

IBM System Storage SAN Volume Controller



# Troubleshooting Guide

*Version 5.1.0*



IBM System Storage SAN Volume Controller



# Troubleshooting Guide

*Version 5.1.0*

**Note:**

Before using this information and the product it supports, read the information in **Notices**.

This edition applies to the IBM System Storage SAN Volume Controller, release 5.1.0, and to all subsequent releases and modifications until otherwise indicated in new editions. This edition replaces GC27-2227-01.

© **Copyright International Business Machines Corporation 2003, 2009.**

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

---

# Contents

<b>Tables</b> . . . . .	vii
<b>About this guide</b> . . . . .	ix
Who should use this guide . . . . .	ix
Summary of changes for GG27-2227-01 and GG27-2227-02 SAN Volume Controller Troubleshooting Guide. . . . .	ix
Emphasis . . . . .	x
SAN Volume Controller library and related publications. . . . .	x
How to order IBM publications . . . . .	xiv
How to send your comments. . . . .	xiv
<b>Chapter 1. SAN Volume Controller overview</b> . . . . .	1
Clusters . . . . .	2
Configuration node . . . . .	3
Configuration node addressing . . . . .	3
Cluster IP failover . . . . .	4
SAN fabric overview . . . . .	5
Service mode overview . . . . .	8
<b>Chapter 2. Introducing the SAN Volume Controller hardware components</b> . . . . .	9
SAN Volume Controller nodes . . . . .	9
SAN Volume Controller front panel controls and indicators . . . . .	9
SAN Volume Controller operator-information panel. . . . .	14
SAN Volume Controller rear-panel indicators and connectors . . . . .	18
Fibre-channel port numbers and worldwide port names . . . . .	32
Requirements for the SAN Volume Controller environment . . . . .	34
Redundant ac-power switch . . . . .	43
Redundant ac-power environment requirements. . . . .	44
Cabling of redundant ac-power switch (example) . . . . .	45
Uninterruptible power supply . . . . .	47
2145 UPS-1U . . . . .	47
Uninterruptible power-supply environment requirements . . . . .	53
Defining the SAN Volume Controller FRUs. . . . .	54
SAN Volume Controller FRUs . . . . .	54
Redundant ac-power switch FRUs. . . . .	61
Uninterruptible power supply FRUs . . . . .	61
<b>Chapter 3. Using the SAN Volume Controller Console and CLI</b> . . . . .	63
Accessing the SAN Volume Controller CLI . . . . .	65
Viewing the node status . . . . .	66
Checking the status of the node using the CLI . . . . .	67
Viewing the state of the fibre-channel ports . . . . .	67
Checking the status of the fibre-channel ports using the CLI . . . . .	68
Viewing the state of Ethernet ports . . . . .	69
Deleting a node from a cluster using the SAN Volume Controller Console . . . . .	69
Deleting a node from a cluster using the CLI . . . . .	71
Adding nodes to a cluster . . . . .	73
Adding a node to the cluster using the CLI. . . . .	76
Shutting down a node . . . . .	77
Shutting down the cluster . . . . .	77
Shutting down a cluster. . . . .	78
Shutting down a cluster using the CLI . . . . .	79
Listing node-dependent VDisks using the CLI . . . . .	80

Viewing the VDisk status . . . . .	81
Verifying and repairing mirrored VDisk copies . . . . .	81
Viewing the progress of VDisk copy verification . . . . .	82
Validating and repairing mirrored VDisk copies using the CLI . . . . .	82
Repairing offline space-efficient VDIs . . . . .	84
Viewing the progress of space-efficient VDisk copy repair . . . . .	84
Repairing a space-efficient VDisk using the CLI . . . . .	85
Recovering from offline VDIs . . . . .	85
Recovering from offline VDIs using the CLI . . . . .	86
Recovering a node and returning it to the cluster using the CLI . . . . .	87
Recovering VDIs . . . . .	88
Recovering offline VDIs using the CLI . . . . .	88
Moving offline VDIs to their original I/O group using the CLI . . . . .	89
Creating MDisk groups . . . . .	90
Adding MDIs to MDisk groups . . . . .	90
Viewing the MDisk status . . . . .	91
Discovering MDIs . . . . .	92
Creating MDisk groups using the CLI . . . . .	93
Viewing the fibre-channel fabric connections . . . . .	101
Displaying all controllers logged in to a node . . . . .	103
Displaying all active devices . . . . .	103
Changing the fibre-channel network speed . . . . .	104
Determining the fibre-channel port speed . . . . .	104
Changing the fibre-channel port speed for a node not in a cluster . . . . .	106
Changing the fibre-channel port speed for a node in a cluster . . . . .	106
Checking disk controller status using the CLI . . . . .	106
Determining the failing enclosure or disk controller using the CLI . . . . .	107
Replacing a faulty node with a spare node . . . . .	108
Replacing a faulty node in the cluster using the CLI . . . . .	112
Updating license settings . . . . .	117
Viewing and updating license settings using the CLI . . . . .	117
Viewing the license settings log . . . . .	118
Collecting SSD dump files using the CLI . . . . .	118
Displaying and saving log and dump files . . . . .	118
Listing log or dump files using the CLI . . . . .	119
Dumping cluster configurations . . . . .	120
Viewing the feature log using the CLI . . . . .	121
Analyzing the error log . . . . .	121
Analyzing the error log using the CLI . . . . .	122
Marking errors as fixed . . . . .	122
Running the cluster maintenance procedure . . . . .	123
Using directed maintenance procedures . . . . .	123
<b>Chapter 4. Viewing the vital product data . . . . .</b>	<b>127</b>
Viewing the vital product data . . . . .	127
Viewing cluster properties . . . . .	127
Displaying the vital product data using the CLI . . . . .	128
Displaying node properties using the CLI . . . . .	128
Displaying cluster properties using the CLI . . . . .	129
Fields for the node VPD . . . . .	130
Fields for the cluster VPD . . . . .	136
<b>Chapter 5. Using the front panel of the SAN Volume Controller . . . . .</b>	<b>139</b>
Boot progress indicator . . . . .	139
Boot failed . . . . .	139
Charging . . . . .	140

I

Error codes . . . . .	140
Hardware boot . . . . .	141
Node rescue request . . . . .	141
Power failure . . . . .	142
Powering off . . . . .	142
Recovering . . . . .	142
Restarting . . . . .	143
Shutting down . . . . .	143
Validate WWNN? . . . . .	143
SAN Volume Controller menu options . . . . .	144
Cluster options . . . . .	146
Node options . . . . .	150
Version option . . . . .	158
Ethernet option . . . . .	158
Fibre-channel port-1 through 4 option . . . . .	159
Select language? option . . . . .	159
Service mode . . . . .	160
Using the power control for the SAN Volume Controller node . . . . .	165
<b>Chapter 6. Diagnosing problems . . . . .</b>	<b>167</b>
Error reporting . . . . .	168
Power-on self-test . . . . .	168
Understanding the error event log . . . . .	169
Managing the error event log . . . . .	170
Viewing the error event log . . . . .	170
Describing the fields in the error event log . . . . .	173
Event notifications . . . . .	176
Call Home and inventory e-mail information . . . . .	178
Understanding the error codes . . . . .	179
Using the error code tables . . . . .	179
Event codes . . . . .	180
SCSI error reporting . . . . .	186
Object types . . . . .	189
Error Codes . . . . .	190
Determining a hardware boot failure . . . . .	199
Understanding the boot codes . . . . .	199
Performing the node rescue . . . . .	217
Understanding the node rescue codes . . . . .	218
Understanding the node error codes . . . . .	223
Understanding the create cluster error codes . . . . .	238
Understanding the cluster recovery codes . . . . .	239
Understanding cluster error codes . . . . .	239
SAN problem determination . . . . .	348
Fibre-channel link failures . . . . .	349
Ethernet iSCSI host link problems . . . . .	349
Servicing storage systems . . . . .	350
Troubleshooting the SAN Volume Controller Console . . . . .	351
Checking connections to the cluster when the status is No Contact . . . . .	351
Cannot access the SAN Volume Controller Console login screen . . . . .	352
Cannot access the cluster from the SAN Volume Controller Console . . . . .	352
Troubleshooting unexpected shutdowns of the SAN Volume Controller Console . . . . .	352
WebSphere Application Server service keeps stopping . . . . .	353
Websphere and CIM Logging . . . . .	354
<b>Chapter 7. Using the maintenance analysis procedures . . . . .</b>	<b>357</b>

MAP 5000: Start . . . . .	357
MAP 5050: Power 2145-CF8, 2145-8G4, 2145-8F4, and 2145-8F2 . . . . .	365
MAP 5060: Power 2145-8A4 . . . . .	372
MAP 5150: 2145 UPS-1U . . . . .	376
MAP 5250: 2145 UPS-1U repair verification . . . . .	381
MAP 5320: Redundant ac power . . . . .	382
MAP 5340: Redundant ac power verification . . . . .	383
MAP 5350: Powering off a SAN Volume Controller node . . . . .	385
Using the SAN Volume Controller Console to power off a node. . . . .	386
Using the SAN Volume Controller CLI to power off a node . . . . .	388
Using the SAN Volume Controller Power control button . . . . .	389
MAP 5400: Front panel . . . . .	390
MAP 5500: Ethernet . . . . .	393
Defining an alternate configuration node . . . . .	396
MAP 5600: Fibre-channel . . . . .	396
MAP 5700: Repair verification . . . . .	404
MAP 5800: Light path . . . . .	406
Light path for SAN Volume Controller 2145-CF8 . . . . .	406
Light path for SAN Volume Controller 2145-8A4 . . . . .	412
Light path for SAN Volume Controller 2145-8G4 . . . . .	415
Light path for SAN Volume Controller 2145-8F2 and SAN Volume Controller 2145-8F4. . . . .	419
MAP 5900: Hardware boot . . . . .	423
MAP 6000: Replace offline SSD . . . . .	428
<b>Appendix. Accessibility.</b> . . . . .	<b>431</b>
<b>Notices</b> . . . . .	<b>433</b>
Trademarks. . . . .	435
Electronic emission notices . . . . .	435
Federal Communications Commission (FCC) statement . . . . .	435
Industry Canada compliance statement . . . . .	436
Avis de conformité à la réglementation d'Industrie Canada . . . . .	436
New Zealand compliance statement. . . . .	436
European Union EMC Directive conformance statement . . . . .	436
Germany compliance statement . . . . .	437
Japanese Voluntary Control Council for Interference (VCCI) statement . . . . .	437
People's Republic of China Class A Electronic Emission Statement . . . . .	438
International Electrotechnical Commission (IEC) statement . . . . .	438
United Kingdom telecommunications requirements . . . . .	438
Korean Class A Electronic Emission Statement . . . . .	438
Taiwan Class A compliance statement . . . . .	438
European Contact Information . . . . .	438
Taiwan Contact Information . . . . .	439
<b>Index</b> . . . . .	<b>441</b>



## Tables

1.	SAN Volume Controller library . . . . .	xii
2.	Other IBM publications . . . . .	xiii
3.	IBM documentation and related Web sites . . . . .	xiii
4.	Link state and activity for bottom LED . . . . .	28
5.	Relative link speed for the top LED . . . . .	29
6.	Actual link speeds . . . . .	29
7.	Actual link speeds . . . . .	29
8.	2145 UPS-1U error indicators . . . . .	49
9.	SAN Volume Controller 2145-CF8 FRU descriptions . . . . .	54
10.	Solid-state drive (SSD) feature FRU descriptions . . . . .	56
11.	2145 UPS-1U FRU descriptions . . . . .	56
12.	SAN Volume Controller 2145-8A4 FRU descriptions . . . . .	57
13.	SAN Volume Controller 2145-8G4 FRU descriptions . . . . .	58
14.	SAN Volume Controller 2145-8F4 FRU descriptions . . . . .	59
15.	SAN Volume Controller 2145-8F2 FRU descriptions . . . . .	60
16.	Uninterruptible power supply FRU descriptions . . . . .	61
17.	Corrective actions when the Availability Status is not OK . . . . .	64
18.	Disk controller attributes for SSDs . . . . .	91
19.	Maximum VDisk capacity by extent size . . . . .	94
20.	Fibre-channel viewing definitions . . . . .	102
21.	Fields for the system board . . . . .	131
22.	Fields for the processors . . . . .	131
23.	Fields for the fans . . . . .	132
24.	Fields that are repeated for each installed memory module . . . . .	132
25.	Fields that are repeated for each adapter card that is installed . . . . .	132
26.	Fields for the direct-attached managed disk . . . . .	133
27.	Fields that are repeated for each SCSI, IDE, SATA, and SAS device that is installed . . . . .	133
28.	Fields that are specific to the node software . . . . .	133
29.	Fields that are provided for the front panel assembly . . . . .	134
30.	Fields that are provided for the Ethernet port . . . . .	134
31.	Fields that are provided for the power supplies in the node . . . . .	134
32.	Fields that are provided for the uninterruptible power supply assembly that is powering the node . . . . .	134
33.	Fields that are provided for the SAS host bus adapter (HBA) . . . . .	135
34.	Fields that are provided for the SAS solid-state drive (SSD) . . . . .	135
35.	Fields that are provided for the small form factor pluggable (SFP) connector . . . . .	136
36.	Fields that are provided for the cluster properties . . . . .	136
37.	Descriptions of log entry icons . . . . .	172
38.	Description of data fields for the error event log . . . . .	173
39.	Description of object types and object IDs for the error event log . . . . .	174
40.	Description of flags for the error-event log . . . . .	174
41.	Reported status for combinations of error-log status flags . . . . .	175
42.	Description error-event types . . . . .	175
43.	SAN Volume Controller notification types . . . . .	176
44.	SAN Volume Controller notification codes and corresponding syslog level codes . . . . .	177
45.	Syslog facility codes and SAN Volume Controller values of user-defined message origin identifiers . . . . .	177
46.	Informational event codes . . . . .	180
47.	Configuration event codes . . . . .	183
48.	SCSI status . . . . .	187
49.	SCSI sense keys, codes, and qualifiers . . . . .	187
50.	Reason codes . . . . .	189
51.	Object types . . . . .	189
52.	Error codes . . . . .	190

I	53. Commands for logging and tracing CIM agent . . . . .	355
	54. 2145 UPS-1U error indicators. . . . .	377
	55. SAN Volume Controller fibre-channel adapter assemblies . . . . .	402
	56. SAN Volume Controller fibre-channel adapter connection hardware . . . . .	402
I	57. Diagnostics panel LED prescribed actions . . . . .	410
	58. SAN Volume Controller 2145-8A4 diagnostics panel LED prescribed actions . . . . .	415
	59. Diagnostics panel LED prescribed actions . . . . .	418
	60. Diagnostics panel LED prescribed actions . . . . .	422

---

## About this guide

This guide describes how to service the IBM® System Storage™ SAN Volume Controller.

The chapters that follow introduce you to the SAN Volume Controller, 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 cluster of nodes through the front panel.

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 uninterruptible power supply, and the IBM System Storage Productivity Center (SSPC) or the master console.

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 357.

**Note:** The *IBM System Storage SAN Volume Controller Troubleshooting Guide* and the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide* were formerly combined in one book that was titled *IBM System Storage SAN Volume Controller Service Guide*.

---

## Who should use this guide

This guide is intended for system administrators or systems services representatives who use and diagnose problems with the SAN Volume Controller, the redundant ac-power switch, the uninterruptible power supply, the IBM System Storage Productivity Center, and the master console.

---

## Summary of changes for GG27-2227-01 and GG27-2227-02 SAN Volume Controller Troubleshooting Guide

The summary of changes provides a list of new and changed information since the last version of the guide.

### New information

This topic describes the changes to this guide since the previous edition, GG27-2227-00. The following sections summarize the changes that have since been implemented from the previous version.

This version includes the following new information:

- Support statements for the SAN Volume Controller 2145-CF8 node and the optional solid-state drive (SSD)
- Support statements for the syslog notification protocol
- Support statements for second cluster IP and service IP addresses
- Support statements for iSCSI 1 Gb support using onboard Ethernet ports
- New error codes
- New event codes

- New information codes

### Changed information

This version includes the following changed information:

- Renaming of the Cache LED to Node Status LED on the front-panel assembly

### Removed information

This version has had the following information removed:

- The glossary from the printable PDF. It remains as a topic in the information center builds and is the last entry in the navigation bar.

---

## Emphasis

Different typefaces are used in this guide to show emphasis.

The following typefaces are used to show emphasis:

<b>Boldface</b>	Text in <b>boldface</b> represents menu items and command names.
<i>Italics</i>	Text in <i>italics</i> 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 cluster.
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.

---

## SAN Volume Controller library and related publications

Product manuals, other publications, and Web sites contain information that relates to SAN Volume Controller.

### SAN Volume Controller Information Center

The IBM System Storage SAN Volume Controller Information Center contains all of the information that is required to install, configure, and manage the SAN Volume Controller. The information center is updated between SAN Volume Controller product releases to provide the most current documentation. The information center is available at the following Web site:

<http://publib.boulder.ibm.com/infocenter/svcic/v3r1m0/index.jsp>

### SAN Volume Controller library

Table 1 on page xi lists and describes the publications that make up the SAN Volume Controller library. Unless otherwise noted, these publications are available in Adobe® portable document format (PDF) from the following Web site:

[www.ibm.com/storage/support/2145](http://www.ibm.com/storage/support/2145)

Table 1. SAN Volume Controller library

Title	Description	Order number
<i>IBM System Storage SAN Volume Controller Planning Guide</i>	This guide introduces the SAN Volume Controller and lists the features that you can order. It also provides guidelines for planning the installation and configuration of the SAN Volume Controller.	GA32-0551
<i>IBM System Storage SAN Volume Controller Model 2145-CF8 Hardware Installation Guide</i>	This guide provides the instructions that the IBM service representative uses to install the hardware for SAN Volume Controller model 2145-CF8.	GC52-1356
<i>IBM System Storage SAN Volume Controller Model 2145-8A4 Hardware Installation Guide</i>	This guide provides the instructions that the IBM service representative uses to install the hardware for SAN Volume Controller model 2145-8A4.	GC27-2219
<i>IBM System Storage SAN Volume Controller Model 2145-8G4 Hardware Installation Guide</i>	This guide provides the instructions that the IBM service representative uses to install the hardware for SAN Volume Controller model 2145-8G4.	GC27-2220
<i>IBM System Storage SAN Volume Controller Models 2145-8F2 and 2145-8F4 Hardware Installation Guide</i>	This guide provides the instructions that the IBM service representative uses to install the hardware for SAN Volume Controller models 2145-8F2 and 2145-8F4.	GC27-2221
<i>IBM System Storage SAN Volume Controller Software Installation and Configuration Guide</i>	This guide provides guidelines for configuring your SAN Volume Controller. Instructions for backing up and restoring the cluster configuration, using and upgrading the SAN Volume Controller Console, using the CLI, upgrading the SAN Volume Controller software, and replacing or adding nodes to a cluster are included.	SC23-6628

Table 1. SAN Volume Controller library (continued)

Title	Description	Order number
<i>IBM System Storage SAN Volume Controller CIM Agent Developer's Guide</i>	This guide describes the concepts of the Common Information Model (CIM) environment. Steps about using the CIM agent object class instances to complete basic storage configuration tasks, establishing new Copy Services relationships, and performing CIM agent maintenance and diagnostic tasks are included.	SC23-6665
<i>IBM System Storage SAN Volume Controller Command-Line Interface User's Guide</i>	This guide describes the commands that you can use from the SAN Volume Controller command-line interface (CLI).	SC26-7903
<i>IBM System Storage SAN Volume Controller Host Attachment Guide</i>	This guide provides guidelines for attaching the SAN Volume Controller to your host system.	SC26-7905
<i>IBM System Storage SAN Volume Controller Troubleshooting Guide</i>	This guide describes the features of each SAN Volume Controller model, explains how to use the front panel, and provides maintenance analysis procedures to help you diagnose and solve problems with the SAN Volume Controller.	GC27-2227
<i>IBM System Storage SAN Volume Controller Hardware Maintenance Guide</i>	This guide provides the instructions that the IBM service representative uses to service the SAN Volume Controller hardware, including the removal and replacement of parts.	GC27-2226
<i>IBM System Storage SAN Volume Controller Master Console Guide</i>	This guide describes how to install, maintain, and service the master console.	GC27-2223
<i>IBM Systems Safety Notices</i>	This guide contains translated caution and danger statements. Each caution and danger statement in the SAN Volume Controller documentation has a number that you can use to locate the corresponding statement in your language in the <i>IBM Systems Safety Notices</i> document.	G229-9054

## Other IBM publications

Table 2 lists IBM publications that contain information related to the SAN Volume Controller.

Table 2. Other IBM publications

Title	Description	Order number
<i>IBM System Storage Productivity Center Introduction and Planning Guide</i>	This guide introduces the IBM System Storage Productivity Center hardware and software.	SC23-8824
<i>Read This First: Installing the IBM System Storage Productivity Center</i>	This guide describes how to install the IBM System Storage Productivity Center hardware.	GI11-8938
<i>IBM System Storage Productivity Center User's Guide</i>	This guide describes how to configure the IBM System Storage Productivity Center software.	SC27-2336
<i>IBM System Storage Multipath Subsystem Device Driver User's Guide</i>	This guide describes the IBM System Storage Multipath Subsystem Device Driver for IBM System Storage products and how to use it with the SAN Volume Controller.	GC52-1309
<i>Implementing the IBM System Storage SAN Volume Controller V4.3</i>	This IBM Redbooks® publication is a detailed technical guide to the IBM System Storage SAN Volume Controller. It provides a high-level overview of storage virtualization and the SAN Volume Controller architecture, discusses implementing and configuring the SAN Volume Controller, tells you how to migrate existing storage to the SAN Volume Controller, and discusses different supported migration activities.	SG24-6423

## IBM documentation and related Web sites

Table 3 lists Web sites that provide publications and other information about the SAN Volume Controller or related products or technologies.

Table 3. IBM documentation and related Web sites

Web site	Address
Support for SAN Volume Controller (2145)	<a href="http://www.ibm.com/storage/support/2145">www.ibm.com/storage/support/2145</a>
Support for IBM System Storage and IBM TotalStorage® products	<a href="http://www.ibm.com/storage/support/">www.ibm.com/storage/support/</a>
IBM Publications Center	<a href="http://www.ibm.com/shop/publications/order/">www.ibm.com/shop/publications/order/</a>
IBM Redbooks publications	<a href="http://www.redbooks.ibm.com/">www.redbooks.ibm.com/</a>

## Related accessibility information

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

[www.adobe.com/support/downloads/main.html](http://www.adobe.com/support/downloads/main.html)

---

## How to order IBM publications

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

The IBM Publications Center offers customized search functions to help you find the publications that you need. Some publications are available for you to view or download at no charge. You can also order publications. The publications center displays prices in your local currency. You can access the IBM Publications Center through the following Web site:

[www.ibm.com/shop/publications/order/](http://www.ibm.com/shop/publications/order/)

---

## How to send your comments

Your feedback is important to help us provide the highest quality information. If you have any comments about this book or any other documentation, you can submit them in one of the following ways:

- E-mail

Submit your comments electronically to the following e-mail address:

[starpubs@us.ibm.com](mailto:starpubs@us.ibm.com)

Be sure to include the name and order number of the book and, if applicable, the specific location of the text you are commenting on, such as a page number or table number.

- Mail

Fill out the Readers' Comments form (RCF) at the back of this book. If the RCF has been removed, you can address your comments to:

International Business Machines Corporation  
RCF Processing Department  
Department 61C  
9032 South Rita Road  
Tucson, Arizona 85775-4401  
U.S.A.



---

## Chapter 1. SAN Volume Controller overview

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

Symmetric virtualization is achieved by creating a pool of managed disks (MDisks) from the attached storage systems. Those storage systems are then mapped to a set of virtual disks (VDisks) for use by attached host systems. System administrators can view and access a common pool of storage on the storage area network (SAN). 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. In a SAN, 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.

### SAN Volume Controller software

The SAN Volume Controller software performs the following functions for the host systems that attach to SAN Volume Controller:

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

The SAN Volume Controller also provides the following functions:

- Large scalable cache
- Copy Services
  - IBM FlashCopy® (point-in-time copy)
  - Metro Mirror (synchronous copy)
  - Global Mirror (asynchronous copy)
  - Data migration
- Space management
  - Mapping that is based on desired performance characteristics
  - Metering of service quality
  - Space-efficient logical volumes (thin provisioning)

### SAN Volume Controller hardware

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

The nodes are always installed in pairs, with a minimum of one and a maximum of four pairs of nodes constituting a *cluster*. Each pair of nodes is known as an *I/O group*. All I/O operations that are managed by the nodes in an I/O group are cached on both nodes.

**Note:** I/O groups take the storage that is presented to the SAN by the storage systems as MDisks and translates the storage into logical disks, known as

VDisks, that are used by applications on the hosts. A node resides in only one I/O group and provides access to the VDisks in that I/O group.

The following nodes are supported in SAN Volume Controller 5.1:

- The new SAN Volume Controller 2145-CF8 node is available for purchase, with up to four of the optional solid-state drives (SSDs).
- The SAN Volume Controller 2145-8A4 node remains available for purchase.
- The SAN Volume Controller 2145-8G4 node is no longer available for purchase, but remains supported in SAN Volume Controller 5.1.
- The SAN Volume Controller 2145-8F4 node is no longer available for purchase, but remains supported in SAN Volume Controller 5.1.
- The SAN Volume Controller 2145-8F2 node is no longer available for purchase, but remains supported in SAN Volume Controller 5.1.

#### **Related reference**

“SAN Volume Controller front panel controls and indicators” on page 9

The controls and indicators are used for power and navigation and to indicate information such as system activity, service and configuration options, service controller failures, and node identification.

“SAN Volume Controller rear-panel indicators and connectors” on page 18

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

“Requirements for the SAN Volume Controller environment” on page 34

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.

---

## **Clusters**

All your configuration, monitoring, and service tasks are performed at the cluster level. Therefore, after configuring your cluster, you can take advantage of the virtualization and the advanced features of the SAN Volume Controller.

A cluster can consist of between two and eight SAN Volume Controller nodes.

All configuration settings are replicated across all nodes in the cluster. Because configuration is performed at the cluster level, management IP addresses are assigned to the cluster instead of to each node. The cluster is configured using the SAN Volume Controller Console, the command-line interface (CLI) or an application developed to access the SAN Volume Controller CIMOM. Each interface accesses the cluster remotely through the Ethernet cluster-management address.

Each node has two Ethernet ports that can be used for management. Ethernet port 1 must be configured and connected on the configuration node. Ethernet port 1 must be connected on all cluster nodes. The use of Ethernet port 2 is optional. At any point in time, only one node in the cluster can operate as the focal point for configuration and monitoring requests. This node is called the *configuration node*. It is the only node that activates the cluster IP addresses. You can use one or more of these addresses to access the cluster through the SAN Volume Controller graphical user interface or the command-line interface (CLI).

Each SAN Volume Controller cluster can have one to four management IP addresses. You can assign up to two IPv4 addresses and up to two IPv6

addresses. When a node has been assigned to a cluster, you can display the cluster IP addresses on the front panel by selecting **Cluster** from the menu.

Each SAN Volume Controller cluster can have optional Small Computer System Interface over Internet Protocol (iSCSI IP) addresses.

**Note:** Management IP addresses that are assigned to a cluster must be different from the iSCSI IP addresses and are used for different purposes. If iSCSI is used, iSCSI addresses are assigned to individual node ports. On the configuration node, a port will have multiple IP addresses active at the same time.

## Configuration node

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

If the configuration node fails, the cluster chooses a new configuration node. This action is called configuration node failover. The new configuration node takes over the cluster IP addresses. Thus you can access the cluster 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 SAN Volume Controller Console.

Figure 1 shows an example cluster containing four nodes. Node 1 has been designated the configuration node. User requests (1) are handled by Node 1.




Figure 1. Configuration node

## Configuration node addressing

At any given time, only one node within a SAN Volume Controller cluster is assigned the cluster IP addresses.

This node then acts as the focal point for all configuration and other requests that are made from the SAN Volume Controller Console 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 cluster determine which node will take on the role of configuration node. The new configuration node binds the cluster IP addresses to its Ethernet ports. It broadcasts this new mapping so that connections to the cluster 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 cluster.

If a device loses its connection to the SAN Volume Controller cluster, it can regenerate the address quickly if the device is on the same subnetwork as the cluster. 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 cluster from a terminal that is on the same subnetwork, and then by starting a secure copy to the device that has lost its connection.

## Cluster IP failover

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

The following changes are performed by the cluster service:

- If software on the failed configuration node is still operational, the software shuts down the cluster IP interfaces. If the software cannot shut down the cluster IP interfaces, the hardware service forces the node to shut down.
- When the cluster IP interfaces shut down, all remaining nodes choose a new node to host the configuration interfaces.
- The new configuration node initializes the configuration daemons, including sshd and httpd, and then binds the cluster IP interfaces to its Ethernet ports.
- The router is configured as the default gateway for the new configuration node.
- The routing tables are established on the new configuration node for the cluster IP addresses. The new configuration node sends five unsolicited address resolution protocol (ARP) packets for each IP address to the local subnet broadcast address. The ARP packets contain the cluster 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. Once 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 into the SAN Volume Controller and starting a secure copy to the affected system. Starting a secure copy forces an update to the ARP cache for all systems connected to the same switch as the affected system.

## Ethernet link failures

If the Ethernet link to the SAN Volume Controller cluster fails because of an event unrelated to the SAN Volume Controller, such as a cable being disconnected or an Ethernet router failure, the SAN Volume Controller does not attempt to failover the configuration node to restore cluster IP access. SAN Volume Controller provides the option for two Ethernet ports, each with its own management IP address, to protect against this type of failure. If you cannot connect through one IP address, attempt to access the cluster through the alternate IP address.

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

## Routing considerations for event notification

SAN Volume Controller supports the following protocols that make outbound connections from the cluster:

- E-mail
- Simple Network Mail Protocol (SNMP)
- Syslog
- Network Time Protocol (NTP)

One or more of these protocols can be configured on the cluster to receive event notifications. When making outbound connections, the SAN Volume Controller uses the following routing decisions:

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

When configuring any of these protocols for event notifications, use these routing decisions to ensure error notification works correctly in the event of a network failure.

---

## 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 using the SAN can communicate only with devices that are included in the same zones that it is in. A SAN Volume Controller cluster requires several distinct types of zones: a cluster zone, host zones, and disk zones. The intercluster zone is optional.

In the host zone, the host systems can identify and address the SAN Volume Controller nodes. You can have more than one host zone and more than one disk zone. The cluster zone contains all ports from all SAN Volume Controller nodes in the cluster, unless you are using a dual-core fabric design. Create one zone for each host fibre-channel port. In a disk zone, the SAN Volume Controller nodes identify the storage systems. Generally, create one zone for each storage system. Host systems cannot operate on the storage systems directly; all data transfer occurs through the SAN Volume Controller nodes. If you are using the Metro Mirror and Global Mirror feature, create a zone with at least one port from each node in each cluster; up to four clusters are supported.

Figure 2 on page 6 shows an example of a host zone. Figure 3 on page 6 shows an example of a cluster zone. Figure 4 on page 7 shows an example of a disk zone.

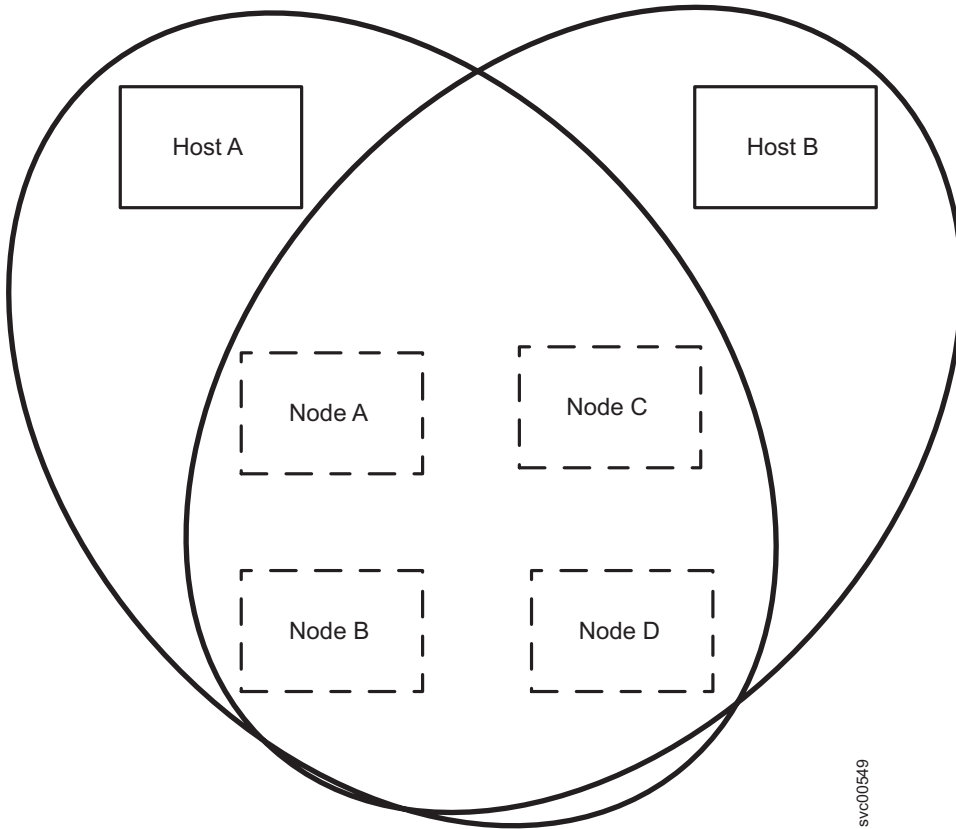


Figure 2. Example of a SAN Volume Controller host zone




Figure 3. Example of a SAN Volume Controller cluster zone




Figure 4. Example of a SAN Volume Controller disk zone

A cluster of SAN Volume Controller nodes is connected to the fibre-channel fabric and presents virtual disks (VDisks) to the host systems. You create these VDisks from units of space within a managed disk (MDisk) group. An MDisk group is a collection of MDisks that are presented by the storage systems (RAID controllers). The MDisk group provides a storage pool. You specify how each group is created, and you can combine MDisks from different manufacturers' controllers in the same MDisk group. However, to optimize the use of resources, ensure that all MDisks in an MDisk group have similar performance characteristics.

**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 SAN Volume Controller nodes is performed through the  
 | SAN. All SAN Volume Controller node configuration and service commands are sent  
 | to the cluster through an Ethernet network.

Each SAN Volume Controller node contains its own vital product data (VPD). Each cluster contains VPD that is common to all the SAN Volume Controller nodes in the cluster, and any system, with the correct access authority, that is connected to the Ethernet network can access this VPD.

---

## Service mode overview

The SAN Volume Controller service mode allows you to access vital product data (VPD), logs, and dump data on the node. It also provides a method of forcing the installation of a different version of software.

There are occasions when a node cannot be managed as part of the cluster. In these instances, the node must be set into service mode. While in service mode, the node does not operate as part of the SAN Volume Controller cluster; therefore, the cluster will be degraded.

### Notes:

1. Use service mode only under the direction of the IBM Support Center.
2. Setting service mode for a node that is in an active cluster might cause data to be lost.
3. Before using service mode, ensure that any faults in the fibre-channel fabric are repaired and any problems that are indicated on the front panel of any node in the cluster are repaired.

The IP address used to access a node that is in server mode is configured as part of the cluster configuration. A service mode IP address must be specified for Ethernet port 1. Optionally, a service mode IP address can also be configured for Ethernet port 2. When in service mode, the node can be managed through any of the Ethernet ports that are configured. Either of the service mode addresses can be configured to use Dynamic Host Configuration Protocol (DHCP), which means that the addresses are not allocated until the node is set into service mode. When in service mode, the IP addresses that are used are displayed through the front panel.

To return the node to normal operation, you can exit service mode through the SAN Volume Controller console GUI, the command-line interface, or by turning the node off and back on.

### Related reference

“Recover cluster navigation” on page 149

The Recover cluster? option is useful if the cluster superuser password has been lost or forgotten.

“Service mode” on page 160

While in service mode, you can use the front panel to view or change a service IPv4 or an IPv6 address. You can also view the version and build level of the SAN Volume Controller software that is installed and active on the node.



## 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

SAN Volume Controller supports five different node types.

The following nodes are supported in SAN Volume Controller 5.1:

- The new SAN Volume Controller 2145-CF8 node is available for purchase, with up to four of the optional solid-state drives (SSDs).
- The SAN Volume Controller 2145-8A4 node remains available for purchase.
- The SAN Volume Controller 2145-8G4 node is no longer available for purchase, but remains supported in SAN Volume Controller 5.1.
- The SAN Volume Controller 2145-8F4 node is no longer available for purchase, but remains supported in SAN Volume Controller 5.1.
- The SAN Volume Controller 2145-8F2 node is no longer available for purchase, but remains supported in SAN Volume Controller 5.1.

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

### SAN Volume Controller front panel controls and indicators

The controls and indicators are used for power and navigation and to indicate information such as system activity, service and configuration options, service controller failures, and node identification.

#### SAN Volume Controller 2145-CF8 controls and indicators

The controls and indicators are used for power and navigation and to indicate information such as system activity, service and configuration options, service controller failures, and node identification.

Figure 5 shows the controls and indicators on the front panel of the SAN Volume Controller 2145-CF8.




Figure 5. SAN Volume Controller 2145-CF8 front panel

- 1 Node-status LED

- | **2** Front-panel display
- | **3** Navigation buttons
- | **4** Operator-information panel
- | **5** Select button
- | **6** Error LED

### SAN Volume Controller 2145-8A4 controls and indicators

The controls and indicators are used for power and navigation and to indicate information such as system activity, service and configuration options, service controller failures, and node identification.

Figure 6 shows the controls and indicators on the front panel of the SAN Volume Controller 2145-8A4.




Figure 6. SAN Volume Controller 2145-8A4 front-panel assembly

- | **1** Operator-information panel
- | **2** Node status LED
- | **3** Front-panel display
- | **4** Navigation buttons
- | **5** Serial number label
- | **6** Select button
- | **7** Node identification label
- | **8** Error LED

### SAN Volume Controller 2145-8G4 controls and indicators

The controls and indicators are used for power and navigation and to indicate information such as system activity, service and configuration options, service controller failures, and node identification.

Figure 7 on page 11 shows the controls and indicators on the front panel of the SAN Volume Controller 2145-8G4.




Figure 7. SAN Volume Controller 2145-8G4 front-panel assembly

- 1** Node status LED
- 2** Front panel display
- 3** Navigation buttons
- 4** Serial number label
- 5** Operator information panel
- 6** Select button
- 7** Node identification label
- 8** Error LED

### SAN Volume Controller 2145-8F4 and SAN Volume Controller 2145-8F2 controls and indicators

The controls and indicators are used for power and navigation and to indicate information such as system activity, service and configuration options, service controller failures, and node identification.

Figure 8 shows the controls and indicators on the front panel of the SAN Volume Controller 2145-8F4 and SAN Volume Controller 2145-8F2.




Figure 8. SAN Volume Controller 2145-8F2 and SAN Volume Controller 2145-8F4 front-panel assembly

- 1** Node status LED
- 2** Front-panel display
- 3** Navigation buttons
- 4** Serial number label

- 5 Operator-information panel
- 6 Select button
- 7 Node identification label
- 8 Error LED

### **Node status LED**

System activity is indicated through the green LED.

The node status LED provides the following system activity indicators:

**Off** The node is not operating as a member of a cluster.

**On** The node is operating as a member of a cluster.

#### **Flashing**

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.

### **Front-panel display**

The front-panel display shows service, configuration, and navigation information.

You can select the language that is displayed on the front panel. The display can show both alphanumeric information and graphical information (progress bars).

The front-panel display shows configuration and service information about the SAN Volume Controller node and the SAN Volume Controller cluster, including the following items:

- Boot progress indicator
- Boot failed
- Charging
- Hardware boot
- Node rescue request
- Power failure
- Powering off
- Recovering
- Restarting
- Shutting down
- Error codes
- Validate WWNN?

#### **Related reference**

Chapter 5, “Using the front panel of the SAN Volume Controller,” on page 139  
The front panel of the SAN Volume Controller has a display, various LEDs, navigation buttons, and a select button that are used when servicing your SAN Volume Controller node.

### **Navigation buttons**

You can use the navigation buttons to move through menus.

There are four navigational buttons that you can use to move throughout a menu: up, down, right, and left.

Each button corresponds to the direction that you can move in a menu. For example, to move right in a menu, press the navigation button that is located on the right side. If you want to move down in a menu, press the navigation button that is located on the bottom.

**Note:** The select button is used in tandem with the navigation buttons.

### **Product serial number**

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

This number is used for warranty and service entitlement checking and is included in the data sent with error reports. It is essential that this number is *not* changed 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.

### **Select button**

Use the select button to select an item from a menu.

The select button and navigation buttons help you to navigate and select menu and boot options, and start a service panel test. The select button is located on the front panel of the SAN Volume Controller, near the navigation buttons.

### **Node identification label**

The node identification label on the front panel displays a six-digit node identification number. Sometimes this number is called the panel name or front panel ID.

The node identification label is the six-digit number that is input to the **svctask addnode** command. It is readable by system software and is used by configuration and service software as a node identifier. The node identification number can also be displayed on the front-panel display when node is selected from the menu.

If the service controller assembly front panel is replaced, the configuration and service software displays the number that is printed on the front of the replacement panel. Future error reports contain the new number. No cluster reconfiguration is necessary when the front panel is replaced.

### **Error LED**

Critical faults on the service controller are indicated through the amber error LED.

The error LED has the following two states:

**OFF** The service controller is functioning correctly.

**ON** A critical service-controller failure was detected and you must replace the service controller.

The error LED can light temporarily when the node is powered on. If the error LED is on, but the front panel display is completely blank, wait five minutes to allow the LED time to turn off before performing any service action.

|  
|  
|  
|

## SAN Volume Controller operator-information panel

The operator-information panel is located on the front panel of the SAN Volume Controller.

### SAN Volume Controller 2145-CF8 operator-information panel

The operator-information panel contains buttons and indicators such as the power-control button, and LEDs that indicate information such as system-board errors, hard-drive activity, and power status.

Figure 9 shows the operator-information panel for the SAN Volume Controller 2145-CF8.




Figure 9. SAN Volume Controller 2145-CF8 operator-information panel

- 1 Power-button cover
- 2 Ethernet 2 activity LED
- 3 Ethernet 1 activity LED
- 4 System-information LED
- 5 System-error LED
- 6 Release latch
- 7 Locator button and LED
- 8 Ethernet 3 activity LED
- 9 Ethernet 4 activity LED
- 10 Power button and LED

### SAN Volume Controller 2145-8A4 operator-information panel

The operator-information panel contains buttons and indicators such as the power-control button, and LEDs that indicate information such as system-board errors, hard-drive activity, and power status.

Figure 10 on page 15 shows the operator-information panel for the SAN Volume Controller 2145-8A4.




Figure 10. SAN Volume Controller 2145-8A4 operator-information panel

- 1 System-error LED (amber)
- 2 Locator LED (blue)
- 3 Hard-disk drive activity LED (green)
- 4 Reset button
- 5 Power-control button
- 6 Power LED (green)

### SAN Volume Controller 2145-8G4 operator information panel

The operator-information panel contains buttons and indicators such as the release latch for the light path diagnostics panel, the power-control button, and LEDs that indicate information such as system-board errors, hard-drive activity, and power status.

Figure 11 shows the operator information panel for the SAN Volume Controller 2145-8G4.




Figure 11. SAN Volume Controller 2145-8G4 operator-information panel

- 1 Release latch for light path diagnostics panel
- 2 System-error LED (amber)
- 3 System-information LED (amber)
- 4 Locator LED (blue)
- 5 Hard disk drive activity LED (green)
- 6 Power LED (green)
- 7 Power-control button

### SAN Volume Controller 2145-8F4 and SAN Volume Controller 2145-8F2 operator information panel

The operator-information panel contains buttons and indicators such as the release latch for the light path diagnostics panel, the power-control button, and LEDs that indicate information such as system-board errors, hard-drive activity, and power status.

Figure 12 shows the operator-information panel that is used by the SAN Volume Controller 2145-8F4 and the SAN Volume Controller 2145-8F2 models.




Figure 12. SAN Volume Controller 2145-8F2 and SAN Volume Controller 2145-8F4 operator-information panel

- 1** Release latch for light path diagnostics panel
- 2** System-error LED (amber)
- 3** Information LED (amber)
- 4** Locator LED (blue)
- 5** Hard disk drive activity LED (green)
- 6** Power control button
- 7** Power LED (green)
- 8** USB connector

### 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 SAN Volume Controller hardware detects a fatal error that requires a new field-replaceable unit (FRU).

**Note:** If you have SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2, see MAP 5800: Light path to help you isolate the faulty FRU.

A system-error LED is also at the rear of the SAN Volume Controller models 2145-CF8, 2145-8G4, 2145-8F4, and 2145-8F2.

### Hard-disk drive activity LED

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

### Reset button

A reset button is available on the SAN Volume Controller 2145-8A4 node, but do not use it.

**Attention:** If you use the reset button, the node restarts immediately without the SAN Volume Controller control data being written to 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” in the *IBM System Storage SAN Volume Controller Troubleshooting Guide*.

**Attention:** When the node is operational and you press and immediately release the power button, the SAN Volume Controller indicates on its front panel that it is turning off and writes its control data to its internal disk. This can take up to five minutes. If you press the power button but do not release it, the node turns off immediately without the SAN Volume Controller control data 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 two seconds.

**Note:** The 2145 UPS-1U does not turn off when the SAN Volume Controller is shut down from the power button.

### **Power LED**

The green power LED indicates the power status of the SAN Volume Controller.

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 SAN Volume Controller node is turned on.

### **Flashing**

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

**Note:** A power LED is also at the rear of the SAN Volume Controller 2145-8F2, 2145-8F4, and 2145-8G4 nodes.

### **Release latch**

The release latch on the SAN Volume Controller models 2145-8G4, 2145-8F4, and 2145-8F2 gives you access to the light path diagnostics panel, which provides a method for determining the location of a problem.

After pressing the release latch on the operator-information panel, you can slide the light path diagnostics panel out to view the lit LEDs. The LEDs indicate the type of error that has occurred. See MAP 5800: Light path for more detail.

To retract the panel, push it back into the node and snap it into place.

### **System-information LED**

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

Check the light path diagnostics panel and the error 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

The Ethernet activity LED identifies Ethernet activity.

When any of these LEDs is lit, they indicate that the node is transmitting to or receiving signals from the Ethernet LAN that is connected to the Ethernet port that corresponds to that LED.

## SAN Volume Controller rear-panel indicators and connectors

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

### SAN Volume Controller 2145-CF8 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 13 shows the rear-panel indicators on the SAN Volume Controller 2145-CF8 back-panel assembly.




Figure 13. SAN Volume Controller 2145-CF8 rear-panel indicators

- 1** Fibre-channel LEDs
- 2** Ac, dc, and power-supply error LEDs
- 3** Power, location, and system-error LEDs
- 4** Ethernet-link LEDs
- 5** Ethernet-activity LED

### SAN Volume Controller 2145-CF8 connectors

External connectors that the SAN Volume Controller 2145-CF8 uses include four fibre-channel ports, a serial port, two Ethernet ports, and two power connectors.

Figure 14 on page 19 shows the external connectors on the SAN Volume Controller 2145-CF8 back panel assembly.




Figure 14. Connectors on the rear of the SAN Volume Controller 2145-CF8

- 1 Fibre-channel port 1
- 2 Fibre-channel port 2
- 3 Fibre-channel port 3
- 4 Fibre-channel port 4
- 5 Power-cord connector for power supply 1
- 6 Power-cord connector for power supply 2
- 7 Serial connection for UPS communication cable
- 8 Ethernet port 2
- 9 Ethernet port 1

Figure 15 shows the type of connector that is located on each power-supply assembly. Use these connectors to connect the SAN Volume Controller 2145-CF8 to the two power cables from the uninterruptible power supply.




Figure 15. Power connector

**SAN Volume Controller 2145-CF8 ports used during service procedures:**

The SAN Volume Controller 2145-CF8 contains a number of ports that are only used during service procedures.

Figure 16 on page 20 shows ports that are used only during service procedures.




Figure 16. Service ports of the SAN Volume Controller 2145-CF8

- 1 System management port
- 2 Two monitor ports, one on the front and one on the rear
- 3 Four USB ports, two on the front and two on the rear

During normal operation, 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-CF8 unused ports:**

The SAN Volume Controller 2145-CF8 can contain one port that is not used.

Figure 17 shows the one port that is not used during service procedures or normal use.




Figure 17. SAN Volume Controller 2145-CF8 port not used

- 1 Serial-attached SCSI (SAS) port

When present, this port is disabled in software to make the port inactive. The SAS port is present when the optional high-speed SAS adapter is installed with one or more solid-state drives (SSDs).

**SAN Volume Controller 2145-8A4 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 18 on page 21 shows the rear-panel indicators on the SAN Volume Controller 2145-8A4 back-panel assembly.




Figure 18. SAN Volume Controller 2145-8A4 rear-panel indicators

- 1** Fibre-channel LEDs
- 2** Ethernet port 1 activity LED
- 3** Ethernet port 1 link LED
- 4** Ethernet port 2 activity LED
- 5** Ethernet port 2 link LED

### SAN Volume Controller 2145-8A4 connectors

The external connectors consist of fibre-channel, serial and Ethernet ports, and the power supply.

Figure 19 shows the external connectors on the SAN Volume Controller 2145-8A4 back-panel assembly.




Figure 19. SAN Volume Controller 2145-8A4 external connectors

- 1** Fibre-channel port 1
- 2** Fibre-channel port 2
- 3** Fibre-channel port 3
- 4** Fibre-channel port 4
- 5** Power supply
- 6** Serial connection
- 7** Ethernet port 2
- 8** Ethernet port 1

Figure 20 on page 22 shows the type of connector that is located on the power supply assembly. The connector enables you to connect the SAN Volume Controller 2145-8A4 to the power source from the uninterruptible power supply.




Figure 20. Power connector

### SAN Volume Controller 2145-8A4 ports used during service procedures

The SAN Volume Controller 2145-8A4 contains a number of ports that are used only during service procedures. These ports are shown in Figure 21.




Figure 21. Service ports of the SAN Volume Controller 2145-8A4

- 1** System management port
- 2** Four USB ports, two on the front and two on the rear
- 3** One video port on the rear

During normal operation, 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 your IBM service representative.

### SAN Volume Controller 2145-8A4 ports not used

The SAN Volume Controller 2145-8A4 has no unused ports.

### SAN Volume Controller 2145-8G4 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 22 on page 23 shows the rear-panel indicators on the SAN Volume Controller 2145-8G4 back-panel assembly.




Figure 22. SAN Volume Controller 2145-8G4 rear-panel indicators

- 1** Fibre-channel LEDs
- 2** Ethernet port 1 activity LED
- 3** Ethernet port 1 link LED
- 4** Ethernet port 2 activity LED
- 5** Ethernet port 2 link LED
- 6** Power, location, and system error LEDs
- 7** Ac and dc LEDs

### SAN Volume Controller 2145-8G4 connectors

The external connectors consist of fibre-channel, serial, and Ethernet ports, and the power supply.

Figure 23 shows the external connectors on the SAN Volume Controller 2145-8G4 back panel assembly.




Figure 23. SAN Volume Controller 2145-8G4 external connectors

- 1** Fibre-channel port 1
- 2** Fibre-channel port 2
- 3** Fibre-channel port 3
- 4** Fibre-channel port 4
- 5** Power supply
- 6** Serial connection
- 7** Ethernet port 2
- 8** Ethernet port 1

Figure 24 on page 24 shows the type of connector that is located on the power supply assembly. The connector enables you to connect the SAN Volume Controller

2145-8G4 to the power source from the uninterruptible power supply.




Figure 24. Power connector

### SAN Volume Controller 2145-8G4 ports used during service procedures

The SAN Volume Controller 2145-8G4 contains a number of ports that are only used during service procedures. These ports are shown in Figure 25.




Figure 25. Service ports of the SAN Volume Controller 2145-8G4

- | **1** System management port
- | **2** Four USB ports, two on the front and two on the rear
- | **3** Two monitor ports, one on the front and one on the rear

During normal operation, 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 your IBM service representative.

### SAN Volume Controller 2145-8G4 ports not used

The SAN Volume Controller 2145-8G4 has no unused ports.

### SAN Volume Controller 2145-8F4 rear-panel indicators

The rear-panel indicators are located on the back-panel assembly.

Figure 26 on page 25 shows the rear-panel indicators on the SAN Volume Controller 2145-8F4 back-panel assembly.






Figure 26. SAN Volume Controller 2145-8F4 rear-panel indicators

- 1** Fibre-channel LEDs
- 2** Ethernet port 1 link LED
- 3** Ethernet port 1 activity LED
- 4** Ethernet port 2 link LED
- 5** Ethernet port 2 activity LED
- 6** Power, location, and system error LEDs
- 7** Ac and dc LEDs

### SAN Volume Controller 2145-8F4 connectors

The external connectors consist of Ethernet, serial, and fibre-channel ports, and the power supply.

Figure 27 shows the external connectors on the SAN Volume Controller 2145-8F4 back panel assembly.




Figure 27. SAN Volume Controller 2145-8F4 external connectors

- 1** Fibre-channel port 1
- 2** Fibre-channel port 2
- 3** Fibre-channel port 3
- 4** Fibre-channel port 4
- 5** Power supply
- 6** Serial connection
- 7** Ethernet port 2
- 8** Ethernet port 1

Figure 28 shows the type of connector that is located on the power supply assembly. The connector enables you to connect the SAN Volume Controller 2145-8F4 to the power source from the uninterruptible power supply.




Figure 28. Power connector

### SAN Volume Controller 2145-8F4 ports used during service procedures

The SAN Volume Controller 2145-8F4 contains the keyboard service port and the monitor service port. These ports are used only during service procedures. Figure 29 provides the locations of the service ports.




Figure 29. Service ports of the SAN Volume Controller 2145-8F4

**1** Keyboard port

**2** Monitor port

### SAN Volume Controller 2145-8F4 ports not used during normal operation

The SAN Volume Controller 2145-8F4 is equipped with several ports that are not used by the SAN Volume Controller during normal operation. Figure 30 and Figure 31 on page 27 show the ports that are not used by the SAN Volume Controller.




Figure 30. Ports not used during normal operation by the SAN Volume Controller 2145-8F4

**1** System management port

**2** Mouse port

**3** Keyboard port

- 4** USB ports
- 5** Monitor port




Figure 31. Ports not used on the front panel of the SAN Volume Controller 2145-8F4

- 1** USB port

### SAN Volume Controller 2145-8F2 rear-panel indicators

The rear-panel indicators are located on the back-panel assembly.

Figure 32 shows the rear-panel indicators on the SAN Volume Controller 2145-8F2 back-panel assembly.




Figure 32. SAN Volume Controller 2145-8F2 rear-panel indicators

- 1** Fibre-channel LEDs
- 2** Ethernet port 1 link LED
- 3** Ethernet port 1 activity LED
- 4** Ethernet port 2 link LED
- 5** Ethernet port 2 activity LED
- 6** Power, location, and system error LEDs
- 7** Ac and dc LEDs

### SAN Volume Controller 2145-8F2 connectors

The external connectors consist of the power supply and Ethernet, fibre-channel, and serial ports.

Figure 33 on page 28 shows the external connectors on the SAN Volume Controller 2145-8F2 back panel assembly.




Figure 33. SAN Volume Controller 2145-8F2 external connectors

- 1** Power supply
- 2** Fibre-channel port 4
- 3** Serial connection
- 4** Fibre-channel port 3
- 5** Fibre-channel port 2
- 6** Fibre-channel port 1
- 7** Ethernet port 2
- 8** Ethernet port 1

Figure 34 shows the type of connector that is located on the power supply assembly. The connector enables you to connect the SAN Volume Controller 2145-8F2 to the power source from the uninterruptible power supply.




Figure 34. Power connector

### Fibre-channel LEDs

The fibre-channel LEDs indicate the status of the fibre-channel ports.

Two LEDs are used to indicate the state and speed of the operation of each fibre-channel port. The bottom LED indicates the link state and activity.

Table 4. Link state and activity for bottom LED

Link activity	Link state
OFF	Link inactive
ON	Link active, no I/O
BLINKING	Link active, I/O active

Each fibre-channel port can operate at one of three speeds. The top LED indicates the relative link speed. The link speed is defined only if the link state is active.

Table 5. Relative link speed for the top LED

Link speed LED	Link speed
OFF	SLOW
ON	FAST
BLINKING	MEDIUM

Table 6 shows the actual link speeds for the SAN Volume Controller models 2145-8A4, 2145-8G4, and 2145-8F4.

Table 6. Actual link speeds

Link speed	Actual link speeds
SLOW	1 Gbps
FAST	4 Gbps
MEDIUM	2 Gbps

Table 7 shows the actual link speeds for the SAN Volume Controller 2145-CF8.

Table 7. Actual link speeds

Link speed	Actual link speeds
SLOW	2 Gbps
FAST	8 Gbps
MEDIUM	4 Gbps

### Related reference

“Fibre-channel port numbers and worldwide port names” on page 32  
Fibre-channel ports are identified by their physical port number and by a worldwide port name (WWPN).

### Ethernet activity LED

The Ethernet activity LED indicates that the SAN Volume Controller 2145-8G4 or SAN Volume Controller 2145-8A4 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 together on the rear of the SAN Volume Controller next to the monitor port. 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.

### Flashing

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 **1** indicates that ac current is present on the node.

### Dc LED

The lower LED **2** 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 35 on page 31 shows the location of the SAN Volume Controller 2145-CF8 ac, dc, and power-supply error LEDs.




Figure 35. SAN Volume Controller 2145-CF8 ac, dc, and power-error LEDs

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

**Ac LED**

The upper LED ( **1** ) on the left side of the power supply, indicates that ac current is present on the node.

**Dc LED**

The middle LED ( **2** ) to the left side of the power supply, indicates that dc current is present on the node.

**Power-supply error LED**

The lower LED ( **3** ) to the left side of the power supply, indicates a problem with the power supply.

**Ac and dc LEDs on the SAN Volume Controller 2145-8G4:**

The ac LED and dc LED are located on the rear of the SAN Volume Controller 2145-8G4.

Figure 36 on page 32 shows the location of the ac and dc LEDs.




Figure 36. SAN Volume Controller 2145-8G4 ac and dc LEDs

**Ac and dc LEDs on the SAN Volume Controller 2145-8F4 and the SAN Volume Controller 2145-8F2:**

The ac LED and dc LED are located on the rear of the SAN Volume Controller 2145-8F4 and the SAN Volume Controller 2145-8F2.

Figure 37 shows the location of the ac and dc LEDs.




Figure 37. SAN Volume Controller 2145-8F4 and SAN Volume Controller 2145-8F2 ac and dc LEDs

**Fibre-channel port numbers and worldwide port names**

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

The physical port numbers identify fibre-channel cards and cable connections when you perform service tasks. The WWPNs are used for tasks such as fibre-channel switch configuration and to uniquely identify the devices on the SAN.



The WWPNs are derived from the worldwide node name (WWNN) of the SAN Volume Controller node in which the card is installed.

The WWNN is in the form 50050768010XXXXX, where XXXXX is initially derived from the unit and is specific to a SAN Volume Controller node. You can change the XXXXX value by using the front panel to facilitate service controller concurrent replacement and to enable some concurrent upgrade operations.

The WWPNs are in the form 5005076801QXXXXX, where XXXXX is as previously stated and Q is related to the port number as follows:

Port	Value of Q
1	4
2	3
3	1
4	2

#### Related reference

“Fibre-channel LEDs” on page 28

The fibre-channel LEDs indicate the status of the fibre-channel ports.

#### SAN Volume Controller 2145-8A4 fibre-channel port numbers

The physical port numbers are 1 - 4, counting from left to right when you view the rear panel of the SAN Volume Controller 2145-8A4.

Figure 38 identifies the fibre-channel port numbers on the SAN Volume Controller 2145-8A4.




Figure 38. The physical port numbers for the SAN Volume Controller 2145-8A4

#### SAN Volume Controller 2145-8G4 fibre-channel port numbers

The physical port numbers are 1 - 4, counting from left to right when you view the rear panel of the SAN Volume Controller 2145-8G4.

Figure 39 provides a view of the rear of the SAN Volume Controller 2145-8G4.




Figure 39. The physical port numbers for the SAN Volume Controller 2145-8G4

## SAN Volume Controller 2145-8F4 fibre-channel port numbers

The physical port numbers are 1 - 4, counting from left to right when you view the rear panel of the SAN Volume Controller 2145-8F4.

Figure 40 provides a view of the rear of the SAN Volume Controller 2145-8F4.




Figure 40. The physical port numbers for the SAN Volume Controller 2145-8F4

## SAN Volume Controller 2145-8F2 fibre-channel port numbers

The physical port numbers are 1 - 4, counting from left to right when you view the rear panel of the SAN Volume Controller 2145-8F2.

Figure 41 provides a view of the rear of the SAN Volume Controller 2145-8F2.




Figure 41. The physical port numbers for the SAN Volume Controller 2145-8F2

## 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-CF8 environment requirements

Before installing a SAN Volume Controller 2145-CF8 node, your 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 following voltage requirements.

Voltage	Frequency
200 to 240 V single phase ac	50 or 60 Hz

**Attention:**

- If the uninterruptible power supply is cascaded from another uninterruptible power supply, the source uninterruptible power supply must have at least three times the capacity per phase and the total harmonic distortion must be less than 5%.
- The uninterruptible power supply also must have input voltage capture that has a slew rate of no more than 3 Hz per second.

**Power requirements for each node**

Ensure that your environment meets the following power requirements.

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

Components	Power requirements
SAN Volume Controller 2145-CF8 node and 2145 UPS-1U power supply	200 W

**Notes:**

- SAN Volume Controller 2145-CF8 nodes will not connect to all revisions of the 2145 UPS-1U power supply unit. The SAN Volume Controller 2145-CF8 nodes require the 2145 UPS-1U power supply unit part number 31P1318. This unit has two power outlets that are accessible. Earlier revisions of the 2145 UPS-1U power supply unit have only one power outlet that is accessible and are not suitable.
- For each redundant ac-power switch, add 20 W to the power requirements.
- For each high-speed SAS adapter with one to four solid-state drives (SSDs), add 50 W to the power requirements.

**Circuit breaker requirements**

The 2145 UPS-1U has an integrated circuit breaker and does not require additional protection.

**Environment requirements without redundant ac power**

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

Environment	Temperature	Altitude	Relative humidity	Maximum wet bulb temperature
Operating in lower altitudes	10°C to 35°C (50°F to 95°F)	0 to 914 m (0 to 2998 ft)	8% to 80% noncondensing	23°C (73°F)
Operating in higher altitudes	10°C to 32°C (50°F to 90°F)	914 to 2133 m (2998 to 6988 ft)	8% to 80% noncondensing	23°C (73°F)
Turned off	10°C to 43°C (50°F to 110°F)	0 to 2133 m (0 to 6988 ft)	8% to 80% noncondensing	27°C (81°F)
Storing	1°C to 60°C (34°F to 140°F)	0 to 2133 m (0 to 6988 ft)	5% to 80% noncondensing	29°C (84°F)

Environment	Temperature	Altitude	Relative humidity	Maximum wet bulb temperature
Shipping	-20°C to 60°C (-4°F to 140°F)	0 to 10668 m (0 to 34991 ft)	5% to 100% condensing, but no precipitation	29°C (84°F)

## Environment requirements with redundant ac power

Ensure that your environment falls within the following ranges if you are using redundant ac power.

Environment	Temperature	Altitude	Relative humidity	Maximum wet bulb temperature
Operating in lower altitudes	15°C to 32°C (59°F to 90°F)	0 to 914 m (0 to 2998 ft)	20% to 80% noncondensing	23°C (73°F)
Operating in higher altitudes	15°C to 32°C (59°F to 90°F)	914 to 2133 m (2998 to 6988 ft)	20% to 80% noncondensing	23°C (73°F)
Turned off	10°C to 43°C (50°F to 110°F)	0 to 2133m (0 to 6988 ft)	20% to 80% noncondensing	27°C (81°F)
Storing	1°C to 60°C (34°F to 140°F)	0 to 2133 m (0 to 6988 ft)	5% to 80% noncondensing	29°C (84°F)
Shipping	-20°C to 60°C (-4°F to 140°F)	0 to 10668 m (0 to 34991 ft)	5% to 100% condensing, but no precipitation	29°C (84°F)

## Preparing your environment

The following tables list the physical characteristics of the SAN Volume Controller 2145-CF8 node.

### Dimensions and weight

Ensure that space is available in a rack that is capable of supporting the node.

Height	Width	Depth	Maximum weight
43 mm (1.69 in.)	440 mm (17.32 in.)	686 mm (27 in.)	12.7 kg (28 lb)

### Additional space requirements

Ensure that space is also available in the rack for the following additional space requirements around the node.

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

## Heat output of each SAN Volume Controller 2145-CF8 node

The node dissipates the following maximum heat output.

Model	Heat output per node
SAN Volume Controller 2145-CF8	160 W (546 Btu per hour)
SAN Volume Controller 2145-CF8 and up to four optional solid-state drives (SSDs)	210 W (717 Btu per hour)
Maximum heat output of 2145 UPS-1U during typical operation:	10 W (34 Btu per hour)
Maximum heat output of 2145 UPS-1U during battery operation:	100 W (341 Btu per hour)

## SAN Volume Controller 2145-8A4 environment requirements

Before the SAN Volume Controller 2145-8A4 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 following voltage requirements.

Voltage	Frequency
200 to 240 V single phase ac	50 or 60 Hz

### Attention:

- If the uninterruptible power supply is cascaded from another uninterruptible power supply, the source uninterruptible power supply must have at least three times the capacity per phase and the total harmonic distortion must be less than 5%.
- The uninterruptible power supply also must have input voltage capture that has a slew rate of no more than 3 Hz per second.

### Power requirements for each node

Ensure that your environment meets the following power requirements.

The power that is required depends on the node type and whether the redundant ac power feature is used.

Components	Power requirements
SAN Volume Controller 2145-8A4 and 2145 UPS-1U	180 W

For each redundant ac-power switch, add 20 W to the power requirements.

### Circuit breaker requirements

The 2145 UPS-1U has an integrated circuit breaker and does not require additional protection.

## Environment requirements without redundant ac power

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

Environment	Temperature	Altitude	Relative humidity	Maximum wet bulb temperature
Operating in lower altitudes	10°C to 35°C (50°F to 95°F)	0 to 914 m (0 to 3000 ft)	8% to 80% noncondensing	23°C (73°F)
Operating in higher altitudes	10°C to 32°C (50°F to 90°F)	914 to 2133 m (3000 to 7000 ft)	8% to 80% noncondensing	23°C (73°F)
Turned off	10°C to 43°C (50°F to 109°F)	0 to 2133 m (0 to 7000 ft)	8% to 80% noncondensing	27°C (81°F)
Storing	1°C to 60°C (34°F to 140°F)	0 to 2133 m (0 to 7000 ft)	5% to 80% noncondensing	29°C (84°F)
Shipping	-20°C to 60°C (-4°F to 140°F)	0 to 10668 m (0 to 34991 ft)	5% to 100% condensing, but no precipitation	29°C (84°F)

## Environment requirements with redundant ac power

Ensure that your environment falls within the following ranges if you are using redundant ac power.

Environment	Temperature	Altitude	Relative humidity	Maximum wet bulb temperature
Operating in lower altitudes	15°C to 32°C (59°F to 90°F)	0 to 914 m (0 to 3000 ft)	20% to 80% noncondensing	23°C (73°F)
Operating in higher altitudes	15°C to 32°C (59°F to 90°F)	914 to 2133 m (3000 to 7000 ft)	20% to 80% noncondensing	23°C (73°F)
Turned off	10°C to 43°C (50°F to 109°F)	0 to 2133 m (0 to 7000 ft)	20% to 80% noncondensing	27°C (81°F)
Storing	1°C to 60°C (34°F to 140°F)	0 to 2133 m (0 to 7000 ft)	5% to 80% noncondensing	29°C (84°F)
Shipping	-20°C to 60°C (-4°F to 140°F)	0 to 10668 m (0 to 34991 ft)	5% to 100% condensing, but no precipitation	29°C (84°F)

## Preparing your environment

The following tables list the physical characteristics of the SAN Volume Controller 2145-8A4 node.

### Dimensions and weight

Ensure that space is available in a rack that is capable of supporting the node.

Height	Width	Depth	Maximum weight
43 mm (1.75 in.)	440 mm (17.32 in.)	559 mm (22 in.)	10.1 kg (22 lb)

## Additional space requirements

Ensure that space is also available in the rack for the following additional space requirements around the node.

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

## Heat output of each SAN Volume Controller 2145-8A4 node

The node dissipates the following maximum heat output.

Model	Heat output per node
SAN Volume Controller 2145-8A4	140 W (478 Btu per hour)

## SAN Volume Controller 2145-8G4 environment requirements

Before the SAN Volume Controller 2145-8G4 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 following voltage requirements.

Voltage	Frequency
200 to 240 V single phase ac	50 or 60 Hz

### Attention:

- If the uninterruptible power supply is cascaded from another uninterruptible power supply, the source uninterruptible power supply must have at least three times the capacity per phase and the total harmonic distortion must be less than 5%.
- The uninterruptible power supply also must have input voltage capture that has a slew rate of no more than 3 Hz per second.

### Power requirements for each node

Ensure that your environment meets the following power requirements.

The power that is required depends on the node type and whether the redundant ac power feature is used.

Components	Power requirements
SAN Volume Controller 2145-8G4 and 2145 UPS-1U	470 W

For each redundant ac-power switch, add 20 W to the power requirements.

## Circuit breaker requirements

The 2145 UPS-1U has an integrated circuit breaker and does not require additional protection.

## Environment requirements without redundant ac power

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

Environment	Temperature	Altitude	Relative humidity	Maximum wet bulb temperature
Operating in lower altitudes	10°C to 35°C (50°F to 95°F)	0 to 914 m (0 to 2998 ft)	8% to 80% noncondensing	23°C (73°F)
Operating in higher altitudes	10°C to 32°C (50°F to 90°F)	914 to 2133 m (2998 to 6988 ft)	8% to 80% noncondensing	23°C (73°F)
Turned off	10°C to 43°C (50°F to 110°F)	0 to 2133 m (0 to 6988 ft)	8% to 80% noncondensing	27°C (81°F)
Storing	1°C to 60°C (34°F to 140°F)	0 to 2133 m (0 to 6988 ft)	5% to 80% noncondensing	29°C (84°F)
Shipping	-20°C to 60°C (-4°F to 140°F)	0 to 10668 m (0 to 34991 ft)	5% to 100% condensing, but no precipitation	29°C (84°F)

## Environment requirements with redundant ac power

Ensure that your environment falls within the following ranges if you are using redundant ac power.

Environment	Temperature	Altitude	Relative humidity	Maximum wet bulb temperature
Operating in lower altitudes	15°C to 32°C (59°F to 90°F)	0 to 914 m (0 to 2998 ft)	20% to 80% noncondensing	23°C (73°F)
Operating in higher altitudes	15°C to 32°C (59°F to 90°F)	914 to 2133 m (2998 to 6988 ft)	20% to 80% noncondensing	23°C (73°F)
Turned off	10°C to 43°C (50°F to 110°F)	0 to 2133m (0 to 6988 ft)	20% to 80% noncondensing	27°C (81°F)
Storing	1°C to 60°C (34°F to 140°F)	0 to 2133 m (0 to 6988 ft)	5% to 80% noncondensing	29°C (84°F)
Shipping	-20°C to 60°C (-4°F to 140°F)	0 to 10668 m (0 to 34991 ft)	5% to 100% condensing, but no precipitation	29°C (84°F)

## Preparing your environment

The following tables list the physical characteristics of the SAN Volume Controller 2145-8G4 node.



## Dimensions and weight

Ensure that space is available in a rack that is capable of supporting the node.

Height	Width	Depth	Maximum weight
43 mm (1.69 in.)	440 mm (17.32 in.)	686 mm (27 in.)	12.7 kg (28 lb)

## Additional space requirements

Ensure that space is also available in the rack for the following additional space requirements around the node.

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

## Heat output of each SAN Volume Controller 2145-8G4 node

The node dissipates the following maximum heat output.

Model	Heat output per node
SAN Volume Controller 2145-8G4	400 W (1350 Btu per hour)

## SAN Volume Controller 2145-8F4 and SAN Volume Controller 2145-8F2 environment requirements

Before the SAN Volume Controller 2145-8F4 or SAN Volume Controller 2145-8F2 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 following voltage requirements.

Voltage	Frequency
200 to 240 V single phase ac	50 or 60 Hz

## Power requirements for each node

Ensure that your environment meets the following power requirements.

The power that is required depends on the node type and whether the redundant ac power feature is used.

Components	Power requirements
SAN Volume Controller 2145-8F4 and 2145 UPS-1U	520 W

Components	Power requirements
SAN Volume Controller 2145-8F2 and 2145 UPS-1U	520 W

For each redundant ac-power switch, add 20 W to the power requirements.

### Circuit breaker requirements

The 2145 UPS-1U has an integrated circuit breaker and does not require additional protection.

### Environment requirements without redundant ac power

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

Environment	Temperature	Altitude	Relative humidity	Maximum wet bulb temperature
Operating in lower altitudes	10°C to 35°C (50°F to 95°F)	0 to 914.4 m (0 to 3000 ft)	8% to 80% noncondensing	23°C (74°F)
Operating in higher altitudes	10°C to 32°C (50°F to 88°F)	914.4 to 2133.6 m (3000 to 7000 ft)	8% to 80% noncondensing	23°C (74°F)
Turned off	10°C to 43°C (50°F to 110°F)	0 to 2133.6 m (3000 to 7000 ft)	8% to 80% noncondensing	27°C (81°F)
Storing	1°C to 60°C (34°F to 140°F)	0 to 2133.6 m (0 to 7000 ft)	5% to 80% noncondensing	29°C (84°F)
Shipping	-20°C to 60°C (-4°F to 140°F)	0 to 10668 m (0 to 34991 ft)	5% to 100% condensing, but no precipitation	29°C (84°F)

### Environment requirements with redundant ac power

Ensure that your environment falls within the following ranges if you are using redundant ac power.

Environment	Temperature	Altitude	Relative humidity	Maximum wet bulb temperature
Operating in lower altitudes	15°C to 32°C (59°F to 89°F)	0 to 914.4 m (0 to 3000 ft)	20% to 80% noncondensing	23°C (74°F)
Operating in higher altitudes	15°C to 32°C (50°F to 88°F)	914.4 to 2133.6 m (3000 to 7000 ft)	20% to 80% noncondensing	23°C (74°F)
Turned off	10°C to 43°C (50°F to 110°F)	0 to 2133.6 m (0 to 7000 ft)	20% to 80% noncondensing	27°C (81°F)
Storing	1°C to 60°C (34°F to 140°F)	0 to 2133.6 m (0 to 7000 ft)	5% to 80% noncondensing	29°C (84°F)
Shipping	-20°C to 60°C (-4°F to 140°F)	0 to 10668 m (0 to 34991 ft)	5% to 100% condensing, but no precipitation	29°C (84°F)

## Preparing your environment

The following tables list the physical characteristics of the SAN Volume Controller 2145-8F4 and SAN Volume Controller 2145-8F2 nodes.

### Dimensions and weight

Ensure that space is available in a rack that is capable of supporting the node.

Height	Width	Depth	Maximum weight
43 mm (1.69 in.)	440 mm (17.32 in.)	686 mm (27 in.)	12.7 kg (28 lb)

### Additional space requirements

Ensure that space is also available in the rack for the following additional space requirements around the node.

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

### Heat output of each SAN Volume Controller 2145-8F4 or SAN Volume Controller 2145-8F2 node

The nodes dissipate the following maximum heat output.

Model	Heat output per node
SAN Volume Controller 2145-8F4	450 W (1540 Btu per hour)
SAN Volume Controller 2145-8F2	450 W (1540 Btu per hour)

---

## Redundant ac-power switch

The redundant ac-power switch is an optional feature that makes the SAN Volume Controller nodes resilient to the failure of a single power circuit. The redundant ac-power switch is not a replacement for an uninterruptible power supply. You must still use a uninterruptible power supply for each node.

You must connect the redundant ac-power switch to two independent power circuits. One power circuit connects to the main power input port and the other power circuit connects to the backup power-input port. If the main power to the SAN Volume Controller node fails for any reason, the redundant ac-power switch automatically uses the backup power source. When power is restored, the redundant ac-power switch automatically changes back to using the main power source.

Place the redundant ac-power switch in the same rack as the SAN Volume Controller node. The redundant ac-power switch logically sits between the rack power distribution unit and the 2145 UPS-1U.

You can use a single redundant ac-power switch to power one or two SAN Volume Controller nodes. If you use the redundant ac-power switch to power two nodes, the nodes must be in different I/O groups. In the event that the redundant ac-power switch fails or requires maintenance, both nodes turn off. Because the nodes are in two different I/O groups, the hosts do not lose access to the back-end disk data.

For maximum resilience to failure, use one redundant ac-power switch to power each SAN Volume Controller node.

Figure 42 shows a redundant ac-power switch.



Figure 42. Photo of the redundant ac-power switch

## Redundant ac-power environment requirements

Ensure that your physical site meets the installation requirements for the redundant ac-power switch.

The redundant ac-power switch requires two independent power sources that are provided through two rack-mounted power distribution units (PDUs). The PDUs must have IEC320-C13 outlets.

The redundant ac-power switch comes with two IEC 320-C19 to C14 power cables to connect to rack PDUs. There are no country-specific cables for the redundant ac-power switch.

The power cable between the redundant ac-power switch and the 2145 UPS-1U is rated at 10 A.

## Redundant ac-power switch specifications

The following tables list the physical characteristics of the redundant ac-power switch.

### Dimensions and weight

Ensure that space is available in a rack that is capable of supporting the redundant ac-power switch.

Height	Width	Depth	Maximum weight
43 mm (1.69 in.)	192 mm (7.56 in.)	240 mm	2.6 kg (5.72 lb)

### Additional space requirements

Ensure that space is also available in the rack for the side mounting plates on either side of the redundant ac-power switch.

Location	Width	Reason
Left side	124 mm (4.89 in.)	Side mounting plate
Right side	124 mm (4.89 in.)	Side mounting plate

### Heat output (maximum)

The maximum heat output that is dissipated inside the redundant ac-power switch is approximately 20 watts (70 Btu per hour).


### Cabling of redundant ac-power switch (example)

You must properly cable the redundant ac-power switch units in your environment.

**Note:** While this topic provides an example of the cable connections, it does not indicate a preferred physical location for the components.

Figure 43 on page 46 shows an example of the main wiring for a SAN Volume Controller cluster with the redundant ac-power switch feature. The four-node cluster consists of two I/O groups:

- I/O group 0 contains nodes A and B
- I/O group 1 contains nodes C and D



svc00358\_cf8

Figure 43. A four-node SAN Volume Controller cluster with the redundant ac-power switch feature

- 1** I/O group 0
- 2** SAN Volume Controller node A
- 3** 2145 UPS-1U A
- 4** SAN Volume Controller node B
- 5** 2145 UPS-1U B
- 6** I/O group 1
- 7** SAN Volume Controller node C
- 8** 2145 UPS-1U C
- 9** SAN Volume Controller node D
- 10** 2145 UPS-1U D
- 11** Redundant ac-power switch 1
- 12** Redundant ac-power switch 2
- 13** Site PDU X (C13 outlets)
- 14** Site PDU Y (C13 outlets)

The site PDUs X and Y ( **13** and **14** ) are powered from two independent power sources.

In this example, only two redundant ac-power switch units are used, and each power switch powers one node in each I/O group. However, for maximum redundancy, use one redundant ac-power switch to power each node in the cluster.

Some SAN Volume Controller node types have two power supply units. Both power supplies must be connected to the same 2145 UPS-1U, as shown by node A and node B. The SAN Volume Controller 2145-CF8 is an example of a node that has two power supplies. The SAN Volume Controller 2145-8A4 is an example of a node that has a single power supply.

---

## Uninterruptible power supply

The uninterruptible power supply protects a SAN Volume Controller node against blackouts, brownouts, and power surges. The uninterruptible power supply contains a power sensor to monitor the supply and a battery to provide power until an orderly shutdown of the system can be performed.

SAN Volume Controller models use the 2145 UPS-1U.

### 2145 UPS-1U

A 2145 UPS-1U is used exclusively to maintain data that is held in the SAN Volume Controller dynamic random access memory (DRAM) in the event of an unexpected loss of external power. This use differs from the traditional uninterruptible power supply that enables continued operation of the device that it supplies when power is lost.

With a 2145 UPS-1U, data is saved to the internal disk of the SAN Volume Controller node. The uninterruptible power supply units are required to power the SAN Volume Controller nodes even when the input power source is considered uninterruptible.

**Note:** The uninterruptible power supply maintains continuous SAN Volume Controller-specific communications with its attached SAN Volume Controller nodes. A SAN Volume Controller node cannot operate without the uninterruptible power supply. The uninterruptible power supply must be used in accordance with documented guidelines and procedures and must not power any equipment other than a SAN Volume Controller node.

#### 2145 UPS-1U operation

Each SAN Volume Controller node monitors the operational state of the uninterruptible power supply to which it is attached.

If the 2145 UPS-1U reports a loss of input power, the SAN Volume Controller node stops all I/O operations and dumps the contents of its dynamic random access memory (DRAM) to the internal disk drive. When input power to the 2145 UPS-1U is restored, the SAN Volume Controller node restarts and restores the original contents of the DRAM from the data saved on the disk drive.

A SAN Volume Controller node is not fully operational until the 2145 UPS-1U battery state indicates that it has sufficient charge to power the SAN Volume Controller node long enough to save all of its memory to the disk drive. In the event of a power loss, the 2145 UPS-1U has sufficient capacity for the SAN Volume

Controller to save all its memory to disk at least twice. For a fully charged 2145 UPS-1U, even after battery charge has been used to power the SAN Volume Controller node while it saves dynamic random access memory (DRAM) data, sufficient battery charge remains so that the SAN Volume Controller node can become fully operational as soon as input power is restored.

**Important:** Do not shut down a 2145 UPS-1U without first shutting down the SAN Volume Controller node that it supports. Data integrity can be compromised by pushing the 2145 UPS-1U on/off button when the node is still operating. However, in the case of an emergency, you can manually shut down the 2145 UPS-1U by pushing the 2145 UPS-1U on/off button when the node is still operating. Service actions must then be performed before the node can resume normal operations. If multiple uninterruptible power supply units are shut down before the nodes they support, data can be corrupted.

### **Connecting the 2145 UPS-1U to the SAN Volume Controller**

To provide redundancy and concurrent maintenance, you must install the SAN Volume Controller nodes in pairs.

For connection to the 2145 UPS-1U, each SAN Volume Controller of a pair must be connected to only one 2145 UPS-1U.

**Note:** A cluster can contain no more than eight SAN Volume Controller nodes. The 2145 UPS-1U must be attached to a source that is both single phase and 200-240 V. The 2145 UPS-1U has an integrated circuit breaker and does not need external protection.

| SAN Volume Controller provides a cable bundle for connecting uninterruptible  
| power supply to a node. For 2145-8F2, 2145-8F4, 2145-8G4, and 2145-8A4, this is  
| a single power cable plus a serial cable. For 2145-CF8, this is a dual-power cable  
| plus serial cable. This cable is used to connect both power supplies of a node to  
| the same uninterruptible power supply.

The SAN Volume Controller software determines whether the input voltage to the uninterruptible power supply is within range and sets an appropriate voltage alarm range on the uninterruptible power supply. The software continues to recheck the input voltage every few minutes. If it changes substantially but remains within the permitted range, the alarm limits are readjusted.

**Note:** The 2145 UPS-1U is equipped with a cable retention bracket that keeps the power cable from disengaging from the rear panel. See the related documentation for more information.

#### **Related reference**

“2145 UPS-1U connectors and switches” on page 52  
The 2145 UPS-1U has external connectors and dip switches.

### **2145 UPS-1U controls and indicators**

All controls and indicators for the 2145 UPS-1U are located on the front-panel assembly.






Figure 44. 2145 UPS-1U front-panel assembly

- 1** Load segment 2 indicator
- 2** Load segment 1 indicator
- 3** Alarm or service indicator
- 4** On-battery indicator
- 5** Overload indicator
- 6** Power-on indicator
- 7** On/off button
- 8** Test and alarm reset button

Table 8 identifies which status and error LEDs that display on the 2145 UPS-1U front-panel assembly relate to the specified error conditions. It also lists the uninterruptible power supply alert-buzzer behavior.

Table 8. 2145 UPS-1U error indicators

[1] Load2	[2] Load1	[3] Alarm	[4] Battery	[5] Overload	[6] Power-on	Buzzer	Error condition
Green (see Note 1)					Green	(see Note 3)	No errors; the 2145 UPS-1U was configured by the SAN Volume Controller
Green	Amber (see Note 2)				Green		No errors; the 2145 UPS-1U is not yet configured by the SAN Volume Controller
Green	Either on or off		Amber		Green	Beeps for two seconds and then stops	The ac power is over or under limit. The uninterruptible power supply has switched to battery mode.
		Flashing red	Flashing amber	Flashing red	Flashing green	Three beeps every ten seconds	Battery undervoltage
Green	Either on or off	Flashing red			Flashing green	Solid on	Battery overvoltage
		Flashing red	Flashing amber		Flashing green	Solid on	Output wave is abnormal when the charger is open, on battery mode

Table 8. 2145 UPS-1U error indicators (continued)

[1] Load2	[2] Load1	[3] Alarm	[4] Battery	[5] Overload	[6] Power-on	Buzzer	Error condition
		Flashing red	Flashing amber			Solid on	The ac-power output wave is under low limit or above high limit on battery mode
Green	Either on or off		Amber			Beeps for four seconds and then stops	On battery (no ac power)
Green	Either on or off		Flashing amber			Beeps for two seconds and then stops	Low battery (no ac power)
Green	Either on or off			Red	Green	Beeps for one second and then stops	Overload while on line
			Amber	Red		Beeps for one second and then stops	Overload while on battery
Either on or off	Either on or off	Flashing red			Green	Solid on	Fan failure
Either on or off	Either on or off	Flashing red	Amber			Solid on	Battery test fail
		Flashing red		Red		Solid on	Overload timeout
		Flashing red	Amber		Green	Solid on	Over temperature
		Flashing red	Amber	Red	Green		Output short circuit

**Notes:**

1. The green Load2 LED ([1]) indicates that power is being supplied to the right pair of ac-power outlets as seen from the rear of the 2145 UPS-1U.
2. The amber Load1 LED ([2]) indicates that power is being supplied to the left pair of ac-power outlets as seen from the rear of the 2145 UPS-1U. These outlets are not used by the SAN Volume Controller.  
This LED might be illuminated during power-on sequences, but it is typically extinguished by the SAN Volume Controller node that is attached to the 2145 UPS-1U.
3. A blank cell indicates that the light or buzzer is off.

**Load segment 2 indicator:**

The load segment 2 indicator on the 2145 UPS-1U is lit (green) when power is available to load segment 2.

When the load segment 2 indicator is green, the 2145 UPS-1U is running normally and power is available to this segment.

**Related reference**

“2145 UPS-1U connectors and switches” on page 52  
The 2145 UPS-1U has external connectors and dip switches.

***Load segment 1 indicator:***

The load segment 1 indicator on the 2145 UPS-1U is not currently used by the SAN Volume Controller.

**Note:** When the 2145 UPS-1U is configured by the SAN Volume Controller, this load segment is disabled. During normal operation, the load segment 1 indicator is off. A “Do not use” label covers the receptacles.

**Related reference**

“2145 UPS-1U connectors and switches” on page 52  
The 2145 UPS-1U has external connectors and dip switches.

***Alarm indicator:***

If the alarm on the 2145 UPS-1U is flashing red, maintenance is required.

If the alarm is on, go to the 2145 UPS-1U MAP to resolve the problem.

***On-battery indicator:***

The amber on-battery indicator is on when the 2145 UPS-1U is powered by the battery. This indicates that the main power source has failed.

If the on-battery indicator is on, go to the 2145 UPS-1U MAP to resolve the problem.

***Overload indicator:***

The overload indicator lights up when the capacity of the 2145 UPS-1U is exceeded.

If the overload indicator is on, go to MAP 5250: 2145 UPS-1U repair verification to resolve the problem.

***Power-on indicator:***

The power-on indicator is displayed when the 2145 UPS-1U is functioning.

When the power-on indicator is a steady green, the 2145 UPS-1U is active.

***On/off button:***

The on/off button turns the power on or off for the 2145 UPS-1U.

**Turning on the 2145 UPS-1U**

After you connect the 2145 UPS-1U to the outlet, it remains in standby mode until you turn it on. Press and hold the on/off button until the power-on indicator is illuminated (approximately five seconds). On some versions of the 2145 UPS-1U, you might need a pointed device, such as a screwdriver, to press the on/off button. A self-test is initiated that takes approximately 10 seconds, during which time the

indicators are turned on and off several times. The 2145 UPS-1U then enters normal mode.

### Turning off the 2145 UPS-1U

Press and hold the on/off button until the power-on light is extinguished (approximately five seconds). On some versions of the 2145 UPS-1U, you might need a pointed device, such as a screwdriver, to press the on/off button. This places the 2145 UPS-1U in standby mode. You must then unplug the 2145 UPS-1U to turn off the unit.

**Attention:** Do not turn off the uninterruptible power supply before you shut down the SAN Volume Controller node that it is connected to. Always follow the instructions that are provided in MAP 5350 to perform an orderly shutdown of a SAN Volume Controller node.

### Test and alarm reset button:

Use the test and alarm reset button to start the self-test.

To start the self-test, press and hold the test and alarm reset button for three seconds. This button also resets the alarm.

**Note:** This button is applicable to both the 2145 UPS-1U and the 2145 UPS.

### 2145 UPS-1U connectors and switches

The 2145 UPS-1U has external connectors and dip switches.

#### Locations for the 2145 UPS-1U connectors and switches

Figure 45 shows the location of the connectors and switches on the 2145 UPS-1U.




Figure 45. 2145 UPS-1U connectors and switches

- 1** Main power connector
- 2** Communication port
- 3** Dip switches
- 4** Load segment 1 receptacles
- 5** Load segment 2 receptacles

### 2145 UPS-1U dip switches

Figure 46 on page 53 shows the dip switches, which can be used to configure the input and output voltage ranges. Because this function is performed by the SAN Volume Controller software, both switches must be left in the OFF position.




Figure 46. 2145 UPS-1U dip switches

### 2145 UPS-1U ports not used

The 2145 UPS-1U is equipped with ports that are not used by the SAN Volume Controller and have not been tested. Use of these ports, in conjunction with the SAN Volume Controller or any other application that might be used with the SAN Volume Controller, is not supported. Figure 47 shows the 2145 UPS-1U ports that are not used.




Figure 47. Ports not used by the 2145 UPS-1U

- 1** USB interface port
- 2** Network ports
- 3** Load segment receptacles

### 2145 UPS-1U power connector

Figure 48 shows the power connector for the 2145 UPS-1U.




Figure 48. Power connector

## Uninterruptible power-supply environment requirements

An uninterruptible power-supply environment requires that certain specifications for the physical site of the SAN Volume Controller must be met.

### 2145 UPS-1U environment

All SAN Volume Controller models are supported with the 2145 UPS-1U.

## 2145 UPS-1U specifications

The following tables describe the physical characteristics of the 2145 UPS-1U.

### 2145 UPS-1U dimensions and weight

Ensure that space is available in a rack that is capable of supporting the 2145 UPS-1U.

Height	Width	Depth	Maximum weight
44 mm (1.73 in.)	439 mm (17.3 in.)	579 mm (22.8 in.)	16 kg (35.3 lb)
<b>Note:</b> The 2145 UPS-1U package, which includes support rails, weighs 18.8 kg (41.4 lb).			

### Heat output

The 2145 UPS-1U unit produces the following approximate heat output.

Model	Heat output during normal operation	Heat output during battery operation
2145 UPS-1U	10 W (34 Btu per hour)	150 W (512 Btu per hour)

---

## Defining the SAN Volume Controller FRUs

The SAN Volume Controller node, redundant ac-power switch, and uninterruptible power supply each consist of one or more field-replaceable units (FRUs).

### SAN Volume Controller FRUs

The SAN Volume Controller nodes each consist of several field-replaceable units (FRUs), such as the fibre-channel adapter, service controller, disk drive, microprocessor, memory module, CMOS battery, power supply assembly, fan assembly, and the operator-information panel.

#### SAN Volume Controller 2145-CF8 FRUs

Table 9 provides a brief description of each SAN Volume Controller 2145-CF8 FRU.

*Table 9. SAN Volume Controller 2145-CF8 FRU descriptions*

FRU	Description
System board	The system board for the SAN Volume Controller 2145-CF8 node.

Table 9. SAN Volume Controller 2145-CF8 FRU descriptions (continued)

FRU	Description
Fibre-channel small form-factor pluggable (SFP) connector	A compact optical transceiver that provides the optical interface to a fibre-channel cable. It is capable of autonegotiating 2, 4, or 8 gigabits-per-second short-wave optical connection on the 4-port fibre-channel adapter. <b>Note:</b> It is possible that SFPs other than those shipped with the product are in use on the fibre-channel 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.
4-port fibre-channel host bus adapter (HBA)	The SAN Volume Controller 2145-CF8 is connected to the fibre-channel fabric through the fibre-channel HBA, which is located in PCI slot 1. The adapter assembly includes the fibre-channel PCI Express adapter, four short-wave SFPs, the riser card, and bracket.
Service controller	The unit that provides the service functions and the front panel display and buttons.
Service controller cable	The USB cable that is used to connect the service controller to the system board.
Disk drive	The serial-attached SCSI (SAS) 2.5" disk drive.
Disk signal cable	200mm SAS disk-signal cable.
Disk power cable	SAS disk-power cable.
Disk controller	A SAS controller card for the SAS 2.5" disk drive.
Disk controller / USB riser card	The riser card that connects the disk controller to the system board and provides the USB port that the service controller cable connects to.
Disk backplane	The hot-swap SAS 2.5" disk drive backplane.
Memory module	A 4 GB DDR3-1333 2RX4 LP RDIMM memory module
Microprocessor	The microprocessor on the system board. 2.40 GHz quad-core microprocessor.
Power supply unit	An assembly that provides dc power to the SAN Volume Controller 2145-CF8 node.
CMOS battery	A 3.0V battery on the system board that maintains power to back up the system BIOS settings.
Operator-information panel	The information panel that includes the power-control button and LEDs that indicate system-board errors, hard drive activity, and power status.
Operator-information panel cable	A cable that connects the operator-information panel to the system board.

Table 9. SAN Volume Controller 2145-CF8 FRU descriptions (continued)

FRU	Description
Fan assembly	A fan assembly that is used in all the fan positions.
Power cable assembly	The cable assembly that connects the SAN Volume Controller and the 2145 UPS-1U. The assembly consists of two power cables and a serial cable bundled together.
Alcohol wipe	A cleaning wipe.
Thermal grease	Grease that is used to provide a thermal seal between a processor and a heat sink.

## Solid-state drive (SSD) feature FRUs

Table 10 provides a brief description of each SSD feature FRU.

Table 10. Solid-state drive (SSD) feature FRU descriptions

FRU	Description
High-speed SAS adapter	An assembly that includes a high-speed SAS adapter card that provides connectivity up to four solid-state drives (SSDs). The assembly also contains riser card, blanking plate, and screws.
High-speed SAS cable	The cable used to connect the high-speed SAS adapter to the disk back plate.
146 GB solid-state drive (SSD)	146 GB solid-state drive (SSD).

## 2145 UPS-1U FRUs

Table 11 provides a brief description of each 2145 UPS-1U FRU.

Table 11. 2145 UPS-1U FRU descriptions

FRU	Description
2145 UPS-1U assembly	An uninterruptible power supply assembly for use with the 2145-CF8.
2145 UPS-1U assembly	An uninterruptible power supply assembly for use with 2145-8F2, 2145-8F4, 2145-8G4, and 2145-8A4.
Battery pack assembly	The battery that provides backup power to the SAN Volume Controller if a power failure occurs.
Power cable, PDU to 2145 UPS-1U	Input power cable for connecting the 2145 UPS-1U to a rack power distribution unit.
Power cable, mains to UPS-1 (US)	Input power cable for connecting the 2145 UPS-1U to mains power (United States only).

## SAN Volume Controller 2145-8A4 FRUs

Table 12 on page 57 provides a brief description of each SAN Volume Controller 2145-8A4 FRU.



Table 12. SAN Volume Controller 2145-8A4 FRU descriptions

FRU	Description
Memory module	A 2 GB PC2–5300 ECC memory module.
Riser card, PCI Express	An interconnection card that provides the interface between the system board and the 4-port fibre-channel adapter.
4-port fibre-channel host bus adapter (HBA)	The SAN Volume Controller 2145-8A4 is connected to the fibre-channel fabric through the fibre-channel HBA, which is located in PCI slot 1.
Fibre-channel small form-factor pluggable (SFP) connector	A compact optical transceiver that provides the optical interface to a fibre-channel cable. It is capable of operating at up to 4 GB/sec.
System board	The system board for the SAN Volume Controller 2145-8A4 node.
Disk drive back plate with cables	A SATA simple-swap hard disk drive back plate with cables.
Power supply	An assembly that provides dc power to the SAN Volume Controller 2145-8A4 node.
Fan	A single fan.
Drive cage	A cage for the SATA simple-swap hard disk drive.
Hard disk drive	A SATA (serial advanced technology attachment) disk drive for the SAN Volume Controller 2145-8A4.
Service controller	The unit that provides the service functions and the front panel display and buttons.
Operator-information panel	The information panel that includes the power-control button and LEDs that indicate system-board errors, hard drive activity, and power status.
Operator-information panel cable	A cable that connects the operator-information panel to the system board.
Air baffle	An apparatus that redirects or contains air flow to keep the computer components cool.
Microprocessor	The microprocessor on the system board.
CMOS battery	A 3.0V battery on the system board that maintains power to backup the system BIOS settings.
Heat-sink assembly retention module	The unit that is used to install the heat-sink assembly in the SAN Volume Controller 2145-8A4 node.
Heat-sink assembly	An apparatus that is used to dissipate the heat that is generated by the microprocessor.
Input-power cable assembly	The cable assembly that provides the power and signal connections between the SAN Volume Controller 2145-8A4 and the 2145 UPS-1U assembly.

## SAN Volume Controller 2145-8G4 FRUs

Table 13 provides a brief description of each SAN Volume Controller 2145-8G4 FRU.

Table 13. SAN Volume Controller 2145-8G4 FRU descriptions

FRU	Description
System board	The planar for the SAN Volume Controller 2145-8G4 node.
4-port fibre-channel host bus adapter (HBA)	The SAN Volume Controller 2145-8G4 is connected to the fibre-channel fabric through the fibre-channel HBA, which is located in PCI slot 1.
Fibre-channel small form-factor pluggable (SFP) connector	A compact optical transceiver that provides the optical interface to a fibre-channel cable. It is capable of operating at up to 4 GB/sec.
Riser card, PCI Express	An interconnection card that provides the interface between the system board and the 4-port fibre-channel adapter.
Service controller	The FRU that provides the service functions and the front panel display and buttons.
Disk drive	A SATA (serial advanced technology attachment) disk drive for the SAN Volume Controller 2145-8G4.
Disk drive cage assembly	A SATA disk drive cage assembly for the SAN Volume Controller 2145-8G4.
Disk-drive backplane	A SATA disk drive cable assembly with back plate.
Memory module	An ECC DRR2 memory module
Microprocessor	The microprocessor on the system board.
Power supply assembly	An assembly that provides dc power to the SAN Volume Controller 2145-8G4.
Power backplane	An assembly that provides a power interface between the system board and the power supply assembly.
CMOS battery	A 3.0V battery on the system board that maintains power to back up the system BIOS settings.
Front panel signal cable	A ribbon cable that connects the operator-information panel to the system board.
Operator-information panel	The information panel that includes the power control button and the light path diagnostics LEDs.
Fan assembly	A fan assembly containing two fans, which is used in all the fan positions.
Input-power cable assembly	The cable assembly that provides the power and signal connections between the SAN Volume Controller 2145-8G4 and the 2145 UPS-1U assembly.

## SAN Volume Controller 2145-8F4 FRUs

Table 14 provides a brief description of each SAN Volume Controller 2145-8F4 FRU.

Table 14. SAN Volume Controller 2145-8F4 FRU descriptions

FRU	Description
Frame assembly	A complete SAN Volume Controller 2145-8F4 with the exception of the fibre-channel cards and the service controller.
4-port fibre-channel host bus adapter (HBA)	The SAN Volume Controller 2145-8F4 is connected to the fibre-channel fabric through the fibre-channel HBA. The card assembly is located in PCI slot 2. It is not permitted to install a fibre-channel card in PCI slot 1 when the card is installed.
Fibre-channel small form-factor pluggable (SFP) connector	A compact optical transceiver that provides the optical interface to a fibre-channel cable. It is capable of operating at up to 4 GB/sec.
Riser card, PCI Express	An interconnection card that provides the interface between the system board and the 4-port fibre-channel adapter.
Service controller	The FRU that provides the service functions and the front panel display and buttons.
Disk drive assembly	A SATA (serial advanced technology attachment) disk drive assembly for the SAN Volume Controller 2145-8F4.
Memory module	A 1 GB ECC DRR2 memory module.
Microprocessor	The microprocessor on the system board.
Voltage regulator module (VRM)	The VRM of the microprocessor.
Power supply assembly	An assembly that provides dc power to the SAN Volume Controller 2145-8F4.
Power backplane	An assembly that provides a power interface between the system board and the power supply assembly.
CMOS battery	A 3.0V battery on the system board that maintains power to backup the system BIOS settings.
Fan power cable	A kit that provides the cables for connecting the fan backplanes to the system board.
Front panel signal cable	A ribbon cable that connects the operator-information panel to the system board.
Fan backplane	A kit that provides all fan holder and fan backplane assemblies.
Operator-information panel	The information panel that includes the power-control button and the light path diagnostics LEDs.
Fan, 40×40×28	The single fan assemblies located in fan positions 1 - 3.
Fan, 40×40×56	The double fan assemblies located in fan positions 4 - 7.

Table 14. SAN Volume Controller 2145-8F4 FRU descriptions (continued)

FRU	Description
Input-power cable assembly	The cable assembly that provides the power and signal connections between the SAN Volume Controller 2145-8F4 and the 2145 UPS-1U assembly.

## SAN Volume Controller 2145-8F2 FRUs

Table 15 provides a brief description of each SAN Volume Controller 2145-8F2 FRU.

Table 15. SAN Volume Controller 2145-8F2 FRU descriptions

FRU	Description
Frame assembly	A complete SAN Volume Controller 2145-8F2 with the exception of the fibre-channel cards and the service controller.
Fibre-channel host bus adapter (HBA) (full height)	The SAN Volume Controller 2145-8F2 is connected to the fibre-channel fabric through the fibre-channel HBA. The full height card assembly is located in PCI slot 2.
Fibre-channel small form-factor pluggable (SFP) connector	A compact optical transceiver that provides the optical interface to a fibre-channel cable. Its maximum speed is limited to 2 GB/sec by the fibre-channel adapter.
Riser card, PCI (full height)	An interconnection card that provides the interface between the system board and the PCI card in slot 2.
Fibre-channel HBA (low profile)	The SAN Volume Controller 2145-8F2 is connected to the fibre-channel fabric through the fibre-channel HBA. The low profile card assembly is located in PCI slot 1.
Riser card, PCI (low profile)	An interconnection card that provides the interface between the system board and the PCI card in slot 1.
Service controller	The FRU that provides the service functions and the front panel display and buttons.
Disk drive assembly	A SATA (serial advanced technology attachment) disk drive assembly for the SAN Volume Controller 2145-8F2.
Memory module	A 1 GB ECC DRR2 memory module.
Microprocessor	The microprocessor on the system board.
Voltage regulator module (VRM)	The VRM of the microprocessor.
Power supply assembly	An assembly that provides DC power to the SAN Volume Controller 2145-8F2
Power backplane	An assembly that provides a power interface between the system board and the power supply assembly.
CMOS battery	A 3.0V battery on the system board that maintains power to backup the system BIOS settings.

Table 15. SAN Volume Controller 2145-8F2 FRU descriptions (continued)

FRU	Description
Fan power cable	A kit that provides the cables for connecting the fan backplanes to the system board.
Front panel signal cable	A ribbon cable that connects the operator-information panel to the system board.
Fan backplane	A kit that provides all fan holder and fan backplane assemblies.
Operator-information panel	The information panel that includes the power control button and the light path diagnostics LEDs.
Fan, 40×40×28	The single fan assemblies located in fan positions 1-3.
Fan, 40×40×56	The double fan assemblies located in fan positions 4-7
Input-power cable assembly	The cable assembly that provides the power and signal connections between the SAN Volume Controller 2145-8F2 and the 2145 UPS-1U assembly.

## Redundant ac-power switch FRUs

The redundant ac-power switch consists of a single field replaceable unit (FRU).

FRU	Description
Redundant ac-power switch assembly	The redundant ac-power switch and its input power cables.

## Uninterruptible power supply FRUs

The uninterruptible power supply consists of several field-replaceable units (FRUs), such as the battery assembly, battery plate, an electronics assembly, the front panel, frame, and input power cable.

Table 16 provides a brief description of each Uninterruptible power supply FRU.

Table 16. Uninterruptible power supply FRU descriptions

FRU	Description
Battery assembly	The battery that provides backup power to the SAN Volume Controller if a power failure occurs. This FRU is part of the uninterruptible power supply.
Battery plate	The cover plate for the 2145 UPS-1U battery pack assembly.
Input power cable, power distribution unit to the uninterruptible power supply	The power cord for the 2145 UPS-1U.
Front panel	A removable FRU for the 2145 UPS-1U.
Uninterruptible power supply	This FRU includes the frame of the uninterruptible power supply and all the FRUs contained within that frame.



## Chapter 3. Using the SAN Volume Controller Console and CLI

The SAN Volume Controller Console is a Web-browser based GUI that is used to manage the cluster. The SAN Volume Controller command-line interface (CLI) is a collection of commands that you can use to manage SAN Volume Controller clusters.

You can use the SAN Volume Controller Console to monitor and maintain the configuration of storage that is associated with SAN Volume Controller clusters. You can also perform service procedures from the SAN Volume Controller Console.

The SAN Volume Controller Console is installed either on an IBM System Storage Productivity Center (SSPC) or, in previous releases, on a master console server. You can access the SAN Volume Controller Console directly from the server where it is installed or remotely by pointing your Web browser to the server where it is installed. Your Web browser must be connected to the same Ethernet that is used by the server where the SAN Volume Controller Console is running.

Perform the following steps to access the SAN Volume Controller Console application from the server where the SAN Volume Controller Console is installed:

1. Launch the SAN Volume Controller Console to display the SAN Volume Controller Console Welcome panel.

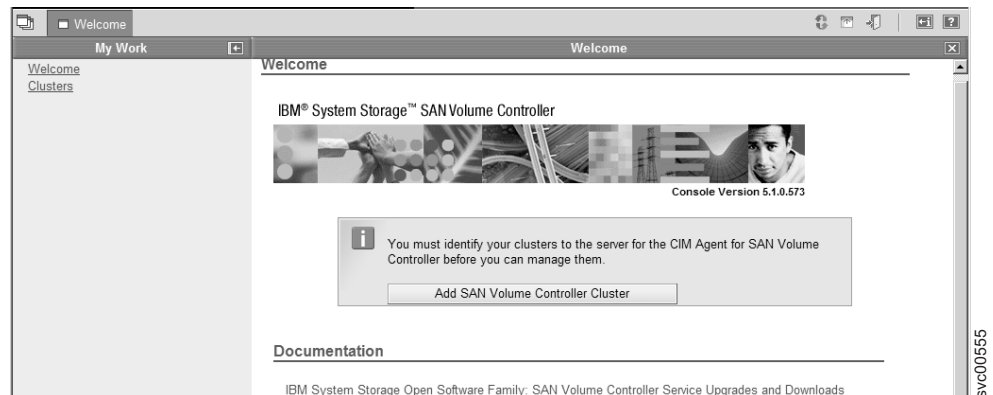


Figure 49. The Add SAN Volume Controller Cluster button on the Welcome panel

2. When you access the SAN Volume Controller Console for the first time, you must identify a cluster that you want to service. Click **Add SAN Volume Controller Cluster** to add the cluster that you want to access. You are shown the Adding a Cluster panel.

**Note:** The next time you access the SAN Volume Controller Console, you do not have to add a cluster. The SAN Volume Controller Console saves the cluster IP address that you typed.

3. Type the IP address of the cluster that you want to service from the SAN Volume Controller Console. Click **OK**. You are shown the sign-on page for the cluster that you added.




Figure 50. Sign on to Cluster panel

4. Sign on with the SAN Volume Controller Console user name and password that is provided by the customer. To analyze and resolve problems, the user name must be authorized for either the service or administrator role. Click **OK**. The Viewing Clusters panel is shown.

**Note:** If the cluster status is not OK, see Table 17 and use the Availability Status message to determine the corrective action that you must take.

Table 17. Corrective actions when the Availability Status is not OK

Availability Status	Action
Unauthenticated	The user is not authenticated to the cluster or the remote authentication server is not available.
No Contact	The IP of the cluster is not reachable, which means the cluster is down or the network card is broken, or the CIMOM service is not available. Verify that the cluster and network are available and that the CIMOM service is running on the cluster. Click <b>Refresh</b> until the status changes to OK.

5. Select the SAN Volume Controller cluster that you want to service and select **Launch the SAN Volume Controller Console** from the **Launch the SAN Volume Controller Console** drop-down list.

**Note:** If the panel does not display, make sure that your Web browser does not have a pop-up blocker enabled.

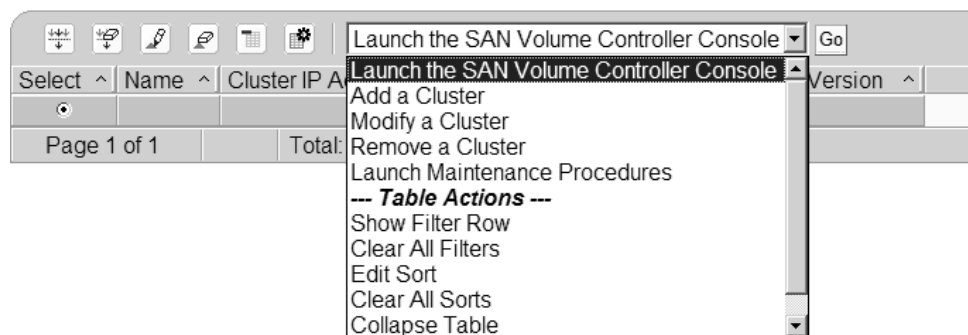


Figure 51. Launch the SAN Volume Controller task

6. Click **Go**. The Welcome panel for the cluster opens in a new window. When the Availability Status for the selected cluster is shown as Unauthenticated, click **Go**. The logon panel for the cluster opens.

From this menu, you can perform several tasks, including the following procedures:



- View the status of a node and the node ports
- Delete a node from a cluster
- Add a node to a cluster
- Shut down a node or a cluster
- View the status of a virtual disk (VDisk)
- Verify and repair VDisk copies
- Repair an offline space-efficient VDisk
- View the status of a managed disk (MDisk)
- View the fibre-channel fabric connections
- View the feature log
- View and update license settings
- List and save dump data
- View and analyze the error log
- Start maintenance

---

## Accessing the SAN Volume Controller CLI

If you must enter and run command-line instructions, you can access the SAN Volume Controller command-line interface (CLI) from the server where the SAN Volume Controller Console is installed.

Perform the following steps to access the CLI:

1. Start PuTTY by clicking **Start** → **All Programs** → **PuTTY** → **PuTTY** or by going to the directory and double-clicking the putty.exe file. If the server is configured to use PuTTY for command-line access, the session settings might be saved. The SAN Volume Controller Console installation package includes PuTTY, which you can use to connect to the SAN Volume Controller cluster over either IPv4 or IPv6. You must use PuTTY 0.60 or a higher version to connect to a SAN Volume Controller cluster that has an IPv6 address.

**Note:** The cluster can have two management IP addresses; therefore, you might see more than one saved session.

If the cluster is not listed under Saved sessions, ask the customer to set up PuTTY for command-line access and then restart this procedure. The PuTTY Configuration panel for the **Session** category contains the details that need to be filled in before PuTTY can open a session at all:

- The **Host Name** field is where you type the name or the IP address of the server you want to connect to.
- The **Connection type** field lets you choose what type of connection you want to make.
- The **Port** field is where you specify which port number on the server to connect to.

Perform the following steps if the cluster that you need to access is listed under Saved Sessions:

- a. Select the menu item for the cluster.
  - b. Click **Load**.
  - c. Click **Open**.
2. If you are using SSH to connect to a server for the first time, you see the PuTTY Security Alert message. Click **Yes**. The command-line window for service is displayed.

```
login as: service
Authenticating with public key "rsa-key-20030213"
IBM_2145:service>
```

3. Type the user name `service` at the login prompt. When the service prompt is displayed, you can use the CLI to issue commands.
4. Issue a command following the service prompt to display information about the SAN Volume Controller. For example, issue the following command to view the current status of the nodes that are used by the SAN Volume Controller cluster:

```
svcinfolsnode -nohdr -delim :
```

The current status of the nodes used by the SAN Volume Controller cluster is displayed.

```
login as: service
Authenticating with public key "rsa-key-20030213"
IBM_2145:service>svcinfolsnode -nohdr -delim :
1:node1:10L3ANP:50050768010000F6:online:0:io_grp0:no:202378101C0D17A0
2:node2:10L3BNZ:5005076801000184:online:0:io_grp0:yes:202378101C0D27AA
3:node3:10L3BNZ:5005076801000101:online:1:io_grp1:no:202378101C0D27AA
4:node4:10L3ANP:5005076801000147:online:1:io_grp1:no:202378101C0D17A0
IBM_2145:service>
```

5. Issue the following command:

```
svcinfolcluster clustername
```

where *clustername* represents the name of the cluster for which you want to list the details.

You are set up to use the CLI from the server where the SAN Volume Controller Console is installed.

---

## Viewing the node status

You can view the properties for a node from the Viewing General Details panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the node properties:

1. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed.
2. Click the name of the node for which you want to view detailed information. The Viewing General Details panel is displayed.

3. In the General navigation area, click **Fibre Channel Ports** to view the worldwide port name (WWPN) details. The Viewing Fibre Channel Port Details panel is displayed. You are shown the attributes and values for the node, which includes the state, WWNN, and I/O group to which it belongs.
4. Click **Vital Product Data** to view the node hardware details. The Viewing Vital Product Data panel is displayed.
5. Click **Ethernet Ports** to view the iSCSI port details. The Viewing Node Ethernet Port panel is displayed.
6. Click **Ethernet Ports–IP** to view the iSCSI port IP address details. The Viewing Node Ethernet Port–IP panel is displayed.
7. Click **Close** to close the panel.

---

## Checking the status of the node using the CLI

You can use the command-line interface (CLI) to check the status of the node.

Issue the following CLI command to check the status of the node:

```
svcinfo lsnode -delim : nodename
```

The following output shows what is displayed when you issue the `svcinfo lsnode -delim : node1` command:

```
id:1
name:node1
UPS_serial_number:10L3ANP
WWNN:5005676801000013
status:online
IO group id:0
IO group name:io_grp0
partner node id:2
partner node name:node2
config node:yes
UPS_unique_id:202378101C0D17A0
port id:5005676801100013
port status:active
port id:5005676801200013
port status:active
port id:5005676801300013
port status:active
port id:5005676801400013
port status:active
```

The characteristics for the node are listed in the output, including its status of *online*. You might also see the following status: *offline*, *adding*, or *deleting*.

For more information about what commands you can use, see the *IBM System Storage SAN Volume Controller Command-Line Interface User's Guide*.

---

## Viewing the state of the fibre-channel ports

You can view the state of the fibre-channel ports from the Viewing Fibre Channel Port Details panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to view the status of the fibre-channel ports:

1. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed.
2. Click the name of the node for which you want to view port status. The Viewing General Details panel is displayed.
3. Click **Fibre Channel Ports** to view the status of the ports and the worldwide port name (WWPN) details. The Viewing Fibre Channel Port Details panel is displayed. The following states are possible:

**Failure**

Indicates that an error condition exists on this port.

**Online**

Indicates that the port is active.

**Offline**

Indicates that the port is inactive or has failed.

**Not Present**

Indicates that the port is not installed.

**Unknown**

Indicates that the node state cannot be determined.

4. Click **Close** to close the panel.

---

## Checking the status of the fibre-channel ports using the CLI

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

Issue the following CLI command to check the status of the fibre-channel ports:

```
svcinfo lsnode -delim : nodename
```

where *nodename* is the name of the node whose ports you want to check.

The following output shows what is displayed when you issue the `svcinfo lsnode -delim : nodename` command:

```
id:1
name:node1
UPS_serial_number:10L3ANP
WWNN:5005676801000013
status:online
IO_group_id:0
IO_group_name:io_grp0
partner_node_id:2
partner_node_name:node2
config_node:yes
UPS_unique_id:202378101C0D17A0
port_id:5005676801100013
port_status:active
port_id:5005676801200013
port_status:active
port_id:5005676801300013
port_status:active
port_id:5005676801400013
port_status:active
```

The previous example shows information for the named node on line two. You can see the port ID and the status of the port in the last eight lines of the example. In this example, the last eight lines show a port status of active. You might also see the following statuses: not installed, failed, or inactive.

For more information about using the command-line interface, see *IBM System Storage SAN Volume Controller Command-Line Interface User's Guide*.

---

## Viewing the state of Ethernet ports

You can view the state of Ethernet ports for a node that connect to iSCSI hosts.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to view the state of the Ethernet ports:

1. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed.
2. Click the name of the node for which you want to view the state of Ethernet ports for the selected node. The Viewing General Details panel is displayed.
3. Click **Ethernet Ports** to view the state and details regarding the Ethernet ports that are associated with the selected node. The Viewing Ethernet Port Details panel is displayed. The following states are possible:

### Online

Indicates that the node Ethernet port is operational, assigned to a cluster, and is connected TCP/IP network.

### Offline

Indicates that the node Ethernet port is not operational, either having lost connection to the cluster of the TCP/IP network.

### Not Configured

Indicates that the node Ethernet port is currently not configured.

4. Click **Close** to close the panel.

---

## Deleting a node from a cluster using the SAN Volume Controller Console

You might have to remove a node from a cluster if the node has failed and is being replaced with a new node or if the repair that has been performed has caused that node to be unrecognizable by the cluster.

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 virtual disks 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. The cache is flushed before the node is deleted to prevent data loss if a failure occurs on the other node in the I/O group.

Before deleting a node from the cluster, record the node serial number, worldwide node name (WWNN), all worldwide port names (WWPNs), and the I/O group that the node is currently part of. Recording this node information can avoid data corruption if the node is re-added to the cluster at a later time.

**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 virtual disks (VDisks) are already degraded before you delete a node, redundancy to the VDisks is degraded. Removing a node might result in a loss of access to data and data loss.
- Removing the last node in the cluster destroys the cluster. Before you delete the last node in the cluster, ensure that you want to destroy the cluster.
- When you delete 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 VDisks assigned to the I/O group, you cannot delete the node from the cluster if the node is online. You must back up or migrate all data that you want to save before you delete the node. If the node is offline, you can delete the node.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to delete a node from a cluster:

1. Unless this is the last node in the cluster, turn the power off to the node that you are removing using the Shut Down a Node option on the SAN Volume Controller Console. This step ensures that the multipathing device driver does not rediscover the paths that are manually removed before you issue the delete node request.

**Attention:**

- When you remove the configuration node, the configuration function moves to a different node within the cluster. This process can take a short time, typically less than a minute. The SAN Volume Controller Console reattaches to the new configuration node transparently.
  - If you turn the power on to the node that has been removed and it is still connected to the same fabric or zone, it attempts to rejoin the cluster. At this point, the cluster tells the node to remove itself from the cluster and the node becomes a candidate for addition to this cluster or another cluster.
  - If you are adding this node into the cluster, 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.
2. In the portfolio, click **Work with Nodes** → **Nodes**. The Viewing Nodes panel is displayed.
  3. Find the node that you want to delete.

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

If the node that you want to delete is shown as `Online`, deleting the node can result in the dependent VDisks to also go offline. Verify whether or not the node has any dependent VDisks.
  4. To check for dependent VDisks before attempting to delete the node, select the node and click **Show Dependent VDisks** from the drop-down menu.

If any VDIs are listed, you should determine why and if access to the VDIs is required while the node is deleted from the cluster. If the VDIs are assigned from MDisk groups that contain solid-state drives (SSDs) that are located in the node, you should check why the VDisk mirror, if it is configured, is not synchronized. There can also be dependent VDIs because the partner node in the I/O group is offline. Fabric issues can also prevent the VDisk from communicating with storage systems. You should resolve these problems before continuing with the node deletion.

5. Select the node that you want to delete and select **Delete a Node** from the task list. Click **Go**. The Deleting Node from Cluster panel is displayed.
6. Click **OK** to delete the node. Before a node is deleted the SAN Volume Controller checks to determine if there are any virtual disks (VDIs) that depend on that node. If the node that you selected contains VDIs within the following situations, VDIs go offline and become unavailable if the node is deleted:
  - The node contains solid-state drives (SSD) and also contains the only synchronized copy of a mirrored VDisk
  - The other node in the I/O group is offline

If you select a node to delete that has these dependencies, another panel displays confirming the deletion. To delete the node in this case, click **Force Delete** message panel that displays.

---

## Deleting a node from a cluster using the CLI

You can use the command-line interface (CLI) to remove a node from a cluster.

After the node is deleted, the other node in the I/O group enters write-through mode until another node is added back into the I/O group.

By default, the **rmnode** command flushes the cache on the specified node before taking the node offline. When operating in a degraded state, the SAN Volume Controller ensures that data loss does not occur as a result of deleting the only node with the cache data.

### Attention:

- If you are removing a single node and the remaining node in the I/O group is online, the data can be exposed to a single point of failure if the remaining node fails.
- If both nodes in the I/O group are online and the VDIs are already degraded prior to deleting the node, redundancy to the VDIs is already degraded and loss of access to data and loss of data might occur if the **-force** option is used.
- Removing the last node in the cluster destroys the cluster. Before you delete the last node in the cluster, ensure that you want to destroy the cluster.
- To take the specified node offline immediately without flushing the cache or ensuring data loss does not occur, run the **rmnode** command with the **-force** parameter. The **force** parameter forces continuation of the command even though any node-dependent VDIs will be taken offline. Use the **force** parameter with caution; access to data on node-dependent VDIs will be lost.

Perform the following steps to delete a node:

1. If you are deleting the last node in an I/O group, determine the VDIs that are still assigned to this I/O group:
  - a. Issue the following CLI command to request a filtered view of the VDIs:

```
svcinfolsvdisk -filtervalue IO_group_name=name
```

Where *name* is the name of the I/O group.

- b. Issue the following CLI command to list the hosts that this VDisk is mapped to:

```
svcinfo lsvdiskhostmap vdiskname/id
```

Where *vdiskname/id* is the name or ID of the VDisk.

- If VDIsks are assigned to this I/O group that contain data that you want to continue to access, back up the data or migrate the VDIsks to a different (online) I/O group.
2. If this is *not* the last node in the cluster, turn the power off to the node that you intend to remove. This step ensures that the multipathing device driver, such as the subsystem device driver (SDD), does not rediscover the paths that are manually removed before you issue the delete node request.

**Attention:**

- a. If you are removing the configuration node, the **rmnode** command causes the configuration node to move to a different node within the cluster. This process may take a short time, typically less than a minute. The cluster IP address remains unchanged, but any SSH client attached to the configuration node might need to reestablish a connection. The SAN Volume Controller Console reattaches to the new configuration node transparently.
  - b. If you turn on the power to the node that has been removed and it is still connected to the same fabric or zone, it attempts to rejoin the cluster. At this point, the cluster causes the node to remove itself from the cluster and the node becomes a candidate for addition to this cluster or another cluster.
  - c. If you are adding this node into the cluster, 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.
  - d. In a service situation, a node should normally be added back into a cluster using the original node name. As long as the partner node in the I/O group has not been deleted too, this is the default name used if **-name** is not specified.
3. Before you delete the node, update the multipathing device driver configuration on the host to remove all device identifiers that are presented by the VDisk that you intend to remove. If you are using the subsystem device driver (SDD), the device identifiers are referred to as virtual paths (vpaths).

**Attention:** Failure to perform this step can result in data corruption.

See the *IBM System Storage Multipath Subsystem Device Driver User's Guide* for details about how to dynamically reconfigure SDD for the given host operating system.

4. Issue the following CLI command to delete a node from the cluster:

**Attention:** Before you delete the node: The **rmnode** command checks for node-dependent VDIsks, which are not mirrored at the time that the command is run. If any node-dependent VDIsks are found, the command stops and returns a message. To continue removing the node despite the potential loss of data, run the **rmnode** command with the **-force** parameter. Alternatively, follow these steps before you remove the node to ensure that all VDIsks are mirrored:

- a. Run the **lsnodedependentvdisks** command.
- b. For each node-dependent VDisk that is returned, run the **lsvdisk** command.
- c. Ensure that each VDisk returns in-sync status.

```
svctask rmnode node_name_or_id
```

Where *node\_name\_or\_id* is the name or ID of the node.



**Note:** Before removing a node, the command checks for any node-dependent VDIs that would go offline. If the node that you selected to delete contains a solid-state drive (SSD) that has dependent VDIs, VDIs that use the SSDs go offline and become unavailable if the node is deleted. To maintain access to VDI data, mirror these VDIs before removing the node. To continue removing the node without mirroring the VDIs, specify the **force** parameter.

#### Related tasks

“Adding a node to the cluster using the CLI” on page 76

You can use the command-line interface (CLI) to add a node that has either been removed or rejected by a cluster, into the cluster.

---

## Adding nodes to a cluster

Before you add a node to a cluster, you must make sure that the switch zoning is configured such that the node being added is in the same zone as all other nodes in the cluster.

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 being added is in the same VSAN/zone.

### Special procedures when adding a node to a cluster

If you are adding a node that has been used previously, either within a different I/O group within this cluster or within a different cluster, consider the following situations before adding the node. If you add a node to the cluster without changing its worldwide node name (WWNN), hosts might detect the node and use it as if it were in its old location. This could cause the hosts to access the wrong virtual disks (VDIs).

- If you are re-adding a node back to the same I/O group after a service action required the node to be deleted from the cluster and the physical node has not changed, then no special procedures are required and the node can be added back to the cluster.
- In a service situation, a node should normally be added back into a cluster using the original node name. As long as the partner node in the I/O group has not been deleted too, this is the default name used if **-name** is not specified.
- If you are replacing a node in a cluster, either because of a node failure or an upgrade, 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 cluster.
- If you are creating a new I/O group in the cluster and are adding new node, then there are no special procedures since this node has never been added to a cluster before and its WWNN that has not existed before.
- If you are creating a new I/O group in the cluster and are adding new node, but this node has been added to a cluster before, then host system might still be configured to the node WWPNs and the node might still be zoned in the fabric. Since you cannot change the WWNN for the node, you must ensure other components in your fabric are configured correctly. Verify that any host that was previously configured to use the node has been correctly updated. Also verify the fabric zoning does not currently include this node in any zones.
- If the node you are adding was previously replaced, either for a node repair or upgrade, you might have used the node's WWNN 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.

**Note:** Consider the following information when 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 SAN Volume Controller virtual disk (VDisk). This association uses an identifier (UID) which is unique to the VDisk and is never reused. The UID allows multipathing device drivers to directly associate vpaths with VDIs.
- Multipathing device drivers operate within a protocol stack that contains disk and fibre channel device drivers that allow it to communicate with the SAN Volume Controller 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 using the same WWNN and LUN numbers that were previously used.
- Multipathing device drivers do not check the association of the VDisk with the vpath on every I/O operation that it performs.

## Adding nodes to a cluster using the SAN Volume Controller Console

### Attention:

1. If you are re-adding a node to the SAN, ensure that you are adding the node to the same I/O group from which it was removed. Failure to do this can result in data corruption. You must use the information that was recorded when the node was originally added to the cluster. If you do not have access to this information, call the IBM Support Center to add the node back into the cluster without corrupting the data.
2. In a service situation, a node should normally be added back into a cluster using the original node name. As long as the partner node in the I/O group has not been deleted too, this is the default name used if **-name** is not specified.
3. 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 cluster. You must ensure that the LUNs are the same before you add the new node to the cluster.
4. LUN masking for each LUN must be identical on all nodes in a cluster. You must ensure that the LUN masking for each LUN is identical before you add the new node to the cluster.
5. You must ensure that the model type of the new node is supported by the SAN Volume Controller software level that is currently installed on the cluster. If the model type is not supported by the SAN Volume Controller software level, upgrade the cluster to a software level that supports the model type of the new node. See the following Web site for the latest supported software levels:  
[www.ibm.com/storage/support/2145](http://www.ibm.com/storage/support/2145)

Each node in an I/O group must:

- Be connected to a different uninterruptible power supply.
- Have a unique name. If you do not provide a name, the cluster assigns a default name to the object.

**Note:** Whenever possible you must provide a meaningful name for objects to make identifying that object easier in the future.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to add a node to a cluster:

1. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed.
2. Select **Add a node** from the task list and click **Go**. The Adding a Node to a Cluster panel is displayed.
3. Select the node that you want to add to the cluster from the **Available Candidate Nodes** list.
4. Select the I/O group from the **I/O Groups** list.
5. In the **Node Name** field, type the name that you want to assign to the node.
6. Click **OK**.
7. If you are adding a node into the cluster for the first time, record the following information:
  - Node serial number
  - All WWPNs

- The I/O group that the node belongs to

**Important:** You need this information to avoid possible data corruption if you have to remove and re-add the node to the cluster.

---

## Adding a node to the cluster using the CLI

You can use the command-line interface (CLI) to add a node that has either been removed or rejected by a cluster, into the cluster.

**Attention:** Before you add a new node to a cluster, make sure that you configure the switch zoning so that the node you are adding is in the same zone as all other nodes in the cluster. If you are replacing a node and the switch is zoned by worldwide port name (WWPN) rather than by switch port, you must follow the service instructions carefully to continue to use the same WWPNs.

Perform the following steps to add a node to a cluster:

1. Issue the following CLI command to list the node candidates:

```
svcinfolnsnodecandidate
```

The following output is an example of what you might see after you issue the `svcinfolnsnodecandidate` command:

id	panel_name	UPS_serial_number	UPS_unique_id	hardware
50050768010037DA	104615	10004BC047	20400001124C0107	8G4
id	panel_name	UPS_serial_number	UPS_unique_id	hardware
5005076801000149	106075	10004BC031	20400001124C00C1	8G4

2. Issue the following CLI command to add the node:

```
svctask addnode -panelname panel_name -name new_name_arg -iogrp  
iogroup_name
```

where *panel\_name* is the name that is noted in step 1 (in this example, the panel name is 000279). This is the number that is printed on the front panel of the node that you are adding back into the cluster; *new\_name\_arg* is optional to specify a name for the new node; *iogroup\_name* is the I/O group that was noted when the previous node was deleted from the cluster.

**Note:** In a service situation, a node should normally be added back into a cluster using the original node name. As long as the partner node in the I/O group has not been deleted too, the default name is used if `-name` is not specified.

The following example shows the command that you might issue:

```
svctask addnode -panelname 000279 -name newnode -iogrp io_grp1
```

The following output is an example of what you might see:

```
Node, id [newnode], successfully added
```

**Attention:** If more than one candidate node exists, ensure that the node that you add into an I/O group is the same node that was deleted from that I/O group. Failure to do so might result in data corruption. If you are uncertain about which candidate node belongs to the I/O group, shut down all host systems that access this cluster before you proceed. Reboot each system when you have added all the nodes back into the cluster.

3. Issue the following CLI command to ensure that the node was added successfully:

```
svcinfolnode
```

The following output is an example of what you might see when you issue the `svcinfolnode` command:

id	name	UPS_serial_number	WWNN	status	IO_group_id	IO_group_name	config_node	UPS_unique_id	hardware
1	node1	1000877059	5005076801000EAA	online	0	io_grp0	yes	20400002071C0149	8F2
2	node2	1000871053	500507680100275D	online	0	io_grp0	no	2040000207040143	8F2

All nodes are now online.

#### Related tasks

“Deleting a node from a cluster using the CLI” on page 71

You can use the command-line interface (CLI) to remove a node from a cluster.

---

## Shutting down a node

You can shut down a SAN Volume Controller node from the Shutting Down Node panel.

If you are shutting down the last SAN Volume Controller node in an I/O group, quiesce all I/O operations that are destined for this SAN Volume Controller node. Failure to do so can result in failed I/O operations being reported to your host operating systems.

This task assumes that you have already launched the SAN Volume Controller Console.

Do not disrupt the access of the hosts to their data when shutting down a node.

Follow MAP 5350 to shut down a node by using either the CLI or the SAN Volume Controller Console.

#### Related tasks

“MAP 5350: Powering off a SAN Volume Controller node” on page 385

MAP 5350: Powering off a SAN Volume Controller node helps you power off a single SAN Volume Controller node to perform a service action without disrupting the host’s access to disks.

---

## Shutting down the cluster

If all input power to a SAN Volume Controller cluster must be removed, you must shut down the cluster before the power is removed. If you do not shut down the cluster before turning off input power to the uninterruptible power supply, the SAN Volume Controller detects the loss of power and continues to run on battery power until all data held in memory is saved to the internal disk drive. This increases the time that is required to make the cluster operational when input power is restored

and severely increases the time that is required to recover from an unexpected loss of power that might occur before the uninterruptible power supply batteries can fully recharge.

Shut down a cluster by using either the SAN Volume Controller Console or the CLI.

When input power is restored, you must press the power button on the uninterruptible power supply units before you press the power buttons on the SAN Volume Controller.

#### **Related tasks**

“MAP 5350: Powering off a SAN Volume Controller node” on page 385  
MAP 5350: Powering off a SAN Volume Controller node helps you power off a single SAN Volume Controller node to perform a service action without disrupting the host’s access to disks.

## **Shutting down a cluster**

You can shut down a SAN Volume Controller cluster from the Shutting Down cluster panel.

If you want to remove all input power to a cluster (for example, the machine room power must be shutdown for maintenance), we you should use the Shutdown Down Cluster operation prior to removing power. Note: If you do not shut down the cluster before turning off input power to the uninterruptible power supply units, the cluster completes an emergency shutdown, which is powered from the uninterruptible power supply battery. This method drains power from the uninterruptible power supply needlessly, and the restart of the cluster is delayed while the uninterruptible power supply charges.

When input power is restored to the uninterruptible power supply units, they start to recharge. However, the SAN Volume Controller do not permit any I/O activity to be performed to the virtual disks (VDisks) until the uninterruptible power supply is charged enough to enable all the data on the SAN Volume Controller nodes to be saved in the event of an unexpected power loss. This process might take as long as two hours. Therefore shutting down the cluster prior to removing input power to the uninterruptible power supply units prevents the battery power from being drained and makes it possible for I/O activity to resume as soon as input power is restored.

Before shutting down a cluster, quiesce all I/O operations that are destined for this cluster. Failure to do so can result in failed I/O operations being reported to your host operating systems.

**Attention:** If you are shutting down the entire cluster, you lose access to all VDisks that are provided by this cluster. Shutting down the cluster also shuts down all SAN Volume Controller nodes. This shutdown causes the hardened data to be dumped to the internal hard drive.

Begin the following process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks that are provided by the cluster.

1. Determine which hosts are using the VDisks that are provided by the cluster.
2. Repeat the previous step for all VDisks.

When input power is restored, you must press the power button on the uninterruptible power supply units before you press the power buttons on the SAN Volume Controller nodes.

Perform the following steps to shut down a cluster:

1. Click **Manage Clusters** → **Shut down Cluster** in the portfolio. The Shutting Down cluster panel is displayed.
2. Click **Yes**.

## Shutting down a cluster using the CLI

You can use the command-line interface (CLI) to shut down a cluster.

If you want to remove all input power to a cluster (for example, the machine room power must be shutdown for maintenance), you must shut down the cluster before the power is removed. If you do not shut down the cluster before turning off input power to the uninterruptible power supply, the SAN Volume Controller nodes detect the loss of power and continue to run on battery power until all data that is held in memory is saved to the internal disk drive. This increases the time that is required to make the cluster operational when input power is restored and severely increases the time that is required to recover from an unexpected loss of power that might occur before the uninterruptible power supply batteries have fully recharged.

When input power is restored to the uninterruptible power supply units, they start to recharge. However, the SAN Volume Controller nodes do not permit any I/O activity to be performed to the virtual disks (VDisks) until the uninterruptible power supply is charged enough to enable all the data on the SAN Volume Controller nodes to be saved in the event of an unexpected power loss. This might take as long as two hours. Shutting down the cluster prior to removing input power to the uninterruptible power supply units prevents the battery power from being drained and makes it possible for I/O activity to resume as soon as input power is restored.

Before shutting down a cluster, quiesce all I/O operations that are destined for this cluster. Failure to do so can result in failed I/O operations being reported to your host operating systems.

**Attention:** If you are shutting down the entire cluster, you lose access to all VDisks that are provided by this cluster. Shutting down the cluster also shuts down all SAN Volume Controller nodes. This shutdown causes the hardened data to be dumped to the internal hard drive.

Begin the following process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks that are provided by the cluster.

1. Determine which hosts are using the VDisks that are provided by the cluster.
2. Repeat the previous step for all VDisks.

If input power is lost and subsequently restored, you must press the power button on the uninterruptible power supply units before you press the power buttons on the SAN Volume Controller nodes.

Perform the following steps to shut down a cluster:

1. Issue the following command to shut down a cluster:  
`svctask stopcluster`

The following output is displayed:

```
Are you sure that you want to continue with the shut down?
```

2. Type `y` to shut down the entire cluster.

---

## Listing node-dependent VDIs using the CLI

You can use the command-line interface (CLI) to list the virtual disks (VDIs) that are dependent on the status of a node.

If a node goes offline or is removed from a cluster, all VDIs that are dependent on the node go offline. Before taking a node offline or removing a node from a cluster, run the `lsnode-dependentvdisk` command to identify any node-dependent VDIs.

By default, the `lsnode-dependentvdisk` command also checks all available quorum disks. If the quorum disks are accessible only through the specified node, the command returns an error.

Various scenarios can produce node-dependent VDIs. The following examples are common scenarios in which the `lsnode-dependentvdisk` command will return node-dependent VDIs:

1. The node contains solid-state drives (SSDs) and also contains the only synchronized copy of a mirrored VDI.
2. The node is the only node that can access an MDisk on the SAN fabric.
3. The other node in the I/O group is offline (all VDIs in the I/O group are returned).
4. Pinned data in the cache is stopping the partner node from joining the I/O group.

To resolve (1), allow VDI mirror synchronizations between SSD MDisks to complete. To resolve (2-4), bring any offline MDisks online and repair any degraded paths.

**Note:** The command lists the node-dependent VDIs at the time the command is run; subsequent changes to a cluster require running the command again.

1. Issue the `svcinfolnode-dependentvdisk` CLI command.

The following example shows the CLI format for listing the VDIs that are dependent on node01:

```
svcinfolnode-dependentvdisk -node01 :
```

The following example shows the output that is displayed:

```
vdisk_id      vdisk_name
0             vdisk0
1             vdisk1
```

2. If the `svcinfolnode-dependentvdisk` command returns an error, you must move your quorum disks to MDisks that are accessible through all nodes. Rerun the command until no errors are returned.
3. Reissue the `svcinfolnode-dependentvdisk` command. When the command returns no VDIs, the cluster is free from any node-dependent VDIs.

The following example shows the command syntax for listing the VDIs that are dependent on node01:

```
svcinfolnode-dependentvdisk -node01 :
```

The following example shows the command output if there are no node-dependent VDIs in the cluster:



| vdisk\_id vdisk\_name  
|

---

## Viewing the VDisk status

You must view the status of virtual disks (VDisks) as part of the repair verification procedures.

This task assumes that you have already launched the SAN Volume Controller Console.

| The repair verification is not complete until all VDisks are shown as online. Any  
| VDisks that remain offline, degraded, or excluded might contain errors or are not  
| recognized because of a problem with the SAN environment. If problems still exist  
| on VDisks after the repair actions on the SAN Volume Controller are complete,  
| resolve the disk drive or SAN problems, and then perform repair verification for the  
| SAN Volume Controller to verify that no other problems exist.

Complete the following steps to view the status of VDisks:

1. Click **Work with Virtual Disks** → **Virtual Disks**. The Viewing Virtual Disks panel is displayed.
2. Ensure that all VDisks are online.

### Notes:

1. To repair offline VDisks, see the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*.
2. If you have a degraded VDisk and all of the associated nodes and MDisks are online, call the IBM Support Center for assistance.

---

## Verifying and repairing mirrored VDisk copies

The virtual disk (VDisk) copy verification process checks if data on mirrored VDisk copies match. You can choose repair options if differences are found during the verification process.

**Attention:** Proceed with this task only if all VDisk copies are synchronized.

Use the Verifying VDisk Copies panel to start the VDisk copy verification process for a selected VDisk. If differences are found during verification, you can choose one of the following actions.

- Stop the process when the first difference is found. Select this option if you only want to verify that the mirrored VDisk copies are identical. You can run this option, starting at a different logical block address (LBA) each time to count the number of differences on a VDisk.
- Automatically repair the copy by overwriting sectors with data from the primary VDisk copy. Select the resync option if you are sure that either the primary VDisk copy data is correct or that your host applications can handle incorrect data.
- Create a virtual medium error at the VDisk level. Select 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.

| If a medium error is encountered on one of the copies, the VDisk copy is  
| automatically repaired if the data can be read from another copy.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to verify mirrored VDisk copies:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
2. Select the VDisk to verify and then select **Verify VDisk Copies** from the task list. Click **Go**. The Verifying VDisk Copies panel is displayed.
3. Select the repair action if errors are found and click **OK**. You can also specify an LBA from which to start the verification. Start at different LBAs to count the number of differences on a VDisk.

## Viewing the progress of VDisk copy verification

You can view the progress of verification of one or more mirror copies for a virtual disk (VDisk) from the Viewing Mirror Copy Verification Progress panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the progress of mirror copy verification:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **VDisk Copy Verification** link. The Viewing Mirror Copy Verification Progress panel is displayed.
3. Click **Close** to close the panel.

---

## Validating and repairing mirrored VDisk copies using the CLI

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

**Attention:** Run the **repairvdiskcopy** command only if all VDisk 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 VDisk 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 if you only want to verify that the mirrored VDisk 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 VDisk.

### **-medium**

Use this parameter to convert sectors on all VDisk 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 VDisk copy to the other VDisk copy. The command corrects any differing sectors by copying the sectors from the primary copy to the copies 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 VDisk 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` using this parameter to list all the differences.

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

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

### **Notes:**

1. Only one **repairvdiskcopy** command can run on a VDisk at a time.
2. Once you start the **repairvdiskcopy** command, you cannot use the command to stop processing.
3. The primary copy of a mirrored VDisk 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 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 VDisk copies using the CLI**

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

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

```
svcinfo 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 offline space-efficient VDIs

When a space-efficient virtual disk (VDisk) is taken offline because its metadata is corrupted, you can use the Repairing Space-Efficient VDisk panel to repair the metadata. The repair operation automatically detects corrupted metadata and performs any necessary repair actions.

This task assumes that you have already launched the SAN Volume Controller Console.

Use the Repairing Space-Efficient VDisk panel when directed through maintenance procedures. When the repair operation completes successfully, the error is automatically marked as fixed and the volume is brought back online. If the repair operation fails, an error is logged (error ID 060003) and the volume remains offline.

Once started, the VDisk remains offline for the duration of the repair, but you can move the VDisk to another I/O group.

**Attention:** You can only use this panel to repair a space-efficient VDisk that has reported corrupt metadata.

Perform the following steps to repair the offline space-efficient VDisk:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
2. Select the VDisk to repair and then select **Repair Space-efficient VDisk** from the task list. Click **Go**. The Repairing Space-Efficient VDIs panel is displayed.
3. Select the VDisk copy to repair and click **OK**.

## Viewing the progress of space-efficient VDisk copy repair

You can view the progress of space-efficient virtual disk (VDisk) copy repair from the Viewing Space-Efficient Copy Repair Progress panel.

The time that is needed to complete a space-efficient VDisk copy repair depends on the amount of data that is currently on the copy. The repair process might complete very quickly.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the progress of space-efficient VDisk copy repair:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **Space-Efficient Copy Repair** link. The Viewing Space-Efficient Copy Repair Progress panel is displayed.
3. Click **Close** to close the panel.

---

## Repairing a space-efficient VDisk using the CLI

You can use the **repairsevdiskcopy** command from the command-line interface to repair the metadata on a space-efficient virtual disk (VDisk).

The **repairsevdiskcopy** command automatically detects and repairs corrupted metadata. The command holds the VDisk 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 virtual disk copies in the configuration.

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

**Attention:** Use this command only to repair a space-efficient VDisk that has reported corrupt metadata.

Issue the following command to repair the metadata on a space-efficient VDisk:

```
svctask 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 space-efficient VDisk using the CLI

Issue the `lsrepairsevdiskcopyprogress` command to list the repair progress for space-efficient VDisk copies of the specified VDisk. If you do not specify a VDisk, the command lists the repair progress for all space-efficient copies in the cluster.

**Note:** Only run this command after you run the `svctask repairsevdiskcopy` command, which you must only run as required by the Directed Maintenance Procedures or by IBM support.

---

## Recovering from offline VDIs

You can use the SAN Volume Controller Console to recover from an offline virtual disk (VDisk) after a node or an I/O group has failed.

If you have lost both nodes in an I/O group and have, therefore, lost access to all the VDIs that are associated with the I/O group, you must perform one of the

following procedures to regain access to your VDisks. Depending on the failure type, you might have lost data that was cached for these VDisks and the VDisks are now offline.

#### Data loss scenario 1

One node in an I/O group has failed and failover has 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 written to hard disk. 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 has lost its hardened data, therefore, the node has no way of recognizing that it is part of the cluster.

Perform the following steps to recover from an offline VDisk when one node has earlier-level hardened data and the other node has lost hardened data:

1. Recover the node and add it back into the cluster.
2. Delete all FlashCopy, Metro Mirror, and Global Mirror mappings and relationships that use the offline VDisks.
3. Run the **recovervdisk**, **recovervdiskbyiogrp** or **recovervdiskbycluster** command.
4. Recreate all FlashCopy, Metro Mirror, and Global Mirror mappings and relationships that use the VDisks.

#### Data loss scenario 2

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

Perform the following steps to recover from an offline VDisk when both nodes have lost their hardened data and cannot be recognized by the cluster:

1. Delete all FlashCopy, Metro Mirror, and Global Mirror mappings and relationships that use the offline VDisks.
2. Run the **recovervdisk**, **recovervdiskbyiogrp** or **recovervdiskbycluster** command.
3. Recreate all FlashCopy, Metro Mirror, and Global Mirror mappings and relationships that use the VDisks.

---

## Recovering from offline VDisks using the CLI

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

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

#### Data loss scenario 1

One node in an I/O group has failed and failover has 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 written to hard disk. 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 has lost its hardened data, therefore, the node has no way of recognizing that it is part of the cluster.

Perform the following steps to recover from an offline VDisk when one node has down-level hardened data and the other node has lost hardened data:

1. Recover the node and add it back into the cluster.
2. Delete all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the offline VDisks.
3. Run the `recovervdisk`, `recovervdiskbyiogrp` or `recovervdiskbycluster` command.
4. Recreate all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the VDisks.

### Data loss scenario 2

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

Perform the following steps to recover from an offline VDisk when both nodes have lost their hardened data and cannot be recognized by the cluster:

1. Delete all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the offline VDisks.
2. Run the `recovervdisk`, `recovervdiskbyiogrp` or `recovervdiskbycluster` command.
3. Recreate all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the VDisks.

## Recovering a node and returning it to the cluster using the CLI

After a node or an I/O group fails, you can use the command-line interface (CLI) to recover a node and return it to the cluster.

Perform the following steps to recover a node and return it to the cluster:

1. Issue the following command to verify that the node is offline:
2. Issue the following command to remove the old instance of the offline node from the cluster:

```
svcinfolnode
```

```
svctask rmnode nodename/id
```

Where *nodename/id* is the name or ID of the node.

3. Issue the following command to verify that the node can be seen on the fabric:

```
svcinfolnodecandidate
```

**Note:** Remember the worldwide node names (WWNNs) for each node because you will need them in the following step.

4. If the nodes are repaired by replacing the service controller, or the node is replaced, be sure to follow the replacement instructions for the specific node or controller. You will be instructed to reset the WWNN of the node to that of the original node. If you do not do that, you may need to reconfigure your SAN fabric, your hosts, and your storage systems.

**Attention:** If more than one I/O group is affected, ensure that you are adding the node to the same I/O group from which it was removed. Failure to do this can result in data corruption. Use the information that was recorded when the node was originally added to the cluster. This can avoid a possible data corruption exposure if the node must be removed from and re-added to the cluster. If you do not have access to this information, call the IBM Support Center to add the node back into the cluster without corrupting the data. If you are adding the node into the cluster for the first time, you must record the following information:

- Node serial number
- WWNN
- All WWPNs
- I/O group that contains the node

5. Issue the following command to add the node back into the cluster:

```
svctask addnode -wwnodename WWNN -iogrp  
IOGRPNAME/ID [-name NODENAME]
```

Where *WWNN* is the worldwide node name, *IOGRPNAME/ID* is the I/O group name or ID and *NODENAME* is the name of the node.

In a service situation, a node should normally be added back into a cluster using the original node name. As long as the partner node in the I/O group has not been deleted too, this is the default name used if **-name** is not specified.

6. Issue the following command to verify that the node is online:

```
svcinfolsnode
```

## Recovering VDisks

Virtual disks (VDisks) or VDisk copies are corrupted if they have lost cached data or space-efficient metadata, usually as a result of hardware failure. A Fast Write State of Corrupt indicates this data loss.

You can recover one or more corrupted VDisks and VDisk copies. This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to recover VDisks and VDisk copies:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
2. Select the offline VDisks and select **Recover VDisk** from the task list and click **Go**. The Recovering VDisks panel is displayed.
3. Verify that the VDisks and VDisk copies have completed recovery by monitoring the VDisk Recovery Results panel.

## Recovering offline VDisks using the CLI

You can recover offline virtual disks (VDisks) using the command-line interface (CLI).

Perform the following steps to recover offline VDisks:

1. Issue the following CLI command to list all VDisks that are offline and belong to an I/O group, enter:

```
svcinfolsvdisk -filtervalue IO_group_name=  
IOGRPNAME/ID:status=offline
```

where *IOGRPNAME/ID* is the name of the I/O group that failed.



2. To acknowledge data loss for a VDisk with a *fast\_write\_state* of **corrupt** and bring the VDisk back online, enter:

```
svctask recovervdisk vdisk_id | vdisk_name
```

where *vdisk\_id* / *vdisk\_name* is the name or ID of the VDisk.

**Notes:**

- If the specified VDisk is space-efficient or has space-efficient copies, the **recovervdisk** command starts the space-efficient repair process.
  - If the specified VDisk is mirrored, the **recovervdisk** command starts the resynchronization process.
3. To acknowledge data loss for all virtual disks in an I/O group with a *fast\_write\_state* of **corrupt** and bring them back online, enter:

```
svctask recovervdiskbyiogrp io_group_id | io_group_name
```

where *io\_group\_id* / *io\_group\_name* is the name or ID of the I/O group.

**Notes:**

- If any VDisk is space-efficient or has space-efficient copies, the **recovervdiskbyiogrp** command starts the space-efficient repair process.
  - If any VDisk is mirrored, the **recovervdiskbyiogrp** command starts the resynchronization process.
4. To acknowledge data loss for all VDIs in the cluster with a *fast\_write\_state* of **corrupt** and bring them back online, enter:

```
svctask recovervdiskbycluster
```

**Notes:**

- If any VDisk is space-efficient or has space-efficient copies, the **recovervdiskbycluster** command starts the space-efficient repair process.
- If any VDisk is mirrored, the **recovervdiskbycluster** command starts the resynchronization process.

## Moving offline VDIs to their original I/O group using the CLI

You can move offline virtual disks (VDIs) to their original I/O group using the command-line interface (CLI).

Beginning with SAN Volume Controller 4.3.1, the recovery I/O group is no longer used for VDisk recovery, but it is possible that VDIs were moved to the I/O group before the upgrade.

After a node or an I/O group fails, you can use the following procedure to move offline VDIs to their original I/O group.

**Attention:** Do not move VDIs to an offline I/O group. Ensure that the I/O group is online before you move the VDIs back to avoid any further data loss.

Perform the following steps to move offline VDIs to their original I/O group:

1. Issue the following command to move the VDisk back into the original I/O group:

```
svctask chvdisk -iogrp IOGRPNAME/ID -force  
vdiskname/ID
```

where *IOGRPNAME/ID* is the name or ID of the original I/O group and *vdiskname/ID* is the name or ID of the offline VDisk.

2. Issue the following command to verify that the VDIs are now online:

```
svcinfolsvdisk -filtervalue IO_group_name=  
IOGRPNAME/ID
```

where *IOGRPNAME/ID* is the name or ID of the original I/O group.

---

## Creating MDisk groups

You can create a new managed disk (MDisk) group using the Create a Managed Disk Group wizard.

If you intend to keep the virtual disk (VDisk) allocation within one disk controller system, ensure that the MDisk group that corresponds with a single disk controller system is presented by that disk controller system. This also enables nondisruptive migration of data from one disk controller system to another disk controller system and simplifies the decommissioning process if you want to decommission a disk controller system at a later time.

Ensure all MDisks that are allocated to a single MDisk group are of the same RAID-type. Using the same RAID-type ensures that a single failure of a physical disk in the disk controller system does not take the entire group offline. For example, if you have three RAID-5 arrays in one group and add a non-RAID disk to this group, you lose access to all the data that is striped across the group if the non-RAID disk fails. Similarly, for performance reasons, you should not mix RAID types.

If you are using a SAN Volume Controller solid-state drive (SSD) managed disk, ensure that you are familiar with the SSD configuration rules.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to create a new MDisk group:

1. Click **Work with Managed Disks** → **Managed Disk Groups** in the portfolio. The Viewing Managed Disk Groups panel is displayed.
2. Select **Create an MDisk Group** from the task list and click **Go**. The Create a Managed Disk Group wizard begins.
3. Complete the Create a Managed Disk Group wizard.

## Adding MDisks to MDisk groups

You can add managed disks (MDisks) to an MDisk group from the Adding Managed Disks to Managed Disk Group panel.

If you are using a SAN Volume Controller solid-state drive (SSD) managed disk, ensure that you are familiar with the SSD configuration rules.

This task assumes that you have already launched the SAN Volume Controller Console.

Solid-state drives (SSDs) that are located in SAN Volume Controller 2145-CF8 nodes are presented to the cluster as MDisks. To determine whether the selected MDisk is an SSD, click the link on the MDisk name to display the Viewing MDisk Details panel. If the selected MDisk is an SSD that is located on a SAN Volume

Controller 2145-CF8 node, the Viewing MDisk Details panel displays values for the Node ID, Node Name, and Node Location attributes. Alternatively, you can select **Work with Managed Disks** → **Disk Controller Systems** from the portfolio. On the Viewing Disk Controller panel, you can match the MDisk to the disk controller system that has the following values for these attributes.

Table 18. Disk controller attributes for SSDs

Attribute	SSD value
Vendor ID	IBM
Product ID Low	2145
Product ID High	Internal

**Note:** The first time that you add a new solid-state drive (SSD) to an MDisk group, the SSD is automatically formatted and set to a block size of 512 bytes.

Perform the following steps to add MDisks to an MDisk group:

1. Click **Work with Managed Disks** → **Managed Disk Groups** in the portfolio. The Viewing Managed Disk Groups panel is displayed.
2. Select the MDisk group to add MDisks to and select **Add MDisks** from the list. Click **Go**. The Adding Managed Disks to Managed Disk Group panel is displayed.
3. Select the MDisks to add and click **OK**.

## Viewing the MDisk status

You must view the status of managed disks (MDisks) as part of the repair verification procedures.

This task assumes that you have already launched the SAN Volume Controller Console.

When all SAN Volume Controller repairs are complete, all MDisks are shown as online. Any MDisks that have the statuses of offline, degraded path, degraded port, or excluded might contain errors or are not recognized because of a problem with the SAN environment. If problems still exist on MDisks after the repair actions on the SAN Volume Controller are complete, resolve the disk drive or SAN problems and then perform repair verification for the SAN Volume Controller to verify that no other problems exist.

Complete the following steps to view the status of MDisks:

1. Click **Work with Managed Disks** → **Managed Disks** from the portfolio. The Viewing Managed Disks panel is displayed.
2. Ensure that all MDisks are online. In addition to the online status, the MDisk can have the following status values:

### Degraded Paths

Indicates that the MDisk is not accessible to one or more nodes in the cluster. Degraded path state is most likely the result of incorrect configuration of either the disk controller or the fibre-channel fabric. However, hardware failures in the disk controller, fibre-channel fabric, or node could also be a contributing factor to this state. Complete the following actions to recover from this state:

- a. Verify that the fabric configuration rules for storage systems are correct.

- b. Ensure that you have configured the storage system properly.
- c. Correct any errors in the error log.

#### Degraded Ports

Indicates that one or more 1220 errors have been logged against the selected MDisk. The 1220 error indicates that the remote fibre-channel port has been excluded from the MDisk. This error may cause reduced performance on the storage controller and usually indicates a hardware problem with the storage controller. To fix this problem you must resolve any hardware problems on the storage controller and fix the 1220 errors in the error log. To resolve these errors in the log, select **Service and Maintenance** → **Run Maintenance Procedures** in the SAN Volume Controller Console. On the Maintenance Procedures panel, select **Start Analysis**. This action displays a list of unfixed errors that are currently in the error log. For these unfixed errors, select the error name to begin a guided maintenance procedure to resolve these errors. Errors are listed in descending order with the highest priority error listed first. Highest priority errors should be resolved first.

#### Excluded

Indicates that the MDisk has been excluded from use by the cluster after repeated access errors. Run the directed maintenance procedures to determine the problem. You can reset an MDisk and include it in the cluster by running the include task.

#### Offline

Indicates that the MDisk cannot be accessed by any of the online nodes. That is, all of the nodes that are currently working members of the cluster cannot access this MDisk. This state can be caused by a failure in the SAN, storage system, or one or more physical disks connected to the storage system. The MDisk is reported as offline if all paths to the disk fail.

## Discovering MDisks

You can have the cluster rescan the fibre-channel network. The rescan discovers any new managed disks (MDisks) that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to discover MDisks:

1. Click **Work with Managed Disks** → **Managed Disks** in the portfolio. The Viewing Managed Disks panel is displayed.
2. Select **Discover MDisks** from the task list and click **Go**. The Discovering Managed Disks panel is displayed. The newly discovered MDisks are displayed in a table on the Discovering Managed Disks panel.
3. Click **Close** to return to the Viewing Managed Disks panel.

#### Related tasks

“Discovering MDisks using the CLI” on page 98

You can use the command-line interface (CLI) to discover managed disks (MDisks).

#### Viewing discovery status

You can view the status of a managed disk (MDisk) discovery from the Viewing Discovery Status panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view status of an MDisk discovery:

1. Click **Work with Managed Disks** → **Discovery Status**. The Viewing Discovery Status panel is displayed. The following status values are possible:

**Active** This status indicates an MDisk discovery is currently in progress.

**Inactive**

This status indicates an MDisk discovery is not currently in progress.

2. Click **Close** to close this panel.

## Creating MDisk groups using the CLI

You can use the command-line interface (CLI) to create a managed disk (MDisk) group.

**Attention:** If you add an MDisk to an MDisk group as an MDisk, any data on the MDisk is lost. If you want to keep the data on an MDisk (for example, because you want to import storage that was previously not managed by SAN Volume Controller), you must create image mode virtual disks (VDisks) instead.

Assume that the cluster has been set up and that a back-end controller has been configured to present new storage to the SAN Volume Controller.

If you are using a SAN Volume Controller solid-state drive (SSD) managed disk, ensure that you are familiar with the SSD configuration rules.

This task assumes that you have already launched the SAN Volume Controller Console.

Consider the following factors as you decide how many MDisk groups to create:

- A VDisk can only be created using the storage from one MDisk group. Therefore, if you create small MDisk groups, you might lose the benefits that are provided by virtualization, namely more efficient management of free space and a more evenly distributed workload for better performance.
- If any MDisk in an MDisk group goes offline, all the VDIs in the MDisk group go offline. Therefore you might want to consider using different MDisk groups for different back-end controllers or for different applications.
- If you anticipate regularly adding and removing back-end controllers or storage, this task is made simpler by grouping all the MDisks that are presented by a back-end controller into one MDisk group.
- All the MDisks in an MDisk group should have similar levels of performance or reliability, or both. If an MDisk group contains MDisks with different levels of performance, the performance of the VDIs in this group is limited by the performance of the slowest MDisk. If an MDisk group contains MDisks with different levels of reliability, the reliability of the VDIs in this group is that of the least reliable MDisk in the group.

**Note:** When you create an MDisk group with a new solid-state drive (SSD), the new SSD is automatically formatted and set to a block size of 512 bytes.

Even with the best planning, circumstances can change and you must reconfigure your MDisk groups after they have been created. The data migration facilities that are provided by the SAN Volume Controller enable you to move data without disrupting I/O.

### Choosing an MDisk group extent size

Consider the following factors as you decide the extent size of each new MDisk group:

- You must specify the extent size when you create a new MDisk group.
- You cannot change the extent size later; it must remain constant throughout the lifetime of the MDisk group.
- MDisk groups can have different extent sizes; however, this places restrictions on the use of data migration.
- The choice of extent size affects the maximum size of a VDisk in the MDisk group.

Table 19 compares the maximum VDisk capacity for each extent size. The maximum is different for space-efficient VDIs. Because the SAN Volume Controller allocates a whole number of extents to each VDisk that is created, using a larger extent size might increase the amount of storage that is wasted at the end of each VDisk. Larger extent sizes also reduces the ability of the SAN Volume Controller to distribute sequential I/O workloads across many MDisks and therefore can reduce the performance benefits of virtualization.

Table 19. Maximum VDisk capacity by extent size

Extent size (MB)	Maximum VDisk capacity in GB (not space-efficient VDIs)	Maximum VDisk capacity in GB (space-efficient VDIs)
16	2048 (2 TB)	2000
32	4096 (4 TB)	4000
64	8192 (8 TB)	8000
128	16,384 (16 TB)	16,000
256	32,768 (32 TB)	32,000
512	65,536 (64 TB)	65,000
1024	131,072 (128 TB)	130,000
2048	262,144 (256 TB)	260,000

**Important:** You can specify different extent sizes for different MDisk groups; however, you cannot migrate VDIs between MDisk groups with different extent sizes. If possible, create all your MDisk groups with the same extent size.

Perform the following steps to create an MDisk group:

Issue the `svctask mkmdiskgrp` CLI command to create an MDisk group. The following is an example of the CLI command you can issue to create an MDisk group:

```
svctask mkmdiskgrp -name maindiskgroup -ext 32
-mdisk mdsk0:mdsk1:mdsk2:mdsk3
```

where *maindiskgroup* is the name of the MDisk group that you want to create, 32 MB is the size of the extent you want to use, and *mdsk0*, *mdsk1*, *mdsk2*, *mdsk3* are the names of the four MDisks that you want to add to the group.

You created and added MDisks to an MDisk group.

The following example provides a scenario where you want to create an MDisk group, but you do not have any MDisks available to add to the group. You plan to add the MDisks at a later time. You use the `svctask mkmdiskgrp` CLI command to create the MDisk group *bkpmdiskgroup* and later used the `svctask addmdisk` CLI command to add *mdsk4*, *mdsk5*, *mdsk6*, *mdsk7* to the MDisk group.

1. Issue `svctask mkmdiskgrp -name bkpmdiskgroup -ext 32`  
where *bkpmdiskgroup* is the name of the MDisk group that you want to create and 32 MB is the size of the extent that you want to use.
2. You find four MDisks that you want to add to the MDisk group.
3. Issue `svctask addmdisk -mdisk msk4:msk5:msk6:msk7 bkpmdiskgroup`  
where *mdsk4*, *mdsk5*, *mdsk6*, *mdsk7* are the names of the MDisks that you want to add to the MDisk group and *bkpmdiskgroup* is the name of the MDisk group for which you want to add MDisks.

### Adding MDisks to MDisk groups using the CLI

You can use the command-line interface (CLI) to add managed disks (MDisks) to MDisk groups.

The MDisks must be in unmanaged mode. Disks that already belong to an MDisk group cannot be added to another MDisk group until they have been deleted from their current MDisk group. You can delete an MDisk from an MDisk group under the following circumstances:

- If the MDisk does not contain any extents in use by a virtual disk (VDisk)
- If you can first migrate the extents in use onto other free extents within the group

**Important:** Do not add an MDisk using this procedure if you are mapping the MDisk to an image mode VDisk. Adding an MDisk to an MDisk group allows the SAN Volume Controller to write new data to the MDisk; therefore, any existing data on the MDisk is lost. If you want to create an image mode VDisk, use the `svctask mkvdisk` command instead of `svctask addmdisk`.

If you are using a SAN Volume Controller solid-state drive (SSD) managed disk, ensure that you are familiar with the SSD configuration rules.

When you are adding MDisks to an MDisk group using the `svctask addmdisk` command or when you are creating an MDisk group using the `svctask mkmdiskgrp -mdisk` command, the SAN Volume Controller performs tests on the MDisks in the list before the MDisks are allowed to become part of an MDisk group. These tests include checks of the MDisk identity, capacity, status and the ability to perform both read and write operations. If these tests fail or exceed the time allowed, the MDisks are not added to the group. However, with the `svctask mkmdiskgrp -mdisk` command, the MDisk group is still created even if the tests fail, but it does not contain any MDisks. If tests fail, confirm that the MDisks are in the correct state and that they have been correctly discovered.

The following events contribute to an MDisk test failure:

- The MDisk is not visible to all SAN Volume Controller nodes in the cluster.

- The MDisk identity has changed from a previous discovery operation.
- The MDisk cannot perform read or write operations.
- The status of the MDisk can be either degraded paths, degraded ports, excluded, or offline.
- The MDisk does not exist.

The following events contribute to an MDisk test timeout:

- The disk controller system on which the MDisk resides is failing.
- A SAN fabric or cable fault condition exists that is preventing reliable communication with the MDisk.

**Note:** The first time that you add a new solid-state drive (SSD) to an MDisk group, the SSD is automatically formatted and set to a block size of 512 bytes.

Perform the following steps to add MDisks to MDisk groups:

1. Issue the **svcinfolsmdiskgrp** CLI command to list the existing MDisk groups.

The following is an example of the CLI command you can issue to list the existing MDisk groups:

```
svcinfolsmdiskgrp -delim :
```

The following is an example of the output that is displayed:

```
id:name:status:mdisk_count:vdisk_count:
capacity:extent_size:free_capacity:virtual_capacity:
used_capacity:real_capacity:overallocation:warning
0:mdiskgrp0:online:3:4:33.3GB:16:32.8GB:64.00MB:64.00MB:64.00MB:0:0
1:mdiskgrp1:online:2:1:26.5GB:16:26.2GB:16.00MB:16.00MB:16.00MB:0:0
2:mdiskgrp2:online:2:0:33.4GB:16:33.4GB:0.00MB:0.00MB:0.00MB:0:0
```

2. Issue the **svctask addmdisk** CLI command to add MDisks to the MDisk group. The following is an example of the CLI command you can issue to add MDisks to an MDisk group:

```
svctask addmdisk -mdisk mdisk4:mdisk5:mdisk6:mdisk7 bkpmdiskgroup
```

Where *mdisk4:mdisk5:mdisk6:mdisk7* are the names of the MDisks that you want to add to the MDisk group and *bkpmdiskgroup* is the name of the MDisk group for which you want to add the MDisks.

## Listing MDisks using the CLI

You can use the command-line interface (CLI) to list the managed disks (MDisks).

Perform the following steps to list the name, status, and mode of an MDisk:

1. Issue the following CLI command:

```
svcinfolsmdisk -delim :
```

The following output is an example of what you might see when you issue the `svcinfolsmdisk -delim :` command:





## Locating a solid-state drive (SSD) using the CLI

You can identify the node slot that a solid-state drive (SSD) is located in using the SAN Volume Controller CLI.

To identify the node slot that an SSD is located in, run the `lsmdisk` command. To identify additional information about an SSD, run the `lsnodevpd` command.

1. To locate an SSD, run the following command :

```
svcinfo lsmdisk mdisk_name | mdisk_id
```

Where *mdisk\_name* / *mdisk\_id* is the name or ID of the MDisk that is an SSD.

Output is similar to the following example. The `node_id` and `node_name` fields identify the node that the SSD is installed in. The `location` field indicates the drive slot that the SSD is located in. If the field is blank, the MDisk is not an SSD in a SAN Volume Controller node.

```
id 0
name mdisk0
status online
mode managed
mdisk_grp_id 1
mdisk_grp_name ssd_n1
capacity 136.7GB
quorum_index
block_size 512
controller_name controller0
ctrl_type 6
ctrl_WWNN 5005076801E00047
controller_id 0
path_count 1
max_path_count 1
ctrl_LUN_# 0000000000000000
UID 5000a720000083910000000000000000\
00000000000000000000000000000000
preferred_WWPN 5000A72A00008391
active_WWPN 5000A72A00008391
node_id 1
node_name node1
location 2
```

2. Run the following command to list additional details about the SSD, including the serial number and firmware level.

```
svcinfo lsnodevpd node_name | node_id
```

Where *node\_name* / *node\_id* is the name or ID of the node that contains the SSD.

Output is similar to the following example. The `drive_location` identifies the specific SSD that is listed:

```
drive_location 2
manufacturer IBM
model Z16IZD2B-73
capacity 00014337400
serial_number S092901FG008
label_serial_number S092901F
supplier_serial_number S092901F
part_number 41Y8476
firmware_level G008
FPGA_revision F5
type SSD
speed
enclosure
connection_type SAS
```

## Discovering MDisks using the CLI

You can use the command-line interface (CLI) to discover managed disks (MDisks).

When back-end controllers are added to the fibre-channel SAN and are included in the same switch zone as a SAN Volume Controller cluster, the cluster automatically discovers the back-end controller and integrates the controller to determine the storage that is presented to the SAN Volume Controller nodes. The SCSI logical units (LUs) that are presented by the back-end controller are displayed as unmanaged MDisks. However, if the configuration of the back-end controller is modified after this has occurred, the SAN Volume Controller cluster might be unaware of these configuration changes. You can request that the SAN Volume Controller cluster rescans the fibre-channel SAN to update the list of unmanaged MDisks.

**Note:** The automatic discovery that is performed by SAN Volume Controller cluster does not write anything to an unmanaged MDisk. You must instruct the SAN Volume Controller cluster to add an MDisk to an MDisk group or use an MDisk to create an image mode virtual disk (VDisk).

Perform the following steps to discover and then view a list of MDisks:

1. Issue the **svctask detectmdisk** CLI command to manually scan the fibre-channel network. The scan discovers any new MDisks that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

**Notes:**

- a. Only issue the **svctask detectmdisk** command when you are sure that all of the disk controller ports are working and correctly configured in the controller and the SAN zoning. Failure to do this can result in errors that are not reported.
  - b. Although it might appear that the **detectmdisk** command has completed, extra time might be required for it to run. The **detectmdisk** is asynchronous and returns a prompt while the command continues to run in the background. You can use the **lsdiscoverystatus** command to view the discovery status.
2. When the detection is complete, issue the **svcinfolsmdiskcandidate** CLI command to show the unmanaged MDisks. These MDisks have not been assigned to an MDisk group.
  3. Issue the **svcinfolsmdisk** CLI command to view all of the MDisks.

You have now seen that the back-end controllers and switches have been set up correctly and that the SAN Volume Controller cluster recognizes the storage that is presented by the back-end controller.

The following example describes a scenario where a single back-end controller is presenting eight SCSI LUs to the SAN Volume Controller cluster:

1. Issue `svctask detectmdisk`.
2. Issue `svcinfolsmdiskcandidate`.

The following output is displayed:



## Checking MDisk group status using the CLI

You can use the command-line interface (CLI) to check the status of a managed disk (MDisk) group.

Issue the following CLI command to display the status of MDisk groups:

```
svcinfolsmdiskgrp -delim :
```

The following output is an example of what you might see:

```
id:name:status:mdisk_count:vdisk_count:capacity:extent_size:free_capacity:
virtual_capacity:used_capacity:real_capacity:overallocation:warning
0:mdiskgrp0:offline:4:0:34.2GB:16:34.2GB:0:0:0:0:0
1:mdiskgrp1:online:4:6:200GB:16:100GB:400GB:75GB:100GB:200:80
```

The characteristics for each MDisk group is listed one line per group. The status of the MDisk group is shown following the name of the MDisk group. In this example, the mdiskgrp0 status is offline. MDisk group status can be offline, online, degraded paths, or degraded ports.

---

## Viewing the fibre-channel fabric connections

Failures of the SAN Volume Controller hardware, fibre-channel cables, fibre-channel switches, fibre-channel hardware in host systems, or disk controllers can cause fibre-channel connectivity problems.

**Note:** The term *fibre-channel fabric* describes all the interconnections between the ports on a fibre-channel network.

Other SAN Volume Controller procedures isolate the hardware errors that cause storage area network (SAN) connectivity problems. You were sent to this topic because it is suspected that a problem exists elsewhere on the SAN. The resolution of problems on the fibre-channel fabric is outside of the scope of current documentation. The procedures here describe how to use the fibre-channel fabric viewing tool to help you understand the SAN Volume Controller view of the SAN. The information you acquire here is valuable when you work with support personnel or with the customer to isolate SAN failures.

Perform the following steps to start the fabric-viewing tool:

1. Log on to the SAN Volume Controller Console.
2. Click **Clusters**. The **Viewing Clusters** panel is displayed.
3. Select the cluster for which you want to view fabric information, and select **Launch the SAN Volume Controller Console** from the task list.
4. Click **Go**.
5. Click **Service and Maintenance** → **Fabrics**. The **Viewing Fabrics** panel is displayed. See the example in Figure 52 on page 102.

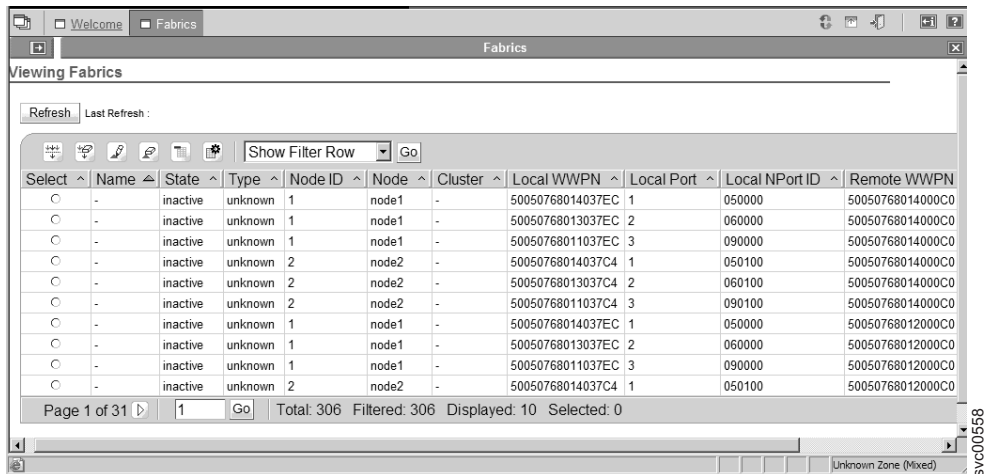


Figure 52. Viewing Fabrics panel

You see the Viewing Fabrics panel. The contents of each displayed field is described in Table 20.

Table 20. Fibre-channel viewing definitions

Fields	Description
Name	This field is the name of the device whose worldwide port name (WWPN) is listed in the <b>Remote WWPN</b> field.
State	Indicates whether the device listed in the <b>Name</b> field is active or inactive.
Type	The type of the device that is listed in the <b>Name</b> field. The expected types are controller (storage controller), node (SAN Volume Controller), or host (system that is using this SAN Volume Controller). If the type of device cannot be determined, unknown is displayed.
Node ID	The ID of the node that is listed in the <b>Node</b> field.
Node	This field is the node name (as displayed on the front panel of the SAN Volume Controller).
Cluster	When the <b>Type</b> field lists a node, the <b>Cluster</b> field displays the name of the cluster to which that node belongs.
Local WWPN	The WWPN of the SAN Volume Controller port that has a connection to the device listed in the <b>Name</b> field, using the WWPN that is listed in the <b>Remote WWPN</b> field.
Local Port	This field is the physical port number on the back of the node listed in the <b>Node</b> field.
Local NPort ID	The NPort number of the local port listed in the <b>Local Port</b> field.
Remote WWPN	The WWPN of the device listed in the <b>Name</b> field.
Remote NPort ID	The NPort number of the device listed in the <b>Name</b> field.

The first line of Figure 52 shows the following connection:




Figure 53. Example of the fabric connections

The fabric viewer provides comprehensive tools to let you display only the information that you need, and format the output in the most convenient form.

**Note:** The following examples show methods of displaying fabric connections. Not all scenarios are represented, and the examples only show two potential methods of finding the information that you seek.

## Displaying all controllers logged in to a node

Perform the following steps to display all the controllers that are logged in to a node called "node1":

1. From the task drop-down menu, select **Show Filter Row** and click **Go**. Click the **Filter** link. In the **Text** field, type `controller`, and then click **OK**.
2. Under the **Name** heading, click the **Filter** link. In the **Text** field, type `controller`, and then click **OK**. This action sorts by controller.
3. Under the **Node** heading, click the **Filter** link. In the **Text** field, type `node1` and then click **OK**. You are shown a display that shows all controllers that are logged in to node1.
4. To restore the original display, select **Clear All Filters** from the task drop-down menu and click **Go**.

## Displaying all active devices

Perform the following steps to display all active devices that are logged in to port 2 of node2:

1. Under the **State** heading, click the **Filter** link. In the **Text** input field, type `active`.
2. From the **Condition** drop-down menu, select **Matches**, and then click **OK**.
3. Under the **Node** heading, click the **Filter** link. In the **Text** input field, type `node2` and then click **OK**.
4. Under the **Local Port** heading, click the **Filter** link. In the **Text** input field, type `2`, and then click **OK**.

You see a display that shows all the devices that are logged in to port 2 on node2.

---

## Changing the fibre-channel network speed

Depending on the model, you can change the speed of the fibre-channel ports on a SAN Volume Controller through the front panel on the SAN Volume Controller node or by a command sent to a SAN Volume Controller cluster using the Ethernet interface.

If you use SAN Volume Controller model 2145-8F4, 2145-8G4, 2145-8A4, or 2145-CF8 the speed of the fibre-channel ports is governed by the maximum speed of the fibre-channel switch to which the ports are connected, up to a maximum of 4 Gbps. The ports operate at the highest speed at which a reliable connection can be established. Different ports and nodes on the SAN Volume Controller can operate at different speeds.

**Note:** You can use the command-line interface, console, or front panel to set the cluster fibre-channel speed for non-negotiating nodes; however, this action affects only non-negotiating node types such as the SAN Volume Controller 2145-8F2. For all other models, the autonegotiated speed applies, and the speed set for the cluster has no effect on the node.

All SAN Volume Controller 2145-8F2 fibre-channel ports and nodes on a SAN Volume Controller must operate at the same speed. The default for the port speed is 2 Gbps. If the fibre-channel fabric is using fibre-channel switches that are unable to operate at 2 Gbps, set the SAN Volume Controller fibre-channel port speed to 1 Gbps during the installation procedure. If you must replace a 2 Gbps fibre-channel switch with a 1 Gbps fibre-channel switch, you must manually switch the SAN Volume Controller fibre-channel port speed before you can use the SAN Volume Controller at 1 Gbps.

If a new SAN Volume Controller 2145-8F2 node is added to an existing SAN Volume Controller configuration, you must change the speed setting if the switch is only capable of running at 1 Gbps. If the SAN Volume Controller internal disk drive fails, the fibre-channel speed setting is lost, but the Node Rescue Procedure that you use to restore the SAN Volume Controller software automatically selects the correct speed. If a 1 Gbps fibre-channel switch is being replaced by a 2 Gbps fibre-channel switch, the existing SAN Volume Controller clusters operate at 1 Gbps. You can switch to 2 Gbps any time using the command-line interface.

## Determining the fibre-channel port speed

You must determine the fibre-channel port speed for several actions. You can determine the fibre-channel port speed using the SAN Volume Controller Console, the CLI, or from the front panel.

### Viewing the fibre-channel port speed from the SAN Volume Controller Console

You can view the fabric speed for the cluster from the View Cluster Properties panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the fibre-channel port speed:

1. Click **Manage Cluster** → **View Cluster Properties** in the portfolio. The Viewing General Properties panel is displayed.
2. Click the **General** tab to display the general properties, including the port speed.



3. Click **Close** to close the panel.

### Determining the fibre-channel port speed from the CLI

If the node is in a cluster that is operational, you can determine the fibre-channel port speed from the cluster vital product data.

To view the fibre-channel port speed of a cluster, issue the following command:

```
svcinfolcluster -delim : cluster1
```

The following output is an example of what you might see. The fibre-channel port speed is shown in bold.

```
id:cluster1
name:rc-cluster-20
location:local partnership:
bandwidth:
cluster_IP_address:9.71.50.32
cluster_service_IP_address:9.71.50.183
total_mdisk_capacity:2976.9GB
space_in_mdisk_grps::2976.9GB
space_allocated_to_vdisks:147.2GB
total_free_space:2828.7GB
statistics_status:on
statistics_frequency:15
required_memory:8192
cluster_locale:en_US
SNMP_setting:none
SNMP_community:
SNMP_server_IP_address:[0.0.0.0]:23
subnet_mask:255.255.254.0
default_gateway:9.71.50.1
time_zone:522
UTC_email_setting:none
email_id:
code_level:4.1.0.12 (build 5.13.0610240000)
FC_port_speed:2Gb
console_IP:9.71.49.176:9080
id_alias:cluster1
gm_link_tolerance:300
gm_inter_cluster_delay_simulation:0
gm_intra_cluster_delay_simulation:0
email_server:8.53.26.131
email_server_port:25
email_reply:manager@mycompany.com
email_contact:manager
email_contact_primary:01202 123456
email_contact_alternate:44-212-786543-4455
email_contact_location:city
email_state:running
email_user_count:2
inventory_mail_interval:0
cluster_IP_address_6:
cluster_service_IP_address_6:
prefix_6:
default_gateway_6:
total_vdiskcopy_capacity:40.00GB
total_used_capacity:22.50GB
total_overallocation:67
total_vdisk_capacity:30.00GB
```

### Determining the fibre-channel port speed from the front panel

Perform the following steps from the front panel to determine the fibre-channel port speed:

1. Select any fibre-channel port from the front panel.

2. Press and hold **Down**.
3. Press and release **Select**.
4. Release **Down**.
5. Press **Select** again to cancel the text display or wait 60 seconds and the text display cancels itself.

## Changing the fibre-channel port speed for a node not in a cluster

You can change the fibre-channel port speed for a 2145-8F2 node that is not in a cluster.

Perform the following steps to change the fibre-channel port speed for a node that is not in a cluster:

1. From the front panel, select any fibre channel port.
2. Press and hold **Down**.
3. Press and release **Select**.
4. Release **Down**.
5. Press **Up** or **Down** until the required speed is displayed.
6. Press **Select** to activate the new speed.

## Changing the fibre-channel port speed for a node in a cluster

The fibre channel port speed of all SAN Volume Controller 2145-8F2 nodes in the cluster can be changed by issuing the `svctask chcluster` command.

This is a customer task that is only available to users who are logged on using the Administrator or Security Administrator role.

**Note:** Changing the fibre-channel port speed causes all nodes in the cluster to simultaneously perform a warm start. This causes any I/O activity through the cluster to fail and consequently might cause applications running on hosts to fail. The SAN Volume Controller Console processing also stops until the cluster connection is re-established.

---

## Checking disk controller status using the CLI

You can use the command-line interface (CLI) to check the status of the disk controllers.

Issue the `svcinfolcontroller -delim :` command to display the example output:

```
id:controller_name:ctrl_s/n:vendor_id:product_id_low:product_id_high
0:controller0::IBM :1742-900:
1:controller1::IBM :2105800 :
```

You can also check the status of a specific disk controller. Issue the following command to check the status of a specific disk controller:

```
svcinfolcontroller -delim = controller_id
```

where *controller\_id* is the ID of the controller whose status you want to check.

```

id=0
controller_name=controller0
WWNN=200600A0B80F80A6
mdisk_link_count=7
max_mdisk_link_count=7
degraded=no
vendor_id=IBM
product_id_low=1742-900
product_id_high=
product_revision=0914
ctrl_s/n=
allow_quorum=yes
WWPN=200600A0B80F80A7
path_count=4
max_path_count=4
WWPN=200600A0B80F80A8
path_count=4
max_path_count=4
WWPN=200700A0B80F80A7
path_count=4
max_path_count=4
WWPN=200700A0B80F80A8
path_count=2
max_path_count=2

```

You can see the status of a specific disk controller (id=0) in the sixth line of the previous example. The value of degraded is defined below:

- no** Specifies that the status is good
- yes** Specifies that the status is undesirable

#### Related reference

“Determining the failing enclosure or disk controller using the CLI”

You can use the command-line interface (CLI) to determine the failing enclosure or disk controller.

## Determining the failing enclosure or disk controller using the CLI

You can use the command-line interface (CLI) to determine the failing enclosure or disk controller.

Issue the following command to list all the managed disks (MDisks):

```
svcinfolsmdisk -nohdr -delim :
```

The following output is an example of what you might see after you issue the `svcinfolsmdisk -nohdr -delim :` command:

```

0:mdisk0:online:managed:0:mdiskgrp0:273.3GB:0000000000000001:controller0:*
1:mdisk1:excluded:managed:0:mdiskgrp0:546.6GB:0000000000000002:controller0:*

* 600a0b80000c5ae4000000093eca105c00000000000000000000000000000000 is appended to this line.
* 600a0b80000f6432000000043ef6b4ff00000000000000000000000000000000 is appended to this line.

```

The MDisks are listed in the object ID order. The first item is the object ID, the third item is the status, and the ninth item is the disk or controller name. In the previous example, mdisk1 has an object ID of 1, is failing with the status excluded, and is part of a disk controller named controller0.

---

## Replacing a faulty node with a spare node

You can use the SAN Volume Controller Console and the SAN Volume Controller front panel to replace a faulty node in a cluster.

Before you attempt to replace a faulty node with a spare node you must ensure that you meet the following requirements:

- SAN Volume Controller version 3.1.0 or higher is installed on the cluster and on the spare node.
- You know the name of the cluster that contains the faulty node.
- A spare node is installed in the same rack as the cluster that contains the faulty node.
- Record of the last five characters of the original worldwide node name (WWNN) of the spare node.

**Note:** A repaired faulty node, which has been successfully replaced in the cluster with a spare node using the original WWPN of the faulty node, must be assigned a new unique WWNN. You can use the original WWNN of the spare node as the new WWNN of the repaired node.

**Attention:** Never connect a node with a WWNN of 00000 to the cluster. If this node is no longer required as a spare and is to be used for normal attachment to a cluster, you must change the WWNN to the number you recorded when a spare was created. Using any other number might cause data corruption.

If a node fails, the cluster continues to operate with degraded performance, until the faulty node is repaired. If the repair operation takes an unacceptable amount of time, it is useful to replace the faulty node with a spare node. However, the appropriate procedures must be followed and precautions must be taken so you do *not* interrupt I/O operations and compromise the integrity of your data.

The following table describes the changes that are made to your configuration when you replace a faulty node in the cluster:

Node attributes	Description
Front panel ID	This is the number that is printed on the front of the node and is used to select the node that is added to a cluster.
Node ID	This is the ID that is assigned to the node. A new node ID is assigned each time a node is added to a cluster; the node name remains the same following service activity on the cluster. You can use the node ID or the node name to perform management tasks on the cluster. However, if you are using scripts to perform those tasks, use the node name rather than the node ID. This ID will change during this procedure.

Node attributes	Description												
Node name	This is the name that is assigned to the node. If you do not specify a name, the SAN Volume Controller assigns a default name. The SAN Volume Controller creates a new default name each time a node is added to a cluster. If you choose to assign your own names, you must type the node name on the Adding a node to a cluster panel. You cannot manually assign a name that matches the naming convention used for names assigned automatically by SAN Volume Controller. If you are using scripts to perform management tasks on the cluster and those scripts use the node name, you can avoid the need to make changes to the scripts by assigning the original name of the node to a spare node. This name might change during this procedure.												
Worldwide node name	This is the WWNN that is assigned to the node. The WWNN is used to uniquely identify the node and the fibre-channel ports. During this procedure, the WWNN of the spare node is changed to that of the faulty node. The node replacement procedures must be followed exactly to avoid any duplication of WWNNs. This name does not change during this procedure.												
Worldwide port names	<p>These are the WWPNS that are assigned to the node. WWPNS are derived from the WWNN that is written to the spare node as part of this procedure. For example, if the WWNN for a node is 50050768010000F6, the four WWPNS for this node are derived as follows:</p> <table data-bbox="740 926 1292 1083"> <tbody> <tr> <td>WWNN</td> <td>50050768010000F6</td> </tr> <tr> <td>WWNN displayed on front panel</td> <td>000F6</td> </tr> <tr> <td>WWPN Port 1</td> <td>50050768014000F6</td> </tr> <tr> <td>WWPN Port 2</td> <td>50050768013000F6</td> </tr> <tr> <td>WWPN Port 3</td> <td>50050768011000F6</td> </tr> <tr> <td>WWPN Port 4</td> <td>50050768012000F6</td> </tr> </tbody> </table> <p>These names do not change during this procedure.</p>	WWNN	50050768010000F6	WWNN displayed on front panel	000F6	WWPN Port 1	50050768014000F6	WWPN Port 2	50050768013000F6	WWPN Port 3	50050768011000F6	WWPN Port 4	50050768012000F6
WWNN	50050768010000F6												
WWNN displayed on front panel	000F6												
WWPN Port 1	50050768014000F6												
WWPN Port 2	50050768013000F6												
WWPN Port 3	50050768011000F6												
WWPN Port 4	50050768012000F6												

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to replace a faulty node in the cluster:

1. Verify the name and ID of the node that you want to replace.

Complete the following steps to verify the name and ID:

- a. Make sure that the SAN Volume Controller Console application is running on the cluster that contains the faulty node.
- b. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed. If the node is faulty, it is shown as offline.
- c. Ensure the partner node in the I/O group is online.
  - If the other node in the I/O group is offline, start the Directed Maintenance Procedures (DMPs) to determine the fault.
  - If you have been directed here by the DMPs, and subsequently the partner node in the I/O group has failed, recover the offline VDisks.
  - If you are replacing the node for other reasons, determine the node that you want to replace and ensure that the partner node in the I/O group is online.
  - If the partner node is offline, you will lose access to the VDisks that belong to this I/O group. Start the DMPs and fix the other node before proceeding to the next step.

2. Click the name of the faulty (offline) node. The Viewing General Details panel is displayed.
3. Click **General** and record the following attributes for the faulty node:
  - ID
  - WWNN
  - I/O Group
  - UPS Serial Number
  - Uninterruptible power supply serial number
4. Click **Close**. Click Fibre Channel Port and record the following attribute for the faulty node:
  - WWPNs
5. Click **Close**. Click Vital Product Data and record the following attribute for the faulty node:
  - System Serial Number
6. Ensure that the faulty node has been powered off.
7. Use the SAN Volume Controller Console to delete the faulty node from the cluster.

**Remember:** You must record the following information to avoid data corruption when this node is re-added to the cluster:

- Node serial number
  - WWNN
  - All WWPNs
  - I/O group that contains the node
8. Disconnect all four fibre-channel cables from the node.

**Important:** Do not plug the fibre-channel cables into the spare node until the spare node is configured with the WWNN of the faulty node.

9. Connect the power and signal cables from the spare node to the uninterruptible power supply that has the serial number you recorded in step 3.

**Note:** For 2145 UPS-1U units, you must disconnect the cables from the faulty node.

10. Power on the spare node.
11. You must change the WWNN of the spare node to that of the faulty node. The procedure for doing this depends on the SAN Volume Controller version that is installed on the spare node. Press and release the down button until the Node: panel displays. Then press and release the right button until the WWNN: panel displays. If repeated pressing of the right button returns you to the Node: panel, without displaying a WWNN: panel, go to step 13 on page 111; otherwise, continue with step 12.
12. Change the WWNN of the spare node (with SAN Volume Controller V4.3 and above installed) to match the WWNN of the faulty node by performing the following steps:
  - a. With the Node WWNN: panel displayed, press and hold the down button, press and release the select button, and then release the down button. The display switches into edit mode. Edit WWNN is displayed on line 1. Line 2 of the display contains the last five numbers of the WWNN.
  - b. Change the WWNN that is displayed to match the last five numbers of the WWNN that you recorded in step 3. To edit the highlighted number, use the

up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.

- c. When the five numbers match the last five numbers of the WWNN that you recorded in step 3 on page 110, press the select button to accept the numbers.
13. Change the WWNN of the spare node (with SAN Volume Controller versions prior to V4.3 installed) to match the WWNN of the faulty node by performing the following steps:
  - a. Press and release the right button until the Status: panel is displayed.
  - b. With the node status displayed on the front panel, press and hold the down button; press and release the select button; release the down button. WWNN is displayed on line 1 of the display. Line 2 of the display contains the last five numbers of the WWNN.
  - c. With the WWNN displayed on the front panel; press and hold the down button; press and release the select button; release the down button. The display switches into edit mode.
  - d. Change the WWNN that is displayed to match the last five numbers of the WWNN that you recorded in step 3 on page 110. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.
  - e. When the five numbers match the last five numbers of the WWNN that you recorded in step 3 on page 110, press the select button to accept the numbers.
  - f. Press the select button to retain the numbers that you have updated and return to the WWNN panel.
14. Connect the four fibre-channel cables that you disconnected from the faulty node and connect them to the spare node.

If the spare node has less Ethernet cables connected than the faulty node, move the Ethernet cables from the faulty node to the spare node. Ensure you connect the cable into the same port on the spare node as it was in on the faulty node.
15. Use the SAN Volume Controller Console to add the spare node to the cluster. If possible, use the same node name that was used for the faulty node. If necessary, the spare node is updated to the same SAN Volume Controller version as the cluster. This update can take up to 20 minutes.
16. Use the tools that are provided with your multipathing device driver on the host systems to verify that all paths are now online. See the documentation that is provided with your multipathing device driver for more information. For example, if you are using the subsystem device driver (SDD), see the *IBM System Storage Multipath Subsystem Device Driver User's Guide* for instructions on how to use the SDD management tool on host systems. It might take up to 30 minutes for the paths to come online.
17. Repair the faulty node.

**Attention:** When the faulty node is repaired, do not connect the fibre-channel cables to it. Connecting the cables might cause data corruption because the spare node is using the same WWNN as the faulty node.

If you want to use the repaired node as a spare node, perform the following steps.

**For SAN Volume Controller V4.3 and above:**

- a. With the Node WWNN: panel displayed, press and hold the down button, press and release the select button, and then release the down button. The display switches into edit mode. Edit WWNN is displayed on line 1. Line 2 of the display contains the last five numbers of the WWNN.
- b. Change the displayed number to 00000. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.
- c. Press the select button to accept the numbers.  
This node can now be used as a spare node.

**For SAN Volume Controller versions prior to V4.3:**

- a. Press and release the right button until the Status: panel is displayed.
- b. With the node status displayed on the front panel, press and hold the down button; press and release the select button; release the down button. WWNN is displayed on line 1 of the display. Line 2 of the display contains the last five numbers of the WWNN.
- c. With the WWNN displayed on the front panel; press and hold the down button; press and release the select button; release the down button. The display switches into edit mode.
- d. Change the displayed number to 00000. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.
- e. Press the select button to accept the numbers.
- f. Press the select button to retain the numbers that you have updated and return to the WWNN panel.  
This node can now be used as a spare node.

**Related concepts**

“SAN Volume Controller menu options” on page 144

During normal operations, menu options are available on the front panel display of the SAN Volume Controller node.

## Replacing a faulty node in the cluster using the CLI

You can use the command-line interface (CLI) and the SAN Volume Controller front panel to replace a faulty node in the cluster.

Before you attempt to replace a faulty node with a spare node you must ensure that you meet the following requirements:

- You know the name of the cluster that contains the faulty node.
- A spare node is installed in the same rack as the cluster that contains the faulty node.
- You must make a record of the last five characters of the original worldwide node name (WWNN) of the spare node. If you repair a faulty node, and you want to make it a spare node, you can use the WWNN of the node. You do not want to duplicate the WWNN because it is unique. It is easier to swap in a node when you use the WWNN.

**Attention:** Never connect a node with a WWNN of 00000 to the cluster. If this node is no longer required as a spare and is to be used for normal attachment to a cluster, you must change the WWNN to the number you recorded when a spare was created. Using any other number might cause data corruption.



If a node fails, the cluster continues to operate with degraded performance until the faulty node is repaired. If the repair operation takes an unacceptable amount of time, it is useful to replace the faulty node with a spare node. However, the appropriate procedures must be followed and precautions must be taken so you do *not* interrupt I/O operations and compromise the integrity of your data.

The following table describes the changes that are made to your configuration when you replace a faulty node in the cluster:

Node attributes	Description
Front panel ID	This is the number that is printed on the front of the node and is used to select the node that is added to a cluster.
Node ID	This is the ID that is assigned to the node. A new node ID is assigned each time a node is added to a cluster; the node name remains the same following service activity on the cluster. You can use the node ID or the node name to perform management tasks on the cluster. However, if you are using scripts to perform those tasks, use the node name rather than the node ID. This ID will change during this procedure.
Node name	This is the name that is assigned to the node. If you are using SAN Volume Controller version 5.1.0 nodes, the SAN Volume Controller automatically re-adds nodes that have failed back to the cluster. If the cluster reports an error for a node missing (error code 1195) and that node has been repaired and restarted, the cluster automatically re-adds the node back into the cluster. For releases prior to 5.1.0, if you do not specify a name, the SAN Volume Controller assigns a default name. The SAN Volume Controller creates a new default name each time a node is added to a cluster. If you choose to assign your own names, you must type the node name on the Adding a node to a cluster panel. You cannot manually assign a name that matches the naming convention used for names assigned automatically by SAN Volume Controller. If you are using scripts to perform management tasks on the cluster and those scripts use the node name, you can avoid the need to make changes to the scripts by assigning the original name of the node to a spare node. This name might change during this procedure.
Worldwide node name	This is the WWNN that is assigned to the node. The WWNN is used to uniquely identify the node and the fibre-channel ports. During this procedure, the WWNN of the spare node changes to that of the faulty node. The node replacement procedures must be followed exactly to avoid any duplication of WWNNs. This name does not change during this procedure.
Worldwide port names	<p>These are the WWPNS that are assigned to the node. WWPNS are derived from the WWNN that is written to the spare node as part of this procedure. For example, if the WWNN for a node is 5005076801000F6, the four WWPNS for this node are derived as follows:</p> <pre> WWNN                               5005076801000F6 WWNN displayed on front panel     000F6 WWPN Port 1                       50050768014000F6 WWPN Port 2                       50050768013000F6 WWPN Port 3                       50050768011000F6 WWPN Port 4                       50050768012000F6 </pre> <p>These names do not change during this procedure.</p>

Complete the following steps to replace a faulty node in the cluster:

1. Verify the name and ID of the node that you want to replace.

Complete the following step to verify the name and ID:

- a. Issue the **svcinfo lsnode** CLI command to ensure that the partner node in the I/O group is online.
  - If the other node in the I/O group is offline, start Directed Maintenance Procedures (DMPs) to determine the fault.
  - If you have been directed here by the DMPs, and subsequently the partner node in the I/O group has failed, see the procedure for recovering from offline VDisks after a node or an I/O group failed.
  - If you are replacing the node for other reasons, determine the node you want to replace and ensure that the partner node in the I/O group is online.
  - If the partner node is offline, you will lose access to the VDisks that belong to this I/O group. Start the DMPs and fix the other node before proceeding to the next step.

2. Find and record the following information about the faulty node using Steps 2a through 2h:

- Node serial number
  - Worldwide node name
  - All of the worldwide port names
  - Name or ID of the I/O group that contains the node
  - Front panel ID
  - Uninterruptible power supply serial number
- a. Issue the **svcinfo lsnode** CLI command to find and record the node name and I/O group name. The faulty node will be offline.
  - b. Issue the following CLI command:

```
svcinfo lsnodevpd nodename
```

Where *nodename* is the name that you recorded in step 2a.
  - c. Find the WWNN field in the output.
  - d. Record the last five characters of the WWNN.
  - e. Find the `front_panel_id` field in the output.
  - f. Record the front panel ID.
  - g. Find the `UPS_serial_number` field in the output.
  - h. Record the uninterruptible power supply serial number.

3. Ensure that the faulty node has been powered off.

4. Issue the following CLI command to remove the faulty node from the cluster:

```
svctask rmnode nodename/id
```

Where *nodename/id* is the name or ID of the faulty node.

5. Disconnect all four fibre-channel cables from the node.

**Important:** Do not plug the fibre-channel cables into the spare node until the spare node is configured with the WWNN of the faulty node.

6. Connect the power and signal cables from the spare node to the uninterruptible power supply that has the serial number you recorded in step 2h.

**Note:** For 2145 UPS-1U units, you must disconnect the cables from the faulty node.

7. Disconnect the faulty node's power and serial cable from the 2145 UPS-1U and connect the new node's power and signal cable in their place.
8. Power on the spare node.
9. Display the node status on the front-panel display.
10. You must change the WWNN of the spare node to that of the faulty node. The procedure for doing this depends on the SAN Volume Controller version that is installed on the spare node. Press and release the down button until the Node: panel displays. Then press and release the right button until the WWNN: panel displays. If repeated pressing of the right button returns you to the Node: panel, without displaying a Node WWNN: panel, go to step 12; otherwise, continue with step 11.
11. Change the WWNN of the spare node (with SAN Volume Controller V4.3 and above installed) to match the WWNN of the faulty node by completing the following steps:
  - a. With the Node WWNN: panel displayed, press and hold the down button, press and release the select button, and then release the down button. The display switches into edit mode. Edit WWNN is displayed on line 1. Line 2 of the display contains the last five numbers of the WWNN.
  - b. Change the WWNN that is displayed to match the last five numbers of the WWNN that you recorded in step 13. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.
  - c. When the five numbers match the last five numbers of the WWNN that you recorded in step 2d on page 114, press the select button to accept the numbers.
12. Change the WWNN of the spare node (with SAN Volume Controller versions prior to V4.3 installed) to match the WWNN of the faulty node by performing the following steps:
  - a. Press and release the right button until the Status: panel is displayed.
  - b. With the node status displayed on the front panel, press and hold the down button; press and release the select button; release the down button. WWNN is displayed on line 1 of the display. Line 2 of the display contains the last five numbers of the WWNN.
  - c. With the WWNN displayed on the front panel; press and hold the down button; press and release the select button; release the down button. The display switches into edit mode.
  - d. Change the WWNN that is displayed to match the last five numbers of the WWNN that you recorded in step 2d on page 114. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.
  - e. When the five numbers match the last five numbers of the WWNN that you recorded in step 2d on page 114, press the select button to retain the numbers that you have updated and return to the WWNN panel.
  - f. Press the select button to apply the numbers as the new WWNN for the node.
13. Connect the four fibre-channel cables that you disconnected from the faulty node to the spare node.

If the spare node has less Ethernet cables connected than the faulty node, move the Ethernet cables from the faulty node to the spare node. Ensure you connect the cable into the same port on the spare node as it was in on the faulty node.

14. Issue the following command to add the spare node to the cluster:

```
svctask addnode -wwnodename WWNN -iogrp iogroupname/id
```

where *WWNN* and *iogroupname/id* are the values that you recorded for the original node.

The SAN Volume Controller V5.1 automatically reassigns the node with the name that was used originally. For versions prior to V5.1, use the **name** parameter with the svctask addnode command to assign a name. If the original node's name was automatically assigned by SAN Volume Controller, it is not possible to reuse the same name. It was automatically assigned if its name starts with node. In this case, either specify a different name that does not start with node or do not use the **name** parameter so that SAN Volume Controller automatically assigns a new name to the node.

If necessary, the new node is updated to the same SAN Volume Controller software version as the cluster. This update can take up to 20 minutes.

15. Use the tools that are provided with your multipathing device driver on the host systems to verify that all paths are now online. See the documentation that is provided with your multipathing device driver for more information. For example, if you are using the subsystem device driver (SDD), see the *IBM System Storage Multipath Subsystem Device Driver User's Guide* for instructions on how to use the SDD management tool on host systems. It might take up to 30 minutes for the paths to come online.

16. Repair the faulty node.

**Attention:** When the faulty node is repaired, do not connect the fibre-channel cables to it. Connecting the cables might cause data corruption because the spare node is using the same WWNN as the faulty node.

If you want to use the repaired node as a spare node, perform the following steps.

**For SAN Volume Controller V4.3 and later versions:**

- a. With the Node WWNN: panel displayed, press and hold the down button, press and release the select button, and then release the down button.
- b. The display switches into edit mode. Edit WWNN is displayed on line 1. Line 2 of the display contains the last five numbers of the WWNN.
- c. Change the displayed number to 00000. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.
- d. Press the select button to accept the numbers.

This node can now be used as a spare node.

**For SAN Volume Controller versions prior to V4.3:**

- a. Press and release the right button until the Status: panel is displayed.
- b. With the node status displayed on the front panel, press and hold the down button; press and release the select button; release the down button. WWNN is displayed on line 1 of the display. Line 2 of the display contains the last five numbers of the WWNN.
- c. With the WWNN displayed on the front panel; press and hold the down button; press and release the select button; release the down button. The display switches into edit mode.

- d. Change the displayed number to 00000. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.
- e. Press the select button to accept the numbers.
- f. Press the select button to retain the numbers that you have updated and return to the WWNN panel.

This node can now be used as a spare node.

#### **Related concepts**

“SAN Volume Controller menu options” on page 144

During normal operations, menu options are available on the front panel display of the SAN Volume Controller node.

---

## **Updating license settings**

You can use the SAN Volume Controller Console to update your license settings.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to update the license settings:

1. In the portfolio, click **Service and Maintenance** → **License Settings**. The License Settings panel is displayed.
2. Choose Capacity Licensing or Physical Disk Licensing and click **Go**.
3. Enter your license settings and click **Update License Settings**.
4. The updated license information is displayed. To confirm that the settings match your license agreement, click **I Agree**.

---

## **Viewing and updating license settings using the CLI**

You can use the command-line interface (CLI) to view and update your license settings.

SAN Volume Controller provides two license options: Physical Disk Licensing and Capacity Licensing. Perform the following steps to view and update your SAN Volume Controller license settings:

1. Issue the **svcinfo lslicense** CLI command to view the current license settings for the cluster.
2. Issue the **svctask chlicense** CLI command to change the licensed settings of the cluster.

#### **Attention:**

- License settings are entered when the cluster is first created; do not update the settings unless you have changed your license.
- To select Physical Disk Licensing, run the **svctask chlicense** command with one or more of the **physical\_disks**, **physical\_flash**, and **physical\_remote** parameters.
- To select Capacity Licensing, run the **svctask chlicense** command with one or more of the **-flash**, **-remote**, and **-virtualization** parameters.

For detailed license command usage information, see the *IBM System Storage SAN Volume Controller Command-Line Interface User's Guide*.

---

## Viewing the license settings log

You can view the license settings log for the cluster from the License Settings panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following step to view the license settings log for the cluster:

Click **Service and Maintenance** → **View License Settings Log** in the portfolio. The License Settings panel is displayed.

---

## Collecting SSD dump files using the CLI

You can use the command-line interface (CLI) to collect dump files from solid-state drives (SSDs).

To collect internal log files from solid-state drive (SSD) MDisks, run the `triggermdiskdump` command. Subsequently, you can list, delete or copy the dump files.

The `triggermdiskdump` command generates a dump file and saves it in the `/dumps/mdisk` directory on the node that contains the SSD.

1. Issue the **`svctask triggermdiskdump`** CLI command.

The following example shows the CLI format for generating a dump file for the specified SSD MDisk:

```
svctask triggermdiskdump mdisk_id | mdisk_name
```

2. Issue the **`svcinfolsmdiskdumps`** command to list files in the `/dumps/drive` directory on the specified node.

The following example shows the CLI format for listing the dump files for the specified node:

```
svcinfolsmdiskdumps node_id | node_name
```

3. Issue the **`svctask clear.dumps`** command to delete all files from the `/dumps` directory and all subdirectories on the specified node. To delete files from a subdirectory of `/dumps` only, specify the `-prefix` parameter.

The following example shows the CLI format for deleting all dump files from the specified node:

```
svctask clear.dumps node_id | node_name
```

The following example shows the CLI format for deleting only the dump files in the specified `/elogs/` directory:

```
svctask clear.dumps -prefix "/dumps/elogs/*"
```

4. Issue the **`svctask cp.dumps`** command to copy dump files to the configuration node. If the `/dumps` directory on the configuration node becomes full before the copy completes, no message is returned. To avoid this scenario, clear the `/dumps` directory after migrating data from the configuration node.

The following example shows the CLI format for copying all dump files from the specified node to the configuration node:

```
svctask cp.dumps -prefix /dumps node_id | node_name
```

---

## Displaying and saving log and dump files

You can save the log and dump files for nodes.

You can save dump data for any node in the cluster. When you use this procedure to display dump data only, the dump files for the configuration node are displayed. An option on the dumps menu allows you to display data from other nodes. If you choose to display or save data from another node, that data is first copied to the configuration node.

The software dump files contain dumps of the SAN Volume Controller memory. Your IBM service representative might ask for these dumps to debug problems. The software dumps are large files (approximately 300 MB). Consider copying these files to your host using secure copy methods.

The **List dumps** option supports the following file types:

- Error logs
- I/O statistic logs
- I/O trace logs
- Feature logs
- Software dumps
- Audit logs
- CIMOM logs
- Managed Disks (MDisks) logs

Complete the following steps to display log and dump files:

This task assumes that you have already launched the SAN Volume Controller Console.

1. Click **Service and Maintenance** → **List Dumps** in the portfolio. The List Dumps panel is displayed.

The List dumps (other nodes) continued panel displays the number of log files or dumps of a particular type that are available on the cluster. If there is more than one node in the cluster, the **Check other nodes** button is displayed. If you click this button, the log files and dumps for all nodes that are part of the cluster are displayed. Dumps and logs on all nodes in the cluster can be deleted or copied to the configuration node.

If you click on one of the file types, all the files of that type are listed in a table.

**Note:** For error logs and software dumps, the file names include the node name and time and date as part of the file name.

2. Copy the files to your local workstation by right-clicking on the filename and using the **Save Link As...** (Netscape) or **Save Target As...** (Internet Explorer) option from the Web browser.

---

## Listing log or dump files using the CLI

You can list log or dump files using the command-line interface (CLI).

Perform the following steps to list log or dump files:

1. Open a command prompt window.
2. You can issue any of the following commands to list error log files:
  - a. **svcinfo lserrlogbydisk**
  - b. **svcinfo lserrlogbydiskgroup**
  - c. **svcinfo lserrlogbyvdisk**

- d. **svcinfolerrlogbyhost**
- e. **svcinfolerrlogbynode**
- f. **svcinfolerrlogbyiogrp**
- g. **svcinfolerrlogbyfcconsistgrp**
- h. **svcinfolerrlogbyfcmap**
- i. **svcinfolerrlogbyrcconsistgrp**
- j. **svcinfolerrlogbyrcrelationship**

These commands will list the error log by type. These commands will return a list of dumps in the appropriate directory. For example, issue the **svcinfolerrlogbydisk** command, displays the error log by MDisk.

You can display the whole log or filter the log so that only errors, events, or unfixed errors are displayed. In addition, you can request the output to be sorted either by error priority or by time. For error priority, the most serious errors are the lowest-numbered errors. They are, therefore, displayed first in the table. For time, either the older or the latest entry can be displayed first in the output.

3. You can issue any of the following command to list dump files.

- a. **svcinfolconfigdumps**
- b. **svcinfolerrlogdumps**
- c. **svcinfolfeaturedumps**
- d. **svcinfoliostatsdumps**
- e. **svcinfoliotracedumps**
- f. **svcinfolsoftwaredumps**
- g. **svcinfol2145dumps**

These commands will list the dump file by type. These commands will return a list of dumps in the appropriate directory. For example, issue the **svcinfolconfigdumps** command, a list of dumps for configurations will be stored in the /dumps/configs destination directory.

The software dump files contain dumps of the SAN Volume Controller memory. Your IBM service representative might ask for these dumps to debug problems. The software dumps are large files (approximately 300 MB). Consider copying these files to your host using secure copy (scp) methods.

---

## Dumping cluster configurations

You can dump the data for the cluster configuration into a file from the Dumping the Cluster Configuration panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to dump the data for the cluster configuration into a file:

1. Click **Service and Maintenance** → **Dump Configuration** in the portfolio. The Dump Configuration panel is displayed.
2. Type the prefix for the name of the file that contains the cluster configuration.
3. Click **OK**.



---

## Viewing the feature log using the CLI

You can use the command-line interface (CLI) to view the feature log.

Perform the following steps to view the feature log:

1. Issue the **svcinfo lsfeaturedumps** command to return a list of dumps in the /dumps/feature destination directory. The feature log is maintained by the cluster. The feature log records events that are generated when license parameters are entered or when the current license settings have been breached.
2. Issue the **svcservicemodeinfo lsfeaturedumps** command to return a list of the files that exist of the type specified on the given node.

---

## Analyzing the error log

You can analyze the error log from the Analyze Error Log panel.

This task assumes that you have already launched the SAN Volume Controller Console.

**Note:** Log files that are copied to the configuration node are *not* automatically deleted by the SAN Volume Controller.

Perform the following steps to analyze the error log:

1. Click **Service and Maintenance** → **Analyze Error Log** in the portfolio. The Analyze Error Log panel is displayed.

The Analyze Error Log panel lets you analyze the cluster error log. You can display the whole log or filter the log so that only errors, events, or unfixed errors are displayed. In addition, you can request that the table is sorted by either error priority or time. For error priority, the most serious errors are the lowest-numbered errors. Therefore, they are displayed first in the table.

Either the oldest or the latest entry can be displayed first in the table. You can also select how many error log entries are displayed on each page of the table. The default is set to 10 and the maximum number of error logs that can be displayed on each page is 99.

2. After selecting the options, click **Process** to display the filtered error log in the table. The Analyze Error Log Continued panel is displayed.

Forward and backward scroll buttons are displayed, depending on the existing page number and the total number of pages that are in the table. If the table contains more than two pages of entries, a **Go to** input area is displayed in the table footer. This input area enables you to skip to a particular page number.

If you click on the sequence number of a table record, more information about that error log entry is displayed. If the record is an error (instead of an event), you can change the fixed or unfixed status of the record; that is, you can mark an unfixed error as fixed or a fixed error as unfixed.

3. Click **Clear log** to erase the entire cluster error log.

**Note:** Clicking **Clear log** does *not* fix the existing errors.

---

## Analyzing the error log using the CLI

You can use the command-line interface (CLI) to analyze the error log.

Perform the following steps to analyze the error log:

Issue any of the following CLI commands to list error log entries by file type:

- **svcinfo lserrlogbymdisk**
- **svcinfo lserrlogbymdiskgroup**
- **svcinfo lserrlogbyvdisk**
- **svcinfo lserrlogbyhost**
- **svcinfo lserrlogbynode**
- **svcinfo lserrlogbyiogrp**
- **svcinfo lserrlogbyfcconsistgrp**
- **svcinfo lserrlogbyfcmap**
- **svcinfo lserrlogbyrcconsistgrp**
- **svcinfo lserrlogbyrcrelationship**

These commands list the error log entries by type. For example, the **svcinfo lserrlogbymdisk** command displays the error log by managed disks (MDisks).

You can display the whole log or filter the log so that only errors, events, or unfixed errors are displayed. You can also request that the output is sorted either by error priority or by time. For error priority, the most serious errors are the lowest-numbered errors. Therefore, the most serious errors are displayed first in the table. For time, either the older or the latest entry can be displayed first in the output.

---

## Marking errors as fixed

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

This task assumes that you have already launched the SAN Volume Controller Console. Perform the following steps to mark errors as fixed:

1. Select a cluster and launch the SAN Volume Controller Console.
2. Click **Service and Maintenance** → **Analyze Error Log** in the portfolio.
3. Select the analysis option from the list of options:
  - Show all errors and events
  - Show all errors
  - Show all events
  - Show all unfixed errors
  - Show all errors or events matching code
4. Select the display option:
  - Sort by error priority
  - Sort by date with the newest first
  - Sort by date with the oldest first
  - Number of entries to display (per page)

5. Click **Process**
6. Click on the sequence number of the error that you have just fixed to display the error log in more detail.
7. Click **Mark Error As Fixed**.  
The error log entry for this error is now marked as fixed and it will not be used as part of future error log analysis.

---

## Running the cluster maintenance procedure

You can use the SAN Volume Controller Console to run the cluster maintenance procedure.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to run the cluster maintenance procedure:

1. Click **Service and Maintenance** → **Run Maintenance Procedures** in the portfolio. The Maintenance Procedures panel is displayed.
2. Click **Start Analysis** to analyze the cluster error log. The Maintenance panel is displayed.

If you click the error code of a error log entry, you are guided through a series of actions that help you estimate the state of the cluster and determine if the error was an isolated event or a component failure. If a component has failed, it might be necessary to exchange that component. Where necessary, images of the failing component are displayed. If a repair is performed successfully, the state of an error record in the error log changes from an unfixed error to a fixed error.

---

## Using directed maintenance procedures

You can use directed maintenance procedures (DMP) to diagnose and resolve problems with the SAN Volume Controller.

For example, to repair a SAN Volume Controller cluster, you might perform the following tasks:

- Analyze the error log
- Replace failed components
- Verify the status of a repaired device
- Restore a device to an operational state in the cluster
- Mark the error as fixed in the error log

Directed maintenance simplifies these procedures by automating as many of the tasks as possible.

The following procedures and panels show examples of what you might see when you use the DMPs. The specific DMPs and panels that you see depend on the procedure that you select.

**Note:** The following procedures and graphics are examples of what you might see when you use the DMPs. The DMPs and graphics that you see depend on the procedure that you select. This task assumes that you have already launched the SAN Volume Controller Console.

Our example uses the SAN Volume Controller Console to repair a SAN Volume Controller cluster. Perform the following steps to start the DMP:

1. Click **Start Analysis** from the Run Maintenance Procedures panel.

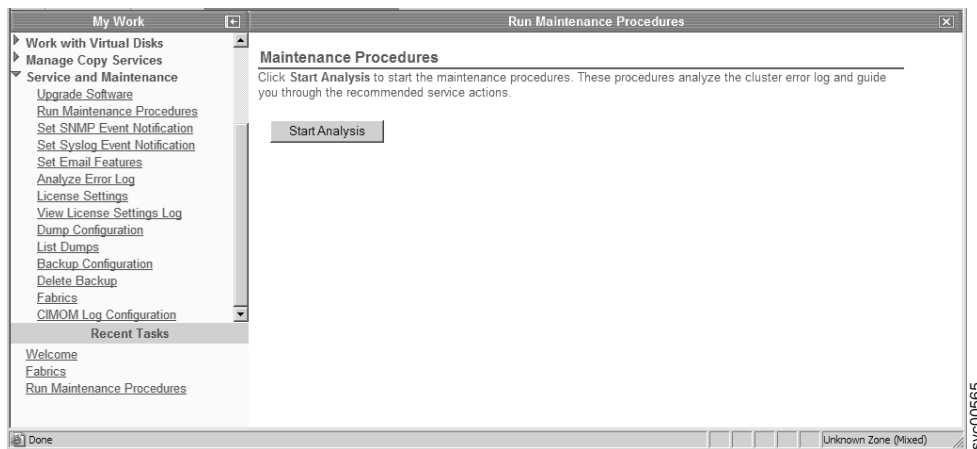
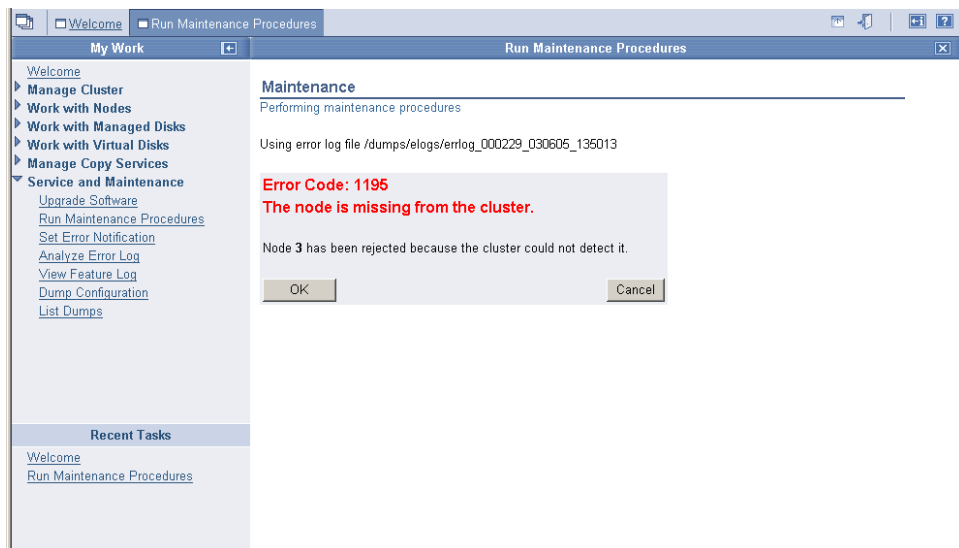


Figure 54. Start Analysis panel

The list might contain any number of errors that must be repaired. If there is more than one error on the list, the error at the top of the list has the highest priority and must always be fixed first. If you do not fix the higher priority errors first, you might not be able to fix the lower priority errors.

2. Click **OK**.



3. From the **Error Code** column, click on the number of the error code. The panel displays the error code and provides a description of the condition. You can click **Continue** to go forward or **Cancel** to return to the previous panel.
4. One or more panels might be displayed with instructions for you to replace parts or perform other repair activity. If you are not able to complete the actions at this time, click **Cancel**, which returns you to the previous panel. Click **Cancel** until you are returned to the Viewing Clusters panel. When you return to the DMPs, the repair can be restarted from step 1. When the actions that you are

| instructed to perform are complete, click **OK**. When the last repair action is  
| completed, the DMPs might attempt to restore failed devices to the cluster.

- | 5. After you complete the fix, you see the statement Click OK to mark the error  
| as fixed. Click **OK**. This action marks the error as fixed in the error log and  
| prevents this instance of the error from being listed again.
- | 6. When you see the statement The repair has been completed., click **Exit**. If  
| other errors need to be fixed, those errors are displayed and the DMPs  
| continue.
- | 7. If no errors remain, you are shown the following statement: There are no  
| unfixed errors in the error log.

| The statement indicates that no further repair procedures are necessary. Verify  
| the repair using the repair verification MAP.

### **Related tasks**

"MAP 5000: Start" on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs)  
for the SAN Volume Controller.

"MAP 5700: Repair verification" on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs)  
that you have exchanged for new FRUs, or repair actions that have been done  
have solved all the problems on the SAN Volume Controller.



---

## Chapter 4. Viewing the vital product data

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

The VPD lists the following types of information:

- Cluster-related values such as the software version, space in MDisk groups, and space allocated to VDisks.
- 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 cluster which 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 cluster VPD and the node VPD. You can also view the VPD through the SAN Volume Controller graphical user interface.

### Related reference

“Fields for the node VPD” on page 130

The node vital product data (VPD) provides information for items such as the system board, processor, fans, memory module, adapter card, devices, software, front panel assembly, the uninterruptible power supply, SAS host bus adapter (HBA), and SAS solid-state drive (SSD).

“Fields for the cluster VPD” on page 136

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

---

## Viewing the vital product data

You can view the vital product data for a node from the Viewing Vital Product Data panel of the SAN Volume Controller Console.

Perform the following steps to view the vital product data for a node:

1. Click **Work With Nodes** in the portfolio.
2. Click **Nodes** in the portfolio. The Nodes panel is displayed.
3. Click on the node for which you want to display the details.
4. Click **Vital Product Data** to view the data.
5. Click **Close** to return to the Viewing Vital Product Data panel.

---

## Viewing cluster properties

You can use the SAN Volume Controller Console to view the properties for a cluster.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the properties of a cluster:

1. Click **Manage Cluster** → **View Cluster Properties** in the portfolio. The Viewing General Properties panel is displayed.

2. Click the following tabs:
  - a. **General** to display the general properties
  - b. **IP Addresses** to view the IP addresses that are used by the cluster
  - c. **Remote Authentication** to view attributes for the remote authentication service, which is used by remote users to access the cluster
  - d. **Space** to view the space and capacity for managed disks (MDisks), MDisk groups, and virtual disks (VDisks)
  - e. **Statistics** to view the cluster statistics details
  - f. **Metro Mirror and Global Mirror** to view the Metro Mirror or Global Mirror properties of the cluster
  - g. **iSCSI** to view the iSCSI properties for the cluster
  - h. **SNMP** to view the SNMP properties for the cluster
  - i. **Syslog** to view the syslog properties for the cluster
  - j. **E-Mail Servers** to view the e-mail server properties for the cluster
  - k. **E-Mail Users** to view the e-mail user properties for the cluster
3. Click **Close** to close the panel.

---

## Displaying the vital product data using the CLI

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

Issue the following CLI commands to display the VPD:

```
svcinfolnodevpd nodename
svcinfolcluster clustername
svcinfolclusterip
```

**Note:** For the SAN Volume Controller 2145-8A4, 2145-8G4, and 2145-8F4 nodes, the `svcinfolnodevpd nodename` command displays the device serial number of the fibre-channel card as "N/A."

For more information about the commands, see the *IBM System Storage SAN Volume Controller Command-Line Interface User's Guide*.

### Related tasks

"Viewing the node status" on page 66

You can view the properties for a node from the Viewing General Details panel.

## Displaying node properties using the CLI

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

Perform the following steps to display the node properties:

1. Issue the `svcinfolnode` CLI command to display a concise list of nodes in the cluster.

The following is an example of the CLI command you can issue to list the nodes in the cluster:

```
svcinfolnode -delim :
```

The following is an example of the output that is displayed:



```
id:name:UPS_serial_number:WWN:status:IO_group_id:IO_group_name:config_node:UPS_unique_id:hardware:iscsi_name:iscsi_alias
1:group1node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202378101C0D18D8:8A4:iqn.1986-03.com.ibm:2145.cluster1:group1node1:
2:group1node2:10L3ANF:5005076801000009:online:0:io_grp0:no:202378101C0D1796:8A4:iqn.1986-03.com.ibm:2145.cluster1:group1node2:
3:group2node1:10L3ASH:5005076801000001:online:1:io_grp1:no:202378101C0D18D8:8A4:iqn.1986-03.com.ibm:2145.cluster1:group2node1:
4:group2node2:10L3ANF:50050768010000F4:online:1:io_grp1:no:202378101C0D1796:8A4:iqn.1986-03.com.ibm:2145.cluster1:group2node2:
```

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

The following is an example of the CLI command you can issue to list detailed output for a node in the cluster:

```
svcinfo lsnode -delim : group1node1
```

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

The following is an example of the output that is displayed:

```
id:1
name:group1node1
UPS_serial_number:10L3ASH
WWN:500507680100002C
status:online
IO_group_id:0
IO_group_name:io_grp0
partner_node_id:2
partner_node_name:group1node2
config_node:yes
UPS_unique_id:202378101C0D18D8
port_id:500507680110002C
port_status:active
port_speed:2GB
port_id:500507680120002C
port_status:active
port_speed:2GB
port_id:500507680130002C
port_status:active
port_speed:2GB
port_id:500507680140003C
port_status:active
port_speed:2GB
hardware:8A4
iscsi_name:iqn.1986-03.com.ibm:2145.ndih11.node2
iscsi_alias
failover_active:no
failover_name:node1
failover_iscsi_name:iqn.1986-03.com.ibm:2145.ndih11.node1
failover_iscsi_alias
```

## Displaying cluster properties using the CLI

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

Perform the following step to display cluster properties:

Issue the **svcinfo lscluster** command to display the properties for a cluster.

The following is an example of the command you can issue:

```
svcinfo lscluster -delim : ldcluster-19
```

where *ldcluster-19* is the name of the cluster.

```

IBM_2145:ldcluster-19:admin>svcinfolcluster -delim : ldcluster-19
id:00000200602052F0
name:ldcluster-19
location:local
partnership:
bandwidth:
total_mdisk_capacity:4205812.3GB
space_in_mdisk_grps:4162044.4GB
space_allocated_to_vdisks:1737236.64GB
total_free_space:2468575.7GB
statistics_status:on
statistics_frequency:10
required_memory:8192
cluster_locale:en_US
time_zone:357 Europe/Athens
code_level:5.1.0.0 (build 16.3.0906260000)
FC_port_speed:2Gb
console_IP:x.xx.xx.xx:xxxx
id_alias:00000200602052F0
gm_link_tolerance:300
gm_intra_cluster_delay_simulation:0
gm_intra_cluster_delay_simulation:0
email_reply:manager@mycompany.com
email_contact:manager
email_contact_primary:01202 123456
email_contact_alternate:
email_contact_location:city
email_state:running
inventory_mail_interval:8
total_vdiskcopy_capacity:2009500.80GB
total_used_capacity:1737157.41GB
total_overallocation:47
total_vdisk_capacity:1322910.92GB
cluster_ntp_IP_address:x.xx.xx.xxx
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

```

---

## Fields for the node VPD

The node vital product data (VPD) provides information for items such as the system board, processor, fans, memory module, adapter card, devices, software, front panel assembly, the uninterruptible power supply, SAS host bus adapter (HBA), and SAS solid-state drive (SSD).

Table 21 shows the fields you see for the system board.

Table 21. 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 cards
	Number of SCSI, IDE, SATA, or SAS devices <b>Note:</b> The service controller is a device.
	Number of power supplies
	Number of high-speed SAS adapters
	Number of local managed disks
	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	

Table 22 shows the fields you see for each processor that is installed.

Table 22. Fields for the processors

Item	Field name
Processor	Part number
	Processor location
	Number of caches
	Manufacturer
	Version
	Speed
	Status
	CPU part number

|

Table 23 shows the fields that you see for each fan that is installed.

*Table 23. Fields for the fans*

<b>Item</b>	<b>Field name</b>
Fan	Part number
	Location

Table 24 shows the fields that are repeated for each installed memory module.

*Table 24. Fields that are repeated for each installed memory module*

<b>Item</b>	<b>Field name</b>
Memory module	Part number
	Device location
	Bank location
	Size (MB)

Table 25 shows the fields that are repeated for each installed adapter card.

*Table 25. Fields that are repeated for each adapter card that is installed*

<b>Item</b>	<b>Field name</b>
Adapter card	Card type
	Part number
	Port numbers
	Location
	Device serial number
	Manufacturer
	Device
	Card revision
	Chip revision

Table 26 on page 133 shows the fields that are repeated for each direct-attached managed disk.

Table 26. Fields for the direct-attached managed disk

Item	Field name
Direct-attached managed disk	Part number
	Manufacturer
	Model
	Serial number
	Firmware
	FPGA firmware
	Capacity
	Type
	Speed
	Expansion tray
	Location
Connection type	

Table 27 shows the fields that are repeated for each device that is installed.

Table 27. 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

Table 28 shows the fields that are specific to the node software.

Table 28. Fields that are specific to the node software

Item	Field name
Software	Code level
	Node name
	Worldwide node name
	ID

Table 29 on page 134 shows the fields that are provided for the front panel assembly.

Table 29. Fields that are provided for the front panel assembly

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

Table 30 shows the fields that are provided for the Ethernet port.

Table 30. Fields that are provided for the Ethernet port

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

Table 31 shows the fields that are provided for the power supplies in the node.

Table 31. Fields that are provided for the power supplies in the node

Item	Field name
Power supplies	Part number
	Serial number
	Location

Table 32 shows the fields that are provided for the uninterruptible power supply assembly that is powering the node.

Table 32. Fields that are provided for the uninterruptible power supply assembly that is powering the node

Item	Field name
Uninterruptible power supply	Electronics assembly part number
	Battery part number
	Frame assembly part number
	Input power cable part number
	UPS serial number
	UPS type
	UPS internal part number
	UPS unique ID
	UPS main firmware
UPS communications firmware	

Table 33 on page 135 shows the fields that are provided for the SAS host bus adapter (HBA).

Table 33. 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
	Card revision
	Chip revision

Table 34 shows the fields that are provided for the SAS solid-state drive (SSD).

Table 34. Fields that are provided for the SAS solid-state drive (SSD)

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 35 on page 136 shows the fields that are provided for the small form factor pluggable (SFP) connector.

Table 35. Fields that are provided for the small form factor pluggable (SFP) connector

Item	Field name
Small form factor pluggable (SFP) connector	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

#### Related reference

“Fields for the cluster VPD”

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

---

## Fields for the cluster VPD

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

Table 36 shows the fields that are provided for the cluster properties as shown by the SAN Volume Controller Console.

Table 36. Fields that are provided for the cluster properties

Item	Field name
General	ID <b>Note:</b> This is the unique identifier for the cluster.
	Name
	Location
	Time Zone
	Required Memory
	Licensed Code Version
	Channel Port Speed  <b>Note:</b> This field represents the speed at which non-negotiating nodes in the cluster will run, for example, the SAN Volume Controller 2145-8F2. All other models that are capable of speed negotiation are not affected by the speed value that is indicated in this field.



Table 36. Fields that are provided for the cluster properties (continued)

Item	Field name
IP Addresses <sup>1</sup>	Cluster Ethernet Port 1 (attributes for both IPv4 and IPv6) <ul style="list-style-type: none"> <li>• IP Address</li> <li>• Service IP Address</li> <li>• Subnet Mask</li> <li>• Prefix</li> <li>• Default Gateway</li> </ul>
	Cluster Ethernet Port 2 (attributes for both IPv4 and IPv6) <ul style="list-style-type: none"> <li>• IP Address</li> <li>• Service IP Address</li> <li>• Subnet Mask</li> <li>• Prefix</li> <li>• Default Gateway</li> </ul>
Remote Authentication	Remote Authentication
	Web Address
	User Name
	Password
	SSL Certificate
Space	Total MDisk Capacity
	Space in MDisk Groups
	Space Allocated to VDIs
	Total Free Space
	Total Used Capacity
	Total Allocation
	Total VDisk Copy Capacity
	Total VDisk Capacity
Statistics	Statistics Status
	Statistics Frequency
Metro and Global Mirror	Link Tolerance
	Inter-Cluster Delay Simulation
	Intra-Cluster Delay Simulation
	Partnership
	Bandwidth
E-mail	SMTP E-mail Server
	E-mail Server Port
	Reply E-mail Address
	Contact Person Name
	Primary Contact Phone Number
	Alternate Contact Phone Number
	Physical Location of the System Reporting Error
	E-mail Status
	Inventory E-mail Interval

Table 36. Fields that are provided for the cluster properties (continued)

Item	Field name
iSCSI	iSNS Server Address
	Supported Authentication Methods
	CHAP Secret
<sup>1</sup> You can also use the <b>svcinfo lsclusterip</b> CLI command to view this data.	

#### Related reference

“Fields for the node VPD” on page 130

The node vital product data (VPD) provides information for items such as the system board, processor, fans, memory module, adapter card, devices, software, front panel assembly, the uninterruptible power supply, SAS host bus adapter (HBA), and SAS solid-state drive (SSD).

---

## Chapter 5. Using the front panel of the SAN Volume Controller

The front panel of the SAN Volume Controller has a display, various LEDs, navigation buttons, and a select button that are used when servicing your SAN Volume Controller node.

Figure 55 shows where the front-panel display **1** is located on the SAN Volume Controller node.




Figure 55. SAN Volume Controller front-panel assembly

### Related reference

“Front-panel display” on page 12

The front-panel display shows service, configuration, and navigation information.

---

## Boot progress indicator

Boot progress is displayed on the front panel of the SAN Volume Controller.

The Boot progress display on the front panel shows that the node is starting.



Figure 56. Example of a boot progress display

During the boot operation, boot progress codes are displayed and the progress bar moves to the right while the boot operation proceeds.

---

## Boot failed

If the boot operation fails, a boot code is displayed.

Failed 180

See "Understanding the boot codes" topic where you can find a description of the failure and the appropriate steps that you must perform to correct the failure.

#### Related information

"Understanding the boot codes" on page 199

The boot codes are displayed on the screen when a node is booting.

---

## Charging

The front panel indicates that the uninterruptible power supply battery is charging.



A node will not start and join a cluster if there is insufficient power in the uninterruptible power supply battery to manage with a power failure. Charging is displayed until it is safe to start the node. This might take up to two hours.

---

## Error codes

Error codes are displayed on the front panel display.

Figure 57 and Figure 58 show how error codes are displayed on the front panel.




Figure 57. Example of a cluster error code



Figure 58. Example of a node error code

For descriptions of the error codes that are displayed on the front panel display, see the various error code topics for a full description of the failure and the actions that you must perform to correct the failure.

**Related reference**

“Understanding the cluster recovery codes” on page 239

Cluster recovery codes indicate that a critical software error has occurred that might corrupt your cluster.

**Related information**

“Understanding the node error codes” on page 223

Node error codes are displayed on the display screen by node software.

“Understanding the create cluster error codes” on page 238

Cluster Create error codes are displayed on the menu screen when you are using the front panel to create a new cluster, but the create operation fails.

“Understanding cluster error codes” on page 239

Every cluster error code includes an error code number, a description, action, and possible field-replaceable units (FRUs).

---

## Hardware boot

The hardware boot display shows system data when power is first applied to the node as the node searches for a disk drive to boot.



If this display remains active for longer than 3 minutes, there might be a problem. The cause might be a hardware failure or the software on the hard disk drive might be missing or damaged.

**Related tasks**

“Determining a hardware boot failure” on page 199

If you see that the hardware boot display stays on the front panel for more than three minutes, the node cannot boot. The cause might be a hardware failure or the software on the hard disk drive is missing or damaged.

---

## Node rescue request

If software is lost, you can use the node rescue process to copy all software from another node.

The node-rescue-request display, which is shown in Figure 59 on page 142, indicates that a request has been made to replace the software on this node. The SAN Volume Controller software is preinstalled on all SAN Volume Controller nodes. This software includes the operating system, the application software, and the SAN Volume Controller publications. It is normally not necessary to replace the software on a node, but if the software is lost for some reason (for example, the hard disk drive in the node fails), it is possible to copy all the software from another node that is connected to the same fibre-channel fabric. This process is known as *node rescue*.



Figure 59. Node rescue display

---

## Power failure

The SAN Volume Controller node uses battery power from the uninterruptible power supply to shut itself down.

The Power failure display shows that the SAN Volume Controller is running on battery power because main power has been lost. All I/O operations have stopped. The node is saving cluster metadata and node cache data to the internal disk drive. When the progress bar reaches zero, the node powers off.

**Note:** When input power is restored to the uninterruptible power supply, the SAN Volume Controller turns on without the front panel power button being pressed.



---

## Powering off

The progress bar on the display shows the progress of the power-off operation.

Powering Off is displayed after the power button has been pressed and while the node is powering off. Powering off might take several minutes.



The progress bar moves to the left when the power is removed.

---

## Recovering

The front panel indicates that the uninterruptible power supply battery is not fully charged.



When a node is active in a cluster but the uninterruptible power supply battery is not fully charged, Recovering is displayed. If the power fails while this message is displayed, the node does not restart until the uninterruptible power supply has charged to a level where it can sustain a second power failure.

---

## Restarting

The front panel indicates when the software on a node is restarting.



The software is restarting for one of the following reasons:

- An internal error was detected.
- The power button was pressed again while the node was powering off.

If you press the power button while powering off, the panel display changes to indicate that the button press was detected; however, the power off continues until the node finishes saving its data. After the data is saved, the node powers off and then automatically restarts. The progress bar moves to the right while the node is restarting.

---

## Shutting down

The front-panel indicator tracks shutdown operations.

The Shutting Down display is shown when you issue a shutdown command to a SAN Volume Controller cluster or a SAN Volume Controller node. The progress bar continues to move to the left until the node turns off.

When the shutdown operation is complete, the node turns off. When you power off a node that is connected to a 2145 UPS-1U, only the node shuts down; the 2145 UPS-1U does not shut down.



---

## Validate WWNN?

The front panel prompts you to validate the WWNN when the worldwide node name (WWNN) that is stored in the service controller (the panel WWNN) does not match the WWNN that is backed up on the SAN Volume Controller disk (the disk WWNN).

Typically, this panel is displayed when the service controller has been replaced. The SAN Volume Controller uses the WWNN that is stored on the service controller. Usually, when the service controller is replaced, you modify the WWNN that is stored on it to match the WWNN on the service controller that it replaced. By doing this, the node maintains its WWNN address, and you do not need to modify the

SAN zoning or host configurations. The WWNN that is stored on disk is the same that was stored on the old service controller.

After it is in this mode, the front panel display will not revert to its normal displays, such as node or cluster options or operational status, until the WWNN is validated. Navigate the Validate WWNN option (shown in Figure 60) to choose which WWNN that you want to use.

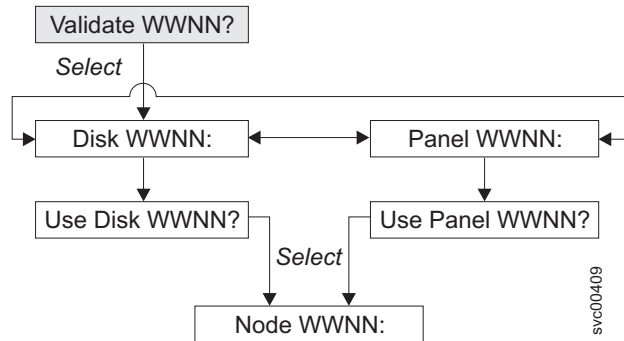


Figure 60. Validate WWNN? navigation

To choose which stored WWNN that you want this node to use, perform the following steps:

1. From the Validate WWNN? panel, press and release the select button. The Disk WWNN: panel is displayed and shows the last five digits of the WWNN that is stored on the disk.
2. To view the WWNN that is stored on the service controller, press and release the right button. The Panel WWNN: panel is displayed and shows the last five numbers of the WWNN that is stored on the service controller.
3. Determine which WWNN that you want to use.
  - a. To use the WWNN that is stored on the disk, perform the following steps:
    - 1) From the Disk WWNN: panel, press and release the down button. The Use Disk WWNN? panel is displayed.
    - 2) Press and release the select button.
  - b. To use the WWNN that is stored on the service controller, perform the following steps:
    - 1) From the Panel WWNN: panel, press and release the down button. The Use Panel WWNN? panel is displayed.
    - 2) Press and release the select button.

The node is now using the selected WWNN. The Node WWNN: panel is displayed and shows the last five numbers of the WWNN that you selected.

If neither WWNN that is stored on the service controller panel or on the disk is suitable, select either WWNN choice and then use the edit WWNN option that is available from the Node WWNN: panel to change the WWNN to the correct value.

---

## SAN Volume Controller menu options

During normal operations, menu options are available on the front panel display of the SAN Volume Controller node.



Menu options enable you to review the operational status of the cluster, node, and external interfaces. They also provide access to the tools that you need to install and to service the node.

Figure 61 shows the sequence of the menu options. Only one option at a time is displayed on the front panel display. For some options, additional data is displayed on line 2. The first option that is displayed is the cluster option.

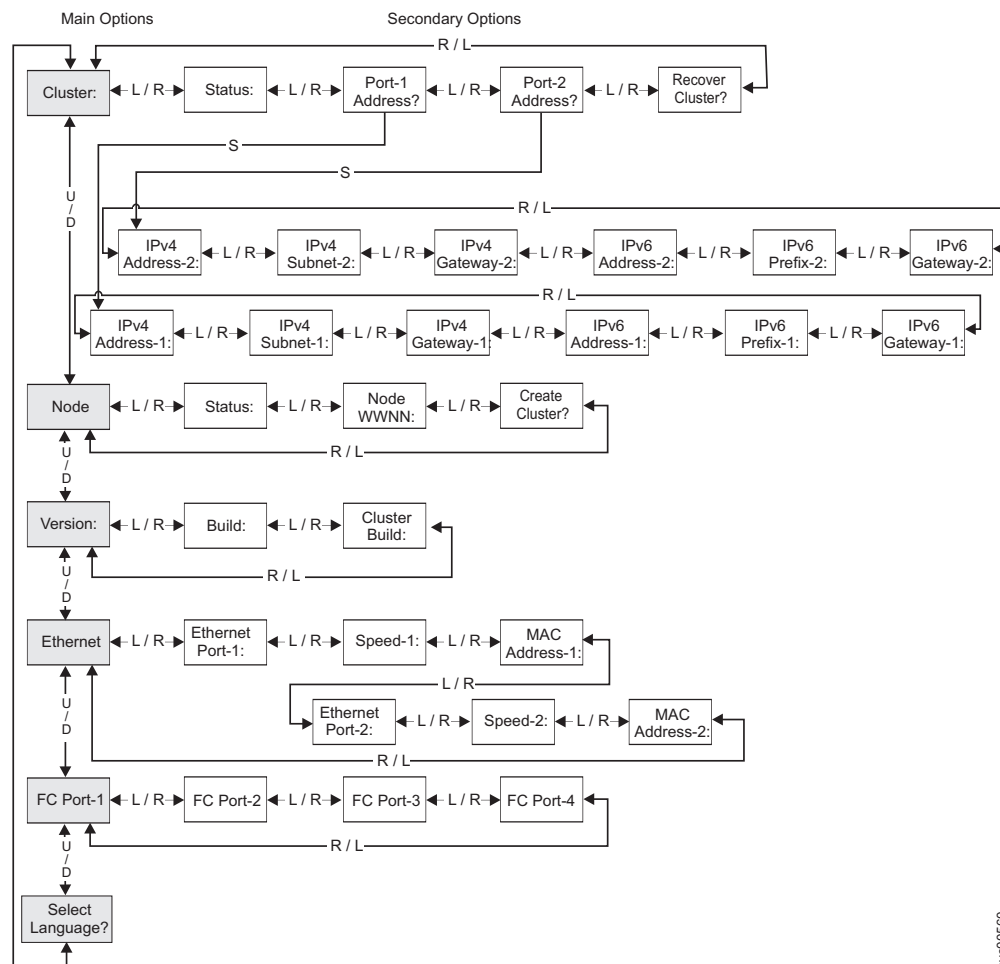


Figure 61. SAN Volume Controller options on the front-panel display

Use the left and right buttons to navigate through the secondary fields that are associated with some of the main fields.

**Note:** Messages might not display fully on the screen. You might see a right angle bracket (>) on the right-hand side of the display screen. If you see a right angle bracket, press the right button to scroll through the display. When there is no more text to display, you can move to the next item in the menu by pressing the right button.

Similarly, you might see a left angle bracket (<) on the left-hand side of the display screen. If you see a left angle bracket, press the left button to scroll through the display. When there is no more text to display, you can move to the previous item in the menu by pressing the left button.

The following main options are available:

- Cluster
- Node
- Version
- Ethernet
- FC port: 1 - 4
- Select language?

#### **Related concepts**

“Cluster options”

The main cluster option from the menu can display the cluster name or can be blank.

“Node options” on page 150

The node option displays the identification number or name of the SAN Volume Controller node.

“Ethernet option” on page 158

The Ethernet option displays the operational state of the Ethernet ports, the speed and duplex information, and their media access control (MAC) addresses.

“Fibre-channel port-1 through 4 option” on page 159

The fibre-channel port-1 through 4 options display the operational status of the fibre-channel ports.

#### **Related tasks**

“Select language? option” on page 159

You can change the language that displays on the front panel.

#### **Related reference**

“Create Cluster?” on page 151

You can create a cluster from the Create Cluster? menu. You also have the option to delete the node from the cluster if the node that is selected is in a cluster.

## **Cluster options**

The main cluster option from the menu can display the cluster name or can be blank.

The main cluster option displays the cluster name that the user has assigned. If a cluster is in the process of being created on the node, and no cluster name has been assigned, a temporary name that is based on the IP address of the cluster is displayed. If this SAN Volume Controller node is not assigned to a cluster, the field is blank.

#### **Related concepts**

“Displaying and editing an IPv6 address” on page 148

After you have set the IPv6 address, you can display the IPv6 addresses and the IPv6 gateway addresses. You can also set the IPv6 addresses for a cluster that you are going to create.

#### **Related reference**

“Recover cluster navigation” on page 149

The Recover cluster? option is useful if the cluster superuser password has been lost or forgotten.

#### **Status**

Status is indicated on the front panel.

This field is blank if this SAN Volume Controller node is not a member of a cluster. If this SAN Volume Controller node is a member of a cluster, the field indicates the operational status of the cluster, as follows:

**Active**

Indicates that this SAN Volume Controller node is an active member of the cluster.

**Inactive**

Indicates that the SAN Volume Controller node is a member of a cluster, but is not now operational. It is not operational because the other SAN Volume Controller nodes that are in the cluster cannot be accessed or because this SAN Volume Controller node was excluded from the cluster.

**Degraded**

Indicates that the cluster is operational, but one or more of the member SAN Volume Controller nodes are missing or have failed.

**IPv4 Address**

A cluster must have either IPv4 addresses or IPv6 addresses that are assigned, or it can have both. If the addresses are assigned, you can use the IPv4 addresses to access the cluster from the command-line tools or the SAN Volume Controller Console.

This field contains the existing IPv4 addresses of the cluster. If this SAN Volume Controller node is not a member of a cluster or IPv4 addresses have not been assigned, this field is blank.

**IPv4 Subnet**

The IPv4 subnet mask addresses are set when the IPv4 addresses are assigned to the cluster.

|  
|  
|

The IPv4 subnet options display the subnet mask addresses when the cluster has IPv4 addresses. If this SAN Volume Controller node is not a member of a cluster or if the IPv4 addresses have not been assigned, this field is blank.

**Related concepts**

“IPv4 Gateway”

The IPv4 gateway addresses are set when the cluster is created.

**IPv4 Gateway**

The IPv4 gateway addresses are set when the cluster is created.

|  
|  
|

The IPv4 gateway options display the gateway addresses for the cluster. If this SAN Volume Controller node is not a member of a cluster, or if the IPv4 addresses have not been assigned, this field is blank.

**Related concepts**

“IPv4 Subnet”

The IPv4 subnet mask addresses are set when the IPv4 addresses are assigned to the cluster.

**IPv6 Address**

You can use the IPv6 addresses to access the cluster from the command line tools or the SAN Volume Controller Console.

|  
|  
|

These options display the existing IPv6 addresses of the cluster. If this SAN Volume Controller node is not a member of a cluster, or if the IPv6 addresses have not been assigned, this field is blank.

### Related concepts

“Displaying and editing an IPv6 address”

After you have set the IPv6 address, you can display the IPv6 addresses and the IPv6 gateway addresses. You can also set the IPv6 addresses for a cluster that you are going to create.

### IPv6 Prefix

The IPv6 prefix is set when a cluster is created.

The IPv6 prefix option displays the network prefix of the cluster and the service IPv6 addresses. The prefix has a value of 0 - 127. If this SAN Volume Controller node is not a member of a cluster, or if the IPv6 addresses have not been assigned, a blank line displays.

### IPv6 Gateway

The IPv6 gateway addresses are set when the cluster is created.

This option displays the IPv6 gateway addresses for the cluster. If this SAN Volume Controller node is not a member of a cluster, or if the IPv6 addresses have not been assigned, a blank line displays.

### Related concepts

“Displaying and editing an IPv6 address”

After you have set the IPv6 address, you can display the IPv6 addresses and the IPv6 gateway addresses. You can also set the IPv6 addresses for a cluster that you are going to create.

### Displaying and editing an IPv6 address

After you have set the IPv6 address, you can display the IPv6 addresses and the IPv6 gateway addresses. You can also set the IPv6 addresses for a cluster that you are going to create.

The IPv6 addresses and the IPv6 gateway addresses consist of eight (4-digit) hexadecimal values that are shown across four panels, as shown in Figure 62. Each panel displays two 4-digit values that are separated by a colon, the address field position (such as 2/4) within the total address, and scroll indicators. Move between the address panels by using the left button or right button.




Figure 62. Viewing the IPv6 address on the front-panel display

You can display or edit the cluster IPv6 addresses or IPv6 gateway addresses when you are creating a cluster from the Create Cluster? menu. You can also display or edit the service mode IPv6 address within service mode.

Perform the following steps to edit an IPv6 address:

1. When the panel is in edit mode, the full address is still shown across four panels as eight (4-digit) hexadecimal values. You edit each digit of the hexadecimal values independently. The current digit is highlighted.
2. Press the up button if you want to increase the value that is highlighted; press the down button if you want to decrease that value.
3. Press the right button or left button to move to the number field that you want to set.

4. Repeat steps 3 and 4 for each number field that you want to set.
5. Press the select button to complete the change or press the right button to display the next secondary option or press the left button to display the previous options.

### Related concepts

“IPv6 Address” on page 147

You can use the IPv6 addresses to access the cluster from the command line tools or the SAN Volume Controller Console.

“IPv6 Gateway” on page 148

The IPv6 gateway addresses are set when the cluster is created.

### Related reference

“Create Cluster?” on page 151

You can create a cluster from the Create Cluster? menu. You also have the option to delete the node from the cluster if the node that is selected is in a cluster.

## Recover cluster navigation

The Recover cluster? option is useful if the cluster superuser password has been lost or forgotten.

Use the recover cluster option (shown in Figure 63) if the user has lost the cluster superuser password or if the user is unable to access the cluster. If it is permitted by the user’s password security policy, use this selection to reset the cluster superuser password. Alternatively, use this selection to enter the node into service mode. This makes the node available through the service IP address.

Navigate to the Service Access? menu by pressing the select button after you access the Recover Cluster? menu.




Figure 63. Recover Cluster? navigation

## Resetting the password

To reset the cluster superuser password on the cluster, complete the following steps from the Service Access? menu:

1. Press and hold the up button.
2. Press and release the select button.
3. Release the up button.

If your password security policy permits password recovery, and if the node is currently a member of a cluster, the cluster superuser password is reset and a new password is displayed for 60 seconds. If your password security policy does not permit password recovery or the node is not a member of a cluster, completing these steps has no effect.

## Entering service mode

Enter service mode only if directed to do so by the IBM Support Center.

This function is capable of degrading the operation of a working cluster. Use it only to recover from a problem that is making the cluster inaccessible.

**Important:** Your cluster can use fixed addresses for its service mode IP addresses. If you do use fixed addresses for your service mode IP addresses, only one node can be in service mode at any time.

To enter service mode, complete the following steps from the Service Access? menu:

1. Press and hold the down button.
2. Press and release the select button.
3. Release the down button.

The node restarts and service mode is enabled. SERVICE MODE is displayed on the screen. You can navigate the front panel to see the current service IP addresses, which can be used to access the node from a Web browser or the CLI. The SERVICE MODE menu continues to be displayed on the front panel until service mode is reset by a command sent to the service IP address, or until the node exits service mode by using the CLI or service Web interface. The node can also be taken out of service mode by turning it off and back on again.

### Related tasks

Chapter 3, “Using the SAN Volume Controller Console and CLI,” on page 63  
The SAN Volume Controller Console is a Web-browser based GUI that is used to manage the cluster. The SAN Volume Controller command-line interface (CLI) is a collection of commands that you can use to manage SAN Volume Controller clusters.

### Related reference

“Service mode” on page 160

While in service mode, you can use the front panel to view or change a service IPv4 or an IPv6 address. You can also view the version and build level of the SAN Volume Controller software that is installed and active on the node.

## Node options

The node option displays the identification number or name of the SAN Volume Controller node.

The main node option displays the identification number of the SAN Volume Controller or the name of the SAN Volume Controller node if the user has assigned a name.

### Related concepts

“Hardware boot” on page 141

The hardware boot display shows system data when power is first applied to the node as the node searches for a disk drive to boot.

### Related reference

“Node identification label” on page 13

The node identification label on the front panel displays a six-digit node identification number. Sometimes this number is called the panel name or front panel ID.

### Status

The node status is indicated on the front panel. The status can be one of the following states:

**Active** The SAN Volume Controller node is operational and assigned to a cluster. It has access to the fibre-channel fabric.

#### Inactive

The SAN Volume Controller node is operational and assigned to a cluster. It has no access to the fibre-channel fabric.

**Free** The SAN Volume Controller node is operational, but has not been assigned to any cluster. It has access to the fibre-channel fabric.

#### Disconnected

The SAN Volume Controller node is operational, but has not been assigned to any cluster. It has no access to the fibre-channel fabric.

**Failed** The SAN Volume Controller node is not operational. A hardware fault is preventing the SAN Volume Controller from being part of a cluster.

### Node WWNN

The node WWNN (worldwide node name) option displays the last five hexadecimal digits of the WWNN that is being used by the SAN Volume Controller node. Only the last five digits of a WWNN vary on a SAN Volume Controller node. The first 11 digits are always 50050768010.

To edit the WWNN, complete the following steps:

**Important:** Only change the WWNN when you are instructed to do so by a service procedure. Nodes must always have a unique WWNN. If you change the WWNN, you might have to reconfigure hosts and the SAN zoning.

1. Press and hold the down button, press and release the select button, and then release the down button. The display switches into edit mode. Edit WWNN is displayed on line 1. Line 2 of the display shows the last five numbers of the WWNN that is currently set. The first number is highlighted.
2. Edit the highlighted number to match the number that is required. Use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.
3. When the highlighted value matches the required number, press and release the select button to activate the change. The Node WWNN: panel displays and the second line shows the last five characters of the changed WWNN.

### Create Cluster?

You can create a cluster from the Create Cluster? menu. You also have the option to delete the node from the cluster if the node that is selected is in a cluster.

The Create Cluster? option allows you to create a new SAN Volume Controller cluster. Press the select button on the Create cluster? menu to start creating a cluster. Figure 64 shows the sequence of the create cluster menu options.

From the front panel, when you create a cluster, you can set either the IPv4 or the IPv6 address for Ethernet port 1. If required, you can add more management IP addresses by using the SAN Volume Controller Console or the CLI.

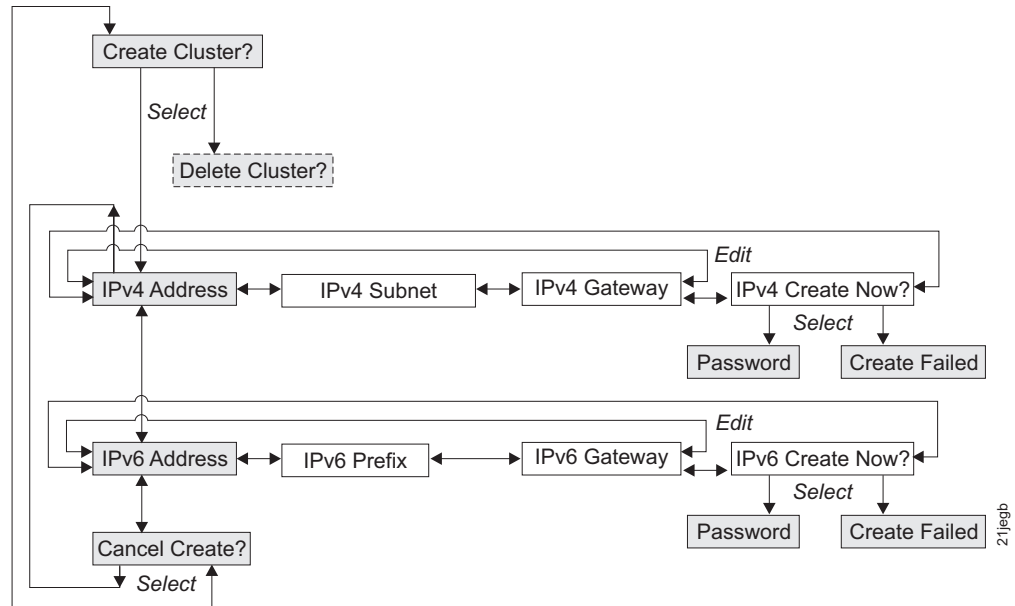


Figure 64. Create Cluster? navigation

Press the up and down buttons to navigate through the secondary options that are associated with the Create Cluster? option. When you have navigated to the desired option, press the select button.

The secondary options that are available include:

- IPv4 Address
- IPv6 Address
- Cancel Create?

If you are creating the cluster with an IPv4 address, complete the following steps:

1. Press and release the up or down button until Node: is displayed
2. Press and release the left or right button until Create Cluster? is displayed.
3. Press and release the select button. The IPv4 Address panel is displayed.
4. Put the panel into edit mode by pressing and releasing the select button.
5. Edit the IPv4 address, the IPv4 subnet, and the IPv4 gateway.
6. Press and release the select button to put the data in view rather than edit mode.
7. Use the right button to navigate to the IPv4 Create Now? panel.
8. Press and release the select button to confirm.

If you are creating the cluster with an IPv6 address, complete the following steps:



1. Press and release the up or down button until Node: is displayed
2. Press and release the left or right button until Create Cluster? is displayed.
3. Press and release the select button and then press the down button. The IPv6 Address panel is displayed.
4. Put the panel into edit mode by pressing and releasing the select button.
5. Edit the IPv6 address, the IPv6 prefix, and the IPv6 gateway.
6. Press and release the select button to put the data in view rather than edit mode.
7. Use the right button to navigate to the IPv6 Create Now? panel.
8. Press and release the select button to confirm.

### IPv4 Address

Using the IPv4 address, you can set the IP address for Ethernet port 1 of the cluster that you are going to create. The cluster can have either an IPv4 or an IPv6 address, or both at the same time. You can set either the IPv4 or IPv6 cluster address for Ethernet port 1 from the front panel when you are creating the cluster. If required, you can add more management IP addresses from the CLI.

**Attention:** When you set the IPv4 address, ensure that you type the correct address. Otherwise, you might not be able to access the cluster using the command-line tools or the SAN Volume Controller Console.

Perform the following steps to set the IPv4 address:

1. Navigate to the IPv4 Address panel.
2. Press the select button. The first IP address number is highlighted.
3. Press the up button if you want to increase the value that is highlighted; press the down button if you want to decrease that value. If you want to quickly increase the highlighted value, hold the up button. If you want to quickly decrease the highlighted value, hold the down button.

**Note:** If you want to disable the fast increase/decrease function, press and hold the down button, press and release the select button, and then release the down button. The disabling of the fast increase/decrease function lasts until cluster creation is completed or until the feature is again enabled. If the up button or down button is pressed and held while the function is disabled, the value increases or decreases once every two seconds. To again enable the fast increase/decrease function, press and hold the up button, press and release the select button, and then release the up button.

4. Press the right button or left button to move to the number field that you want to set.
5. Repeat steps 3 and 4 for each number field that you want to set.
6. Press the select button to confirm the settings. Otherwise, press the right button to display the next secondary option or press the left button to display the previous options.

Press the right button to display the next secondary option or press the left button to display the previous options.

## IPv4 Subnet

Using this option, you can set the IPv4 subnet mask for Ethernet port 1.

**Attention:** When you set the IPv4 subnet mask address, ensure that you type the correct address. Otherwise, you might not be able to access the cluster using the command-line tools or the SAN Volume Controller Console.

Perform the following steps to set the subnet mask:

1. Navigate to the IPv4 Subnet panel.
2. Press the select button. The first subnet mask number is highlighted.
3. Press the up button if you want to increase the value that is highlighted; press the down button if you want to decrease that value. If you want to quickly increase the highlighted value, hold the up button. If you want to quickly decrease the highlighted value, hold the down button.

**Note:** If you want to disable the fast increase/decrease function, press and hold the down button, press and release the select button, and then release the down button. The disabling of the fast increase/decrease function lasts until cluster creation is completed or until the feature is again enabled. If the up button or down button is pressed and held while the function is disabled, the value increases or decreases once every two seconds. To again enable the fast increase/decrease function, press and hold the up button, press and release the select button, and then release the up button.

4. Press the right button or left button to move to the number field that you want to set.
5. Repeat steps 3 and 4 for each number field that you want to set.
6. Press the select button to confirm the settings. Otherwise, press the right button to display the next secondary option or press the left button to display the previous options.

## IPv4 Gateway

Using this option, you can set the IPv4 gateway address for Ethernet port 1.

**Attention:** When you set the IPv4 gateway address, ensure that you type the correct address. Otherwise, you might not be able to access the cluster using the command-line tools or the SAN Volume Controller Console.

Perform the following steps to set the IPv4 gateway address:

1. Navigate to the IPv4 Gateway panel.
2. Press the select button. The first gateway address number field is highlighted.
3. Press the up button if you want to increase the value that is highlighted; press the down button if you want to decrease that value. If you want to quickly increase the highlighted value, hold the up button. If you want to quickly decrease the highlighted value, hold the down button.

**Note:** If you want to disable the fast increase/decrease function, press and hold the down button, press and release the select button, and then release the down button. The disabling of the fast increase/decrease function lasts until cluster creation is completed or until the feature is again enabled. If the up button or down button is pressed and held while the function is disabled, the value increases or decreases once every two

seconds. To again enable the fast increase/decrease function, press and hold the up button, press and release the select button, and then release the up button.

4. Press the right button or left button to move to the number field that you want to set.
5. Repeat steps 3 and 4 for each number field that you want to set.
6. Press the select button to confirm the settings. Otherwise, press the right button to display the next secondary option or press the left button to display the previous options.

### IPv4 Create Now?

Using this option, you can start an operation to create a cluster with an IPv4 address.

1. Navigate to the IPv4Create Now? panel.
2. Press the select button to start the operation.

If the create operation is successful, Password is displayed on line 1. The password that you can use to access the cluster is displayed on line 2. Be sure to immediately record the password; it is required on the first attempt to manage the cluster from the SAN Volume Controller Console.

**Attention:** The password displays for only 60 seconds, or until a front panel button is pressed. The cluster is created only after the password display is cleared.

If the create operation fails, Create Failed: is displayed on line 1 of the front-panel display screen. Line 2 displays one of two possible error codes that you can use to isolate the cause of the failure.

### IPv6 Address

Using this option, you can set the IPv6 address for Ethernet port 1 of the cluster that you are going to create. The cluster can have either an IPv4 or an IPv6 address, or both at the same time. You can set either the IPv4 or IPv6 cluster address for Ethernet port 1 from the front panel when you are creating the cluster. If required, you can add more management IP addresses from the CLI.

**Attention:** When you set the IPv6 address, ensure that you type the correct address. Otherwise, you might not be able to access the cluster using the command-line tools or the SAN Volume Controller Console.

Perform the following steps to set the IPv6 address:

1. From the Create Cluster? option, press the select button, and then press the down button. The IPv6 Address option is displayed.
2. Press the select button again. The first IPv6 address number is highlighted. .
3. Move between the address panels by using the left button or right button. The IPv6 addresses and the IPv6 gateway addresses consist of eight (4-digit) hexadecimal values that are shown across four panels
4. You can change each number in the address independently. Press the up button if you want to increase the value that is highlighted; press the down button if you want to decrease that value.
5. Press the right button or left button to move to the number field that you want to set.
6. Repeat steps 3 and 4 for each number field that you want to set.

7. Press the select button to confirm the settings. Otherwise, press the right button to display the next secondary option or press the left button to display the previous options.

### IPv6 Prefix

Using this option, you can set the IPv6 prefix for Ethernet port 1.

**Attention:** When you set the IPv6 prefix, ensure that you type the correct network prefix. Otherwise, you might not be able to access the cluster using the command-line tools or the SAN Volume Controller Console.

Perform the following steps to set the IPv6 prefix:

**Note:** This option is restricted to a value between 0 and 127.

1. Navigate to the IPv6 Prefix panel.
2. Press the select button. The first prefix number field is highlighted.
3. Press the up button if you want to increase the value that is highlighted; press the down button if you want to decrease that value. If you want to quickly increase the highlighted value, hold the up button. If you want to quickly decrease the highlighted value, hold the down button.

**Note:** If you want to disable the fast increase/decrease function, press and hold the down button, press and release the select button, and then release the down button. The disabling of the fast increase/decrease function lasts until cluster creation is completed or until the feature is again enabled. If the up button or down button is pressed and held while the function is disabled, the value increases or decreases once every two seconds. To again enable the fast increase/decrease function, press and hold the up button, press and release the select button, and then release the up button.

4. Press the select button to confirm the settings. Otherwise, press the right button to display the next secondary option or press the left button to display the previous options.

### IPv6 Gateway

Using this option, you can set the IPv6 gateway for Ethernet port 1.

**Attention:** When you set the IPv6 gateway address, ensure that you type the correct address. Otherwise, you might not be able to access the cluster using the command-line tools or the SAN Volume Controller Console.

Perform the following steps to set the IPv6 gateway address:

1. Navigate to the IPv6 Gateway panel.
2. Press the select button. The first gateway address number is highlighted. The IPv6 addresses and the IPv6 gateway addresses consist of eight (4-digit) hexadecimal values that are shown across four panels.
3. You can change each number in the address independently. Press the up button if you want to increase the value that is highlighted; press the down button if you want to decrease that value.
4. Press the right button or left button to move to the number field that you want to set.
5. Repeat steps 3 and 4 for each number field that you want to set.

6. Press the select button to confirm the settings. Otherwise, press the right button to display the next secondary option or press the left button to display the previous options.

### IPv6 Create Now?

Using this option, you can start an operation to create a cluster with an IPv6 address.

1. Navigate to the IPv6 Create Now? panel.
2. Press the select button to start the operation.

If the create operation is successful, Password is displayed on line 1. The password that you can use to access the cluster is displayed on line 2. Be sure to immediately record the password; it is required on the first attempt to manage the cluster from the SAN Volume Controller Console.

**Attention:** The password displays for only 60 seconds, or until a front panel button is pressed. The cluster is created only after the password display is cleared.

If the create operation fails, Create Failed: is displayed on line 1 of the front-panel display screen. Line 2 displays one of two possible error codes that you can use to isolate the cause of the failure.

### Delete Cluster?

The Delete Cluster? option deletes the cluster state data from the node. This option is displayed only if you select the Create Cluster? option on a SAN Volume Controller node that is already a member of a cluster.

Use this option as the final step in decommissioning a cluster after the other nodes have been removed from the cluster using the command-line interface (CLI) or the graphical user interface (GUI).

**Attention:** Use the front panel to remove cluster state data from a single node cluster. To remove a node from a multi-node cluster, always use the CLI or the remove node options from the GUI.

From the Delete Cluster? panel, perform the following steps to delete the cluster state data from the node:

1. Press and hold the up button.
2. Press and release the select button.
3. Release the up button.

After the option is run, the node shows "Cluster:" with no cluster name. If this option is performed on a node that is still a member of a cluster, the cluster will show error 1195, "Node missing," and the node will still appear in the list of nodes in the cluster. The node should be removed by using the GUI or CLI.

Use the up button to return to the Create Cluster? option.

### Related concepts

“Displaying and editing an IPv6 address” on page 148

After you have set the IPv6 address, you can display the IPv6 addresses and the IPv6 gateway addresses. You can also set the IPv6 addresses for a cluster that you are going to create.

### Related information

“Understanding the create cluster error codes” on page 238

Cluster Create error codes are displayed on the menu screen when you are using the front panel to create a new cluster, but the create operation fails.

## Version option

The version option displays the version of the SAN Volume Controller software that is active on the node. The version consists of four fields that are separated by full stops. The fields are the version, release, modification, and fix level; for example, 4.3.2.1.

### Build

The Build: panel displays the level of the SAN Volume Controller software that is currently active on this node.

### Cluster Build

The Cluster Build: panel displays the level of the SAN Volume Controller software that is currently active on the cluster that this node is operating in.

## Ethernet option

The Ethernet option displays the operational state of the Ethernet ports, the speed and duplex information, and their media access control (MAC) addresses.

The Ethernet panel shows one of the following states:

#### Config - Yes

This node is the configuration node.

#### Config - No

This node is not the configuration node.

#### No Cluster

This node is not a member of a cluster.

Press the right button to view the details of the individual Ethernet ports.

### Ethernet Port-1 and Ethernet Port-2

The Ethernet port option displays the state of the links and indicates whether or not there is an active link with an Ethernet network.

#### Link Online

An Ethernet cable is attached to this port.

#### Link Offline

No Ethernet cable is attached to this port or the link has failed.

## Speed-1 or Speed-2

The speed option displays the speed and duplex information for the Ethernet port. The speed information can be one of the following values:

**\*10** The speed is 10 MBps.

**\*100** The speed is 100 MBps.

**\*1000** The speed is 1000 MBps.

The duplex information can be one of the following values:

**\*Full** Data can be sent and received at the same time.

**\*Half** Data can be sent and received in one direction at a time.

## MAC Address-1 or MAC Address-2

The MAC address option displays the media access control (MAC) address of the Ethernet port.

## Fibre-channel port-1 through 4 option

The fibre-channel port-1 through 4 options display the operational status of the fibre-channel ports.

**Active** The port is operational and can access the fibre-channel fabric.

### Inactive

The port is operational but cannot access the fibre-channel fabric. One of the following conditions caused this result:

- The fibre-channel cable has failed
- The fibre-channel cable is not installed
- The device that is at the other end of the cable has failed

**Failed** The port is not operational because of a hardware failure.

### Not installed

This port is not installed.

To display the current fibre-channel port speed, press and hold the down button, press and release the select button, and release the down button. For the SAN Volume Controller 2145-8F2, you can also use this action to change the fibre-channel port speed of a node that is not participating in a cluster, by pressing the up or down buttons.

## Select language? option

You can change the language that displays on the front panel.

The Select language? option allows you to change the language that is displayed on the menu. Figure 65 on page 160 shows the Select language? option sequence.




Figure 65. Select Language? navigation

The following languages are available:

- English
- Japanese

To select the language that you want to be used on the front panel, perform the following steps:

1. Navigate to the Select language? panel.
2. Press and release the select button.
3. Use the left and right buttons to move to the desired language. The translated language names are displayed in their own character set. If you do not understand the language that is displayed, wait for at least 60 seconds for the menu to reset to the default option.
4. Press and release the select button to select the language that is displayed.

If the selected language uses the Latin alphabet, the front panel display shows two lines. The panel text is displayed on the first line and additional data is displayed on the second line.

If the selected language does not use the Latin alphabet, the display shows only one line at a time to clearly display the character font. For those languages, you can switch between the panel text and the additional data by pressing and releasing the select button.

Additional data is not available when the front panel displays a menu option, which ends with a question mark (?). In this case, press and release the select button to choose the menu option.

**Note:** You cannot select another language when the node is displaying a boot error.

---

## Service mode

While in service mode, you can use the front panel to view or change a service IPv4 or an IPv6 address. You can also view the version and build level of the SAN Volume Controller software that is installed and active on the node.

Enter service mode only if directed to do so by the IBM Support Center.

By default, all nodes in a cluster are configured to use the same service IP addresses by using the SAN Volume Controller Console or the `svctask chclusterip` CLI command. This means that you can place only one node in service mode at a time. However, to make it possible to have more than one node in service mode at a time, you have the following options:

- Modify the service mode IP addresses from the front panel after the node is in service mode.



- Configure the nodes to use DHCP to allocate service mode IP addresses.

**Note:** Any changes made to the service addresses from the service mode front panel menus affect only the one node. The changes apply only to this service mode session.

To access a node that is in service mode, point your Web browser to the following Web address, where *serviceipaddress* is one of the IPv4 or IPv6 addresses that is shown on the front panel display:

`https://serviceipaddress`

The service mode panel is displayed when you enter service mode. You can navigate to the IPv4 Address, IPv6 Address, or the Version panels, as shown in Figure 66 on page 162.

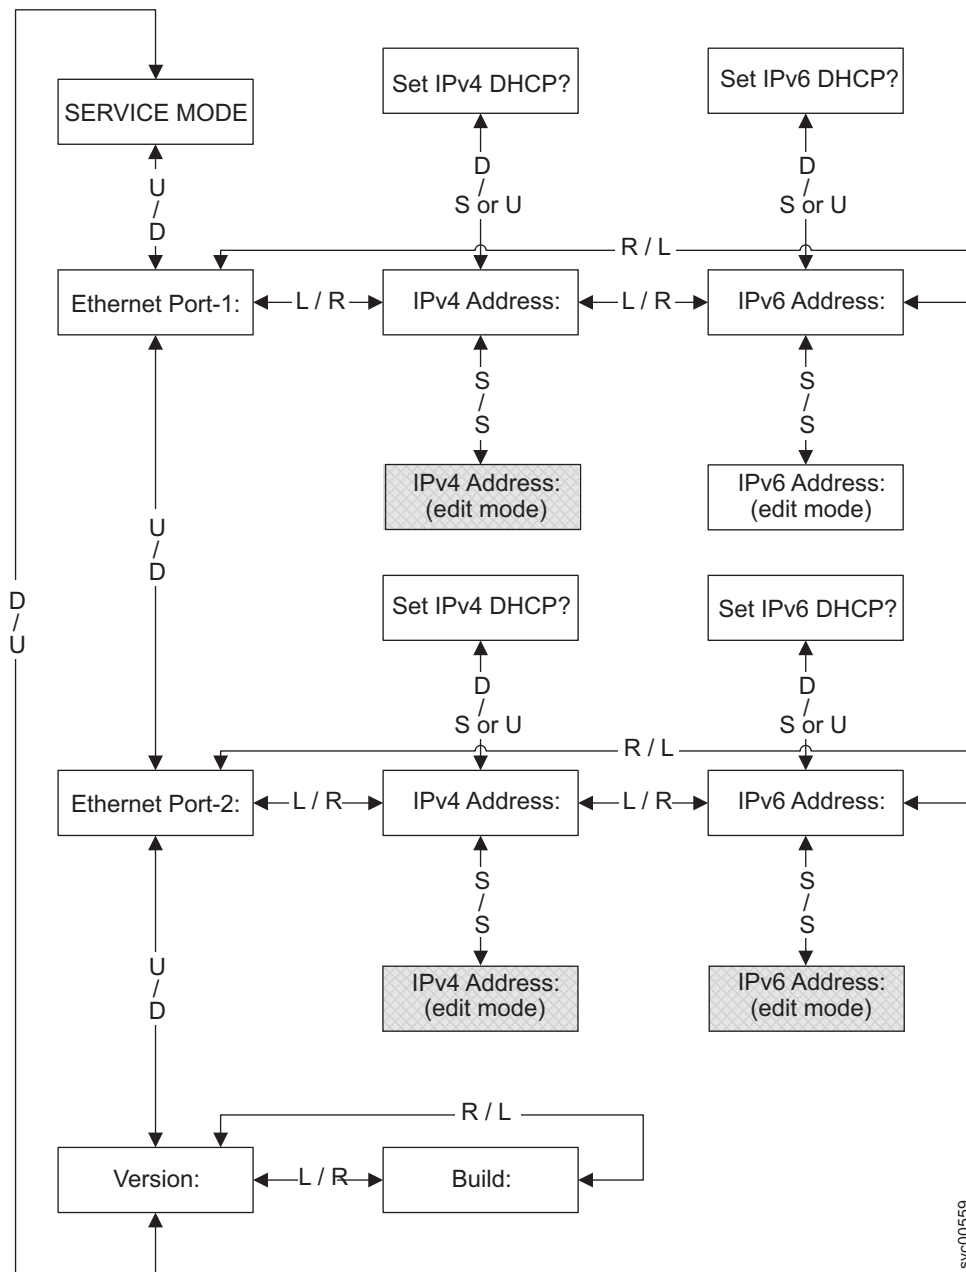


Figure 66. Service mode navigation

You can exit service mode through the Web browser, by using the CLI, or by turning the node off and then on.

### IPv4 Address

The IPv4 Address panels show one of the following items for the selected Ethernet port:

- The active service mode address if the cluster has an IPv4 address. This can be either a configured or fixed address, or it can be an address obtained through DHCP.
- DHCP Failed if the IPv4 service address is configured for DHCP but the node was unable to obtain an IP address.

- DHCP Configuring if the IPv4 service address is configured for DHCP while the node attempts to obtain an IP address. This will change to the IPv4 address automatically if a DHCP address is allocated and activated.
- A blank line if the cluster does not have an IPv4 address.

If the service IPv4 address was not set correctly or a DHCP address was not allocated, you have the option of correcting the IPv4 address from this panel. Note that the service IP address must be in the same subnet as the cluster IP address.

To set a fixed service IPv4 address from the IPv4 Address: panel, perform the following steps:

1. Press and release the select button to put the panel in edit mode.
2. Press the right button or left button to move to the number field that you want to set.
3. Press the up button if you want to increase the value that is highlighted; press the down button if you want to decrease that value. If you want to quickly increase the highlighted value, hold the up button. If you want to quickly decrease the highlighted value, hold the down button.

**Note:** If you want to disable the fast increase/decrease function, press and hold the down button, press and release the select button, and then release the down button. The disabling of the fast increase/decrease function lasts until cluster creation is completed or until the feature is again enabled. If the up button or down button is pressed and held while the function is disabled, the value increases or decreases once every two seconds. To again enable the fast increase/decrease function, press and hold the up button, press and release the select button, and then release the up button.

4. When all the fields are set as required, press and release the select button to activate the new IPv4 address.

The IPv4 Address: panel is displayed. The new service IPv4 address is not displayed until it has become active. If the new address has not been displayed after two minutes, check that the selected address is valid on the subnetwork and that the Ethernet switch is working correctly.

To set the service IPv4 address to use DHCP, perform the following steps:

1. Navigate to the IPv4 Address: panel.
2. Press and release the down button. Set IPv4 DHCP? is displayed on the front panel.
3. Press and release the select button to activate DHCP, or you can press and release the up button to keep the existing address.
4. If you activate DHCP, DHCP Configuring is displayed while the node attempts to obtain a DHCP address. It changes automatically to show the allocated address if a DHCP address is allocated and activated, or it changes to DHCP Failed if a DHCP address is not allocated.

You can exit service mode through the CLI, or by turning the node off and then on.

## IPv6 Address

The IPv6 Address panels show one of the following for the selected Ethernet port:

- The active service mode address if the cluster has an IPv6 address. This can be either a configured or fixed address, or it can be an address obtained through DHCP.
- DHCP Failed if the IPv6 service address is configured for DHCP but the node was unable to obtain an IP address.
- DHCP Configuring if the IPv6 service address is configured for DHCP while the node attempts to obtain an IP address. This changes to the IPv6 address automatically if a DHCP address is allocated and activated.
- A blank line if the cluster does not have an IPv6 address.

If the service IPv6 address was not set correctly or a DHCP address was not allocated, you have the option of correcting the IPv6 address from this panel. Note that the service IP address must be in the same subnet as the cluster IP address.

To set a fixed service IPv6 address from the IPv6 Address: panel, perform the following steps:

1. Press and release the select button to put the panel in edit mode. When the panel is in edit mode, the full address is still shown across four panels as eight (4-digit) hexadecimal values. You edit each digit of the hexadecimal values independently. The current digit is highlighted.
2. Press the right button or left button to move to the number field that you want to set.
3. Press the up button if you want to increase the value that is highlighted; press the down button if you want to decrease that value.
4. When all the fields are set as required, press and release the select button to activate the new IPv6 address.

The IPv6 Address: panel is displayed. The new service IPv6 address is not displayed until it has become active. If the new address has not been displayed after two minutes, check that the selected address is valid on the subnetwork and that the Ethernet switch is working correctly.

To set the service IPv6 address to use DHCP, perform the following steps:

1. Navigate to the IPv6 Address: panel.
2. Press and release the down button. Set IPv6 DHCP? is displayed on the front panel.
3. Press and release the select button to activate DHCP, or you can press and release the up button to keep the existing address.
4. If you activate DHCP, DHCP Configuring is displayed while the node attempts to obtain a DHCP address. It changes automatically to show the allocated address if a DHCP address is allocated and activated, or it changes to DHCP Failed if a DHCP address is not allocated.

**Note:** If an IPv6 router is present on the local network, SAN Volume Controller does not differentiate between an autoconfigured address and a DHCP address. Therefore, SAN Volume Controller uses the first address that is detected.

## Version

The version option displays the version of the SAN Volume Controller software that is active on the node. The version consists of four fields that are separated by full stops. The fields are the version, release, modification, and fix level; for example, 4.3.2.1.

## Build

The Build: panel displays the level of the SAN Volume Controller software that is currently active on this node.

### Related concepts

“Displaying and editing an IPv6 address” on page 148

After you have set the IPv6 address, you can display the IPv6 addresses and the IPv6 gateway addresses. You can also set the IPv6 addresses for a cluster that you are going to create.

### Related reference

“Service mode overview” on page 8

The SAN Volume Controller service mode allows you to access vital product data (VPD), logs, and dump data on the node. It also provides a method of forcing the installation of a different version of software.

---

## Using the power control for the SAN Volume Controller node

SAN Volume Controller nodes are powered by an uninterruptible power supply that is located in the same rack as the SAN Volume Controller nodes.

The power state of the SAN Volume Controller is displayed by a power indicator on the front panel. If the uninterruptible power supply battery is not sufficiently charged to enable the SAN Volume Controller to become fully operational, its charge state is displayed on the front panel display of the SAN Volume Controller node.

The power to a SAN Volume Controller is controlled by the power button on the front panel of the SAN Volume Controller node or by commands sent by the Ethernet interface. *Never* turn off the SAN Volume Controller node by removing the power cable. You might lose data. For more information about how to power off the SAN Volume Controller node, see “MAP 5350: Powering off a SAN Volume Controller node” on page 385.

If the SAN Volume Controller software is running and you request it to power off from the SAN Volume Controller Console, CLI, or power button, the SAN Volume Controller node starts its power off processing. During this time, the SAN Volume Controller node indicates the progress of the power-off operation on the front panel display. After the power-off processing is complete, the front panel becomes blank and the front panel power light flashes. It is safe for you to remove the power cable from the rear of the SAN Volume Controller node. If the power button on the front panel is pressed during power-off processing, the front panel display changes to indicate that the SAN Volume Controller node is being restarted, but the power-off process completes before the restart is performed.

If the SAN Volume Controller software is not running when the front panel power button is pressed, the SAN Volume Controller node immediately powers off.

If you turn off a SAN Volume Controller node using the power button or by a command, the SAN Volume Controller node is put into a power-off state. The SAN Volume Controller remains in this state until the power cable is connected to the rear of the SAN Volume Controller node and the power button is pressed.

During the SAN Volume Controller startup sequence, the SAN Volume Controller tries to detect the status of the uninterruptible power supply through the uninterruptible power supply signal cable. If an uninterruptible power supply is not detected, the SAN Volume Controller node pauses and an error is shown on the

front panel display. If the uninterruptible power supply is detected, the software monitors the operational state of the uninterruptible power supply. If no uninterruptible power supply errors are reported and the uninterruptible power supply battery is sufficiently charged, the SAN Volume Controller becomes operational. If the uninterruptible power supply battery is not sufficiently charged, the charge state is indicated by a progress bar on the front panel display. When an uninterruptible power supply is first turned on, it might take up to two hours before the battery is sufficiently charged for the SAN Volume Controller node to become operational.

If input power to the uninterruptible power supply is lost, the SAN Volume Controller node immediately stops all I/O operations and saves the contents of its dynamic random access memory (DRAM) to the internal disk drive. While data is being saved to the disk drive, a Power Failure message is shown on the front panel and is accompanied by a descending progress bar that indicates the quantity of data that remains to be saved. After all the data is saved, the SAN Volume Controller node is turned off and the power light on the front panel turns off.

**Note:** The SAN Volume Controller node is now in standby state. If the input power to the uninterruptible power supply unit is restored, the SAN Volume Controller node restarts. If the uninterruptible power supply battery was fully discharged, Charging is displayed and the boot process waits for the battery to charge. When the battery is sufficiently charged, Booting is displayed, the node is tested, and the software is loaded. When the boot process is complete, Recovering is displayed while the uninterruptible power supply finalizes its charge. While Recovering is displayed, the cluster can function normally. However, when the power is restored after a second power failure, there is a delay (with Charging displayed) before the node can complete its boot process.

#### **Related concepts**

“Powering off” on page 142

The progress bar on the display shows the progress of the power-off operation.

---

## Chapter 6. Diagnosing problems

You can diagnose problems with SAN Volume Controller, the uninterruptible power supply, and the IBM System Storage Productivity Center, or the master console server using either the command-line interface (CLI) or the SAN Volume Controller Console. The diagnostic LEDs on the SAN Volume Controller nodes and uninterruptible power supply units also help you diagnose hardware problems.

### Error event logs

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

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

### Error codes

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

- Error reporting
- Understanding the error event log
- Understanding the error codes
- Understanding the cluster error codes
- Determining a hardware boot failure
- Understanding the boot error codes
- Performing the node rescue
- Understanding the node rescue error codes
- Understanding the create cluster error codes
- Marking errors as fixed
- Checking the status of a node

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 SAN Volume Controller node to see if 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 SAN Volume Controller node” on page 385.
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.

### Related tasks

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

### Related information

“Websphere and CIM Logging” on page 354

You can obtain log files for the Websphere Application Server and the Common Information Model (CIM). Instructions are supplied in the following topics.

---

## Error reporting

Errors detected by the SAN Volume Controller are saved in an error log. As soon as an entry is made in this error log, the error condition is analyzed. If any service activity is required, the user is notified of the error.

### Error reporting process

The following methods are used to notify the user and the IBM Support Center of a new error:

- The most serious cluster code is displayed on the front panel of each node in the cluster.
- If you enabled simple network management protocol (SNMP), an SNMP trap is sent to an SNMP manager that is configured by the customer.  
The SNMP manager might be IBM Director, if it is installed, or another SNMP manager.
- If enabled, log messages can be forwarded from a sender to a receiver on an IP network by using the syslog protocol.
- If enabled, errors and other event notifications can be forwarded from a sender to a receiver through Call Home e-mail.
- If you enabled Call Home, critical faults are reported directly to the IBM Support Center by e-mail.

### Related tasks

“Using directed maintenance procedures” on page 123

You can use directed maintenance procedures (DMP) to diagnose and resolve problems with the SAN Volume Controller.

### Related information

“Understanding cluster error codes” on page 239

Every cluster error code includes an error code number, a description, action, and possible field-replaceable units (FRUs).

## Power-on self-test

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

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

If a critical failure is detected during the POST, the SAN Volume Controller software is not loaded and the system error LED on the operator information panel is illuminated. If this occurs, use “MAP 5000: Start” on page 357 to help isolate the cause of the failure.



When the SAN Volume Controller software is loaded, additional testing takes place, which ensures that all of the required hardware and software components are installed and functioning correctly. During the additional testing, the word `Booting` is displayed on the front panel along with a boot progress code and a progress bar.

If a test failure occurs, the word `Failed` is displayed on the front panel along with a boot progress code. Some parts of the boot process take several minutes to complete. If the progress bar fails to move and the boot code number does not change for several minutes, see the boot progress code section to understand the situation.

The service controller performs internal checks and is vital to the operation of the SAN Volume Controller. If the error (check) LED is illuminated on the service controller front panel, the front-panel display might not be functioning correctly and you can ignore any message displayed.

The uninterruptible power supply also performs internal tests. If the uninterruptible power supply is capable of reporting the failure condition, the SAN Volume Controller displays critical failure information on the front-panel display or sends noncritical failure information to the cluster error log. If the SAN Volume Controller cannot communicate with the uninterruptible power supply, it displays a boot failure error message on the front-panel display. Further problem determination information might also be displayed on the front panel of the uninterruptible power supply.

---

## Understanding the error event log

The SAN Volume Controller error event log contains details of error events and informational events that have occurred on the system. You also receive log data for the commands that have been run.

### Error data

Error data consists of two types of events:

- An error event is logged when a failure condition is detected. An error event has an associated error code. Each error code has a service action that must be performed as soon as possible. The service actions are automated through the Directed Maintenance Procedures.
- An informational event is logged when a significant change of state occurs, for instance, a FlashCopy completes. Whether or not any action is required when an informational event is logged depends on the environment in which the SAN Volume Controller is deployed in. In some environments, some informational events are considered a condition that must be rectified, for example, the space-efficient VDisk space warning.

### Configuration data

Configuration data is logged when a configuration event has occurred.

### Audit log

An audit log records configuration command actions that are issued through a Secure Shell (SSH) session or through the SAN Volume Controller Console. The log entries can provide information such as the user who issued the action command, the name of the command, and a time stamp of when the command was issued on the configuration node. The audit log records only successful changes.

### Related tasks

Chapter 7, “Using the maintenance analysis procedures,” on page 357  
The maintenance analysis procedures (MAPs) inform you how to analyze a failure that occurs with a SAN Volume Controller node.

## Managing the error event log

The error event log has a limited size. After it is full, newer entries replace the oldest entries. If the old entry has not been fixed, it is not replaced by newer entries.

To avoid the possibility of an error condition causing the log to be flooded by a single error, some errors of the same type are recorded in the same space in the error event log. When error-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.

### Related tasks

“MAP 5000: Start” on page 357  
MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

## Viewing the error event log

You can view the error event log by using the SAN Volume Controller command-line interface (CLI) or the SAN Volume Controller Console.

Perform the following steps to view the full contents of each error-event log entry using the CLI:

1. Dump the contents of the error log to a file.
  - a. Issue the `svctask dumperrlog` command to create a dump file that contains the current error-event log data.
  - b. Issue the `svcinfo lserrlogdumps` command to determine the name of the dump file that you have just created.
  - c. Issue the secure copy command (`scp`) to copy the dump file to the IBM System Storage Productivity Center.
2. View the file with a text viewer after the dump file is extracted from the cluster.

Figure 67 on page 171 shows an example of an error-event log entry that might be displayed:

```

Error Log Entry 21
Node Identifier      : node3
Object Type         : adaptor
Object ID           : 1
Sequence Number     : 174
Root Sequence Number : 174
First Error Timestamp : Tue Aug 23 16:02:18 2005
                   : Epoch + 1051027338
Last Error Timestamp  : Tue Aug 23 16:02:18 2005
                   : Epoch + 1051027338
Error Count          : 1
Error ID             : 73003 : More/Less fibre channel ports operational
Error Code           : 1060 : Fibre Channel ports not operational
Status Flag          : UNFIXED
Type Flag            : ERROR CAT 1

02 02 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00


```

Figure 67. Example of an error-event log entry

You can also view the error event log using the SAN Volume Controller Console. The error event log contains a large number of entries, but by using this method of viewing the log you can select only the type of information that you need. For example, if you are repairing a fault, you might only want to select **Show unfixed errors**.

Figure 68 on page 172 shows an example of an error-event log summary that is displayed when you select the type of information that you want.

Whole error log (sorted by error priority)

Node Identifier	Object Type	Object ID	Sequence Number	Time Stamp	Error Code	Type	Message
Node2	adapter	1	661	03-Jul-2003 11:48:50	1060	 error	Fibre Channel ports not operational
Node2	device	1	662	03-Jul-2003 11:48:05	1630	 error	Number of device logins reduced
Node2	cluster	0	667	03-Jul-2003 11:53:50	981001	 info	Cluster Fabric View updated by fabric discovery
Node2	cluster	0	666	03-Jul-2003 11:53:43	981001	 info	Cluster Fabric View updated by fabric discovery
Node1	cluster	0	665	03-Jul-2003 11:52:35	981001	 info	Cluster Fabric View updated by fabric discovery
Node2	cluster	0	664	03-Jul-2003 11:52:23	981001	 info	Cluster Fabric View updated by fabric discovery
Node2	cluster	0	663	03-Jul-2003 11:50:35	981001	 info	Cluster Fabric View updated by fabric discovery
Node1	cluster	0	662	03-Jul-2003 11:50:35	981001	 info	Cluster Fabric View updated by fabric discovery
Node2	cluster	0	660	03-Jul-2003 11:48:05	981001	 info	Cluster Fabric View updated by fabric discovery
Node2	cluster	0	658	03-Jul-2003 11:47:50	981001	 info	Cluster Fabric View updated by fabric discovery

Page 1 of 3        Total = 27    Filtered = 27    Displaying records 1 to 10

svc00567

Figure 68. Example of an error-event log summary

Details of each listed error can be displayed by clicking on the sequence number of any record. The **Type** field contains an icon and a text message to indicate the cause of the log entry. Table 37 describes the meaning of the information in the type field.

Table 37. Descriptions of log entry icons





Icon	Description
	The error icon indicates that this log entry requires service activity. Select <b>Run Maintenance Procedures</b> from the <b>Service and Maintenance</b> menu to start the repair activity.
	The fixed icon indicates that a problem existed, but it has now been resolved. It might have been resolved as a result of service activity, or it might have been resolved as a result of some other action, for example, powering on a missing node.

Table 37. Descriptions of log entry icons (continued)

Icon	Description
	The warn icon indicates that some condition has occurred that might have been caused by a temporary problem or by a problem that is external to the SAN Volume Controller, such as an error in a RAID controller. If a specified number of these events occurs in 25 hours, the warning converts to an error. No service action is required on the SAN Volume Controller for this log entry.
	The info icon indicates that the log entry provides information about a configuration change or the state of a command. In some cases, the SAN Volume Controller user might need to take some action based on this information.

## Describing the fields in the error event log

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

Table 38 describes the fields that you see when you use the command-line interface to produce an error event log.

Table 38. Description of data fields for the error event log

Data field	Description
Node identifier	The name of the node that the error or event occurred on. If the event applies to a specific node, for example, it indicates that the fibre-channel port of the node is not working, then the node identifier is indicated here. If the event applies to the cluster rather than to a node, for example, a license limit is exceeded, then this field is set to 0.
Object type	The object type to which the error event log relates. See Table 39 on page 174.
Object ID	A number that uniquely identifies the instance of the object.
Sequence number	A number that identifies the event. Where appropriate, sense data is returned to host systems references the sequence number.
Root sequence number	The sequence number of another log entry that enables all errors that have a single source to be marked as fixed by a single action.
First error timestamp	The time when this error event was reported. If events of a similar type are being coalesced together, so one error event log record represents more than one event, this is the time the first error event was logged.
Last error timestamp	The time when the last instance of this error event was recorded in the log.
Error count	The number of error events coalesced into this error event log record.
Error ID	This number is a unique identifier for the error or event.
Error code	This number is used as the starting point for service procedures.
Status flag	For details on the status flag, see Table 40 on page 174.
Type flag	For details on the type flag, see Table 42 on page 175.

Table 38. Description of data fields for the error event log (continued)

Data field	Description
Additional sense data	Data specific to this error or event code. This is a binary data record. When the error event log is viewed using the command line tools, this data is shown in hex. When the data is viewed using the Web interface, this data is translated to ASCII characters on the right side of the page. You are not normally required to interpret this data as part of the service procedures. However, any procedures that do refer to the data describe the ASCII format.

Table 39 describes the types of error-event log objects.

Table 39. Description of object types and object IDs for the error event log

Object type	Object ID
mdisk	Managed disk number
mdiskgrp	Managed disk group number
vdisk	Virtual disk
node	Node ID
host	Host identifier
iogroup	I/O group identifier
fcgrp	FlashCopy consistency group number
rcgrp	Metro or Global Mirror consistency group number
fcmap	FlashCopy mapping number
rcmap	Metro or Global Mirror mapping number
wwpn	Worldwide port name
cluster	Cluster ID (shown in decimal)
device	Device number
SCSI lun	SCSI logical unit number (LUN) identifier
quorum	Quorum identifier
clusterip	Cluster IP
fc_adapter	Fibre-channel adapter
Emailuser	E-mail user
Syslog	Syslog server
Snmp	SNMP server
Smtpt	E-mail server
Auth_Group	User group
SAS_Adapter	SAS adapter

Table 40 shows the types of error-event log flags.

**Note:** Informational events only have the SNMP trap-raised flag on when configured to do so.

Table 40. Description of flags for the error-event log

Flag	Description
Unfixed	This log entry requires a service action.

Table 40. Description of flags for the error-event log (continued)

Flag	Description
Fixed	This entry is marked as fixed. It remains in the error event log until it becomes the oldest record in the log, at which point it is overwritten by the next log entry.
Expired	Some events require a certain number of occurrences in 25 hours before they are shown as unfixed. If they do not reach this threshold in 25 hours, they are flagged as expired. Any further events of this type are then placed in a new event log entry.

Table 41 shows the various combinations of flags that might be logged and the resulting status that is reported by the user interfaces. The Analyze Error Log panel in the SAN Volume Controller Console summarizes the status flags in the error event log record. The following table lists the status values shown, and the status flag combination that value represents.

Table 41. Reported status for combinations of error-log status flags

Reported Status	UNFIXED	ERROR_FIXED	ERROR_EXPIRED
BELOW_THRESHOLD	0	0	0
EXPIRED	0	0	1
UNFIXED	1	0	0
FIXED	x <sup>1</sup>	1	0
<b>Note:</b> <sup>1</sup> This flag has no effect on the status.			

Table 42 shows the error-event types.

Table 42. Description error-event types

Type	Description
Error Cat 1	These errors require a service action. A Cat 1 error normally indicates a problem with a hardware component. A FRU or list of FRUs are included with the trap data sent with the error record.
Error Cat 2	These errors require a service action. A Cat 2 error normally indicates a configuration or environment problem rather than a hardware problem.
Transient error	Errors of type <i>transient</i> have been recovered by an error recovery procedure.
Configuration event	This entry is from the configuration event log. This flag is useful when displaying both logs in a seamless display as an aid to relating logged error conditions to configuration events.
Informational	This type indicates that the log entry is an informational event. Information events can be used to warn the user about an unexpected configuration result or prompt a user to initiate further configuration actions.

## Related tasks

“Viewing the error event log” on page 170

You can view the error event log by using the SAN Volume Controller command-line interface (CLI) or the SAN Volume Controller Console.

---

## Event notifications

SAN Volume Controller can use Simple Network Management Protocol (SNMP) traps, syslog messages, and Call Home e-mail to notify you and the IBM Support Center when significant events are detected. Any combination of these notification methods can be used simultaneously.

Each event that SAN Volume Controller detects is assigned a notification type of Error, Warning, or Information. You can configure SAN Volume Controller to send each type of notification to specific recipients.

Table 43 describes the types of event notifications.

*Table 43. SAN Volume Controller notification types*

Notification type	Description
Error	<p>An error notification is sent to indicate a problem that must be corrected as soon as possible.</p> <p>This notification indicates a serious problem with the SAN Volume Controller. 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 type. Error notifications can be configured to be sent as a Call Home e-mail to the IBM Support Center.</p>
Warning	<p>A warning notification is sent to indicate a problem or unexpected condition with the SAN Volume Controller. Always immediately investigate this type of notification to determine the effect that it might have on your operation, and make any necessary corrections.</p> <p>A warning notification does not require any replacement parts and therefore should not require IBM Support Center involvement. However, the event being reported might indicate a condition that could be fatal to your operating environment: such as, for example, a critical FlashCopy operation has failed.</p>
Information	<p>An informational notification is sent to indicate that an expected event has occurred. For example, a FlashCopy operation has completed. No remedial action is required when these notifications are sent.</p>

## SNMP traps

SNMP is a standard protocol for managing networks and exchanging messages. SAN Volume Controller can send SNMP messages that notify personnel about an event. You can use an SNMP manager to view the SNMP messages that SAN Volume Controller sends. You can use the SAN Volume Controller Console or the SAN Volume Controller command-line interface to configure and modify your SNMP settings.



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 SAN Volume Controller. This file can be used with SNMP messages from all versions of SAN Volume Controller software. More information about the MIB file for SNMP is available at the Support for SAN Volume Controller (2145) Web site:

[www.ibm.com/storage/support/2145](http://www.ibm.com/storage/support/2145)

Search for **SAN Volume Controller MIB**. Go to the downloads results to find **Management Information Base (MIB) file for SNMP**. Click this link to find download options. The name of this file is SVC\_MIB\_<release>.MIB such as in SVC\_MIB\_4.3.1.MIB.

## Syslog messages

The syslog protocol is a standard protocol for forwarding log messages from a sender to a receiver on an IP network. The IP network can be either IPv4 or IPv6. SAN Volume Controller can send syslog messages that notify personnel about an event. SAN Volume Controller can transmit syslog messages in either expanded or concise format. You can use a syslog manager to view the syslog messages that SAN Volume Controller sends. SAN Volume Controller uses the User Datagram Protocol (UDP) to transmit the syslog message. You can use the SAN Volume Controller Console or the SAN Volume Controller command-line interface to configure and modify your syslog settings.

Table 44 shows how SAN Volume Controller notification codes map to syslog security-level codes.

*Table 44. SAN Volume Controller notification codes and corresponding syslog level codes*

SAN Volume Controller notification code	Syslog level code	Description
SS_EID_UNKNOWN	Not mapped	
SS_EID_ERROR	LOG_ALERT	Error that needs immediate attention
SS_EID_WARNING	LOG_ERROR	Warning that needs attention
SS_EID_INFO	LOG_INFO	Informational messages
SS_EID_TEST	LOG_DEBUG	Test message

Table 45 shows how syslog facility codes map to SAN Volume Controller values of user-defined message origin identifiers.

*Table 45. Syslog facility codes and SAN Volume Controller values of user-defined message origin identifiers*

Syslog facility code	Syslog value	SAN Volume Controller value
LOG_LOCAL0	16	0
LOG_LOCAL1	17	1
LOG_LOCAL2	18	2
LOG_LOCAL3	19	3
LOG_LOCAL4	20	4
LOG_LOCAL5	21	5
LOG_LOCAL6	22	6

Table 45. Syslog facility codes and SAN Volume Controller values of user-defined message origin identifiers (continued)

Syslog facility code	Syslog value	SAN Volume Controller value
LOG_LOCAL7	23	7

## Call Home e-mail

The Call Home feature transmits operational and error-related data to you and IBM through a Simple Mail Transfer Protocol (SMTP) server connection in the form of an event notification e-mail. When configured, this function alerts IBM service personnel about hardware failures and potentially serious configuration or environmental issues.

To send e-mail, you must configure at least one SMTP server. You can specify as many as five additional SMTP servers for backup purposes. The SMTP server must accept the relaying of e-mail from the SAN Volume Controller cluster IP address. You can then use the SAN Volume Controller Console or the SAN Volume Controller command-line interface to configure the e-mail settings, including contact information and e-mail recipients. Set the reply address to a valid e-mail address. Send a test e-mail to check that all connections and infrastructure are set up correctly. You can disable the Call Home function at any time using the SAN Volume Controller Console or the SAN Volume Controller command-line interface.

---

## Call Home and inventory e-mail information

SAN Volume Controller can use Call Home e-mail and inventory information e-mail to provide data and event notifications to you and to the IBM Support Center.

### Call Home e-mail

Call Home support is initiated for the following reasons or types of data:

- **Event notifications:** Data is sent to the specified e-mail address when events occur. There are three types of events—error, warning, and information. All these types of events can cause an e-mail notification to be sent, depending on how e-mail settings are configured.
- **Communication tests:** You can test for the successful installation and communication infrastructure.
- **Inventory information:** A notification is sent to provide the necessary status and hardware information to IBM service personnel.

To send data and notifications to IBM service personnel, use one of the following e-mail addresses:

- For SAN Volume Controller nodes located in North America, Latin America, South America or the Caribbean Islands, use `callhome1@de.ibm.com`
- For SAN Volume Controller nodes located anywhere else in the world, use `callhome0@de.ibm.com`

Call Home e-mail can contain any combination of the following types of information:

- Contact name
- Contact phone number
- Offshift phone number

- Contact e-mail
- Machine location
- Record type
- Machine type
- Machine serial number
- Error ID
- Error code
- Software version
- FRU part number
- Cluster name
- Node ID
- Error sequence number
- Time stamp
- Object type
- Object ID
- Problem data

### Inventory information e-mail

Inventory information e-mail is a type of Call Home notification. Inventory information can be sent to IBM to assist IBM service personnel in evaluating your SAN Volume Controller system. Because inventory information is sent using the Call Home e-mail function, you must meet the Call Home function requirements and enable the Call Home e-mail function before you can attempt to send inventory information e-mail. You can adjust the contact information, adjust the frequency of inventory e-mail, or manually send an inventory e-mail using the SAN Volume Controller Console or the SAN Volume Controller command-line interface. Inventory information is automatically reported to IBM when you activate error reporting.

Inventory information that is sent to IBM includes the following information about the cluster on which the Call Home function is enabled:

- The output from the **svcinfo lscluster** command
- The output from the **svcinfo lsnodevpd** command (once for each node).
- The output from the **svcinfo lslicense** command

---

## Understanding the error codes

Error codes are generated for the SAN Volume Controller by the system error-log analysis and system configuration code.

Error codes help you to identify the cause of a problem, the failing field-replaceable units (FRUs), 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 displays on the front panel. The lower the number for the error code, the higher the priority. For example, cluster error code 1020 has a higher priority than cluster error code 1370.

## Using the error code tables

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

Perform the following steps to use the error code tables:

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 perform 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.

### Related tasks

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## Event codes

The SAN Volume Controller software generates events, such as informational and configuration. An event code 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 error event log, and depending on the configuration, can be notified through e-mail, SNMP, or syslog.

Configuration events are generated when configuration parameters are set. Configuration events are recorded in a separate log. Configuration events are not notified through e-mail, SNMP, or syslog.

### Informational event codes

The informational event codes provide information on the status of an operation.

Informational events are recorded in the error log and, depending on the configuration, can be notified through e-mail, SNMP, or syslog.

Informational event codes can be either notification type I (information) or notification type W (warning). An informational event report of type (W) might require user attention. Table 46 provides a list of informational event codes, the notification type, and the meaning of the event code.

Table 46. Informational event codes

Event code	Notification type	Description
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.
980310	I	A degraded or offline managed disk group is now online.
980349	I	A node has been successfully added to the cluster.
980352	I	Attempt to automatically recover offline node starting.
980371	I	One node in the I/O group is unavailable.

Table 46. Informational event codes (continued)

Event code	Notification type	Description
980372	W	Both nodes in the I/O group are unavailable.
980435	W	Failed to obtain directory listing from remote node.
980440	W	Failed to transfer file from remote node.
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.
981001	I	The cluster fabric view has been updated by a multiphase discovery.
981007	W	The managed disk is not on the preferred path.
981014	W	The LUN discovery has failed. The cluster 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 the first 2 TB of the disk can be accessed.
981020	W	The managed disk error count warning threshold has been met.
981022	I	Managed disk view smoothing start
981025	I	Drive firmware download started
981026	I	Drive FPGA download started
982003	W	Insufficient virtual extents.
982004	W	The migration suspended because of insufficient virtual extents or too many media errors on the source managed disk.
982007	W	Migration has stopped.
982009	I	Migration is complete.
982010	W	Copied disk I/O medium error.
983001	I	The FlashCopy is prepared.
983002	I	The FlashCopy is complete.
983003	W	The FlashCopy has stopped.
984001	W	First customer data being pinned in a virtual disk working set.
984002	I	All customer data in a virtual disk working set is now unpinned.
984003	W	The virtual disk working set cache mode is in the process of changing to synchronous destage because the virtual disk working set has too much pinned data.
984004	I	Virtual disk working set cache mode updated to allow asynchronous destage because enough customer data has been unpinned for the virtual disk working set.

Table 46. Informational event codes (continued)

Event code	Notification type	Description
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 within the timeout period.
986001	W	The space-efficient virtual disk copy data in a node is pinned.
986002	I	All space-efficient virtual disk copy data in a node is unpinned.
986010	I	The space-efficient virtual disk (VDisk) copy import has failed and the new VDisk is offline; either upgrade the SAN Volume Controller software to the required version or delete the VDisk.
986011	I	The space-efficient virtual disk copy import is successful.
986020	W	A space-efficient virtual disk copy space warning has occurred.
986030	I	A space-efficient virtual disk copy repair has started.
986031	I	A space-efficient virtual disk copy repair is successful.
986032	I	A space-efficient virtual disk copy validation is started.
986033	I	A space-efficient virtual disk copy validation is successful.
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 set to 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.
987301	W	The connection to a configured remote cluster has been lost.
987400	W	The node unexpectedly lost power but has now been restored to the cluster.

I  
I  
I

Table 46. Informational event codes (continued)

Event code	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 SAN Volume Controller cluster. If the problem persists, contact your IBM service representative for assistance.
989001	W	A managed disk group space warning has occurred.

## Configuration event codes

Configuration event codes are generated when configuration parameters are set.

Configuration event codes are recorded in a separate log. They do not raise notification types or send e-mails. Their error fixed flags are ignored. Table 47 provides a list of the configuration event codes and their meanings.

Table 47. Configuration event codes

Event code	Description
990101	Modify cluster (attributes in the <b>svctask chcluster</b> command)
990102	The e-mail test completed successfully
990103	The e-mail test failed
990105	Delete node from cluster (attributes in the <b>svctask rmnode</b> command)
990106	Create host (attributes in the <b>svctask mkhost</b> command)
990112	Cluster configuration dumped to file (attributes from the <b>svcluster -x dumpconfig</b> command)
990117	Create cluster (attributes in the <b>svctask mkcluster</b> command)
990118	Modify node (attributes in the <b>svctask chnode</b> command)
990119	Configure set controller name
990120	Shut down node (attributes in the <b>svctask stopcluster</b> command)
990128	Modify host (attributes in the <b>svctask chhost</b> command)
990129	Delete node (attributes in the <b>svctask rmnode</b> command)
990138	Virtual disk modify (attributes in the <b>svctask chvdisk</b> command)
990140	Virtual disk delete (attributes in the <b>svctask rmvdisk</b> command)
990144	Modify managed disk group (attributes in the <b>svctask chmdiskgrp</b> command)
990145	Delete managed disk group (attributes in the <b>svctask rmdiskgrp</b> command)
990148	Create managed disk group (attributes in the <b>svctask mkmdiskgrp</b> command)
990149	Modify managed disk (attributes in the <b>svctask chmdisk</b> command)
990150	Modify managed disk
990158	Managed disk included
990159	Quorum created

Table 47. Configuration event codes (continued)

Event code	Description
990160	Quorum destroy
990168	Modify the I/O group a virtual disk is assigned to
990169	Create a new virtual disk (attributes in the <b>svctask mkvdisk</b> command)
990173	Add a managed disk to managed disk group (attributes in the <b>svctask addmdisk</b> command)
990174	Delete a managed disk from managed disk group (attributes in the <b>svctask rmdisk</b> command)
990178	Add a port to a Host (attributes in the <b>svctask addhostport</b> command)
990179	Delete a port from a Host (attributes in the <b>svctask rmhostport</b> command)
990182	Create a virtual disk to Host SCSI mapping (attributes in the <b>svctask mkvdiskhostmap</b> command)
990183	Delete an virtual disk to Host SCSI mapping (attributes in the <b>svctask rmdiskhostmap</b> command)
990184	Create a FlashCopy mapping (attributes in the <b>svctask mkfcmap</b> command)
990185	Modify a FlashCopy mapping (attributes in the <b>svctask chfcmap</b> command)
990186	Delete a FlashCopy mapping (attributes in the <b>svctask rmfcmap</b> command)
990187	Prepare a FlashCopy mapping (attributes in the <b>svctask prestartfcmap</b> command)
990188	Prepare a FlashCopy consistency group (attributes in the <b>svctask prestartfcconsistgrp</b> command)
990189	Trigger a FlashCopy mapping (attributes in the <b>svctask startfcmap</b> command)
990190	Trigger a FlashCopy consistency group (attributes in the <b>svctask startfcconsistgrp</b> command)
990191	Stop a FlashCopy mapping (attributes in the <b>svctask stopfcmap</b> command)
990192	Stop a FlashCopy consistency group (attributes in the <b>svctask stopfcconsistgrp</b> command)
990193	FlashCopy set name
990194	Delete a list of ports from a Host (attributes in the <b>svctask rmhostport</b> command)
990196	Shrink a virtual disk.
990197	Expand a virtual disk (attributes in the <b>svctask expandvdisksize</b> command)
990198	Expand single extent a virtual disk
990199	Modify the I/O governing rate for a virtual disk
990203	Initiate manual managed disk discovery (attributes in the <b>svctask detectmdisk</b> command)
990204	Create FlashCopy consistency group (attributes in the <b>svctask mkfcconsistgrp</b> command)



Table 47. Configuration event codes (continued)

Event code	Description
990205	Modify FlashCopy consistency group (attributes in the <b>svctask chfcconsistgrp</b> command)
990206	Delete FlashCopy consistency group (attributes in the <b>svctask rmfcconsistgrp</b> command)
990207	Delete a list of Hosts (attributes in the <b>svctask rmhost</b> command)
990213	Change the I/O group a node belongs to (attributes in the <b>svctask chiogrp</b> command)
990216	Apply software upgrade (attributes in the <b>svctask applysoftware</b> command)
990219	Analyze error log (attributes in the <b>svctask finderr</b> command)
990220	Dump error log (attributes in the <b>svctask dumperrlog</b> command)
990222	Fix error log entry (attributes in the <b>svctask cherrstate</b> command)
990223	Migrate a single extent (attributes in the <b>svctask migrateexts</b> command)
990224	Migrate a number of extents
990225	Create a Metro Mirror or Global Mirror or Global Mirror relationship (attributes in the <b>svctask mkrrelationship</b> command)
990226	Modify a Metro Mirror or Global Mirror relationship (attributes in the <b>svctask chrrelationship</b> command)
990227	Delete a Metro Mirror or Global Mirror relationship (attributes in the <b>svctask rmrrelationship</b> command)
990229	Start a Metro Mirror or Global Mirror relationship (attributes in the <b>svctask startcrrelationship</b> command)
990230	Stop a Metro Mirror or Global Mirror relationship (attributes in the <b>svctask stopcrrelationship</b> command)
990231	Switch a Metro Mirror or Global Mirror relationship (attributes in the <b>svctask switchcrrelationship</b> command)
990232	Start a Metro Mirror or Global Mirror consistency group (attributes in the <b>svctask startcrconsistgrp</b> command)
990233	Stop a Metro Mirror or Global Mirror consistency group (attributes in the <b>svctask stopcrconsistgrp</b> command)
990234	Switch a Metro Mirror or Global Mirror consistency group (attributes in the <b>svctask switchcrconsistgrp</b> command)
990235	Managed disk migrated to a managed disk group
990236	Virtual disk migrated to a new managed disk
990237	Create partnership with remote cluster (attributes in the <b>svctask mkpartnership</b> command)
990238	Modify partnership with remote cluster (attributes in the <b>svctask chpartnership</b> command)
990239	Delete partnership with remote cluster (attributes in the <b>svctask rmpartnership</b> command)
990240	Create a Metro Mirror or Global Mirror consistency group (attributes in the <b>svctask mkrccconsistgrp</b> command)
990241	Modify a Metro Mirror or Global Mirror consistency group (attributes in <b>svctask chrccconsistgrp</b> )

Table 47. Configuration event codes (continued)

Event code	Description
990242	Delete a Metro Mirror or Global Mirror consistency group (attributes in the <b>svctask rmrcconsistgrp</b> command)
990245	Node shutdown imminent
990246	Node remove
990247	Node unpend
990380	Time zone changed (attributes in the <b>svctask settimezone</b> command)
990383	Change cluster time (attributes in the <b>svctask setclustertime</b> command)
990385	System time changed
990386	SSH key added (attributes in the <b>svctask addsshkey</b> command)
990387	SSH key removed (attributes in the <b>svctask rmsshkey</b> command)
990388	All SSH keys removed (attributes in the <b>svctask rmallsshkeys</b> command)
990390	Add node to the cluster
990395	Shutdown or reset node
990410	The software installation has started.
990415	The software installation has completed.
990420	The software installation has failed.
990423	The software installation has stalled.
990425	The software installation has stopped.
990430	The Planar Serial Number has changed.
990501	The licensed feature has changed. See the license settings log for details.
990510	The configuration limits have been changed.
991024	I/O tracing has finished and the managed disk has been triggered.
991025	The autoexpand setting of the VDisk has been modified.
991026	The primary copy of the VDisk has been modified.
991027	The VDisk synchronization rate has been modified.
991028	The space-efficient VDisk warning capacity has been modified.
991029	A mirrored copy has been added to a VDisk.
991030	A repair of mirrored VDisk copies has started.
991031	A VDisk copy has been split from a mirrored VDisk.
991032	A VDisk copy has been removed from a mirrored VDisk.

## SCSI error reporting

SAN Volume Controller nodes can notify their hosts of errors for SCSI commands that are issued.

### SCSI status

Some errors are part of the SCSI architecture and are handled by the host application or device drivers without reporting an error. Some errors, such as read

and write I/O errors and errors 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 errors, 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 error in the SAN Volume Controller cluster error log.

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

Table 48 describes the SCSI status and codes that are returned by the SAN Volume Controller nodes.

*Table 48. SCSI status*

Status	Code	Description
Good	00h	The command was successful.
Check condition	02h	The command failed and sense data is available.
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 SAM2 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 is reported as specified in SAM-2.
Task aborted	40h	This is returned if TAS is set in the control mode page 0Ch. The SAN Volume Controller node has a default setting of TAS=0 , which is cannot be changed; therefore, the SAN Volume Controller node does not report this status.

## SCSI Sense

SAN Volume Controller nodes notify the hosts of errors on SCSI commands. Table 49 defines the SCSI sense keys, codes and qualifiers that are returned by the SAN Volume Controller nodes.

*Table 49. 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 cluster and cannot perform I/O operations. The additional sense does not have additional information.

Table 49. SCSI sense keys, codes, and qualifiers (continued)

Key	Code	Qualifier	Definition	Description
2h	04h	0Ch	Not Ready. The target port is in the state of unavailable.	The following conditions are possible: <ul style="list-style-type: none"> <li>The node lost sight of the cluster and cannot perform I/O operations. The additional sense does not have additional information.</li> <li>The node is in contact with the cluster 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 virtual disks (VDisks).</li> </ul>
3h	00h	00h	Medium error	This is only returned for read or write I/Os. The I/O suffered an error at a specific LBA within its scope. The location of the error is reported within the sense data. The additional sense also includes a reason code that relates the error to the corresponding error log entry. For example, a RAID controller error or a migrated medium error.
4h	08h	00h	Hardware error. A command to logical unit communication failure has occurred.	The I/O suffered an error that is associated with an I/O error 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 error is also returned from FlashCopy target VDisks 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 SAN Volume Controller node specific log entry. The field is a 32-bit unsigned number that is presented with the most significant byte first. Table 50 on page 189 lists the reason codes and their definitions.

If the reason code is not listed in Table 50 on page 189, the code refers to a specific error in the SAN Volume Controller cluster error log that corresponds to the sequence number of the relevant error log entry.

Table 50. 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 has been recreated.
73	Running on a node that has been contracted out for some reason that is not attributable to any path going offline.
80	Wait for the repair to complete, or delete the virtual disk.
81	Wait for the validation to complete, or delete the virtual disk.
82	An offline space-efficient VDisk has caused data to be pinned in the directory cache. Adequate performance cannot be achieved for other space-efficient VDisks, so they have been taken offline.
85	The VDisk has been taken offline because checkpointing to the quorum disk failed.
86	The <b>svctask repairvdiskcopy -medium</b> command has created a virtual medium error where the copies differed.

## Object types

You can use the object code to determine the object type.

Table 51 lists the object codes and corresponding object types.

Table 51. Object types

Object code	Object type
1	mdisk
2	mdiskgrp
3	vdisk
4	node
5	host
7	iogroup
8	fcgrp
9	rcgrp
10	fcmap
11	rcmap
12	wwpn
13	cluster
16	device
17	SCSI lun
18	quorum

Table 51. Object types (continued)

Object code	Object type
34	Fibre-channel adapter
38	VDisk copy
39	Syslog server
40	SNMP server
41	E-mail server
42	User group
44	Cluster IP
46	SAS adapter

## Error Codes

Error codes provide a unique entry to service procedures. Each error code has an error ID that uniquely identifies the condition that caused the error.

Error IDs are recorded in the error event log. When the number of error IDs of a specific type for a specific resource exceeds a predetermined threshold, a notification type might be raised and an e-mail might be sent, depending on the cluster configuration. Different notification types are signaled, depending on the severity of the error event. The notification type is used by the Call Home e-mail service to determine the recipients, the title, and the contents of the e-mail. The following notification types are possible:

**Error** This type identifies unexpected conditions that might be the result of a system failure. If configured, this type causes a notification type to be sent to the monitoring application. An e-mail can also be sent to the IBM Support Center and the system administrator.

**Warning** This type identifies unexpected conditions that might be experienced during user operations. These conditions can result from device errors or incorrect user actions. If configured, this type causes a notification type to be sent to the monitoring application. An e-mail can also be sent to the system administrator.

**Information** This type identifies conditions where a user might want to be notified of the completion of an operation. If configured, this type causes a notification type to be sent to the monitoring application. An e-mail can also be sent to the system administrator.

Table 52 lists the error codes and corresponding error IDs.

Table 52. Error codes

Error ID	Notification type	Condition	Error Code
009020	E	An automatic cluster recovery has started. All configuration commands are blocked.	1001
009040	E	The error event log is full.	1002

Table 52. Error codes (continued)

Error ID	Notification type	Condition	Error Code
009052	E	The following causes are possible: <ul style="list-style-type: none"> <li>• The node is missing</li> <li>• The node is no longer a functional member of the cluster</li> <li>• One or more nodes are not available</li> <li>• The node has probably had a hardware or software failure</li> </ul>	1195
009100	W	The software install process has failed.	2010
009101	W	The software upgrade package delivery has failed.	2010
009110	W	Software installation process stalled due to lack of redundancy.	2010
009115	W	A software downgrade has failed.	2008
009150	W	Unable to connect to the SMTP (e-mail) server	2600
009151	W	Unable to send mail through the SMTP (e-mail) server	2601
009170	W	The Metro Mirror or Global Mirror 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	The Metro Mirror or Global Mirror feature has exceeded the amount that is licensed.	3032
009176	W	The value set for the virtualization feature capacity is not valid.	3029
009177	W	A physical disk FlashCopy feature license is required.	3035
009178	W	A physical disk Metro Mirror and Global Mirror feature license is required.	3036
009179	W	A virtualization feature license is required.	3025
009180	E	Automatic recovery of offline node failed.	1194
009181	W	Unable to send e-mail to any of the configured e-mail servers.	3081
010002	E	The node ran out of base event sources. As a result, the node has stopped and exited the cluster.	2030
010003	W	The number of device logins has reduced.	1630
010006	E	A software error has occurred.	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

Table 52. Error codes (continued)

Error ID	Notification type	Condition	Error Code
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 I/O error has occurred.	1310
010020	E	The managed disk error count threshold has exceeded.	1310
010021	E	There are too many devices presented to the cluster.	1200
010022	E	There are too many managed disks presented to the cluster.	1200
010023	E	There are too many LUNs presented to a node.	1200
010025	W	A disk I/O medium error has occurred.	1320
010026	E	There are no managed disks that can be used as a quorum disk.	1330
010027	E	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
010030	E	A managed disk error recovery procedure (ERP) has occurred. The node or controller reported the following: <ul style="list-style-type: none"> <li>• Sense</li> <li>• Key</li> <li>• Code</li> <li>• Qualifier</li> </ul>	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	E	Persistent unsupported controller configuration.	1695
010035	W	A quorum disk is configured on the controller that is set to not allow quorum.	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 cluster nodes through a single target port.	1627
010043	E	The controller system device is only connected to the cluster nodes through half of the expected target ports.	1627
010044	E	The controller system device has disconnected all target ports to the cluster nodes.	1627
010050	E	A solid-state drive (SSD) failed. A rebuild is required.	1201
010051	E	A solid-state drive (SSD) is missing.	1202
010053	E	A solid-state drive (SSD) is reporting a PFA.	1215
010054	E	A solid-state drive (SSD) is reporting too many errors.	1215
010055	E	An unrecognized SAS device.	1665

I  
I  
I  
I  
I



Table 52. Error codes (continued)

Error ID	Notification type	Condition	Error Code
010056	E	SAS error counts exceeded the warning thresholds.	1216
010057	E	SAS errors exceeded critical thresholds.	1216
010058	E	A solid-state drive (SSD) initialization failed.	1661
010059	E	A solid-state drive (SSD) is offline due to excessive errors.	1311
010060	E	A solid-state drive (SSD) exceeded the warning temperature threshold.	1218
010061	E	A solid-state drive (SSD) exceeded the offline temperature threshold.	1217
020001	E	There are too many medium errors on the managed disk.	1610
020002	E	A managed disk group is offline.	1620
020003	W	There are insufficient virtual extents.	2030
030000	W	The trigger prepare command has failed because of a cache flush failure.	1900
030010	W	The mapping is stopped because of the error that is indicated in the data.	1910
030020	W	The mapping is stopped because of a cluster or complete I/O group failure, and the current state of the relationship could not be recovered.	1895
050001	W	The relationship is stopped because of a cluster or complete I/O group failure, and the current state of the mapping could not be recovered.	1700
050002	W	A Metro Mirror or Global Mirror relationship or consistency group exists within a cluster, but its partnership has been deleted.	3080
050010	W	A Metro Mirror or Global Mirror relationship has stopped because of a persistent I/O error.	1920
050020	W	A Metro Mirror or Global Mirror relationship has stopped because of an error that is not a persistent I/O error.	1720
050030	W	There are too many cluster partnerships. The number of cluster partnerships has been reduced.	1710
050031	W	There are too many cluster partnerships. The cluster has been excluded.	1710
060001	W	The space-efficient virtual disk copy is offline because there is insufficient space.	1865
060002	W	The space-efficient virtual disk copy is offline because the metadata is corrupt.	1862
060003	W	The space-efficient virtual disk copy is offline because the repair has failed.	1860
062001	W	Unable to mirror medium error during VDisk copy synchronization	1950
062002	W	The mirrored VDisk is offline because the data cannot be synchronized.	1870
062003	W	The repair process for the mirrored disk has stopped because there is a difference between the copies.	1600

Table 52. Error codes (continued)

Error ID	Notification type	Condition	Error Code
072004	E	A CMOS battery failure has occurred. This error applies to the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1670
072005	E	A CMOS battery failure has occurred. This error applies to only the SAN Volume Controller 2145-8G4 model.	1670
072006	E	A CMOS battery failure has occurred. This error applies to only the SAN Volume Controller 2145-8A4 model.	1670
072007	E	A CMOS battery failure has occurred. This error applies to only the SAN Volume Controller 2145-CF8 model.	1670
072101	E	The processor is missing. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1025
072102	E	The processor is missing. This error applies to only the SAN Volume Controller 2145-8G4 model.	1025
073003	E	The fibre-channel ports are not operational.	1060
073005	E	A cluster path failure has occurred.	1550
073006	W	The SAN is not correctly zoned. As a result, more than 512 ports on the SAN have logged into one SAN Volume Controller port.	1800
073101	E	The 2-port fibre-channel adapter card in slot 1 is missing. This error applies to only the SAN Volume Controller 2145-8F2 model.	1014
073102	E	The 2-port fibre-channel adapter in slot 1 has failed. This error applies to only the SAN Volume Controller 2145-8F2 model.	1054
073104	E	The 2-port fibre-channel adapter in slot 1 has detected a PCI bus error. This error applies to only the SAN Volume Controller 2145-8F2 model.	1017
073201	E	The 2-port fibre-channel adapter in slot 2 is missing. This error applies to only the SAN Volume Controller 2145-8F2 model.	1015
073202	E	The 2-port fibre-channel adapter in slot 2 has failed. This error applies to only the SAN Volume Controller 2145-8F2 model.	1056
073204	E	The 2-port fibre-channel adapter in slot 2 has detected a PCI bus error. This error applies to only the SAN Volume Controller 2145-8F2 model.	1018
073251	E	The 4-port fibre-channel adapter in slot 1 is missing. This error applies to only the SAN Volume Controller 2145-8G4 model.	1011
073252	E	The 4-port fibre-channel adapter in slot 1 has failed. This error applies to only the SAN Volume Controller 2145-8G4 model.	1055
073258	E	The 4-port fibre-channel adapter in slot 1 has detected a PCI bus error. This error applies to only the SAN Volume Controller 2145-8G4 model.	1013

Table 52. Error codes (continued)

Error ID	Notification type	Condition	Error Code
073261	E	The 4-port fibre-channel adapter in slot 1 has detected a PCI bus error. This error applies to only the SAN Volume Controller 2145-8A4 model.	1011
073262	E	The 4-port fibre-channel adapter in slot 1 has detected a PCI bus error. This error applies to only the SAN Volume Controller 2145-8A4 model.	1055
073268	E	The 4-port fibre-channel adapter in slot 1 has detected a PCI bus error. This error applies to only the SAN Volume Controller 2145-8A4 model.	1013
073271	E	The 4-port fibre-channel adapter in slot 1 has detected a PCI bus error. This error applies to only the SAN Volume Controller 2145-CF8 model.	1011
073272	E	The 4-port fibre-channel adapter in slot 1 has detected a PCI bus error. This error applies to only the SAN Volume Controller 2145-CF8 model.	1055
073278	E	The 4-port fibre-channel adapter in slot 1 has detected a PCI bus error. This error applies to only the SAN Volume Controller 2145-CF8 model.	1013
073301	E	The 4-port fibre-channel adapter in slot 2 is missing. This error applies to only the SAN Volume Controller 2145-8F4 model.	1016
073302	E	The 4-port fibre-channel adapter in slot 2 has failed. This error applies to only the SAN Volume Controller 2145-8F4 model.	1057
073304	E	The 4-port fibre-channel adapter in slot 2 has detected a PCI bus error. This error applies to only the SAN Volume Controller 2145-8F4 model.	1019
073305	E	One or more fibre-channel ports are running at a speed that is lower than the last saved speed.	1065
073310	E	A duplicate fibre-channel frame has been detected, which indicates that there is an issue with the fibre-channel fabric. Other fibre-channel errors might also be generated.	1203
074001	W	Unable to determine the vital product data (VPD) for an FRU. This is probably because a new FRU has been installed and the software does not recognize that FRU. The cluster continues to operate; however, you must upgrade the software to fix this warning.	2040
074002	E	The node warm started after a software error.	2030
075001	E	The flash boot device has failed.	1040
075002	E	The flash boot device has recovered.	1040
075005	E	A service controller read failure has occurred.	1044
075011	E	The flash boot device has failed. This error applies to only the SAN Volume Controller 2145-8G4 model.	1040
075012	E	The flash boot device has recovered. This error applies to only the SAN Volume Controller 2145-8G4 model.	1040

Table 52. Error codes (continued)

Error ID	Notification type	Condition	Error Code
075015	E	A service controller read failure has occurred. This error applies to only the SAN Volume Controller 2145-8G4 model.	1044
075021	E	The flash boot device has failed. This error applies to only the SAN Volume Controller 2145-8A4 model.	1040
075022	E	The flash boot device has recovered. This error applies to only the SAN Volume Controller 2145-8A4 model.	1040
075025	E	A service controller read failure has occurred. This error applies to only the SAN Volume Controller 2145-8A4 model.	1044
075031	E	The flash boot device has failed. This error applies to only the SAN Volume Controller 2145-CF8 model.	1040
075032	E	The flash boot device has recovered. This error applies to only the SAN Volume Controller 2145-CF8 model.	1040
075035	E	A service controller read failure has occurred. This error applies to only the SAN Volume Controller 2145-CF8 model.	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
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	Degraded PCIe lanes on a high-speed SAS adapter.	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	Vital product data (VPD) is corrupt on high-speed SAS adapter.	1121
077101	E	The service processor shows a fan 40×40×28 failure. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1090
077102	E	The service processor shows a fan 40×40×56 failure. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1091
077105	E	The service processor shows a fan failure. This error applies to only the SAN Volume Controller 2145-8G4 model.	1089
077106	E	The service processor shows a fan failure. This error applies to only the SAN Volume Controller 2145-8A4 model.	1089
077107	E	The service processor shows a fan failure. This error applies to only the SAN Volume Controller 2145-CF8 model.	1089

Table 52. Error codes (continued)

Error ID	Notification type	Condition	Error Code
077111	E	The node ambient temperature threshold has exceeded. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1094
077112	E	The node processor warning temperature threshold has exceeded. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1093
077113	E	The node processor or ambient critical threshold has exceeded. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1092
077121	E	System board - any voltage high. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1100
077124	E	System board - any voltage low. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1105
077128	E	A power management board voltage failure has occurred. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1110
077161	E	The node ambient temperature threshold has exceeded. This error applies to only the SAN Volume Controller 2145-8G4 model.	1094
077162	E	The node processor warning temperature threshold has exceeded. This error applies to only the SAN Volume Controller 2145-8G4 model.	1093
077163	E	The node processor or ambient critical threshold has exceeded. This error applies to only the SAN Volume Controller 2145-8G4 model.	1092
077165	E	The node ambient temperature threshold has exceeded. This error applies to only the SAN Volume Controller 2145-8A4 model.	1094
077166	E	The node processor warning temperature threshold has exceeded. This error applies to only the SAN Volume Controller 2145-8A4 model.	1093
077167	E	The node processor or ambient critical threshold has exceeded. This error applies to only the SAN Volume Controller 2145-8A4 model.	1092
077171	E	System board - any voltage high. This error applies to only the SAN Volume Controller 2145-8G4 model.	1101
077172	E	System board - any voltage high. This error applies to only the SAN Volume Controller 2145-8A4 model.	1101
077173	E	System board - any voltage high. This error applies to only the SAN Volume Controller 2145-CF8 model.	1101
077174	E	System board - any voltage low. This error applies to only the SAN Volume Controller 2145-8G4 model.	1106

I  
I

Table 52. Error codes (continued)

Error ID	Notification type	Condition	Error Code
077175	E	System board - any voltage low. This error applies to only the SAN Volume Controller 2145-8A4 model.	1106
077176	E	System board - any voltage low. This error applies to only the SAN Volume Controller 2145-CF8 model.	1106
077178	E	A power management board voltage failure has occurred. This error applies to only the SAN Volume Controller 2145-8G4 model.	1110
077185	E	The node ambient temperature threshold has exceeded. This error applies to only the SAN Volume Controller 2145-CF8 model.	1094
077186	E	The node processor warning temperature threshold has exceeded. This error applies to only the SAN Volume Controller 2145-CF8 model.	1093
077187	E	The node processor or ambient critical threshold has exceeded. This error applies to only the SAN Volume Controller 2145-CF8 model.	1092
077188	E	A power management board voltage failure has occurred. This error applies to only the SAN Volume Controller 2145-CF8 model.	1110
078001	E	A power domain error has occurred. Both nodes in a pair are powered by the same uninterruptible power supply.	1155
079000	W	Data has not been recovered on virtual disks (VDisks).	1850
079500	W	The limit on the number of cluster secure shell (SSH) sessions has been reached.	2500
079501	W	Unable to access the Network Time Protocol (NTP) network time server.	2700
081001	E	An Ethernet port failure has occurred.	1400
082001	E	A server error has occurred.	2100
083101	E	An uninterruptible power supply communications failure has occurred. The RS232 connection between a node and its uninterruptible power supply is faulty. This error applies to only the 2145 UPS-1U model.	1146
083102	E	The uninterruptible power supply output is unexpectedly high. The uninterruptible power supply is probably connected to a non-SAN Volume Controller load. This error applies to only the 2145 UPS-1U model.	1166
083103	E	The uninterruptible power supply battery has reached end of life. This error applies to only the 2145 UPS-1U model.	1191
083104	E	An uninterruptible power supply battery failure has occurred. This error applies to only the 2145 UPS-1U model.	1181
083105	E	An uninterruptible power supply electronics failure has occurred. This error applies to only the 2145 UPS-1U model.	1171
083107	E	Uninterruptible power supply overcurrent. This error applies to only the 2145 UPS-1U model.	1161
083108	E	An uninterruptible power supply failure has occurred. This error applies to only the 2145 UPS-1U model.	1186

Table 52. Error codes (continued)

Error ID	Notification type	Condition	Error Code
083109	E	Uninterruptible power supply ac input power fault. This error applies to only the 2145 UPS-1U model.	1141
083110	E	An uninterruptible power supply configuration error has occurred. This error applies to only the 2145 UPS-1U model.	1151
083111	E	Uninterruptible power supply ambient over temperature. This error applies to only the 2145 UPS-1U model.	1136
083112	E	Uninterruptible power supply over temperature warning. This error applies to only the 2145 UPS-1U model.	3001
083113	E	An uninterruptible power supply software error has occurred. This error applies to only the 2145 UPS-1U model.	3011

## Determining a hardware boot failure

If you see that the hardware boot display stays on the front panel for more than three minutes, the node cannot boot. The cause might be a hardware failure or the software on the hard disk drive is missing or damaged.

The following figure shows an example of a hardware boot display.



Figure 69. Example of a boot error code

Perform the following steps to determine a hardware boot failure:

1. Attempt to restore the software by using the node rescue procedure.
2. If node rescue fails, perform the actions that are described for any failing node rescue code or procedure.

### Related concepts

“Hardware boot” on page 141

The hardware boot display shows system data when power is first applied to the node as the node searches for a disk drive to boot.

### Related tasks

“Performing the node rescue” on page 217

If it is necessary 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 SAN Volume Controller software.

## Understanding the boot codes

The boot codes are displayed on the screen when a node is booting.

The codes indicate the progress of the boot operation or the condition that has caused the node to fail to boot. The codes can isolate failures when the boot hangs or when the boot detects an unrecoverable error. Line 1 of the front panel displays the message `Booting` that is followed by the boot code. Line 2 of the display shows a boot progress indicator. If the boot hangs, the progress bar stops and you can

use the code to isolate the fault. If the boot code detects a hardware error, Failed is displayed and you can use the error code to isolate the failure. In some cases, the same code that displays with the Booting message is also displayed as the Failed error code. Figure 70 provides a view of the boot progress display and Figure 71 provides a view of the boot failed display.



Figure 70. Example of a boot progress display



Figure 71. Example of a boot error code

#### **Related concepts**

“Boot progress indicator” on page 139

Boot progress is displayed on the front panel of the SAN Volume Controller.

#### **100 Boot is running.**

##### **Explanation**

If the progress bar does not advance for two minutes, the boot process is hung.

##### **Action**

Go to the hardware boot MAP to resolve the problem.

Possible Cause-FRUs or other:

2145-CF8

- Service controller (47%)
- Service controller cable (47%)
- System board assembly (6%)

2145-8G4 or 2145-8A4

- Service controller (95%)
- System board (5%)

2145-8F2 or 2145-8F4

- Service controller (95%)
- Frame assembly (5%)

2145-4F2

- Service controller (95%)
- System board (5%)



**Related tasks**

“MAP 5900: Hardware boot” on page 423

MAP 5900: Hardware boot helps you solve problems that are preventing the node from starting its boot sequence.

**110 The 2145 is loading kernel code.****Explanation**

The progress bar has stopped.

**Action**

If the progress bar has been stopped for two 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.

**Related tasks**

“Performing the node rescue” on page 217

If it is necessary 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 SAN Volume Controller software.

**120 A disk drive hardware error has occurred.****Explanation**

A disk drive hardware error has occurred.

**Action**

Exchange the FRU for a new FRU. (See “Possible Cause-FRUs or other.”)

Possible Cause-FRUs or other:

2145-CF8

- Disk drive (50%)
- Disk controller (30%)
- Disk backplane (10%)
- Disk signal cable (8%)
- Disk power cable (1%)
- System board (1%)

2145-8G4 or 2145-8A4

- Disk drive assembly (95%)
- Disk cable assembly (4%)
- System board (1%)

2145-8F2 or 2145-8F4

- Disk drive assembly (98%)
- Frame assembly (2%)

2145-4F2

- Disk drive assembly (95%)
- Disk drive cables (5%)

### **130 The 2145 is checking the file systems.**

#### **Explanation**

The progress bar is stopped.

#### **Action**

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.

#### **Related tasks**

“Performing the node rescue” on page 217

If it is necessary 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 SAN Volume Controller software.

### **132 The 2145 is updating the node’s BIOS CMOS settings.**

#### **Explanation**

When the update completes, the node reboots.

#### **Action**

If the progress bar has stopped for more than ten minutes, or if the display has shown codes 100 and 132 three times or more, go to MAP 5900: Hardware boot to resolve the problem.

#### **Related tasks**

“MAP 5900: Hardware boot” on page 423

MAP 5900: Hardware boot helps you solve problems that are preventing the node from starting its boot sequence.

### **135 The 2145 is verifying the software.**

#### **Explanation**

The 2145 is verifying the software.

#### **Action**

If the progress bar has stopped for at least 90 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.

## **137 Updating the system board service processor firmware.**

### **Explanation**

The process of updating the system board service processor firmware might take up to 90 minutes.

### **Action**

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.

Possible Cause-FRUs or other:

2145-8G4 or 2145-8A4 or 2145-CF8

- System board (100%)

2145-8F2 or 2145-8F4

- Frame assembly (100%)

### **Related tasks**

“Performing the node rescue” on page 217

If it is necessary 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 SAN Volume Controller software.

## **140 The 2145 software is damaged.**

### **Explanation**

The 2145 software is damaged.

### **Action**

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.

### **Related tasks**

“Performing the node rescue” on page 217

If it is necessary 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 SAN Volume Controller software.

## **145 Unable to detect fibre-channel adapter**

### **Explanation**

The 2145 cannot detect any fibre-channel adapter cards.

### **Action**

Ensure that a fibre-channel adapter card has been installed. Ensure that the fibre-channel card is seated correctly in the riser card. Ensure that the riser card is seated correctly on the system board. If the problem persists, exchange FRUs for new FRUs in the order shown.

Possible Cause-FRUs or other:

2145-CF8

- 4-port fibre-channel host bus adapter assembly (95%)
- System board assembly (5%)

2145-8G4 or 2145-8A4

- 4-port fibre-channel host bus adapter (80%)
- Riser card (19%)
- System board (1%)

2145-8F4

- 4-port fibre-channel host bus adapter (99%)
- Frame assembly (1%)

2145-8F2

- fibre-channel host bus adapter (full height) (40%)
- fibre-channel host bus adapter (low profile) (40%)
- Riser card, PCI (full height) (9%)
- Riser card, PCI (low profile) (9%)
- Frame assembly (2%)

2145-4F2

- fibre-channel adapter assembly (98%)
- System board (2%)

### **150 The 2145 is loading the cluster code.**

#### **Explanation**

The 2145 is loading the cluster code.

#### **Action**

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.

#### **Related tasks**

“Performing the node rescue” on page 217

If it is necessary 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 SAN Volume Controller software.

### **155 The 2145 is loading the cluster data.**

#### **Explanation**

The 2145 is loading the cluster data.

## Action

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.

## Related tasks

“Performing the node rescue” on page 217

If it is necessary 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 SAN Volume Controller software.

## 160 The 2145 is recovering flash disk.

### Explanation

The flash disk recovery operation will take more time when a node rescue is occurring than when a node rescue is not occurring.

## Action

When a node rescue is occurring, if the progress bar has been stopped for at least thirty minutes, exchange the FRU for a new FRU. When a node rescue is not occurring, if the progress bar has been stopped for at least fifteen minutes, exchange the FRU for a new FRU. See “Possible Cause-FRUs or other” at the end of this section.

Possible Cause-FRUs or other:

2145-CF8

- Service controller (95%)
- Service controller cable (5%)

All previous 2145 models

- Service Controller (100%)

## 170 A flash module hardware error has occurred.

### Explanation

A flash module hardware error has occurred.

## Action

Exchange the FRU for a new FRU. (See “Possible Cause-FRUs or other.”)

Possible Cause-FRUs or other:

2145-CF8

- Service controller (95%)
- Service controller cable (5%)

All previous 2145 models

- Service controller (100%)

## **174 The service processor on the system board has failed.**

### **Explanation**

The service processor on the system board has failed.

### **Action**

Exchange the FRU for a new FRU. (See "Possible Cause-FRUs or other.")

Possible Cause-FRUs or other:

2145-8G4 or 2145-8A4 or 2145-CF8

- System board assembly (100%)

2145-8F2 or 2145-8F4

- Frame assembly (100%)

2145-4F2

- System board assembly (100%)

## **175 The service processor has indicated an overheating.**

### **Explanation**

The ambient temperature of the node is too high.

The node pauses in this condition. When the ambient temperature returns to being within the acceptable range, the start sequence continues.

If the temperature continues to rise, the node might exceed a critical temperature threshold. If the threshold is exceeded, the node will power off.

### **Action**

1. Ensure that the air temperature within the rack meets requirements.
2. Clear the vents in the rack and both the front and the back of the node. Remove any heat sources from around the node. Ensure that the airflow around the node is not restricted.
3. Go to the light path diagnostic MAP and perform the light path diagnostic procedures.
4. If the temperature around the node is within range but the error is still reported, assume that the sensor has failed and replace the system board or frame assembly depending on the model.

Possible Cause-FRUs or other:

2145-8G4 or 2145-8A4 or 2145-CF8

- FRU indicated by light path diagnostic (5%)
- System board (5%)

2145-8F2 or 2145-8F4

- FRU indicated by light path diagnostic (5%)
- Frame assembly (5%)

2145-4F2

- System board assembly (10%)

Other:

- System environment (90%)

#### **Related tasks**

“Performing the node rescue” on page 217

If it is necessary 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 SAN Volume Controller software.

“MAP 5800: Light path” on page 406

MAP 5800: Light path helps you to solve hardware problems on the SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2 that are preventing the node from booting.

“MAP 5900: Hardware boot” on page 423

MAP 5900: Hardware boot helps you solve problems that are preventing the node from starting its boot sequence.

#### **Related reference**

“SAN Volume Controller 2145-CF8 environment requirements” on page 34

Before installing a SAN Volume Controller 2145-CF8 node, your physical environment must meet certain requirements. This includes verifying that adequate space is available and that requirements for power and environmental conditions are met.

“SAN Volume Controller 2145-8A4 environment requirements” on page 37

Before the SAN Volume Controller 2145-8A4 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.

“SAN Volume Controller 2145-8G4 environment requirements” on page 39

Before the SAN Volume Controller 2145-8G4 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.

### **180 There is a fault in the communications cable, the serial interface in the uninterruptible power supply 2145 UPS, or the 2145.**

#### **Explanation**

There is a fault in the communications cable, the serial interface in the uninterruptible power supply 2145 UPS, or 2145.

#### **Action**

Check that the communications cable is correctly plugged in to the 2145 and the 2145 UPS. If the cable is plugged in correctly, replace the FRUs in the order shown.

Possible Cause-FRUs or other:

2145-4F2

- 2145 power cable assembly (40%)
- 2145 UPS electronics assembly (30%)
- 2145 system board assembly (25%)
- 2145 disk drive assembly (5%)

**181 There is a fault in the communications cable, the serial interface in the uninterruptible power supply 2145 UPS-1U, or the 2145.**

**Explanation**

There is a fault in the communications cable, the serial interface in the uninterruptible power supply 2145 UPS-1U, or 2145.

**Action**

Check that the communications cable is correctly plugged in to the 2145 and the 2145 UPS-1U. If the cable is plugged in correctly, replace the FRUs in the order shown.

Possible Cause-FRUs or other:

2145-8G4 or 2145-8A4 or 2145-CF8

- 2145 power cable assembly (40%)
- 2145 UPS-1U assembly (30%)
- 2145 system board (30%)

2145-8F2 or 2145-8F4

- 2145 power cable assembly (40%)
- 2145 UPS-1U assembly (30%)
- 2145 frame assembly (30%)

2145-4F2

- 2145 power cable assembly (40%)
- 2145 UPS-1U assembly (30%)
- 2145 system board assembly (25%)
- 2145 disk drive assembly (5%)

**185 The uninterruptible power supply battery has reached its end of life.**

**Explanation**

The maximum available capacity can no longer support four 2145s.

**Action**

Exchange the FRU for a new FRU. (See "Possible Cause-FRUs or other.")

Possible Cause-FRUs or other:

- 2145 UPS battery assembly (100%)

**186 The uninterruptible power supply battery has reached its end of life.**

**Explanation**

The maximum available capacity can no longer support one 2145.



### **Action**

Exchange the FRU for a new FRU. (See "Possible Cause-FRUs or other.") After replacing the battery assembly, if the 2145 UPS-1U service indicator is on, press and hold the 2145 UPS-1U Test button for three seconds to start the self-test and verify the repair.

Possible Cause-FRUs or other:

- 2145 UPS-1U battery assembly (100%)

### **190 A problem has occurred with the uninterruptible power supply battery.**

#### **Explanation**

A problem has occurred with the 2145 UPS battery.

### **Action**

Exchange the FRU for a new FRU. (See "Possible Cause-FRUs or other.")

Possible Cause-FRUs or other:

- 2145 UPS battery assembly (100%)

### **191 A problem has occurred with the uninterruptible power supply battery.**

#### **Explanation**

A problem has occurred with the uninterruptible power supply 2145 UPS-1U battery.

### **Action**

Exchange the FRU for a new FRU. (See "Possible Cause-FRUs or other.") After replacing the battery assembly, if the 2145 UPS-1U service indicator is on, press and hold the 2145 UPS-1U Test button for three seconds to start the self-test and verify the repair.

Possible Cause-FRUs or other:

- 2145 UPS-1U battery assembly
- 2145 UPS-1U assembly

### **195 A problem has occurred with the uninterruptible power supply electronics.**

#### **Explanation**

A problem has occurred with the 2145 UPS electronics.

### **Action**

Exchange the FRU for a new FRU. (See "Possible Cause-FRUs or other.")

Possible Cause-FRUs or other:

- 2145 UPS electronics assembly (100%)

**196 A problem has occurred with the uninterruptible power supply electronics.**

**Explanation**

A problem has occurred with the 2145 UPS-1U electronics.

**Action**

Exchange the FRU for a new FRU. (See "Possible Cause-FRUs or other.")

Possible Cause-FRUs or other:

- 2145 UPS-1U assembly

**200 A problem has occurred with the uninterruptible power supply.**

**Explanation**

A problem has occurred with the 2145 UPS.

**Action**

Exchange the FRU for a new FRU. (See "Possible Cause-FRUs or other.")

Possible Cause-FRUs or other:

- 2145 UPS (100%)

**205 A problem with output overload was reported by the uninterruptible power supply.**

**Explanation**

A problem with output overload has been reported by the uninterruptible power supply 2145 UPS. The Mode Indicator on the 2145 UPS front panel is illuminated red.

**Action**

Ensure that only one or two 2145s are receiving power from the 2145 UPS. Also ensure that no other devices are connected to the 2145 UPS.

Disconnect the 2145-4F2s from the 2145 UPS. If the Mode Indicator is now green on the disconnected 2145 UPS, reconnect the 2145-4F2s one at a time to determine which one causes the overload. Then, on the 2145-4F2 which caused the problem, in the sequence shown, exchange the FRUs for new FRUs. See "Possible Cause-FRUs or other" after the last action in this section.

If the Mode Indicator is still red with all outputs disconnected, replace the 2145 UPS electronics assembly. If the Mode Indicator is still red, replace the 2145 UPS assembly.

Possible Cause-FRUs or other:

- 2145-4F2 power cable assembly (45%)
- 2145-4F2 power supply assembly (45%)
- 2145 UPS electronics assembly (9%)
- 2145 UPS assembly (1%)

**206 A problem with output overload has been reported by the uninterruptible power supply 2145 UPS-1U. The Overload Indicator on the 2145 UPS-1U front panel is illuminated red.**  
**Explanation**

A problem with output overload has been reported by the uninterruptible power supply 2145 UPS-1U. The Overload Indicator on the 2145 UPS-1U front panel is illuminated red.

**Action**

1. Ensure that only one 2145 is receiving power from the 2145 UPS-1U. Also ensure that no other devices are connected to the 2145 UPS-1U.
2. Disconnect the 2145 from the 2145 UPS-1U. If the Overload Indicator is still illuminated, on the disconnected 2145 replace the 2145 UPS-1U.
3. If the Overload Indicator is now off, and the node is a 2145-8F2, 2145-8F4, 2145-8G4 or 2145-8A4, on the disconnected 2145, with all outputs disconnected, in the sequence shown, exchange the FRUs for new FRUs. See "Possible Cause-FRUs or other" after the last action in this section.
4. If the Overload Indicator is now off, and the node is a 2145-CF8, on the disconnected 2145, with all outputs disconnected, determine whether it is one of the two power supplies or the power cable assembly that must be replaced. Plug just one power cable into the left hand power supply and start the node and see if the error is reported. Then shut down the node and connect the other power cable into the left hand power supply and start the node and see if the error is repeated. Then repeat the two tests for the right hand power supply. If the error is repeated for both cables on one power supply but not the other, replace the power supply that showed the error; otherwise, replace the power cable assembly.

Possible Cause-FRUs or other:

- 2145 power cable assembly (45%)
- 2145 power supply assembly (45%)
- 2145 UPS-1U assembly (10%)

**210 A problem has occurred in the uninterruptible power supply.**  
**Explanation**

No specific FRU has been identified.

**Action**

In the sequence shown, exchange the FRUs for new FRUs. See "Possible Cause-FRUs or other."

Possible Cause-FRUs or other:

- 2145 UPS electronics assembly (40%)
- 2145 UPS battery assembly (40%)
- 2145 UPS (20%)

Other:

- None.

**211 A problem has occurred in the uninterruptible power supply.  
Explanation**

A problem has occurred in the 2145 UPS-1U.

**Action**

In the sequence shown, exchange the FRUs for new FRUs. See "Possible Cause-FRUs or other."

Possible Cause-FRUs or other:

- 2145 UPS-1U assembly (100%)

Other:

- None.

**215 A problem has occurred with the uninterruptible power supply load.  
Explanation**

The 2145 detected that the current of the 2145 UPS exceeds the current that four 2145s require.

**Action**

1. Ensure also that only 2145s are receiving power from the 2145 UPS; that is, no switches or disk controllers are connected to the 2145 UPS.
2. If only one 2145 is connected to the 2145 UPS, exchange the FRU for a new FRU. See "Possible Cause-FRU or other." If more than one 2145 is connected to the 2145 UPS, disconnect the 2145s from the 2145 UPS and reconnect them one-at-a-time. While the problem persists, the nodes fail to start with boot error code 215 displayed on the 2145 front panel. When the first failure occurs, exchange the FRU for a new FRU. See "Possible Cause-FRU or other."

Possible Cause-FRUs or other:

- 2145 UPS electronics assembly (40%)
- 2145 power supply assembly (10%)
- Other: Configuration error (50%)

**216 A problem has occurred with the uninterruptible power supply load.  
Explanation**

The 2145 has detected that the 2145 UPS-1U current exceeds the current that one 2145 requires.

**Action**

Ensure that only one 2145 is receiving power from the 2145 UPS-1U; that is, no other devices are connected to the 2145 UPS-1U.

Possible Cause-FRUs or other:

- None.

**220 The uninterruptible power supply is receiving input power that might be unstable or in low voltage conditions.**

**Explanation**

The 2145 UPS is receiving input power that might be unstable or in low voltage conditions.

**Action**

Ask the customer to check the site power to the 2145 UPS providing power to this 2145. Check the connection, voltage and frequency. If the input power is sufficient, exchange the FRUs for new FRUs. (See "Possible Cause-FRUs or other.")

Possible Cause-FRUs or other:

- 2145 UPS input power cable (10%)
- 2145 UPS electronics assembly (10%)

Other:

- AC input power (80%)

**221 The uninterruptible power supply is receiving input power that might be unstable in low or high voltage conditions.**

**Explanation**

The 2145 UPS-1U is receiving input power that might be unstable in low or high voltage conditions.

**Action**

Ask the customer to check the site power to the 2145 UPS-1U providing power to this 2145. Check the connection, voltage, and frequency. If the input power is sufficient, exchange the FRUs for new FRUs. (See "Possible Cause-FRUs or other.")

Possible Cause-FRUs or other:

- 2145 UPS-1U input power cable (10%)
- 2145 UPS-1U assembly (10%)

Other:

- AC input power (80%)

**225 An incorrect type of uninterruptible power supply was installed.**

**Explanation**

The 2145 UPS installed is not compatible. If it is connected to a 2145 4F2, the incorrect model type is installed. The 2145 UPS cannot be used with other 2145 model types.

**Action**

Exchange the 2145 UPS for one of the correct type, or replace the 2145 UPS with a 2145 UPS-1U.

Possible Cause-FRUs or other:

- 2145 UPS (100%)

**226 An incorrect type of uninterruptible power supply was installed.**

**Explanation**

An incorrect type of 2145 UPS-1U was installed.

**Action**

Exchange the 2145 UPS-1U for one of the correct type.

Possible Cause-FRUs or other:

- 2145 UPS-1U (100%)

**230 An uninterruptible power supply is not configured 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 assemblies.

**Action**

Connect the cables correctly.

Possible Cause-FRUs or other:

- None.

Other:

- Cabling error (100%)

**231 An uninterruptible power supply is not configured 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.

**Action**

Connect the cables correctly.

Possible Cause-FRUs or other:

- None.

Other:

- Cabling error (100%)

**235 A 2145 is powered on, but the uninterruptible power supply has been instructed by another 2145 to power off because a loss of AC input power has occurred.**

**Explanation**

Although the AC input power has now returned, the 2145 still powers off. It then powers on again.

**Action**

Wait for the 2145 to power off.

Possible Cause-FRUs or other:

- None.

**236 A 2145 is powered on, but the uninterruptible power supply has been instructed by the 2145 to power off because a loss of AC input power has occurred.**

**Explanation**

Although the AC input power has now returned, the 2145 still powers off. It then powers on again.

**Action**

Wait for the 2145 to power off.

Possible Cause-FRUs or other:

- None.

**240 The ambient temperature threshold for the uninterruptible power supply has been exceeded.**

**Explanation**

The 2145 UPS shows a red warning light, and an alarm sounds. The 2145 UPS switches to bypass mode to lower the temperature.

**Action**

1. Turn off the 2145 UPS and unplug it from the power source.
2. Clear the vents and remove any heat sources.
3. Ensure that the air flow around the 2145 UPS is not restricted.
4. Wait at least five minutes, and then restart the 2145 UPS.
5. If the problem remains, exchange, in the sequence shown, the FRUs for new FRUs. (See "Possible Cause-FRUs or other.")

Possible Cause-FRUs or other:

- 2145 UPS electronics assembly (60%)
- 2145 UPS battery assembly (20%)
- 2145 UPS (20%)

## **241 The ambient temperature threshold for the uninterruptible power supply has been exceeded.**

### **Explanation**

The ambient temperature threshold for the 2145 UPS-1U has been exceeded.

### **Action**

1. Turn off the 2145 UPS-1U and unplug it from the power source.
2. Clear the vents and remove any heat sources.
3. Ensure that the air flow around the 2145 UPS-1U is not restricted.
4. Wait at least five minutes, and then restart the 2145 UPS-1U.
5. If the problem remains, exchange, in the sequence shown, the FRUs for new FRUs. (See "Possible Cause-FRUs or other.")

Possible Cause-FRUs or other:

- 2145 UPS-1U assembly (100%)

## **245 Repetitive node restarts due to uninterruptible power supply errors.**

### **Explanation**

Multiple node restarts have occurred due to 2145 UPS errors.

### **Action**

The uninterruptible power supply has been repeatedly restarted because 2145 UPS errors are being detected. Verify that the room temperature is within specified limits and that the input power is stable. Verify that the 2145 UPS signal cable is fastened securely at both ends.

The condition will be reset by powering off the node from the node front panel. If a reset does not fix the problem, replace FRUs in the order shown.

Possible Cause-FRUs or other:

- 2145 UPS electronics assembly (85%)
- 2145 UPS battery assembly (5%)
- 2145 UPS frame assembly (5%)
- Power/signal cable (5%)

## **246 Repetitive node restarts due to uninterruptible power supply errors.**

### **Explanation**

Multiple node restarts have occurred due to 2145 UPS-1U errors.

### **Action**

The uninterruptible power supply has been repeatedly restarted because 2145 UPS-1U errors are being detected. Verify that the room temperature is within specified limits and that the input power is stable. Verify that the 2145 UPS-1U signal cable is fastened securely at both ends.

The condition will be reset by powering off the node from the node front panel.



Possible Cause-FRUs or other:

- 2145 UPS-1U assembly (95%)
- Power/signal cable (5%)

#### **Related tasks**

“MAP 5900: Hardware boot” on page 423

MAP 5900: Hardware boot helps you solve problems that are preventing the node from starting its boot sequence.

## **250 The Ethernet ports did not start correctly.**

### **Explanation**

The Ethernet ports on the node did not start correctly. This is a transient error that only occurs during node start up. It does not indicate a hardware error.

### **Action**

Perform the following actions in order:

1. Power off the node using the front panel power control button.
2. Wait 30 seconds, then press the front power control button again to restart the node.
3. This error might occur again. If the error recurs, repeat the sequence up to four times. If this error still occurs after four repetitions of the step 1 and step 2 sequence, contact your Service Support Center.

Possible Cause-FRUs or other:

- None.

## **Performing the node rescue**

If it is necessary 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 SAN Volume Controller software.

Similarly, if you have replaced the service controller, you should use the node rescue procedure to ensure that the service controller has the correct software.

**Attention:** If you recently replaced both the service controller and the disk drive as part of the same repair operation, node rescue fails.

To provide an alternate boot device, a minimal operating system is also available in nonvolatile memory on the service controller. If it is necessary to replace the hard disk drive or the software on the hard disk drive has become corrupted, the node cannot boot and the hardware boot indicator remains on the front panel display or the boot operation does not progress. If this occurs, use the node rescue procedure to reinstall the SAN Volume Controller software.

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 running node rescue operations, only run one node rescue operation on the same SAN, at any one time. Wait for one node rescue operation to complete before starting another.

Perform the following steps to complete the node rescue:

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. This is particularly 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 72).



Figure 72. 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 two 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 cluster, the node is now offline. Delete the offline node from the cluster and then add the node back into the cluster. If node recovery was used to recover a node that failed during a software upgrade process, it is not possible to add the node back into the cluster until the upgrade or downgrade process has completed. This can take up to four hours for an eight-node cluster.

#### **Related tasks**

“Deleting a node from a cluster using the SAN Volume Controller Console” on page 69

You might have to remove a node from a cluster if the node has failed and is being replaced with a new node or if the repair that has been performed has caused that node to be unrecognizable by the cluster.

#### **Related information**

“Adding nodes to a cluster” on page 73

Before you add a node to a cluster, you must make sure that the switch zoning is configured such that the node being added is in the same zone as all other nodes in the cluster.

## **Understanding the node rescue codes**

The node rescue codes are displayed on the menu screen during node rescue.

Start node rescue if the boot image on the hard disk is missing or corrupted. Corrupted code is indicated during the boot process either by the display of an error code or by a hang condition.

To start node rescue, press and hold the **left** and **right** buttons on the front panel during a power-on cycle. The menu screen displays the Node rescue request. See the node rescue request topic. The hard disk is formatted and, if the format completes without error, the software image is downloaded from any available node. During node recovery, Line 1 of the menu screen displays the message **Booting** followed by one of the node rescue codes. Line 2 of the menu screen displays a **boot progress indicator**. Figure 73 shows an example of a displayed node rescue code.



Figure 73. Example of a node-rescue error code

The three-digit code that is shown in Figure 73 represents a node rescue code.

**Attention:** If the 2145 UPS is only connected to this SAN Volume Controller, the 2145 UPS turns off within five minutes of a node-rescue process failure. For example, if a donor node cannot be found. When the problem that is preventing node rescue has been resolved, the 2145 UPS must be turned on before turning on the SAN Volume Controller.

**Note:** The 2145 UPS-1U will not power off following a node rescue failure.

#### **Related concepts**

“Node rescue request” on page 141

If software is lost, you can use the node rescue process to copy all software from another node.

### **300 The 2145 is running node rescue.**

#### **Explanation**

The 2145 is running node rescue.

#### **Action**

If the progress bar has been stopped for at least two minutes, exchange the FRU for a new FRU. See “Possible Cause-FRUs or other”.

Possible Cause-FRUs or other:

2145-CF8

- Service controller (95%)
- Service controller cable (5%)

2145-8F2 or 2145-8F4 or 2145-8G4 or 2145-8A4

- Service controller (100%)

### **310 The 2145 is running a format operation.**

#### **Explanation**

The 2145 is running a format operation.

## Action

If the progress bar has been stopped for two minutes, exchange the FRU for a new FRU. See "Possible Cause-FRUs or other".

Possible Cause-FRUs or other:

### 2145-CF8

- Disk drive (50%)
- Disk controller (30%)
- Disk backplane (10%)
- Disk signal cable (8%)
- Disk power cable (1%)
- System board (1%)

### 2145-8G4 or 2145-8A4

- Disk drive assembly (90%)
- Disk cable assembly (10%)

### 2145-8F2 or 2145-8F4

- Disk drive assembly (100%)

### 2145-4F2

- Disk drive assembly (95%)
- Disk drive cables (5%)

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

### **Explanation**

A 2145 format operation has failed.

## Action

Exchange the FRU for a new FRU. See "Possible Cause-FRUs or other".

Possible Cause-FRUs or other:

### 2145-CF8

- Disk drive (50%)
- Disk controller (30%)
- Disk backplane (10%)
- Disk signal cable (8%)
- Disk power cable (1%)
- System board (1%)

### 2145-8G4 or 2145-8A4

- Disk drive assembly (90%)
- Disk cable assembly (10%)

### 2145-8F2 or 2145-8F4

- Disk drive assembly (95%)

- Frame assembly (5%)

2145-4F2

- Disk drive assembly (95%)
- Disk drive cables (5%)

### **330 The 2145 is partitioning its disk drive.**

#### **Explanation**

The 2145 is partitioning its disk drive.

#### **Action**

If the progress bar has been stopped for two minutes, exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

2145-CF8

- Disk drive (50%)
- Disk controller (30%)
- Disk backplane (10%)
- Disk signal cable (8%)
- Disk power cable (1%)
- System board (1%)

2145-8G4 or 2145-8A4

- Disk drive assembly (90%)
- Disk cable assembly (10%)

2145-8F2 or 2145-8F4

- Disk drive assembly (95%)
- Frame assembly (5%)

2145-4F2

- Disk drive assembly (95%)
- Disk drive cables (5%)

Other:

- Configuration problem
- Software error

### **340 The 2145 is searching for donor node.**

#### **Explanation**

The 2145 is searching for donor node.

#### **Action**

If the progress bar has been stopped for more than two minutes, exchange the FRU for a new FRU. See "Possible Cause-FRUs or other".

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.

#### **Action**

If the progress bar has stopped for more than two minutes, exchange the FRU for a new FRU. See "Possible Cause-FRUs or other".

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.

#### **Action**

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

#### **Related tasks**

"MAP 5600: Fibre-channel" on page 396

MAP 5600: Fibre-channel helps you to solve problems that have occurred on the SAN Volume Controller fibre-channel ports.

### **360 The 2145 is loading software from the donor.**

#### **Explanation**

The 2145 is loading software from the donor.

## Action

If the progress bar has been stopped for at least two minutes, restart the node rescue procedure.

Possible Cause-FRUs or other:

- None

## 370 The 2145 is installing software.

### Explanation

The 2145 is installing software.

### Action

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 upgrade operation continues.
4. Report this problem immediately to your Software Support Center.

Possible Cause-FRUs or other:

- None

## Understanding the node error codes

Node error codes are displayed on the display screen by node software.

Each code indicates that a critical error was detected that prevents the node from becoming a member of a cluster. Line 1 of the menu screen contains the message Node Error.

Line 2 contains either the error code or the error code and additional data. In errors that involve a node with more than one power supply, the error code is followed by two numbers. The first number indicates the power supply that has a problem (either a 1 or a 2). The second number indicates the problem that has been detected.

Figure 74 provides an example of a node error code. This data might exceed the maximum width of the menu screen. You can press the Right navigation to scroll the display.



Node Error:  
550 000125

Figure 74. Example of a displayed node error code

The additional data is unique for any error code. It provides necessary information that enables you to isolate the problem in an offline environment. Examples of additional data are disk serial numbers and field replaceable unit (FRU) location codes. When these codes are displayed, you can do additional fault isolation by navigating the default menu to determine the node and fibre-channel port status.

## **510 The detected memory size for this 2145 does not match the expected memory size for the cluster.**

### **Explanation**

The detected memory size, in MB, is the first number following the error code. The expected memory size for the cluster is the second number following the error code. This problem might have occurred because a memory module has failed or because failing memory modules were exchanged and the wrong size modules were installed.

### **Action**

Check the memory size of another 2145 that is in the same cluster. For the 2145-4F2, exchange the memory modules in this 2145-4F2 for modules of the correct size. For the 2145-8F2, 2145-8F4, 2145-8G4, 2145-8A4 and 2145-CF8, if you have just replaced a memory module, check that the module that you have installed is the correct size, then go to the light path MAP to isolate any possible failed memory modules.

Possible Cause-FRUs or other:

- Memory module (100%)

### **Related tasks**

“MAP 5800: Light path” on page 406

MAP 5800: Light path helps you to solve hardware problems on the SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2 that are preventing the node from booting.

## **511 Memory bank 1 of the 2145 is failing.**

### **Explanation**

Memory bank 1 of the 2145 is failing.

### **Action**

For the 2145-8F2, 2145-8F4, 2145-8G4 and 2145-8A4, go to the light path MAP to resolve this problem.

For the 2145-4F2, exchange both memory modules of bank 1 for new modules.

Possible Cause-FRUs or other:

- Memory module (100%)

### **Related tasks**

“MAP 5800: Light path” on page 406

MAP 5800: Light path helps you to solve hardware problems on the SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2 that are preventing the node from booting.

## **513 Memory bank 2 of the 2145 is failing.**

### **Explanation**

Memory bank 2 of the 2145 is failing.



## Action

For the 2145-8F2, 2145-8F4, 2145-8G4 and 2145-8A4, go to the light path MAP to resolve this problem.

For the 2145-4F2, exchange both memory modules of bank 2 for new modules.

Possible Cause-FRUs or other:

- Memory module (100%)

## Related tasks

“MAP 5800: Light path” on page 406

MAP 5800: Light path helps you to solve hardware problems on the SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2 that are preventing the node from booting.

## 514 Memory bank 3 of the 2145 is failing.

### Explanation

This cannot occur on the 2145-4F2.

## Action

For the 2145-8F2, 2145-8F4, 2145-8G4 and 2145-8A4, go to the light path MAP to resolve this problem.

Possible Cause-FRUs or other:

- Memory module (100%)

## Related tasks

“MAP 5800: Light path” on page 406

MAP 5800: Light path helps you to solve hardware problems on the SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2 that are preventing the node from booting.

## 515 Memory bank 4 of the 2145 is failing.

### Explanation

This cannot occur on the 2145-4F2.

## Action

For the 2145-8F2, 2145-8F4, 2145-8G4 and 2145-8A4, go to the light path MAP to resolve this problem.

Possible Cause-FRUs or other:

- Memory module (100%)

### Related tasks

“MAP 5800: Light path” on page 406

MAP 5800: Light path helps you to solve hardware problems on the SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2 that are preventing the node from booting.

### 520 A fibre-channel adapter is failing.

#### Explanation

A failure has been detected on a fibre-channel adapter.

If the failure has been isolated to one or more fibre-channel ports, the first failing fibre-channel adapter port is shown by the additional data for the displayed error code, for example 520 2.

For the 2145-4F2 and 2145-8F2, port number 1 or 2 indicates adapter one and port number 3 or 4 indicates adapter two. For the 2145-8F4, 2145-8G4 and 2145-CF8 there is only one adapter.

For the 2145-8F4, 2145-8G4 and 2145-CF8, if the failure has not been isolated to a port, no port number is shown after the error code.

#### Action

1. If the node has more than one fibre-channel adapter, determine which adapter is failing.
2. Ensure that the adapter is correctly seated in its socket.
3. Exchange the failing FRU for a new FRU.

Possible Cause-FRUs or other:

- Fibre-channel adapter assembly (100%)

Note: Although this error might be reported on a single port, the fibre-channel Small Form-factor Pluggable (SFP) connector is not a suitable replacement part.

### 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 indicates that the power supply is not detected. 2 indicates that the power supply has failed. 3 indicates that there is no input power to the power supply.

If the node is a member of a cluster, the cluster will report error code 1096 or 1097, depending on the error reason.

The error will automatically clear when the problem is fixed.

#### Action

1. Ensure that the power supply is seated correctly and that the power cable is attached correctly to both the node and to the 2145 UPS-1U.

2. If the error has not been automatically marked fixed after two minutes, note the status of the three LEDs on the back of the power supply. For the 2145-CF8, the AC LED is the top green LED, the DC LED is the middle green LED and the error LED is the bottom amber LED.
3. If the power supply error LED is off and the AC and DC power LEDs are both on, this is the normal condition. If the error has not been automatically fixed after two minutes, replace the system board.
4. Follow the action specified for the LED states noted in the table below.
5. If the error has not been automatically fixed after two minutes, contact support.

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 whether the 2145 UPS-1U is showing any errors. Follow MAP 5150 2145 UPS-1U if the UPS-1U 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:

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%)

### **540 The 2145 cannot connect to the Ethernet. Explanation**

This error is reported if either both Ethernet ports are offline, or Ethernet port 1 is offline and Ethernet port 2 is not configured as a management port.

## Action

Go to Ethernet MAP.

Possible Cause-FRUs or other:

2145-8G4 or 2145-8A4 or 2145-CF8

- Ethernet cable (60%)
- System board (5%)

2145-8F2 or 2145-8F4

- Ethernet cable (60%)
- Frame assembly (5%)

2145-4F2

- Ethernet cable (60%)
- System board assembly (5%)

Other:

- The Ethernet cable is disconnected (30%)
- Ethernet hub (5%)

## Related tasks

“MAP 5500: Ethernet” on page 393

MAP 5500: Ethernet helps you solve problems that have occurred on the SAN Volume Controller Ethernet.

## **550 Cannot form a cluster due to a lack of cluster resources. Explanation**

Supplemental data displayed with this error code list the missing IDs for the 2145s and the quorum disk controller. Each missing node is listed by its node ID. A missing quorum disk is listed as WWWWWWWWWWWWWWWWW/LL, where WWWWWWWWWWWWWWWWW is a worldwide port name (WWPN) on the disk controller that contains the missing quorum disk and LL is the Logical Unit Number (LUN) of the missing quorum disk on that controller.

## Action

1. Ensure that the other 2145s in the cluster are powered on and operational.
2. From the front panel, display the fibre-channel port status. If any port is not active, perform the fibre-channel port problem determination procedures.
3. Ensure that fibre-channel network zoning changes have not restricted communication between nodes, or between the nodes and the quorum disk.
4. Do the problem determination procedures for the network.
5. The quorum disk failed or cannot be accessed. Perform the problem determination procedures for the disk controller.

Possible Cause-FRUs or other:

- None

## **555 Power Domain error.**

### **Explanation**

Both 2145s in an I/O group are being powered by the same uninterruptible power supply. The other 2145's ID is displayed with the node error code on the front panel.

### **Action**

Ensure that the configuration is correct and that each 2145 in an I/O group is connected from a separate uninterruptible power supply.

Possible Cause-FRUs or other:

- None

Other:

- Configuration problem.

## **556 Duplicate WWNN detected.**

### **Explanation**

The node has detected another device on the fibre-channel network with the same World Wide Node Name (WWNN) as itself. The last five digits of the node's WWNN are shown as the additional data for the error.

The node is not an active member of a cluster. It has disabled its fibre-channel ports to prevent disrupting the operation of the fibre-channel network.

Either just one node, or both nodes, with the same WWNN will show the error.

Due to the way WWNNs are allocated, the device with a duplicate WWNN should be another SAN Volume Controller node.

Only the last five digits of the WWNN are shown on the front panel displays. A WWNN is 16 hexadecimal digits long. For a SAN Volume Controller the first 11 digits are always 50050768010.

### **Action**

1. Find the SAN Volume Controller node with the same WWNN as the node reporting the error. The WWNN for a SAN Volume Controller node can be found from the node Vital Product Data (VPD) or from the Node menu on the front panel. The node with the duplicate WWNN need not be part of the same cluster as the node reporting the error; it could be remote from the node reporting the error on a part of the fabric connected through an inter-switch link. The node's WWNN is stored with the service controller, so the duplication is most likely caused by a node that has recently had its service controller replaced and/or its WWNN changed.
2. If a SAN Volume Controller node with a duplicate WWNN is found, determine whether it, or the node reporting the error, has the incorrect WWNN. Normally it is the node that has had its service controller replaced, and/or its WWNN changed, that is incorrect. Also consider how the SAN is zoned when making your decision.
3. If both nodes with the same WWNN reported the error, the node with the correct WWNN can be restarted using the front panel power control button.

4. Determine the correct WWNN for the node with the incorrect WWNN. If the complete node or the service controller has been replaced as part of a service action, the WWNN for the node should have been written down. If the correct WWNN cannot be determined contact your support center for assistance.
5. Use the front panel menus 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 VDisks that it is managing will be lost. You should ensure that the host systems are in the correct state before you change the WWNN.
6. 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.
7. If you are unable to find a SAN Volume Controller node with the same WWNN as the node showing the error, use the SAN monitoring tools to determine if there is another device on the SAN with the same WWNN. This device should not be using a WWNN assigned to a SAN Volume Controller, so you should follow the service procedures for the device to change its WWNN.

Possible Cause-FRUs or other:

- None.

#### **Related concepts**

“Node WWNN” on page 151

The node WWNN (worldwide node name) option displays the last five hexadecimal digits of the WWNN that is being used by the SAN Volume Controller node. Only the last five digits of a WWNN vary on a SAN Volume Controller node. The first 11 digits are always 50050768010.

“Validate WWNN?” on page 143

The front panel prompts you to validate the WWNN when the worldwide node name (WWNN) that is stored in the service controller (the panel WWNN) does not match the WWNN that is backed up on the SAN Volume Controller disk (the disk WWNN).

### **558 The 2145 cannot see the fibre-channel fabric or the fibre-channel card port speed might be set to a different speed than the fibre channel fabric.**

#### **Explanation**

The 2145 cannot see the fibre-channel fabric or the fibre-channel card port speed might be set to a different speed than the fibre channel fabric.

#### **Action**

Ensure that:

1. The fibre-channel network fabric switch is powered-on.
2. At least one fibre-channel cable connects the 2145 to the fibre-channel network fabric.
3. The fibre-channel card port speed is equal to the fibre-channel fabric.
4. At least one fibre-channel adapter is installed in the 2145.
5. Go to the Fibre-channel MAP.

Possible Cause-FRUs or other:

- None

**Related tasks**

“MAP 5600: Fibre-channel” on page 396

MAP 5600: Fibre-channel helps you to solve problems that have occurred on the SAN Volume Controller fibre-channel ports.

**Related reference**

“Changing the fibre-channel network speed” on page 104

Depending on the model, you can change the speed of the fibre-channel ports on a SAN Volume Controller through the front panel on the SAN Volume Controller node or by a command sent to a SAN Volume Controller cluster using the Ethernet interface.

**560 The fibre-channel network fabric is too big.****Explanation**

The configuration is not valid.

**Action**

1. Ensure that all the fibre-channel connections are correct.
2. Restart the 2145.

Possible Cause-FRUs or other:

- None

Other:

- See your product’s configuration guide

**562 The hardware configuration is not valid.****Explanation**

This error has probably been caused by a service action error when replacing FRUs.

**Action**

1. Ensure that the 2145 hardware is correct.
2. Restart the 2145.

Possible Cause-FRUs or other:

- None

**564 This 2145 node is repeatedly crashing because of a software failure.****Explanation**

Software dump data is available in the Dumps directory.

**Action**

- Call your software support center for assistance.

Possible Cause-FRUs or other:

- None

Other:

- Software error.

### **Related tasks**

“Deleting a node from a cluster using the SAN Volume Controller Console” on page 69

You might have to remove a node from a cluster if the node has failed and is being replaced with a new node or if the repair that has been performed has caused that node to be unrecognizable by the cluster.

### **Related information**

“Adding nodes to a cluster” on page 73

Before you add a node to a cluster, you must make sure that the switch zoning is configured such that the node being added is in the same zone as all other nodes in the cluster.

## **565 The node’s internal disk is failing.**

### **Explanation**

An excessive number of errors has been reported by the node’s internal disk drive. It is no longer safe to rely on the integrity of the disk drive.

The node has removed itself from the cluster.

### **Action**

Perform the following actions in order:

1. Replace the node’s disk drive assembly.
2. Restart the node using the node rescue procedure.
3. Delete the node from the cluster, then add the node back into the cluster.
4. This error will not recur immediately when the node is restarted. If the error has occurred previously on this node and the disk drive has been replaced, replace FRUs in the order shown.

FRU list:

2145-CF8

- Disk drive (50%)
- Disk controller (30%)
- Disk backplane (10%)
- Disk signal cable (8%)
- Disk power cable (1%)
- System board (1%)

2145-8G4 or 2145-8A4

- Disk drive assembly (95%)
- Disk drive cables (5%)

2145-8F2 or 2145-8F4

- Disk drive assembly (100%)

2145-4F2

- Disk drive assembly (95%)
- SCSI signal cable (5%)



## **570 The 2145 data is readable, but damaged.**

### **Explanation**

The node has restarted but it is not able to participate in a 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 570 or 578 when it is in this state.

If this node was a member of a cluster, the cluster shows the node as offline while the node is in this state and reports an error code 1195 "Node missing". A node that is showing node error 570 or 578 will appear as a candidate node for addition to the cluster if the node's model is supported.

There are circumstances when a cluster automatically attempts to add a node back into the cluster. The cluster checks for any offline nodes every 15 minutes. If the cluster discovers that a node is offline, the cluster searches for candidate nodes that match the characteristics of the offline node. If a matching candidate node is found, the cluster attempts to delete the offline node from the cluster and add the matching candidate node into the cluster. The characteristics that are used to determine a matching candidate node are the WWNN and the front panel name. Both of these characteristics must match to determine that the candidate node is a match for the offline node.

The cluster logs event 980352 "Attempt to automatically recover offline node starting" when it begins to automatically add a node into the cluster. If the attempt succeeds, an event 980349 "Node added" is logged to indicate that a node was added to the cluster. If the attempt fails, an event is not logged and if the node is still a matching candidate node after 15 minutes the cluster will start another attempt to recover the node and log another event 980352.

If three attempts to automatically add a matching candidate node to a cluster have been made, but the node has not returned online and remained offline for 24 hours, the cluster stops automatic attempts to add the node and logs error code 1194 "Automatic recovery of offline node failed".

### **Action**

1. If the node is offline to the cluster, delete the offline node from the cluster and then add the node back into the cluster.
2. Determine whether this error has occurred because a power was removed from this node in a manner that prevented saving all of the state data to its internal disk. If you have determined the cause of the error, continue with MAP 5700: Repair verification.
3. If the cluster is now showing error 1194 "Automatic recovery of offline node failed" or if you cannot identify an action or event that caused a temporary loss of power to the node, ensure that the input power cable is securely connected to both the node and to the uninterruptible power supply. You must manually delete the offline node from the cluster and add the candidate node into the cluster.
4. If you still cannot determine the cause of the temporary loss of power to the node, call IBM technical support for assistance. You might be asked to recover the dump and trace data records to help analyze the root cause of the problem. If the problem persists, exchange the FRUs with new FRUs in the sequence shown.

Possible Cause-FRUs or other:

#### 2145-CF8

- disk drive (5%)
- Disk controller (1%)
- Disk backplane (1%)
- Disk signal cable (1%)
- Disk power cable (1%)
- System board (1%)

#### 2145-8A4

- Disk drive assembly (6%)
- Disk cable assembly (3%)
- System board (1%)

#### 2145-8G4

- Disk drive assembly (6%)
- Disk drive cables (3%)
- System board (1%)

#### 2145-8F2 or 2145-8F4

- Disk drive assembly (9%)
- Frame assembly (1%)

#### 2145-4F2

- Disk drive assembly (6%)
- Disk drive cables (3%)
- System board assembly (1%)

Other:

- Incorrect power off (80%)
- Software problem (10%)

#### **Related tasks**

“Performing the node rescue” on page 217

If it is necessary 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 SAN Volume Controller software.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **572 The 2145 cannot determine the VPD for a FRU.**

#### **Explanation**

A FRU in the 2145 has been changed, and the VPD is unreadable or unrecognized.

#### **Action**

1. Update the 2145 software to the latest level.
2. Exchange the most-recently replaced FRU for a new FRU.

Possible Cause-FRUs or other:

- None

Other:

- Software problem.

### **574 The 2145 software on this node is damaged.**

#### **Explanation**

The 2145 software on this node is damaged. Recovery is required.

If this is the only node with this problem and if you can still access the data on the virtual disks (VDisks), perform the following actions. If more than one node has this problem or if you cannot access the data on the virtual disks, call your support center for assistance.

#### **Action**

- Perform the node rescue procedure.

Possible Cause-FRUs or other:

- None

Other:

- Software problem.

#### **Related tasks**

“Performing the node rescue” on page 217

If it is necessary 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 SAN Volume Controller software.

### **576 The 2145 data cannot be read.**

#### **Explanation**

The 2145 has been rejected from the cluster.

#### **Action**

In the sequence shown, exchange the FRUs for new FRUs.

Possible Cause-FRUs or other:

2145-CF8

- disk drive (50%)
- Disk controller (30%)
- Disk backplane (10%)
- Disk signal cable (8%)
- Disk power cable (1%)
- System board (1%)

2145-8A4

- Disk drive assembly (80%)
- Disk cable assembly (15%)

- System board (5%)

#### 2145-8G4

- Disk drive assembly (80%)
- Disk drive cables (10%)
- System board (10%)

#### 2145-8F2 or 2145-8F4

- Disk drive assembly (90%)
- Frame assembly (10%)

#### 2145-4F2

- Disk drive assembly (45%)
- System board assembly (50%)
- Disk drive cables (5%)

### **578 This node has lost power without saving its current state data to its internal disk.**

#### **Explanation**

The node has shut down without saving its current state to its internal disk. The node has restarted but it is not able to participate in the cluster.

The usual cause of this error is a user action that has removed power from the node in a manner that does not allow the node to save its state data to its internal disk. Actions that might cause this error include:

- Removing the node power cable from the node or uninterruptible power supply.
- Pressing the power button on the node for more than four seconds. This action causes an immediate power off.
- Powering off the uninterruptible power supply while the node is active or while the node is in the process of shutting down.

A software error or hardware error might also cause the node to power off without saving the state data. If the cause is a hardware error, the hardware problem is intermittent because the node has restarted. In this case, the node message and trace files might indicate the cause of the power off and your support center might request these files for analysis.

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 570 or 578 when it is in this state.

If this node was a member of a cluster, the cluster shows the node as offline while the node is in this state and reports an error code 1195 "Node missing". A node that is showing node error 570 or 578 will appear as a candidate node for addition to the cluster if the node's model is supported.

There are circumstances when a cluster automatically attempts to add a node back into the cluster. The cluster checks for any offline nodes every 15 minutes. If the cluster discovers that a node is offline, the cluster searches for candidate nodes that match the characteristics of the offline node. If a matching candidate node is found, the cluster attempts to delete the offline node from the cluster and add the

matching candidate node into the cluster. The characteristics that are used to determine a matching candidate node are the WWNN and the front panel name. Both of these characteristics must match to determine that the candidate node is a match for the offline node.

The cluster logs event 980352 "Attempt to automatically recover offline node starting" when it begins to automatically add a node into the cluster. If the attempt succeeds, an event 980349 "Node added" is logged to indicate that a node was added to the cluster. If the attempt fails, an event is not logged and if the node is still a matching candidate node after 15 minutes the cluster will start another attempt to recover the node and log another event 980352.

If three attempts to automatically add a matching candidate node to a cluster have been made, but the node has not returned online and remained offline for 24 hours, the cluster stops automatic attempts to add the node and logs error code 1194 "Automatic recovery of offline node failed".

### **Action**

1. If the node is offline to the cluster, delete the offline node from the cluster and then add the node back into the cluster.
2. Determine whether this error has occurred because power was removed from this node in a manner that prevented saving all of the state data to its internal disk. If the cause of the error has been determined, continue with MAP 5700: Repair verification.
3. If the cluster is now showing error 1194 "Automatic recovery of offline node failed" or if you cannot identify an action or event that caused a temporary loss of power to the node, ensure that the input power cable is securely connected to both the node and to the uninterruptible power supply.
4. If you still cannot determine the cause of the temporary loss of power to the node, call IBM technical support for assistance. You might be asked to recover the dump and trace data records to help analyze the root cause of the problem.

Possible Cause-FRUs or other:

Your support center might indicate a FRU based on their problem analysis (2%)

Other:

- User action (98%)

## **580 Unable to read service controller ID.**

### **Explanation**

The 2145 cannot read the unique ID from the service controller, so the fibre-channel adapters cannot be started.

### **Action**

In the sequence shown, exchange the following FRUs for new FRUs. See "Possible Cause-FRUs or other".

Possible Cause-FRUs or other:

2145-CF8

- Service controller (70%)

- Service controller cable (30%)

2145-8F2 or 2145-8F4 or 2145-8G4 or 2145-8A4

Service controller (100%)

2145-4F2

- Front panel assembly (50%)
- Service controller (50%)

Other:

- None

## Understanding the create cluster error codes

Cluster Create error codes are displayed on the menu screen when you are using the front panel to create a new cluster, but the create operation fails.

Figure 75 provides an example of a create-cluster error code.



The image shows a black rectangular box with white text. The text is arranged in two lines: the first line reads "Create Failed:" and the second line reads "700".

Figure 75. Example of a create-cluster error code

Line 1 of the menu screen contains the message Create Failed. Line 2 shows the error code and, where necessary, additional data.

### Related reference

“Create Cluster?” on page 151

You can create a cluster from the Create Cluster? menu. You also have the option to delete the node from the cluster if the node that is selected is in a cluster.

### **700 All the available unique identifications have been used.** **Explanation**

Each time a node creates a new cluster a unique ID is generated by the node’s Service Controller. Once 255 clusters have been created the Service Controller must be replaced.

### **Action**

Use a different node to create the cluster.

Possible Cause-FRUs or other:

- Service controller (100%)

### **710 The service controller cannot increase the cluster ID counter.** **Explanation**

When a new cluster ID is requested from the service controller, the service controller is told to increase the ID counter. The new ID is read back for verification. If the ID counter has not been increased, this error code is displayed. This error has

occurred because the service controller failed.

### Action

Exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

- Service controller (100%)

## Understanding the cluster recovery codes

Cluster recovery codes indicate that a critical software error has occurred that might corrupt your cluster.

You must perform software problem analysis before you can perform further cluster operations to avoid the possibility of corrupting your cluster configuration.

Figure 76 provides an example of a cluster-recovery error code.

A screenshot of a terminal window with a black background and white text. The text reads "Node Error:" on the first line and "901" on the second line.

Figure 76. Example of a cluster-recovery error code

### 9xx Explanation

A cluster recovery action is required.

### Action

Contact the software support center for assistance in performing the software problem analysis.

Possible Cause-FRUs or other:

- None

## Understanding cluster error codes

Every cluster error code includes an error code number, a description, action, and possible field-replaceable units (FRUs).

Figure 77 provides an example of a cluster error code.

A screenshot of a terminal window with a dark background and light-colored text. The text reads "Cluster Error:" on the first line and "1001" on the second line. To the right of the terminal window, the text "svc00433" is visible vertically.

Figure 77. Example of a cluster error code

### 1001 Automatic cluster recovery has run. Explanation

All cluster configuration commands are blocked.

## Action

Call your software support center.

Caution:

You can unblock the configuration commands by issuing the `svctask enablecli` command, but you must first consult with your software support to avoid corrupting your cluster configuration.

Possible Cause-FRUs or other:

- None

## 1002 Error log full.

### Explanation

Error log full.

## Action

To fix the errors in the error log, go to the start MAP.

Possible Cause-FRUs or other:

- Unfixed errors in the log.

### Related tasks

"MAP 5000: Start" on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

## 1011 Fibre-channel adapter (4 port) in slot 1 is missing.

### Explanation

Fibre-channel adapter (4 port) in slot 1 is missing.

## Action

1. In the sequence shown, exchange the FRUs for new FRUs. See "Possible Cause-FRUs or other" after the last action in this section.
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 with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-CF8

- 4-port fibre-channel host bus adapter (98%)
- System board (2%)

2145-8G4 or 2145-8A4

- 4-port fibre-channel host bus adapter (90%)
- PCI Express riser card (8%)
- System board (2%)



2145-8F4

N/A

2145-8F2

N/A

2145-4F2

N/A

### **Related tasks**

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## **1013 Fibre-channel adapter (4-port) in slot 1 PCI fault.**

### **Explanation**

Fibre-channel adapter (4-port) in slot 1 PCI fault.

### **Action**

1. In the sequence shown, exchange the FRUs for new FRUs. See “Possible Cause-FRUs or other” after the last action in this section.
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 with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-CF8

- 4-port fibre-channel host bus adapter (98%)
- System board (2%)

2145-8G4 or 2145-8A4

- 4-port fibre-channel host bus adapter (80%)
- PCI Express riser card (10%)
- System board (10%)

2145-8F4

N/A

2145-8F2

N/A

2145-4F2

N/A

**Related tasks**

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

**1014 Fibre-channel adapter in slot 1 is missing.**

**Explanation**

Fibre-channel adapter in slot 1 is missing.

**Action**

1. In the sequence shown, exchange the FRUs for new FRUs. See “Possible Cause-FRUs or other” after the last action in this section.
2. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F2

- Dual port fibre-channel HBA - low profile (90%)
- PCI riser card - low profile (8%)
- Frame assembly (2%)

2145-8G4

N/A

2145-8F4

N/A

2145-4F2

N/A

### Related tasks

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## 1015 Fibre-channel adapter in slot 2 is missing.

### Explanation

Fibre-channel adapter in slot 2 is missing.

### Action

1. In the sequence shown, exchange the FRUs for new FRUs. See “Possible Cause-FRUs or other” after the last action in this section.
2. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F2

- Dual port fibre-channel host bus adapter - full height (90%)
- PCI riser card (8%)
- Frame assembly (2%)

2145-8G4

N/A

2145-8F4

N/A

2145-4F2

N/A

### Related tasks

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## 1016 Fibre-channel adapter (4 port) in slot 2 is missing.

### Explanation

Fibre-channel adapter (4 port) in slot 2 is missing.

### Action

1. In the sequence shown, exchange the FRUs for new FRUs. See “Possible Cause-FRUs or other” after the last action in this section.
2. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F4

- 4-port fibre-channel host bus adapter (90%)
- PCI Express riser card (8%)
- Frame assembly (2%)

2145-8G4

N/A

2145-8F2

N/A

2145-4F2

N/A

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **Related reference**

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1017 Fibre-channel adapter in slot 1 PCI bus error.**

### **Explanation**

Fibre-channel adapter in slot 1 PCI bus error.

### **Action**

1. In the sequence shown, exchange the FRUs for new FRUs. See “Possible Cause-FRUs or other” after the last action in this section.
2. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F2

- Dual port fibre-channel host bus adapter - low profile (80%)
- PCI riser card (10%)
- Frame assembly (10%)

2145-8G4

N/A

2145-8F4

N/A

2145-4F2

N/A

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **Related reference**

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1018 Fibre-channel adapter in slot 2 PCI fault.**

### **Explanation**

Fibre-channel adapter in slot 2 PCI fault.

### **Action**

1. In the sequence shown, exchange the FRUs for new FRUs. See “Possible Cause-FRUs or other” after the last action in this section.
2. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F2

- Dual port fibre-channel host bus adapter - full height (80%)
- PCI riser card (10%)
- Frame assembly (10%)

2145-8G4

N/A

2145-8F4

N/A

2145-4F2

N/A

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **Related reference**

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1019 Fibre-channel adapter (4-port) in slot 2 PCI fault. Explanation**

Fibre-channel adapter (4-port) in slot 2 PCI fault.

### **Action**

1. In the sequence shown, exchange the FRUs for new FRUs. See “Possible Cause-FRUs or other” after the last action in this section.
2. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F4

- 4-port fibre-channel host bus adapter (80%)
- PCI Express riser card (10%)
- Frame assembly (10%)

2145-8G4

N/A

2145-8F2

N/A

2145-4F2

N/A

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **Related reference**

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1025 The 2145 system assembly is failing.**

### **Explanation**

The 2145 system assembly is failing.

### **Action**

1. Go to the light path diagnostic MAP and perform the light path diagnostic procedures.
2. If the light path diagnostic procedure isolates the FRU, mark this error as “fixed” and go to the repair verification MAP. If you have just replaced a FRU but it has not corrected the problem, ensure that the FRU is installed correctly and go to the next step.
3. Replace the system board or frame assembly as indicated in the Possible Cause list below.
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 the start MAP. If you return to this step, contact your support center to resolve the problem with the 2145.
5. Go to the repair verification MAP.

Possible Cause-FRUs or other:

2145-8G4 or 2145-CF8

- The FRUs that are indicated by the Light path diagnostics (98%)
- System board (2%)

2145-8F2 or 2145-8F4

- The FRUs that are indicated by the Light path diagnostics (98%)
- Frame assembly (2%)



2145-4F2

N/A

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

“MAP 5800: Light path” on page 406

MAP 5800: Light path helps you to solve hardware problems on the SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2 that are preventing the node from booting.

### **Related reference**

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1030 A node's internal disk 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.

### **Action**

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:

#### **2145-CF8**

- disk drive (50%)
- Disk controller (30%)
- Disk backplane (10%)
- Disk signal cable (8%)
- Disk power cable (1%)
- System board (1%)

#### **2145-8A4**

- disk drive (90%)

- disk cable assembly (10%)

2145-8G4

- disk drive assembly (90%)
- disk drive cables (10%)

2145-8F4 or 2145-8F2

- disk drive assembly (100%)

2145-4F2

- disk drive assembly (90%)
- disk drive cables (10%)

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

## **1040 A flash module error has occurred after a successful start of a 2145.**

### **Explanation**

Note: The node containing the flash module has not been rejected by the cluster.

### **Action**

1. replace the FRUs below in the order listed
2. Check node status. If all nodes show a status of Online, mark the error that you have just repaired "fixed". If any nodes do not show a status of Online, go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-CF8

- Service controller (50%)
- Service controller cable (50%)

2145-8F2 or 2145-8F4 or 2145-8G4 or 2145-8A4

Service controller (100%)

2145-4F2

Service controller (100%)

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## 1044 A service controller read failure occurred.

### Explanation

A service controller read failure occurred.

### Action

1. Replace the FRUs below in the order listed.
2. Check node status. If all nodes show a status of Online, mark the error that you have just repaired "fixed". If any nodes do not show a status of Online, go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-CF8

- Service controller (50%)
- Service controller cable (50%)

2145-8F2 or 2145-8F4 or 2145-8G4 or 2145-8A4

Service controller (100%)

2145-4F2

Service controller (100%)

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1054 Fibre-channel adapter in slot 1 adapter present but failed.**

### **Explanation**

Fibre-channel adapter in slot 1 adapter present but failed.

### **Action**

1. Replace the fibre-channel adapter.
2. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F2

Dual port fibre-channel host bus adapter - low profile (100%)

2145-8G4

N/A

2145-8F4

N/A

2145-4F2

N/A

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **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.

### **Action**

1. Exchange the FRU for new FRU. See “Possible Cause-FRUs or other” after the last action in this section.
2. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8G4 or 2145-8A4 or 2145-CF8

- 4-port fibre-channel host bus adapter (100%)

2145-8F4

N/A

2145-8F2

N/A

2145-4F2

N/A

## **1056 Fibre-channel adapter in slot 2 adapter present but failed.**

### **Explanation**

Fibre-channel adapter in slot 2 adapter present but failed.

### **Action**

1. Replace the fibre-channel adapter.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F2

Dual port fibre-channel host bus adapter - full height (100%)

2145-8G4

N/A

2145-8F4

N/A

2145-4F2

N/A

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1057 Fibre-channel adapter (4 port) in slot 2 adapter present but failed.**

### **Explanation**

Fibre-channel adapter (4 port) in slot 2 adapter present but failed.

### **Action**

1. Exchange the FRU for new FRU. See “Possible Cause-FRUs or other” after the last action in this section.
2. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F4

- 4-port fibre-channel host bus adapter (100%)

2145-8G4

N/A

2145-8F2

N/A

2145-4F2

N/A

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1060 One or more fibre-channel ports on the 2145 are not operational.**

### **Explanation**

One or more fibre-channel ports on the 2145 are not operational.

### **Action**

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:

2145-8F4 or 2145-8G4 or 2145-CF8

- Fibre-channel cable (80%)
- Small Form-factor Pluggable (SFP) connector (5%)
- 4-port fibre-channel host bus adapter (5%)

2145-8F2

- Fibre-channel cable (80%)
- Small Form-factor Pluggable (SFP) connector (5%)
- Dual port fibre-channel host bus adapter (fibre-channel MAP isolates to the correct type) (5%)

2145-4F2

- Fibre-channel cable (80%)
- Small Form-factor Pluggable (SFP) connector (5%)
- Fibre-channel adapter port (5%)

Other:



- Fibre-channel network fabric (10%)

#### **Related tasks**

“MAP 5600: Fibre-channel” on page 396

MAP 5600: Fibre-channel helps you to solve problems that have occurred on the SAN Volume Controller fibre-channel ports.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **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.

#### **Action**

- Go to MAP 5600: Fibre Channel to resolve the problem.

Possible Cause-FRUs or other:

2145-8F4 or 2145-8G4 or 2145-8A4 or 2145-CF8

- Fibre-channel cable (50%)
- Small Form-factor Pluggable (SFP) connector (20%)
- 4-port fibre-channel host bus adapter (5%)

Other:

- Fibre-channel switch, SFP or GBIC (25%)

#### **Related tasks**

“MAP 5600: Fibre-channel” on page 396

MAP 5600: Fibre-channel helps you to solve problems that have occurred on the SAN Volume Controller fibre-channel ports.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **1089 One or more fans are failing.**

#### **Explanation**

One or more fans are failing.

#### **Action**

1. Determine the failing fan(s) from the fan indicator on the system board or from the text of the error data in the log. The reported fan for the 2145-8A4 or the 2145-CF8 matches the fan assembly position. For the 2145-8G4, if you have determined the failing fan number from the error data in the log, use the following list to determine the position of the fan assembly to replace. Each fan assembly contains two fans.

2. Exchange the FRU for a new FRU. See "Possible Cause-FRUs or other" after the last action in this section.
3. Go to repair verification MAP.
  - Fan number:Fan assembly 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:

2145-8G4 or 2145-8A4 or 2145-CF8

- Fan assembly (100%)

2145-4F2 or 2145-4F2 or 2145-8F4

N/A

#### **Related tasks**

"MAP 5700: Repair verification" on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **1090 One or more fans (40×40×28) are failing.**

#### **Explanation**

One or more fans (40×40×28) are failing.

#### **Action**

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.
2. If all fans on the fan backplane are failing or if no fan fault lights are illuminated, verify that the cable between the fan backplane and the system board is connected.
3. Exchange the FRU for a new FRU. See "Possible Cause-FRUs or other" after the last action in this section.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F2 or 2145-8F4

- Fan 40×40×28 (98%)
- Fan power cable assembly (2%)

2145-4F2 or 2145-8G4

N/A

### Related tasks

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## **1091 One or more fans (40×40×56) are failing.**

### Explanation

One or more fans (40×40×56) are failing.

### Action

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.
2. If all fans on the fan backplane are failing or if no fan fault lights are illuminated, verify that the cable between the fan backplane and the system board is connected.
3. Exchange the FRU for a new FRU. See “Possible Cause-FRUs or other” after the last action in this section.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F2 or 2145-8F4

- Fan 40×40×56 (98%)
- Fan power cable assembly (2%)

2145-4F2 or 2145-8G4

N/A

### Related tasks

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## **1092 The temperature soft or hard shutdown threshold of the 2145 has been exceeded. The 2145 has automatically powered off.**

### Explanation

The temperature soft or hard shutdown threshold of the 2145 has been exceeded. The 2145 has automatically powered off.

### Action

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 with the 2145.

6. Go to the repair verification MAP.

Possible Cause-FRUs or other:

2145-8G4 or 2145-8A4 or 2145-CF8

- The FRU that is indicated by the Light path diagnostics (25%)
- System board (5%)

2145-8F2 or 2145-8F4

- The FRU that is indicated by the Light path diagnostics (25%)
- Frame assembly (5%)

2145-4F2

N/A

Other:

System environment or airflow blockage (70%)

#### **Related tasks**

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

#### **Related reference**

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

### **1093 The internal temperature sensor of the 2145 has reported that the temperature warning threshold has been exceeded.**

#### **Explanation**

The internal temperature sensor of the 2145 has reported that the temperature warning threshold has been exceeded.

#### **Action**

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 with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8G4 or 2145-8A4 or 2145-CF8

- Fan assembly (25%)

- System board (5%)

2145-8F2 or 2145-8F4

- Fan assembly (25%)
- Frame assembly (5%)

2145-4F2

N/A

Other:

Airflow blockage (70%)

#### **Related tasks**

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

#### **Related reference**

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

### **1094 The ambient temperature threshold has been exceeded. Explanation**

The ambient temperature threshold has been exceeded.

#### **Action**

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%)

## Related tasks

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## 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.

### Action

Error code 1096 is reported when the power supply either cannot be detected or reports an error.

1. Ensure that the power supply is seated correctly and that the power cable is attached correctly to both the node and to the 2145 UPS-1U.
2. If the error has not been automatically marked fixed after two minutes, note the status of the three LEDs on the back of the power supply. For the 2145-CF8, the AC LED is the top green LED, the DC LED is the middle green LED and the error LED is the bottom amber LED.
3. If the power supply error LED is off and the AC and DC power LEDs are both on, this is the normal condition. If the error has not been automatically fixed after two minutes, replace the system board.
4. Follow the action specified for the LED states noted in the table below.
5. If the error has not been automatically fixed after two minutes, contact support.
6. Go to repair verification MAP.

Error,AC,DC>Action

ON,ON or OFF,ON or OFF:The power supply has a fault. Replace the power supply.

OFF,OFF,OFF:There is no power detected. Ensure that the power cable is connected at the node and 2145 UPS-1U. If the AC LED does not light, check the status of the 2145 UPS-1U to which the power supply is connected. Follow MAP 5150 2145 UPS-1U if the UPS-1U is showing no power or an error; otherwise, replace the power cable. If the AC LED still does not light, replace the power supply.

OFF,OFF,ON:The power supply has a fault. Replace the power supply.

OFF,ON,OFF:Ensure that the power supply is installed correctly. If the DC LED does not light, replace the power supply.

Possible Cause-FRUs or other:

Failed PSU:

- Power supply (90%)
- Power cable assembly (5%)

- System board (5%)

Missing PSU:

- Power supply (19%)
- System board (1%)
- Other: Power supply not correctly installed (80%)

#### **Related tasks**

“MAP 5150: 2145 UPS-1U” on page 376

MAP 5150: 2145 UPS-1U helps you solve problems that have occurred in the 2145 UPS-1U systems that are used on a SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **1097 A Power Supply Unit reports no A/C power.**

#### **Explanation**

One of the two power supply units in the node is reporting that no main power is detected.

#### **Action**

1. Ensure that the power supply is attached correctly to both the node and to the 2145 UPS-1U.
2. If the error has not been automatically marked fixed after two minutes, note the status of the three LEDs on the back of the power supply. For the 2145-CF8, the AC LED is the top green LED, the DC LED is the middle green LED and the error LED is the bottom amber LED.
3. If the power supply error LED is off and the AC and DC power LEDs are both on, this is the normal condition. If the error has not been automatically fixed after two minutes, replace the system board.
4. Follow the action specified for the LED states noted in the table below.
5. If the error has not been automatically fixed after two minutes, contact support.
6. Go to repair verification MAP.

Error,AC,DC>Action

ON,ON or OFF,ON or OFF:The power supply has a fault. Replace the power supply.

OFF,OFF,OFF:There is no power detected. Ensure that the power cable is connected at the node and 2145 UPS-1U. If the AC LED does not light, check whether the 2145 UPS-1U is showing any errors. Follow MAP 5150 2145 UPS-1U if the UPS-1U 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%)

#### **Related tasks**

“MAP 5150: 2145 UPS-1U” on page 376

MAP 5150: 2145 UPS-1U helps you solve problems that have occurred in the 2145 UPS-1U systems that are used on a SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **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.

#### **Action**

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 with the 2145.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F2 or 2145-8F4

- Light path diagnostic MAP FRUs (98%)
- Frame assembly (2%)

2145-4F2

N/A



### Related tasks

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

“MAP 5800: Light path” on page 406

MAP 5800: Light path helps you to solve hardware problems on the SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2 that are preventing the node from booting.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

## **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.

### **Action**

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 with the 2145.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8G4 or 2145-8A4 or 2145-CF8

- Light path diagnostic MAP FRUs (98%)
- System board (2%)

### Related tasks

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

“MAP 5800: Light path” on page 406

MAP 5800: Light path helps you to solve hardware problems on the SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2 that are preventing the node from booting.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

## **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.

### Action

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 with the 2145.
5. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F2 or 2145-8F4

- Light path diagnostic MAP FRUs (98%)
- Frame assembly (2%)

2145-4F2

N/A

### Related tasks

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

“MAP 5800: Light path” on page 406

MAP 5800: Light path helps you to solve hardware problems on the SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2 that are preventing the node from booting.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

## **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.

### **Action**

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 with the 2145.
5. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8G4 or 2145-8A4 or 2145-CF8

- Light path diagnostic MAP FRUs (98%)
- System board (2%)

### Related tasks

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

“MAP 5800: Light path” on page 406

MAP 5800: Light path helps you to solve hardware problems on the SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2 that are preventing the node from booting.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

## 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.

### Action

1. In the sequence shown, exchange the FRUs for new FRUs. See “Possible Cause-FRUs or other”.
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 with the 2145.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

#### 2145-CF8

- Power supply unit (50%)
- System board (50%)

#### 2145-8G4

- Power backplane (90%)
- Power supply assembly (5%)
- System board (5%)

#### 2145-8F2 or 2145-8F4

- Power backplane (90%)
- Power supply assembly (5%)
- Frame assembly (5%)

#### 2145-4F2

N/A

#### **Related tasks**

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

#### **Related reference**

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

### **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.

#### **Action**

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. See “Possible Cause-FRUs or other” after the last action in this section.

Go to the repair verification MAP.

Possible Cause-FRUs or other:

1. High speed SAS adapter (90%)
2. System board (10%)

#### **Related tasks**

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **1121 A high speed SAS adapter has failed.**

#### **Explanation**

A fault has been detected on a high speed SAS adapter.

#### **Action**

In the sequence shown, exchange the FRUs for new FRUs. See “Possible Cause-FRUs or other” after the last action in this section.

Go to the repair verification MAP.

Possible Cause-FRUs or other:

1. High speed SAS adapter (90%)
2. System board (10%)

**Related tasks**

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

**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 solid-state drives that were being accessed through this adapter to go Offline.

**Action**

If this is the first time that this error has occurred on this node:

1. Power off the node.
2. Reseat the high speed SAS adapter card.
3. Power on the node.
4. Submit the 'svcinfo lsmdisk' task and ensure that all of the solid-state 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:

1. In the sequence shown, exchange the FRUs for new FRUs. See “Possible Cause-FRUs or other” after the last action in this section.
2. Submit the 'svcinfo lsmdisk' task and ensure that all of the solid-state 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%)

**Related tasks**

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

**1135 The 2145 UPS has reported an ambient over temperature.****Explanation**

The 2145 UPS has reported an ambient over temperature. The uninterruptible power supply switches to Bypass mode to allow the 2145 UPS to cool.

**Action**

1. Power off the nodes attached to the 2145 UPS.
2. Turn off the 2145 UPS, and then unplug the 2145 UPS from the main power source.
3. Ensure that the air vents of the 2145 UPS are not obstructed.

4. Ensure that the air flow around the 2145 UPS is not restricted.
5. Wait for at least five minutes, and then restart the 2145 UPS. If the problem remains, check the ambient temperature. Correct the problem. Otherwise, exchange the FRU for a new FRU. See "Possible Cause-FRUs or other" after the last action in this section.
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 electronics unit (50%)

Other:

The system ambient temperature is outside the specification (50%)

### **Related tasks**

"Marking errors as fixed" on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

"Viewing the node status" on page 66

You can view the properties for a node from the Viewing General Details panel.

"MAP 5000: Start" on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

"MAP 5700: Repair verification" on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **Related reference**

"Checking the status of the node using the CLI" on page 67

You can use the command-line interface (CLI) to check the status of the node.

"Checking the status of the fibre-channel ports using the CLI" on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

"Uninterruptible power-supply environment requirements" on page 53

An uninterruptible power-supply environment requires that certain specifications for the physical site of the SAN Volume Controller must be met.

## **1136 The 2145 UPS-1U has reported an ambient over temperature. Explanation**

The 2145 UPS-1U has reported an ambient over temperature.

### **Action**

1. Power off the node attached to the 2145 UPS-1U.
2. Turn off the 2145 UPS-1U, and then unplug the 2145 UPS-1U from the main power source.

3. Ensure that the air vents of the 2145 UPS-1U are not obstructed.
4. Ensure that the air flow around the 2145 UPS-1U is not restricted.
5. Wait for at least five minutes, and then restart the 2145 UPS-1U. If the problem remains, check the ambient temperature. Correct the problem. Otherwise, exchange the FRU for a new FRU. See "Possible Cause-FRUs or other" after the last action in this section.
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%)

#### **Related tasks**

"Marking errors as fixed" on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

"Viewing the node status" on page 66

You can view the properties for a node from the Viewing General Details panel.

"MAP 5000: Start" on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

"MAP 5700: Repair verification" on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

#### **Related reference**

"Checking the status of the node using the CLI" on page 67

You can use the command-line interface (CLI) to check the status of the node.

"Checking the status of the fibre-channel ports using the CLI" on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

### **1140 The 2145 UPS has reported that it has a problem with the input AC power.**

#### **Explanation**

The 2145 UPS has reported that it has a problem with the input AC power.

#### **Action**

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. See "Possible Cause-FRUs or other" after the last action in this section.
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:

- 2145 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%)

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **Related reference**

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1141 The 2145 UPS-1U has reported that it has a problem with the input AC power.**

### **Explanation**

The 2145 UPS-1U has reported that it has a problem with the input AC power.

### **Action**

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. See “Possible Cause-FRUs or other” after the last action in this section.
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:

- 2145 UPS-1U input power cable (10%)
- 2145 UPS-1U assembly (10%)

Other:

- The input AC power is missing (40%)
- The input AC power is not in specification (40%)

#### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

#### **Related reference**

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1145 The signal connection between a 2145 and its 2145 UPS is failing.**

### **Explanation**

The signal connection between a 2145 and its 2145 UPS is failing.

### **Action**

1. If other 2145s that are using this uninterruptible power supply are reporting this error, exchange the 2145 UPS electronics unit for a new one.
2. If only this 2145 is reporting the problem, check the signal cable, exchange the FRUs for new FRUs in the sequence shown. See “Possible Cause-FRUs or other” after the last action in this section.
3. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F2 or 2145-8F4 or 2145-8G4

N/A

2145-4F2

- Power cable assembly (40%)
- 2145 UPS electronics assembly (30%)

- System board assembly (25%)
- 2145 disk drive assembly (5%)

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## 1146 The signal connection between a 2145 and its 2145 UPS-1U is failing.

### Explanation

The signal connection between a 2145 and its 2145 UPS-1U is failing.

### Action

1. Exchange the FRUs for new FRUs in the sequence shown. See “Possible Cause-FRU or other” after the last action in this section.
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 with the uninterruptible power supply.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8G4

- Power cable assembly (40%)
- 2145 UPS-1U assembly (30%)
- System board (30%)

2145-8F2 or 2145-8F4

- Power cable assembly (40%)
- 2145 UPS-1U assembly (30%)
- Frame assembly (30%)

2145-4F2

- Power cable assembly (40%)
- 2145 UPS-1U assembly (30%)
- System board assembly (25%)
- 2145 disk drive assembly (5%)

#### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

#### **Related reference**

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

**1150 Data that the 2145 has received from the 2145 UPS suggests the 2145 UPS power cable, the signal cable, or both, are not connected correctly.**

#### **Explanation**

Data that the 2145 has received from the 2145 UPS suggests the 2145 UPS power cable, the signal cable, or both, are not connected correctly.

#### **Action**

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

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1151 Data that the 2145 has received from the 2145 UPS-1U suggests the 2145 UPS-1U power cable, the signal cable, or both, are not connected correctly.**

### **Explanation**

Data that the 2145 has received from the 2145 UPS-1U suggests the 2145 UPS-1U power cable, the signal cable, or both, are not connected correctly.

### **Action**

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

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## 1155 A power domain error has occurred.

### Explanation

Both 2145s of a pair are powered by the same uninterruptible power supply.

### Action

1. List the cluster's 2145s 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

### Related tasks

“Viewing the vital product data” on page 127

You can view the vital product data for a node from the Viewing Vital Product Data panel of the SAN Volume Controller Console.

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

## 1160 The output load on the 2145 UPS exceeds the specification.

### Explanation

The 2145 UPS is reporting that too much power is being drawn from it. The power overload warning LED, which is above the load level indicators, on the 2145 UPS will be on.

### Action

1. Determine the 2145 UPS that is reporting the error from the error event data. Perform the following steps on just this uninterruptible power supply.
2. Check that the 2145 UPS is still reporting the error. If the power overload warning LED is no longer on, go to step 6.
3. Ensure that only 2145s are receiving power from the uninterruptible power supply. Ensure that there are no switches or disk controllers that are connected to the 2145 UPS.
4. Remove each connected 2145 input power in turn, until the output overload is removed.
5. Exchange the FRUs for new FRUs in the sequence shown, on the overcurrent 2145. See “Possible Cause-FRU or other” after the last action in this section.
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 2145 UPS.
7. Go to repair verification MAP.

Possible Cause-FRUs or other:

- Power cable assembly (50%)
- Power supply assembly (40%)
- 2145 UPS electronics assembly (10%)

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1161 The output load on the 2145 UPS-1U exceeds the specifications (reported by 2145 UPS-1U alarm bits).**

### **Explanation**

The output load on the 2145 UPS-1U exceeds the specifications (reported by 2145 UPS-1U alarm bits).

### **Action**

1. Ensure that only 2145s are receiving power from the uninterruptible power supply. Also, ensure that no other devices are connected to the 2145 UPS-1U.
2. Exchange, in the sequence shown, the FRUs for new FRUs. See “Possible Cause-FRUs or other” after the last action in this section. If the Overload Indicator is still illuminated with all outputs disconnected, replace the 2145 UPS-1U.
3. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145 UPS-1U.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

- Power cable assembly (50%)
- Power supply assembly (40%)
- 2145 UPS-1U assembly (10%)



### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1165 The 2145 UPS output load is unexpectedly high. The 2145 UPS output is possibly connected to an extra non-2145 load.**

### **Explanation**

The 2145 UPS output load is unexpectedly high. The 2145 UPS output is possibly connected to an extra non-2145 load.

### **Action**

1. Ensure that only 2145s are receiving power from the uninterruptible power supply. Ensure that there are no switches or disk controllers that are connected to the 2145 UPS.
2. Check node status. If all nodes show a status of “online”, the problem no longer exists. Mark the error that you have just repaired “fixed” and go to the repair verification MAP.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

None

Other:

- Configuration error

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## 1166 The 2145 UPS-1U output load is unexpectedly high.

### Explanation

The uninterruptible power supply output is possibly connected to an extra non-2145 load.

### Action

1. Ensure that there are no other devices that are connected to the 2145 UPS-1U.
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:

- 2145 UPS-1U assembly (5%)

Other:

- Configuration error (95%)

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1170 2145 UPS electronics fault (reported by the 2145 UPS alarm bits).**

### **Explanation**

2145 UPS electronics fault (reported by the 2145 UPS alarm bits).

### **Action**

1. Replace the uninterruptible power supply electronics assembly.
2. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the UPS.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145 UPS electronics assembly (100%)

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1171 2145 UPS-1U electronics fault (reported by the 2145 UPS-1U alarm bits).**

### **Explanation**

2145 UPS-1U electronics fault (reported by the 2145 UPS-1U alarm bits).

### **Action**

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 2145 UPS-1U.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145 UPS-1U assembly (100%)

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **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).

### **Action**

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%)

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1180 2145 UPS battery fault (reported by 2145 UPS alarm bits). Explanation**

2145 UPS battery fault (reported by 2145 UPS alarm bits).

### Action

1. Replace the 2145 UPS battery assembly.
2. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145 UPS battery assembly (100%)

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1181 2145 UPS-1U battery fault (reported by 2145 UPS-1U alarm bits).**

### **Explanation**

2145 UPS-1U battery fault (reported by 2145 UPS-1U alarm bits).

### **Action**

1. Replace the 2145 UPS-1U battery assembly.
2. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145 UPS-1U battery assembly (100%)

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1185 2145 UPS fault, with no specific FRU identified (reported by uninterruptible power supply alarm bits).**

### **Explanation**

2145 UPS fault, with no specific FRU identified (reported by 2145 UPS alarm bits).

### **Action**

1. In the sequence shown, exchange the FRU for a new FRU. See “Possible Cause-FRUs or other” after the last action in this section.
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.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- 2145 UPS electronics assembly (60%)
- 2145 UPS battery assembly (20%)
- 2145 UPS assembly (20%)



### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1186 A problem has occurred in the 2145 UPS-1U, with no specific FRU identified (reported by 2145 UPS-1U alarm bits).**

### **Explanation**

A problem has occurred in the 2145 UPS-1U, with no specific FRU identified (reported by 2145 UPS-1U alarm bits).

### **Action**

1. In the sequence shown, exchange the FRU for a new FRU. See “Possible Cause-FRUs or other” after the last action in this section.
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:

2145 UPS-1U assembly (100%)

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **Related reference**

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1190 The 2145 UPS battery has reached its end of life.**

### **Explanation**

The 2145 UPS battery has reached its end of life.

### **Action**

1. Replace the 2145 UPS battery assembly.
2. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145 UPS battery assembly (100%)

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1191 The 2145 UPS-1U battery has reached its end of life.**

### **Explanation**

The 2145 UPS-1U battery has reached its end of life.

### **Action**

1. Replace the 2145 UPS-1U battery assembly.
2. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145 UPS-1U battery assembly (100%)

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

“Checking the status of the fibre-channel ports using the CLI” on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## 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.

### Action

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 error 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 error 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 SAN Volume Controller 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.

### Related tasks

"MAP 5700: Repair verification" on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### 1195 A 2145 is missing from the cluster.

#### Explanation

You can resolve this problem by repairing the failure on the missing 2145.

## Action

1. If it is not obvious which node in the cluster has failed, check the status of the nodes and find the 2145 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 2145.
5. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

## Related tasks

"Marking errors as fixed" on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

"Viewing the node status" on page 66

You can view the properties for a node from the Viewing General Details panel.

"MAP 5000: Start" on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

"MAP 5700: Repair verification" on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## Related reference

"Checking the status of the node using the CLI" on page 67

You can use the command-line interface (CLI) to check the status of the node.

"Checking the status of the fibre-channel ports using the CLI" on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

## **1200 The configuration is not valid. Too many devices have been presented to the cluster or 2145.**

### **Explanation**

The configuration is not valid. Too many devices have been presented to the cluster or 2145.

## Action

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%)

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Listing MDisks using the CLI” on page 96

You can use the command-line interface (CLI) to list the managed disks (MDisks).

“Discovering MDisks” on page 92

You can have the cluster rescan the fibre-channel network. The rescan discovers any new managed disks (MDisks) that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

“Discovering MDisks using the CLI” on page 98

You can use the command-line interface (CLI) to discover managed disks (MDisks).

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## **1201 A solid-state drive requires a recovery.**

### **Explanation**

The solid-state drive identified by this error has suffered an error which requires the drive to be reformatted before it can be used again. A corruption in the drive metadata has caused all of the data on the drive to be lost. The drive is still usable, but it must be reformatted to resolve the issue.

### **Action**

The following actions delete all the VDIs and VDisk host maps for the managed disk group that contains the offline managed disk. Therefore, it is strongly recommended that you back up the SVC configuration with the `svconfig backup` command so that you can use the backup file later.

If the MDisk is a managed mode MDisk:

1. Record all of the properties of all of the VDisk copies, including unmirrored VDIs, that use storage from the managed disk group that contains the offline managed disk. Record all of the properties of the managed disk group and the list of all of the managed disks in that managed disk group.
2. Delete the MDisk group that contains the affected MDisk. The affected MDisk will become unmanaged.
3. Submit the command `'svctask includemdisk (mdisk)'`, where (mdisk) is the name or the ID of the offline MDisk.
4. Create a new managed disk group with the same properties as the deleted managed disk group and add the drive back into the managed disk group.

5. Adding the disk into the managed disk group in step 4 reformats the drive. The reformat process might take some time to complete. Warning: Do not remove the solid-state drive from the node, or power off the node, while the disk is being formatted.
6. If the reformat process does not complete quickly, the command will fail with the message 'CMMVC6296E One or more managed disks (MDisks) have failed validation tests. The first failing MDisk ID is (managed disk ID)'. If this error is returned, the managed disk might disappear from the list of managed disks in the cluster while the format process is occurring. When the format has completed, the managed disk will be rediscovered as a new managed disk with a default name and possibly a new managed disk ID. Once the managed disk has been rediscovered, resubmit the command to add the drive back into the managed disk group.
7. Add all of the other managed disks that were in the deleted managed disk group into the new managed disk group.
8. Add a VDisk copy for all of the VDIs that were mirrored in the original managed disk group.
9. Create all of the standard VDIs that were members of the original managed disk group and restore data from backups using the information recorded in step 1.

If the MDisk is an image mode MDisk:

1. Record the properties the image mode VDisk.
2. Delete the image mode VDisk for the affected MDisk. The MDisk will become unmanaged.
3. Submit the command 'svctask includemdisk (mdisk)', where (mdisk) is the name or the ID of the offline MDisk.
4. Create a new, temporary, managed disk group and add the MDisk into this new managed disk group.
5. Adding the disk into the managed disk group in step 4 reformats the drive. The reformat process might take some time to complete. Warning: Do not remove the solid-state drive from the node, or power off the node, while the disk is being formatted.
6. If the reformat process does not complete quickly, the command will fail with the message 'CMMVC6296E One or more managed disks (MDisks) have failed validation tests. The first failing MDisk ID is (managed disk ID)'. If this error is returned, the managed disk might disappear from the list of managed disks in the cluster while the format process is occurring. When the format has completed, the managed disk will be rediscovered as a new managed disk with a default name and possibly a new managed disk ID. Once the managed disk has been rediscovered, resubmit the command to add the drive back into the managed disk group.
7. Delete the temporary managed disk group that was created in step 4.
8. Either recreate the image mode VDisk or add the image mode VDisk copy, depending on the original configuration of the VDisk.
9. NOTE: If the original configuration of the VDisk was an image mode VDisk, ALL of the data has been lost and must be recreated.

Possible Cause-FRUs or other:

- None



## 1202 A solid-state drive is offline.

### Explanation

The offline solid-state drive identified by this error must be repaired. The SAN Volume Controller (SVC) error log will identify a managed disk ID. This can be determined by running the maintenance procedure on the 1202 error.

### Action

Determine whether the managed disk referred to in the error is currently offline using the SVC GUI or command-line interface (CLI). If the drive is online then the problem was transient, and should be investigated by IBM technical support before you replace hardware.

If the fault is still present, the managed disk will be offline. If the managed disk that is indicated in the error is still offline, perform the following steps:

1. Submit the command `'svcinfo lsmdisk -filtervalue status=offline'` to identify all of the offline solid-state drives.
2. Submit the command `'svcinfo lsmdisk (mdisk id)'`, where (mdisk id) is the ID of the offline MDisk. Record the 'controller\_name', 'node\_name' and 'location' properties of the managed disk.
3. Submit the command `'svcinfo lsnodevpd (node_name)'`, where (node\_name) is the value of that property variable recorded in step 2. Record the front\_panel\_id property of this node.
4. If there are multiple solid-state drives in a single node and all of the solid-state drives in the node are offline, it is likely that the faulty component will not be the solid-state drive. Determine how many solid-state drives are in the affected node, first ensure that the node is online. Submit the command `'svcinfo lsnode (node_name)'`, where (node\_name) is the name of the node identified in step 2, and ensure that the 'status' property value is Online. If the node is offline, follow standard service procedures to resolve the node offline status.
5. Submit the command `'svcinfo lsmdisk -filtervalue controller_name=(controller name)'`, where (controller name) is the value of that property variable recorded in step 2.
6. If the command in step 5 displays multiple managed disks and all of them are 'offline', replace the following components in sequence: High speed SAS adapter and SAS Cable, PCIe Riser card, SAS drive backplane.
7. If there is a spare drive slot in any of slots 0-3 in another SVC node in the same cluster that contains the high speed SAS adapter, swap the solid-state drive into a spare drive slot using the solid-state drive remove/replace instructions in the note below. If the managed disk is also offline in the new node, the solid-state drive must be replaced by following the procedure in MAP 6000 Replace offline SSD. If the managed disk comes online in this new drive bay, the drive has not failed. Swap the solid-state drive back into its original location to determine whether the SAS components in the original node have failed. If the drive stays offline in the original node, the faulty component is either the high speed SAS adapter, the SAS cable or the disk drive backplane. Otherwise, the problem has been resolved by reseating the drive.

Note: A solid-state drive can be swapped into a spare drive bay on any node that contains a high speed SAS adapter, but installing the drive into a different node will introduce a performance penalty because I/Os must be forwarded between nodes. To restore performance, the drive should be returned to its original node as soon as possible once the issue has been resolved.

Possible Cause-FRUs or other:

- High speed SAS adapter (30%)
- SAS Cable (30%)
- solid-state drive (30%)
- Disk drive backplane (10%)

#### **Related tasks**

“MAP 6000: Replace offline SSD” on page 428

MAP 6000: This procedure replaces a solid-state drive (SSD) that has failed while it is still a member of a managed disk (MDisk) group.

### **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.

#### **Action**

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%)

#### **Related tasks**

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **1210 A local fibre-channel port has been excluded.**

#### **Explanation**

A local fibre-channel port has been excluded.

#### **Action**

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%)

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“MAP 5600: Fibre-channel” on page 396

MAP 5600: Fibre-channel helps you to solve problems that have occurred on the SAN Volume Controller fibre-channel ports.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

“Fibre-channel link failures” on page 349

When a failure occurs on a single fibre channel link, the small form-factor pluggable (SFP) connector might need to be replaced.

## 1215 A solid-state drive is failing.

### Explanation

The solid-state drive has detected faults that indicate that the drive is likely to fail soon. The drive should be replaced. The SAN Volume controller (SVC) error log will identify a managed disk ID for the solid-state drive that caused the error.

### Action

Perform the maintenance procedure for the 1215 error to identify the managed disk ID for the solid-state drive that caused the error.

If the managed disk has gone offline since this error has occurred, the managed disk has failed and you must follow the solid-state drive replacement procedure in MAP 6000 Replace offline SSD.

If the managed disk is still online, perform the following procedure to replace the solid-state drive without any data loss:

1. Submit the command `'svctask rmmdisk -force (mdisk name/id)'`, where (mdisk name/id) is the name or ID of the managed disk identified in the error log. This command migrates all of the data from the failing managed disk into the free extents in the rest of the managed disk group. If the command fails with a message that indicates that there are not enough free extents, create more free extents in the managed disk group and resubmit the command. If you cannot create a sufficient number of free extents so that the command completes without error, you must use MAP 6000 Replace offline SSD to replace the drive. You can use any of the following three options to increase the number of free extents: The first option is to remove some of the VDisk copies that exist in this managed disk group. The second option is to migrate some of the VDisk copies into other managed disk groups. The third option is to temporarily add more managed disks to the managed disk group.

2. Wait until the status of the managed disk that must be replaced is 'Unmanaged'.
3. Submit the command 'svcinfo lsmdisk (mdisk id)', where (mdisk id) is the name or the ID of the managed disk that is identified in the error log. Record the 'controller\_name', 'node\_name' and 'location' properties of the MDisk.
4. Submit the command 'svcinfo lsnodevpd (node name/id)', where (node name/id) is the node name displayed by the lsmdisk command in step 3. Record the front\_panel\_id property of this node.
5. Perform the solid-state drive remove/replace instructions in the Hardware Maintenance guide to replace the solid-state drive. To identify the correct SSD to replace, use the following information: The 'front\_panel\_id' is on a label on the front of the node, and 'location' identifies the specific drive bay of the node. The drive bays are labeled in red numerals to the right of the drive slot.
6. Submit the command 'svctask detectmdisk' to discover the new solid-state drive. Verify that a new managed disk is discovered and that it has the correct slot number.
7. Add the new managed disk into the managed disk group.
8. Reverse any procedures that might have been performed in step 1 to create free extents in the managed disk group.
9. This step is optional. Run the 'balance.pl' script that is available from the alphaworks package called SVCTools to redistribute the VDisk extents equally across all of the managed disks in the managed disk group.

Possible Cause-FRUs or other:

- solid-state drive (100%)

#### **Related tasks**

"MAP 6000: Replace offline SSD" on page 428

MAP 6000: This procedure replaces a solid-state drive (SSD) that has failed while it is still a member of a managed disk (MDisk) group.

## **1216 SAS errors have exceeded thresholds.**

### **Explanation**

The SAN Volume Controller has experienced a large number of SAS communication errors, which indicates a faulty SAS component that must be replaced.

### **Action**

In the sequence shown, exchange the FRUs for new FRUs. See "Possible Cause-FRUs or other" after the last action in this section.

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. solid-state drive (5%)

### **Related tasks**

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **1217 A solid-state drive has exceeded the temperature warning threshold.**

#### **Explanation**

The solid-state drive identified by this error has reported that its temperature is higher than the warning threshold.

#### **Action**

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:

- Solid-state 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.

#### **Action**

1. View the error 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%)

#### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Listing MDisks using the CLI” on page 96

You can use the command-line interface (CLI) to list the managed disks (MDisks).

“Discovering MDisks” on page 92

You can have the cluster rescan the fibre-channel network. The rescan discovers any new managed disks (MDisks) that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

“Discovering MDisks using the CLI” on page 98

You can use the command-line interface (CLI) to discover managed disks (MDisks).

“Viewing the error event log” on page 170

You can view the error event log by using the SAN Volume Controller command-line interface (CLI) or the SAN Volume Controller Console.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

#### **Related reference**

“Checking disk controller status using the CLI” on page 106

You can use the command-line interface (CLI) to check the status of the disk controllers.

“Determining the failing enclosure or disk controller using the CLI” on page 107

You can use the command-line interface (CLI) to determine the failing enclosure or disk controller.

### **1230 A login has been excluded.**

#### **Explanation**

A port to port fabric connection, or login, between the SAN Volume Controller node and either a controller or another SAN Volume Controller cluster has had excessive errors. The login has therefore been excluded, and will not be used for I/O operations.

#### **Action**

Determine the remote system, which might be either a controller or a SAN Volume Controller cluster. Check the error 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 or adapter, (30%)
- Fibre-channel cable, local port to switch, (30%)
- SAN Volume controller SFP, (9%)
- SAN Volume controller fibre channel adapter, (1%)

Note: The first two FRUs are not SAN Volume Controller FRUs.

#### **Related tasks**

"Marking errors as fixed" on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

"Discovering MDisks" on page 92

You can have the cluster rescan the fibre-channel network. The rescan discovers any new managed disks (MDisks) that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

"Discovering MDisks using the CLI" on page 98

You can use the command-line interface (CLI) to discover managed disks (MDisks).

"MAP 5000: Start" on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

"MAP 5700: Repair verification" on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

#### **Related reference**

"Checking disk controller status using the CLI" on page 106

You can use the command-line interface (CLI) to check the status of the disk controllers.

### **1310 A managed disk is reporting excessive errors.**

#### **Explanation**

A managed disk is reporting excessive errors.

#### **Action**

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%)

#### **Related tasks**

"Marking errors as fixed" on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

"Listing MDisks using the CLI" on page 96

You can use the command-line interface (CLI) to list the managed disks (MDisks).

"MAP 5700: Repair verification" on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **1311 A solid-state drive is reporting excessive errors.**

#### **Explanation**

The drive that is reporting excessive errors has been taken offline.

#### **Action**

If the managed disk that is referenced in the error is no longer offline then the error has been resolved; therefore, mark the error as fixed.

It is likely that there are higher priority SAS or solid-state drive errors in the error log. These log entries will provide more detailed information about the cause of the errors. If any higher priority errors exist for this node or managed disk, resolve those errors first and mark this error as fixed.

If there are no other errors in the error log related to a solid-state drive or any such errors have been recently resolved, and the managed disk is offline, replace the FRUs below in the order listed. After replacing each FRU, submit the command 'svctask includemdisk (mdisk id / name)'. This command attempts to bring the drive online. If replacing the FRU has not resolved the problem, the managed disk will be returned to the offline state.

NOTE: If you replace the solid-state drive, use the procedure defined in MAP 6000.

Possible Cause-FRUs or other:

1. SAS cable (30%)
2. High speed SAS adapter (30%)
3. SAS drive backplane (30%)
4. Solid-state drive (10%)



Other:

None

### **1320 A disk I/O medium error has occurred.**

#### **Explanation**

A disk I/O medium error has occurred.

#### **Action**

1. Check whether the VDisk the error is reported against is mirrored. If it is, check if there is a "1870 Mirrored VDisk offline because a hardware read error has occurred" error relating to this VDisk in the error log. Also check if one of the mirror copies is synchronizing. If all these tests are true then you must delete the VDisk Copy that is not synchronized from the VDisk. Check that the VDisk is online before continuing with the following actions. Wait until the medium error is corrected until trying to recreate the VDisk 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 VDisk 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 VDisk 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 VDisk content that is being used to a new VDisk. Depending on the circumstances, this new VDisk can be kept and mirrored, or the original VDisk 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%)

## Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Listing MDisks using the CLI” on page 96

You can use the command-line interface (CLI) to list the managed disks (MDisks).

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## **1330 A suitable Managed Disk (MDisk) 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 SAN Volume Controller 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 an MDisk group.
- 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.

### **Action**

1. Resolve any known fibre-channel network problems.
2. Ask the customer to confirm that MDisks have been added to MDisk groups 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%)

#### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Listing MDisks using the CLI” on page 96

You can use the command-line interface (CLI) to list the managed disks (MDisks).

“SAN problem determination” on page 348

The procedures to service the SAN Volume Controller that are provided here help you solve problems on the SAN Volume Controller and its connection to the storage area network (SAN).

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **1335 Quorum disk not available.**

#### **Explanation**

Quorum disk not available.

#### **Action**

1. View the error 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%)

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Listing MDisks using the CLI” on page 96

You can use the command-line interface (CLI) to list the managed disks (MDisks).

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **Related reference**

“Determining the failing enclosure or disk controller using the CLI” on page 107

You can use the command-line interface (CLI) to determine the failing enclosure or disk controller.

## **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.

### **Action**

1. Repair problems on all enclosures/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 error 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 switch

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **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 a component of the SAN.

#### **Action**

1. View the error 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 switch problem determination and repair procedures for the switches connected to the 2145 node or controller.
3. Perform fibre-channel 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 error log dump. Call your hardware support center.
6. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Fibre-channel switch
- Fibre-channel cabling

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **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.

#### **Action**

1. View the error 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 switches connected to the 2145 and any other fibre-channel 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 error 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 switch

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## **1400 The 2145 cannot detect an Ethernet connection.**

### **Explanation**

The 2145 cannot detect an Ethernet connection.

### **Action**

1. Go to the Ethernet MAP.
2. Go to the repair verification MAP.

Possible Cause-FRUs or other:

2145-8G4 or 2145-8A4 or 2145-CF8

- Ethernet cable (25%)
- System board (25%)

2145-8F2 or 2145-8F4

- Ethernet cable (25%)
- Frame assembly (25%)

2145-4F2

- Ethernet cable (25%)
- System board assembly (25%)

Other:

- Ethernet cable is disconnected or damaged (25%)
- Ethernet hub fault (25%)

### **Related tasks**

“MAP 5500: Ethernet” on page 393

MAP 5500: Ethernet helps you solve problems that have occurred on the SAN Volume Controller Ethernet.

## **1550 A cluster path has failed.**

### **Explanation**

One of the 2145 fibre-channel ports is unable to communicate with all the other 2145s in the cluster.

### **Action**

1. Check for incorrect switch zoning.
2. Repair the fault in the fibre-channel network fabric.

3. Check the status of the node ports. If the status of the node ports shows as active, mark the error that you have just 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 with the 2145.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

Fibre-channel network fabric fault (100%)

#### **Related tasks**

"Marking errors as fixed" on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

"MAP 5000: Start" on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

"MAP 5700: Repair verification" on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

#### **Related reference**

"Checking the status of the fibre-channel ports using the CLI" on page 68

You can use the command-line interface (CLI) to check the status of the fibre-channel ports.

### **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 SAN Volume Controller, although multiple component controller definitions are created, the SAN Volume Controller 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 SAN Volume Controller 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.

#### **Action**

- 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 the SAN Volume Controller support Web site [www.ibm.com/storage/support/2145](http://www.ibm.com/storage/support/2145) 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%)

### **1600 Mirrored disk repair halted because of difference. Explanation**

During the repair of a mirrored VDisk 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 either VDisk Copy's data. Therefore it is important not to use the VDisk 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.

### **Action**

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 VDisk. 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 VDisk from a backup.
- Decide which disk has correct data, then delete the copy that is different and recreate 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 SAN Volume Controller 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 VDisk 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 VDisk validate and repair process has the option to create virtual medium errors on sectors that do not match on all VDisk 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.

## Action

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 due to a mirrored disk validate and repair operation that created errors for differences, or whether the errors were created due to a copy operation. Follow the corresponding option shown below.

1. If the virtual medium errors occurred due to a mirrored disk validate and repair operation that created medium errors for differences, then also ensure that the VDisk 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 VDisk 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 VDisk will have to be recovered from a backup using a process that rewrites all of the VDisk's sectors.
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 VDisk and to not propagate the medium errors to copies of the VDisk. Fixing higher priority errors in the error log would have corrected the medium error on the source VDisk. Once the medium errors have been fixed, you must run the copy operation again to clear the virtual medium errors from the target VDisk. 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 VDIs on the target managed disk that have the virtual medium errors. This VDisk deletion reduces the number of virtual medium error entries in the MDisk table. Migrating the VDisk 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 VDisk is migrated.

Possible Cause-FRUs or other:

- None

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Listing MDisks using the CLI” on page 96

You can use the command-line interface (CLI) to list the managed disks (MDisks).

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Determining the failing enclosure or disk controller using the CLI” on page 107

You can use the command-line interface (CLI) to determine the failing enclosure or disk controller.

## 1620 An MDisk group is offline.

### Explanation

An MDisk group is offline.

### Action

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%)

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Listing MDisks using the CLI” on page 96

You can use the command-line interface (CLI) to list the managed disks (MDisks).

“Discovering MDisks” on page 92

You can have the cluster rescan the fibre-channel network. The rescan discovers any new managed disks (MDisks) that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

“Discovering MDisks using the CLI” on page 98

You can use the command-line interface (CLI) to discover managed disks (MDisks).

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Determining the failing enclosure or disk controller using the CLI” on page 107

You can use the command-line interface (CLI) to determine the failing enclosure or disk controller.

## **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.

### **Action**

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**

SAN Volume Controller has detected that an IBM DS series disk controller's configuration is not supported by SAN Volume Controller. The disk controller is operating in RDAC mode. The disk controller might appear to be operating with SAN Volume Controller; however, the configuration is unsupported because it is known to not work with SAN Volume Controller.

#### **Action**

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 SAN Volume Controller has detected that a disk controller's configuration is not supported by SAN Volume Controller. The disk controller might appear to be operating with SAN Volume Controller; however, the configuration detected can potentially cause issues and should not be used. The unsupported configuration is shown in the event data.

#### **Action**

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

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Discovering MDisks” on page 92

You can have the cluster rescan the fibre-channel network. The rescan discovers any new managed disks (MDisks) that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

“Discovering MDisks using the CLI” on page 98

You can use the command-line interface (CLI) to discover managed disks (MDisks).

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## **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 error 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.

### Action

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 `svcinfo 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 `svctask 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

### Related tasks

"MAP 5700: Repair verification" on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## 1630 The number of device logins was reduced.

### Explanation

The number of port to port fabric connections, or logins, between the SAN Volume Controller node and a storage controller has decreased. This might be caused by a problem on the SAN or by a deliberate reconfiguration of the SAN.

### Action

1. Check the error in the cluster error log to identify the object ID associated with the error.
2. Check the availability of the failing device using the following command line: `svcinfo 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 error 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 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%)

### **Related tasks**

"Marking errors as fixed" on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

"Discovering MDisks" on page 92

You can have the cluster rescan the fibre-channel network. The rescan discovers any new managed disks (MDisks) that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

"Discovering MDisks using the CLI" on page 98

You can use the command-line interface (CLI) to discover managed disks (MDisks).

"MAP 5000: Start" on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

"MAP 5700: Repair verification" on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **Related reference**

"Checking disk controller status using the CLI" on page 106

You can use the command-line interface (CLI) to check the status of the disk controllers.

## **1660 The initialization of the managed disk has failed.**

### **Explanation**

The initialization of the managed disk has failed.

### **Action**

1. View the error 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%)

#### **Related tasks**

"Marking errors as fixed" on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

"Listing MDisks using the CLI" on page 96

You can use the command-line interface (CLI) to list the managed disks (MDisks).

"Discovering MDisks" on page 92

You can have the cluster rescan the fibre-channel network. The rescan discovers any new managed disks (MDisks) that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

"Discovering MDisks using the CLI" on page 98

You can use the command-line interface (CLI) to discover managed disks (MDisks).

"MAP 5000: Start" on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

"MAP 5700: Repair verification" on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

#### **Related reference**

"Determining the failing enclosure or disk controller using the CLI" on page 107

You can use the command-line interface (CLI) to determine the failing enclosure or disk controller.

### **1661 The solid-state drive initialization process has failed. Explanation**

The drive has not been initialized.

#### **Action**

If there was data on the drive, contact IBM technical support for assistance.

If the drive is currently in a managed disk group and must be replaced, follow the procedure in MAP 6000 Replace offline SSD to replace the hardware. If the drive is not currently in a managed disk group and must be replaced, use the remove/replace instructions in the Hardware Maintenance guide to replace the solid-state drive (SSD).

Possible Cause-FRUs or other:

- solid-state drive (100%)

**Related tasks**

“MAP 6000: Replace offline SSD” on page 428

MAP 6000: This procedure replaces a solid-state drive (SSD) that has failed while it is still a member of a managed disk (MDisk) group.

**1665 A storage device is not recognized.**

**Explanation**

An unrecognized storage device has been inserted into one of the drive bays that is reserved for solid-state drives. The unrecognized device will not appear as a managed disk.

**Action**

If a drive that is not a solid-state drive has been inserted in the slot, remove the drive from the slot.

If you believe that a solid-state drive that was purchased from IBM was inserted into the slot and that the drive is supported for the SAN Volume Controller, contact IBM technical support for assistance.

Possible Cause-FRUs or other:

- None.

**1670 The CMOS battery on the 2145 system board failed.**

**Explanation**

The CMOS battery on the 2145 system board failed.

**Action**

1. Replace the CMOS battery.
2. Mark the error that you have just repaired as “fixed”.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

CMOS battery (100%)

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### 1695 Persistent unsupported disk controller configuration.

#### Explanation

A disk controller configuration that might prevent failover for SAN Volume Controller has persisted for more than four hours. The problem was originally logged through a 010032 event, service error code 1625.

#### Action

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 Metro Mirror or Global Mirror relationship

#### 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 Metro Mirror or Global Mirror relationships, whose control data is stored by the I/O group, were active at the time of the failure and the current state of the relationship could not be recovered.

#### Action

To fix this error it is necessary to delete all of the relationships that could not be recovered and then recreate the relationships.

1. Note the I/O group index against which the error is logged.
2. List all of the Metro Mirror and Global Mirror relationships that have either a master or an auxiliary VDisk in this I/O group. Use the VDisk view to determine which VDIs in the I/O group you noted have a relationship defined.
3. Note the details of the Metro Mirror and Global Mirror relationships that are listed so that they can be recreated.

4. Delete all of the Metro Mirror and Global Mirror relationships that are listed.  
Note: The error will automatically be marked as "fixed" once the last relationship on the I/O group is deleted. New relationships should not be created until the error is fixed.
5. Using the details noted in step 3, recreate all of the Metro Mirror and Global Mirror relationships that you just deleted. Note: You are able to delete a Metro Mirror or Global Mirror relationship from either the master or auxiliary cluster; however, you must recreate the relationship on the master cluster. Therefore, it might be necessary to go to another cluster 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 directly partnered with each other or are indirectly partnered 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.

The following 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.

The following 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 SAN Volume Controller 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.

### Action

To fix this error it is necessary to delete all of the relationships that could not be recovered and then recreate 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 SAN Volume Controller 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

### Related tasks

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

**1720 In a Metro Mirror or Global Mirror operation, the relationship has stopped and lost synchronization, for a reason other than a persistent I/O error.**

### Explanation

In a Metro Mirror or Global Mirror operation, the relationship has stopped and lost synchronization, for a reason other than a persistent I/O error.

### Action

1. Restart the relationship after fixing errors of higher priority.
2. Mark the error that you have just repaired as “fixed”.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## **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.

### **Action**

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 switch configuration error
- Fibre-channel switch

### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Listing MDisks using the CLI” on page 96

You can use the command-line interface (CLI) to list the managed disks (MDisks).

“Discovering MDisks” on page 92

You can have the cluster rescan the fibre-channel network. The rescan discovers any new managed disks (MDisks) that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

“Discovering MDisks using the CLI” on page 98

You can use the command-line interface (CLI) to discover managed disks (MDisks).

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Determining the failing enclosure or disk controller using the CLI” on page 107

You can use the command-line interface (CLI) to determine the failing enclosure or disk controller.

## **1850 A cluster recovery operation was performed but data on one or more VDIsks has not been recovered.**

### **Explanation**

A cluster recovery operation was performed but data on one or more VDIsks has not been recovered.

### **Action**

1. The support center will direct the user to restore the data on the affected virtual disks (VDIsks).
2. When the VDisk 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



### Related tasks

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Listing MDisks using the CLI” on page 96

You can use the command-line interface (CLI) to list the managed disks (MDisks).

“Discovering MDisks” on page 92

You can have the cluster rescan the fibre-channel network. The rescan discovers any new managed disks (MDisks) that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

“Discovering MDisks using the CLI” on page 98

You can use the command-line interface (CLI) to discover managed disks (MDisks).

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### Related reference

“Determining the failing enclosure or disk controller using the CLI” on page 107

You can use the command-line interface (CLI) to determine the failing enclosure or disk controller.

## **1860 Space-efficient VDisk Copy offline due to failed repair. Explanation**

The attempt to repair a space-efficient volume’s metadata that describes the disk contents has failed due to problems with the automatically maintained backup copy of this data. The error event data describes the problem.

### Action

Delete the space-efficient VDisk 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 Space-efficient VDisk Copy offline due to corrupt metadata.**

### Explanation

A space-efficient VDisk has been taken offline because there is an inconsistency in the SAN Volume Controller 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 T3 recovery) or because of a software error. The event data gives information on the reason.

The SAN Volume Controller maintains backup copies of the metadata and it might be possible to repair the space-efficient VDisk using this data.

## Action

SAN Volume Controller is able to repair the inconsistency in some circumstances. Run the repair VDisk 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 space-efficient VDisk 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

## Related tasks

"Repairing offline space-efficient VDIs" on page 84

When a space-efficient virtual disk (VDisk) is taken offline because its metadata is corrupted, you can use the Repairing Space-Efficient VDisk panel to repair the metadata. The repair operation automatically detects corrupted metadata and performs any necessary repair actions.

## Related reference

"Repairing a space-efficient VDisk using the CLI" on page 85

You can use the **repairsevdiskcopy** command from the command-line interface to repair the metadata on a space-efficient virtual disk (VDisk).

## 1865 Space-efficient VDisk Copy offline due to insufficient space.

### Explanation

A space-efficient VDisk has been taken offline because there is insufficient allocated real capacity available on the VDisk for the used space to increase further. If the space-efficient VDisk is auto-expand enabled, then the managed disk group it is in also has no free space.

## Action

The service action differs depending on whether the space-efficient VDisk 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 VDisk Copy is auto-expand enabled, perform one or more of the following actions. When you have performed all of the actions that you intend to perform, mark the error as "fixed"; the VDisk Copy will then return online.

- Determine why the managed disk group free space has been depleted. Any of the space-efficient VDisk Copies, with auto-expand enabled, in this managed disk group might have expanded at an unexpected rate; this could indicate an application error. New VDisk Copies might have been created in, or migrated to, the managed disk group.
- Increase the capacity of the managed disk group that is associated with the space-efficient VDisk Copy by adding more MDisks to the group.
- Provide some free capacity in the managed disk group by reducing the used space. VDisk Copies that are no longer required can be deleted, the size of VDisk Copies can be reduced or VDisk Copies can be migrated to a different managed disk group.

- Migrate the space-efficient VDisk Copy to a managed disk group that has sufficient unused capacity.
- Consider reducing the value of the managed disk group warning threshold to give more time to allocate extra space.

If the VDisk Copy is not auto-expand enabled, perform one or more of the following actions. In this case the error will automatically be marked as "fixed", and the VDisk Copy will return online when space is available.

- Determine why the space-efficient VDisk Copy used space has grown at the rate that it has. There might be an application error.
- Increase the real capacity of the VDisk Copy.
- Enable auto-expand for the space-efficient VDisk Copy.
- Consider reducing the value of the space-efficient VDisk Copy warning threshold to give more time to allocate more real space.

Possible Cause-FRUs or other:

- None

### **1870 Mirrored VDisk offline because a hardware read error has occurred.**

#### **Explanation**

While attempting to maintain the VDisk mirror, a hardware read error occurred on all of the synchronized Virtualized Disk copies.

The Virtual Disk Copies might be inconsistent, so the VDisk is now offline.

#### **Action**

- 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 Virtual Disk Copies have been fixed, mark this error as "fixed" to run without the mirror. You can then delete the Virtual Disk Copy that cannot read data and recreate 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.

#### **Action**

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 recreated.
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, recreate 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.

#### **Action**

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%)

#### **Related tasks**

"MAP 5700: Repair verification" on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **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 VDisks in the same I/O group. Preparing the stopped FlashCopy operations as soon as possible is advised.

#### **Action**

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

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## **1920 Global and Metro Mirror persistent error.**

### **Explanation**

This error might be caused by a problem on the primary cluster, a problem on the secondary cluster, or a problem on the inter-cluster link. The problem might be a failure of a component, a component becoming unavailable or having reduced performance due to a service action or it might be that the performance of a component has dropped to a level where the Metro Mirror or Global Mirror relationship cannot be maintained. Alternatively the error might be caused by a change in the performance requirements of the applications using Metro Mirror or Global Mirror.

This error is reported on the primary cluster when the copy relationship has not progressed sufficiently over a period of time. Therefore, if the relationship is restarted before all of the problems are fixed, the error might be reported again when the time period next expires (the default period is five minutes).

This error might also be reported because the primary cluster has encountered read errors.

You might need to refer to the Copy Services features chapter of the “IBM System Storage SAN Volume Controller Software Installation and Configuration Guide” while diagnosing this error.

### **Action**

1. If the 1920 error has occurred previously on Metro Mirror or Global Mirror between the same clusters and all the following actions have been attempted, contact your product support center to resolve the problem.
2. On the primary cluster reporting the error, correct any higher priority errors.
3. On the secondary cluster, review the maintenance logs to determine if the cluster was operating with reduced capability at the time the error was reported. The reduced capability might be due to a software upgrade, hardware maintenance to a 2145 node, maintenance to a backend disk system or maintenance to the SAN.
4. On the secondary 2145 cluster, correct any errors that are not fixed.
5. On the intercluster link, review the logs of each link component for any incidents that would cause reduced capability at the time of the error. Ensure the problems are fixed.
6. If a reason for the error has been found and corrected, go to Action 10.
7. On the primary cluster reporting the error, examine the 2145 statistics using a SAN productivity monitoring tool and confirm that all the Metro Mirror and Global Mirror requirements described in the “IBM System Storage SAN Volume

Controller Planning Guide" are met. Ensure that any changes to the applications using Metro Mirror or Global Mirror have been taken into account. Resolve any issues.

8. On the secondary cluster, examine the 2145 statistics using a SAN productivity monitoring tool and confirm that all the Metro Mirror and Global Mirror requirements described in the "IBM System Storage SAN Volume Controller Software Installation and Configuration Guide" document are met. Resolve any issues.
9. On the intercluster link, examine the performance of each component using an appropriate SAN productivity monitoring tool to ensure that they are operating as expected. Resolve any issues.
10. 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 will be an initial period during which Metro Mirror or Global Mirror performs a background copy to resynchronize the VDisk data on the primary and secondary clusters. During this period the data on the Metro Mirror or Global Mirror auxiliary VDIs on the secondary cluster is inconsistent and the VDIs could not be used as backup disks by your applications.

Note: To ensure the system has the capacity to handle the background copy load you may want to delay restarting the Metro Mirror or Global Mirror relationship until there is a quiet period when the secondary cluster and the SAN fabric (including the intercluster link) have the required capacity. If the required capacity is not available you might experience another 1920 error and the Metro Mirror or Global Mirror relationship will stop in an inconsistent state.

Note: If the Metro Mirror or Global Mirror relationship has stopped in a consistent state ("consistent-stopped") it is possible to use the data on the Metro Mirror or Global Mirror auxiliary VDIs on the secondary cluster as backup disks by your applications. You might therefore want to start a Flash Copy of your Metro Mirror or Global Mirror auxiliary disks on the secondary system before restarting the Metro Mirror or Global Mirror relationship. This means you maintain the current, consistent, image until the time when the Metro Mirror or Global Mirror relationship is again synchronized and in a consistent state.

Possible Cause-FRUs or other:

- None

Other:

- Primary 2145 cluster or SAN fabric problem (10%)
- Primary 2145 cluster or SAN fabric configuration (10%)
- Secondary 2145 cluster or SAN fabric problem (15%)
- Secondary 2145 cluster or SAN fabric configuration (25%)
- Intercluster link problem (15%)
- Intercluster link configuration (25%)

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## **1930 Migration suspended.**

### **Explanation**

Migration suspended.

### **Action**

1. Ensure that all error codes of a higher priority have already been fixed.
2. Ask the customer to ensure that all MDisk groups 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

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## **1950 Unable to mirror medium error.**

### **Explanation**

During the synchronization of a mirrored Virtual Disk Copy it was necessary to duplicate the record of a medium error onto the Virtual Disk 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 Virtual Disk Copy is in an inconsistent state and has been taken offline.

### **Action**

Three different approaches can be taken to resolving this problem: 1) the source Virtual Disk 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 Virtual Disk 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 Virtual Disk 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 Virtual Disk Copy. Determine where the virtual medium errors are using the error log events and re-write the block or volume from backup.

Approach 3) - Delete the offline Virtual Disk Copy and create a new one either forcing the use of different MDisks in the managed disk group or using a completely different managed disk group.

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.

#### **Action**

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:

2145 software (100%)

### **2010 A software upgrade has failed.**

#### **Explanation**

Cluster configuration changes are restricted until the upgrade is completed or rolled back. The cluster upgrade process waits for user intervention when this error is logged.

#### **Action**

The action required to recover from a stalled upgrade depends on the current state of the cluster being upgraded. Call IBM Support for an action plan to resolve this problem.

Possible Cause-FRUs or other:



- None

Other:

2145 software (100%)

#### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“Viewing the node status” on page 66

You can view the properties for a node from the Viewing General Details panel.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

#### **Related reference**

“Checking the status of the node using the CLI” on page 67

You can use the command-line interface (CLI) to check the status of the node.

### **2030 Software error.**

#### **Explanation**

The 2145 software has restarted due to a problem in the cluster, on a disk system or on the fibre-channel fabric.

#### **Action**

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:

- 2145 software (48%)
- Enclosure/controller software (25%)
- Fibre-channel switch or switch configuration (25%)

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## **2040 A software upgrade 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.

### **Action**

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:

2145 software (100%)

### **Related tasks**

“Marking errors as fixed” on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## **2100 A software error has occurred.**

### **Explanation**

One of the 2145 server software components (sshd, crond, or httpd) has failed and reported an error.

### **Action**

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:

2145 software (100%)

### **Related tasks**

"Marking errors as fixed" on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

"MAP 5000: Start" on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

"MAP 5700: Repair verification" on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## **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.

A separate SSH session limit allocation is maintained for each user type. Examples of user type are "admin" and "service".

### **Action**

- 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.
- Select Launch Maintenance Procedures for the cluster on the Viewing Clusters panel of the IBM System Storage SAN Volume Controller Console.
- Select List the current sessions.
- Either use the maintenance procedures option to terminate all sessions, or close individual sessions that are no longer required. Identify sessions based on the host system and application using the session.

- If you are unable to resolve the error using one of the above options, stop and restart the configuration node to terminate all SSH sessions.
- Mark the error as "fixed".

#### **Related tasks**

Chapter 3, "Using the SAN Volume Controller Console and CLI," on page 63  
 The SAN Volume Controller Console is a Web-browser based GUI that is used to manage the cluster. The SAN Volume Controller command-line interface (CLI) is a collection of commands that you can use to manage SAN Volume Controller clusters.

"Marking errors as fixed" on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

### **2600 The cluster was unable to send an e-mail.**

#### **Explanation**

The cluster has attempted to send an e-mail 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 e-mail 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 e-mail function because it responds immediately with a result code.

#### **Action**

- Ensure that the SMTP e-mail server is active.
- Ensure that the SMTP server TCP/IP address and port are correctly configured in the cluster e-mail configuration.
- Send a test e-mail 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

#### **Related tasks**

"MAP 5700: Repair verification" on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

### **2601 Error detected while sending an e-mail.**

#### **Explanation**

An error has occurred while the cluster was attempting to send an e-mail in response to an event. The cluster is unable to determine if the e-mail has been sent and will attempt to resend it. The problem might be with the SMTP server or with the cluster e-mail configuration. The problem might also be caused by a failover of the configuration node. This error is not logged by the test e-mail function because it responds immediately with a result code.

### **Action**

- If there are higher-priority unfixed errors in the log, fix those errors first.
- Ensure that the SMTP e-mail server is active.
- Ensure that the SMTP server TCP/IP address and port are correctly configured in the cluster e-mail configuration.
- Send a test e-mail 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

### **Related tasks**

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## **2700 Unable to access NTP network time server**

### **Explanation**

Cluster time cannot be synchronized with the NTP network time server that is configured.

### **Action**

There are three main causes to examine:

- The SVC 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 SAN Volume Controller cluster can access the NTP network time server and that the NTP protocol is permitted.

The error will automatically fix when the SAN Volume Controller is able to synchronize its time with the NTP network time server.

Possible Cause-FRUs or other:

- None

## **3000 The 2145 UPS temperature is close to its upper limit. If the temperature continues to rise the 2145 UPS will power off.**

### **Explanation**

The temperature sensor in the 2145 UPS is reporting a temperature that is close to the operational limit of the unit. If the temperature continues to rise the 2145 UPS will power off for safety reasons. The sensor is probably reporting an excessively high temperature because the environment in which the 2145 UPS is operating is too hot.

### **Action**

1. Ensure that the room ambient temperature is within the permitted limits.

2. Ensure that the air vents at the front and back of the 2145 UPS are not obstructed.
3. Ensure that other devices in the same rack are not overheating.
4. When you are satisfied that the cause of the overheating has been resolved, mark the error "fixed".

#### **Related tasks**

"Marking errors as fixed" on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

### **3001 The 2145 UPS-1U temperature is close to its upper limit. If the temperature continues to rise the 2145 UPS-1U will power off.**

#### **Explanation**

The temperature sensor in the 2145 UPS-1U is reporting a temperature that is close to the operational limit of the unit. If the temperature continues to rise the 2145 UPS-1U will power off for safety reasons. The sensor is probably reporting an excessively high temperature because the environment in which the 2145 UPS-1U is operating is too hot.

#### **Action**

1. Ensure that the room ambient temperature is within the permitted limits.
2. Ensure that the air vents at the front and back of the 2145 UPS-1U are not obstructed.
3. Ensure that other devices in the same rack are not overheating.
4. When you are satisfied that the cause of the overheating has been resolved, mark the error "fixed".

#### **Related tasks**

"Marking errors as fixed" on page 122

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

### **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.

#### **Action**

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 down, 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%)

### **Related tasks**

“MAP 5350: Powering off a SAN Volume Controller node” on page 385  
MAP 5350: Powering off a SAN Volume Controller node helps you power off a single SAN Volume Controller node to perform a service action without disrupting the host’s access to disks.

## **3025 A virtualization feature license is required.**

### **Explanation**

The SAN Volume Controller has no virtualization feature license registered. You should have either a SAN Volume Controller Entry Edition Physical Disk virtualization feature license or a SAN Volume Controller Capacity virtualization feature license that covers the cluster.

The SAN Volume Controller will continue to operate, but it might be violating the license conditions.

### **Action**

- 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 SAN Volume Controller 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.

### **Action**

- 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 cluster, but the capacity has not been set.

This error event is created when a SAN Volume Controller 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 feature can only be set to On or Off; with version 4.3.0 and later the licensed capacity for the feature must also be set.

### **Action**

Perform one of the following actions:

- Change the Global and Metro Mirror license settings for the cluster either to the licensed Global and Metro Mirror capacity, or if the license applies to more than one cluster, to the portion of the license allocated to this cluster. 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 cluster, but the capacity has not been set.

This error event is created when a SAN Volume Controller 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 feature can only be set to On or Off; with version 4.3.0 and later the licensed capacity for the feature must also be set.



## Action

Perform one of the following actions:

- Change the FlashCopy license settings for the cluster either to the licensed FlashCopy capacity, or if the license applies to more than one cluster, to the portion of the license allocated to this cluster. 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 SAN Volume Controller feature is being exceeded.

The feature that is being exceeded might be:

- Virtualization feature - event identifier 009172
- FlashCopy feature - event identifier 009173
- Global and Metro Mirror feature - event identifier 009174

The SAN Volume Controller will continue to operate, but it might be violating the license conditions.

## Action

- Determine which feature license limit has been exceeded. This might be:
- Virtualization feature - event identifier 009172
- FlashCopy feature - event identifier 009173
- Global and Metro Mirror feature - event identifier 009174
- 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 VDisks. 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 VDisks that are the source VDisk 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 VDisks that are in a Metro Mirror or Global Mirror relationship; both master and auxiliary VDisks are counted.
- 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 SAN Volume Controller Entry Edition cluster has some FlashCopy mappings defined. There is, however, no SAN Volume Controller Physical Disk FlashCopy license registered on the cluster. The SAN Volume Controller will continue to operate, but it might be violating the license conditions.

#### **Action**

- Check if you have a SAN Volume Controller 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 SAN Volume Controller Entry Edition cluster has some Global Mirror or Metro Mirror relationships defined. There is, however, no SAN Volume Controller Physical Disk Global and Metro Mirror license registered on the cluster. The SAN Volume Controller will continue to operate, but it might be violating the license conditions.

#### **Action**

- Check if you have a SAN Volume Controller 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

### **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 SAN Volume Controller cluster whose partnership is deleted.

Beginning with SAN Volume Controller version 4.3.1 this configuration is not supported and should be resolved. This condition can occur as a result of an upgrade to SAN Volume Controller version 4.3.1 or later.

#### **Action**

The issue can be resolved either by deleting all of the Global Mirror or Metro Mirror relationships or consistency groups that exist with a SAN Volume Controller 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.

- 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.
- 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.
- 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 SAN Volume Controller limit. In version 4.3.1 this limit is 1. 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.
- Re-establish any selected partnerships.
- 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.
- 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 e-mail to any of the configured e-mail servers.**

### **Explanation**

Either the system was not able to connect to any of the SMTP e-mail servers, or the e-mail transmission has failed. A maximum of six e-mail servers can be configured. Error event 2600 or 2601 is raised when an individual e-mail server is found to be not working. This error indicates that all of the e-mail servers were found to be not working.

### **Action**

- Check the error 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 e-mail function to test that an e-mail server is operating properly.

Possible Cause-FRUs or other:

- None

---

## **SAN problem determination**

The procedures to service the SAN Volume Controller that are provided here help you solve problems on the SAN Volume Controller and its connection to the storage area network (SAN).

SAN failures might cause the SAN Volume Controller cluster to be unable to form or they might cause SAN Volume Controller disks 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 switch
- Application-specific integrated circuits
- Installed small form-factor pluggable (SFP) connector
- Fiber-optic cables

Perform the following steps if you were sent here from either the Maintenance Analysis Procedures or the error codes:

1. If the customer has changed the SAN configuration by changing the fibre-channel cable connections or switch zoning, ask the customer to verify that the changes were correct and, if necessary, reverse those changes.
2. Verify that the power is turned on to all switches and redundant array of independent disk (RAID) controllers that the SAN Volume Controller uses and that they are not reporting any hardware failures. If problems are found, resolve those problems before proceeding further.
3. Verify that the fibre-channel cables that connect the SAN Volume Controllers to the switches are securely connected.
4. If the customer is running a SAN management tool that you are familiar with and that you have access to, you can use that tool to view the SAN topology

and isolate the failing component. You can also review “Viewing the fibre-channel fabric connections” on page 101 to understand the SAN Volume Controller view of the SAN.

#### **Related tasks**

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“Viewing the fibre-channel fabric connections” on page 101

Failures of the SAN Volume Controller hardware, fibre-channel cables, fibre-channel switches, fibre-channel hardware in host systems, or disk controllers can cause fibre-channel connectivity problems.

---

## **Fibre-channel link failures**

When a failure occurs on a single fibre channel link, the small form-factor pluggable (SFP) connector might need to be replaced.

The following items can indicate that a single fibre-channel link has failed:

- The customer’s SAN monitoring tools
- The fibre-channel port status on the front panel of the SAN Volume Controller node
- The fibre-channel status LEDs at the rear of the SAN Volume Controller node
- A SAN Volume Controller cluster, node, or boot error that indicates a single port has failed

Attempt each of the following actions, in the following order, until the failure is fixed:

1. Ensure that the fibre channel cable is securely connected at each end.
2. Replace the fibre channel cable.
3. Replace the SFP connector for the failing port on the SAN Volume Controller node.

**Note:** SAN Volume Controller nodes are supported with both long-wave SFPs and short-wave SFPs. You must replace an SFP with the same type of SFP that you are replacing. If the SFP to replace is a long wave SFP, for example, you must provide a suitable replacement. Removing the wrong SFP connector could result in loss of data access.

4. Have the customer perform the fibre-channel switch service procedures for a failing fibre channel link. This might involve replacing the SFP connector at the switch.
5. Replace the fibre-channel adapter on the SAN Volume Controller 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 SAN Volume Controller, or the host.

For network problems, you can attempt any of the following items:

- Test your connectivity between the host and SAN Volume Controller ports.
- Try to ping the SAN Volume Controller 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 SAN Volume Controller host configuration.

For SAN Volume Controller problems, you can attempt any of the following items:

- View the configured node port IP addresses by using the `svcinfolshost` CLI command.
- View the list of VDisks mapped to a host by using the `svcinfolshostvdiskmap` command to ensure that the VDisk host mappings are correct.
- Verify that the VDisk is online by using the `svcinfolsvdisk` command.

For host problems, you can attempt any of the following items:

- Verify that the host 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.

---

## Servicing storage systems

Storage systems that are supported for attachment to the SAN Volume Controller 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 guidelines apply to all storage systems that are attached to the SAN Volume Controller:

- Always follow the service instructions that are provided in the documentation for your storage system.
- Ensure that there are no unfixed errors in the SAN Volume Controller error log before you perform any maintenance procedures.
- After you perform a maintenance procedure, check the SAN Volume Controller error log and fix any errors. Expect to see the following types of errors:
  - MDisk error recovery procedures (ERPs)
  - Reduced paths

The following categories represent the types of service actions for storage systems:

- Controller code upgrade
- Field replaceable unit (FRU) replacement

### Controller code upgrade

Ensure that you are familiar with the following guidelines for upgrading controller code:

- Check to see if the SAN Volume Controller supports concurrent maintenance for your storage system.
- Allow the storage system to coordinate the entire upgrade process.
- If it is not possible to allow the storage system to coordinate the entire upgrade process, perform the following steps:
  1. Reduce the storage system workload by 50%.
  2. Use the configuration tools for the storage system to manually failover all logical units (LUs) from the controller that you want to upgrade.
  3. Upgrade 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 upgrade. To disable external data paths, disconnect or disable the appropriate ports on the fabric switch. The SAN Volume Controller 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 disk 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 additional steps.

---

## Troubleshooting the SAN Volume Controller Console

These topics provide information that can help you troubleshoot and resolve problems with the SAN Volume Controller Console, which is installed on the IBM System Storage Productivity Center (SSPC) or, in previous releases, on the master console server.

In addition to troubleshooting on your own, you can also request an Assist On-site session with an IBM service representative.

Use the following topics to resolve problems with the SAN Volume Controller Console.

## Checking connections to the cluster when the status is No Contact

If the Availability Status on the Viewing Clusters page is No Contact, the cluster either does not exist or is offline, or there is a network problem.

Also, node error 540 might be displayed on the front panel on the SAN Volume Controller node.

Perform the following actions to make the cluster available:

**Note:** Click **Refresh** after each step.

1. Restart the IBM System Storage Productivity Center (SSPC) or master console server.
2. If the Availability status is still No Contact, open a command prompt from the SSPC or master console server and ping the IP address of the cluster. You can view the IP address either from the SAN Volume Controller Console screens or the front panel of the SAN Volume Controller node.
3. If you can successfully ping the SSPC or master console server, run the CLI command to check the cluster status. Otherwise, see “MAP 5500: Ethernet”. If those procedures do not solve the problem, contact the IBM Support Center for assistance.
4. If the Availability status is still No Contact, the connection between the SAN Volume Controller Console and the cluster might have been corrupted. Contact the IBM Support Center.

## Cannot access the SAN Volume Controller Console login screen

When you cannot access the SAN Volume Controller Console login screen, but you can access the cluster through PuTTY, a default IP address might be set in the hosts file.

Perform the following actions to access the SAN Volume Controller Console login screen:

1. From the IBM System Storage Productivity Center (SSPC) or master console server, go to **Computer management** → **Manage** → **Services and applications** → **Services**. To access **Computer management**, right click on My computer.
2. Check the status of Tivoli NetView services. Ensure that this service has a status of Stopped and is configured for Manual Startup. If the service was started, the host file on the SSPC or master console server might have a default IP address.
3. Check the hosts file to ensure that the SSPC or master console IP address has not changed to a default IP address.
4. Edit the hosts file if it lists a default IP address. The hosts file is located in the Microsoft Windows \system32\drivers\etc directory.
5. Open a new browser from the SAN Volume Controller Console and try to access the login screen again.

## Cannot access the cluster from the SAN Volume Controller Console

When you cannot access the SAN Volume Controller cluster from the SAN Volume Controller Console, but you can access the cluster through PuTTY, some of the required services on the IBM System Storage Productivity Center (SSPC) or master console server might be stopped.

Check the messages that begin with the CMMU prefix in the SAN Volume Controller Information Center. These messages might provide additional information and solutions for this problem.

Perform the following actions to access the cluster from the SAN Volume Controller Console:

1. From the SSPC or master console server, right-click **My Computer** on your desktop.
2. Click **Computer management** → **Manage** → **Services and applications** → **Services**.
3. Ensure that the following services are configured for Automatic Startup with a status of Started:
  - Service Location Protocol
  - IBM System Storage SAN Volume Controller Pegasus Server
  - IBM WebSphere® Application Server V6 - SVC

**Note:** If the status for any of these services changes from Starting to Stopped, contact the IBM Support Center for assistance.

## Troubleshooting unexpected shutdowns of the SAN Volume Controller Console

If you are working with the SAN Volume Controller Console and you receive a You have signed off message before the SAN Volume Controller Console closes unexpectedly, use these instructions to help you troubleshoot the problem.



You can perform any of the following actions to troubleshoot an unexpected shutdown of the SAN Volume Controller Console:

- Open a new browser window and try to reconnect to the SAN Volume Controller Console. The logoff message is typically caused when an open session times out. This often happens if the browser window was left open from a previous session.
- Check Windows Task Manager to ensure that the cimserver.exe process is running.
- Ensure that the Websphere Application Server (WAS) service is still running in the Windows Service Manager.
- Ensure that the disk on the server is not full.
- Ensure that the server is not pegged.
- Determine if the IP address of the server where the SAN Volume Controller Console is running has changed since the last time that the server was restarted. If it has changed, restart the server to correct the problem.

## WebSphere Application Server service keeps stopping

The IBM WebSphere Application Server V6 - SVC service might stop shortly after it is started if a conflict is detected because another application that is running on the IBM System Storage Productivity Center or master console server is using the same port.

To determine if another application is also using port 9100, issue the following command from a Windows command prompt:

```
netstat -an | find "9100"
```

If any output is returned, another application is using the default WebSphere Application Server (WAS) port.

To resolve the problem, complete the following steps:

1. Identify the application that is conflicting with WAS.
2. Either shut down the conflicting application or change the port that the WAS service uses.
3. Complete the following steps if you decide to change the WAS configuration and use a different port:
  - a. Change to the `\Program Files\IBM\svconsole\console\embeddedWAS\profiles\SVCProfile` directory.
  - b. Replace all entries that have 9100 with an alternate port number in the following files:  
`\config\cells\SVCCell\nodes\SVCCell\serverindex.xml`  
`\config\templates\servertypes\APPLICATION_SERVER\serverindex.xml`  
`\properties\portdef.props`
  - c. Save your changes and then restart the WAS service:
    - 1) Right-click on **My Computer** and then click **Manage**. The Computer Management window is displayed.
    - 2) Expand **Services and Applications** and then click **Services**. The list of services is displayed.
    - 3) Right-click on **IBM WebSphere Application Server V6 - SVC** and then click **Stop**. After it stops, right-click on the service again and click **Start**.

## Websphere and CIM Logging

You can obtain log files for the Websphere Application Server and the Common Information Model (CIM). Instructions are supplied in the following topics.

### Websphere Application Server logging

The Websphere Application Server (WAS) produces log files from the IBM System Storage Productivity Center or master console that can help with problem determination.

The WAS collects trace data and writes the information to log files that are stored in the `svcconsole\console\embeddedWAS\profiles\SVCPProfile\logs\server1` directory. The default installation directory is `Program Files\IBM\svcconsole`.

### Enabling Websphere Application Server logging

You can enable WAS logging and create the `trace.log` file by performing the following steps:

1. Open the command-line application.
2. Issue the following command:

```
enableConsoleTrace
```

You can run the `enableConsoleTrace` command from any directory as long as `svcconsole\console\embeddedWAS\bin` is in your path.

A successfully enabled logging session is indicated by an entry in the `trace.log` file that is similar to the following:

```
[5/21/03 14:31:13:874 PDT] 2ff3581b ManagerAdmin I TRAS0018I: The trace state has changed. The new trace state is ConsoleTrace=all=enabled.
```

### Disabling Websphere Application Server logging

You can disable the WAS logging by following these steps:

1. Open the command-line application and change to the WAS bin directory:

```
cd svcconsole\console\embeddedWAS\bin
```

2. Issue either one of the following commands:

- `disableConsoleTrace.bat`
- `wsadmin -connType NONE -c "$AdminControl setAttribute [$AdminControl completeObjectName type=TraceService,process=server1,*] traceSpecification ConsoleTrace=event=disabled:ConsoleTrace=debug=disabled"`

A successfully disabled logging session is indicated with an entry in the `SystemOut.log` file similar to the following:

```
[5/21/03 14:38:57:400 PDT] 2ff3581b ManagerAdmin I TRAS0018I: The trace state has changed. The new trace state is *=all=disabled
```

**Note:** Logging has minimal impact on the performance of the graphical user interface application.

### Common information model provider logging

The Common Information Model (CIM) can produce log files that can help with problem determination.

### **Collecting CIM log files:**

The SAN Volume Controller CIM agent command-line interface (CLI) supports the collection of CIM agent logging information for maintenance and diagnostic reporting to the IBM Support Center.

To collect logs from the CIMOM that resides on the cluster, typical cluster log collection is used. Log levels can be set using the SAN Volume Controller Console by selecting **Service and Maintenance > CIMOM Log Configuration**.

To collect all cluster logs and CIMOM logs, issue the `svctask cpdumps` CLI command. You can also use the SAN Volume Controller Console to list the CIMOM logs by selecting **CIMOM Logs > Service and Maintenance > List Dumps**.

1. Increase the tracing levels:
  - a. Launch the SAN Volume Controller Console
  - b. Select **Service and Maintenance > CIMOM Log Configuration**
  - c. To enable CIMOM trace log at the highest level, select **Max Level**
2. Reproduce the error.
3. Collect CIMOM logs:
  - a. Launch the SAN Volume Controller Console
  - b. Select **Service and Maintenance > List Dumps**
  - c. Click **CIMOM Logs** and save all CIMOM log files in the list
4. After you collect the information, decrease the tracing levels and return the server to its typical performance level:
  - a. Launch the SAN Volume Controller Console
  - b. Select **Service and Maintenance > CIMOM Log Configuration**
  - c. Choose the **Default Level** to decrease the CIMOM trace log level

Table 53 describes the commands that are used to log and trace CIM agent activity.

*Table 53. Commands for logging and tracing CIM agent*

<b>Command</b>	<b>Description</b>
collectlogs	Collects the SAN Volume Controller GUI logs



---

## Chapter 7. 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.

SAN Volume Controller nodes must be configured in pairs so you can perform concurrent maintenance.

When you service one SAN Volume Controller 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 SAN Volume Controller 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 SAN Volume Controller nodes unless instructed to do so. When you need to remove power, see “MAP 5350: Powering off a SAN Volume Controller node” on page 385.

- To isolate the FRUs in the failing SAN Volume Controller node, complete the actions and answer the questions 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 have been exchanged.
  5. Complete the next action indicated by the MAP.
  6. If you are using one or more MAPs because of a cluster error code, mark the error as fixed in the cluster error log after the repair, but before you verify the repair.

**Note:** Start all problem determination procedures and repair procedures with “MAP 5000: Start.”

### Related tasks

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

“MAP 5800: Light path” on page 406

MAP 5800: Light path helps you to solve hardware problems on the SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2 that are preventing the node from booting.

---

## MAP 5000: Start

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 7, “Using the maintenance analysis procedures,” on page 357.

This MAP applies to all SAN Volume Controller 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 have been sent here for one of the following reasons:

- The web-based Directed Maintenance procedure sent you here
- A problem occurred during the installation of a SAN Volume Controller
- Another MAP sent you here
- A user observed a problem that was not detected by the system

SAN Volume Controller nodes are configured in pairs. While you service one SAN Volume Controller 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 SAN Volume Controller 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 SAN Volume Controller nodes unless instructed to do so.
- If a recommended action in these procedures involves removing or replacing a part, use the applicable procedure in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- If the problem persists after performing the recommended actions in this procedure, return to step 1 of the MAP to try again to fix the problem.

Perform the following steps:

1. **Were you sent here from a Directed Maintenance procedure?**

**NO** Go to step 2

**YES** Go to step 8 on page 359

2. (from step 1)

Find the IBM System Storage Productivity Center (SSPC) that is close to and is set up to manage the SAN Volume Controller cluster. The SSPC is normally located in the same rack as the SAN Volume Controller cluster.

3. (from step 2)

Log in to the SSPC or master console server using the user ID and password that is provided by the user.

4. (from step 3)

Log into the SAN Volume Controller Console using the user ID and password that is provided by the user and launch the SAN Volume Controller Console for the cluster that you are repairing.

5. (from step 4)

**Does the SAN Volume Controller Console start?**

**NO** Go to step 8 on page 359. See “Troubleshooting the SAN Volume Controller Console” on page 351.

**YES** Go to step 6.

6. (from step 5)

**When the SAN Volume Controller cluster that you want to service is selected, is the Welcome panel displayed?**

**NO** Go to step 8.

**YES** Go to step 7.

7. (from step 6 on page 358)

Start the Directed Maintenance Procedures.

**Did the maintenance procedures find an error that needs to be fixed?**

**NO** Go to step 8.

**YES** Follow the Directed Maintenance procedures.

8. (from steps 1 on page 358, 5 on page 358, 6 on page 358, and 7)

**Is the power indicator on any SAN Volume Controller front panel off?** On the SAN Volume Controller models 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, and 2145-8F2 check to see if the power LED **1** on the operator-information panel is off.




Figure 78. Power LED on the SAN Volume Controller models 2145-CF8, 2145-8A4, 2145-8G4, and 2145-8F4 or 2145-8F2 operator-information panel

**1** SAN Volume Controller models 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, and 2145-8F2 Power LED

**NO** Go to step 9.

**YES** Try to turn on the SAN Volume Controller nodes. See “Using the power control for the SAN Volume Controller node” on page 165.

**Note:** The uninterruptible power supply unit that supplies power to the SAN Volume Controller node might also be turned off. This must be turned on before the SAN Volume Controller node is turned on.

If the SAN Volume Controller nodes are turned on, go to step 9; otherwise, go to the appropriate Power MAP: “MAP 5060: Power 2145-8A4” on page 372, “MAP 5050: Power 2145-CF8, 2145-8G4, 2145-8F4, and 2145-8F2” on page 365.

9. (from step 8)

**Does the front panel of any SAN Volume Controller node show a hardware error?** There is a hardware error if any of the following conditions are true for the node:

- None of the LEDs on the front panel are on and the front-panel display is blank.

- The node is a SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2 and the error LED, which is the bottom LED on the front panel, is on.

Figure 79 shows the location of the service controller error light.




Figure 79. SAN Volume Controller service controller error light

**1** SAN Volume Controller models 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, and 2145-8F2 service-controller error light (the bottom LED)

**NO** Go to step 10.

**YES** The service controller for the SAN Volume Controller has failed.

- Check that the service controller that is indicating an error is correctly installed. If it is, replace the service controller.
- Go to “MAP 5700: Repair verification” on page 404.

10. Are you working on the SAN Volume Controller 2145-4F2?

**NO** Go to step 11.

**YES** Go to step 12.

11. (from step 10)

**Is the operator-information panel error LED **1** that you see in Figure 80 illuminated or flashing?**




Figure 80. Error LED on the SAN Volume Controller models 2145-CF8, 2145-8A4, 2145-8G4, and 2145-8F4 or 2145-8F2 operator-information panel

**NO** Go to step 12.

**YES** Go to “MAP 5800: Light path” on page 406.

12. (from step 10 and step 11)

**Is the hardware boot display that you see in Figure 81 on page 361 displayed on any of the SAN Volume Controller nodes?**





Figure 81. Hardware boot display

- NO** Go to step 14.
- YES** Go to step 13.
13. (from step 12 on page 360)  
**Has the hardware boot display that you see in Figure 81 displayed for more than three minutes?**
- NO** Go to step 14.
- YES** Perform the following:
- Go to “MAP 5900: Hardware boot” on page 423.
  - Go to “MAP 5700: Repair verification” on page 404.
14. (from steps 12 on page 360 and 13)  
**Is Failed displayed on the top line of the front-panel display of any SAN Volume Controller node?**
- NO** Go to step 15.
- YES** Perform the following:
- Note the failure code and go to boot code “Understanding the boot codes” on page 199 to perform the repair actions.
  - Go to “MAP 5700: Repair verification” on page 404.
15. (from step 14)  
**Is Booting displayed on the top line of the front-panel display of any SAN Volume Controller node?**
- NO** Go to step 17.
- YES** Go to step 16.
16. (from step 15)  
A progress bar and a boot code are displayed. If the progress bar does not advance for more than three minutes, it has stalled.  
**Has the progress bar stalled?**
- NO** Go to step 17.
- YES** Perform the following:
- Note the failure code and go to boot code “Understanding the boot codes” on page 199 to perform the repair actions.
  - Go to “MAP 5700: Repair verification” on page 404.
17. (from steps 15 and step 16)  
If you pressed any of the navigation buttons on the front panel, wait for 60 seconds to ensure that the display has switched to its default display.  
**Is Node Error displayed on the top line of the front-panel display of any SAN Volume Controller node?**
- NO** Go to step 18 on page 362.
- YES** Perform the following steps:

- a. Note the failure code and go to node error code “Understanding the node error codes” on page 223 to perform the repair actions.
  - b. Go to “MAP 5700: Repair verification” on page 404.
18. (from step 17 on page 361)  
**Is Cluster Error displayed on the top line of the front-panel display of any SAN Volume Controller node?**
  - NO** Go to step 19.
  - YES** A cluster error was detected. This error code is displayed on all the operational nodes in the cluster. This type of error is normally repaired using the Directed Maintenance procedures. Perform the following steps:
    - a. Go to step 2 on page 358 to perform the Directed Maintenance procedure. If you return here, go to cluster error code “Understanding cluster error codes” on page 239 to perform the repair actions.
    - b. Go to “MAP 5700: Repair verification” on page 404.
19. (from step 18)  
**Is Powering Off, Restarting, Shutting Down, or Power Failure displayed in the top line of the front-panel display?**
  - NO** Go to step 21.
  - YES** The progress bar moves every few seconds. Wait for the operation to complete and then return to step 1 on page 358 in this MAP. If the progress bar does not move for three minutes, press the power button and go to step 20.
20. (from step 19)  
**Did the SAN Volume Controller node power off?**
  - NO** Perform 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 the power button to power-on the SAN Volume Controller node and then return to step 1 on page 358 in this MAP.
  - YES** Perform the following steps:
    - a. Wait 60 seconds.
    - b. Click the power button to turn on the SAN Volume Controller node and then return to step 1 on page 358 in this MAP.

**Note:** The 2145 UPS-1U turns off only when its power button is pressed, input power has been lost for more than five minutes, or the SAN Volume Controller node has shut it down following a reported loss of input power.
21. (from step 20)  
**Is Charging or Recovering displayed in the top line of the front-panel display of any SAN Volume Controller node?**
  - NO** Go to step 22 on page 363.
  - YES**

- If Charging is displayed, the uninterruptible power supply battery is not yet charged sufficiently to support the SAN Volume Controller node. If Charging is displayed for more than two hours, go to “MAP 5150: 2145 UPS-1U” on page 376.
- If Recovering is displayed, the uninterruptible power supply battery is not yet charged sufficiently to be able to support the SAN Volume Controller node immediately following a power supply failure. However, if Recovering is displayed, the SAN Volume Controller node can be used normally.
- If Recovering is displayed for more than two hours, go to “MAP 5150: 2145 UPS-1U” on page 376.

22. (from step 21 on page 362)

**Is Validate WWNN? displayed on the front-panel display of any SAN Volume Controller node?**

**NO** Go to step 23 on page 364.

**YES** The node is indicating that its WWNN might need changing. It enters this mode when the node service controller or disk has been changed but the required service procedures have not been followed.

**Note:** Do not validate the WWNN until you read the following information to ensure that you choose the correct value. If you choose an incorrect value, you might find that the SAN zoning for the node is also not correct and more than one node is using the same WWNN. Therefore, it is important to establish the correct WWNN before you continue.

- Determine which WWNN that you want to use.
  - If the service controller has been replaced, the correct value is probably the WWNN that is stored on disk (the disk WWNN).
  - If the disk has been replaced, perhaps as part of a frame replacement procedure, but has not been re-initialized, the correct value is probably the WWNN that is stored on the service controller (the panel WWNN).
- Select the stored WWNN that you want this node to use:
  - To use the WWNN that is stored on the disk, perform the following steps:
    - 1) From the Validate WWNN? panel, press and release the select button. The Disk WWNN: panel is displayed and shows the last five digits of the WWNN that is stored on the disk.
    - 2) From the Disk WWNN: panel, press and release the down button. The Use Disk WWNN? panel is displayed.
    - 3) Press and release the select button.
  - To use the WWNN that is stored on the service controller, perform the following steps:
    - 1) From the Validate WWNN? panel, press and release the select button. The Disk WWNN: panel is displayed.
    - 2) From the Disk WWNN: panel, press and release the right button. The Panel WWNN: panel is displayed and shows the last five numbers of the WWNN that is stored on the service controller.

- 3) From the Panel WWNN: panel, press and release the down button. The Use Panel WWNN? panel is displayed.
  - 4) Press and release the select button.
- c. After you set the WWNN, check the front-panel display:
- If the Node WWNN: panel is displayed on the front panel, the node is now using the selected WWNN. The Node WWNN: panel shows the last five numbers of the WWNN that you selected.
  - If the front panel shows Cluster: but does not show a cluster name, you must use the console application to delete the node from the cluster and add the node back into the cluster.

23. (from step 22 on page 363)

**Is there a node that is not a member of a cluster?** You can tell if a node is not a member of a cluster by checking the front panel cluster menu. If Cluster: is displayed but no cluster name is shown, the node is not a member of a cluster. (The cluster name is on the second line of the front-panel display if the current language font allows a two-line display. Otherwise, you can press the select button to display the cluster name.)

**NO** Go to step 24.

**YES** The node is not a member of a cluster. The node might have been deleted from the cluster during a maintenance procedure and has not been added back into the cluster. Make sure that each I/O group in the cluster contains two nodes. If an I/O group in the cluster has only one node, add the node back into that cluster and ensure that the node is restored to the same I/O group that it was deleted from.

24. (from step 23)

**Is the front-panel display unreadable?**

**NO** Go to step 25.

**YES** Perform the following steps:

- a. Check the language. The display might be set to another language.
- b. If the language is set correctly, go to "MAP 5400: Front panel" on page 390.

25. (from step 24)

No errors were detected by the SAN Volume Controller. If you suspect that the problem that is reported by the customer is a hardware problem, perform the following tasks:

- a. Perform Problem Determination procedures on your host systems, disk controllers, and fibre-channel switches.
- b. Ask your hardware support center for assistance.

If you suspect that the problem is a software problem, see "Upgrading the SAN Volume Controller software" in the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*.

If the problem is still not fixed, collect diagnostic information and contact the IBM support center.

### **Related concepts**

“Validate WWNN?” on page 143

The front panel prompts you to validate the WWNN when the worldwide node name (WWNN) that is stored in the service controller (the panel WWNN) does not match the WWNN that is backed up on the SAN Volume Controller disk (the disk WWNN).

### **Related tasks**

“Using directed maintenance procedures” on page 123

You can use directed maintenance procedures (DMP) to diagnose and resolve problems with the SAN Volume Controller.

“Select language? option” on page 159

You can change the language that displays on the front panel.

“Determining a hardware boot failure” on page 199

If you see that the hardware boot display stays on the front panel for more than three minutes, the node cannot boot. The cause might be a hardware failure or the software on the hard disk drive is missing or damaged.

“Fibre-channel link failures” on page 349

When a failure occurs on a single fibre channel link, the small form-factor pluggable (SFP) connector might need to be replaced.

Chapter 7, “Using the maintenance analysis procedures,” on page 357

The maintenance analysis procedures (MAPs) inform you how to analyze a failure that occurs with a SAN Volume Controller node.

“Viewing the vital product data” on page 127

You can view the vital product data for a node from the Viewing Vital Product Data panel of the SAN Volume Controller Console.

### **Related reference**

“Displaying the vital product data using the CLI” on page 128

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

“Troubleshooting the SAN Volume Controller Console” on page 351

These topics provide information that can help you troubleshoot and resolve problems with the SAN Volume Controller Console, which is installed on the IBM System Storage Productivity Center (SSPC) or, in previous releases, on the master console server.

---

## **MAP 5050: Power 2145-CF8, 2145-8G4, 2145-8F4, and 2145-8F2**

MAP 5050: Power 2145-CF8, 2145-8G4, 2145-8F4, and 2145-8F2 helps you to solve power problems that have occurred on SAN Volume Controller models 2145-CF8, 2145-8G4, 2145-8F4, and 2145-8F2. If you are using a SAN Volume Controller 2145-8A4, see the Power MAP for that SAN Volume Controller model.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 7, “Using the maintenance analysis procedures,” on page 357.

You might have been sent here for one of the following reasons:

- A problem occurred during the installation of a SAN Volume Controller 2145-CF8, 2145-8G4, 2145-8F4, or 2145-8F2 node.
- The power switch failed to turn the node on
- The power switch failed to turn the node off
- Another MAP sent you here

Perform the following steps:

1. **Are you here because the node is not powered on?**

**NO** Go to step 10 on page 370.

**YES** Go to step 2.

2. (from step 1)

**Is the power LED on the operator-information panel continuously illuminated?** Figure 82 shows the location of the power LED **1** on the operator-information panel.




Figure 82. Power LED on the SAN Volume Controller models 2145-CF8, 2145-8G4, and 2145-8F4 or 2145-8F2 operator-information panel

**NO** Go to step 3.

**YES** The node is powered on correctly. Reassess the symptoms and return to “MAP 5000: Start” on page 357 or go to “MAP 5700: Repair verification” on page 404 to verify the correct operation.

3. (from step 2)

**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 approximately once per second, then go to step 5 on page 367.

If this behavior persists for more than three minutes, perform the following procedure:

- a. Remove all input power from the SAN Volume Controller node by removing the power retention brackets and the power cords from the back of the node. See “Removing the cable-retention brackets” in the *IBM System Storage SAN Volume Controller Model 2145-CF8 Hardware Installation Guide*.
- b. Wait one minute and then verify that all power LEDs on the node are extinguished.
- c. Reinsert the power cords and power retention brackets.
- d. Wait for the flashing rate of the power LED to slow down to one flash per second. Go to step 5 on page 367.
- 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” on page 404.

4. (from step 3)

**Is the Power LED on the operator-information panel flashing approximately once per second?**

**YES** The node is in standby mode. Input power is present. Go to step 5.

**NO** Go to step 5.

5. (from step 3 on page 366 and step 4 on page 366)

Press the power-on button 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 you are working on a SAN Volume Controller 2145-CF8 and the node still fails to power on, replace parts in the following sequence:

- a. Operator-information panel assembly
- b. System board

If you are working on a SAN Volume Controller 2145-8G4, verify that the operator-information panel cable is correctly 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

If the SAN Volume Controller 2145-8F4 or SAN Volume Controller 2145-8F2 node still fails to power on, replace parts in the following sequence:

- a. Operator-information panel
- b. Cable, signal, front panel
- c. Frame assembly

Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

**YES** The power-on indicator on the operator-information panel shows that the node has successfully powered on. Continue with “MAP 5700: Repair verification” on page 404 to verify the correct operation.

6. (from step 3 on page 366 and step 4 on page 366)

**Is the rear panel power LED on or flashing?** Figure 83 on page 368 shows the location of the power LED **1** on the rear panel of the 2145-8G4, 2145-8F4, or 2145-8F2 nodes. Figure 84 on page 368 shows the location of the power LED **1** on the 2145-CF8.




Figure 83. Power LED on the SAN Volume Controller models 2145-8G4, 2145-8F4, and 2145-8F2 rear panel




Figure 84. Power LED indicator on the rear panel of the SAN Volume Controller 2145-CF8

**NO** Go to step 7 on page 369.

**YES** The operator-information panel is failing.

Verify that the operator-information panel cable is seated on the system board.

If you are working on a SAN Volume Controller 2145-CF8, and the node still fails to power on, replace parts in the following sequence:

- a. Operator-information panel assembly
- b. System board

If you are working on a SAN Volume Controller 2145-8G4, verify that the operator-information panel cable is correctly seated on the system board. If the SAN Volume Controller 2145-8G4 still fails to power on, replace parts in the following sequence:

- a. Operator-information panel assembly
- b. System board



If you are working on a SAN Volume Controller 2145-8F4 or SAN Volume Controller 2145-8F2, verify that the operator-information panel cable is correctly seated at both ends. If the cable is correctly seated and the operator-information panel power light is still not on or blinking, replace the parts in the following sequence:

- a. Operator-information panel
- b. Cable, signal, front panel
- c. Frame assembly

7. (from step 6 on page 367)

Locate the 2145 UPS-1U (2145 UPS-1U) that is connected to this node.

**Does the 2145 UPS-1U that is powering this node have its power on and is its load segment 2 indicator a solid green?**

**NO** Go to “MAP 5150: 2145 UPS-1U” on page 376.

**YES** Go to step 8.

8. (from step 7)

**Are the ac LED indicators on the rear of the power supply assemblies illuminated?** Figure 85 shows the location of the ac LED **1** and the dc LED **2** on the rear of the power supply assembly that is on the rear panel of the 2145-8G4, 2145-8F4, or 2145-8F2 nodes. Figure 86 shows the location of the ac LED **1** and the dc LED **2** on the rear of the power supply assembly that is on the rear panel of the 2145-CF8.




Figure 85. SAN Volume Controller models 2145-8G4 and 2145-8F4 or 2145-8F2 ac and dc LED indicators on the rear panel




Figure 86. Power LED indicator and ac and dc indicators on the rear panel of the SAN Volume Controller 2145-CF8

**NO** Verify that the input power cable or cables are securely connected at both ends and show no sign of damage; otherwise, if the cable or cables are faulty or damaged, replace them. If the node still fails to power on, replace the specified parts based on the SAN Volume Controller model type.

Replace the SAN Volume Controller 2145-CF8 parts in the following sequence:

a. Power supply 675W

Replace the SAN Volume Controller 2145-8G4 parts in the following sequence:

a. Power supply 670W

b. Power backplane

Replace the SAN Volume Controller 2145-8F4 or SAN Volume Controller 2145-8F2 parts in the following sequence:

a. Power supply, 585W

b. Power backplane

Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

**YES** Go to step 10.

9. (from step 8 on page 369)

**Is the power supply error LED on the rear of the SAN Volume Controller 2145-CF8 power supply assemblies illuminated?** Figure 84 on page 368 shows the location of the power LED **1** on the 2145-CF8.

**YES** Replace the power supply unit.

**NO** Go to step 10

10. (from step 1 on page 366, step 8 on page 369, or step 9)

**Are the dc LED indicators on the rear of the power supply assemblies illuminated?**

**NO** Replace the SAN Volume Controller 2145-CF8 parts in the following sequence:

a. Power supply 675W

b. System board

Replace the SAN Volume Controller 2145-8G4 parts in the following sequence:

a. Power backplane

b. Power supply 670W

c. System board

Replace the SAN Volume Controller 2145-8F4 or SAN Volume Controller 2145-8F2 parts in the following sequence:

a. Power backplane

b. Power supply, 585W

c. Frame assembly

Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

**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, front panel

c. System board (if the node is a SAN Volume Controller 2145-CF8 or a SAN Volume Controller 2145-8G4)

- d. Frame assembly (if the node is a SAN Volume Controller 2145-8F4 or SAN Volume Controller 2145-8F2)

Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

11. (from step 1 on page 366)

The node will not power off immediately when the power button is pressed. When the node is fully booted, power-off is performed under the control of the SAN Volume Controller software. The power-off operation can take up to five minutes to complete.

**Is Powering Off displayed on the front panel?**

**NO** Go to step 12.

**YES** Wait for the node to power off. If the node fails to power off after 5 minutes, go to step 12.

12. (from step 11)

**Attention:** Turning off the node by any means other than using the SAN Volume Controller Console might cause a loss of data in the node cache. If you are performing concurrent maintenance, this node must be deleted from the cluster before you proceed. Ask the customer to delete the node from the cluster 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 should turn off within five seconds.

**Did the node turn off?**

**NO** Turn off the 2145 UPS-1U that is connected to this node.

**Attention:** Be sure that you are turning off the correct 2145 UPS-1U. If necessary, trace the cables back to the 2145 UPS-1U assembly. Turning off the wrong 2145 UPS-1U might cause customer data loss.

Go to step 13.

**YES** Go to step 13.

13. (from step 12)

If necessary, turn on the 2145 UPS-1U that is connected to this node and then press the power button to turn the node on.

**Did the node turn on and boot correctly?**

**NO** Go to “MAP 5000: Start” on page 357 to resolve the problem.

**YES** Go to step 14.

14. (from step 13)

The node has probably suffered a software failure. Dump data might have been captured that will help resolve the problem. Call your support center for assistance.

### Related tasks

Chapter 7, “Using the maintenance analysis procedures,” on page 357

The maintenance analysis procedures (MAPs) inform you how to analyze a failure that occurs with a SAN Volume Controller node.

“MAP 5150: 2145 UPS-1U” on page 376

MAP 5150: 2145 UPS-1U helps you solve problems that have occurred in the 2145 UPS-1U systems that are used on a SAN Volume Controller.

“MAP 5250: 2145 UPS-1U repair verification” on page 381

MAP 5250: 2145 UPS-1U repair verification helps you to verify that field replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that were done, have solved all the problems on the SAN Volume Controller 2145 UPS-1U.

### Related reference

“SAN Volume Controller front panel controls and indicators” on page 9

The controls and indicators are used for power and navigation and to indicate information such as system activity, service and configuration options, service controller failures, and node identification.

“Fields for the node VPD” on page 130

The node vital product data (VPD) provides information for items such as the system board, processor, fans, memory module, adapter card, devices, software, front panel assembly, the uninterruptible power supply, SAS host bus adapter (HBA), and SAS solid-state drive (SSD).

---

## MAP 5060: Power 2145-8A4

MAP 5060: Power 2145-8A4 helps you to solve power problems that have occurred on the SAN Volume Controller 2145-8A4 node. If you are using any other SAN Volume Controller model, see the Power MAP for that SAN Volume Controller model.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 7, “Using the maintenance analysis procedures,” on page 357.

You might have been sent here for one of the following reasons:

- A problem occurred during the installation of a 2145-8A4 node.
- The power switch failed to turn the node on.
- The power switch failed to turn the node off.
- Another MAP sent you here.

Perform the following steps:

1. **Are you here because the node is not turned on?**

**NO** Go to step 9 on page 375.

**YES** Go to step 2.

2. (from step 1)

**Is the power LED on the operator-information panel continuously illuminated?** Figure 87 on page 373 shows the location of the power LED **1** on the operator-information panel.




Figure 87. Power LED on the SAN Volume Controller 2145-8A4 operator-information panel

- NO** Go to step 3.
- YES** The node turned on correctly. Reassess the symptoms and return to “MAP 5000: Start” on page 357 or go to “MAP 5700: Repair verification” on page 404 to verify the correct operation.
3. (from step 2 on page 372)
- Is the power LED on the operator-information panel flashing?**
- NO** Go to step 5.
- YES** The node is in standby mode. Input power is present. Go to step 4.
4. (from step 3)
- Press the power-on button 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 turn on, replace parts in the following sequence:
- Operator-information panel
  - Operator-information panel cable
  - System board
- Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.
- YES** The power-on indicator on the operator-information panel shows that the node has successfully turned on. Continue with “MAP 5700: Repair verification” on page 404 to verify the correct operation.
5. (from step 3)
- Locate the 2145 UPS-1U that is connected to this node.
- Does the 2145 UPS-1U that is powering this node have its power on and is its load segment 2 indicator a solid green?**
- NO** Go to “MAP 5150: 2145 UPS-1U” on page 376.
- YES** Verify that the input-power cable is securely connected at both ends and shows no sign of damage; otherwise, if the cable is faulty or damaged, replace it. If the node still fails to turn on, go to step 6. If the node turns on, continue with “MAP 5700: Repair verification” on page 404.
6. (from step 5)
- Remove the node from the rack and remove the top cover. Reconnect the power cable, which is still connected to the 2145 UPS-1U, to the node. **Is the standby power LED that is on the system board illuminated?** Figure 88 on page 374 shows where the diagnostics LEDs are located on the system board.

**NO** Go to step 7 on page 375.

**YES** Replace the SAN Volume Controller 2145-8A4 parts in the following sequence:

- a. Operator-information panel
- b. Operator-information panel cable




Figure 88. SAN Volume Controller 2145-8A4 system board LEDs

- 1** Fan 1 error LED
- 2** Fan 2 error LED
- 3** Fan 3 error LED
- 4** DIMM 1 error LED
- 5** DIMM 2 error LED
- 6** DIMM 3 error LED
- 7** DIMM 4 error LED
- 8** PCI Express slot 2 error LED
- 9** PCI Express slot 1 error LED
- 10** Fan 4 error LED
- 11** Fan 5 error LED
- 12** Voltage regulator error LED
- 13** Standby power LED
- 14** Power good LED

**15** Baseboard management controller heartbeat LED

**16** SAS/SATA controller error LED

7. (from step 6 on page 373)

**Is the voltage regulator LED that is on the system board illuminated?**

**NO** Go to step 8.

**YES** Replace the system board.

8. (from step 7)

Replace the SAN Volume Controller 2145-8A4 parts in the following sequence:

- a. Input-power cable (or the 2145 UPS-1U to SAN Volume Controller node power cable)
- b. Power supply

**Are you now able to turn on the node?**

**NO** Contact your IBM service representative for assistance.

**YES** The power-on indicator on the front panel shows that the node has successfully turned on. Continue with “MAP 5700: Repair verification” on page 404 to verify the correct operation.

9. (from step 1 on page 372)

The node does not turn off when the power button is pressed. When the node is fully booted, power-off is performed under the control of the SAN Volume Controller software. The power-off operation can take up to five minutes to complete.

**Is Powering Off displayed on the front panel?**

**NO** Go to step 10.

**YES** Wait for the node to turn off. If the node fails to turn off after 5 minutes, go to step 10.

10. (from step 9)

**Attention:** Turning off the node by any means other than using the SAN Volume Controller Console might cause a loss of data in the node cache. If you are performing concurrent maintenance, this node must be deleted from the cluster before you proceed. Ask the customer to delete the node from the cluster now. If they are unable to delete the node, contact your IBM service representative 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 should turn off within five seconds.

**Did the node turn off?**

**NO** Turn off the 2145 UPS-1U that is connected to this node.

**Attention:** Be sure that you are turning off the correct 2145 UPS-1U. If necessary, trace the cables back to the 2145 UPS-1U assembly. Turning off the wrong 2145 UPS-1U might cause customer data loss.

Go to step 11.

**YES** Go to step 11.

11. (from step 8)

If necessary, turn on the 2145 UPS-1U that is connected to this node and then 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 357 to resolve the problem.

**YES** Go to step 12.

12. (from step 11 on page 375)

The node has probably suffered a software failure. Dump data might have been captured that will help resolve the problem. Contact your IBM service representative for assistance.

---

## MAP 5150: 2145 UPS-1U

MAP 5150: 2145 UPS-1U helps you solve problems that have occurred in the 2145 UPS-1U systems that are used on a SAN Volume Controller.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 7, “Using the maintenance analysis procedures,” on page 357.

You may have been sent here for one of the following reasons:

- The system problem determination procedures sent you here
- A problem occurred during the installation of a SAN Volume Controller
- Another MAP sent you here
- A customer observed a problem that was not detected by the system problem determination procedures

Figure 89 shows an illustration of the front of the panel for the 2145 UPS-1U.




Figure 89. 2145 UPS-1U front-panel assembly

- 1** Load segment 2 indicator
- 2** Load segment 1 indicator
- 3** Alarm
- 4** On-battery indicator
- 5** Overload indicator
- 6** Power-on indicator
- 7** On/off button
- 8** Test and alarm reset button

Table 54 on page 377 identifies which status and error LEDs that display on the 2145 UPS-1U front-panel assembly relate to the specified error conditions. It also lists the uninterruptible power supply alert-buzzer behavior.



Table 54. 2145 UPS-1U error indicators

[1] Load2	[2] Load1	[3] Alarm	[4] Battery	[5] Overload	[6] Power-on	Buzzer	Error condition
Green (see Note 1)					Green	(see Note 3)	No errors; the 2145 UPS-1U was configured by the SAN Volume Controller
Green	Amber (see Note 2)				Green		No errors; the 2145 UPS-1U is not yet configured by the SAN Volume Controller
Green	Either on or off		Amber		Green	Beeps for two seconds and then stops	The ac power is over or under limit. The uninterruptible power supply has switched to battery mode.
		Flashing red	Flashing amber	Flashing red	Flashing green	Three beeps every ten seconds	Battery undervoltage
Green	Either on or off	Flashing red			Flashing green	Solid on	Battery overvoltage
		Flashing red	Flashing amber		Flashing green	Solid on	Output wave is abnormal when the charger is open, on battery mode
		Flashing red	Flashing amber			Solid on	The ac-power output wave is under low limit or above high limit on battery mode
Green	Either on or off		Amber			Beeps for four seconds and then stops	On battery (no ac power)
Green	Either on or off		Flashing amber			Beeps for two seconds and then stops	Low battery (no ac power)
Green	Either on or off			Red	Green	Beeps for one second and then stops	Overload while on line
			Amber	Red		Beeps for one second and then stops	Overload while on battery
Either on or off	Either on or off	Flashing red			Green	Solid on	Fan failure
Either on or off	Either on or off	Flashing red	Amber			Solid on	Battery test fail
		Flashing red		Red		Solid on	Overload timeout

Table 54. 2145 UPS-1U error indicators (continued)

[1] Load2	[2] Load1	[3] Alarm	[4] Battery	[5] Overload	[6] Power-on	Buzzer	Error condition
		Flashing red	Amber		Green	Solid on	Over temperature
		Flashing red	Amber	Red	Green		Output short circuit

**Notes:**

- The green Load2 LED ([1]) indicates that power is being supplied to the right pair of ac-power outlets as seen from the rear of the 2145 UPS-1U.
- The amber Load1 LED ([2]) indicates that power is being supplied to the left pair of ac-power outlets as seen from the rear of the 2145 UPS-1U. These outlets are not used by the SAN Volume Controller.  
This LED might be illuminated during power-on sequences, but it is typically extinguished by the SAN Volume Controller node that is attached to the 2145 UPS-1U.
- A blank cell indicates that the light or buzzer is off.

- 1. Is the power-on indicator for the 2145 UPS-1U that is connected to the failing SAN Volume Controller off?**

**NO** Go to step 3.

**YES** Go to step 2.
- (from step 1)

**Are other 2145 UPS-1U units showing the power-on indicator as off?**

**NO** The 2145 UPS-1U might be in standby mode. This can be because the on/off button on this 2145 UPS-1U was pressed, input power has been missing for more than five minutes, or because the SAN Volume Controller shut it down following a reported loss of input power. Press and hold the on/off button until the 2145 UPS-1U power-on indicator is illuminated (approximately five seconds). On some versions of the 2145 UPS-1U, you need a pointed device, such as a screwdriver, to press the on/off button.

Go to step 3.

**YES** Either main power is missing from the installation or a redundant ac-power switch has failed. If the 2145 UPS-1U units are connected to a redundant ac-power switch, go to “MAP 5320: Redundant ac power” on page 382. Otherwise, complete these steps:

  - Restore main power to installation.
  - Verify the repair by continuing with “MAP 5250: 2145 UPS-1U repair verification” on page 381.
- (from step 1 and step 2)

**Are the power-on and load segment 2 indicators for the 2145 UPS-1U illuminated solid green, with service, on-battery, and overload indicators off?**

**NO** Go to step 4.

**YES** The 2145 UPS-1U is no longer showing a fault. Verify the repair by continuing with “MAP 5250: 2145 UPS-1U repair verification” on page 381.
- (from step 3)

**Is the 2145 UPS-1U on-battery indicator illuminated yellow (solid or flashing), with service and overload indicators off?**

**NO** Go to step 5.

**YES** The input power supply to this 2145 UPS-1U is not working or is not correctly connected, or the 2145 UPS-1U is receiving input power that might be unstable or outside the specified voltage or frequency range. (The voltage should be between 200V and 240V and the frequency should be either 50 Hz or 60 Hz.) The SAN Volume Controller automatically adjusts the 2145 UPS-1U voltage range. If the input voltage has recently changed, the alarm condition might be present until the SAN Volume Controller has adjusted the alarm setting. Power on the SAN Volume Controller that is connected to the 2145 UPS-1U. If the SAN Volume Controller starts the on-battery indicator should go off within five minutes. If the SAN Volume Controller powers off again or if the condition persists for at least five minutes, do the following:

- a. Check the input circuit protector on the 2145 UPS-1U rear panel, and press it, if it is open.
- b. If redundant ac power is used for the 2145 UPS-1U, check the voltage and frequency at the redundant ac-power switch output receptacle connected to this 2145 UPS-1U. If there is no power, go to “MAP 5340: Redundant ac power verification” on page 383. If the power is not within specification, ask the customer to resolve the issue. If redundant ac power is not used for this uninterruptible power supply, check the site power outlet for the 2145 UPS-1U providing power to this SAN Volume Controller. Check the connection, voltage, and frequency. If the power is not within specification, ask the customer to resolve the issue.
- c. If the input power is within specification and the input circuit protector is stable, replace the field-replaceable units (FRUs) in the following sequence:
  - 1) 2145 UPS-1U power cord
  - 2) 2145 UPS-1U
- d. Verify the repair by continuing with “MAP 5250: 2145 UPS-1U repair verification” on page 381.

5. (from step 4 on page 378)

**Is the 2145 UPS-1U overload indicator illuminated solid red?**

**NO** Go to step 6 on page 380.

**YES** The 2145 UPS-1U output power requirement has exceeded the 2145 UPS-1U capacity.

- a. Check that only one SAN Volume Controller node is connected to the 2145 UPS-1U.
- b. Check that no other loads are connected to the 2145 UPS-1U.
- c. After ensuring that the output loading is correct, turn off the 2145 UPS-1U by pressing the on/off button until the power-on indicator goes off. Then unplug the input power from the 2145 UPS-1U. Wait at least five seconds until all LEDs are off and restart the 2145 UPS-1U by reconnecting it to input power and pressing the on/off button until the 2145 UPS-1U power-on indicator is illuminated (approximately five seconds). On some versions of the 2145 UPS-1U, you need a pointed device, such as a screwdriver, to press the on/off button.
- d. If the condition persists, replace the 2145 UPS-1U.

**Note:** If the condition recurs, replace the power supply or power supplies in the node.

e. Verify the repair by continuing with “MAP 5250: 2145 UPS-1U repair verification” on page 381.

6. (from step 5 on page 379)

**Is the 2145 UPS-1U service indicator illuminated flashing red and the on-battery indicator illuminated solid yellow, with the power-on and overload indicators off?**

**NO** Go to step 7.

**YES** The 2145 UPS-1U battery might be fully discharged or faulty.

- a. Check that the 2145 UPS-1U has been connected to a power outlet for at least two hours to charge the battery. After charging the battery, press and hold the test/alarm reset button for three seconds; and then check the service indicator.
- b. If the service indicator is still flashing, replace the 2145 UPS-1U.
- c. Verify the repair by continuing with “MAP 5250: 2145 UPS-1U repair verification” on page 381.

7. (from step 6)

**Is the 2145 UPS-1U service indicator illuminated flashing red, the on-battery indicator illuminated solid yellow, and the power-on illuminated solid green, with the overload indicator off?**

**NO** Go to step 8.

**YES** The 2145 UPS-1U internal temperature is too high.

- a. Turn off turn off the 2145 UPS-1U by pressing the on/off button until the power-on indicator goes off. Then unplug the 2145 UPS-1U. Clear vents at the front and rear of the 2145 UPS-1U. Remove any heat sources. Ensure the airflow around the 2145 UPS-1U is not restricted.
- b. Wait at least five minutes and restart the 2145 UPS-1U by reconnecting to input power and pressing the on/off button until the 2145 UPS-1U power-on indicator is illuminated (approximately five seconds).
- c. If the condition persists, replace the 2145 UPS-1U.
- d. Verify the repair by continuing with “MAP 5250: 2145 UPS-1U repair verification” on page 381.

8. (from step 7)

**Is the 2145 UPS-1U, service, on-battery, overload, and power-on indicators illuminated and flashing?**

**NO** The 2145 UPS-1U has an internal fault.

- a. Replace the 2145 UPS-1U.
- b. Verify the repair by continuing with “MAP 5250: 2145 UPS-1U repair verification” on page 381.

**YES** The 2145 UPS-1U battery might be fully discharged or faulty.

- a. Check that the 2145 UPS-1U has been connected to a power outlet for at least two hours to charge the battery. After charging the battery, press and hold the test/alarm reset button for three seconds and then check the service indicator.
- b. If the service indicator is still flashing, replace the 2145 UPS-1U.

- c. Verify the repair by continuing with “MAP 5250: 2145 UPS-1U repair verification.”

### Related tasks

Chapter 7, “Using the maintenance analysis procedures,” on page 357

The maintenance analysis procedures (MAPs) inform you how to analyze a failure that occurs with a SAN Volume Controller node.

“MAP 5250: 2145 UPS-1U repair verification”

MAP 5250: 2145 UPS-1U repair verification helps you to verify that field replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that were done, have solved all the problems on the SAN Volume Controller 2145 UPS-1U.

---

## MAP 5250: 2145 UPS-1U repair verification

MAP 5250: 2145 UPS-1U repair verification helps you to verify that field replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that were done, have solved all the problems on the SAN Volume Controller 2145 UPS-1U.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 7, “Using the maintenance analysis procedures,” on page 357.

You may have been sent here because you have performed a repair and want to confirm that no other problems exist on the machine.

Perform the following steps:

1. **Are the power-on and load segment 2 indicators for the repaired 2145 UPS-1U illuminated solid green, with service, on-battery, and overload indicators off?**  
**NO** Continue with “MAP 5000: Start” on page 357.  
**YES** Go to step 2.
2. (from step 1)  
**Is the SAN Volume Controller node powered by this 2145 UPS-1U powered on?**  
**NO** Press power-on on the SAN Volume Controller node that is connected to this 2145 UPS-1U and is powered off. Go to step 3.  
**YES** Go to step 3.
3. (from step 2)  
**Is the node that is connected to this 2145 UPS-1U still not powered on or showing error codes in the front panel display?**  
**NO** Go to step 4.  
**YES** Continue with “MAP 5000: Start” on page 357.
4. (from step 3)  
**Does the SAN Volume Controller node that is connected to this 2145 UPS-1U show “Charging” on the front panel display?**  
**NO** Go to step 5.  
**YES** Wait for the “Charging” display to finish (this might take up to two hours). Go to step 5.
5. (from step 4)

Press and hold the test/alarm reset button on the repaired 2145 UPS-1U for three seconds to initiate a self-test. During the test, individual indicators illuminate as various parts of the 2145 UPS-1U are checked.

**Does the 2145 UPS-1U service, on-battery, or overload indicator stay on?**

**NO** 2145 UPS-1U repair verification has completed successfully. Continue with “MAP 5700: Repair verification” on page 404.

**YES** Continue with “MAP 5000: Start” on page 357.

**Related tasks**

Chapter 7, “Using the maintenance analysis procedures,” on page 357

The maintenance analysis procedures (MAPs) inform you how to analyze a failure that occurs with a SAN Volume Controller node.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

---

## MAP 5320: Redundant ac power

MAP 5320: Redundant ac power helps you solve problems that have occurred in the redundant ac-power switches used on a SAN Volume Controller. Use this MAP when a 2145 UPS-1U that is connected to a redundant ac-power switch does not appear to have input power.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 7, “Using the maintenance analysis procedures,” on page 357.

You might have been sent here for one of the following reasons:

- A problem occurred during the installation of a SAN Volume Controller.
- “MAP 5150: 2145 UPS-1U” on page 376 sent you here.

Perform the following steps to solve problems that have occurred in the redundant ac-power switches:

1. One or two 2145 UPS-1Us might be connected to the redundant ac-power switch. **Is the power-on indicator on any of the connected 2145 UPS-1Us on?**

**NO** Go to step 3 on page 383.

**YES** The redundant ac-power switch is powered. Go to step 2.

2. (from step 1)

Measure the voltage at the redundant ac-power switch output socket connected to the 2145 UPS-1U that is not showing power-on.

**CAUTION:**

**Ensure that you do not remove the power cable of any powered uninterruptible power supply units**

**Is there power at the output socket?**

**NO** One redundant ac-power switch output is working while the other is not. Replace the redundant ac-power switch.

**CAUTION:**

You might need to power-off an operational node to replace the redundant ac-power switch assembly. If this is the case, consult with the customer to determine a suitable time to perform the replacement. See “MAP 5350: Powering off a SAN Volume Controller node” on page 385. After you replace the redundant ac-power switch, continue with “MAP 5340: Redundant ac power verification.”

**YES** The redundant ac-power switch is working. There is a problem with the 2145 UPS-1U power cord or the 2145 UPS-1U . Return to the procedure that called this MAP and continue from where you were within that procedure. It will help you analyze the problem with the 2145 UPS-1U power cord or the 2145 UPS-1U.

3. (from step 1 on page 382)

None of the used redundant ac-power switch outputs appears to have power.

**Are the two input power cables for the redundant ac-power switches correctly connected to the redundant ac-power switch and to different mains circuits?**

**NO** Correctly connect the cables. Go to “MAP 5340: Redundant ac power verification.”

**YES** Verify that there is main power at both the site’s power distribution units that are providing power to this redundant ac-power switch. Go to step 4.

4. (from step 3)

Is power available at one or more of the site’s power distribution units that are providing power to this redundant ac-power switch?

**NO** Have the customer fix the mains circuits. Return to the procedure that called this MAP and continue from where you were within that procedure.

**YES** The redundant ac-power switch should operate in this situation. Replace the redundant ac-power switch assembly. After you replace the redundant ac-power switch, continue with “MAP 5340: Redundant ac power verification.”

---

## MAP 5340: Redundant ac power verification

MAP 5340: Redundant ac power verification helps you verify that a redundant ac-power switch is functioning correctly.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 7, “Using the maintenance analysis procedures,” on page 357.

You might have been sent here because you have replaced a redundant ac-power switch or corrected the cabling of a redundant ac-power switch. You can also use this MAP if you think a redundant ac-power switch might not be working correctly, because it is connected to nodes that have lost power when only one ac power circuit lost power.

In this MAP, you will be asked to confirm that power is available at the redundant ac-power switch output sockets 1 and 2. If the redundant ac-power switch is connected to nodes that are not powered on, use a voltage meter to confirm that power is available.

If the redundant ac-power switch is powering nodes that are powered on (so the nodes are operational), take some precautions before continuing with these tests. Although you do not have to power off the nodes to conduct the test, the nodes will power off if the redundant ac-power switch is not functioning correctly.

For each of the powered-on nodes connected to this redundant ac-power switch, perform the following steps:

1. Use the graphical user interface (GUI) or the command-line interface (CLI) to confirm that the other node in the same I/O group as this node is online.
2. Use the GUI or the CLI to confirm that all virtual disks connected to this I/O group are online.
3. Check the redundant ac-power switch output cables to confirm that the redundant ac-power switch is not connected to two nodes in the same I/O group.

If any of these tests fail, correct any failures before continuing with this MAP. If you are performing the verification using powered-on nodes, understand that power is no longer available if the following is true:

- The on-battery indicator on the 2145 UPS-1U that connects the redundant ac-power switch to the node lights for more than five seconds.
- The SAN Volume Controller node display shows Power Failure.

When the instructions say “remove power,” you can switch the power off if the sitepower distribution unit has outputs that are individually switched; otherwise, remove the specified redundant ac-power switch power cable from the site power distribution unit’s outlet.

Perform the following steps:

1. Are the two site power distribution units providing power to this redundant ac-power switch connected to different power circuits?  
**NO** Correct the problem and then return to this MAP.  
**YES** Go to step 2.
2. (from step 1)  
Are both of the site power distribution units providing power to this redundant ac-power switch powered?  
**NO** Correct the problem and then return to the start of this MAP.  
**YES** Go to step 3.
3. (from step 2)  
Are the two cables that are connecting the site power distribution units to the redundant ac-power switch connected?  
**NO** Correct the problem and then return to the start of this MAP.  
**YES** Go to step 4.
4. (from step 3)  
Is there power at the redundant ac-power switch output socket 2?  
**NO** Go to step 8 on page 385.  
**YES** Go to step 5.
5. (from step 4)  
Is there power at the redundant ac-power switch output socket 1?



- NO** Go to step 8.
- YES** Go to step 6.
6. (from step 5 on page 384)  
Remove power from the Main power cable to the redundant ac-power switch. Is there power at the redundant ac-power switch output socket 1?
- NO** Go to step 8.
- YES** Go to step 7.
7. (from step 6)  
Reconnect the Main power cable. Remove power from the Backup power cable to the redundant ac-power switch. Is there power at the redundant ac-power switch output socket 1?
- NO** Go to step 8.
- YES** Reconnect the Backup power cable. The redundant ac power verification has been successfully completed. Continue with “MAP 5700: Repair verification” on page 404.
8. (from steps 4 on page 384, 5 on page 384, 6, and 7)  
The redundant ac-power switch has not functioned as expected. Replace the redundant ac-power switch assembly. Return to the start of this MAP.

---

## MAP 5350: Powering off a SAN Volume Controller node

MAP 5350: Powering off a SAN Volume Controller node helps you power off a single SAN Volume Controller node to perform a service action without disrupting the host's access to disks.

Powering off a single node will not normally disrupt the operation of a SAN Volume Controller cluster. This is because, within a SAN Volume Controller cluster, nodes operate in pairs called an I/O group. An I/O group will continue to handle I/O to the disks it manages with only a single node powered on. There will, however, be degraded performance and reduced resilience to error.

Care must be taken when powering off a node to ensure the cluster is not impacted more than it need be. If the procedures outlined here are not followed, it is possible your application hosts will lose access to their data or, in the worst case, data will be lost.

You can use the following preferred methods to power off a node that is a member of a cluster and not offline:

1. Use the **Shut Down a Node** option on the SAN Volume Controller Console
2. Use the CLI command **svctask stopcluster -node *name***

It is preferable to use either the SAN Volume Controller Console or the command-line interface (CLI) to power off a node, as these methods provide a controlled handover to the partner node and provide better resilience to other faults in the system.

If a node is offline or not a member of a cluster, it must be powered off using the power button.

To provide the least disruption when powering off a node, the following should all apply:

- The other node in the I/O group should be powered on and active in the cluster.
- The other node in the I/O group should have SAN fibre channel connections to all the hosts and disk controllers managed by the I/O group.
- All the virtual disks handled by this I/O group should be online.
- The host multipathing is online to the other node in the I/O group.

In some circumstances, the reason you are powering off the node might make meeting these conditions impossible; for instance, if you are replacing a broken fibre channel card, the virtual disks will not be showing an online status. You should use your judgment to decide when it is safe to proceed when a condition has not been met. Always check with the system administrator before proceeding with a power off that you know will disrupt I/O access, as they might prefer to either wait until a more suitable time or suspend the host applications

To ensure a smooth restart, a node must save the data structures it cannot recreate to its local, internal, disk drive. The amount of data it 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, you should NOT power off a node using these methods:

- Removing the power cable between the node and the uninterruptible power supply. Normally the uninterruptible power supply provides sufficient power to allow the write to local disk in the event of a power failure, but obviously it is unable to provide power in this case.
- Holding down the node's power button. When the power button is pressed and released, the SAN Volume Controller indicates this to the software and the node can write its data to local disk before it powers off. If the power button is held down, the SAN Volume Controller hardware interprets this as an emergency power off and shuts down immediately without giving you the opportunity to save the data to a local disk. The emergency power off occurs approximately four seconds after the power button is pressed and held down.
- Pressing the reset button on the light path diagnostics panel.

The following topics describe the methods for powering off a node:

- "Using the SAN Volume Controller Console to power off a node"
- "Using the SAN Volume Controller CLI to power off a node" on page 388
- "Using the SAN Volume Controller Power control button" on page 389

## Using the SAN Volume Controller Console to power off a node

This topic describes how to power off a node using the SAN Volume Controller Console.

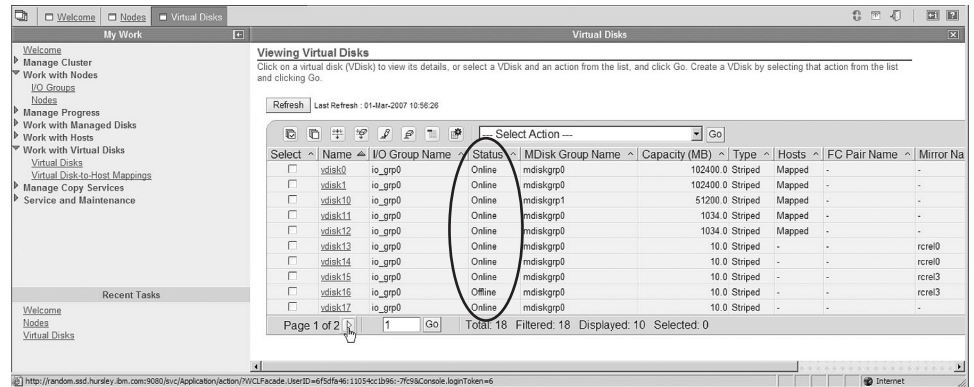
Perform the following steps to use the SAN Volume Controller Console to power off a node:

1. Sign on to the IBM System Storage Productivity Center or master console as an administrator and then launch the SAN Volume Controller Console for the cluster that you are servicing.
2. Click **Work with Nodes** → **Nodes** in the My Work pane. The Viewing Nodes panel is displayed.
3. Find the node that you are about to shut down.

If the node that you want to power off is shown as *Offline*, then the node is not participating in the cluster. In these circumstances, you must use the power button on the node to power off the node.

If the node that you want to power off is shown as *Online*, powering off the node can result in the dependent VDisks to also go offline. Verify whether or not the node has any dependent VDisks.

4. Select the node and click **Show Dependent VDisks** from the drop-down menu.
5. Make sure that the status of each virtual disk in the I/O group is *Online*. You might need to view more than one page.



If any VDisks are shown as degraded, only one node in the I/O is processing I/O requests for that VDisk. If that node is powered off, it impacts all the hosts that are submitting I/O requests to the degraded VDisks.

If any virtual disks are degraded and you believe this might be because the partner node in the I/O group has been powered off recently, wait until a refresh of the screen shows all the virtual disks online. All the virtual disks should be online within thirty minutes of the partner node being powered off.

**Note:** If, after waiting 30 minutes, you have a degraded VDisk and all of the associated nodes and MDisks are online, contact the IBM Support Center for assistance.

Ensure that all VDisks that are being used by hosts are online before you continue.

6. If possible, check that all the hosts that access VDisks 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.

Perform this check using the host system's multipathing device driver software. The commands to use differ, depending on the multipathing device driver being used. If you are using 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 being powered on or you might lose access to VDisks.

7. If you have decided it is okay to continue and power off the node, select the node that you want to power off, and then select **Shut Down a Node** from the drop-down menu.

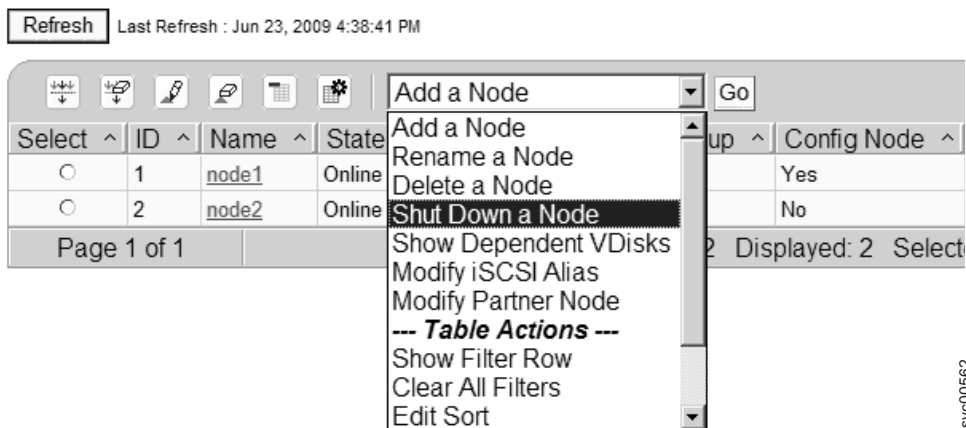


Figure 90. Shut Down a Node option

8. Click **OK**. If you have selected a node that is the last remaining node that provides access to a VDisk, for example, a node that contains solid-state drives (SSDs) with unmirrored VDisks, the Shutting Down a Node-Force panel is displayed with a list of VDisks that will go offline if this node is shut down.
9. Check that no host applications are accessing the VDisks that will go offline; only continue with the shut down if the loss of access to these VDisks is acceptable. To continue with shutting down the node, click **Force Shutdown**.

During the shut down, the node saves its data structures to its local disk and destages all the write data held in cache to the SAN disks; this processing can take several minutes.

At the end of this process, the node powers off.

## Using the SAN Volume Controller CLI to power off a node

This topic describes how to power off a node using the SAN Volume Controller CLI.

1. Issue the **svcinfolnode** CLI command to display a list of nodes in the cluster and their properties. Find the node that you are about to shut down and write down the name of the I/O group it belongs to. Confirm that the other node in the I/O group is online.

```
svcinfolnode -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:202378101C0D18D8
2:group1node2:10L3ANF:5005076801000009:online:0:io_grp0:no:202378101C0D1796
3:group2node1:10L3ASH:5005076801000001:online:1:io_grp1:no:202378101C0D18D8
4:group2node2:10L3ANF:50050768010000F4:online:1:io_grp1:no:202378101C0D1796
```

If the node that you want to power off is shown as `Offline`, the node is not participating in the cluster and is not processing I/O requests. In these circumstances, you must use the power button on the node to power off the node.

If the node that you want 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 the hosts that are submitting I/O requests to the VDisks 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 **svcinfolnodedependentvdisks** CLI command to list the VDisks that are dependent on the status of a specified node.

```
svcinfolsnodedependentvdisks group1node1
```

```
vdisk_id      vdisk_name
0             vdisk0
1             vdisk1
```

If the node goes offline or is removed from the cluster, the dependent VDisks also go offline. Before taking a node offline or removing it from the cluster, you can use the command to ensure that you do not lose access to any VDisks.

3. If you have decided that it is okay to continue and that you can power off the node, issue the **svctask stopcluster -node <name>** CLI command to power off the node. Ensure that you use the **-node** parameter, because you do not want to power off the whole cluster:

```
svctask stopcluster -node group1node1
Are you sure that you want to continue with the shut down? yes
```

**Note:** If there are dependent VDisks and you want to shut down the node anyway, add the **-force** parameter to the **svctask stopcluster** command. The **force** parameter forces continuation of the command even though any node-dependent VDisks will be taken offline. Use the **force** parameter with caution; access to data on node-dependent VDisks will be lost.

During the shut down, the node saves its data structures to its local disk and destages all the write data held in the cache to the SAN disks; this process can take several minutes.

At the end of this process, the node powers off.

## Using the SAN Volume Controller Power control button

Do not use the power control button to power off a node unless it is an emergency or you have been directed to do so by another procedure.

With this method, you cannot check the cluster status from the front panel, so you cannot tell if the power off is liable to cause excessive disruption to the cluster. Instead, use the SAN Volume Controller Console or the CLI commands, described in the previous topics, to power off an active node.

If you must use this method, notice in Figure 91 that each SAN Volume Controller model type has a power control button **1** on the front.




Figure 91. SAN Volume Controller models 2145-CF8, 2145-8A4, 2145-8G4, and 2145-8F4 or 2145-8F2 power control button

When you have determined it is safe to do so, press and immediately release the power button. The front panel display changes to display *Powering Off*, and a progress bar is displayed.

The 2145-CF8 requires that you remove a power button cover before you can press the power button. The 2145-8A4, the 2145-8G4, the 2145-8F4, or 2145-8F2 might require you to use a pointed device to press the power button.

If you press the power button for too long, the node cannot write all the data to its local disk. An extended service procedure is required to restart the node, which involves deleting the node from the cluster and adding it back into the cluster.



Powering Off

The node saves its data structures to disk while powering off. The power off process can take up to five 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 time taken by this destage depends on the speed and utilization of the disk controllers; it should complete in less than 15 minutes, but it could be longer, and it cannot complete if there is data waiting to be written to a disk that is offline.

If a node powers off and restarts while its partner node continues to process I/O, it might not be able to become an active member of the I/O group immediately. It has to wait until the partner node completes its destage of the cache. If the partner node is powered 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, for example, because the partner node in the I/O group is still flushing its write cache, the VDisks that are managed by that I/O group will have a status of Degraded.

---

## MAP 5400: Front panel

MAP 5400: Front panel helps you to solve problems that have occurred on the SAN Volume Controller front panel.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 7, “Using the maintenance analysis procedures,” on page 357.

This MAP applies to all SAN Volume Controller 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 have been sent here because:

- A problem occurred during the installation of a SAN Volume Controller system, the front panel display test failed, or the correct node number failed to be displayed
- Another MAP sent you here

Perform the following steps:

1. **Is the power LED on the operator-information panel illuminated and showing a solid green?**

**NO** Continue with the power MAP. See “MAP 5050: Power 2145-CF8, 2145-8G4, 2145-8F4, and 2145-8F2” on page 365 or “MAP 5060: Power 2145-8A4” on page 372.

**YES** Go to step 2.

2. (from step 1)

**Is the service controller error light that you see in Figure 92 illuminated and showing a solid amber?**




Figure 92. SAN Volume Controller service controller error light

**1** SAN Volume Controller models 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, and 2145-8F2 service controller error light

**NO** Start the front panel tests by pressing and holding the select button for five seconds. Go to step 3.

**Attention:** Do not start this test until the node is powered on for at least two minutes. You might receive unexpected results.

**YES** The SAN Volume Controller service controller has failed. Replace the parts in the following sequence:

SAN Volume Controller models 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, and 2145-8F2	Service controller
---	--------------------

Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

3. (from step 2)

The front panel check light illuminates and the display test of all display bits turns on for 3 seconds and then turns off for 3 seconds, then a vertical line travels from left to right, followed by a horizontal line travelling from top to bottom. The test completes with the switch test display of a single rectangle in the center of the display.

**Did the front panel lights and display operate as described?**

**NO** SAN Volume Controller front panel has failed its display test.

- Replace the parts in the following sequence:

SAN Volume Controller models 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, and 2145-8F2	Service controller

- Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

**YES** Go to step 4.

4. (from step 3 on page 391)

Figure 93 provides four examples of what the front panel display shows before you press any button and then when you press the up button, the left and right buttons, and the select button. To perform the front panel switch test, press any button in any sequence or any combination. The display indicates which buttons you pressed.




Figure 93. Front-panel display when push buttons are pressed

Check each switch in turn. Did the service panel switches and display operate as described in Figure 93?

**NO** The SAN Volume Controller front panel has failed its switch test.

- Replace the parts in the following sequence:

SAN Volume Controller models 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, and 2145-8F2	Service controller
---	--------------------

- Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

**YES** Press and hold the select button for five seconds to exit the test. Go to step 5.

5. **Is front panel display now showing Cluster:?**

**NO** Continue with “MAP 5000: Start” on page 357.

**YES** Keep pressing and releasing the down button until Node is displayed in line 1 of the menu screen. Go to step 6.

6. (from step 5)

**Is this MAP being used as part of the installation of a new node?**

**NO** Front panel tests have completed with no fault found. Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

**YES** Go to step 7.

7. (from step 6)

**Is the node number that is displayed in line 2 of the menu screen the same as the node number that is printed on the front panel of the node?**



**NO** Node number stored in front panel electronics is not the same as that printed on the front panel.

SAN Volume Controller models 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, and 2145-8F2	Service controller
---	--------------------

**YES** Front panel tests have completed with no fault found. Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

#### **Related concepts**

“SAN Volume Controller menu options” on page 144

During normal operations, menu options are available on the front panel display of the SAN Volume Controller node.

#### **Related tasks**

Chapter 7, “Using the maintenance analysis procedures,” on page 357

The maintenance analysis procedures (MAPs) inform you how to analyze a failure that occurs with a SAN Volume Controller node.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

---

## **MAP 5500: Ethernet**

MAP 5500: Ethernet helps you solve problems that have occurred on the SAN Volume Controller Ethernet.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 7, “Using the maintenance analysis procedures,” on page 357.

This MAP applies to all SAN Volume Controller 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 have been sent here for one of the following reasons:

- A problem occurred during the installation of a SAN Volume Controller system and the Ethernet checks failed
- Another MAP sent you here
- The customer needs immediate access to the cluster by using an alternate configuration node. See “Defining an alternate configuration node” on page 396

Perform the following steps:

1. **Is the front panel of any node in the cluster displaying Node Error with error code 540?**

**YES** Go to step 6 on page 394.

**NO** Go to step 2.

2. **Is the cluster reporting cluster error 1400 either on the front panel or in the error 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 395.
- NO** Go to step 10 on page 395.
4. (from step 2 on page 393) **On all nodes perform the following actions:**
- Press the down button until the top line of the display shows Ethernet. If the second line of the display shows Active or Inactive, refer to the *IBM System Storage SAN Volume Controller Troubleshooting Guide*, Release 4.3.1.
  - Press right until the top line displays Ethernet port 1.
  - If the second line of the display shows link offline, record this port as one that requires fixing.
  - If the cluster is configured with two Ethernet cables per node, press the right button until the top line of the display shows Ethernet port 2 and repeat the previous step.
  - 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 395.
6. (from step 5) **Do the SAN Volume Controller nodes have one or two cables connected?**
- One** Go to step 7.
- Two** Go to step 8 on page 395.
7. (from step 6) **Perform the following actions:**
- Plug the Ethernet cable from that node into the Ethernet port 2 from a different node.
  - 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 94. Port 2 Ethernet link LED on the SAN Volume Controller rear panel

- 1** SAN Volume Controller 2145-CF8 port 2 (upper right) Ethernet link LED
- 2** SAN Volume Controller 2145-8F2 or SAN Volume Controller 2145-8F4 port 2 (lower right) Ethernet link LED
- 3** SAN Volume Controller 2145-8G4 port 2 (center) Ethernet link LED

**4** SAN Volume Controller 2145-8A4 port 2 (upper right) Ethernet link LED

- 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 404.
8. (from step 5 on page 394 or step 6 on page 394) **Perform 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 404.
9. (from step 3 on page 394) **Perform 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) Press the down button until the top line of the display shows Ethernet. If the second line of the display shows Active or Inactive, refer to the *IBM System Storage SAN Volume Controller Troubleshooting Guide*, Release 4.3.1.
    - 2) Press the right button until the top line displays Speed 1.
    - 3) If the second line of the display shows link offline, record this port as one that requires fixing.
    - 4) If the cluster is configured with two Ethernet cables per node, press the right button until the top line of the display shows Speed 2 and repeat the previous step.
  - b. Check that the SAN Volume Controller port has negotiated at the highest speed available on the switch. All nodes have gigabit Ethernet network ports.
  - c. If the Duplex setting is half, perform 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 393)
- 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. Perform these steps in the following sequence until the problem is resolved:
- 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 404.
  - Determine if similar Ethernet connection problems have occurred recently on this node. If they have, replace the system board.
  - Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.
- YES** Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

#### Related tasks

Chapter 7, “Using the maintenance analysis procedures,” on page 357

The maintenance analysis procedures (MAPs) inform you how to analyze a failure that occurs with a SAN Volume Controller node.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

## Defining an alternate configuration node

A situation can arise where the customer needs immediate access to the cluster by using an alternate configuration node.

If all Ethernet connections to the configuration node have failed, the cluster is unable to report failure conditions, and the SAN Volume Controller Console is unable to access the cluster to perform administrative or service tasks. If this is the case and the customer needs immediate access to the cluster, you can make the cluster use an alternate configuration node.

If only one node is displaying Node Error 540 on the front panel, perform the following steps:

- Press and release the power button on the node that is displaying Node Error 540.
- When Powering off is displayed on the front panel display, press the power button again.
- Restarting is displayed.

The cluster will select a new configuration node. The SAN Volume Controller Console is able to access the cluster again.

---

## MAP 5600: Fibre-channel

MAP 5600: Fibre-channel helps you to solve problems that have occurred on the SAN Volume Controller fibre-channel ports.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 7, “Using the maintenance analysis procedures,” on page 357.

This MAP applies to all SAN Volume Controller 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 have been sent here for one of the following reasons:

- A problem occurred during the installation of a SAN Volume Controller system and the fibre-channel checks failed
- Another MAP sent you here

Perform the following steps to solve problems caused by the fibre-channel ports:

1. **Are you here to diagnose a problem on a SAN Volume Controller 2145-8F2?**

**NO** Go to step 2.

**YES** Go to step 3.

2. **Are you trying to resolve a fibre-channel port speed problem?**

**NO** Go to step 3.

**YES** Go to step 12 on page 402.

3. (from step 1 and step 2) Display fibre-channel port 1 status on the SAN Volume Controller front-panel display. For more information, see Chapter 5, "Using the front panel of the SAN Volume Controller," on page 139.

**Is the front-panel display on the SAN Volume Controller showing fibre-channel port-1 active?**

**NO** A fibre-channel port is not working correctly. Check the port status on the second line of the display.

- **Inactive:** The port is operational but cannot access the fibre-channel fabric. The fibre-channel adapter is not configured correctly, the fibre-channel SFP has failed, the fibre-channel cable has either failed or is not installed, or the device at the other end of the cable has failed. Make a note of port-1. Go to step 8 on page 399.
- **Failed:** The port is not operational because of a hardware failure. Make a note of port-1. Go to step 10 on page 401.
- **Not installed:** This port is not installed. Make a note of port-1. Go to step 11 on page 402.

**YES** Press and release the right button to display fibre-channel port-2 . Go to step 4.

4. (from step 3)

**Is the front panel display on the SAN Volume Controller showing fibre-channel port-2 active?**

**NO** A fibre-channel port is not working correctly. Check the port status on the second line of the display.

- **Inactive:** The port is operational but cannot access the fibre-channel fabric. The fibre-channel adapter is not configured correctly, the fibre-channel SFP has failed, the fibre-channel cable has either failed or is not installed, or the device at the other end of the cable has failed. Make a note of port-2. Go to step 8 on page 399.

- **Failed:** The port is not operational because of a hardware failure. Make a note of port-2. Go to step 10 on page 401.
- **Not installed:** This port is not installed. Make a note of port-2. Go to step 11 on page 402.

**YES** Press and release the right button to display fibre-channel port-3. Go to step 5.

5. (from step 4 on page 397)

**Is the front panel display on the SAN Volume Controller showing fibre-channel port-3 active?**

**NO** A fibre-channel port is not working correctly. Check the port status on the second line of the display.

- **Inactive:** The port is operational but cannot access the fibre-channel fabric. The fibre-channel adapter is not configured correctly, the fibre-channel SFP has failed, the fibre-channel cable has either failed or is not installed, or the device at the other end of the cable has failed. Make a note of port-3. Go to step 8 on page 399.
- **Failed:** The port is not operational because of a hardware failure. Make a note of port-3. Go to step 10 on page 401.
- **Not installed:** This port is not installed. Make a note of port-3. Go to step 11 on page 402.

**YES** Press and release the right button to display fibre-channel port-4. Go to step 6.

6. (from step 5)

**Is the front panel display on the SAN Volume Controller showing fibre-channel port-4 active?**

**NO** A fibre-channel port is not working correctly. Check the port status on the second line of the display.

- **Inactive:** The port is operational but cannot access the fibre-channel fabric. The fibre-channel adapter is not configured correctly, the fibre-channel SFP has failed, the fibre-channel cable has either failed or is not installed, or the device at the other end of the cable has failed. Make a note of port-4. Go to step 8 on page 399.
- **Failed:** The port is not operational because of a hardware failure. Make a note of port-4. Go to step 9 on page 400.
- **Not installed:** This port is not installed. Make a note of port-4. Go to step 11 on page 402.

**YES** Go to step 7.

7. (from step 6)

A previously reported fault with a fibre-channel port is no longer being shown. A problem with the SAN fibre-channel fabric might have been fixed or there might be an intermittent problem.

Check with the customer to see if any fibre-channel ports have been disconnected or if any component of the SAN fibre-channel fabric has failed and has been fixed recently.

**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 404.
- b. Check if similar fibre-channel errors have occurred recently on the same port on this SAN Volume Controller node. If they have, replace the fibre-channel cable, unless it has already been replaced.
- c. Replace the fibre-channel SFP connector, unless it has already been replaced.

**Note:** SAN Volume Controller nodes are supported with both long-wave SFPs and short-wave SFPs. You must replace an SFP with the same type of SFP that you are replacing. If the SFP to replace is a long wave SFP, for example, you must provide a suitable replacement. Removing the wrong SFP connector could result in loss of data access. See “Removing and replacing the fibre-channel SFP connector on a SAN Volume Controller node” in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.

- d. Replace the fibre-channel adapter assembly shown in the following table.

SAN Volume Controller 2145-CF8 port 1, 2, 3, or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8A4 port 1, 2, 3, or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8G4 port 1, 2, 3, or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8F4 port 1, 2, 3, or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8F2 port 1 or 2	Dual port fibre-channel host bus adapter (HBA) - Low profile
SAN Volume Controller 2145-8F2 port 3 or 4	Dual port fibre-channel HBA - Full height

- e. Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

**YES** Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

8. (from steps 3 on page 397, 4 on page 397, 5 on page 398, and 6 on page 398)

The noted port on the SAN Volume Controller is showing a status of inactive. For certain models, this might occur when the fibre-channel speed is not set correctly.

**Are you diagnosing a problem on a SAN Volume Controller 2145-8F2?**

**NO** Go to step 9 on page 400.

**YES** All SAN Volume Controller ports on the SAN Volume Controller 2145-8F2 nodes must run at the same speed. This speed is set by one of the cluster properties; therefore the cluster speed must be set to a speed that all SAN Volume Controller ports can use.

If the node or nodes are currently online in the cluster, change the cluster speed property to a speed that all SAN Volume Controller 2145-8F2 ports can use.

**Attention:** Changing the SAN Volume Controller speed setting causes an I/O outage on the cluster. Ensure that all host operations are stopped before performing these steps.

- a. Press the down button until the top line of the display shows Ethernet. If the second line of the display shows Active or Inactive, refer to the *IBM System Storage SAN Volume Controller Troubleshooting Guide*, Release 4.3.1.
- b. Press the right button until the top line displays Speed 1.
- c. If the second line of the display shows link offline, record this port as one that requires fixing.
- d. If the cluster is configured with two Ethernet cables per node, press the right button until the top line of the display shows Speed 2 and repeat the previous step.
- e. Go to step 9.

If the node is not currently online in the cluster, you might need to set the speed of the node to a different speed setting before the SAN Volume Controller node can join the cluster. To temporarily set the speed of the node, perform the following steps:

**Note:** After the SAN Volume Controller joins the cluster, the node's fibre-channel port speed will be changed to match the cluster setting. Check the cluster setting before changing the node.

- a. Press and hold the down button.
- b. Press and release the select button.
- c. Release the down button.  
The fibre-channel speed setting is shown on the display. If this value does not match the speed of the SAN, use the down and up buttons to set it correctly.
- d. Press the select button to accept any changes and return to the fibre-channel status display.
- e. If the status shows active, continue with "MAP 5700: Repair verification" on page 404. Otherwise, go to step 9.

9. (from step 8 on page 399)

The noted port on the SAN Volume Controller 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 SAN Volume Controller to fibre-channel network.
- b. Faulty fibre-channel fabric connections, particularly the SFP connector at the fibre-channel switch. Use the SAN problem determination procedure to resolve any fibre-channel fabric connection problem.
- c. SAN Volume Controller fibre-channel SFP connector.

**Note:** SAN Volume Controller nodes are supported with both long-wave SFPs and short-wave SFPs. You must replace an SFP with the same type of SFP that you are replacing. If the SFP to replace is a long wave SFP, for example, you must provide a suitable replacement. Removing the wrong SFP connector could result in



loss of data access. See "Removing and replacing the fibre-channel SFP connector on a SAN Volume Controller node" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.

d. Fibre-channel adapter assemblies.

SAN Volume Controller 2145-CF8 port 1, 2, 3, or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8A4 port 1, 2, 3, or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8G4 port 1, 2, 3, or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8F4 port 1, 2, 3, or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8F2 port 1 or 2	Dual port fibre-channel host bus adapter (HBA) - Low profile
SAN Volume Controller 2145-8F2 port 3 or 4	Dual port fibre-channel HBA - Full height

e. Verify the repair by continuing with "MAP 5700: Repair verification" on page 404.

10. (from steps 3 on page 397, 4 on page 397, 5 on page 398, and 6 on page 398)

The noted port on the SAN Volume Controller displays a status of failed. Verify that the fibre-channel cables that connect the SAN Volume Controller 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 connector.

**Note:** SAN Volume Controller nodes are supported with both long-wave SFPs and short-wave SFPs. You must replace an SFP with the same type of SFP that you are replacing. If the SFP to replace is a long wave SFP, for example, you must provide a suitable replacement. Removing the wrong SFP connector could result in loss of data access. See "Removing and replacing the fibre-channel SFP connector on a SAN Volume Controller node" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.

b. Fibre-channel adapter assemblies.

SAN Volume Controller 2145-CF8 port 1, 2, 3, or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8A4 port 1, 2, 3, or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8G4 port 1, 2, 3, or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8F4 port 1, 2, 3 or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8F2 port 1 or 2	Dual port fibre-channel host bus adapter (HBA) - Low profile
SAN Volume Controller 2145-8F2 port 3 or 4	Dual port fibre-channel HBA - Full height

- c. Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.
11. (from steps 3 on page 397, 4 on page 397, 5 on page 398, and 6 on page 398)

The noted port on the SAN Volume Controller displays a status of not installed. If you have just replaced the fibre-channel adapter, make sure that it is installed correctly. If you have replaced any other system board components, make sure that the fibre-channel adapter has not been disturbed.

**Is the fibre-channel adapter failure explained by the previous checks?**

**NO**

- a. Replace the fibre-channel adapter assembly for the node, as listed in Table 55.

*Table 55. SAN Volume Controller fibre-channel adapter assemblies*

SAN Volume Controller 2145-CF8 port 1, 2, 3, or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8A4 port 1, 2, 3, or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8G4 port 1, 2, 3, or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8F4 port 1, 2, 3, or 4	4-port fibre channel HBA
SAN Volume Controller 2145-8F2 port 3 or 4	Dual port fibre-channel HBA - Full height
SAN Volume Controller 2145-8F2 port 1 or 2	Dual port fibre-channel host bus adapter (HBA) - Low profile

- b. If the problem is not fixed, replace the fibre-channel connection hardware in the order that is shown in Table 56.

*Table 56. SAN Volume Controller fibre-channel adapter connection hardware*

SAN Volume Controller 2145-8A4 port 1, 2, 3, or 4	1. Riser card, PCI Express 2. System board
SAN Volume Controller 2145-8G4 port 1, 2, 3, or 4	1. Riser card, PCI Express 2. System board
SAN Volume Controller 2145-8F4 port 1, 2, 3, or 4	1. Riser card, PCI Express 2. Frame assembly
SAN Volume Controller 2145-8F2 port 1 or 2	1. Riser card, PCI Low profile 2. Frame assembly
SAN Volume Controller 2145-8F2 port 3 or 4	1. Riser card, PCI 2. Frame assembly

- c. Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

**YES** Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

12. (from step 2 on page 397)

For the SAN Volume Controller models 2145-8A4, 2145-8G4, and 2145-8F4, each fibre-channel port autonegotiates its operating speed with the switch to which it is connected. If the speed at which it is operating is lower than the

operating speed that is supported by the switch, this indicates that a high number of link errors are being detected.

To display the current speed of the link, perform the following steps:

- a. Press the up or down button on the front panel until FC Port-1: is displayed on the first line of the service panel.
- b. Press the left or right button until the required port is displayed.
- c. Press and hold the down button.
- d. Press and release the select button.
- e. Release the down button.

The second line of the front-panel display shows the current fibre-channel speed of the port.

**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 404.

**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. Any bend should have 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. This will cause the fibre-channel adapter to renegotiate its operating speed.
- c. Recheck the speed of the fibre-channel port. If it is now correct, you have resolved the problem. Otherwise, the problem might be caused by one of the following:
  - 4-port fibre channel HBA
  - SAN Volume Controller SFP connector
  - Fibre-channel switch GBIC or SFP
  - Fibre-channel switch

Recheck the speed after changing any component until the problem is resolved and then verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

### Related concepts

“SAN Volume Controller menu options” on page 144

During normal operations, menu options are available on the front panel display of the SAN Volume Controller node.

“Fibre-channel port-1 through 4 option” on page 159

The fibre-channel port-1 through 4 options display the operational status of the fibre-channel ports.

### Related tasks

“SAN problem determination” on page 348

The procedures to service the SAN Volume Controller that are provided here help you solve problems on the SAN Volume Controller and its connection to the storage area network (SAN).

Chapter 7, “Using the maintenance analysis procedures,” on page 357

The maintenance analysis procedures (MAPs) inform you how to analyze a failure that occurs with a SAN Volume Controller node.

“MAP 5700: Repair verification”

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

---

## MAP 5700: Repair verification

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 7, “Using the maintenance analysis procedures,” on page 357.

You might have been sent here because you performed a repair and want to confirm that no other problems exists on the machine.

Perform the following steps to verify your repair:

1. (from step )

**Are all the SAN Volume Controller nodes displaying Cluster: on the top line of the front panel display with the second line blank or displaying a cluster name?**

**NO** Go to “MAP 5000: Start” on page 357.

**YES** Go to step 2.

2. (from step 1)

Using the SAN Volume Controller application for the cluster you have just repaired, 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. See “Determining the failing enclosure or disk controller using the CLI” on page 107 to locate the disk controller with the offline MDisk. Use the problem determination procedure for the disk controller to repair the MDisk faults before returning 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 returning to this MAP.

If any MDisks show a status of excluded, include MDisks before returning to this MAP.

Go to “MAP 5000: Start” on page 357.

**YES** Go to step 3.

3. (from step 2 on page 404)

Using the SAN Volume Controller application for the cluster you have just repaired, check the status of all configured virtual disks (VDisks). **Do all VDisks have a status of online?**

**NO** Go to step 4.

**YES** Go to step 5.

4. (from step 3)

Following a repair of the SAN Volume Controller, a number of VDisks are showing a status of offline. VDisks will be held offline if SAN Volume Controller cannot confirm the integrity of the data. The VDisk might be the target of a copy that did not complete, or cache write data that was not written back to disk might have been lost. Determine why the VDisk is offline. If the VDisk was the target of a copy that did not complete, you can start the copy again. Otherwise, write data might not have been written to the disk, so its state cannot be verified. Your site procedures will determine how data is restored to a known state.

To bring the VDisks 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. See the topic about recovering from offline VDisks in the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide* for details on how to resolve this problem.

Go to “MAP 5000: Start” on page 357.

5. (from step 3)

You have successfully repaired the SAN Volume Controller.

### Related tasks

“SAN problem determination” on page 348

The procedures to service the SAN Volume Controller that are provided here help you solve problems on the SAN Volume Controller and its connection to the storage area network (SAN).

Chapter 7, “Using the maintenance analysis procedures,” on page 357

The maintenance analysis procedures (MAPs) inform you how to analyze a failure that occurs with a SAN Volume Controller node.

“MAP 5000: Start” on page 357

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

### Related reference

“Determining the failing enclosure or disk controller using the CLI” on page 107

You can use the command-line interface (CLI) to determine the failing enclosure or disk controller.

Chapter 6, “Diagnosing problems,” on page 167

You can diagnose problems with SAN Volume Controller, the uninterruptible power supply, and the IBM System Storage Productivity Center, or the master console server using either the command-line interface (CLI) or the SAN Volume Controller Console. The diagnostic LEDs on the SAN Volume Controller nodes and uninterruptible power supply units also help you diagnose hardware problems.

### Related information

“Understanding cluster error codes” on page 239

Every cluster error code includes an error code number, a description, action, and possible field-replaceable units (FRUs).

---

## MAP 5800: Light path

MAP 5800: Light path helps you to solve hardware problems on the SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2 that are preventing the node from booting.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 7, “Using the maintenance analysis procedures,” on page 357.

You might have been sent here because of the following:

- The Error LED on the operator-information panel is on or flashing
- Another MAP sent you here

### Related tasks

Chapter 7, “Using the maintenance analysis procedures,” on page 357

The maintenance analysis procedures (MAPs) inform you how to analyze a failure that occurs with a SAN Volume Controller node.

## Light path for SAN Volume Controller 2145-CF8

Use the diagnostics LEDs that are located on the system board to solve hardware problems with the SAN Volume Controller 2145-CF8 node.

Ensure that the node is turned on and then perform the following steps to resolve any hardware errors that are indicated by the Error LED and light path LEDs:

1. **Is the Error LED, shown in Figure 95 on page 407, on the SAN Volume Controller 2145-CF8 operator-information panel on or flashing?**




Figure 95. SAN Volume Controller 2145-CF8 operator-information panel

- 5** System error LED
- 6** Release latch

**NO** Reassess your symptoms and return to “MAP 5000: Start” on page 357.

**YES** Go to step 2.

2. (from step 1 on page 406)

Press the release latch and open the light path diagnostics panel, which is shown in Figure 96.

**Are one or more LEDs on the light path diagnostics panel on or flashing?**

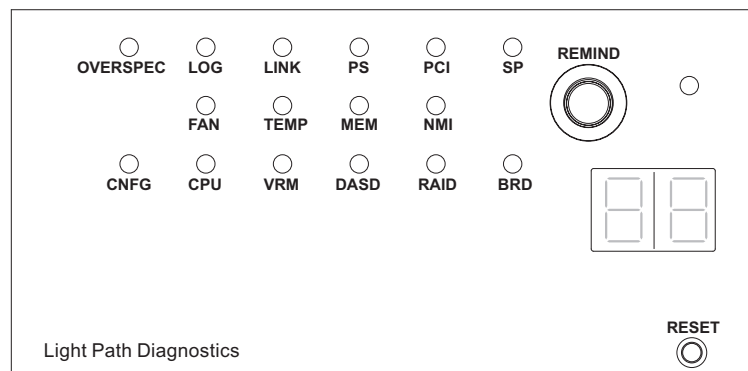


Figure 96. SAN Volume Controller 2145-CF8 light path diagnostics panel

**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 404.

**YES** See Table 57 on page 410 and perform the action specified for the specific light path diagnostics LEDs. Then go to step 3 on page 412. Some actions will require that you observe the state of LEDs on the system board. Figure 97 on page 408 shows the location of the system board LEDs. The fan LEDs are located adjacent to each FAN. To view the LEDs you will need to do the following:

- a. Turn off the node while ensuring that its data is mirrored and synchronized. See “MAP 5350: Powering off a SAN Volume Controller node” on page 385 for more information.

- b. (Optional) Identify and label all the cables that are attached to the node so they can be replaced in the same port. Remove the node from the rack and place it on a flat, static-protective surface. See "Removing the SAN Volume Controller from a rack" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- c. Remove the top cover.
- d. See Table 57 on page 410 and perform the action specified for the specific light path diagnostics LEDs. Then go to step 3 on page 412.




Figure 97. SAN Volume Controller 2145-CF8 system board LEDs diagnostics panel

- 1** Slot 2 missing PCI riser card LED
- 2** Enclosure manager heartbeat LED
- 3** Battery LED
- 4** IMM heartbeat LED
- 5** Slot 1 missing PCI riser card LED
- 6** System error LED
- 7** Microprocessor 1 error LED



- | **8** DIMM 1-8 error LEDs
- | **9** Fan 1 error LED
- | **10** Fan 2 error LED
- | **11** Fan 3 error LED
- | **12** Fan 4 error LED
- | **13** Fan 5 error LED
- | **14** Fan 6 error LED
- | **15** 240 VA error LED
- | **16** Power channel A error LED
- | **17** Power channel B error LED
- | **18** Power channel C error LED
- | **19** Power channel D error LED
- | **20** Power channel E error LED
- | **21** AUX power channel error LED
- | **22** SAS/SATA RAID error LED
- | **23** Microprocessor 2 error LED
- | **24** DIMM 9-16 error LEDs

Table 57. Diagnostics panel LED prescribed actions

Diagnostics panel LED	Action
OVER SPEC	<p>The power supplies are using more power than their maximum rating. If the OVER SPEC LED is lit, one or more of the six 12V channel error LEDs (A, B, C, D, E, or AUX) is also lit on the system board. Perform the following actions to resolve the problem:</p> <ol style="list-style-type: none"> <li>1. Turn off the node, pull the node forward in the rack, and remove the cover. Do not disconnect power from the node.</li> <li>2. Check which 12V channel error LED is lit on the system board, and remove the components listed for that LED: <ul style="list-style-type: none"> <li>• LED A: fans, disk drive, any solid-state drives (SSDs), or disk backplane</li> <li>• LED B: fibre-channel adapter and riser, all memory</li> <li>• LED C: disk controller, all memory</li> <li>• LED D: go to f.</li> <li>• LED E: High-speed SAS adapter and riser, if installed</li> <li>• LED AUX: fibre-channel adapter and high-speed SAS adapter, if installed</li> </ul> </li> <li>3. Restart the node to see whether the problem remains.</li> <li>4. Reinstall each device that you removed in substep b one at a time. Start the node each time to isolate the failing device.</li> <li>5. Replace any failing device.</li> <li>6. If no device was isolated, and if LED C or LED D is lit, turn off the node and remove the microprocessor. You need alcohol wipes and thermal grease to replace the microprocessor. Toggle switch block 3 (SW3) bit 6 to allow the server to power on. Restart the server. If the problem has resolved, replace the microprocessor; otherwise, reinstall the microprocessor. In either case, toggle switch block 3 (SW3) bit 6 back to its original position.</li> <li>7. If no device was isolated, and if LED AUX is lit, turn off the node and remove the operator-information panel. Toggle switch block 3 (SW3) bit 6 to allow the server to power on. Restart the server. If the problem was resolved, replace the operator-information panel; otherwise, reinstall the operator-information panel. In either case, toggle switch block 3 (SW3) bit 6 back to its original position.</li> <li>8. If no failing device is isolated, replace the system board.</li> </ol>
LOG	<p>An error occurred. Connect a keyboard and a monitor. Check the IMM system event log and the system error log for information about the error. Replace any components that are identified in the error logs.</p>
LINK	<p>This is not used on the SAN Volume Controller 2145-CF8. Replace the system board.</p>
PS	<p>Power supply 1 or power supply 2 has failed. Perform the following actions to resolve the problem:</p> <ol style="list-style-type: none"> <li>1. Check the power supply that has a lit amber LED.</li> <li>2. Make sure that the power supplies are seated correctly.</li> <li>3. Remove one of the power supplies to isolate the failed power supply.</li> <li>4. Replace the failed power supply.</li> </ol>

Table 57. Diagnostics panel LED prescribed actions (continued)

Diagnostics panel LED	Action
PCI	<p>An error occurred on a PCI bus or on the system board. An additional LED is lit next to a failing PCI slot. Perform the following actions to resolve the problem:</p> <ol style="list-style-type: none"> <li>1. Identify the failing adapter by checking the LEDs on the PCI slots.</li> <li>2. If PCI slot 1 shows an error, replace the 4-port fibre-channel adapter assembly.</li> <li>3. If PCI slot 2 shows an error, replace the high-speed SAS adapter assembly.</li> <li>4. If the error is not resolved, replace the system board.</li> </ol>
SP	<p>A service processor error was detected. Perform the following actions to resolve the problem:</p> <ol style="list-style-type: none"> <li>1. Remove power from the node. Reconnect the server to the power, and restart the node.</li> <li>2. If the problem remains, replace the system board.</li> </ol>
FAN	<p>A fan has failed, is operating too slowly, or was removed. A failing fan can also cause the TEMP LED to be lit. Perform the following actions to resolve the problem:</p> <ol style="list-style-type: none"> <li>1. Reseat the failing fan, which is indicated by a lit LED near the fan connector on the system board.</li> <li>2. If the problem remains, replace the failing fan.</li> </ol>
TEMP	<p>The system temperature exceeded a threshold level. A failing fan can cause the TEMP LED to be lit. Perform the following actions to resolve the problem:</p> <ol style="list-style-type: none"> <li>1. Make sure that the heat sink is seated correctly.</li> <li>2. Determine whether a fan has failed. If it has, replace it.</li> <li>3. Verify that the ambient temperature is within normal operating specifications.</li> <li>4. Make sure that airflow in and around the SAN Volume Controller 2145-CF8 is not obstructed.</li> </ol>
MEM	<p>A memory configuration or a memory error that is not valid has occurred. Both the MEM LED and CNFG LED might be lit. Perform the following actions to resolve the problem:</p> <ol style="list-style-type: none"> <li>1. Check that all the memory DIMMs are correctly installed.</li> <li>2. If any memory error LEDs are lit, replace the indicated memory module.</li> <li>3. If the MEM LED and the CNFG LED are lit, adjust the memory so that DIMM slots 2, 3, 5, 6, 7, and 8 are the only ones used.</li> </ol>
NMI	<p>A non-maskable interrupt occurred or the NMI button was pressed. This situation should not occur. If the NMI button on the light path diagnostic panel was pressed by mistake, restart the node; otherwise, call your support center.</p>
CNFG	<p>A hardware configuration error occurred. If the MEM LED is also lit, follow the actions shown for MEM LED. If the CPU LED is lit, check to see if a microprocessor is installed in CPU 2. If one is installed, remove it because the configuration is not supported. If no other light path LEDs are lit, replace the FRUs in the order shown until the problem is resolved:</p> <ol style="list-style-type: none"> <li>1. Operator-information panel</li> <li>2. Operator-information panel cable</li> <li>3. System board</li> </ol>

Table 57. Diagnostics panel LED prescribed actions (continued)

Diagnostics panel LED	Action
CPU	<p>A microprocessor failed or a microprocessor configuration is not valid. Both the CPU LED and the CNFG LED might be lit. Perform the following actions:</p> <ol style="list-style-type: none"> <li>1. Check the system board error LEDs.</li> <li>2. If CPU 1 error LED is lit, check that the microprocessor is correctly installed.</li> <li>3. If the error persists, replace the microprocessor.</li> <li>4. If the error persists, replace the system board.</li> </ol>
VRM	This is not used on the SAN Volume Controller 2145-CF8.
DASD	<p>A disk drive failed or is missing. A SAN Volume Controller 2145-CF8 must have its system hard disk drive installed in drive slot 4. Up to four optional solid-state drives (SSDs) can be installed in drive slots 0 to 3.</p> <p>If an SSD has been deliberately removed from a slot, the system error LED and the DASD diagnostics panel LED will light. The error is maintained even if the SSD is replaced in a different slot. If an SSD has been removed or moved, the error is cleared by powering off the node using MAP 5350, removing both the power cables, replacing the power cables, and then restarting the node.</p> <p>Resolve any node or cluster errors that relate to SSDs or the system disk drive.</p> <p>If an error is still shown, power off the node and reseal all the drives.</p> <p>If the error remains, replace the following components in the order listed:</p> <ol style="list-style-type: none"> <li>1. The system disk drive.</li> <li>2. The disk backplane.</li> </ol>
RAID	This is not used on the SAN Volume Controller 2145-CF8.
BRD	<p>An error occurred on the system board. Perform the following actions to resolve the problem:</p> <ol style="list-style-type: none"> <li>1. Check the LEDs on the system board to identify the component that caused the error. The BRD LED can be lit because of any of the following reasons: <ul style="list-style-type: none"> <li>• Battery</li> <li>• Missing PCI riser-card assembly. There must be a riser card in PCI slot 2 even if the optional adapter is not present.</li> <li>• Failed voltage regulator</li> </ul> </li> <li>2. Replace any failed or missing replacement components, such as the battery or PCI riser-card assembly.</li> <li>3. If a voltage regulator fails, replace the system board.</li> </ol>

3. Continue with “MAP 5700: Repair verification” on page 404 to verify the correct operation.

## Light path for SAN Volume Controller 2145-8A4

Use the diagnostics LEDs that are located on the system board to solve hardware problems with the SAN Volume Controller 2145-8A4 node.

Ensure that the node is turned on and then perform the following steps to resolve any hardware errors that are indicated by the Error LED and light path LEDs:

1. Is the Error LED, shown in Figure 98, on the SAN Volume Controller 2145-8A4 operator-information panel on or flashing?




Figure 98. SAN Volume Controller 2145-8A4 operator-information panel

**1** Error LED

**NO** Reassess your symptoms and return to “MAP 5000: Start” on page 357.

**YES** Go to step 2.

2. (from step 1)

**Observe the state of the diagnostic LEDs on the system board.** To view the LEDs, follow these steps:

- a. Turn off the node while ensuring that its data is mirrored and synchronized. See “MAP 5350: Powering off a SAN Volume Controller node” on page 385 for more information.
- b. Identify and label all the cables that are attached to the node so they can be replaced in the same port. Remove the node from the rack and place it on a flat, static-protective surface.
- c. Remove the top cover.
- d. Turn on the node.

3. (from step 2)

**Other than the Standby Power, Power good, and the Baseboard management controller heartbeat LEDs, are one or more LEDs on the system board on or flashing?**

**NO** Verify that the operator-information panel cable is correctly seated at both ends. If the error LED is still on but no error LEDs are illuminated on the system board, replace parts in the following sequence:

- a. Operator-information panel
- b. Operator-information panel cable
- c. System board

Go to step 5 on page 415.

**YES** Identify any diagnostic LEDs on the system board that are on. Figure 99 on page 414 shows the location of the system board LEDs. The fan LEDs are located adjacent to each fan. You can ignore the three LEDs that do not indicate an error: **13**, **14**, and **15**.




Figure 99. SAN Volume Controller 2145-8A4 system board LEDs

- 1** Fan 1 error LED
- 2** Fan 2 error LED
- 3** Fan 3 error LED
- 4** DIMM 1 error LED
- 5** DIMM 2 error LED
- 6** DIMM 3 error LED
- 7** DIMM 4 error LED
- 8** PCI Express slot 2 error LED
- 9** PCI Express slot 1 error LED
- 10** Fan 4 error LED
- 11** Fan 5 error LED
- 12** Voltage regulator error LED
- 13** Standby power LED
- 14** Power good LED
- 15** Baseboard management controller heartbeat LED
- 16** SAS/SATA controller error LED

4. (from step 3 on page 413)

**Are any diagnostic LEDs other than **13** , **14** , and **15** on the system board illuminated?**

**NO** Go to step 5 on page 415.

**YES** Refer to Table 58 and replace the parts specified for the specific LEDs one-at-a-time in the following order until the error is repaired. Then go to step 5.

Table 58. SAN Volume Controller 2145-8A4 diagnostics panel LED prescribed actions

Diagnosics panel LED	Action
DIMM error LEDs (1 through 4)	Replace parts in the following sequence: 1. Indicated DIMM 2. System board
Fan error LEDs (1 through 5)	Replace parts in the following sequence: 1. Indicated fan 2. System board
PCI Express™ slot 1 error LED	Replace parts in the following sequence: 1. PCI riser card 2. System board 3. Fibre-channel adapter
PCI Express slot 2 error LED	This is not used on the SAN Volume Controller 2145-8A4. Replace the system board.
Voltage regulator error LED	Replace the system board.
SAS/SATA controller error LED	This is not used on the SAN Volume Controller 2145-8A4. Replace the system board.

5. (from step 4 on page 414)

**Replace the top cover and place the node in the rack.** See "Replacing the SAN Volume Controller in a rack" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*. Then continue with "MAP 5700: Repair verification" on page 404 to verify the correct operation.

## Light path for SAN Volume Controller 2145-8G4

Use light path diagnostics to solve hardware problems with the SAN Volume Controller 2145-8G4 node.

Ensure that the node is turned on and then perform the following steps to resolve any hardware errors indicated by the Error LED and light path LEDs:

1. **Is the Error LED, shown in Figure 100, on the SAN Volume Controller 2145-8G4 operator-information panel illuminated or flashing?**



Figure 100. SAN Volume Controller 2145-8G4 operator-information panel

- 1 Release latch

## 2 Error LED

**NO** Reassess your symptoms and return to “MAP 5000: Start” on page 357.

**YES** Go to step 2.

### 2. (from step 1 on page 415)

Press the release latch and open the light path diagnostics panel, which is shown in Figure 101.




Figure 101. SAN Volume Controller 2145-8G4 light path diagnostics panel

### Are one or more LEDs on the light path diagnostics panel on or flashing?

**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 404.

**YES** See Table 59 on page 418 and perform the action specified for the specific light path diagnostics LEDs. Then go to step 3 on page 419. Some actions will require that you observe the state of LEDs on the system board. Figure 102 on page 417 shows the location of the system board LEDs. The fan LEDs are located adjacent to each FAN. To view the LEDs you will need to do the following:

- a. Turn off the node while ensuring that its data is mirrored and synchronized. See “MAP 5350: Powering off a SAN Volume Controller node” on page 385 for more information.
- b. Identify and label all the cables that are attached to the node so they can be replaced in the same port. Remove the node from the rack and place it on a flat, static-protective surface. See “Removing the SAN Volume Controller from a rack” in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- c. Remove the top cover and open the fan doors.



d. Press the light path diagnostics button ( **7** in Figure 102).

**Note:** The light path diagnostics button is used to illuminate the light path diagnostics LEDs when power is disconnected from the SAN Volume Controller 2145-8G4 node.




Figure 102. SAN Volume Controller 2145-8G4 system board LEDs

- 1** System-board battery error LED
- 2** DIMM 5 error LED
- 3** DIMM 6 error LED
- 4** DIMM 7 error LED
- 5** DIMM 8 error LED
- 6** Light path diagnostics active LED
- 7** Light path diagnostics button
- 8** Microprocessor 2 error LED
- 9** Microprocessor 1 error LED
- 10** DIMM 1 error LED
- 11** DIMM 2 error LED
- 12** DIMM 3 error LED
- 13** DIMM 4 error LED
- 14** System-board fault LED
- 15** Power B error LED

**16** Power A error LED

**17** Power C error LED

**18** Power D error LED

Table 59. Diagnostics panel LED prescribed actions

Diagnostics panel LED	Action
OVER SPEC	Replace parts in the following sequence: <ol style="list-style-type: none"><li>1. Power supply</li><li>2. Power backplane</li><li>3. System board</li></ol>
PS1	If you have just replaced the power supply, check that it is correctly installed. If it is correctly installed, replace parts in the following sequence: <ol style="list-style-type: none"><li>1. Power supply</li><li>2. Power backplane</li></ol>
PS2	This is not used on the SAN Volume Controller 2145-8G4. This is a false indication. A sensor has failed or the system board service processor firmware is not functioning correctly. Contact your support center to see if a firmware update is available. If not, replace parts in the following sequence: <ol style="list-style-type: none"><li>1. Power backplane</li><li>2. Operator-information panel</li><li>3. System board</li></ol>
CPU	A microprocessor has failed. Make sure that the failing microprocessor, which is indicated by a lit LED on the system board, is installed correctly. If it is installed correctly, replace the microprocessor.
VRM	This is not used on the SAN Volume Controller 2145-8G4.
CNFG	Microprocessor configuration error. Check the installed microprocessors for compatibility.
MEM	Observe the DIMM LEDs on the system board. If any DIMM LED is flashing, make sure that the correct type of DIMM is installed in every slot. Replace parts in the following sequence: <ol style="list-style-type: none"><li>1. Failing DIMM</li><li>2. System board</li></ol> <p><b>Note:</b> If more than one DIMM is indicated by the light path diagnostics, replace the DIMMs one-at-a-time, starting at the lowest-numbered DIMM slot that the diagnostics indicated.</p>
NMI	A non-maskable interrupt occurred. Call your support center and check if any software updates need to be applied to this SAN Volume Controller 2145-8G4. If this node will not join the cluster, run node recovery. If node recovery does not resolve the problem, replace the system board assembly.
S ERR	A soft error occurred. Call your support center and check if any software updates need to be applied to this SAN Volume Controller 2145-8G4. If this node will not join the cluster, run node recovery. If node recovery does not resolve the problem, replace the system board assembly.
SP	The Service processor has failed. Replace the system board assembly.

Table 59. Diagnostics panel LED prescribed actions (continued)

Diagnostics panel LED	Action
DASD	This is not used on the SAN Volume Controller 2145-8G4. A sensor has failed or the system board service processor firmware is not functioning correctly. Contact your support center to see if a firmware update is available. If not, replace parts in the following sequence: <ol style="list-style-type: none"> <li>Operator-information panel</li> <li>System board</li> </ol>
BRD	Observe the battery LED and the system board LED. If the battery LED is illuminated, replace the battery. If the system board LED is illuminated, replace the system board.
FAN	A fan has failed, is operating too slowly, or has been removed. A failing fan can also cause the TEMP LED to be lit. Replace the failing fan, which is indicated by a lit LED near the fan connector on the system board.
TEMP	If any fan failures exist, repair those before attempting this procedure. Verify that the ambient temperature is within normal operating specifications. Make sure that airflow in and around the SAN Volume Controller 2145-8G4 is not obstructed. If the error persists, replace the system board.
RAID	This is not used on the SAN Volume Controller 2145-8G4.
PCI	The fibre-channel card might be failing. Ensure the fibre-channel card and the riser card are correctly installed. If the error persists, replace the fibre-channel card.

- Continue with “MAP 5700: Repair verification” on page 404 to verify the correct operation.

## Light path for SAN Volume Controller 2145-8F2 and SAN Volume Controller 2145-8F4

Use light path diagnostics to solve hardware problems with the SAN Volume Controller 2145-8F2 and SAN Volume Controller 2145-8F4 nodes.

Ensure that the node is turned on and then perform the following steps to resolve any hardware errors indicated by the Error LED and light path LEDs:

- Is the Error LED, shown in Figure 103, on the SAN Volume Controller 2145-8F2 or the SAN Volume Controller 2145-8F4 operator-information panel illuminated or flashing?**




Figure 103. SAN Volume Controller 2145-8F4 operator-information panel

- 1** Error LED
- 2** Release latch

**NO** Reassess your symptoms and return to “MAP 5000: Start” on page 357.

**YES** Go to step 2.

2. (from step 1 on page 419)

Press the release latch and open the light path diagnostics panel, which is shown in Figure 104.

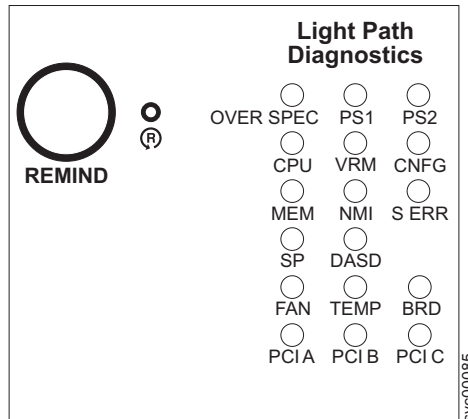


Figure 104. SAN Volume Controller 2145-8F2 and SAN Volume Controller 2145-8F4 light path diagnostics panel

**Are one or more LEDs on the light path diagnostics panel on or flashing?**

**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. Cable, signal, front panel
- c. Frame assembly

Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

**YES** Refer to Table 60 on page 422 and perform the action specified for the specific light path diagnostics LEDs, then go to step 3 on page 423. Some actions will require that you observe the state of LEDs on the system board or on the fan backplanes. The location of the system board LEDs are shown in Figure 105 on page 421. The fan LEDs are located adjacent to each FAN. To view the LEDs you will need to do the following:

- a. Turn off the node while ensuring that its data is mirrored and synchronized. See “MAP 5350: Powering off a SAN Volume Controller node” on page 385 for more information.
- b. Identify and label all the cables that are attached to the node so they can be replaced in the same port. Remove the node from the rack and place it on a flat, static-protective surface. See “Removing the SAN Volume Controller from a rack” in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- c. Remove the top cover and open the fan doors.
- d. Press the light path diagnostics button **1**. See Figure 105 on page 421.

**Note:** The light path diagnostics button is used to illuminate the light path diagnostics LEDs when power is disconnected from the SAN Volume Controller 2145-8F2 or SAN Volume Controller 2145-8F4 node.




Figure 105. SAN Volume Controller 2145-8F2 and SAN Volume Controller 2145-8F4 system board LEDs

- 1** Light path diagnostics button
- 2** System board fault LED
- 3** Light path activity LED
- 4** VRM 2 Error LED
- 5** CPU 2 Error LED
- 6** CPU 1 Error LED
- 7** VRM 1 Error LED
- 8** Battery LED
- 9** DIMM 1 error LED
- 10** DIMM 2 error LED
- 11** DIMM 3 error LED
- 12** DIMM 4 error LED
- 13** DIMM 5 error LED
- 14** DIMM 6 error LED
- 15** DIMM 7 error LED
- 16** DIMM 8 error LED

Table 60. Diagnostics panel LED prescribed actions

Diagnostics panel LED	Action
OVER SPEC	Replace the power supply
PS1	<p>If you have just replaced the power supply, check that it is correctly installed. If it is correctly installed, replace parts in the following sequence:</p> <ol style="list-style-type: none"> <li>1. Power supply</li> <li>2. Power backplane</li> </ol>
PS2	<p>This is not used on the SAN Volume Controller 2145-8F2 nor the SAN Volume Controller 2145-8F4. A sensor has failed or the system board service processor firmware is not functioning correctly. Contact your support center to see if a firmware update is available. If not, replace parts in the following sequence:</p> <ol style="list-style-type: none"> <li>1. Power backplane</li> <li>2. Operator-information panel</li> <li>3. Frame assembly</li> </ol>
CPU	<p>Observe the CPU indicators on the system board. The microprocessor adjacent to the illuminated LED is failing. If you have installed the incorrect type of microprocessor, the LED will be flashing. Replace parts in the following sequence:</p> <ol style="list-style-type: none"> <li>1. Microprocessor</li> <li>2. Frame assembly</li> </ol>
VRM	<p>Observe the VRM indicators on the system board. The VRM adjacent to the illuminated LED is failing. Verify that the VRM is correctly installed. Replace parts in the following sequence:</p> <ol style="list-style-type: none"> <li>1. VRM</li> <li>2. Frame assembly</li> </ol>
CNFG	<p>Observe all system board LEDs. Make sure that DIMMs, microprocessors, and VRMs are installed correctly and are of the correct type. Replace parts in the following sequence:</p> <ol style="list-style-type: none"> <li>1. Component adjacent to the illuminated LED</li> <li>2. Frame assembly</li> </ol>
MEM	<p>Observe the DIMM LEDs on the system board. If any DIMM LED is flashing, make sure that the correct type of DIMM is installed in every slot. Replace parts in the following sequence:</p> <ol style="list-style-type: none"> <li>1. Failing DIMM</li> <li>2. Frame assembly</li> </ol> <p><b>Note:</b> If more than one DIMM is indicated by the light path diagnostics, replace the DIMMs one-at-a-time, starting at the lowest-numbered DIMM slot that the diagnostics indicated.</p>
NMI	<p>A non-maskable interrupt occurred. Call your support center and check if any software updates need to be applied to this SAN Volume Controller 2145-8F2 or SAN Volume Controller 2145-8F4. If this node will not join the cluster, run node recovery. If node recovery does not resolve the problem, replace the frame assembly.</p>
S ERR	<p>A soft error occurred. Call your support center and check if any software updates need to be applied to this SAN Volume Controller 2145-8F2 or SAN Volume Controller 2145-8F4. If this node will not join the cluster, run node recovery. If node recovery does not resolve the problem, replace the frame assembly.</p>
SP	<p>The Service processor has failed. Replace the frame assembly.</p>

Table 60. Diagnostics panel LED prescribed actions (continued)

Diagnostics panel LED	Action
DASD	<p>This is not used on the SAN Volume Controller 2145-8F2 or SAN Volume Controller 2145-8F4. This is a false indication. A sensor has failed or the system board service processor firmware is not functioning correctly. Contact your support center to see if a firmware update is available. If not, replace parts in the following sequence:</p> <ol style="list-style-type: none"> <li>1. Operator-information panel</li> <li>2. Frame assembly</li> </ol>
FAN	<p>Observe the LEDs on the fan backplanes. The fan adjacent to the failing LED is failing. Replace parts in the following sequence:</p> <ol style="list-style-type: none"> <li>1. Fan</li> <li>2. Fan backplane</li> </ol>
TEMP	<p>If any fan failures exist, repair those before attempting this procedure. Verify that the ambient temperature is within normal operating specifications. Make sure that airflow in and around the SAN Volume Controller 2145-8F2 or SAN Volume Controller 2145-8F4 is not obstructed. Replace the frame assembly.</p>
BRD	<p>Observe the battery LED and the system board LED. If the battery LED is illuminated, replace the battery. If the system board LED is illuminated, replace the frame assembly.</p>
PCI A	<p>This is not used on the SAN Volume Controller 2145-8F2 nor SAN Volume Controller 2145-8F4. This is a false indication. A sensor has failed or the system board service processor firmware is not functioning correctly. Contact your support center to see if a firmware update is available. If not, replace parts in the following sequence:</p> <ol style="list-style-type: none"> <li>1. Operator-information panel</li> <li>2. Frame assembly</li> </ol>
PCI B	<p>One of the fibre-channel adapter cards connected to this bus might be failing. Ensure that both adapters are correctly installed and that the riser card latches are fully closed. If possible, display the fibre-channel card status on the SAN Volume Controller 2145-8F2 or SAN Volume Controller 2145-8F4 front panel to determine the failing card. Otherwise, remove the fibre-channel cards one-at-a-time to determine the failing card. Replace parts in the following sequence:</p> <ol style="list-style-type: none"> <li>1. Fibre-channel adapter card</li> <li>2. Frame assembly</li> </ol>
PCI C	<p>Replace the frame assembly.</p>

3. Continue with “MAP 5700: Repair verification” on page 404 to verify the correct operation.

---

## MAP 5900: Hardware boot

MAP 5900: Hardware boot helps you solve problems that are preventing the node from starting its boot sequence.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 7, “Using the maintenance analysis procedures,” on page 357.

This MAP applies to all SAN Volume Controller 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 have been sent here for one of the following reasons:

- The hardware boot display, shown in Figure 106, is displayed continuously.



Figure 106. Hardware boot display

- The node rescue display, shown in Figure 107, is displayed continuously.



Figure 107. Node rescue display

- The boot progress is hung and Booting 100 is displayed on the front panel
- Another MAP sent you here

Perform the following steps to allow the node to start its boot sequence:

1. **Is this a SAN Volume Controller 2145-8A4?**

**NO** Go to step 2.

**YES** Go to step 3.

2. (From step 1)

**Is the Error LED on the operator-information panel illuminated or flashing?**

**NO** Go to step 3.

**YES** Go to "MAP 5800: Light path" on page 406 to resolve the problem.

3. (From steps 1 and 2)

**If you have just installed the SAN Volume Controller node or have just replaced a field replaceable unit (FRU) inside the node, perform the following steps:**

- a. For a SAN Volume Controller 2145-8A4 node, ensure that the correct power cable assembly from the 2145 UPS-1U to the node is installed. The correct power cable assembly has red tape that binds the cables together.
- b. Identify and label all the cables that are attached to the node so they can be replaced in the same port. Remove the node from the rack and place it on a flat, static-protective surface. See "Removing the SAN Volume Controller from a rack" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- c. Remove the top cover. See "Removing the top cover" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- d. If you have just replaced a FRU, ensure that the FRU is correctly placed and that all connections to the FRU are secure.



- e. Ensure that all memory modules are correctly installed and that the latches are fully closed. See "Replacing the memory modules (DIMM)" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- f. Ensure that the fibre-channel adapter cards are correctly installed. See "Replacing the fibre-channel adapter assembly" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- g. Ensure that the disk drive and its connectors are correctly installed. See "Replacing the disk drive" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- h. Ensure that the service controller is correctly installed. See "Replacing the service controller" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- i. Replace the top cover. See "Replacing the top cover" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- j. Place the node in the rack. See "Replacing the SAN Volume Controller in a rack" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- k. Turn on the node.

**Does the boot operation still hang?**

**NO** Verify the repair by continuing with "MAP 5700: Repair verification" on page 404.

**YES** Go to step 4.

4. (from step 3 on page 424)

Check if the system BIOS is reporting any errors. You need to attach a display and keyboard to see the BIOS output. The customer should be able to supply a suitable display and keyboard.

- a. Turn off the node while ensuring that its data is mirrored and synchronized. See "MAP 5350: Powering off a SAN Volume Controller node" on page 385.
- b. Connect the keyboard **1** and the display **2**. Figure 108 shows the location of the keyboard and monitor ports. Figure 109 on page 426 shows the location of the keyboard and monitor ports on the 2145-CF8.




Figure 108. Keyboard and monitor ports on the SAN Volume Controller models 2145-8A4, 2145-8G4, and 2145-8F4 or 2145-8F2




Figure 109. Keyboard and monitor ports on the SAN Volume Controller 2145-CF8

- c. Turn on the node.
- d. Watch the display.
  - If the POST sequence indicates an error, or if the BIOS Configuration/Setup Utility program indicates an error during startup, you need to resolve the error.
  - If it indicates an error with a specific hardware item, power off the node and remove it from the rack. Ensure the item specified is correctly installed, replace the node, and then restart the node. If the error is still reported, replace the specified item.
  - If a configuration error is reported, run the Configuration/Setup Utility program option to reset the BIOS to its default (factory) settings.
- e. Turn off the node and remove the keyboard and display.
- f. Turn on the node.

**Does the boot operation still hang?**

**NO** Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

**YES** Go to step 5.

5. (from step 4 on page 425)
  - a. Turn off the node while ensuring that its data is mirrored and synchronized. See “MAP 5350: Powering off a SAN Volume Controller node” on page 385.
  - b. Identify and label all the cables that are attached to the node so they can be replaced in the same port. Remove the node from the rack and place it on a flat, static-protective surface. See “Removing the SAN Volume Controller from a rack” in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
  - c. Remove the top cover. See “Removing the top cover” in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
  - d. Remove some of the memory modules:
    - If you are using the SAN Volume Controller 2145-CF8, remove the memory modules in slots 2, 5, 7, and 8.
    - If you are using the SAN Volume Controller 2145-8A4, remove the memory modules in slots 2 through 4.
    - If you are using the SAN Volume Controller 2145-8G4, remove the memory modules in slots 2 and 4 through 8.
    - If you are using the SAN Volume Controller 2145-8F4 or the SAN Volume Controller 2145-8F2, remove the memory modules in slots 3 through 8.
  - e. Remove all installed fibre-channel cards.
  - f. Remove the disk drive.
  - g. Replace the top cover. See “Replacing the top cover” in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.

- h. Place the node in the rack. See "Replacing the SAN Volume Controller in a rack" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- i. Turn on the node.

**Does the boot operation still hang with the booting display or is Booting 100 displayed on the front panel?**

**Note:** With the FRUs removed, the boot will hang with a different boot failure code.

**NO** Replace the FRUs, one-at-a-time, until the failing FRU is isolated.

**YES** Go to step 6.

6. (from step 5 on page 426)

- a. Turn off the node while ensuring that its data is mirrored and synchronized. See "MAP 5350: Powering off a SAN Volume Controller node" on page 385 for more information.
- b. Identify and label all the cables that are attached to the node so they can be replaced in the same port. Remove the node from the rack and place it on a flat, static-protective surface. See "Removing the SAN Volume Controller from a rack" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*
- c. Remove the top cover. See "Removing the top cover" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- d. Replace the fibre-channel cards and the disk drive.
- e. Replace the memory modules:
  - If you are using the SAN Volume Controller 2145-CF8, replace the memory modules in slots 2, 5, 7, and 8.
  - If you are using the SAN Volume Controller 2145-8A4, replace the memory module in slot 1 with any of the removed memory modules from slots 2 through 4.
  - If you are using the SAN Volume Controller 2145-8G4, replace the memory modules in slots 1 and 3 with any two of the removed memory modules from slots 2 and 4 through 8.
  - If you are using the SAN Volume Controller 2145-8F4 or the SAN Volume Controller 2145-8F2, replace the memory modules in slots 1 and 2 with any two of the removed memory modules from slots 3 through 8.
- f. Replace the top cover. See "Replacing the top cover" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- g. Place the node in the rack. See "Replacing the SAN Volume Controller in a rack" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*.
- h. Turn on the node.

**Does the boot operation still hang with the booting display or is Booting 100 displayed on the front panel?**

**NO** Exchange the failing memory modules for new FRUs and verify the repair by continuing with "MAP 5700: Repair verification" on page 404.

**YES** Replace the parts in the following sequence:

- For the SAN Volume Controller 2145-CF8:
  - a. Service controller
  - b. System board

- For the SAN Volume Controller 2145-8A4 and SAN Volume Controller 2145-8G4:
  - a. Service controller
  - b. System board
- For the SAN Volume Controller 2145-8F4 and SAN Volume Controller 2145-8F2:
  - a. Service controller
  - b. Frame assembly

Verify the repair by continuing with “MAP 5700: Repair verification” on page 404.

#### **Related tasks**

Chapter 7, “Using the maintenance analysis procedures,” on page 357

The maintenance analysis procedures (MAPs) inform you how to analyze a failure that occurs with a SAN Volume Controller node.

“MAP 5800: Light path” on page 406

MAP 5800: Light path helps you to solve hardware problems on the SAN Volume Controller model 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, or 2145-8F2 that are preventing the node from booting.

“MAP 5700: Repair verification” on page 404

MAP 5700: Repair verification helps you to verify that field-replaceable units (FRUs) that you have exchanged for new FRUs, or repair actions that have been done have solved all the problems on the SAN Volume Controller.

#### **Related information**

“Understanding the boot codes” on page 199

The boot codes are displayed on the screen when a node is booting.

## **MAP 6000: Replace offline SSD**

MAP 6000: This procedure replaces a solid-state drive (SSD) that has failed while it is still a member of a managed disk (MDisk) group.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 7, “Using the maintenance analysis procedures,” on page 357.

This MAP applies to all SAN Volume Controller 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.

Perform the following steps to delete all the VDisks and VDisk host mapping or the managed disk group that contains the failed managed disk.

**Attention:** Back up your SAN Volume Controller configuration before you begin these steps.

### **1. Is the MDisk a managed mode MDisk or an image mode MDisk?**

#### **Managed mode MDisk**

Go to step 2.

#### **Image mode MDisk**

Go to step 3 on page 429.

### **2. (From step 1)**

- a. Record the properties of all VDisk copies, including unmirrored VDIsks, that use storage from the managed disk group that contains the offline managed disk.
- Note:** You can recover the copied VDisk data from the copy. All the data on the unmirrored VDIsks will be lost and will need to be restored from backup.
- b. Record the properties of the managed disk group and the list of managed disks in that managed disk group. Make particular note of the *node\_name*, *controller\_name*, and *location* of the SSD managed disk that needs to be replaced.
  - c. Submit the command `svcinfolnodevpd (node_name)`, where *(node\_name)* is the value of that property variable recorded in step b. Record the *front\_panel\_id* property of this node.
  - d. Delete the MDisk group that contains the offline MDisk. Use the force parameter to remove the MDisk group because it will still contain VDIsks.
  - e. Follow the instructions to replace an SSD or to remove an SSD. See "Replacing a SAN Volume Controller 2145-CF8 solid-state drive (SSD)" or "Removing a SAN Volume Controller 2145-CF8 solid-state drive (SSD)" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*. Use the *front\_panel\_id* to identify the correct node from the label on the front panel of the node and the *location* of the MDisk to identify the correct drive slot.
  - f. Run the `svctask detectmdisk` command to detect the new managed disk.
  - g. Create a new managed disk group with the same properties as the deleted MDisk group and add the new managed disk to the group.
  - h. Add all other managed disks that had been present in the deleted MDisk group into the new managed disk group.
  - i. Use the information that you recorded in substeps a and b to:
    - Add a VDisk copy for all VDIsks that had been mirrored in the original managed disk group.
    - Re-create all standard VDIsks and VDisk host mapping that had been members of the original managed disk group and restore the data from the backups.
3. (From step 1 on page 428)
    - a. Record the properties of the image mode VDisk or VDisk copy that corresponds to the offline image mode managed disk.
    - b. Record the properties of the offline managed disk, especially the *node\_name*, *controller\_name*, and *location*.
    - c. Submit the command `svcinfolnodevpd (node_name)`, where *(node\_name)* is the value of that property variable recorded in step b. Record the *front\_panel\_id* property of this node.
    - d. If the image mode VDisk is mirrored, remove the VDisk copy that is residing on the offline MDisk; otherwise, delete the image mode VDisk that corresponds to the offline MDisk.
    - e. Follow the instructions to replace an SSD or to remove an SSD. See "Replacing a SAN Volume Controller 2145-CF8 solid-state drive (SSD)" or "Removing a SAN Volume Controller 2145-CF8 solid-state drive (SSD)" in the *IBM System Storage SAN Volume Controller Hardware Maintenance Guide*. Use the *front\_panel\_id* to identify the correct node from the label on the front panel of the node and the *location* of the MDisk to identify the correct drive slot.

- f. Re-create the image mode VDisk or re-add the image mode VDisk copy on the new managed disk using the information that you recorded in substeps a and b.

**Related tasks**

Chapter 7, “Using the maintenance analysis procedures,” on page 357  
The maintenance analysis procedures (MAPs) inform you how to analyze a failure that occurs with a SAN Volume Controller node.

---

## Appendix. Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

### Features

These are the major accessibility features in the SAN Volume Controller Console:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. The following screen reader has been tested: Window-Eyes v6.1.
- You can operate all features using the keyboard instead of the mouse.
- When setting or changing an IP address on the SAN Volume Controller front panel, you can disable the fast increase and decrease address scrolling speed function of the up and down buttons to two seconds. This feature is documented in the topic that discusses initiating cluster creation from the front panel, which is located in the IBM System Storage SAN Volume Controller Information Center and the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*.

### Navigating by keyboard

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. You can navigate the SAN Volume Controller Console and help system from the keyboard by using the following key combinations:

- To traverse to the next link, button, or topic, press Tab inside a frame (page).
- To expand or collapse a tree node, press → or ←, respectively.
- To move to the next topic node, press V or Tab.
- To move to the previous topic node, press ^ or Shift+Tab.
- To scroll all the way up or down, press Home or End, respectively.
- To go back, press Alt+←.
- To go forward, press Alt+→.
- To go to the next frame, press Ctrl+Tab.
- To move to the previous frame, press Shift+Ctrl+Tab.
- To print the current page or active frame, press Ctrl+P.
- To select, press Enter.

### Accessing the publications

You can find the HTML version of the IBM System Storage SAN Volume Controller information at the following Web site:

<http://publib.boulder.ibm.com/infocenter/svcic/v3r1m0/index.jsp>

You can access this information using screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. JAWS version 10 has been tested.





---

## Notices

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

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 give 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 (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.  
3-2-12, Roppongi, Minato-ku, Tokyo 106-8711 Japan*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:**

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply 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 Corporation  
Almaden Research  
650 Harry Road  
Bldg 80, D3-304, Department 277  
San Jose, CA 95120-6099  
U.S.A.*

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.

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

Information concerning 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 may be addressed to the suppliers of those products.

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

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](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at Copyright and trademark information at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Adobe and the Adobe logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Intel logo, Intel Xeon, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

---

## Electronic emission notices

The following electronic emission statements apply to this product. The statements for other products that are intended for use with this product are included in their accompanying documentation.

## Federal Communications Commission (FCC) statement

Ensure that you are familiar with the Federal Communications Commission (FCC) statement.

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, might cause 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. Neither the provider nor the manufacturer is 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 FCC Rules. Operation is subject to the following two conditions: (1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

## Industry Canada compliance statement

This Class A digital apparatus complies with IECS-003.

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

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

## New Zealand compliance statement

Ensure that you are familiar with the New Zealand compliance statement.

This is a Class A product. In a domestic environment this product might cause radio interference, in which event the user might be required to take adequate measures.

## European Union EMC Directive conformance statement

Ensure that you are familiar with the European Union (EU) statement.

This product is in conformity with the protection requirements of EU council directive 2004/108/EC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a nonrecommended modification of the product, including the fitting of non-IBM option cards.

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

**Attention:** This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Ensure that you use properly shielded and grounded cables and connectors in order to reduce interference to radio and TV communications and to other electrical or electronic equipment. Such cables and connectors are available from IBM authorized dealers. IBM cannot accept responsibility for any interference caused by using other than recommended cables and connectors.

European community contact:

IBM Technical Regulations  
Pascalstr. 100, Stuttgart, Germany 70569

Telephone: 0049 (0)711 785 1176  
Fax: 0049 (0)711 785 1283  
E-mail: tjahn@de.ibm.com

## Germany compliance statement

### Deutschsprachiger EU Hinweis:

#### Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2004/108/EG zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55022 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung der IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung der IBM gesteckt/eingebaut werden.

EN 55022 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden: "Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

#### Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem "Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)". Dies ist die Umsetzung der EU-Richtlinie 2004/108/EG in der Bundesrepublik Deutschland.

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

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

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

Generelle Informationen:

**Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 Klasse A.**

## Japanese Voluntary Control Council for Interference (VCCI) statement

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

vcci

## People's Republic of China Class A Electronic Emission Statement

中华人民共和国“A类”警告声明

声明

此为A级产品，在生活环境中，该产品可能会造成无线电干扰。在这种情况下，可能需要用户对其干扰采取切实可行的措施。

## International Electrotechnical Commission (IEC) statement

This product has been designed and built to comply with (IEC) Standard 950.

## United Kingdom telecommunications requirements

This apparatus is manufactured to the International Safety Standard EN60950 and as such is approved in the U.K. under approval number NS/G/1234/J/100003 for indirect connection to public telecommunications systems in the United Kingdom.

## Korean Class A Electronic Emission Statement

이 기기는 업무용으로 전자파 적합등록을 받은 기기 이오니, 판매자 또는 사용자는 이점을 주의하시기 바라며, 만약 잘못 구입하셨을 때에는 구입한 곳에서 비업무용으로 교환하시기 바랍니다.

## Taiwan Class A compliance statement

警告使用者:

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

---

## European Contact Information

This topic contains the product service contact information for Europe.

European Community contact:  
IBM Technical Regulations  
Pascalstr. 100, Stuttgart, Germany 70569  
Tele: 0049 (0)711 785 1176  
Fax: 0049 (0)711 785 1283  
e-mail: <mailto:tjahn@de.ibm.com>

---

## Taiwan Contact Information

This topic contains the product service contact information for Taiwan.

IBM Taiwan Product Service Contact Information:  
IBM Taiwan Corporation  
3F, No 7, Song Ren Rd., Taipei Taiwan  
Tel: 0800-016-888

台灣IBM 產品服務聯絡方式：  
台灣國際商業機器股份有限公司  
台北市松仁路7號3樓  
電話：0800-016-888

f2c00790





---

# Index

## Numerics

- 2145 UPS
  - test and alarm-reset button 52
- 2145 UPS-1U
  - alarm 51
  - circuit breakers 52
  - connecting 48
  - connectors 52
  - controls and indicators on the front panel 48
  - description of parts 52
  - dip switches 52
  - environment 54
  - Load segment 1 indicator 51
  - Load segment 2 indicator 50
  - MAP
    - 5150: 2145 UPS-1U 376
    - 5250: repair verification 381
  - on-battery indicator 51
  - on/off button 51
  - operation 47
  - overload indicator 51
  - ports not used 52
  - power-on indicator 51
  - service indicator 51
  - test and alarm-reset button 52
  - unused ports 52
- 2145-8A4 node
  - SAN Volume Controller 2145-8A4 37
- 2145-8F2 node
  - SAN Volume Controller 41
- 2145-8F4 node
  - SAN Volume Controller 41
- 2145-8G4 node
  - SAN Volume Controller 39
- 2145-CF8 node
  - SAN Volume Controller 34

## A

- ac and dc LEDs 30
- ac power switch, cabling 45
- accessibility
  - keyboard 431
  - repeat rate of up and down buttons 149, 431
  - shortcut keys 431
- accessing
  - command-line interface 65
  - dump files 8
  - log files 8
  - vital product data (VPD) 8
- active status 147
- adding
  - managed disks (MDisks) 90
  - nodes 73
  - nodes to a cluster 76
- address
  - MAC 158

- Address Resolution Protocol (ARP) 3
- addressing
  - configuration node 3
- analyzing error logs 121
- Availability status
  - No Contact 351

## B

- back panel assembly
  - SAN Volume Controller 2145-8A4
    - connectors 21
    - indicators 20
  - SAN Volume Controller 2145-8F2
    - connectors 27
  - SAN Volume Controller 2145-8F4
    - connectors 25
  - SAN Volume Controller 2145-8G4
    - connectors 23
  - SAN Volume Controller 2145-CF8
    - connectors 18
    - indicators 18
- back-panel assembly
  - SAN Volume Controller 2145-8F2
    - indicators 27
  - SAN Volume Controller 2145-8F4
    - indicators 24
  - SAN Volume Controller 2145-8G4
    - indicators 22
- battery
  - Charging, front panel display 140
  - power 142
- boot
  - codes, understanding 199
  - failed 139
  - progress indicator 139
- buttons, navigation 12

## C

- Call Home 176, 178
- Canadian electronic emission notice 436
- changing
  - fibre-channel port speed 106
- charging 140
- checking
  - state of node Ethernet ports 69
  - state of node fibre-channel ports 67
  - status of fibre-channel ports 68
  - status of node ports 66
  - status of the node 67
- CIM log files 355
- circuit breakers
  - 2145 UPS-1U 52
  - requirements
    - SAN Volume Controller 2145-8A4 37
    - SAN Volume Controller 2145-8G4 40

- circuit breakers *(continued)*
  - requirements *(continued)*
    - SAN Volume Controller 2145-CF8 35
- CLI (command-line interface)
  - using to update cluster license 117
- CLI commands
  - svcinfo lscluster
    - displaying cluster properties 129
  - svcinfo lslicense 117
  - svctask chlicense 117
- cluster
  - error codes 239
- clusters
  - adding nodes 73, 76
  - Call Home e-mail 176, 178
  - cannot access 352
  - changing the fibre-channel port speed 106
  - deleting nodes 69, 71
  - diagnose failures 151
  - dumping configurations 120
  - error codes 238, 239
  - error logs 122
  - IP address
    - configuration node 3
  - IP failover 4
  - IPv4 address 147
  - IPv6 address 147, 148
  - login screen 352
  - logs 121
  - maintaining 123
  - metadata, saving 142
  - No Contact 351
  - options 146
  - overview 2
  - properties 129
  - recovering nodes 87
  - recovery codes 239
  - removing nodes 69, 71
  - shutting down 78, 79
  - updating
    - license 117
  - viewing
    - license 117
  - viewing feature logs 121
  - viewing properties 127
- codes
  - configuration events 183
  - event 180
  - informational event 180
  - node error 223
  - node rescue 218
- command-line interface (CLI)
  - accessing 65
  - using to update cluster license 117
  - using to view cluster license 117
- common information model (CIM) logging 355
- configuration
  - event codes 183
  - node failover 4
- configuration node 3
- connecting
  - 2145 UPS-1U 48
- connectors
  - 2145 UPS-1U 52
  - SAN Volume Controller 2145-8A4 21
  - SAN Volume Controller 2145-8F2 27
  - SAN Volume Controller 2145-8F4 25
  - SAN Volume Controller 2145-8G4 23
  - SAN Volume Controller 2145-CF8 18
- console
  - cannot access cluster 352
  - cannot access login screen 352
  - log files 354
  - No Contact status 351
  - unexpected shutdowns 353
  - WAS service keeps stopping 353
- contact information
  - European 439
  - Taiwan 439
- controls and indicators on the front panel
  - 2145 UPS
    - test and alarm-reset button 52
  - 2145 UPS-1U
    - alarm 51
    - illustration 48
    - Load segment 1 indicator 51
    - Load segment 2 indicator 50
    - on-battery indicator 51
    - on/off button 51
    - overload indicator 51
    - power-on indicator 51
    - test and alarm-reset button 52
- front-panel display 12
- SAN Volume Controller
  - navigation buttons 12
  - node status LED 12
  - select button 13
- SAN Volume Controller 2145-8A4
  - illustration 10
  - operator-information panel 14
- SAN Volume Controller 2145-8F2
  - error LED 13
  - illustration 11
  - operator information panel 16
- SAN Volume Controller 2145-8F4
  - illustration 11
  - operator information panel 16
- SAN Volume Controller 2145-8G4
  - illustration 10
  - operator information panel 15
- SAN Volume Controller 2145-CF8
  - illustration 9
  - operator-information panel 14
- status indicators
  - boot failed 139
  - boot progress 139
  - charging 140
  - error codes 140
  - hardware boot 141
  - menu options 145
  - node rescue request 141

- controls and indicators on the front panel *(continued)*
  - status indicators *(continued)*
    - power failure 142
    - powering off 142
    - recovering 142
    - restarting 143
    - shutting down 143
- create cluster
  - error codes 238
  - menu option 152
- creating
  - managed disk (MDisk) groups 90

## D

- defining FRUs
  - for the redundant ac-power switch 61
  - for the SAN Volume Controller 54
  - for the uninterruptible power supply 61
- degraded status 147
- deleting
  - nodes 69, 71
- determining
  - disk enclosure 107
  - failing enclosure 107
  - fibre-channel port speed 104
  - hardware boot failure 199
  - SAN problem 348
- Deutschsprachiger EU Hinweis 437
- diagnosing problems
  - through error codes 167
  - through error logs 167
  - through light path diagnostics 167
  - with SAN Volume Controller 167
  - with the master console 167
  - with the uninterruptible power supply 167
- directed maintenance procedures 123
- discovering
  - managed disks 92, 99
  - MDisks 92
- disk controllers
  - determining the failing 107
  - status 106
- display on front panel
  - IPv6 address 148
  - Node rescue request 217
  - node WWNN 151
  - overview 12
  - status indicators
    - boot failed 139
    - boot progress 139
    - charging 140
    - error codes 140
    - hardware boot 141
    - menu options 145
    - node rescue request 141
    - power failure 142
    - powering off 142
    - recovering 142
    - restarting 143
    - shutting down 143

- display on front panel *(continued)*
  - validate WWNN? 143
  - version 158
- displaying vital product data 128
- dump files
  - accessing 8
  - SSDs
    - collecting 118
- dumping cluster configurations 120

## E

- e-mail
  - Call Home 178
  - inventory information 179
- editing
  - IPv6 address 148
- electronic emission notices
  - Avis de conformité à la réglementation d'Industrie Canada 436
  - Deutschsprachiger EU Hinweis 437
  - European Union (EU) 436
  - Federal Communications Commission (FCC) 435
  - French Canadian 436
  - Germany 437
  - Industry Canada 436
  - International Electrotechnical Commission (IEC) 438
  - Japanese Voluntary Control Council for Interference (VCCI) 437
  - Korean 438
  - New Zealand 436
  - People's Republic of China 438
  - Taiwan 438
  - United Kingdom 438
- EMC statement, People's Republic of China 438
- enclosure 107
- error codes 190
  - front panel display 140
  - understanding 179
- error events 169
- error ID 190
- error LED 13
- errors
  - logs
    - describing the fields 173
    - error events 169
    - managing 170
    - understanding 169
    - viewing 170
  - marking as fixed 122
  - node 223
  - reporting 168
- Ethernet
  - activity LED 29
  - link failures 4, 393
  - link LED 29
  - MAP 5500 393
  - port 158
- Ethernet activity LED 18

- Ethernet ports
  - state 69
- European contact information 439
- European Union (EU), EMC Directive conformance statement 436
- event
  - codes 180
    - informational 180
- event notification 176, 178
- events
  - codes
    - configuration 183
- examples
  - redundant ac power switch cabling 45
  - SAN Volume Controller cluster in a SAN fabric 5

## F

- fabric, SAN 5
- failover, configuration node 3
- FCC (Federal Communications Commission) electronic emission notice 435
- Federal Communications Commission (FCC) electronic emission notice 435
- fibre channel
  - changing the port speed for a node in a cluster 106
  - changing the port speed for node not in a cluster 106
  - fabric, connections 101
  - network speed 104
  - port speed 104
- fibre-channel
  - LEDs 28
  - link failures 349
  - MAP 396
  - network, rescanning 92
  - port menu option 159
  - port numbers 32
  - SFP connector 349
- fibre-channel ports
  - state 67
- field replaceable units
  - redundant ac-power switch
    - describing 61
  - SAN Volume Controller
    - describing 54
    - disk drive assembly 54
    - disk drive cables 54
    - Ethernet cable 54
    - fan assembly 54
    - fibre-channel cable 54
    - fibre-channel SFP connector 54
    - frame assembly 54
    - front panel 54
    - operator-information panel 54
    - power cable assembly 54
    - service controller 54
    - system board assembly 54
  - uninterruptible power supply
    - battery 61
    - describing 61

- field replaceable units *(continued)*
  - uninterruptible power supply *(continued)*
    - electronics 61
    - frame assembly 61
    - power supply 61
- fields
  - cluster 136
  - description for the cluster vital product data 136
  - description for the node vital product data 131
  - device 131
  - error event log 173
  - fibre-adapter card 131
  - front panel 131
  - memory module 131
  - processor 131
  - processor cache 131
  - software 131
  - system board 131
  - uninterruptible power supply 131
- French Canadian electronic emission notice 436
- front panel
  - 2145 UPS-1U 48
  - booting 165
  - buttons and indicators 139
  - charging 165
  - display 12
  - ID 13
  - menu options 145
    - Ethernet 158
    - fibre-channel port-1 through 4 159
    - IPv4 address 147
    - IPv6 address 147, 148
    - node 150
    - select language 159
    - version 158
  - power failure 165
  - powering off the SAN Volume Controller 165
  - recovering 165
  - SAN Volume Controller 139

## G

- gateway
  - menu option 147, 148
  - node option 154, 156
- Germany electronic emission compliance statement 437

## H

- hard-disk drive activity LED 16
- hardware
  - boot 141, 423
  - boot failure 199
  - components 9
  - failure 141
  - node 9

## I

- I/O operations, stopped 142

- IBM System Storage Productivity Center
  - troubleshooting 351
- identification
  - label, node 13
  - name 150
  - number 150
- IEC (International Electrotechnical Commission)
  - electronic emission notice 438
- inactive status 147
- including
  - MDisks 100
- indicators and controls on the front panel
  - 2145 UPS
    - test and alarm-reset button 52
  - 2145 UPS-1U
    - alarm 51
    - illustration 48
    - Load segment 1 indicator 51
    - Load segment 2 indicator 50
    - on-battery indicator 51
    - on/off button 51
    - overload indicator 51
    - power-on indicator 51
    - test and alarm-reset button 52
  - SAN Volume Controller
    - navigation buttons 12
    - node status LED 12
    - select button 13
  - SAN Volume Controller 2145-8A4
    - illustration 10
    - operator-information panel 14
  - SAN Volume Controller 2145-8F2
    - error LED 13
    - illustration 11
    - operator information panel 16
  - SAN Volume Controller 2145-8F4
    - illustration 11
    - operator information panel 16
  - SAN Volume Controller 2145-8G4
    - illustration 10
    - operator information panel 15
  - SAN Volume Controller 2145-CF8
    - illustration 9
    - operator-information panel 14
- status indicators
  - boot failed 139
  - boot progress 139
  - charging 140
  - error codes 140
  - hardware boot 141
  - menu options 145
  - node rescue request 141
  - power failure 142
  - powering off 142
  - recovering 142
  - restarting 143
  - shutting down 143
- indicators on the rear panel
  - ac and dc LEDs 30, 31, 32
  - Ethernet
    - activity LED 29

- indicators on the rear panel *(continued)*
  - Ethernet *(continued)*
    - link LED 29
  - fibre-channel LEDs 28
  - power-supply error LED 30
  - power, location, and system-error LEDs 30
  - SAN Volume Controller 2145-8A4
    - Ethernet activity LED 29
  - SAN Volume Controller 2145-8G4
    - Ethernet activity LED 29
- information
  - center x
- information, system
  - LED 17
- informational event 169
- International Electrotechnical Commission (IEC)
  - electronic emission notice 438
- inventory information 176, 178
- IP address
  - cluster 146, 147, 148
  - IPv6 147, 148
  - service 160
- IPv4 address 147
- IPv6
  - address 147, 148
  - gateway menu option 148
  - prefix mask menu option 148
- iSCSI
  - link problems 349

## J

- Japanese electronic emission notice 437

## K

- keyboard 431
- Korean electronic emission statement 438

## L

- language menu selection options 159
- LEDs
  - ac and dc 30, 31, 32
  - diagnostics 406
  - Ethernet
    - activity 29
    - link 29
  - Ethernet activity LED 18
  - fibre-channel 28
  - hard-disk drive activity 16
  - location 18, 30
  - power 17, 30
  - power-supply error 30
  - rear-panel indicators 18, 20, 22, 24, 27
  - SAN Volume Controller 2145-8A4 20
  - SAN Volume Controller 2145-8F2 27
  - SAN Volume Controller 2145-8F4 24
  - SAN Volume Controller 2145-8G4 22
  - SAN Volume Controller 2145-CF8 18
  - system information 17

- LEDs (*continued*)
  - system-error 16, 30
- legal notices 433
- license
  - disabling features 117
  - enabling features 117
  - updating
    - using the CLI (command-line interface) 117
- license settings
  - viewing log 118
- light path MAP 406
- link failures
  - fibre-channel 349
- link problems
  - iSCSI 349
- listing
  - dump files 119
  - log files 119
  - managed disks 96
- Load segment 1 indicator 51
- Load segment 2 indicator 50
- locator LED 18
- log files
  - accessing 8
  - CIM 355
  - viewing 170
  - Websphere Application Server 354
- logging
  - CIM agent 354
  - Websphere 354

## M

- MAC address 158
- maintenance analysis procedures (MAPs)
  - 2145 UPS-1U 376
  - Ethernet 393
  - fibre-channel 396
  - front panel 390
  - hardware boot 423
  - light path 406
  - overview 357
  - power
    - SAN Volume Controller 2145-8A4 372
    - SAN Volume Controller 2145-8F2 365
    - SAN Volume Controller 2145-8F4 365
    - SAN Volume Controller 2145-8G4 365
  - repair verification 404
  - SSD failure 428
  - start 358
- maintenance procedures, clusters 123
- managed disk (MDisk) groups
  - adding
    - managed disks 90
  - creating 90
  - creating using the CLI 93
  - status 101
- managed disks (MDisks)
  - adding 95
  - discovering 92, 99
  - including 100

- managed disks (MDisks) (*continued*)
  - rebalancing access 99
- managing
  - error event log 170
- MAP
  - 5000: Start 358
  - 5050: Power 2145-CF8, 2145-8G4, 2145-8F4, and 2145-8F2 365
  - 5060: Power 2145-8A4 372
  - 5150: 2145 UPS-1U 376
  - 5250: 2145 UPS-1U repair verification 381
  - 5320: Redundant ac power 382
  - 5340: Redundant ac power verification 383
  - 5400: Front panel 390
  - 5500: Ethernet 393
  - 5600: Fibre-channel 396
  - 5700: Repair verification 404
  - 5800: Light path 406
  - 5900: Hardware boot 423
  - 6000: Replace offline SSD 428
  - power off SAN Volume Controller node 385
- MAPs (maintenance analysis procedures)
  - 2145 UPS-1U 376
  - 2145 UPS-1U repair verification 381
  - Ethernet 393
  - fibre-channel 396
  - front panel 390
  - hardware boot 423
  - light path 406
  - power
    - SAN Volume Controller 2145-8A4 372
    - SAN Volume Controller 2145-8F2 365
    - SAN Volume Controller 2145-8F4 365
    - SAN Volume Controller 2145-8G4 365
    - SAN Volume Controller 2145-CF8 365
  - power off 385
  - redundant ac power 382, 383
  - repair verification 404
  - SSD failure 428
  - start 358
  - using 357
- marking errors as fixed 122
- MDisks (managed disks)
  - adding 95
  - listing 96
- media access control (MAC) address 158
- menu options
  - clusters
    - gateway 148
    - IPv4 address 147
    - IPv4 gateway 147
    - IPv4 subnet 147
    - IPv6 address 147, 148
    - IPv6 prefix 148
    - options 146
    - recover cluster? 149
    - status 147, 151
  - Ethernet
    - MAC address 158
    - port 158
    - speed 158

- menu options (*continued*)
  - fibre-channel port-1 through 4 159
  - front panel display 145
  - IPv4 gateway 147
  - IPv6 gateway 148
  - IPv6 prefix 148
  - node
    - create cluster 152
    - options 150
    - status 151
  - SAN Volume Controller
    - active 147
    - degraded 147
    - inactive 147
    - IPv4 subnet 147
  - select language 159
  - sequence 145

## N

- navigation
  - buttons 12
  - create cluster 152
  - recover cluster 149
  - select language 159
- network speed, fibre-channel ports 104
- New Zealand electronic emission statement 436
- node
  - create cluster 152
  - failover 4
  - options
    - create cluster? 152
    - delete cluster? 157
    - gateway 156
    - IPv4 address 153
    - IPv4 create now? 155
    - IPv4 gateway 154
    - IPv4 subnet mask 154
    - IPv6 address 155
    - IPv6 create now? 157
    - IPv6 prefix 156
    - status 151
    - subnet mask 154
  - rescue request 141
  - software failure 365, 372
- node status LED 12
- nodes
  - adding 73
  - adding to a cluster 76
  - cache data, saving 142
  - configuration 3
    - addressing 3
    - failover 3
  - deleting 69, 71
  - error codes 223
  - hard disk drive failure 141
  - identification label 13
  - options
    - main 150
  - removing 69, 71
  - replacing 108, 112

- nodes (*continued*)
  - rescue
    - performing 217
  - rescue codes
    - understanding 218
  - returning to cluster 87
  - shutting down 77
  - status 66, 67
  - understanding 223
  - viewing
    - Ethernet port state 69
    - fibre-channel port state 67
    - general details 66, 128
    - vital product data 127
- not used
  - 2145 UPS-1U ports 52
  - location LED 30
- notifications
  - Call Home information 178
  - inventory information 179
  - sending 176

## O

- object classes and instances 189
- object codes 189
- object types 189
- on/off button 51
- operator information panel
  - locator LED 18
  - SAN Volume Controller 2145-8F2 16
  - SAN Volume Controller 2145-8F4 16
  - SAN Volume Controller 2145-8G4 15
  - system-information LED 17
- operator-information panel
  - Ethernet activity LED 18
  - hard-disk drive activity LED 16
  - power button 17
  - power LED 17
  - release latch 17
  - reset button 16
  - SAN Volume Controller 2145-8A4 14
  - SAN Volume Controller 2145-CF8 14
  - system-error LED 16
- overload indicator 51
- overview
  - redundant ac-power switch 43
  - SAN fabric 5
  - SAN Volume Controller 1
  - service mode 8
  - vital product data 127

## P

- panel
  - front 12
  - name 13
  - operator information
    - SAN Volume Controller 2145-8A4 14
    - SAN Volume Controller 2145-8F2 16
    - SAN Volume Controller 2145-8F4 16

- panel (*continued*)
  - operator information (*continued*)
    - SAN Volume Controller 2145-8G4 15
    - SAN Volume Controller 2145-CF8 14
  - rear
    - SAN Volume Controller 2145-8A4 20
    - SAN Volume Controller 2145-8F2 27
    - SAN Volume Controller 2145-8F4 24
    - SAN Volume Controller 2145-8G4 22
    - SAN Volume Controller 2145-CF8 18
- passwords
  - resetting 149
- People's Republic of China, electronic emission statement 438
- physical characteristics
  - 2145 UPS-1U 54
  - redundant ac-power switch 44
  - SAN Volume Controller 2145-8A4
    - connectors 21
  - SAN Volume Controller 2145-8F2
    - connectors 27
  - SAN Volume Controller 2145-8F4
    - connectors 25
  - SAN Volume Controller 2145-8G4
    - connectors 23
  - SAN Volume Controller 2145-CF8
    - connectors 18
    - service ports 19
    - unused ports 20
- port speed
  - determining 104
  - fibre-channel 159
  - for a node not in a cluster 106
  - node in a cluster 106
- ports
  - Ethernet 29
  - not used
    - 2145 UPS-1U 52
    - SAN Volume Controller 2145-8A4 21
    - SAN Volume Controller 2145-8F4 25
    - SAN Volume Controller 2145-8G4 23
  - port names, worldwide 32
  - port numbers, fibre-channel 32
  - SAN Volume Controller 2145-CF8 18
  - status 68
- POST (power-on self-test) 168
- power
  - button 17
  - controls 165
  - failure 142
  - off
    - operation 142
  - requirements
    - SAN Volume Controller 2145-8A4 37
    - SAN Volume Controller 2145-8F2 41
    - SAN Volume Controller 2145-8F4 41
    - SAN Volume Controller 2145-8G4 39
    - SAN Volume Controller 2145-CF8 35
  - restored 142
  - switch, failure 365, 372
  - uninterruptible power supply 165
- power LED 17
- Power MAP 2145-8A4 372
- Power MAP 2145-CF8, 2145-8G4, 2145-8F4, and 2145-8F2 365
- power off
  - SAN Volume Controller 385
- power-supply error LED 30
- preparing
  - SAN Volume Controller environment 34
  - uninterruptible power supply environment 53
- procedures
  - directed maintenance 123
- properties 127

## R

- rear-panel indicators
  - SAN Volume Controller 2145-8A4 20
  - SAN Volume Controller 2145-8F2 27
  - SAN Volume Controller 2145-8F4 24
  - SAN Volume Controller 2145-8G4 22
  - SAN Volume Controller 2145-CF8 18
- rebalancing
  - managed disks (MDisks) access 99
- recover cluster menu option
  - entering service mode 149
  - navigation 149
  - resetting the password 149
- recovering
  - front panel display 142
  - offline virtual disks (VDisks) 85
    - using CLI 86
- redundant ac power switch
  - cabling 45
  - examples 45
- redundant ac-power switch
  - environment preparation 44
  - field replaceable units 61
  - MAP 382, 383
  - overview 43
  - problems 382
  - specifications 44
  - verifying 383
- related information x
- release latch 17
- removing
  - node from a cluster 157
  - nodes 69, 71
- Repair verification MAP 404
- repairing
  - space-efficient VDisk 85
  - space-efficient VDIsks 84
  - VDisk copies 81
- replacing
  - nodes 108, 112
- reporting
  - errors 168
- requirements
  - 2145 UPS-1U 47
  - ac voltage 34, 36, 37, 38, 39, 40
  - circuit breakers 35, 37, 40



- requirements *(continued)*
  - electrical 34, 35, 37, 39
  - power 35, 37, 39
  - SAN Volume Controller 2145-8A4 37
  - SAN Volume Controller 2145-8G4 39
  - SAN Volume Controller 2145-CF8 34
- rescanning the fibre-channel network 92
- rescue
  - node
    - codes, understanding 218
    - performing 217
- reset button 16
- resetting passwords 149
- running
  - cluster maintenance procedure 123

## S

- SAN (storage area network)
  - fabric overview 5
  - problem determination 348
- SAN Volume Controller
  - 2145 UPS-1U 48
  - field replaceable units
    - 4-port fibre-channel adapter 54
    - 40×40×28 fan 54
    - 40×40×56 fan 54
    - alcohol wipe 54
    - CMOS battery 54
    - disk backplane 54
    - disk controller 54
    - disk drive assembly 54
    - disk drive cables 54
    - disk power cable 54
    - disk signal cable 54
    - Ethernet cable 54
    - fan assembly 54
    - fan power cable 54
    - fibre-channel adapter assembly 54
    - fibre-channel cable 54
    - fibre-channel HBA 54
    - frame assembly 54
    - front panel 54
    - memory module 54
    - microprocessor 54
    - operator-information panel 54
    - power backplane 54
    - power cable assembly 54
    - power supply assembly 54
    - riser card, PCI 54
    - riser card, PCI Express 54
    - service controller 54
    - service controller cable 54
    - system board 54
    - thermal grease 54
    - voltage regulator module 54
  - front-panel display 139
  - hardware 1
  - hardware components 9
  - menu options
    - active 147

- SAN Volume Controller *(continued)*
  - menu options *(continued)*
    - create cluster 152
    - degraded 147
    - Ethernet 158
    - fibre-channel port-1 through 4 159
    - inactive 147
    - IPv4 subnet 147
    - IPv6 prefix 148
    - node 150
    - select language? 159
  - node 9
  - overview 1
  - power control 165
  - power off 385
  - power-on self-test 168
  - preparing environment 34
  - properties 128
  - shutting down 77
  - shutting down the cluster 78
  - software
    - overview 1
- SAN Volume Controller 2145-8A4
  - additional space requirements 39
  - air temperature without redundant ac power 38
  - circuit breaker requirements 37
  - connectors 21
  - controls and indicators on the front panel 10
  - dimensions and weight 38
  - fibre-channel
    - port number 33
  - heat output of node 39
  - humidity with redundant ac power 38
  - humidity without redundant ac power 38
  - indicators and controls on the front panel 10
  - input-voltage requirements 37
  - light path MAP 412
  - MAP 5800: Light path 412
  - nodes
    - heat output 39
  - not used, service ports 21
  - operator-information panel 14
  - ports 21
  - power requirements for each node 37
  - product characteristics 37
  - rear-panel indicators 20
  - requirements 37
  - specifications 37
  - temperature with redundant ac power 38
  - weight and dimensions 38
- SAN Volume Controller 2145-8F2
  - air temperature 41
  - connectors 27
  - controls and indicators on the front panel 11
  - dimensions and weight 41
  - fibre-channel port number 34
  - heat output 41
  - humidity 41
  - indicators and controls on the front panel 11
  - light path MAP 419
  - MAP 5800: Light path 419

- SAN Volume Controller 2145-8F2 *(continued)*
  - operator information panel 16
  - product characteristics 41
  - rear-panel indicators 27
  - specifications 41
  - weight and dimensions 41
- SAN Volume Controller 2145-8F4
  - air temperature 41
  - connectors 25
  - controls and indicators on the front panel 11
  - dimensions and weight 41
  - fibre-channel
    - port number 34
  - heat output 41
  - humidity 41
  - indicators and controls on the front panel 11
  - light path MAP 419
  - MAP 5800: Light path 419
  - operator information panel 16
  - product characteristics 41
  - rear-panel indicators 24
  - specifications 41
  - weight and dimensions 41
- SAN Volume Controller 2145-8G4
  - additional space requirements 41
  - air temperature without redundant ac power 40
  - circuit breaker requirements 40
  - connectors 23
  - controls and indicators on the front panel 10
  - dimensions and weight 41
  - fibre-channel
    - port number 33
  - heat output of node 41
  - humidity with redundant ac power 40
  - humidity without redundant ac power 40
  - indicators and controls on the front panel 10
  - input-voltage requirements 39
  - light path MAP 415
  - MAP 5800: Light path 415
  - nodes
    - heat output 41
  - not used, service ports 23
  - operator information panel 15
  - ports 23
  - power requirements for each node 39
  - product characteristics 39
  - rear-panel indicators 22
  - requirements 39
  - specifications 39
  - temperature with redundant ac power 40
  - weight and dimensions 41
- SAN Volume Controller 2145-CF8
  - additional space requirements 36
  - air temperature without redundant ac power 35
  - circuit breaker requirements 35
  - connectors 18
  - controls and indicators on the front panel 9
  - dimensions and weight 36
  - heat output of node 37
  - humidity with redundant ac power 36
  - humidity without redundant ac power 35
- SAN Volume Controller 2145-CF8 *(continued)*
  - indicators and controls on the front panel 9
  - input-voltage requirements 34
  - light path MAP 406
  - MAP 5800: Light path 406
  - nodes
    - heat output 37
  - operator-information panel 14
  - ports 18
  - power requirements for each node 35
  - product characteristics 34
  - rear-panel indicators 18
  - requirements 34
  - service ports 19
  - specifications 34
  - temperature with redundant ac power 36
  - unused ports 20
  - weight and dimensions 36
- SAN Volume Controller Console
  - shut down a node 385
  - troubleshooting 351, 353
- SAN Volume Controller library
  - related publications x
- scanning
  - fibre-channel network 99
  - rebalancing MDisk access 99
- self-test, power-on 168
- serial number 13
- service
  - actions, uninterruptible power supply 47
- service controller
  - replacing
    - validate WWNN 143
- service mode
  - navigation 160
  - options 160
  - overview 8
  - setting 149
- service ports
  - SAN Volume Controller 2145-CF8 19
- shortcut keys 431
- shutting down
  - clusters 78
  - front panel display 143
  - nodes 77
- SNMP traps 176
- software
  - failure, MAP 5050 365
  - failure, MAP 5060 372
  - overview 1
  - version
    - display 158
- solid-state drive (SSD)
  - locating 98
- space requirements
  - SAN Volume Controller 2145-8A4 39
  - SAN Volume Controller 2145-8G4 41
  - SAN Volume Controller 2145-CF8 36
- space-efficient VDisks
  - repairing 84

- specifications
  - redundant ac-power switch 44
- speed
  - fibre-channel network 104
  - fibre-channel port 159
- Start MAP 358
- state
  - of node Ethernet ports 69
  - of node fibre-channel ports 67
- status
  - active 147
  - degraded 147
  - disk controller 106
  - inactive 147
  - managed disk group 101
  - of fibre-channel ports 68
  - of node 67
  - of node ports 66
  - operational 147, 151
- storage area network (SAN)
  - fabric overview 5
  - problem determination 348
- storage systems
  - servicing 350
- subnet
  - menu option 147
- subnet mask
  - node option 154
- summary of changes ix
- switches
  - 2145 UPS-1U 52
  - redundant ac power 43
- syslog
  - messages 177
- syslog messages 176
- system-error LED 16

## T

- Taiwan
  - contact information 439
  - electronic emission notice 438
- test and alarm-reset button 52
- trademarks 435
- troubleshooting
  - cluster 351, 352
  - event notification e-mail 176, 178
  - SAN failures 348
  - SAN Volume Controller Console 353
  - using error logs 140
  - using the front panel 139
  - Web page cannot be displayed 352
- troubleshootingIBM System Storage Productivity Center
  - SAN Volume Controller Console 351

## U

- understanding
  - cluster error codes 239
  - cluster recovery codes 239
  - create cluster error codes 238

- understanding (*continued*)
  - error codes 179
  - error event log 169
  - fields for the cluster vital product data 136
  - fields for the node vital product data 131
  - node error codes 223
  - node rescue codes 218
- uninterruptible power supply
  - 2145 UPS-1U
    - controls and indicators 48
    - environment 54
    - operation 47
    - overview 47
  - field replaceable units
    - battery 61
    - electronics 61
    - power supply 61
  - front panel MAP 390
  - operation 47
  - overview 47
  - preparing environment 53
- United Kingdom electronic emission notice 438
- unused ports
  - 2145 UPS-1U 52
  - SAN Volume Controller 2145-8A4 21
  - SAN Volume Controller 2145-8F4 25
  - SAN Volume Controller 2145-8G4 23
  - SAN Volume Controller 2145-CF8 20
- Updating
  - license
    - using the CLI (command-line interface) 117
- using
  - directed maintenance procedures 123
  - error code tables 180
  - SAN Volume Controller CLI 63
  - SAN Volume Controller Console 63

## V

- validating
  - VDisk copies 82
- VDisk copies
  - validating 82
- VDisks (virtual disks)
  - listing node dependent 80
  - offline 88
  - recovering 88
  - recovering from offline 85
    - using CLI 86
    - using the CLI 88
- verifying
  - VDisk copies 81
- viewing
  - error event log 170
  - virtual disk status 81
- Viewing
  - license
    - using the CLI (command-line interface) 117
- virtual disks (VDisks)
  - copies, repairing 81
  - copies, verifying 81

- virtual disks (VDisks) *(continued)*
  - offline 88
  - recovering 88
  - recovering from offline 85
    - using CLI 86
  - status 81
  - using the CLI 88
- vital product data (VPD)
  - accessing 8
  - displaying 128
  - overview 127
  - understanding the fields for the cluster 136
  - understanding the fields for the node 131
  - viewing
    - nodes 127
- VPD (vital product data)
  - accessing 8
  - displaying 128
  - overview 127
  - understanding the fields for the cluster 136
  - understanding the fields for the node 131

## W

- Websphere Application Server
  - logging 354
- WebSphere Application Server
  - service keeps stopping 353
- worldwide node names
  - choose 143
  - display 151
  - node, front panel display 151
  - validate, front panel display 143
- worldwide port names (WWPNs)
  - description 32

---

# Readers' Comments — We'd Like to Hear from You

**IBM System Storage SAN Volume Controller  
Troubleshooting Guide  
Version 5.1.0**

**Publication No. GC27-2227-02**

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Send your comments to the address on the reverse side of this form.

If you would like a response from IBM, please fill in the following information:

\_\_\_\_\_  
Name

\_\_\_\_\_  
Address

\_\_\_\_\_  
Company or Organization

\_\_\_\_\_  
Phone No.

\_\_\_\_\_  
E-mail address



Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE  
NECESSARY  
IF MAILED IN THE  
UNITED STATES

# BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation  
Information Development  
Department 61C  
9032 South Rita Road  
Tucson, Arizona  
USA 85775-4401



Fold and Tape

Please do not staple

Fold and Tape





Part Number: 31P1331

Printed in USA

GC27-2227-02



(1P) P/N: 31P1331





Spine information:



IBM System Storage SAN Volume  
Controller

**SAN Volume Controller Troubleshooting Guide**

Version 5.1.0