or delete the volume.
81	Wait for the validation to complete, or delete the volume.
82	An offline thin-provisioned volume that caused data to be pinned in the directory cache. Adequate performance cannot be achieved for other thin-provisioned volumes, so they are taken offline.
85	The volume that is taken offline because checkpointing to the quorum disk failed.
86	The **repairvdiskcopy -medium** command that created a virtual medium error where the copies differed.
93	An offline RAID-5 or RAID-6 array that caused in-flight-write data to be pinned. Good performance cannot be achieved for other arrays and so they are taken offline.
94	An array MDisk that is part of the volume that is taken offline because checkpointing to the quorum disk failed.

Table 71. Reason codes (continued)

Reason code (decimal)	Description
95	This reason code is used in MDisk bad block dump files to indicate that the data loss was caused by having to resync parity with rebuilding strips or some other RAID algorithm reason due to multiple failures.
96	A RAID-6 array MDisk that is part of the volume that is taken offline because an internal metadata table is full.

Object types

You can use the object code to determine the type of the object the event is logged against.

Table 72 lists the object codes and corresponding object types.

Table 72. Object types

Object code	Object type
1	mdisk
2	mdiskgrp
3	volume
4	node
5	host
7	iogroup
8	fcgrp
9	rcgrp
10	fcmap
11	rcmap
12	wwpn
13	cluster (system)
16	device
17	SCSI lun
18	quorum
34	Fibre Channel adapter
38	volume copy
39	Syslog server
40	SNMP server
41	Email server
42	User group
44	Cluster (management) IP
46	SAS adapter

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 software-based products, such as IBM Spectrum Virtualize.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 73 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 111.

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 73. 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 73. 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 73. 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
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 73. 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
010051		A Solid state drive is missing from the configuration	1202
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
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
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

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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.	1680
010102	E	A drive is reported as continuously slow with contributory factors.	1680
010103	E	Too many long IOs to drive (Mercury drives).	1680
010104	E	A drive is reported as continuously slow with contributory factors (Mercury drives).	1680
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 202
010126	W	A Flash drive with high write endurance usage rate.	“2560” on page 231
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

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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
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 204
045106	E	Enclosure secondary expander module FRU identity is not valid	"1266" on page 204
045107	E	Enclosure secondary expander module temperature sensor cannot be read	"1267" on page 204
045108	E	Enclosure secondary expander module temperature has passed warning threshold	"1098" on page 191
045109	E	Enclosure secondary expander module temperature has passed critical threshold	"1095" on page 189
045110	E	Enclosure display panel is not installed	"1268" on page 205
045111	E	Enclosure display panel temperature sensor cannot be read	"1268" on page 205
045112	E	Enclosure display panel temperature has passed warning threshold	"1098" on page 191
045113	E	Enclosure display panel temperature has passed critical threshold	"1095" on page 189
045114	E	Enclosure secondary expander module connector excluded due to too many change events	"1267" on page 204
045119	E	Enclosure display panel VPD cannot be read	"1268" on page 205
045120	E	Enclosure secondary expander module is missing	"1267" on page 204
045121	E	Enclosure secondary expander module connector excluded due to dropped frames	"1267" on page 204
045122	E	Enclosure secondary expander module connector is excluded and cannot be unexcluded	"1267" on page 204
045123	E	Enclosure secondary expander module connectors excluded as the cause of single ported drives	"1267" on page 204
045124	E	Enclosure secondary expander module leaf expander connector excluded as the cause of single ported drives	"1267" on page 204
050001	W	The Metro Mirror or Global Mirror relationship cannot be recovered.	1700

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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
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

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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
060008	W	Data reduction pool meta data corrupt	1862
060009	W	Pool's virtual disk copies offline due to failed data reduction pool repair	1860
060010	W	Virtual Disk Copies offline due to insufficient space in Data Reduction Pool	1865
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
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

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
070547	E	Pluggable TPM is missing or broken	1051
070548	E	The node has compression hardware configured but no compression hardware is available.	1046
070549	E	The node's compression hardware has failed.	1046
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
070560	E	Battery cabling fault.	1108
070561	E	Battery backplane or cabling fault.	1109
070562	E	The node hardware does not meet minimum requirements.	1183
070564	E	Too many software failures.	1188
070565	E	The internal drive of the node is failing.	1030
070569	E	CPU temperature breached critical threshold.	1093
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
070579	E	Battery subsystem has insufficient charge to save system data.	1107
070580	E	Unable to read the service controller ID.	1044
070581	E	UPS battery fault	1181
070582	E	UPS battery fault	1181
070583	E	UPS electronics fault	1171
070584	E	UPS output load high	1166
070585	E	UPS electronics fault	1171
070586	E	UPS AC input power fault	1141
070587	E	Incorrect type of uninterruptible power supply detected.	1152
070588	E	UPS configuration error	1151
070589	E	UPS ambient temperature threshold exceeded	1136
070590	E	UPS fault	1186

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
070670	W	Insufficient uninterruptible power supply charge to allow node to start.	1193
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
070710	W	A high speed SAS adapter is missing	1120
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
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
070736	E	A system board device breached warning temperature threshold.	1084
070737	E	A power supply breached temperature threshold.	1212
070738	E	A PCI Riser breached warning temperature threshold.	1085
070743	E	Boot drive missing or out of sync or failed.	1213
070744	W	A boot drive is in the wrong location.	1214
070745	W	Boot drive in unsupported slot.	1472
070746	W	Technician port connection is not valid.	3024
070747	W	Technician connected.	747
070760	E	Voltage fault	1110
070761	E	Voltage high	1100

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
070762	E	Voltage low	1105
070765	E	Fan error	1089
070766	E	CMOS battery has failed.	1670
070768	W	Ambient temperature warning	1094
070769	W	CPU temperature warning	1093
070770	W	Shutdown temperature reached	1092
070775	E	Power supply has a problem.	1097
070776	W	Power supply mains cable is unplugged.	1097
070777	E	Power supply is missing.	1097
070779	W	Battery is missing.	1129
070780	E	Battery has failed.	1130
070810	W	Battery is below the minimum operating temperature.	1476
070782	W	Battery is above the maximum operating temperature.	1475
070783	E	Battery has a communications error.	1109
070784	W	Battery is nearing end of life.	1474
070786	E	Battery VPD has a checksum error.	1130
070787	E	Battery is at a hardware revision level not supported by the current code level.	1473
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
070840	W	Detected hardware is not a valid configuration.	1198
070841	W	Detected hardware needs activation.	1199
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

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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
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
071548	E	The node has compression hardware configured but no compression hardware is available.	1046
071549	E	The node's compression hardware has failed.	1046
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

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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
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
071747	W	Technician connected.	747
071766	E	CMOS error	1670
071768	W	Ambient temperature warning	1094
071769	W	CPU temperature warning	1093
071810	W	Battery cold	1156
071782	W	Battery hot	1157
071786	E	Battery VPD checksum	1154
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
071850	W	Canister battery is nearing end of life	1159
072005	E	CMOS battery has a failure.	1670
072007	E	CMOS battery has a failure.	1670

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
072008	E	CMOS battery has a failure.	1032
072101		System board has more or fewer processors detected.	1025
072102		System board has more or fewer processors detected.	1025
072103		System board has more or lesss processors detected.	1032
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
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

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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
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

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
072860	W	Fibre-channel network fabric is too big.	1800
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
073006	W	The SAN is not correctly zoned. As a result, more than 512 ports on the SAN have logged into one system port.	1800
073251	E	More or less Fibre Channel adapters are detected.	1011
073252	E	Fibre Channel adapter is faulty.	1055
073258	E	Fibre Channel adapter has detected PCI bus error.	1013
073261	E	More or less Fibre Channel adapters are detected.	1011
073262	E	Fibre Channel adapter is faulty.	1055
073268	E	Fibre Channel adapter has detected PCI bus error.	1013
073271	E	More or less Fibre Channel adapters are detected.	1011
073272	E	Fibre Channel adapter is faulty.	1055
073278	E	Fibre Channel adapter has detected PCI bus error.	1013
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
073512	E	Enclosure VPD is inconsistent.	1008
073522	E	The system board service processor has failed.	1034
073524	E	Unable to update BIOS settings	1034
073528	E	Ambient temperature is too high during system startup.	1098
073541	E	Multiple hardware failures	1184
073551	W	Cannot form cluster due to lack of cluster resources, overridequorum possible	1192
073564	W	Too many software failures	1188
073569	E	CPU over temp.	1032
073576	E	Cluster state and configuration data cannot be read	1032
073650	W	The canister battery is not supported.	1149

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
073690	W	Node held in service state	1189
073715	W	Fewer SAS ports operational	1046
073717	W	SAS ports degraded	1046
073718	W	SASA port has unsupported SAS device	1669
073766	E	CMOS error	1670
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
073850	W	Canister battery is nearing end of life	1159
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
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

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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
074655	E	Canister battery communications error	1192
074656	W	Canister battery has insufficient charge to support a fire hose 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

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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
075011	E	Flash boot device has a failure.	1040
075012	E	Flash boot device has recovered.	1040
075015	E	Service controller has a read failure.	1044
075021	E	Flash boot device has a failure.	1040
075022	E	Flash boot device has recovered.	1040
075025	E	Service controller has a read failure.	1044
075031	E	Flash boot device has a failure.	1040
075032	E	Flash boot device has recovered.	1040
075035	E	A service controller read failure occurred	1044
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 73. 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
076501	E	A high-speed SAS adapter is missing.	1120
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
077105	E	The node service processor indicated a fan failure.	1089
077106	E	The node service processor indicated a fan failure.	1089
077107	E	The node service processor indicated a fan failure.	1089
077161	E	Node ambient temperature threshold has exceeded.	1094
077162	E	The node processor indicated a temperature warning.	1093
077163	E	The node service processor or ambient critical temperature threshold has exceeded.	1092
077165	E	Node ambient temperature threshold has exceeded.	1094
077166	E	Node processor temperature has a warning.	1093
077167	E	Node processor or ambient critical temperature threshold has exceeded.	1092
077171	E	System board voltage is high.	1101
077172	E	System board voltage is high.	1101
077173	E	System board voltage is high.	1101
077174	E	System board voltage is low.	1106
077175	E	System board voltage is low.	1106
077176	E	System board voltage is low.	1106
077178	E	Power management board has a voltage fault.	1110
077185	E	Node ambient temperature threshold has exceeded.	1094
077186	E	Temperature warning threshold exceeded	1093
077187	E	Temperature critical threshold exceeded	1092
077188	E	Power management board voltage has a fault.	1110
078001	E	Power domain error. Both nodes in the I/O group are powered by the same UPS.	1155

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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
079506	I	Technician port connection is not active.	3024
079507	I	Technician port connection is active.	3024
079508	W	Performance not optimised for V9000 variants without managed enclosures.	3300
079509	W	Performance not optimised for V9000 variants with managed enclosures.	3300
081001	E	Ethernet interface has a failure.	1400
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
083101	E	System failed to communicate with UPS.	1146
083102	E	UPS output loading was unexpectedly high.	1166
083103	E	Battery has reached end of life.	1191
083104	E	UPS has a battery fault.	1181
083105	E	UPS has an electronics fault.	1171
083107	E	UPS is overcurrent.	1161

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
083108	E	UPS has a fault but no specific FRU is identified.	1186
083109	E	The UPS has detected an input power fault.	1141
083110	E	UPS has a cabling error.	1151
083111	E	UPS ambient temperature threshold has exceeded.	1136
083112	E	UPS ambient temperature is high.	3001
083113	E	UPS crossed-cable test is bypassed because of an internal UPS software error.	3011
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 235
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
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

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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
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 219
086007	E	Key Server reported vendor information error	“1785” on page 219
086008	E	Failed to connect to Key Server	“1785” on page 219
086009	W	Key Server reported misconfigured primary	“1785” on page 219
087001	E	Cloud gateway service restarted	“2031” on page 227
087002	E	Cloud gateway service restarted too often	“1404” on page 208
087003	W	Cloud account SSL certificate will expire within the next 30 days	“3140” on page 238

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
087004	W	Cloud account not available, cannot resolve hostname	"1580" on page 210
087005	W	Cloud account not available, cannot contact cloud provider	"2310" on page 230
087006	W	Cloud account not available, cannot communicate with cloud provider	"2320" on page 230
087007	W	Cloud account not available, no matching CA certificate	"2300" on page 229
087008	W	Cloud account not available, no matching CA certificate	"2300" on page 229
087009	W	Cloud account not available, cannot establish secure connection with cloud provider	"3100" on page 236
087010	W	Cloud account not available, cannot authenticate with cloud provider	"2330" on page 230
087011	W	Cloud account not available, cannot obtain permission to use cloud storage	"2330" on page 230 "2305" on page 230
087012	W	Cloud account not available, cannot complete cloud storage operation	"3100" on page 236
087013	W	Cloud account not available, cannot access cloud object storage	"2105" on page 228
087014	W	Cloud account not available, incompatible object data format	"3135" on page 238
087016	W	Cloud account not available, cloud object storage encrypted	"1656 " on page 215
087017	W	Cloud account not available, cloud object storage not encrypted	"1656 " on page 215
087018	W	Cloud account not available, cloud object storage encrypted with the wrong key	"1657" on page 215
087019	W	No permission to use cloud storage snapshot operation	"2305" on page 230
087020	W	Cloud account out of space during cloud storage snapshot operation	"2125" on page 229
087021	W	Cannot create container object to cloud object storage during cloud snapshot operation	"2305" on page 230
087022	W	A cloud object could not be found during cloud snapshot operation.	"3108" on page 237
087023	W	A cloud object was found to be corrupt during cloud snapshot operation.	"3108" on page 237
087024	W	A cloud object was found to be corrupt during cloud snapshot decompression operation.	"3108" on page 237
087025	W	Etag integrity error during cloud snapshot operation	"3108" on page 237
087026	W	Internal Read error during cloud snapshot operation	"2120" on page 229

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
087027	W	Unexpected error occurred, cannot complete cloud snapshot operation	"3108" on page 237
087028	W	No permission to use cloud snapshot restore operation	"2305" on page 230
087029	W	A cloud object could not be found during a cloud snapshot restore operation	"3108" on page 237
087030	W	A cloud object was found to be corrupt during a cloud snapshot restore operation	"3108" on page 237
087031	W	A cloud object was found to be corrupt during a cloud snapshot restore decompression operation	"3108" on page 237
087032	W	Etag integrity error during cloud snapshot restore operation	"3108" on page 237
087033	W	Internal write error during cloud snapshot operation	"2120" on page 229
087034	W	Cannot create bad blocks on a managed disk during cloud snapshot restore operation.	"3108" on page 237
087035	W	Unexpected error occurred, cannot complete cloud snapshot restore operation	"3108" on page 237
087036	W	No permission to use cloud snapshot delete operation	"2305" on page 230
087037	W	A cloud object could not be found during a cloud snapshot delete operation	"3108" on page 237
087038	W	A cloud object was found to be corrupt during cloud snapshot delete operation	"3108" on page 237
087039	W	A cloud object was found to be corrupt during cloud snapshot delete decompression operation	"3108" on page 237
087040	W	Unexpected error occurred, cannot complete cloud snapshot delete operation	"3108" on page 237
087044	W	Cloud account out of space during cloud snapshot restore commit operation	"2125" on page 229
087045	W	Cloud account out of space during cloud snapshot delete operation	"2125" on page 229
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 fibre channel 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

Table 73. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
088004	W	A spare node could not be automatically removed from the cluster	3180
089001	W	Single PSU failure in bare metal server	1810
089002	W	Node IP missing, only single path connection available between nodes	1811
089003	W	The IP connections between nodes were broken	1812
089004	W	A node rejoined cluster with identity changed	1813
981110	I	iSCSI discovery occurred, configuration changes pending.	
981111	I	iSCSI discovery occurred, configuration changes complete.	
981112	I	iSCSI discovery occurred, no configuration changes were detected.	
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.	

Resolving a problem with the SAN Volume Controller boot drives

Complete the following steps to resolve most problems with SAN Volume Controller boot drives.

Before you begin

The node serial number (also known as the product or machine serial number) is on the MT-M S/N label (Machine Type - Model and Serial Number label) on the front (left side) of the node. The node serial number is written to the system board and to each of the two boot drives during the manufacturing process.

When the SAN Volume Controller software starts, it reads the node serial number from the system board (by using the node serial number for the panel name) and compares it with the node serial numbers that are stored on the two boot drives.

Specific node errors are produced under the following conditions:

- Unrecoverable node error 543: This error indicates that none of the node serial numbers that are stored in the three locations match. The node serial number

from the system board must match with at least one of the two boot drives for the SAN Volume Controller software to assume that node serial number is good.

- Unrecoverable node error 545: This error indicates that the node serial numbers on each boot drive match each other but are not the same as the node serial number from the system board. In this case, the node serial number on the system board might be wrong or the node serial number on the boot drives might be wrong. For example, the system board that is changed or the boot drives come from another node.
- Node error 743: This error indicates that the node serial number cannot be read from one of the two boot drives because that drive failed, is missing, or is out of sync with the other boot drive.
- Node error 744: This error indicates that the node serial number from one of the boot drives identifies as belonging to a different node. If boot drives were swapped between drive slots 1 and 2, node error 744 is produced.
- Node error 745: This error indicates that a boot drive is found in an unsupported slot. This error occurs when at least one of the first two drives is online and at least one invalid slot (3-8) is occupied.

About this task

An event is displayed in the Monitoring > Events panel of the management GUI if the problem produces node error 743, 744, or 745. Run the fix procedure for that event. Otherwise, connect to the technician port to use the MT-M S/N label on the node to see the boot drive slot information and determine the problem.

Attention: If a drive slot has `Yes` in the Active column, the operating system depends on that drive. Do not remove that drive without first shutting down the node.

- Do not swap boot drives between slots.
- Each boot drive has a copy of the VPD on the system board.
- Software upgrading is to one boot drive at a time to prevent failures during CCU.

Procedure

To resolve a problem with a boot drive, complete the following steps in order:

1. Remove any drive that is in an unsupported slot. Move the drive to the correct slot if you can.
2. If possible, replace any drive that is shown as missing from a slot. Otherwise, reseat the drive or replace it with a drive from FRU stock.
3. Move any drive that is in the wrong node back to the correct node.

   Note: If the node serial number does not match the node serial number on the system board, a drive slot has a status of `wrong_node`. If the serial number on the MT-M S/N label matches the node serial number on the drive, you can ignore this status.
4. Move any drive that is in the wrong slot back to the correct slot.
5. Reseat the drive in any slot that has a status of `failed`. If the status remains `failed`, replace the drive with one from FRU stock.
6. If the drive slot has status `out of sync` and `Yes` in the can_sync column, then:
   - Use the service assistant GUI to synchronize boot drives, or
   - Use the command-line interface (CLI) command **`satask chbootdrive -sync`**.

- If `No` is displayed in the can_sync column, you must resolve another boot drive problem first.

Replacing the system board:

7. Replace the SAN Volume Controller 2145-DH8 or SAN Volume Controller 2145-SV1 main board.

When neither of the boot drives have usable SAN Volume Controller software:

For example, if you replace both of the boot drives from FRU stock at the same time, neither boot drive has usable SAN Volume Controller software. If the SAN Volume Controller software is not running, the node status, node fault, battery status, and battery fault LEDs remain off.

8. If you cannot replace at least one of the original boot drives with a drive that contains usable SAN Volume Controller software and has a node serial number that matches the MT-M S/N label on the front of the node, contact IBM Remote Technical support. IBM Remote Technical support can help you install the SAN Volume Controller software with a bootable USB flash drive.
   - Field-based USB installation also repairs the node serial number and WWNN stored on each boot drive by finding values that are stored on the system board during manufacturing.
   - If the WWNN of this node that is changed in the past, you must change the WWNN again after you complete the SAN Volume Controller software installation. For example, if the node replaced an earlier SAN Volume Controller node, you must change the WWNN to that of the earlier node. You can repeat the change to the WWNN after the SAN Volume Controller software installation with the service assistant GUI or by command.

When every copy of the node serial number is lost:

For example, if you replace the system board and both of the boot drives with FRU stock at the same time, every copy of the node serial number is lost.

9. If you cannot replace one of the original boot drives or the original system board so that at least one copy of the original node serial number is present, you cannot repair the node in the field. Return the node to IBM for repair.

Results

The status of a drive slot is uninitialized only if the SAN Volume Controller software might not automatically initialize the FRU drive. This status can happen if the node serial number on the other boot drive does not match the node serial number on the system board. If the node serial number on the other boot drive matches the MT-M S/N label on the front that is left of the node, you can rescue the uninitialized boot drive from the other boot drive safely. Use the service assistant GUI or the **`satask rescuenode`** command to rescue the drive.

Resolving a problem with failure to boot

Light path LEDs might indicate a hardware failure on SAN Volume Controller 2145-DH8. SAN Volume Controller 2145-SV1 does not have light path LEDs, but it does have some diagnostic LEDs.Diagnostic LEDs might indicate a hardware failure on a SAN Volume Controller 2145-SV1.

Before you begin

If the SAN Volume Controller software is not running, then the node status and battery status LEDs are off. The service interfaces such as the technician port and statask.txt on a USB flash drive do not work.

Note: The SAN Volume Controller 2145-SV1 node fault LED can flash when a warning or critical error shows in the BMC event log (SEL). The warning or critical error prevents the SAN Volume Controller code from booting.

If the SAN Volume Controller software is running, then the node error LED might be on. The node error code and error data can be seen by connecting to the technician port or by using the other service interfaces. Look up the node error code in the IBM SAN Volume Controller Knowledge Center.

About this task

Complete the following steps if the SAN Volume Controller software is not running.

Procedure

1. Connect a monitor to the VGA port and a keyboard to a USB port. Consider any error messages on the monitor.

   For example, was it unable to find a device from which to boot? (Check that the SAS cables between the boot drives and the main system board are connected correctly.)
2. If no useful messages display on the monitor, complete the following steps.
   a. Power off the system by using the power button.
   b. Disconnect the power cables.
   c. Wait for 1 minute.
   d. Reconnect the power cables. The node attempts to power on.
   e. If the power LED comes on green, then watch the VGA monitor for any useful messages.
3. Attempt to access the UEFI setup utility on the VGA monitor by powering off, and then powering on by using the power button while you hold down the **ESC** or **Delete** key for the SAN Volume Controller 2145-SV1, or the **F1** key for the SAN Volume Controller 2145-DH8. If the **Setup Utility** displays, complete the following steps.
   a. If the node fault LED is flashing, access the **Bmc self test log** from the **Server Mgmt** tab to look for a cause.
   b. Access the **System Event Log** from the **Server Mgmt** tab. Events in this log might help to pinpoint the problem.
4. If by using the setup utility you are unable to pinpoint a broken component, or if the setup utility does not start, complete the following steps. It is best to initially investigate a fault with the DIMMs.
   a. Power off the system by using the power button.
   b. Disconnect the power cables.
   c. Remove the DIMMs but leave in one DIMM per microprocessor (CPU). For example, leave the DIMM in the first DIMM slot of each CPU.
   d. Reconnect the power cables. The node attempts to power on.

e. If the SAN Volume Controller software now boots and the node fault LED comes on, then one of the DIMMs that you removed might be broken. Repeat the steps with a different DIMM until you find the broken DIMM.
f. If all the DIMMs work when only one is fitted per CPU, then refit the DIMMs.

5. If the SAN Volume Controller software does not load with all of the tested good DIMMs fitted, complete the following steps. It is best to investigate a fault with the CPUs before you consider replacing the system board.
   a. Power off the system by using the power button.
   b. Disconnect the power cables.
   c. Remove the CPU labeled as CPU 1 on the system board.
   d. Reconnect the power cables. The node attempts to power on. If the SAN Volume Controller software now boots and the node fault LED comes on, then the CPU that you removed might be broken.
   e. If the SAN Volume Controller software does not boot, swap the CPUs. If the SAN Volume Controller software now boots and the node fault LED comes on, then the CPU that you removed might be broken.
6. If you did not find any evidence of a broken DIMM or CPU, contact IBM Remote Technical Support. They might want to know the state of the SAN Volume Controller 2145-SV1 system board LEDs.

Node error code overview

Node error codes describe failures that relate to a specific node.

Connect to the technician port so that you can use the service assistant GUI to view node errors and other error data.

Because node errors are specific to a node, for example, memory failures, the errors might be reported only on that node. However, if the node can communicate with the configuration node, then it is reported in the system event log.

When the node error code indicates that a critical error was detected that prevents the node from becoming a member of a clustered system, the Node fault LED is on.

The following example shows a node error:

```
Node Error
550  000125
```

The additional data is unique for any error code. It provides the necessary information to isolate the problem in an offline environment. Examples of extra data are disk serial numbers and field replaceable unit (FRU) location codes. For more information, refer to the help for the specific three-digit node error.

Node errors can be divided into 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 fails or the system detects that the software is corrupted. 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. Resolve the errors in priority order. The range of critical errors is 500 - 699.

Noncritical errors

A noncritical error code is logged when a hardware or code failure is related to 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 range of errors that are reserved for noncritical errors are 800 - 899.

Error code range

This topic shows the number range for each message classification.

Table 74 lists the number range for each message classification.

Table 74. Message classification number range

Message classification	Range	
Booting codes (no longer used)	100-299	
Node errors	Node rescue errors (no longer used)	300-399
	Log-only node errors (no longer used)	400-499
	Critical node errors	500-699
	Noncritical node errors	800-899
Error codes when creating a clustered system (no longer used)	700, 710	
Error codes when recovering a clustered system (no longer used)	920, 990	
Error codes for a clustered system	1001-3081	

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.

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.

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: If the progress bar has stopped for

more than 10 minutes, or if the display has shown codes 100 and 132 three times or more, go to "Resolving a problem with failure to boot" on page 147 to resolve the problem.

---

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 system 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.

---

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.

---

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.

---

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.

---

320 **A 2145 format operation has failed.**

Explanation: A 2145 format operation has failed.

User response: Exchange the FRU for a new FRU.

---

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.

---

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.
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: 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 differs from the amount required for the node 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. The next 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: Check the memory size of another 2145 that is in the same cluster.

Possible Cause-FRUs or other:

- Memory module (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%)

---

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, one at a time.

522 The system board service processor has failed.

Explanation: The service processor on the system board failed.

User response: For the 2145-DH8 only:

1. Shutdown the node.
2. Remove mains power cable.
3. Wait for lights to stop blinking.
4. Plug in power, and then wait for the node to boot.
5. If that fails, replace the system board.

Exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

2145-DH8

- System board assembly (100%)

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 Procedure: Rescuing node canister machine code from another node (node rescue).
2. If the rescue node does not succeed, use the hardware remove and replace procedures.

Possible Cause—FRUs or other:

- Node canister (80%)
- Other (20%)

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. The update might take 90 minutes.

User response:

1. If the progress bar has been stopped for more than 90 minutes, power off and reboot the node. If the boot progress bar stops again on this code, replace the FRU shown.
2. If the power off or restart does not work, try removing the power cords and then restarting.

528 Ambient temperature is too high during system startup.

Explanation: The ambient temperature read during the node startup procedures is too high for the node 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.
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.

---

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 system 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 system 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 cannot become active because it is unable to connect to enough system resources. The system resources are the node 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 system. This connection prevents the 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 nodes 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).

If the system topology is stretched and the number of operational nodes is less than half, then node error 550 is displayed. In this case, the Site Disaster Recovery feature cannot be used as the number of operational

nodes is less than the quorum required to create the system that uses the Site Disaster Recovery feature.

User response: Follow troubleshooting procedures to correct connectivity issues between the 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 system are powered on and operational.
3. Check the Fibre Channel port status. If any port is not active, run 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. Run the problem determination procedures for the network.
6. The quorum disk failed or cannot be accessed. Run the problem determination procedures for the disk controller.

---

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 SAN Volume Controller system, the first 11 digits are 500507680C0 for DH8 and 500507680F0 for SV1. The last 5 digits of the WWNN are given in the additional data of the error. For more information, see "Service assistant interface." 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 SAN Volume Controller node.

User response: Follow troubleshooting procedures to configure the WWNN of the node:

1. Find the cluster node with the same WWNN as the node reporting the error. The WWNN for a cluster 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.
2. If a cluster node with a duplicate WWNN is found, determine whether it, or the node reporting the error, has the incorrect WWNN. 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 correct WWNN cannot be determined contact your support representative 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 I/O group, 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 front panel power control button, after the node with the duplicate WWNN is updated.
6. If you are unable to find a cluster 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 cluster, 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 192.

---

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. View node VPD information, to see whether anything looks inconsistent. Compare the failing node VPD with the VPD of a working node of the same type. Pay particular attention to the number and type of CPUs and memory.
2. Replace any incorrect parts.

---

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. View hardware information.
2. Replace parts (canister or disk).

---

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: Exchange the FRUs for new FRUs, one at a time.

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.

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%)

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. The adapter might not be correctly installed, or it might have failed.

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. The location indicates an adapter slot, see the node canister description for the definition of the adapter slot locations

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. 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. If the previous steps have not isolated the problem, use the remove and replace procedures to replace the adapter, if this does not fix the problem replace the system board.

Possible Cause—FRUs or other cause:

- Fibre Channel adapter
- System board

---

701 A Fibre Channel adapter has failed.

Explanation: A Fibre Channel adapter has failed.

This node error does not, in itself, stop the node becoming active in the system. However, the Fibre Channel network might be being used to communicate between the nodes 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.

Data:

- A number indicating the adapter location. The location indicates an adapter slot. See the node description for the definition of the adapter slot locations.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Use the remove and replace procedures to replace the adapter. If this does not fix the problem, replace the system board.

Possible Cause-FRUs or other cause:

- Fibre Channel adapter
- System board

---

702 A Fibre Channel adapter has a PCI error.

Explanation: A Fibre Channel adapter has a PCI error.

This node error does not, in itself, stop the node from becoming active in the system. However, the Fibre Channel network might be being used to communicate between the nodes 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.

Data:

- A number indicating the adapter location. The location indicates an adapter slot. See the node description for the definition of the adapter slot locations.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Use the remove and replace procedures to replace the adapter. If this does not fix the problem, replace the system board.

Possible Cause-FRUs or other cause:

- Fibre Channel adapter
- System board

---

703 A Fibre Channel adapter is degraded.

Explanation: A Fibre Channel adapter is degraded.

This node error does not, in itself, stop the node becoming active in the system. However, the Fibre Channel network might be being used to communicate between the nodes 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.

Data:

- A number indicating the adapter location. The location indicates an adapter slot. See the node description for the definition of the adapter slot locations.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Use the remove and replace procedures to replace the adapter. If this does not fix the problem, replace the system board.

Possible Cause FRUs or other cause:

- Fibre Channel adapter
- System board

---

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 becoming active in the system. However, the Fibre Channel network might be being used to communicate between the nodes 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.

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 2145 node and the SFP transceiver in the connected switch or device.
   - Use the remove and replace procedures to replace the adapter.

Possible Cause-FRUs or other cause:

- Fibre Channel cable
- SFP transceiver
- Fibre Channel adapter

---

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 nodes 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 nodes.
- 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 nodes.
   - 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 might not be correctly installed or it might have failed.

Data:

- A number indicating the adapter location. The location indicates an adapter slot. See the node description for the definition of the adapter slot locations.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Possibilities:
   - If the adapter has been intentionally removed, use the management GUI recommended actions for the service error code, to acknowledge the change.
   - Use the remove and replace procedures to remove and open the node and check the adapter is fully installed.
   - If the previous steps have not isolated the problem, use the remove and replace procedures to replace the adapter. If this does not fix the problem, replace the system board.

Possible Cause-FRUs or other cause:

- High-speed SAS adapter
- System board

---

711 A SAS adapter has failed.

Explanation: A SAS adapter has failed.

Data:

- A number indicating the adapter location. The location indicates an adapter slot. See the node description for the definition of the adapter slot locations.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Use the remove and replace procedures to replace the adapter. If this does not fix the problem, replace the system board.

Possible Cause-FRUs or other cause:

- High-speed SAS adapter
- System board

---

712 A SAS adapter has a PCI error.

Explanation: A SAS adapter has a PCI error.

Data:

- A number indicating the adapter location. The location indicates an adapter slot. See the node description for the definition of the adapter slot locations.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Replace the adapter using the remove and replace procedures. If this does not fix the problem, replace the system board.

Possible Cause-FRUs or other cause:

- SAS adapter
- System board

---

713 A SAS adapter is degraded.

Explanation: A SAS adapter is degraded.

Data:

- A number indicating the adapter location. The location indicates an adapter slot. See the node description for the definition of the adapter slot locations.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.

2. Use the remove and replace procedures to replace the adapter. If this does not fix the problem, replace the system board.

Possible Cause-FRUs or other cause:

- High-speed SAS adapter
- System board

---

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 adapter might not be correctly installed or it might have failed.

Data:

- A number indicating the adapter location. The location indicates an adapter slot. See the node 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. If the adapter location is 0, use the remove and replace procedures to replace the Ethernet edge board, if there is one, or the system board.
3. If the location is not 0, there are a number of possibilities:
   a. Use the remove and replace procedures to remove and open the node and check that the adapter is fully installed.
   b. If the previous steps have not located and isolated the problem, use the remove and replace procedures to replace the adapter. If this does not fix the problem, replace the system board.

Possible Cause-FRUs or other cause:

- Ethernet adapter
- System board

---

721 An Ethernet adapter has failed.

Explanation: An Ethernet adapter failed.

Data:

- A number indicating the adapter location. The location indicates an adapter slot. See the node 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. If the adapter location is 0, use the remove and replace procedures to replace the system board.
3. If the adapter location is not 0, use the remove and replace procedures to replace the adapter. If this does not fix the problem, replace the system board.

Possible Cause—FRUs or other cause:

- Ethernet adapter
- System board

---

722 An Ethernet adapter has a PCI error.

Explanation: An Ethernet adapter has a PCI error.

Data:

- A number indicating the adapter location. The location indicates an adapter slot. See the node 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. If the adapter location is 0, use the remove and replace procedures to replace the system board.
3. If the adapter location is not 0, use the remove and replace procedures to replace the adapter. If this does not fix the problem, replace the system board.

Possible Cause—FRUs or other cause:

- Ethernet adapter
- System board

723 An Ethernet adapter is degraded.

Explanation: An Ethernet adapter is degraded.

Data:

- A number indicating the adapter location. The location indicates an adapter slot. See the node 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. If the adapter location is 0, use the remove and replace procedures to replace the system board.
3. If the adapter location is not 0, use the remove and replace procedures to replace the adapter. If this does not fix the problem, replace the system board.

Possible Cause—FRUs or other cause:

- Ethernet adapter
- System board

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. Check that 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 1 Gbps port, use the remove and replace procedures to replace the SFP transceiver in the system and the SFP transceiver in the connected switch or device.
   e. Replace the adapter or the system board (depending on the port location) by using the remove and replace procedures.

Possible Cause—FRUs or other cause:

- Ethernet cable
- Ethernet SFP transceiver
- Ethernet adapter
- System board

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: The ambient temperature of the node is close to the point where it stops performing I/O and enters a service state. The node is currently continuing to operate.

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 to make sure no vents are blocked.

Possible Cause-FRUs or other cause:

- None

---

769 CPU temperature warning.

Explanation: The temperature of the CPU within the node is close to the point where the node stops performing I/O and enters service state. The node is currently continuing to operate. This is most likely an ambient temperature problem, but it might be a hardware problem.

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. Ensure no vents are blocked.
4. Make sure the node fans are operational.
5. If the error is still reported, replace the node's CPU.

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 192.

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: Follow troubleshooting procedures to fix the hardware.

806 Node IP missing

Explanation: When the **sainfo lsnodeip** command was run, no IP addresses were found for the node. This error is caused if node IP addresses were not specified during installation or all node IP addresses were deleted.

User response:

1. Verify that the node IP addresses are missing by running the **sainfo lsnodeip** command.
2. Run the **satask chnodeip** command to set node IP addresses. Configure at least two node IP addresses.

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 a node canister with hardware type 500 in an enclosure with MTM 2076-624.

This is an expected condition when a control enclosure is being upgraded to a different type of node canister.

User response:

1. Check that all the upgrade instructions have been followed completely.
2. Use the management GUI to run the recommended actions for the associated service error code.

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 corrupted.

This node error clears after the USB flash drive that contains 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 that contains an encryption key is present but the expected file cannot be located. This error can occur if a key for a different system or an old key for this system was 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 clears when the USB flash drive identified is removed.

User response: Remove the USB flash drive from the port.

833 Unsupported USB device.

Explanation: An unsupported device was connected to a USB port.

Only USB flash drives are supported and this node error is 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: A change has been detected to the hardware configuration for this node. The new configuration is not supported by the node software. User action is required to repair the hardware or update the software.

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 getting node and clustered-system information. A change to the hardware configuration is expected.
3. If the hardware configuration is unexpectedly reduced, make sure the component has not been unseated. Hardware replacement might be necessary.
4. If a new hardware component is shown as unsupported, check the software version required to support the hardware component. Update the software to a version that supports the hardware.

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: A change has been detected in the node hardware configuration. The new configuration is supported by the node software. The new configuration does not become active until it is activated.

A node configuration is remembered only while it is active in a system. This node error is therefore resolved using the management GUI.

User response: 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.

Important: If you update your system software to version 8.1.1 or later from a version earlier than 8.1.0, on a system where you have already installed more than 64 GB of RAM, all nodes return from the update with an error code of 841. Versions 8.1.0 and later allocate memory in a different way than previous versions, so the RAM must be "accepted" again. To resolve the error, complete the following steps:

1. On a single node, run the **svctask chnodehw** command. Do not run the command on more than one node at a time.
2. Wait for the node to restart and return without the error.
3. Wait an additional 30 minutes for multipath drives to recover on the host.
4. Repeat this process for each node individually until you clear the error on all nodes.

---

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

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 exceeds the allowed limit. The node 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 cannot be determined, so the node's FC connectivity might vary after each restart. The connection might be with host systems, other storage systems, or with other nodes.

This error might be the reason the node 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 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 systems have been created on this node.

Data:

- None

User response:

1. Try to create the clustered system on a different node.
2. Contact your service representative.

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. Contact your service representative.

875 Request to cluster rejected.

Explanation: A candidate node could not be added to the clustered system. The node 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 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 cannot read its state data. It reports this error while waiting to be added back into a clustered system. If the node 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. 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.

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. Exchange 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.

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.

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: An unrecoverable canister error has occurred. Contact your support representative for assistance in replacing the canister.

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 the faulty canister. It is preferable that this canister shut 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 occurred after a successful system start. Note: The node that contains the flash module was not rejected by the cluster.

User response:

1. Replace the FRUs.
2. Check node status. If all nodes show a status of Online, mark the error that you 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 support to resolve the problem.
3. Go to repair verification MAP.

1046 Adapter has failed

Explanation: The node has hardware that is configured but no hardware is available, or the hardware failed.

User response:

1. In the management GUI, select **Monitoring** > **Events**. Click **Run Fix** on the associated service error for this issue. **Run Fix** starts a guided fix procedure that helps you resolve this issue.
2. Complete the suggested tasks that are provided by the fix procedure. You might be required to complete the following tasks based on the adapter location and specifics of your configuration:
   - If the adapter location is 0, use the remove and replace procedures to replace the system board.
   - If the adapter location is not 0, use the remove and replace procedures to replace the adapter. If this replacement does not fix the problem, replace the system board.

Possible Cause-FRUs or other cause:

- Adapter
- System board

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 Incorrect type of uninterruptible power supply detected

Explanation: The cluster is reporting that a node is not operational because of critical node error 587. For more information, see the details of node error 587.

User response: See node error 587.

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 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 support to resolve the problem.
3. Go to repair verification MAP.

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

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.

For the 2145-DH8, a fan has a fault condition.

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. Each fan module contains two fans.
2. For the 2145-DH8, mechanically stop fan or remove fan. If a fan is not installed, shut down the node, open it, and install the fan. If a fan is installed, replace fan FRU indicated by the FAN identifier that is supplied in the Extra data.
3. Exchange the FRU for a new FRU.
4. 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 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.

Possible Cause-FRUs or other:

2145-DH8

- 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.
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.

For the 2145-DH8, a power supply has a fault condition.

User response:

1. For the 2145-DH8, replace the power supply FRU.

   For all other models, complete the following steps.
2. Ensure that the power supply is attached correctly to both the node and to the UPS.
3. 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.
4. 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.
5. Follow the action that is specified for the LED states noted in the following list.
6. If the error is not automatically fixed after 2 minutes, contact support.
7. 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%)

- For the 2145-DH8: power supply (100%)

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.

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.

User response: This error can be corrected on its own. For example, if the partner node comes online, the reconditioning begins.

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.

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, one at a time.
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.

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.

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

1153 Canister battery is missing

Explanation: The canister battery cannot be detected.

User response:

1. Select the error code on the events page of the management GUI, and then run the fix procedure. For more information, see node error 651.

Possible Cause-FRUs or other:

- Canister battery

Other:

- Configuration error

1154 The canister battery has failed

Explanation: The canister battery failed. The battery might be showing an error state, it might have reached the end of life, or it might have failed to charge.

User response:

1. Select the error code on the events page of the management GUI, and then run the fix procedure. For more information, see node error 652.

Possible Cause-FRUs or other:

- Canister battery

Other:

- Configuration error

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

1156 The canister battery’s temperature is too low

Explanation: The canister battery’s temperature is below its minimum operating temperature.

User response:

1. Select the error code on the events page of the management GUI, and then run the fix procedure. For more information, see node error 653.

Possible Cause-FRUs or other:

- Canister battery

Other:

- Configuration error

1157 The canister battery’s temperature is too high

Explanation: The canister battery’s temperature is above its safe operating temperature.

User response:

1. Select the error code on the events page of the management GUI, and then run the fix procedure. For more information, see node error 654.

Possible Cause-FRUs or other:

- Canister battery

Other:

• Configuration error

1158 Canister battery communications fault

Explanation: The canister cannot communicate with the battery.

User response:

1. Select the error code on the events page of the management GUI, and then run the fix procedure. For more information, see node error 655.

Possible Cause-FRUs or other:

• Canister battery

Other:

• Configuration error

1159 The canister battery is reaching the end of its useful life.

Explanation: The canister battery is reaching the end of its useful life. Replace the canister battery within a week of the node error first being reported.

User response:

1. Select the error code on the events page of the management GUI, and then run the fix procedure. For more information, see node error 850.

Possible Cause-FRUs or other:

• Canister battery

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%)

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%)

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.

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.

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.

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.

1245 Array storage is critically low on space

Explanation: When available space goes below a predetermined critical threshold, the error is displayed, along with the associated event code:

085081 Array storage is critically low on available physical space

The array is automatically write-protected when the error is displayed; no further data can be written to the array until the situation is corrected.

The exact value of the critical threshold is not user-configurable, and is subject to change.

User response: Use the **rmvdisk** command to delete unwanted volumes.

Space is not immediately available after you remove a volume. To reclaim the space, run the **recoverarray -trim** command. This command might impact performance.

Note: Operating systems that do not have the unmap capability cannot delete data on a FlashSystem array.

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.

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 195. 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 195.

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.

1480 Array storage is low on space

Explanation: When available space goes below a predetermined warning threshold, the error is displayed, along with the associated event code:

085080 Array storage is low on available physical space

The exact value of the critical threshold is not user-configurable, and is subject to change.

User response: Use the **rmvdisk** command to delete unwanted volumes.

Space is not immediately available after you remove a volume. To reclaim the space, run the **recoverarray -trim** command. This command might impact performance.

Note: Operating systems that do not have the unmap capability cannot delete data on a FlashSystem array.

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 -srcip4 source_ip_address target_ip_address`.
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.

086009 Key Server reported misconfigured primary

An SKLM key server reported a server type that conflicted with the value defined on the system. The key server reported it is not the primary, but the server is defined to be the primary on the system.

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 179 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

1810 The bare metal server which runs SV_Cloud lost 1 power supply

Explanation: One of the two power supplies for the bare metal server that runs the IBM Spectrum Virtualize for Public Cloud software is not functioning.

User response: If the other power supply fails, you might lose the contents of the volume cache. To prevent this problem, complete one of the following actions:

- Turn off the IBM Spectrum Virtualize for Public Cloud software on the bare server. This forces the volumes in that I/O group to run in write-through mode, so no customer data is cached on the server. When the software stops, the cache is flushed to backend storage.
- Use the **chvdisk** to disable the cache for each volume in the I/O group. No customer data will be cached, so no data is lost if the second power supply fails.

1811 Node IP missing

Explanation: No IP addresses were found for a node in the system.

User response: Complete the following steps:

1. Run the **sainfo lsnodeip** command to determine the port that has no IP addresses.
2. Run the **satask chnodeip** command to set node IP addresses. Configure at least two node IP addresses.

1812 The connection between one pair of nodes is disconnected.

Explanation: A node is disconnected.

User response: Complete the following steps:

1. Run the `lseventlog` *`sequence_number`* command and note the values for the following attributes:

   reporting_node_id
   The ID for the node that reported the error.

   sense Among the other sense data, locate the destination_ip, which is the IP address of the disconnected node.

   object_id
   The port ID for the connection.
2. Run the following command:

   ```
   sainfo lsnodeip
   ```

   Note the node IP address, which is in same row with the port ID from the previous step.
3. As superuser, ping the disconnected node from the reporting node:

   ```
   ping -srcip4 --reporting_ip destination_ip
   ```

4. If the ping is successful, contact your support representative. If the ping fails, look for an issue with the network or with the IP configuration.

1813 Node identity changed

Explanation: The ID of the node was changed.

User response: Consult logs and the history of operations for the system to see if a valid reason exists for the change. If not, investigate the possibility of a security breach. You might want to change the backend storage passwords.

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 is 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 complete all of the actions that you intend to perform, mark the error as "fixed"; the volume copy then returns online.

- Determine why the storage pool free space is depleted. Any of the thin-provisioned volume copies, with auto-expand enabled, in this storage pool might have expanded at an unexpected rate. It might 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.

  Note: Migration is not supported for thin-provisioned or compressed volume copies in data reduction storage pools.
- 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 is automatically marked as "fixed", and the volume copy returns 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.

Remember: If the volume is thin-provisioned or compressed, the **-autoexpand** parameter must be enabled or the **mkvdisk** command fails.
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 that are 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. 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 5 minutes).

This error might also be reported because the primary system 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 occurred previously on Metro Mirror or Global Mirror between the same systems and all the following actions were attempted, contact your product support center to resolve the problem.
2. On both systems, check the partner Fibre Channel port mask to ensure that sufficient connectivity is available. If the partner Fibre Channel 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 was found and corrected, go to Action 11.
8. On the primary system that is reporting the error, examine the statistics by 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 are accounted for. 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 that are 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 that 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 stops in an inconsistent state.

Note: If the Metro Mirror or Global Mirror relationship 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. Therefore, you might want to start a FlashCopy of your Metro Mirror or Global Mirror auxiliary disks on the secondary system before you restart the Metro Mirror or Global Mirror relationship. This means that 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 system 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 Reconditioning of batteries required.

Explanation: Reconditioning 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 the software on your node.

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 software 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 can occur for one of the following reasons:

- The encryption key on the USB flash drive is corrupted.
- The expected encryption key cannot be found on the USB flash drive. This error can occur if a key for a different system or an old key for this system was 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.
- An unsupported device is connected to a USB port. Only USB flash drives are supported.

User response: Remove the USB flash drive or the unsupported device 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.

2561 Node IP is missing

Explanation: At least two IP addresses are required for each node.

User response: Use the **satask chnodeip** command to add the required IP addresses.

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

2650 Remote support application is unable to connect to IBM

Explanation: The remote support assistance feature could not establish a connection with the IBM support network.

User response: Complete the following steps:

1. Run the **lssystemsupportcenter** command to list the defined support centers.
2. If no proxy is defined (all of the support centers in the list show proxy=no), verify that all IP addresses and port numbers are correct. This information is pre-configured by IBM or defined by IBM.
3. If any proxy is defined (any of the support centers in the list shows proxy=yes), complete the following steps:
   a. Make sure that the IP addresses and port numbers are correct for all defined proxies.
   b. Verify the proxy configurations. For more information, refer to your remote-support proxy installation and configuration instructions.
4. Check your network firewall settings to ensure that the proxy (if configured) or the system ports (if no proxy is configured) can communicate with external IP addresses.
5. Run a connectivity test by entering the following command:

   ```
   chsystemsupportcenter -test
   ```

   If the test succeeds, the event is automatically fixed.
6. If the connectivity test fails, contact your support representative.

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.

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.

3300 Performance not optimised for configuration.

Explanation: A V9000 cluster can operate with the fibre queue switch set to ON or OFF. The optimum setting is determined automatically by the system based on whether you are managing any AE2 enclosures. If so, the switch must be ON for optimal performance. If the cluster detects that it is not in the correct performance mode, the 3300 error is displayed. This situation typically occurs when the fibre queue switch is manually changed by using the management GUI or the **chenclosure** command.

User response: Enter the following command for each node in the system, in turn, to restart the I/O process:

```
satask stopnode -warmstart
```

This command clears the error.

Procedure: SAN problem determination

You can solve problems on the system and its connection to the storage area network (SAN).

About this task

SAN failures might cause system 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 either the maintenance analysis procedures or the error codes sent you here, complete the following steps:

Procedure

1. If the SAN configuration was changed by changing the Fibre Channel cable connections or switch zoning, verify that the changes are correct and, if necessary, reverse those changes.
2. Verify that the power is turned on to all switches and storage controllers that the system uses, and that they are not reporting any hardware failures. If problems are found, resolve those problems before you proceed further.
3. Verify that the Fibre Channel cables that connect the systems to the switches are securely connected.
4. If you have a SAN management tool, use that tool to view the SAN topology and isolate the failing component.

Resolving a problem with SSL/TLS clients

Changing the security level of the system might cause the web interface, CIM clients, and other SSL/TLS clients to stop working. If any clients stop working, complete the following procedure.

Procedure

1. Wait 5 minutes and try again. The clients might still need to wait for the services to restart.

2. Confirm that the SSL/TLS implementation of the client (for example, the web browser or CIM management tool) is up to date and supports the level of security that is being enforced. If necessary, revert to a weaker SSL/TLS security level in the system and see whether this action resolves the issue.
3. If the problem is a browser problem, check the exact error message that is reported by the browser.

   If the error message is `cipher error`, `SSL error`, `TLS error`, or `handshake error`, then the error implies that there is a problem with the secure connection. In this case, confirm that the browser is up to date. All of the supported browsers (Internet Explorer, Firefox, Firefox ESR, and Chrome) support TLS 1.2 at the latest version.

   If there is only a blank screen, it is likely that either the web service needs to restart, or there is a problem unrelated to the security level.

Procedure: Making drives support protection information

You can use this procedure to migrate drives and arrays to pick up support for protection information.

About this task

Drives cannot start by using protection information for I/O requests on demand. They must be validated as having a correct format and general support for the function within the code. The system can validate the format and general support when the drive object is first discovered by the system. The requirement for system validation means that no drive that exists can use protection information on an update from version 730 regardless of use in the configuration. The system can reject a request to make a drive a candidate if the media is not formatted correctly for use with protection information. The process to use protection information on an existing drive is to use the system interface (GUI/CLI) and involves unmanaging and rediscovering the drive to allow the software to reacquire the drive characteristics.

The **`lsdrive`** view contains the `protection_enabled` field that shows whether a drive is using protection information. Drives and arrays that exist on an update to version 740 do not automatically pick up support for protection information. All newly discovered drives at this code level support protection information. If the system has spare capacity, then migration can proceed an MDisk at a time. Otherwise, the migration to using protection information on drives proceeds drive by drive.

To migrate a MDisk that is using spare storage capacity, complete the following procedure.

Procedure

1. Migrate data off the MDisk. The data migration can be accomplished by MDisk migration as part of MDisk delete (**`rmmdisk`**, **`lsmigrate`**) within a storage pool. You can also use volume mirroring to create an in-sync mirrored copy of each volume in another pool (**`addvdiskcopy`**). When it is copied (**`lsvdisksyncprogress`**), delete the original volume copies (**`rmvdiskcopy`**), and then delete the MDisk (**`rmmdisk`**) that has no data.
2. Follow the instructions in step 5 on page 241 for all the drives that are now candidates when the MDisk is deleted (see **`lsmigrate`**).

3. Re-create the array by using the system interface when all old members adopt protection information.
4. If the drive is a member, complete the following steps to adopt protection information on an individual drive.
   a. Run the **charraymember** command to eject the drive from the array (either immediately with redundancy loss or after an exchange).
   b. When the drive is no longer a member, follow the instructions in step 5 for candidates or spares.
   c. Repeat for the next member.
5. If the drive is a spare or candidate, complete the following steps:
   a. Use the management GUI to take the drive offline.
   b. When the drive is offline, use the system interfaces to change the drive's use to unused.
   c. The system reacquires the drive and brings it back online, possibly changing the drive ID.
   d. Attempt to make the drive a candidate.

      Depending on the drive, this step might generate error CMMVC6624E, The command cannot be initiated because the drive is not in an appropriate state to perform the task. This step is necessary to run the format command in the next step.
   e. Run the following format command.

      svctask chdrive -task format *drive_id*
   f. Wait approximately 3 minutes until the drive is online again. Use **lsdrive** *drive_id* to see the drive's online/offline status.
   g. Use the system interface to change the drive's use to `candidate`. If required, use the system interface to change the drive's use to `spare`.
   h. Enter **lsdrive** *drive_id* and check that the `protection_enabled` field is *yes*. This drive can now be used in an array.

Resolving a problem with new expansion enclosures

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. SAN Volume Controller 2145-24F and 2145-92F enclosures have two LEDs per Serial-attached SCSI (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 SAN Volume Controller 2145-24F expansion canister SAS ports and indicators and *SAN Volume Controller2145-92F expansion enclosure LEDs*.
- Verify that the SAS cabling to the expansion enclosure is correctly installed. To review the requirements, see Connecting the optional 2U SAS expansion enclosures to the 2145-DH8, Connecting the optional 2U SAS expansion enclosures to the 2145-SV1, and Connecting the optional 2145-92F SAS expansion enclosures.

Optical link failures

You might need to replace the small form-factor pluggable (SFP) transceiver when a failure occurs on a single Fibre Channel or 10G Ethernet link (applicable to Fibre Channel over Ethernet personality enabled 10G Ethernet link).

Before you begin

The following items can indicate that a single Fibre Channel or 10G Ethernet link failed:

- The Fibre Channel port status on the front panel of the node
- The Fibre Channel status light-emitting diodes (LEDs) at the rear of the node
- An error that indicates that a single port failed (703, 723).

About this task

Use only IBM supported 10 Gb SFP transceivers with the SAN Volume Controller 2145-DH8. Using any other SFP transceivers can lead to unexpected system behavior. Copper DAC is not supported by these 10 Gb ports. The SFP transceiver replacement in a 10 Gbps Ethernet adapter port is governed by the following rules:

- An existing 10 Gb SFP transceiver that is replaced with a new 10 Gb SFP transceiver: The 10 Gbps Ethernet adapter port detects a new SFP transceiver and becomes operational immediately.
- If the 10 Gbps Ethernet adapter port detects a new SFP transceiver and becomes operational immediately, the port has an incorrect SFP transceiver since the last reboot. The SFP transceiver is then replaced with the correct 10 Gb SFP transceiver. This situation can occur with an incompatible SFP transceiver (8 Gb SFP or 4 Gb SFP) that is inserted in the 10 Gbps Ethernet adapter port.
  - The node requires a reboot for detecting the new SFP transceiver. The new SFP transceiver will be operational only after reboot (no DMP is produced).
- The 10 Gbps Ethernet adapter port contains no SFP transceiver since the last reboot and the correct 10 Gb SFP transceiver is installed:
  - System reboot is required for detecting the new SFP transceiver.

Procedure

Attempt each of these actions, in the following order, until the failure is fixed.

1. Ensure that the Fibre Channel or 10G Ethernet cable is securely connected at each end.
2. Replace the Fibre Channel or 10G Ethernet cable.
3. Replace the SFP transceiver for the failing port on the node.

   Note: The system is 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 or Fibre Channel over Ethernet adapter on the node.

Ethernet iSCSI host-link problems

If you are having problems attaching to the Ethernet hosts, your problem might be related to the network, the system, or the host.

Note: The system and Host IP should be on the same VLAN. Host and system 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 system ports.
- Try to ping the 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 system host configuration.

Using the management GUI for system 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, system 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 ID`s can cause network connectivity problems.

Fibre Channel over Ethernet host-link problems

Problems attaching to the Fibre Channel over Ethernet (FCoE) hosts might be related to the network, the system, or the host.

Before you begin

If error code 705 on node is displayed, this code means the Fibre Channel (FC) I/O port is inactive. Fibre Channel over Ethernet (FCoE) uses Fibre Channel (FC) as a protocol and Ethernet as an inter-connect.

Note: Concerning a Fibre Channel over Ethernet (FCoE) enabled port: either the Fibre Channel forwarder (FCF) is not seen, or the Fibre Channel over Ethernet (FCoE) feature is not configured on switch.

- Verify that the Fibre Channel over Ethernet (FCoE) feature is enabled on the Fibre Channel forwarder (FCF).
- Verify the remote port (switch port) properties on the Fibre Channel forwarder (FCF).

If you connect the host through a Converged Enhanced Ethernet (CEE) Switch:

- Test your connectivity between the host and Converged Enhanced Ethernet (CEE) Switch.

- Ask the Ethernet network administrator to check the firewall and router settings to verify the settings.

Run **lsfabric**, and verify that the host is seen as a remote port in the output. If the host is not seen, in order:

- Verify that system and host get a Fibre Channel ID (FCID) on the Fibre Channel forwarder (FCF). If unable to verify, check the VLAN configuration.
- Verify that system and host port are part of a zone and that zone is in force.
- Verify that the volumes are mapped to host and are online. For more information, see **lshostvdiskmap** and **lsvdisk** in the description in the IBM Knowledge Center.

What to do next

If the problem is not resolved, verify the state of the host adapter.

- Unload and load the device driver.
- Use the operating system utilities (for example, Windows Device Manager) to verify that the device driver is installed, loaded, and operating correctly.

Servicing storage systems

Storage systems that are supported for attachment to the 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.

The following categories represent the types of service actions for storage systems:

- Controller code update
- Field replaceable unit (FRU) replacement

Controller code update

Ensure that you are familiar with the following guidelines for updating controller code:

- Check to see if the system supports concurrent maintenance for your storage system.
- Allow the storage system to coordinate the entire update process.
- If it is not possible to allow the storage system to coordinate the entire update process, complete the following steps:
  1. Reduce the storage system workload by 50%.
  2. Use the configuration tools for the storage system to manually fail over all logical units (LUs) from the controller that you want to update.
  3. Update the controller code.
  4. Restart the controller.
  5. Manually failback the LUs to their original controller.
  6. Repeat for all controllers.

FRU replacement

Ensure that you are familiar with the following guidelines for replacing FRUs:

- If the component that you want to replace is directly in the host-side data path (for example, cable, Fibre Channel port, or controller), disable the external data

paths to prepare for update. To disable external data paths, disconnect or disable the appropriate ports on the fabric switch. The system ERPs reroute access over the alternate path.

- If the component that you want to replace is in the internal data path (for example, cache, or drive) and did not completely fail, ensure that the data is backed up before you attempt to replace the component.
- If the component that you want to replace is not in the data path (for example, uninterruptible power supply units, fans, or batteries), the component is generally dual-redundant and can be replaced without more steps.

Chapter 7. Disaster recovery

Use these disaster recovery solutions for HyperSwap, Metro Mirror, Global Mirror, and Stretched System, where access to storage is still possible after the failure of a site.

HyperSwap

Active-active volume access is always provided while there is an up-to-date consistent copy. If there is an out-of-date consistent copy, there is not an automatic failover to it, nor is read-only access given to it. Use the **stoprcrelationship-access** or **stoprcconsistgrp-access** command to make it accessible. The relationship is then in the Idling state. After you enable access with the **stoprcrelationship-access** or **stoprcconsistgrp-access** command, use the **startrcrelationship -primary** <*master/aux*> or **startrcconsistgrp -primary** <*master/aux*> command to make the relationship leave the Idling state and resume HyperSwap replication. If you previously ran **overridequorum**, the **startrcrelationship** or **startrcconsistgrp** command fails.

When you resume HyperSwap replication, consider whether you want to continue by using the out-of-date consistent copy or revert to the up-to-date copy. To identify whether the master or auxiliary volume has access, look at the `primary` field that is shown by the **lsrcrelationship** or **lsrcconsistgrp** command. To continue to use the out-of-date copy, provide that value as the argument to the **-primary** parameter of the **startrcrelationship** or **startrcconsistgrp** command. To revert to the up-to-date copy, specify the opposite value as the argument to the **-primary** parameter. For example, if `master` is shown in the primary field of **lsrcconsistgrp** for an active-active consistency group in the Idling state, to revert to the up-to-date copy, use **startrcconsistgrp -primary** *aux*.

Metro Mirror and Global Mirror

Note: Inappropriate use of these procedures can allow host systems to make independent modifications to both the primary and secondary copies of data. You are responsible for ensuring that no host systems are continuing to use the primary copy of the data before you enable access to the secondary copy.

In a Metro Mirror or Global Mirror configuration a system is configured at each site. Relationships are configured between the systems to mirror data from storage at the primary site to storage at the secondary site. If an outage occurs at the secondary site the primary site continues operation without any intervention. If an outage occurs at the primary site, then it is necessary to enable access to storage at the secondary site.

Use the **stoprcrelationship-access** or **stoprcconsistgrp-access** command to enable access to the storage at the secondary site.

Stretched System

In a stretched system (formerly split-site) configuration, a system is configured with half the nodes at each site and a quorum device at a third location. If an outage occurs at either site, then the other nodes at the other site access the quorum device and continue operation without any intervention. If connectivity

between the two sites is lost, then whichever nodes access the quorum device first continues operation. For disaster recovery purposes you might want to enable access to the storage at the site that lost the race to access the quorum device.

Use the **satask overridequorum** command to enable access to the storage at the secondary site. This feature is only available if the system was configured by assigning sites to nodes and storage controllers, and changing the system topology to stretched.

Important: If you run disaster recovery on one site and then power up the remaining, failed site (which contained the configuration node at the time of the disaster), the cluster asserts itself as designed. This procedure would start a second, identical cluster in parallel, which can cause data corruption.You must follow these steps:

Example

1. Remove the connectivity of the nodes from the site that is experiencing the outage.
2. Power up or recover those nodes.
3. Run **satask leavecluster-force** or **svctask rmnode** command for all the nodes in the cluster.
4. Bring the nodes into candidate state.
5. Connect them to the site on which the site disaster recovery feature was run.

Other configurations

To recover access to the storage in other configurations, use “Recover system procedure” on page 249.

Chapter 8. Recovery procedures

This topic describes these recovery procedures: recover a system and back up and restore a system configuration. This topic also contains information about performing the node rescue.

Recover system procedure

The recover system procedure recovers the entire system if the system state is lost from all nodes. The procedure re-creates the 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**

  Do not attempt to initiate a node rescue (which the user can initiate either by using the SAN Volume Controller front panel, the service assistant GUI, or the `satask rescuenode` service CLI command). 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 configuration backup files.

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 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 are met in "When to run the recover system procedure" on page 250.
- All hardware errors are fixed. See "Fix hardware errors" on page 251
- All nodes have candidate status. Otherwise, see step 1.
- All nodes must be at the same level of code that the system had before the failure. If any nodes 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 nodes in the system. For more information, see "Removing system information for nodes with error code 550 or error code 578 using the service assistant" on page 252.

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 about 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 nodes with error code 550 or error code 578 by using the service assistant, but only if the recommended user response for these node errors are followed.
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**.
e. Remove hot spare nodes from the system and set them into candidate mode before starting the recovery process. Run the following CLI command to remove the node from the system.

   satask leavecluster -force *spare-node-panel-name*

   Once the node returns in service mode, run the following CLI command to set it into candidate mode.

   satask stopservice *spare-node-panel-name*

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 nodes 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 remote technical support. 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:

1. All enclosures and external storage systems are powered up and can communicate with each other.
2. Check that all nodes in the system are shown in the service assistant tool or using the service command: **sainfo lsservicenodes**. Investigate any missing nodes.
3. 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.
4. 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" on page 251.
5. Ensure all backend storage that is administered by the system is present before you run the recover system procedure.

6. 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.

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 that can be easily resolved:

- The node is powered off or the power cords were unplugged.
- Check the node status of every node that is a member of the system. Resolve all errors.
  - All nodes must be reporting either a node error 578, or no cluster name is shown on the `Cluster:` display. These error codes indicate that the system 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 node error 550. The 550 error can also indicate that a node is not able to join a system.

    Note: If any of the buttons on the front panel are pressed after these two error codes are reported, the report for the node returns to the 578 node error. The change in the report happens after approximately 60 seconds. Also, if the node was rebooted or if hardware service actions were taken, the node might show no cluster name on the `Cluster:` display.
  - If any nodes show `Node Error: 550`, record the data from the second line of the display. If the last character on the second line of the display is `>`, use the right button to scroll the display to the right.
    - In addition to the `Node Error: 550`, the second line of the display can show a list of node front panel IDs (7 digits) that are separated by spaces. The list can also show the WWPN/LUN ID (16 hexadecimal digits followed by a forward slash and a decimal number).
    - If the error data contains any front panel IDs, ensure that the node referred to by that front panel ID is showing `Node Error 578:`. If it is not reporting node error 578, ensure that the two nodes can communicate with each other. Verify the SAN connectivity and restart one of the two nodes by pressing the front panel power button twice.
    - If the error data contains a WWPN/LUN ID, verify the SAN connectivity between this node and that WWPN. Check the storage system to ensure that the LUN referred to is online. After verifying, restart the node by pressing the front panel power button twice.

    Note: If (after you resolve all these scenarios) half or greater than half of the nodes are reporting `Node Error: 578`, it is appropriate to run the recovery procedure.
  - 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 are not 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 nodes with error code 550 or error code 578 using the service assistant

The system recovery procedure works only when all nodes in the system of nodes to be recovered are in candidate status. If there are any nodes that display error code 550 or error code 578, 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.

Having used the service assistant to identify the system status and specific error, you will continue to use the service assistant to complete this procedure.

Selecting Change Node in the service assistant tool lists all of the Spectrum Virtualize nodes that have logged in to the node that is running the tool. Follow these guidelines when performing the recovery procedure:

- The system column of the node table identifies any nodes that are **not** in the system of nodes that must be recovered. Do not remove the system data for these nodes.
- Do not remove system information from any node that has online status, unless directed to do so by remote technical support.
- Do not remove the system data from the first node until you ensure that the following conditions are met:
  - All nodes in the system of nodes are listed in the Change Node part of the service assistant and are in service status with error 550 or 578
  - You have checked the extra node error data for each node to ensure that no other communication or hardware problem is causing the node error.

Procedure

1. In the change node part of the service assistant tool, select the radio button of the node with status service and error 550 or 578.
2. Select **Manage System**.
3. Click **Remove System Data**.

   Note: Spare nodes do not go into the 878/578 state that active nodes do. As such, the **Manage System** screen does not have the **Remove System Data** button for spare nodes. To remove system data on spare nodes, ssh onto any spare node and run the following commands.

   ```
   satask leavecluster -force

   satask stopservice
   ```

   Failure to remove the cluster state from the spare nodes results in the T3 failing, as the new cluster is unable to find the spare nodes as available candidates.
4. Confirm that you want to remove the system data when prompted.
5. 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.

6. Resolve any hardware errors until the error condition for all nodes in the system is **None**.
7. Ensure that all nodes in the system of nodes to be recovered display a status of candidate.

Results

When all nodes display a status of candidate and all error conditions are **None**, you can run the system recovery procedure.

Running system recovery by using the service assistant

You can use the service assistant to start recovery when all nodes that were members of the system are online and are in candidate status. If any nodes display error code 550 or 578, remove system information to place them into candidate status. Do not run the recovery procedure on different nodes in the same system; this restriction includes remote systems.

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 "Recover system procedure" on page 249.

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 nodes in the system; the nodes must not participate in any other system.

If the system has USB encryption, run the recovery from any node in the system that has a USB flash drive that is inserted which contains the encryption key.

If the system contains an encrypted cloud account that uses USB encryption, a USB flash drive with the system 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 system 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 that is replaced or node that is 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.

If the system uses both USB and key server encryption, providing either a USB flash drive or a connection to the key server (only one is needed, but both will work also) will unlock the system.

If you use USB flash drives to manage encryption keys, the T3 recovery causes the connection to a cloud service provider to go offline if the USB flash drive is not inserted into the system. To fix this issue, insert the USB flash drive with the current keys into the system.

If you use key servers to manage encryption keys, the T3 recovery causes the connection to a cloud service provider to go offline if the key server is offline. To fix this issue, ensure that the key server is online and available during T3 recovery.

If you use both key servers and USB flash drives to manage encryption keys, the T3 recovery causes the connection to a cloud service provider to go offline if none of the key providers are available. To fix this issue, ensure that either the key server is online or a USB flash drive is inserted into the system (only one is needed, but both will work also) during T3 recovery.

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 nodes.

   If you do not know the IP address or if it was not configured, configure the service address in the following way:

   - Use the technician port to connect to the service assistant and configure a service address on the node.
2. Log on to the service assistant.
3. Select **Recover System** from the navigation.
4. Follow the online instructions to complete the recovery procedure.
   a. Click **Prepare for Recovery**. The system searches for the most recent backup file and scans quorum disk. If this step is successful, **Preparation Status: Prepare complete** is displayed on the bottom of the page.
   b. 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.
   c. 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.
   d. If the quorum time and backup date are correct, click **Recover** to recreate the system.

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 that is completed with errors: One or more of the volumes are offline because fast write data was in the cache. To bring the volumes online, see "Recovering from offline volumes by 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 256.

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 by using the CLI."

Recovering from offline volumes by using the CLI

If a Tier 3 recovery procedure completes with offline volumes, then it is likely that the data that is in the write-cache of the node canisters is 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 that is 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 you fix the offline volume errors.

Important: For systems that are using data reduction pools, contact IBM support for assistance in recovering offline volumes.

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 there are corrupted volumes in a data reduction pool, the user must run the **recovervdiskbysystem** command to recover all volumes.

   Note: Use this command only under the supervision of IBM Support personnel.
3. If there are corrupted volumes in a pool, and the volumes are space efficient or compressed, run the following command:

   repairsevdiskcopy *vdisk_name* | *vdisk_id*

   This command brings the volume back online so that you can attempt to deal with the data loss.

   Note: If running the **repairsevdiskcopy** command does not start the repair operation, then use the **recovervdisk** command.
4. If the volume is not a space efficient or compressed volume, and it is outside of a data reduction pool, then run the **recovervdiskbysystem** command. This brings all corrupted volumes back online so that you can attempt to deal with the data loss.
5. 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.
6. 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 action 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 that requires system recovery, then the situation is more severe. Under this situation you need to check the internal storage:

- Parity arrays are likely syncing to restore parity; they do not have redundancy when this operation proceeds.
- Because there is no redundancy in this process, bad blocks might are 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 about the recovered configuration:

- FlashCopy mappings are restored as "idle_or_copied" with 0% progress. Both volumes must are 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 are restored.
- Any Global Mirror secondary volumes on the recovered system might have inconsistent data if there was replication I/O from the primary volume that is 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, which are 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.

  Note: Any data that was in the system write cache at the time of the failure is lost.
- 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 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:

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 system configuration. If you had to replace a faulty node, the new node must use the same worldwide node name (WWNN) as the faulty node that it replaced.

  Note: You can add new hardware, but you must not remove any hardware because the removal can cause the restore process to fail.
- No zoning changes were made on the Fibre Channel fabric that would prevent communication between the system and any storage controllers that are present in the configuration.
- 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 and panel_name.
- Use the CLI **sainfo lsservicenodes** command and the data to determine which nodes previously belonged in each I/O group.

Restoring the system configuration must be performed by one of the nodes previously in I/O group zero. For example, **property name="IO_group_id" value="0"** . The remaining nodes must be added, as required, in the appropriate order based on the previous **IO_group_id** of its nodes.

The system 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, system nodes, 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 by 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 75 on page 261 describes the three files that are created by the backup process:

Table 75. 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 system procedure fails or if the data that is stored on the volumes is not required. This procedure is also known as Tier 4 (T4) recovery. For directions on the recover procedure, see "Recover system procedure" on page 249.

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 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 system) 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 reencrypted 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 all key servers 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 system'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 a system that contains nodes with more than four Fibre Channel ports, the system **localfcportmask** and **partnerfcportmask** settings are manually reapplied before you restore your data. See step 8 on page 264.
- For a system with nodes that are connected to expansion enclosures, all nodes must be added into the system before you restore your data. See step 9 on page 264.
- For 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 10 on page 265.
- 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).
- If the system has USB encryption, run the recovery from any node in the system that has a USB flash drive inserted which contains the encryption key.
- If the system has key server encryption, run the recovery on a node that is attached to the key server. The keys are fetched remotely from the key server.
- If the system uses both USB and key server encryption, providing either a USB flash drive or a connection to the key server (only one is needed, but both will work also) will unlock the system.
- For systems with a 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 the system contains an encrypted cloud account that is configured with both USB and key server encryption, the master keys from both need to be available at the time of a T4 recovery.
- If you use USB flash drives to manage encryption keys, the T4 recovery causes the connection to a cloud service provider to go offline if the USB flash drive is not inserted into the system. To fix this issue, insert the USB flash drive with the current keys into the system.
- If you use key servers to manage encryption keys, the T4 recovery causes the connection to a cloud service provider to go offline if the key server is offline. To fix this issue, ensure that the key server is online and available during T4 recovery.
- If you use both key servers and USB flash drives to manage encryption keys, the T4 recovery causes the connection to a cloud service provider to go offline if the key server is offline. To fix this issue, ensure that both the key server is online and a USB flash drive is inserted into the system during T4 recovery.
- If the system contains an encrypted cloud account that uses USB encryption, a USB flash drive with the system 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 system 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.
- After a T4 recovery, volumes with cloud snapshots that were enabled before the recovery need to have the cloud snapshots manually reenabled.

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.
2. Create a system. If possible, use the node that was originally in I/O group `0`.
   - For SAN Volume Controller 2145-DH8 and SAN Volume Controller 2145-SV1 systems, use the technician port.
3. In a supported browser, enter the IP address that you used to initialize the system and the default superuser password (`passw0rd`).
4. Issue the following CLI command to ensure that only the configuration node is online:

   ```
   svcinfo lsnode
   ```

   The following output is an example of what is displayed:

   ```
   id name status IO_group_id IO_group_name config_node
   1 node1 online 0 io_grp0 yes
   ```
5. Using the command-line interface, 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 system for which you want to restore the configuration.

Note: Because the RSA host key changed, a warning message might display when you connect to the system by using SSH.

6. 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** > **Support Package**.
   b. Expand **Manual Upload Instructions** and select **Download Support Package**.
   c. On the **Download New Support Package or Log File** page, select **Download Existing Package**.
   d. 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) Select the files and click **Download** to download them to your computer.
   e. If a recent configuration file is not present on this node, configure service IP addresses for other nodes and connect to the service assistant to look for configuration files on other nodes. For more information, see the **Service IPv4 or Service IPv6 options** topic at Service IPv4 or Service IPv6 options.

   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`.

7. 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
   ```

8. If the system contains any nodes with a 10 GB interface adapter or a second Fibre Channel interface adapter that is installed and non-default **localfcportmask** and **partnerfcportmask** settings were previously configured, then manually reconfigure these settings before you restore your data.

9. If the system uses a stretched or HyperSwap topology with nodes that are at two sites, or if the system contains any nodes with internal flash drives (including nodes that are connected to expansion enclosures), these nodes must be added to the system now. 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 addnode -panelname panel_name -iogrp iogrp_name_or_id -name node_name
   ```

   Where *panel_name* is the name that is displayed on the panel, *iogrp_name_or_id* is the name or ID of the I/O group to which you want to add this node, and *node_name* is the name of the node.

10. 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 addnode -panelname *panel_name* **-iogrp** *iogrp_name_or_id* **-name** *node_name*

       Where *panel_name* is the name that is displayed on the panel, *iogrp_name_or_id* is the name or ID of the I/O group to which you want to add this node, and *node_name* is the name of the node.
    b. 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** *remote_copy_port_group_id* **-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** *remote_copy_port_group_id* **-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.
    c. 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), i*scsi_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_secr*

   Where *iogrp_id* is the I/O group ID or name that is added, *iscsi_user_name* is the target controller user name that is 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 that is added.

4) If the configuration is a HyperSwap or stretched system, the controller name and site needs to be restored. To restore the controller name and site, determine c*ontroller_name* and controller *site_id/name* 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 or 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.

11. 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.

12. 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`

13. Open the log file from the server where the copy is now stored.
14. 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 15 on page 267.

- If you need assistance, contact the support center.

15. 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.

16. 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*

17. Open the log file from the server where the copy is now stored.
18. 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_2145:your_cluster_name:superuser>
    ```

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 by 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

Completing the node rescue when the node boots

On SAN Volume Controller 2145-CG8 or 2145-CF8, you might have to replace the hard disk drive. Or, if the software on the hard disk drive is corrupted, you can use the node rescue procedure to reinstall the software across the Fibre Channel fabric from its partner node in the same I/O group.

Before you begin

Similarly, if you replace the service controller, use the node rescue procedure to ensure that the service controller has the correct software.

About this task

Attention: If you recently replaced both the service controller and the disk drive as part of the same repair operation, node rescue fails.

Node rescue works by booting the operating system from the service controller and running a program that copies all the SAN Volume Controller software from any other node that can be found on the Fibre Channel fabric.

Attention: When you run node rescue operations, run only one node rescue operation on the same SAN, at any one time. Wait for one node rescue operation to complete before you start another.

Perform the following steps to complete the node rescue:

Procedure

1. Ensure that the Fibre Channel cables are connected.
2. Ensure that at least one other node is connected to the Fibre Channel fabric.
3. Ensure that the SAN zoning allows a connection between at least one port of this node and one port of another node. It is better if multiple ports can connect, which is important if the zoning is by worldwide port name (WWPN) and you are using a new service controller. In this case, you might need to use SAN monitoring tools to determine the WWPNs of the node. If you need to change the zoning, remember to set it back when the service procedure is complete.
4. Turn off the node.
5. Press and hold the left and right buttons on the front panel.
6. Press the power button.
7. Continue to hold the left and right buttons until the node-rescue-request symbol is displayed on the front panel (Figure 34).

Results

Figure 34. Node rescue display

The node rescue request symbol displays on the front panel display until the node starts to boot from the service controller. If the node rescue request symbol displays for more than 2 minutes, go to the hardware boot MAP to resolve the problem. When the node rescue starts, the service display shows the progress or failure of the node rescue operation.

Note: If the recovered node was part of a clustered system, the node is now offline. Delete the offline node from the system and then add the node back into the system. If node recovery was used to recover a node that failed during a

software update process, it is not possible to add the node back into the system until the code update process completes. This process can take up to 4 hours for an eight-node clustered system.

Chapter 9. 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 system 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 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 that is 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 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** outputs the virtual medium errors that is 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 76 lists the bad block error codes.

Table 76. Bad block errors

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

Table 76. Bad block errors (continued)

Error code	Description
1225	The system fails 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. However, it is possible 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 is not 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 10. Using the maintenance analysis procedures

The maintenance analysis procedures (MAPs) inform you how to analyze a failure that occurs with a SAN Volume Controller node.

About this task

SAN Volume Controller nodes must be configured in pairs so you can perform concurrent maintenance.

When you service one node, the other node keeps the storage area network (SAN) operational. With concurrent maintenance, you can remove, replace, and test all field replaceable units (FRUs) on one node while the SAN and host systems are powered on and doing productive work.

Note: Unless you have a particular reason, do not remove the power from both nodes unless instructed to do so. When you need to remove power, see "MAP 5350: Powering off a node" on page 284.

Procedure

- To isolate the FRUs in the failing node, complete the actions and answer the questions that are given in these maintenance analysis procedures (MAPs).
- When instructed to exchange two or more FRUs in sequence:
  1. Exchange the first FRU in the list for a new one.
  2. Verify that the problem is solved.
  3. If the problem remains:
     a. Reinstall the original FRU.
     b. Exchange the next FRU in the list for a new one.
  4. Repeat steps 2 and 3 until either the problem is solved, or all the related FRUs are exchanged.
  5. Complete the next action that is indicated by the MAP.
  6. If you are using one or more MAPs because of a system error code, mark the error as fixed in the event log after the repair, but before you verify the repair.

  Note: Start all problem determination procedures and repair procedures with "MAP 5000: Start."

MAP 5000: Start

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the system.

Before you begin

Note: The service assistant interface must be used if there is no front panel display, for example on the SAN Volume Controller 2145-DH8.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 10, "Using the maintenance analysis procedures."

This MAP applies to all system models. Be sure that you know which model you are using before you start this procedure. To determine which model you are working with, look for the label that identifies the model type on the front of the node.

You might be sent here for one of the following reasons:

- The fix procedures sent you here
- A problem occurred during the installation of the system
- Another MAP sent you here
- A user observed a problem that was not detected by the system

System nodes are configured in pairs. While you service one node, you can access all the storage managed by the pair from the other node. With concurrent maintenance, you can remove, replace, and test all FRUs on one system while the SAN and host systems are powered on and doing productive work.

Notes:

- Unless you have a particular reason, do not remove the power from both nodes unless instructed to do so.
- If an action in these procedures involves removing or replacing a part, use the applicable procedure.
- If the problem persists after you complete the actions in this procedure, return to step 1 of the MAP to try again to fix the problem.

Procedure

1. **Were you sent here from a fix procedure?**

   NO Go to step 2

   YES Go to step 6 on page 275

2. (from step 1)

   Access the management GUI. See "Accessing the management GUI" on page 57

3. (from step 2)

   Does the management GUI start?

   NO Go to step 6 on page 275.

   YES Go to step 4.

4. (from step 3)

   Is the Welcome window displayed?

   NO Go to step 6 on page 275.

   YES Go to step 5.

5. (from step 4)

   Log in to the management GUI. Use the user ID and password that is provided by the user.

   Go to the Events page.

   Start the fix procedure for the recommended action.

   Did the fix procedures find an error that is to be fixed?

   NO Go to step 6 on page 275.

   YES Follow the fix procedures.

6. (from steps 1 on page 274, 3 on page 274, 4 on page 274, and 5 on page 274)
   Is the power indicator off? Check to see whether the power LED is off.

   NO Go to step 7.

   YES Try to turn on the nodes.

   Note:

   SAN Volume Controller 2145-DH8 does not have an external uninterruptible power supply unit. This system has battery modules in its front panel instead.
   If the nodes are turned on, go to step 7; otherwise, go to "MAP 5040: Power SAN Volume Controller 2145-DH8" on page 279.
7. (from step 6)
   Does the node show a hardware error?

   NO Go to step 8.

   YES The service controller for the system failed. (The SAN Volume Controller 2145-DH8 does not have a service controller.)
   a. Check that the service controller that is indicating an error is correctly installed. If it is, replace the service controller.
   b. Go to "MAP 5700: Repair verification" on page 300.
8. (from step 7)
   Is the operator-information panel error LED (4 in Figure 35 or 7 in Figure 36 on page 276) illuminated or flashing? Or, is the check log LED (6 in Figure 36 on page 276) illuminated or flashing?
   Figure 35 shows the operator-information panel for the SAN Volume Controller 2145-SV1.

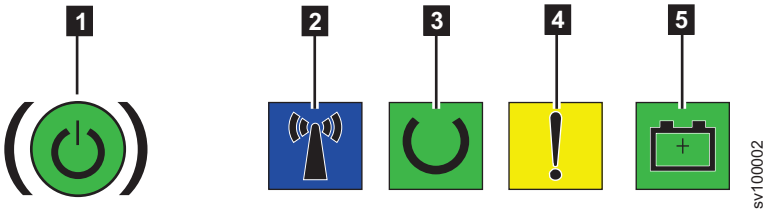

1 Power-control button and power-on LED
2 Identify LED
3 Node status LED
4 Node fault LED
5 Battery status LED

Figure 35. SAN Volume Controller 2145-SV1 operator-information panel

Figure 36 on page 276 shows the operator-information panel for the SAN Volume Controller 2145-DH8.

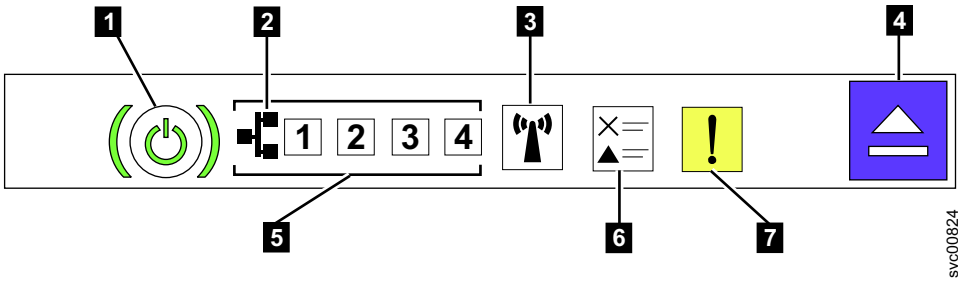

1 Power-control button and power-on LED
2 Ethernet icon
3 System-locator button and LED
4 Release latch for the light path diagnostics panel
5 Ethernet activity LEDs
6 Check log LED
7 System-error LED

Note: If the node has more than four Ethernet ports, activity for ports 5 and above is not indicated by the Ethernet activity LEDs on the operator-information panel.

Figure 36. SAN Volume Controller 2145-DH8 operator-information panel

NO Go to step 9.

YES Go to "MAP 5800: Light path" on page 301.

9. (from step 8 on page 275)

   For a 2145-DH8 model, is the node status LED, node fault LED, and battery status LED that you see in Figure 37 on page 277 all off?

   NO Go to step 11.

   YES Go to step 10.

10. (from step 9)

    For 2145-DH8, has the node status LED, node fault LED, and battery status LED that you see in Figure 37 on page 277 all been off for more than 3 minutes?

    NO Go to step 11.

    YES For 2145-DH8, go to step 20 on page 278. Otherwise:

    a. Go to "Resolving a problem with failure to boot" in the *IBM SAN Volume Controller Troubleshooting Guide*
    b. Go to "MAP 5700: Repair verification".

11. (from step 9)

    Is the node fault LED (8 in Figure 37 on page 277) on the front panel of a SAN Volume Controller 2145-DH8 on?

    Figure 37 on page 277 shows the node fault LED.

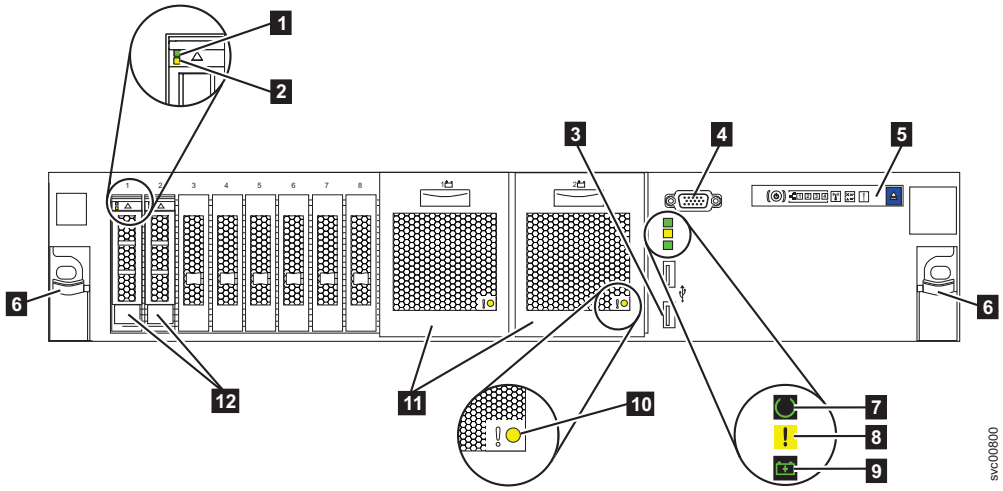

7 Node status LED
8 Node fault LED
9 Battery status LED

Figure 37. SAN Volume Controller 2145-DH8 front panel

NO Go to step 12.

YES Complete these steps:

a. Access the service assistant interface via the Technician port for node access and follow the service recommendation presented.
b. Go to "MAP 5700: Repair verification" on page 300.

12. (from step 11 on page 276)

Is `Booting` indicated on the node?

NO Go to step 14.

YES Go to step 13.

13. (from step 12)

If the boot progress does not advance for more than 3 minutes, the progress is stalled.

Has the boot progress stalled?

NO Go to step 14.

YES

a. Go to "MAP 5700: Repair verification" on page 300.

14. (from step 12 and step 13)

Is the node fault LED, which is the middle of the three status LEDs on the front panel of a SAN Volume Controller 2145-DH8, on? Figure 37 shows the node fault LED.

NO Go to step 15 on page 278.

YES Complete these steps:

a. Note the failure code and go to "Node error code overview" on page 149 to complete the repair actions.
b. If the node does not have a front panel display, access the service assistant interface via the Technician port for node access and follow the service recommendation presented.

c. Go to "MAP 5700: Repair verification" on page 300.

15. (from step 14 on page 277)

    Is `Cluster Error` reported on the node?

    NO Go to step 16.

    YES A cluster error was detected. This error code is displayed on all the operational nodes in the system. The fix procedures normally repair this type of error. Follow these steps:

    a. Complete the error code repair actions.
    b. Go to "MAP 5700: Repair verification" on page 300.

16. (from step 15)

    Is `Powering Off`, `Restarting`, `Shutting Down`, or `Power Failure` reported on the node?

    NO Go to step 17.

    YES Wait for the operation to complete and then return to step 1 on page 274 in this MAP. If the progress is stalled after 3 minutes, press **power** and go to step 17.

17. (from step 16)

    Did the node power off?

    NO Complete the following steps:

    a. Remove the power cord from the rear of the box.
    b. Wait 60 seconds.
    c. Replace the power cord.
    d. If the node does not power on, press **power** to power on the node and then return to step 1 on page 274 in this MAP.

    YES Complete the following steps:

    a. Wait 60 seconds.
    b. Click **power** to turn on the node and then return to step 1 on page 274 in this MAP.

18. **Is there a node that is not a member of a clustered system?** You can tell if a node is not a member of a system because the node status LED is off or blinking for SAN Volume Controller 2145-DH8.

    NO Go to step 19.

    YES The node is not a member of a system. The node might have been deleted during a maintenance procedure and was not added back into the system. Make sure that each I/O group in the system contains two nodes. If an I/O group has only one node, add the node back into that system. Then, ensure that the node is restored to the same I/O group from which it was deleted.

19. No errors were detected by the system. If you suspect that the problem that is reported by the customer is a hardware problem, follow these tasks:

    a. Complete Problem Determination procedures on your host systems, disk controllers, and Fibre Channel switches.
    b. Ask IBM remote technical support for assistance.

20. (from step 10 on page 276)

Can you access the service assistant interface through the 2145-DH8 technician port or service IP address, or use a USB flash drive to get `satask_results.html`?

NO The system software might not be running. Attach a USB keyboard and VGA monitor to the 2145-DH8 to see whether the node is stuck booting.

YES Go to step 21.

21. (from step 20 on page 278)

Can node error 561 be seen?

NO Follow the recommended action for any node error that can be seen.

YES The system software might not be able to communicate with the battery backplane.

Check the connections between the system board and the battery backplane. Then, follow the recommended action for node error 561.

Results

If you suspect that the problem is a software problem, see "Updating the system" documentation for details about how to update your entire system environment.

If the problem is still not fixed, collect diagnostic information and contact IBM Remote Technical Support.

MAP 5040: Power SAN Volume Controller 2145-DH8

It might become necessary to solve problems that are associated with power on the SAN Volume Controller 2145-DH8.

Before you begin

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 10, "Using the maintenance analysis procedures," on page 273.

Power problems might be associated with any of the following reasons:

- A problem occurred during the installation of a SAN Volume Controller node
- The power switch failed to turn on the node
- The power switch failed to turn off the node
- Another MAP sent you here

Procedure

1. **Are you here because the node is not powered on?**

   NO Go to step 10 on page 283.

   YES Go to step 2.

2. (from step 1)

   Is the power LED on the operator-information panel continuously illuminated? Figure 38 on page 280 shows the location of the power LED **1** on the operator-information panel.

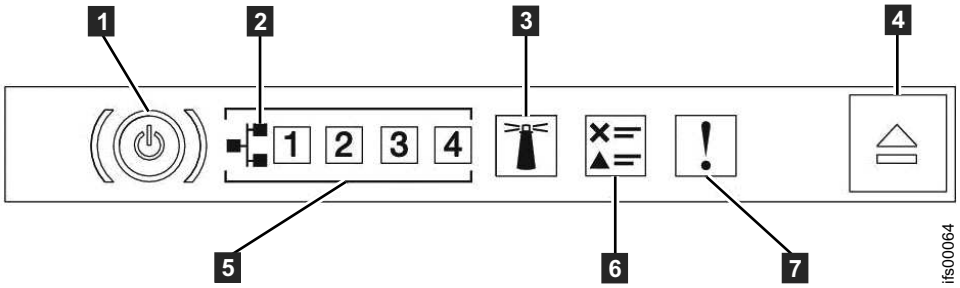

1 **Power** button and power LED (green)

Figure 38. Power LED on the SAN Volume Controller 2145-DH8

NO Go to step 3.

YES The node is powered on correctly. Reassess the symptoms and return to MAP 5000: Start or go to MAP 5700: Repair verification to verify the correct operation.

3. (from step 2 on page 279)

   Is the power LED on the operator-information panel flashing approximately four times per second?

   NO Go to step 4.

   YES The node is turned off and is not ready to be turned on. Wait until the power LED flashes at a rate of once per second, then go to step 5.

   If this behavior persists for more than 3 minutes, complete the following procedure:

   a. Remove all input power from the SAN Volume Controller node by removing the power supply from the back of the node. See "Removing a SAN Volume Controller 2145-DH8 power supply" when you are removing the power cords from the node.
   b. Wait 1 minute and then verify that all power LEDs on the node are extinguished.
   c. Reinsert the power supply.
   d. Wait for the flashing rate of the power LED to slow down to one flash per second. Go to step 5.
   e. If the power LED keeps flashing at a rate of four flashes per second for a second time, replace the parts in the following sequence:
      - System board

   Verify the repair by continuing with MAP 5700: Repair verification.

4. (from step 3)

   Is the Power LED on the operator-information panel flashing once per second?

   YES The node is in standby mode. Input power is present. Go to step 5.

   NO Go to step 6 on page 281.

5. (from step 3 and step 4)

   Press **Power** on the operator-information panel of the node.

   Is the Power LED on the operator-information panel illuminated a solid green?

NO Verify that the operator-information panel cable is correctly seated at both ends.

If the node still fails to power on, replace parts in the following sequence:

a. Operator-information panel assembly

b. System board

Verify the repair by continuing with MAP 5700: Repair verification.

YES The power LED on the operator-information panel shows that the node successfully powered on. Verify the correct operation by continuing with MAP 5700: Repair verification.

6. (from step 4 on page 280)

   Is the rear panel power LED on or flashing?Figure 39 shows the location of the power LED **1** on the SAN Volume Controller 2145-DH8.

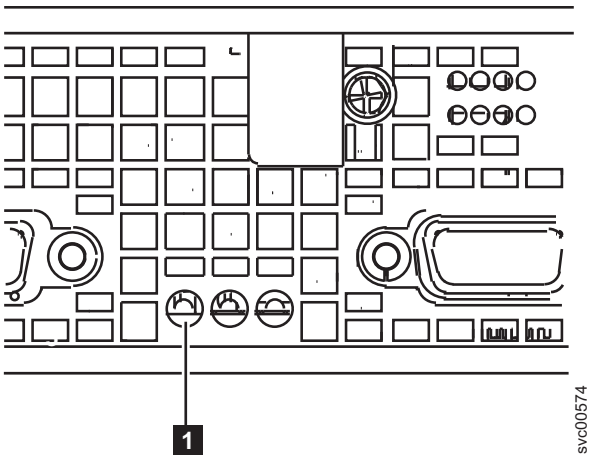

Figure 39. Power LED indicator on the rear panel of the SAN Volume Controller 2145-DH8

NO Go to step 7.

YES The operator-information panel is failing.

Verify that the operator-information panel cable is seated on the system board.

If the node still fails to power on, replace parts in the following sequence:

a. Operator-information panel assembly

b. System board

7. (from step 6)

   Are the ac LED indicators on the rear of the power supply assemblies illuminated? Figure 40 on page 282 shows the location of the ac LED **1**, the dc LED **2**, and the power-supply error LED **3** on the rear of the power supply assembly that is on the rear panel of the SAN Volume Controller 2145-DH8.

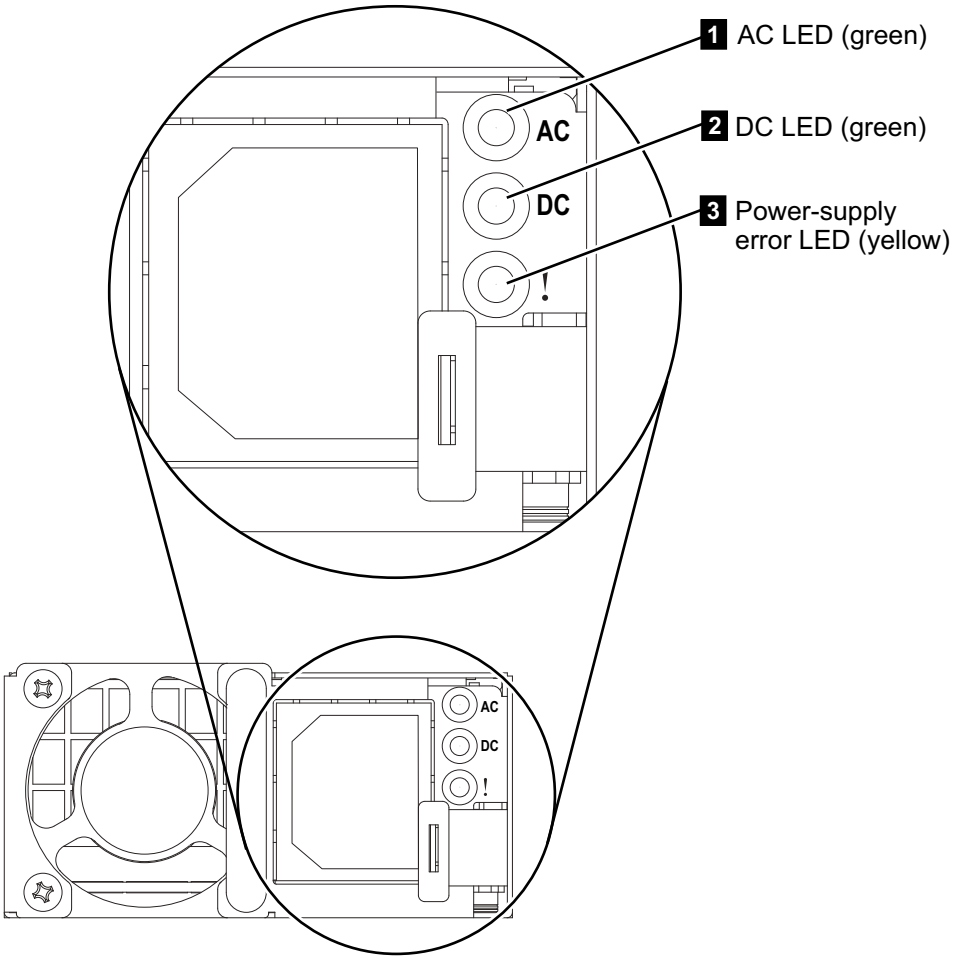

Figure 40. AC, dc, and power-supply error LED indicators on the rear panel of the SAN Volume Controller 2145-DH8

NO Verify that the input power cable or cables are securely connected at both ends and show no sign of damage; replace damaged cables. If the node still fails to power on, replace the specified parts that are based on the SAN Volume Controller model type.

Replace the SAN Volume Controller 2145-DH8 parts in the following sequence:

a. Power supply 750 W

YES Go to step 8.

8. (from step 7 on page 281)

Is the power-supply error LED on the rear of the SAN Volume Controller 2145-DH8 power supply illuminated? Figure 40 shows the location of the power-supply error LED **3** .

YES Replace the power supply unit.

NO Go to step 9

9. (from step 8)

Are the dc LED indicators on the rear of the power supply assemblies illuminated?

NO Replace the SAN Volume Controller 2145-DH8 parts in the following sequence:

a. Power supply 750 W

b. System board

YES Verify that the operator-information panel cable is correctly seated at both ends. If the node still fails to power on, replace parts in the following sequence:

a. Operator-information panel

b. Cable, signal

c. System board

Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

10. (from step 1 on page 279)

The node does not power off immediately when the power button is pressed. When the node fully boots, the node powers-off under the control of the SAN Volume Controller software. The power-off operation can take up to 5 minutes to complete.

Is the power LED on the operator-information panel flashing approximately four times per second?

NO Go to step 11.

YES Wait for the node to power off. If the node fails to power off after 5 minutes, go to step 11.

11. (from step 10)

Attention: Turning off the node by any means other than using the management GUI might cause a loss of data in the node cache. If you are performing concurrent maintenance, this node must be deleted from the system before you proceed. Ask the customer to delete the node from the system now. If they are unable to delete the node, call your support center for assistance before you proceed.
The node cannot be turned off either because of a software fault or a hardware failure. Press and hold the power button. The node can turn off within 5 seconds.

Did the node turn off?

NO Determine whether you are using an Advanced Configuration and Power Interface (ACPI) or a non-ACPI operating system.

If you are using a non-ACPI operating system, complete the following steps:

Press **Ctrl+Alt+Delete**.

Turn off the server by pressing and holding **Power** for 5 seconds.

Restart the server.

If the server fails POST and pressing **Power** does not work, disconnect the power cord for 20 seconds

Reconnect the power cord and restart the server.

If the problem remains or if you are using an ACPI-aware operating system, suspect the system board.

Go to step 12 on page 284

YES Go to step 12.

12. (from step 11 on page 283)

    Press the power button to turn on the node.

    Did the node turn on and boot correctly?

    NO Go to "MAP 5000: Start" on page 273 to resolve the problem.

    YES Go to step 13.

13. (from step 12)

    The node probably suffered a software failure. Memory dump data might be captured that helps resolve the problem. Call your support center for assistance.

MAP 5350: Powering off a node

MAP 5350: Powering off a node helps you power off a single node to complete a service action without disrupting host access to volumes.

Before you begin

If the solution is set up correctly, powering off a single node does not disrupt the normal operation of a system. A system has nodes in pairs called I/O groups. An I/O group continues to handle I/O to the disks it manages with only a single node that is powered on. However, performance degrades and resilience to error is reduced.

Be careful when you power off a system node to impact the system no more than necessary.

Note: If you do not follow the procedures that are outlined here, your application hosts might lose access to their data or they might lose data in the worst case.

You can use the following preferred methods to power off a node that is a member of a system and not offline:

1. Use the **Power off** option in the management GUI or in the service assistant interface.
2. Use the CLI command **stopsystem –node** *name*.

It is preferable to use either the management GUI or the command-line interface (CLI) to power off a node. These methods provide a controlled handover to the partner node and provide better resilience to other faults in the system.

Only if a node is offline or not a member of a system must you power it off using the power button.

About this task

To provide the least disruption when you power off a node, all of the following conditions must apply:

- The other node in the I/O group is powered on and active in the system.
- The other node in the I/O group has SAN Fibre Channel connections to all hosts and disk controllers that are managed by the I/O group.
- All volumes that are handled by this I/O group are online.
- Host multipathing is online to the other node in the I/O group.

In some circumstances, the reason you power off the node might make these conditions impossible. For instance, if you replace a failed Fibre Channel adapter, volumes do not show an online status. Use your judgment to decide that it is safe to proceed when a condition is not met. Always check with the system administrator before you proceed with power off as that might disrupt I/O access. The system administrator might prefer to wait for a more suitable time or suspend host applications.

To ensure a smooth restart, a node must save data structures that it cannot re-create to its local, internal disk drive. The amount of data the node saves to local disk can be high, so this operation might take several minutes. Do not attempt to interrupt the controlled power off.

Attention: The following actions do not allow the node to save data to its local disk. Therefore, do not power off a node by using the following methods:

- Holding down the power button on the node (unless it is a SAN Volume Controller 2145-SV1).

  When you press and release the power button, the node indicates this action to the software so the node can write its data to local disk before the node powers off.

  When you hold down the power button, the hardware interprets this action as an emergency power off indication and shuts down immediately. The hardware does not save the data to a local disk before you power down. The emergency power off occurs approximately 4 seconds after you press and hold down the power button.
- Pressing the reset button on the light path diagnostics panel.

Important: If you power off a SAN Volume Controller 2145-DH8 node and might not power it back on the same day, follow these steps to prevent the batteries from being discharged too much while the node is connected to power but not powered on:

1. Pull both batteries out of the node. Keep them out until you're ready to power on the node.
2. Push the batteries in just before you press the power button to power on the node.

If you disconnect the power from a SAN Volume Controller 2145-DH8 node and might not reconnect power to it again within the next 24 hours, follow these steps to prevent the batteries from being discharged too much while the node is not connected to power:

1. After both power cords are disconnected from the node, pull both batteries out of the node. This step completely turns off the battery backplane.
2. Push the batteries back in again.

Using the management GUI to power off a system

Use the management GUI to power off a system.

Procedure

To use the management GUI to power off a system, complete the following steps:

1. Start the management GUI for the system that you are servicing.
2. Select **Monitoring** > **System**.

If the nodes to power off are shown as `Offline`, the nodes are not participating in the system. In such circumstances, use the power button on the offline nodes to power off the nodes.

If the nodes to power off are shown as `Online`, powering off the nodes can result in their dependent volumes also going offline:

a. Select the node and click **Show Dependent Volumes**.
b. Make sure the status of each volume in the I/O group is `Online`. You might need to view more than one page. You might need to view more than one page.

   If any volumes are `Degraded`, only one node in the I/O is processing I/O requests for that volume. If that node is powered off, it impacts all the hosts that are submitting I/O requests to the degraded volume.

   If any volumes are degraded and you believe that it might be because the partner node in the I/O group is powered off recently, wait until a refresh of the screen shows all volumes online. All the volumes must be online within 30 minutes of the partner node that is being powered off.

   Note: After you wait 30 minutes, if you have a degraded volume and all of the associated nodes and MDisks are online, contact support for assistance.

   Ensure that all volumes that are used by hosts are online before you continue.
c. If possible, check that all hosts that access volumes that are managed by this I/O group are able to fail over to use paths that are provided by the other node in the group.

   Complete this check by using the multipathing device driver software of the host system. Commands to use differ, depending on the multipathing device driver that is being used.

   If you use the System Storage Multipath Subsystem Device Driver (SDD), the command to query paths is **`datapath query device`**.

   It can take some time for the multipathing device drivers to rediscover paths after a node is powered on. If you are unable to check on the host that all paths to both nodes in the I/O group are available, do not power off a node within 30 minutes of the partner node that is being powered on or you might lose access to the volume.
d. If you decide that it is okay to continue with powering off the nodes, select the node to power off and click **Shut Down System**.
e. Click **OK**. If the node that you select is the last remaining node that provides access to a volume, for example a node that contains flash drives with unmirrored volumes, the Shutting Down a Node-Force panel is displayed with a list of volumes that go offline if the node is shut down.
f. Check that no host applications access the volumes that are going offline. Continue with the shutdown only if the loss of access to these volumes is acceptable. To continue with shutting down the node, click **Force Shutdown**.

What to do next

During the shutdown procedure, the node saves its data structures to its local disk and destages all write data that is held in cache to the SAN disks. Such processing can take several minutes.

At the end of this processing, the system powers off.

Using the system CLI to power off a node

Use the command-line interface (CLI) to power off a node.

Procedure

1. Issue the **lsnode** CLI command to display a list of nodes in the system and their properties. Find the node to shut down and write down the name of its I/O group. Confirm that the other node in the I/O group is `online`.

   ```
   lsnode -delim :

   id:name:UPS_serial_number:WWNN:status:IO_group_id: IO_group_name:config_node:
   UPS_unique_id
   1:group1node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202381001C0D18D8
   2:group1node2:10L3ANF:5005076801000009:online:0:io_grp0:no:202381001C0D1796
   3:group2node1:10L3ASH:5005076801000001:online:1:io_grp1:no:202381001C0D18D8
   4:group2node2:10L3ANF:50050768010000F4:online:1:io_grp1:no:202381001C0D1796
   ```

   If the node to power off is shown as `Offline`, the node is not participating in the system and is not processing I/O requests. In such circumstances, use the power button on the node to power off the node.

   If the node to power off is shown as `Online`, but the other node in the I/O group is not online, powering off the node impacts all hosts that are submitting I/O requests to the volumes that are managed by the I/O group. Ensure that the other node in the I/O group is online before you continue.

2. Issue the **lsdependentvdisks** CLI command to list the volumes that depend on the status of a specified node.

   ```
   lsdependentvdisks group1node1

   vdisk_id       vdisk_name
   0              vdisk0
   1              vdisk1
   ```

   If the node goes offline or is removed from the system, the dependent volumes also go offline. Before you take a node offline or remove it from the system, you can use the command to ensure that you do not lose access to any volumes.

3. If you decide that it is okay to continue powering off the node, enter the **stopsystem –node <name>** CLI command to power off the node. Use the **–node** parameter to avoid powering off the whole system:

   ```
   stopsystem –node group1node1
   Are you sure that you want to continue with the shut down? yes
   ```

   Note: To shut down the node even though there are dependent volumes, add the **-force** parameter to the **stopsystem** command. The **force** parameter forces continuation of the command even though any node-dependent volumes will be taken offline. Use the **force** parameter with caution; access to data on node-dependent volumes will be lost.

   During the shutdown procedure, the node saves its data structures to its local disk and destages all write data that is held in the cache to the SAN disks, which can take several minutes.

   At the end of this process, the node powers off.

Using the system power control button

Do not use the power control button to power off a node unless an emergency exists or another procedure directs you to do so.

Before you begin

With this method, you cannot check the system status from the front panel, so you cannot tell if the power off is liable to cause excessive disruption to the system. Instead, use the management GUI or the CLI commands, described in the previous topics to power off an active node.

About this task

If you must use this method, notice in Figure 41 and Figure 42 that each model type has a power control button **1** on the front.

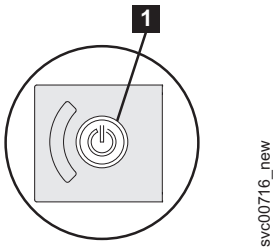

Figure 41. Power control button on the SAN Volume Controller 2145-DH8 model

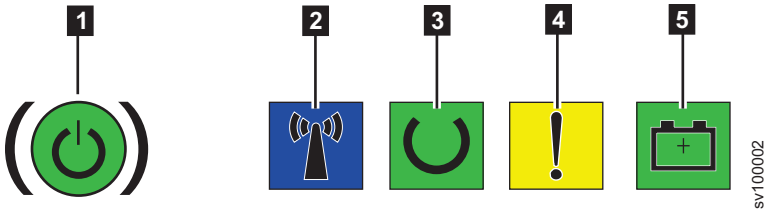

Figure 42. Power control button and LED lights on the SAN Volume Controller 2145-SV1 model

- **1** Power-control button and power-on LED
- **2** Identify LED
- **3** Node status LED
- **4** Node fault LED
- **5** Battery status LED

When you determine it is safe to do so, press and immediately release the power button. On models other than the 2145-DH8 and 2145-SV1, the front panel display changes to display `Powering Off` and displays a progress bar.

Note: The 2145-DH8 and 2145-SV1 do not have a front panel display, but status LED **2**, **3**, **4**, and **5** in Figure 42 all turn off, and the power-on LED **1** goes from on to flashing.

Results

The node saves its data structures to disk while it is powering off. The power off process can take up to 5 minutes.

When a node is powered off by using the power button (or because of a power failure), the partner node in its I/O group immediately stops using its cache for new write data and destages any write data already in its cache to the SAN-attached disks.

The destaging duration depends on the speed and utilization of the disk controllers. The time to complete is less than 15 minutes, but it might be longer. If data is waiting to be written to a disk that is offline, the destaging cannot complete.

A node that powers off and restarts while its partner node continues to process I/O might not be able to become an active member of the I/O group immediately. The node must wait until the partner node completes destaging the cache.

If the partner node powers off during this period, access to the SAN storage that is managed by this I/O group is lost. If one of the nodes in the I/O group is unable to service any I/O, volumes that are managed by that I/O group have a status of `Degraded`. For example, if the partner node in the I/O group is still flushing its write cache, it has a status of `Degraded`.

MAP 5500: Ethernet

MAP 5500: Ethernet helps you solve problems that occur on the system Ethernet connections.

Before you begin

Note: The management GUI or service assistant GUI must be used to see the status of the Ethernet ports from the software's point of view.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 10, “Using the maintenance analysis procedures,” on page 273.

If you encounter problems with the 10 Gbps Ethernet feature, see “MAP 5550: 10G Ethernet and Fibre Channel over Ethernet personality enabled adapter port” on page 292.

If you encounter problems with the 25 Gbps Ethernet ports, see 25 Gbps Ethernet link failures.

You might have been sent here for one of the following reasons:

- A problem occurred during the installation of a system and the Ethernet checks failed.
- Another MAP sent you here.
- The customer needs immediate access to the system by using an alternate configuration node. See “Defining an alternate configuration node” on page 292.

About this task

Complete the following steps:

Procedure

1. **Is any node in the system reporting error code 805?**

   YES Go to step 6 on page 290.

NO Go to step 2.

2. **Is the system reporting error 1400 in the event log?**

   YES Go to step 4.

   NO Go to step 3.

3. **Are you experiencing Ethernet performance issues?**

   YES Go to step 9 on page 291.

   NO Go to step 10 on page 291.

4. (from step 2) **On all nodes, complete the following actions:**
   a. Check `Ethernet port 1`.
   b. If `Ethernet port 1` shows `link offline`, record this port as one that requires fixing.
   c. If the system is configured with two Ethernet cables per node, check `Ethernet port 2` and repeat the previous step.
   d. Go to step 5.

5. (from step 4) **Are any Ethernet ports that have cables attached to them reporting `link offline`?**

   YES Go to step 6.

   NO Go to step 10 on page 291.

6. (from step 5) **Do the system nodes have one or two cables connected?**

   One Go to step 7.

   Two Go to step 8.

7. (from step 6) **Complete the following actions:**
   a. Plug the Ethernet cable from that node into the Ethernet port 2 from a different node, as shown in Figure 43.
   b. If the Ethernet link light is illuminated when the cable is plugged into Ethernet port 2 of the other node, replace the system board of the original node.

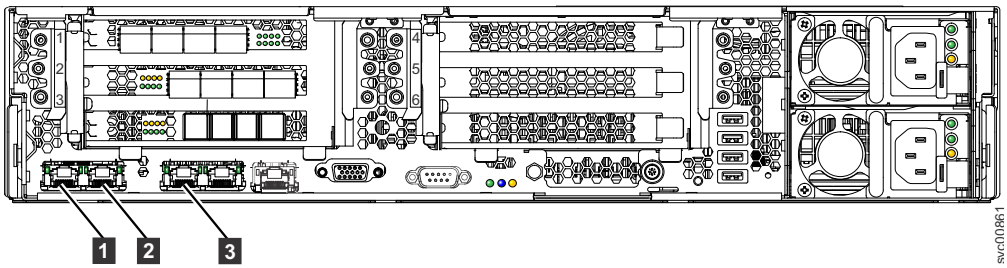

Figure 43. Ethernet ports on the rear of the SAN Volume Controller 2145-DH8

1 1 Gbps Ethernet port 1
2 1 Gbps Ethernet port 2
3 1 Gbps Ethernet port 3

   c. If the Ethernet link light does not illuminate, check the Ethernet switch or hub port and cable to resolve the problem.
   d. Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

8. (from step 5 or step 6) **Complete the following actions:**

a. Plug the Ethernet cable from that node into another device, for example, the SSPC.
b. If the Ethernet link light is illuminated when the cable is plugged into the other Ethernet device, replace the system board of the original node.
c. If the Ethernet link light does not illuminate, check the Ethernet switch/hub port and cable to resolve the problem.
d. Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

9. (from step 3 on page 290) **Complete the following actions:**
   a. Check all Speed port 1 and Speed port 2 panels for the speed and duplex settings. The format is: <Speed>/<Duplex>.
      1) Check `Speed 1`.
      2) If `Speed 1` shows `link offline`, record this port as one that requires fixing.
      3) If the system is configured with two Ethernet cables per node, check `Speed 2` and repeat the previous step.
   b. Check that the system port negotiated at the highest speed available on the switch. All nodes have gigabit Ethernet network ports.
   c. If the Duplex setting is `half`, complete the following steps:
      1) There is a known problem with gigabit Ethernet when one side of the link is set to a fixed speed and duplex and the other side is set to autonegotiate. The problem can cause the fixed side of the link to run at full duplex and the negotiated side of the link to run at half duplex. The duplex mismatch can cause significant Ethernet performance degradation.
      2) If the switch is set to full duplex, set the switch to autonegotiate to prevent the problem described previously.
      3) If the switch is set to half duplex, set it to autonegotiate to allow the link to run at the higher bandwidth available on the full duplex link.
   d. If none of the above are true, call your support center for assistance.

10. (from step 2 on page 290)

    A previously reported fault with the Ethernet interface is no longer present. A problem with the Ethernet might have been fixed, or there might be an intermittent problem. Check with the customer to determine that the Ethernet interface has not been intentionally disconnected. Also, check that there is no recent history of fixed Ethernet problems with other components of the Ethernet network.

    Is the Ethernet failure explained by the previous checks?

    NO There might be an intermittent Ethernet error. Complete these steps in the following sequence until the problem is resolved:
    a. Use the Ethernet hub problem determination procedure to check for and resolve an Ethernet network connection problem. If you resolve a problem, continue with "MAP 5700: Repair verification" on page 300.
    b. Determine whether similar Ethernet connection problems occurred recently on this node. If they have, replace the system board.
    c. Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

    YES Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

Defining an alternate configuration node

A situation can arise where the customer needs immediate access to the system by using an alternate configuration node.

About this task

If all Ethernet connections to the configuration node failed, the system is unable to report failure conditions, and the management GUI is unable to access the system to complete administrative or service tasks. If this is the case and the customer needs immediate access to the system, you can make the system use an alternate configuration node by using the service assistant GUI. The service assistant is accessed via the technician port.

Note: If the system has no front panel display such as on SAN Volume Controller 2145-DH8, use the service assistant GUI. The service assistant is accessed through the technician port.

If only one node is reporting `Node Error 805`, complete the following steps:

Procedure

1. Press and release the power button on the node that is reporting `Node Error 805`.
2. When `Powering off` is displayed, press the power button again.
3. `Restarting` is displayed.

Results

The system selects a new configuration node. The management GUI is able to access the system again.

MAP 5550: 10G Ethernet and Fibre Channel over Ethernet personality enabled adapter port

MAP 5550: 10G Ethernet helps you solve problems that occur on a node with 10G Ethernet capability, and Fibre Channel over Ethernet personality enabled.

Before you begin

Note: The service assistant GUI might be used if there is no front panel display, for example on the SAN Volume Controller 2145-DH8.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 10, "Using the maintenance analysis procedures," on page 273.

This MAP applies to system models with the 10G Ethernet feature installed. Be sure that you know which model you are using before you start this procedure. To determine which model you are working with, look for the label that identifies the model type on the front of the node. Check that the 10G Ethernet adapter is installed and that an optical cable is attached to each port.

If you experience a problem with error code 805, go to "MAP 5500: Ethernet" on page 289.

If you experience a problem with error code 703 or 723, go to "Optical link failures" on page 242.

You might be sent here for one of the following reasons:

- A problem occurred during the installation of a system and the Ethernet checks fail.
- Another MAP sent you to this location.

About this task

Perform the following steps:

Procedure

1. **Is node error 720 or 721 displayed on the front panel of the affected node or is service error code 1072 shown in the event log?**

   YES Go to step 11 on page 294.

   NO Go to step 2.

2. (from step 1) **Perform the following actions from the front panel of the affected node:**

   a. Press and release the up or down button until `Ethernet` is shown.

   b. Press and release the left or right button until `Ethernet port 3` is shown.

   Was Ethernet port 3 found?

   No Go to step 11 on page 294

   Yes Go to step 3

3. (from step 2) **Perform the following actions from the front panel of the affected node:**

   a. Press and release the up or down button until `Ethernet` is shown.

   b. Press and release the up or down button until `Ethernet port 3` is shown.

   c. Record if the second line of the display shows `Link offline`, `Link online`, or `Not configured`.

   d. Press and release the up or down button until `Ethernet port 4` is shown.

   e. Record if the second line of the display shows `Link offline`, `Link online`, or `Not configured`.

   f. Go to step 4.

4. (from step 3) **What was the state of the 10G Ethernet ports that were seen in step 3?**

   Both ports show `Link online`
   The 10G link is working now. Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

   One or more ports show `Link offline`
   Go to step 5 on page 294.

   One or more ports show `Not configured`
   For information about the port configuration, see the CLI command **`cfgportip`** description in the SAN Volume Controller Information Center for iSCSI.

   For Fibre Channel over Ethernet information, see the CLI command **`lsportfc`** description in the SAN Volume Controller Information

Center. This command provides connection properties and status to help determine whether the Fibre Channel over Ethernet is a part of a correctly configured VLAN.

5. (from step 4 on page 293) **Is the amber 10G Ethernet link LED off for the offline port?**

   YES Go to step 6

   NO The physical link is operational. The problem might be with the system configuration. See the configuration topic "iSCSI configuration details" and "Fibre Channel over Ethernet configuration details" in the SAN Volume Controller Information Center.

6. (from step 5) **Perform the following actions:**
   a. Check that the 10G Ethernet ports are connected to a 10G Ethernet fabric.
   b. Check that the 10G Ethernet fabric is configured.
   c. Pull out the small form-factor pluggable (SFP) transceiver and plug it back in.
   d. Pull out the optical cable and plug it back in
   e. Clean contacts with a small blast of air, if available.
   f. Go to step 7.

7. (from step 6) **Did the amber link LED light?**

   YES The physical link is operational. Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

   NO Go to step 8.

8. (from step 7) Swap the 10G SFPs in port 3 and port 4, but keep the optical cables connected to the same port.

   Is the amber link LED on the other port off now?

   YES Go to step 10.

   NO Go to step 9.

9. (from step 8) Swap the 10G Ethernet optical cables in port 3 and port 4.

   Observe how the amber link LED changes. Swap the cables back.

   Did the amber link LED on the other port go off?

   YES Check the 10G Ethernet optical link and fabric that is connected to the port that now has the amber LED off. The problem is associated with the cable. The problem is either in the optical cable or the Ethernet switch. Check that the Ethernet switch shows that the port is operational. If it does not show that the port is operational, replace the optical cable. Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

   NO Go to step 11.

10. (from step 8) **Perform the following actions:**
    a. Replace the SFP that now has the amber link LED off.
    b. Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

11. (from steps 1 on page 293, 2 on page 293, and 9) **Have you already removed and replaced the 10G Ethernet adapter?**

    YES Go to step 12 on page 295.

    NO Perform the following actions:

a. Remove and replace the 10G Ethernet adapter.
b. Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

12. (from steps 11 on page 294) **Replace the 10G Ethernet adapter with a new one.**
    a. Replace the 10G Ethernet adapter.
    b. Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

MAP 5600: Fibre Channel

MAP 5600: Fibre Channel helps you to solve problems that occur on the system Fibre Channel ports.

Before you begin

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 10, "Using the maintenance analysis procedures," on page 273.

This MAP applies to all system models. Be sure that you know which model you are using before you start this procedure. To determine which model you are working with, look for the label that identifies the model type on the front of the node.

You might be sent here for one of the following reasons:

- A problem occurred during the installation of a system and the Fibre Channel checks failed
- Another MAP sent you here

About this task

Complete the following steps to solve problems that are caused by the Fibre Channel ports. You can use the technician port on the system to access the service assistant.

Procedure

1. **Are you trying to resolve a Fibre Channel port speed problem?**

   NO Go to step 2.

   YES Go to step 11 on page 299.

2. (from step 1) Display the Fibre Channel port 1 status on the service assistant GUI.

   Is the service assistant GUI on the system showing Fibre Channel port-1 active?

   NO A Fibre Channel port is not working correctly. Check the port status on the service assistant GUI.

   - **Inactive:** The port is operational but cannot access the Fibre Channel fabric. The Fibre Channel adapter is not configured correctly; the Fibre Channel small form-factor pluggable (SFP) transceiver failed; the Fibre Channel cable that is either failed or is not installed; or the device at the other end of the cable failed. Make a note of port-1. Go to step 7 on page 298.

- **Failed:** The port is not operational because of a hardware failure. Make a note of port-1. Go to step 9 on page 298.
- **Not installed:** This port is not installed. Make a note of port-1. Go to step 10 on page 298.

YES Press and release the right button to display Fibre Channel port-2. Go to step 3.

3. (from step 2 on page 295)

   Is the service assistant GUI on the system showing Fibre Channel port-2 active?

   NO A Fibre Channel port is not working correctly. Check the port status.
   - **Inactive:** The port is operational but cannot access the Fibre Channel fabric. The Fibre Channel adapter is not configured correctly; the Fibre Channel small form-factor pluggable (SFP) transceiver failed; the Fibre Channel cable that is either failed or is not installed; or the device at the other end of the cable failed. Make a note of port-2. Go to step 7 on page 298.
   - **Failed:** The port is not operational because of a hardware failure. Make a note of port-2. Go to step 9 on page 298.
   - **Not installed:** This port is not installed. Make a note of port-2. Go to step 10 on page 298.

   YES Go to step 4.

4. (from step 3)

   Is the service assistant GUI on the system showing Fibre Channel port-3 active?

   NO A Fibre Channel port is not working correctly. Check the port status.
   - **Inactive:** The port is operational but cannot access the Fibre Channel fabric. The Fibre Channel adapter is not configured correctly; the Fibre Channel small form-factor pluggable (SFP) transceiver failed; the Fibre Channel cable that is either failed or is not installed; or the device at the other end of the cable failed. Make a note of port-3. Go to step 7 on page 298.
   - **Failed:** The port is not operational because of a hardware failure. Make a note of port-3. Go to step 9 on page 298.
   - **Not installed:** This port is not installed. Make a note of port-3. Go to step 10 on page 298.

   YES Go to step 5.

5. (from step 4)

   Is the service assistant GUI on the system showing Fibre Channel port-4 active?

   NO A Fibre Channel port is not working correctly. Check the port status.
   - **Inactive:** The port is operational but cannot access the Fibre Channel fabric. The Fibre Channel adapter is not configured correctly; the Fibre Channel small form-factor pluggable (SFP) transceiver failed; the Fibre Channel cable that is either failed or is not installed; or the device at the other end of the cable failed. Make a note of port-4. Go to step 7 on page 298.
   - **Failed:** The port is not operational because of a hardware failure. Make a note of port-4. Go to step 8 on page 298.

- **Not installed:** This port is not installed. Make a note of port-4. Go to step 10 on page 298.

YES If there are more than four Fibre Channel ports on the node, repeat step 5 on page 296 for each additional Fibre Channel port that uses the service assistant.

Go to step 6.

6. (from step 5 on page 296)

   A previously reported fault with a Fibre Channel port is no longer being shown. A problem with the SAN Fibre Channel fabric might be fixed or there might be an intermittent problem.

   Check with the customer to see whether any Fibre Channel ports are disconnected or if any component of the SAN Fibre Channel fabric failed and was recently fixed.

   Is the Fibre Channel port failure explained by the previous checks?

   NO There might be an intermittent Fibre Channel error.

   a. Use the SAN problem determination procedure to check for and resolve any Fibre Channel fabric connection problems. If you resolve a problem, continue with "MAP 5700: Repair verification" on page 300.
   b. Check whether similar Fibre Channel errors occurred recently on the same port on this system node. If they have, replace the Fibre Channel cable, unless it was replaced.
   c. Replace the Fibre Channel SFP transceiver, unless it was replaced.

      Note: System 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. See the "Removing and replacing the Fibre Channel SFP transceiver on a node" documentation to find out how to replace an SFP transceiver.
   d. Replace the Fibre Channel adapter assembly that is shown in Table 77.

Table 77. Fibre Channel assemblies

Node	Adapter assembly
SAN Volume Controller 2145-DH8 port 1, 2, 3, or 4 (slot 1 mandatory; the first FC adapter)	Four-port Fibre Channel adapter
SAN Volume Controller 2145-DH8 port 5, 6, 7, or 8 (slot 2 optional; the second FC adapter)	Four-port Fibre Channel adapter
SAN Volume Controller 2145-DH8 port 9, 10, 11, or 12 (slot 5 optional; the third FC adapter)	Four-port Fibre Channel adapter

   e. Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

YES Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

7. (from steps 2 on page 295, 3 on page 296, 4 on page 296, and 5 on page 296)

   The port noted on the system is showing a status of inactive. For certain models, this inactive status might occur when the Fibre Channel speed is not set correctly.

8. (from step 7)

   The noted port on the system displays a status of inactive. If the noted port still displays a status of inactive, replace the parts that are associated with the noted port until the problem is fixed in the following order:

   a. Fibre Channel cables from the system to Fibre Channel network.
   b. Faulty Fibre Channel fabric connections, particularly the SFP transceiver at the Fibre Channel switch. Use the SAN problem determination procedure to resolve any Fibre Channel fabric connection problem.
   c. System Fibre Channel SFP transceiver.

      Note: System nodes are supported by both longwave SFPs and shortwave SFPs. You must replace an SFP 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 provide a suitable replacement. Removing the wrong SFP transceiver might result in loss of data access. See the "Removing and replacing the Fibre Channel SFP transceiver on a system node" documentation to find out how to replace an SFP transceiver.

   d. Replace the Fibre Channel adapter assembly as shown in Table 77 on page 297.
   e. Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

9. (from steps 2 on page 295, 3 on page 296, 4 on page 296, and 5 on page 296)

   The noted port on the system displays a status of failed. Verify that the Fibre Channel cables that connect the system nodes to the switches are securely connected. Replace the parts that are associated with the noted port until the problem is fixed in the following order:

   a. Fibre Channel SFP transceiver.

      Note: system 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. See the "Removing and replacing the Fibre Channel SFP transceiver on a node" documentation to find out how to replace an SFP transceiver.

   b. Replace the Fibre Channel adapter assembly as shown in Table 77 on page 297.
   c. Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

10. (from steps 2 on page 295, 3 on page 296, 4 on page 296, and 5 on page 296)

    The noted port on the system displays a status of not installed. If you replaced the Fibre Channel adapter, make sure that it is installed correctly. If you replaced any other system board components, make sure that the Fibre Channel adapter was not disturbed.

Is the Fibre Channel adapter failure explained by the previous checks?

NO

a. Replace the Fibre Channel adapter assembly as shown in Table 77 on page 297.

b. If the problem is not fixed, replace the Fibre Channel connection hardware in the order that is shown in Table 78.

Table 78. System Fibre Channel adapter connection hardware

Node	Adapter connection hardware
SAN Volume Controller 2145-DH8 port 1-8	1. PCI Express® Riser card assembly 1 2. System board
SAN Volume Controller 2145-DH8 port 9-12	1. PCI Express® Riser card assembly 2 2. System board

c. Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

YES Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

11. (from step 1 on page 295)

If the operating speed is lower than the operating speed that is supported by the switch, a high number of link errors are being detected.

To display the current speed of the link, see http://www-01.ibm.com/support/knowledgecenter/STPVGU_7.6.0/com.ibm.storage.svc.console.760.doc/svc_svcdetfibrenetspeed_23eeaf.html

Is the port operating at lower than the expected speed?

NO Repeat the check with the other Fibre Channel ports until the failing port is located. If no failing port is located, the problem no longer exists. Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

YES Perform the following steps:

a. Check the routing of the Fibre Channel cable to ensure that no damage exists and that the cable route contains no tight bends (no less than a 3-inch radius). Either reroute or replace the Fibre Channel cable.

b. Remove the Fibre Channel cable for 2 seconds and then reinsert it to force the Fibre Channel adapter to renegotiate its operating speed.

c. Recheck the speed of the Fibre Channel port. If it is now correct, the problem is resolved. Otherwise, the problem might be caused by one of the following conditions:

- Four-port Fibre Channel HBA
- System SFP transceiver
- Fibre Channel switch gigabit interface converter (GBIC) or SFP transceiver
- Fibre Channel switch

Recheck the speed after you change any component until the problem is resolved and then verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

MAP 5700: Repair verification

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you exchange for new FRUs, or repair actions that are completed solve all the problems on the SAN Volume Controller .

Before you begin

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 10, "Using the maintenance analysis procedures," on page 273.

You might have been sent here because you performed a repair and want to confirm that no other problems exist on the machine.

Procedure

1. **Are the Power LEDs on all the nodes on?** For more information about this LED, see "Power LED" on page 27.

   NO Go to "MAP 5000: Start" on page 273.

   YES Go to step 2.

2. (from step 1)

   Are all the nodes displaying `Cluster:` or is the node status LED on?

   NO Go to "MAP 5000: Start" on page 273.

   YES Go to step 3.

3. (from step 2)

   Using the SAN Volume Controller application for the system you repair, check the status of all configured managed disks (MDisks).

   Do all MDisks have a status of `online`?

   NO If any MDisks have a status of `offline`, repair the MDisks. Use the problem determination procedure for the disk controller to repair the MDisk faults before you return to this MAP.

   If any MDisks have a status of `degraded paths` or `degraded ports`, repair any storage area network (SAN) and MDisk faults before you return to this MAP.

   If any MDisks show a status of `excluded`, include MDisks before you return to this MAP.

   Go to "MAP 5000: Start" on page 273.

   YES Go to step 4.

4. (from step 3)

   Using the SAN Volume Controller application on the repaired system, check the status of all configured volumes. **Do all volumes have a status of `online`?**

   NO Go to step 5.

   YES Go to step 6 on page 301.

5. (from step 4)

   Following a repair of the SAN Volume Controller , a number of volumes are showing a status of offline. Volumes are held offline if SAN Volume Controller cannot confirm the integrity of the data. The volumes might be the target of a copy that did not complete, or cache write data that was not written back to disk might be lost. Determine why the volume is offline. If the volume was the

target of a copy that did not complete, you can start the copy again. Otherwise, write data might not be written to the disk, so its state cannot be verified. Your site procedures determine how data is restored to a known state.

To bring the volume online, you must move all the offline disks to the recovery I/O group and then move them back to an active I/O group.

Go to "MAP 5000: Start" on page 273.

6. (from step 4 on page 300)

   You successfully repair the SAN Volume Controller .

MAP 5800: Light path

MAP 5800: Light path helps you to solve hardware problems that prevent the SAN Volume Controller 2145-DH8 from booting.

Before you begin

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 10, "Using the maintenance analysis procedures," on page 273.

You might be sent here because of the following situations:

- The Error LED on the operator-information panel is on or flashing.
- Another MAP sent you here:
  - "Light path for SAN Volume Controller 2145-DH8"

Light path for SAN Volume Controller 2145-DH8

Light path diagnostics is a system of LEDs on top of the operator-information panel of the SAN Volume Controller 2145-DH8 node, which leads you to the failed component.

About this task

When an error occurs, LEDs are lit along the front of the operator-information panel, the light path diagnostics panel, then on the failed component. By viewing the LEDs in a particular order, you can often identify the source of the error.

LEDs that are lit to indicate an error, remain lit when the server is turned off, if the node is connected to an operating power supply.

Ensure that the node is turned on, and then resolve any hardware errors that are indicated by the Error LED and light path LEDs:

Procedure

1. **Is the System error LED 7 , shown in Figure 44 on page 302, on the SAN Volume Controller 2145-DH8 operator-information panel on or flashing?**

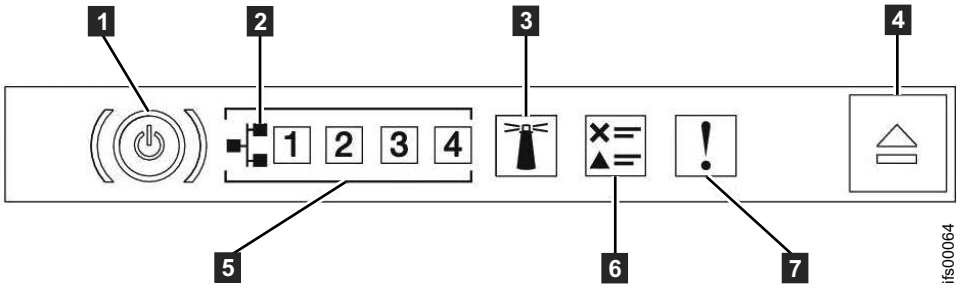

Figure 44. SAN Volume Controller 2145-DH8 operator-information panel

1. **Power control** button and LED.
2. Ethernet LED.
3. **Locator** button and LED.
4. Release latch.
5. Ethernet activity LEDs.
6. Check log LED.
7. System error LED.

NO Reassess your symptoms and return to "MAP 5000: Start" on page 273.

YES Go to step 2.

2. (from step 1 on page 301)

   Press the release latch, as shown in Figure 45, and open the light path diagnostics panel, which is shown in Figure 46 on page 303.

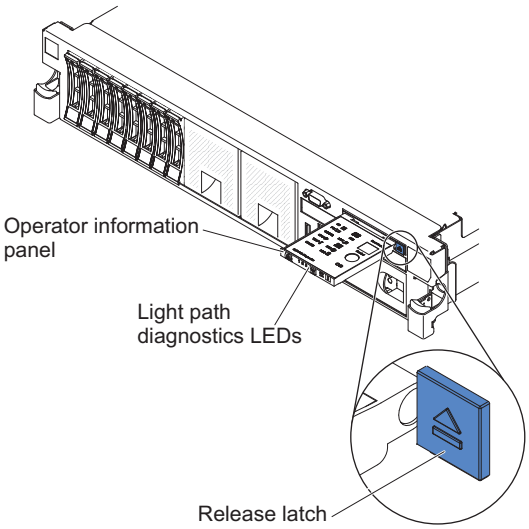

Figure 45. Press the release latch

Are one or more LEDs on the light path diagnostics panel on or flashing?

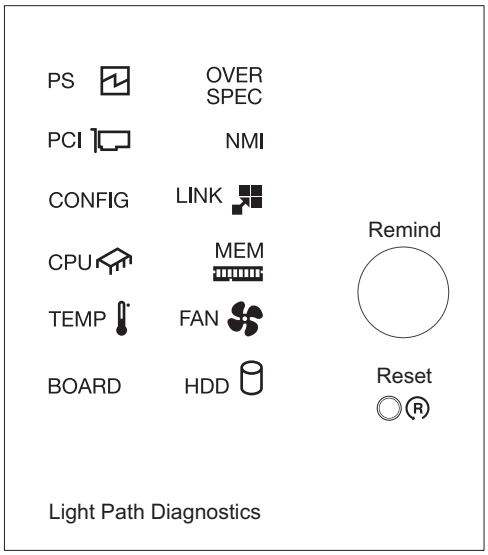

Figure 46. SAN Volume Controller 2145-DH8 light path diagnostics panel

NO Verify that the operator-information panel cable is correctly seated at both ends. If the error LED is still illuminated but no LEDs are illuminated on the light path diagnostics panel, replace parts in the following sequence:

a. Operator-information panel
b. System board

Verify the repair by continuing with "MAP 5700: Repair verification" on page 300.

YES See Table 79 on page 304 and complete the action that is specified for the specific light-path-diagnostics LEDs. Then, go to step 3 on page 308. Some actions require that you observe the state of LEDs on the system board. Figure 47 on page 304 shows the location of the system board LEDs. The fan LEDs are located next to each FAN. To view the LEDs, complete the following actions:

a. Before you turn off the node, ensure that its data is mirrored and synchronized.
b. Identify and label all the cables that are attached to the node so that they can be replaced in the same port. Remove the node from the rack and place it on a flat, static-protective surface. For more information, see "Removing the node from a rack".
c. Remove the top cover.
d. See Table 79 on page 304 and complete the action that is specified for the specific light-path-diagnostics LEDs. Then, go to step 3 on page 308.

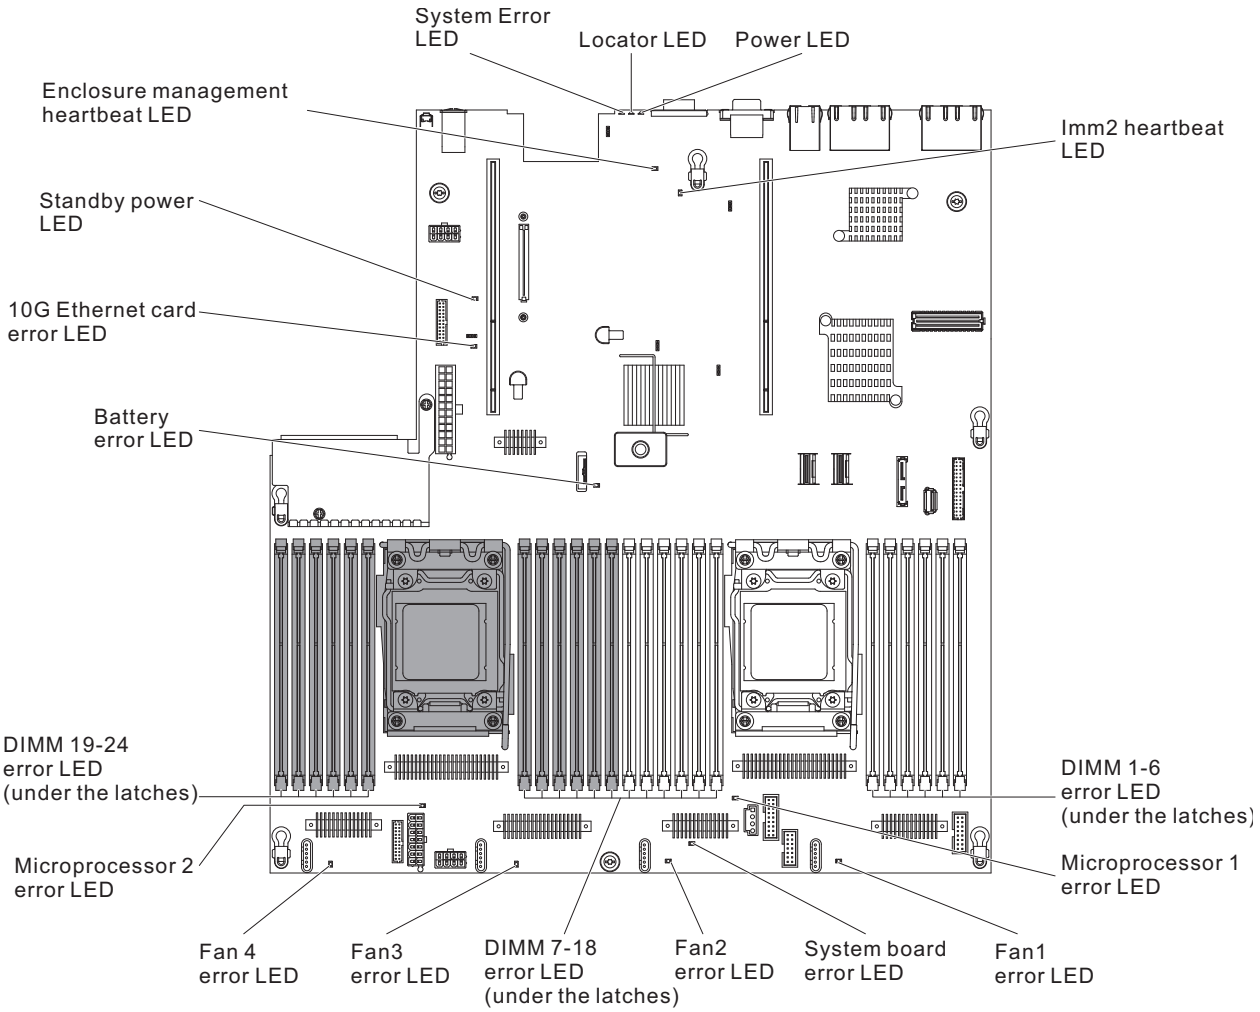

Figure 47. SAN Volume Controller 2145-DH8 system board LEDs.

Table 79. Diagnostics panel LEDs

LED	Description	Action
The Error log or Check log LED operator-information panel	An error occurs and cannot be isolated without completing certain procedures.	1. Plug in the VGA screen and the USB keyboard. 2. Check the IMM2 system event log and the system-error log for information about the error. 3. Save the log if necessary and clear the log afterward.
System-error LED operator-information panel	An error occurred.	1. Check the light-path-diagnostics LEDs and follow the instructions. 2. Check the IMM2 system event log and the system-error log for information about the error. 3. Save the log if necessary and clear the log afterward.

Table 79. Diagnostics panel LEDs (continued)

LED	Description	Action
PS	When only the PS LED is lit, a power supply failed.	The system might detect a power supply error. Complete the following steps to correct the problem: 1. Check the power-supply with a lit yellow LED. 2. Make sure that the power supplies are seated correctly and plugged in a good AC outlet. 3. Remove a power supply to isolate the failed power supply. 4. Make sure that both power supplies installed in the server are of the same AC input voltage. 5. Replace the failed power supply.
	PS + CONFIG When both the PS and CONFIG LEDs are lit, the power supply configuration is not valid.	If the PS LED and the CONFIG LED are lit, the system logs an invalid power configuration error. Make sure that both power supplies installed in the node are of the same rating or wattage.
OVER SPEC	The system consumption reaches the power supply over-current protection point or the power supplies are damaged.	1. If the power rail (A, B, C, D, E, F, G, and H) error was not detected, complete the following steps: a. Use the IBM Systems Energy Estimator to determine the current system power consumption. For more information, go to the following website: https://www-947.ibm.com/systems/support/tools/estimator/energy/index.html b. Replace the failed power supply. 2. If the power rail (A, B, C, D, E, F, G, and H) error was also detected, follow actions that are listed in MAP 5040: Power.
PCI	An error occurred on a PCI bus or on the system board. Another LED is lit next to a failing PCI slot.	1. Check the riser-card LEDs, the ServeRAID error LED, and the dual-port network adapter error LED to identify the component that caused the error. 2. Check the system-error log for information about the error. 3. If you cannot isolate the failing component by using the LEDs and the information in the system-error log, remove one component at a time. Then, restart the server after each component is removed. 4. Replace the following components, in the order that is shown, restarting the server each time: • PCI riser cards • ServeRAID adapter • Network adapter • (Trained technician only) System board. 5. If the failure remains, contact your IBM service representative.

Table 79. Diagnostics panel LEDs (continued)

LED	Description	Action
NMI	A nonmaskable interrupt occurred, or the **NMI** button was pressed.	1. Check the system-error log for information about the error. 2. Restart the server.
CONFIG	CONFIG + PS An invalid power configuration error occurred.	If the CONFIG LED and the PS LED are lit, the system logs an invalid power configuration error. Make sure that both power supplies installed in the server are of the same rating or wattage.
	CONFIG + CPU A hardware configuration error occurred.	If the CONFIG LED and the CPU LED are lit, complete the following steps to correct the problem: 1. Check the microprocessors that were installed to make sure that they are compatible with each other. 2. (Trained technician only) Replace the incompatible microprocessor. 3. Check the system-error logs for information about the error. Replace any component that is identified in the error log.
	CONFIG + MEM A hardware configuration error occurred.	If the CONFIG LED and the MEM LED are lit, check the system-event log in the Setup utility or IMM2 error messages.
	CONFIG + PCI A hardware configuration error occurred.	If the CONFIG LED and the PCI LED are lit, check the system-error logs for information about the error. Replace any component that is identified in the error log.
	CONFIG + HDD A disk drive error occurred.	If the CONFIG LED and the HDD LED are lit, check the system-error logs for information about the error. Replace any component that is identified in the error log.
LINK	Reserved.	

Table 79. Diagnostics panel LEDs (continued)

LED	Description	Action
CPU	When only the CPU LED is lit, a microprocessor failed. When both the CPU and CONFIG LEDs are lit, the microprocessor configuration is invalid.	1. If the CONFIG LED is not lit, a microprocessor failure occurs, complete the following steps: a. (Trained technician only) Make sure that the failing microprocessor and its heat sink, which are indicated by a lit LED on the system board, are installed correctly. b. (Trained technician only) Replace the failing microprocessor. c. For more information, contact your IBM service representative. 2. If the CONFIG LED and the CPU LED are lit, the system logs an invalid microprocessor configuration error. Complete the following steps to correct the problem: a. Check recently installed microprocessors to ensure that they are compatible with each other. b. (Trained technician only) Replace any incompatible microprocessor. c. Check the system-error logs for information about the error. Replace any component that is identified in the error log.
MEM	When only the MEM LED is lit, a memory error occurs.	**Note:** Note: Each time that you install or remove a DIMM, you must disconnect the node from the power source; then, wait 10 seconds before you restart the server. If the CONFIG LED is not lit, the system might detect a memory error. Complete the following steps to correct the problem: 1. Update the node firmware. 2. Reseat or swap the DIMMs with lit LED. 3. Check the system-event log in the Setup utility or IMM error messages. 4. Replace the failing DIMM.
	MEM + CONFIG When both the MEM and CONFIG LEDs are lit, the memory configuration is not valid.	If the MEM LED and the CONFIG LED are lit, check the system-event log in the Setup utility or IMM2 error messages.
TEMP	The system or the system component temperature exceeded a threshold level. A failing fan can cause the TEMP LED to be lit.	1. Make sure that the heat sink is seated correctly. 2. Determine whether a fan failed and replace the fan if necessary. 3. Make sure that the room temperature is not too high. See the environment requirements for the server temperature information. 4. Make sure that the air vents are not blocked. 5. Make sure that the heat sink or the fan on the adapter, or any other network adapter is seated correctly. If the fan failed, replace it. 6. For more information, contact your IBM service representative.

Table 79. Diagnostics panel LEDs (continued)

LED	Description	Action
FAN	A fan is either failed, operating too slowly, or is removed. The TEMP LED might also be lit.	1. Check whether your node is installed with the dual-port network adapter. If yes, make sure that your node compiles with the configuration with four fans installed. 2. Reseat the failing fan, which is indicated by a lit LED near the fan connector on the system board. 3. Replace the failing fan.
BOARD	An error occurred on the system board or the system battery.	1. Check the LEDs on the system board to identify the component that caused the error. The BOARD LED can be lit due to any of the following reasons: • Battery • (Trained technician only) System board 2. Check the system-error log for information about the error. 3. Replace the failing component.
HDD	A hard disk drive that is failed or is missing.	1. Check the LEDs on the hard disk drives for the drive with a lit status LED and reseat the hard disk drive. 2. Reseat the hard disk drive backplane. 3. If the error remains, replace the following components one at a time, in the order that is listed, restarting the server after each: a. Replace the hard disk drive. b. Replace the hard disk drive backplane. 4. If the problem remains, contact your IBM service representative.

3. Continue with "MAP 5700: Repair verification" on page 300 to verify the correct operation.

Chapter 11. iSCSI performance analysis and tuning

This procedure provides a solution for Internet Small Computer Systems Interface (iSCSI) host performance problems while connected to a 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 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 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 system 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 to be effective. Send a ping packet to be delivered without fragmentation to verify that the network supports jumbo frames. 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 system 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 system 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 system. system 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 must not be blank.
      - For IP address type IPv6, check the `vlan_6` field in the **lsportip** output. It must not be blank.

- If the `vlan` and `vlan_6` fields are blank, use Configuring VLAN for iSCSI to set the VLAN for the IP type.

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 must be `yes`.
   - For IP address type IPv6, check the `host_6` field in the **`lsportip`** output. It must be `yes`.
   - If the `host` and `host_6` fields are not `yes`, use the **`cfgportip`** CLI command to set the host flag for the IP type .

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.

   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.

Appendix A. Accessibility features for the system

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 system:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. HTML documents are tested by using JAWS version 15.0.
- This product uses standard Windows navigation keys.
- Interfaces are commonly used by screen readers.
- Keys are discernible by touch, but do not activate just by touching them.
- Industry-standard devices, ports, and connectors.
- You can attach alternative input and output devices.

The system 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 for operations and to initiate menu actions that can also be done through mouse actions. You can go to the system online documentation from the keyboard by using the keyboard shortcuts for your browser or screen-reader software. See your browser or screen-reader software Help for a list of keyboard shortcuts 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.

Appendix B. Where to find the Statement of Limited Warranty

The Statement of Limited Warranty is available in both hardcopy format and in the SAN Volume Controller IBM Knowledge Center.

The *Statement of Limited Warranty* is included (in hardcopy form) with your product. It can also be ordered from IBM (see Table 2 on page x for the part number).

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.

Linux and the Linux logo is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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.

Product support statement

If you have an operating system, Hypervisor, platform or host attachment card in your environment, check the IBM System Storage Interoperation Center (SSIC) to confirm the support status for this product.

SSIC can be found at http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss.

Homologation statement

This product may not be certified in your country for connection by any means whatsoever to interfaces of public telecommunications networks. Further certification may be required by law prior to making any such connection. Contact an IBM representative or reseller for any questions.

Electromagnetic compatibility notices

The following Class A statements apply to IBM products and their features unless designated as electromagnetic compatibility (EMC) Class B in the feature information.

When attaching a monitor to the equipment, you must use the designated monitor cable and any interference suppression devices that are supplied with the monitor.

Canada Notice

CAN ICES-3 (A)/NMB-3(A)

European Community and Morocco Notice

This product is in conformity with the protection requirements of Directive 2014/30/EU of the European Parliament and of the Council on the harmonization of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

This product may cause interference if used in residential areas. Such use must be avoided unless the user takes special measures to reduce electromagnetic emissions to prevent interference to the reception of radio and television broadcasts.

Warning: This equipment is compliant with Class A of CISPR 32. In a residential environment this equipment may cause radio interference.

Germany Notice

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2014/30/EU zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaatenund hält die Grenzwerte der EN 55032 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung von IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung von IBM gesteckt/eingebaut werden.

EN 55032 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden: “Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen.”

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten
Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG).” Dies ist die Umsetzung der EU-Richtlinie 2014/30/EU in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC Richtlinie 2014/30/EU) für Geräte der Klasse A

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

Verantwortlich für die Einhaltung der EMV-Vorschriften ist der Hersteller:

International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:

IBM Deutschland GmbH
Technical Relations Europe, Abteilung M456
IBM-Allee 1, 71139 Ehningen, Germany
Tel: +49 800 225 5426
e-mail: Halloibm@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55032 Klasse A.

Japan Electronics and Information Technology Industries Association (JEITA) Notice

（一社）電子情報技術産業協会　高調波電流抑制対策実施
要領に基づく定格入力電力値： Knowledge Centerの各製品の
仕様ページ参照

This statement applies to products less than or equal to 20 A per phase.

高調波電流規格 JIS C 61000-3-2 適合品

This statement applies to products greater than 20 A, single phase.

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
- 回路分類：６（単相、ＰＦＣ回路付）
- 換算係数：０

This statement applies to products greater than 20 A per phase, three-phase.

高調波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。
- 回路分類：５（３相、ＰＦＣ回路付）
- 換算係数：０

Japan Voluntary Control Council for Interference (VCCI) Notice

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磁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

Korea Notice

이 기기는 업무용 환경에서 사용할 목적으로 적합성평가를 받은 기기로서 가정용 환경에서 사용하는 경우 전파간섭의 우려가 있습니다.

People's Republic of China Notice

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。

Russia Notice

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать
радиопомехи, для снижения которых необходимы
дополнительные меры

rusemi

Taiwan Notice

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。 taiemi

IBM Taiwan Contact Information:

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888 f2c00790

United States Federal Communications Commission (FCC) Notice

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors, or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions:

(1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

Index

Numerics

A

B

C

D

E

F

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W

Printed in USA