



IBM System Storage SAN Volume Controller Troubleshooting Guide

Version 6.1.0

GC27-2284-00





IBM System Storage SAN Volume Controller Troubleshooting Guide

Version 6.1.0

GC27-2284-00

Note

Before using this information and the product it supports, read the information in “Notices” on page 263.

This edition applies to the IBM System Storage SAN Volume Controller System Storage SAN Volume Controller and Storwize V7000, Version 6.1.0, and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces GC27-2227-02.

© **Copyright IBM Corporation 2003, 2010.**

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

Contents

Figures v

Tables vii

About this guide ix

Who should use this guide ix

| Summary of changes for GC27-2284-00 SAN Volume

| Controller Troubleshooting Guide ix

Summary of changes for GG27-2227-01 and

GG27-2227-02 SAN Volume Controller

Troubleshooting Guide x

Emphasis xi

SAN Volume Controller library and related

publications xi

How to order IBM publications xii

Sending your comments xiii

**Chapter 1. SAN Volume Controller
overview 1**

Clusters 4

 Configuration node 4

 Configuration node addressing 5

 Cluster IP failover 5

SAN fabric overview 6

**Chapter 2. Introducing the SAN Volume
Controller hardware components 9**

SAN Volume Controller nodes 9

 SAN Volume Controller front panel controls and
 indicators 9

 SAN Volume Controller operator-information
 panel 13

 SAN Volume Controller rear-panel indicators and
 connectors 17

 Fibre Channel port numbers and worldwide port
 names 29

 Requirements for the SAN Volume Controller
 environment 31

Redundant ac-power switch 40

 Redundant ac-power environment requirements 40

 Cabling of redundant ac-power switch (example) 41

Uninterruptible power supply 43

 2145 UPS-1U 43

 Uninterruptible power-supply environment
 requirements 47

Defining the SAN Volume Controller FRUs 48

 SAN Volume Controller FRUs 48

 Redundant ac-power switch FRUs 54

**Chapter 3. SAN Volume Controller user
interfaces for servicing your system 55**

Management GUI interface 55

 When to use the management GUI 55

 Accessing the management GUI 56

 Deleting a node from a cluster using the
 management GUI 56

 Adding nodes to a cluster 58

Service assistant interface 60

 When to use the service assistant 60

 Accessing the service assistant 60

 Adding a node to a cluster 61

Cluster command-line interface 61

 When to use the cluster CLI 61

 Accessing the cluster CLI 62

Service command-line interface 62

 When to use the service CLI 62

 Accessing the service CLI 62

**Chapter 4. Performing recovery actions
using the SAN Volume Controller CLI 63**

Validating and repairing mirrored VDisk copies
using the CLI 63

Repairing a space-efficient VDisk using the CLI 64

Recovering from offline VDIsks using the CLI 65

Replacing nodes nondisruptively 66

**Chapter 5. Viewing the vital product
data 71**

Viewing the vital product data using the
management GUI 71

Displaying the vital product data using the CLI 71

 Displaying node properties using the CLI 71

 Displaying cluster properties using the CLI 72

Fields for the node VPD 73

Fields for the cluster VPD 78

**Chapter 6. Using the front panel of the
SAN Volume Controller. 81**

Boot progress indicator 81

Boot failed. 81

Charging 82

Error codes 82

Hardware boot 82

Node rescue request 82

Power failure. 83

Powering off 83

Recovering 83

Restarting 84

Shutting down 84

Validate WWNN? 84

SAN Volume Controller menu options 85

 Cluster options 87

 Node options. 88

 Version options 89

 Ethernet options. 89

 Fibre-channel port-1 through 4 options 90

 Actions options 91

 Language? option 103

| | |
|--|-----|
| Using the power control for the SAN Volume Controller node | 104 |
|--|-----|

Chapter 7. Diagnosing problems 107

| | |
|--|-----|
| Event reporting. | 107 |
| Power-on self-test | 108 |
| Understanding events | 108 |
| Managing the event log | 108 |
| Viewing the event log | 108 |
| Describing the fields in the event log | 109 |
| Event notifications. | 109 |
| Inventory information email | 112 |
| Understanding the error codes. | 113 |
| Using the error code tables | 113 |
| Event codes | 113 |
| SCSI event reporting | 119 |
| Object types | 122 |
| Error codes | 122 |
| Determining a hardware boot failure | 131 |
| Boot code reference | 131 |
| Node error code overview | 132 |
| Cluster code overview | 133 |
| Error code range | 134 |
| Bootting codes | 134 |
| Create cluster errors | 136 |
| Node errors | 136 |
| Cluster recovery and states | 145 |
| Cluster error codes | 146 |
| SAN problem determination | 186 |
| Fibre Channel link failures | 186 |
| Ethernet iSCSI host-link problems | 187 |
| Servicing storage systems | 187 |

Chapter 8. Recovery procedures 189

| | |
|--|-----|
| Recover cluster procedure | 189 |
| When to run the recover cluster procedure | 189 |
| Fix hardware errors | 190 |
| Removing cluster information for nodes with error code 550 or error code 578 using the front panel | 190 |
| Removing cluster information for nodes with error code 550 or error code 578 using the service assistant | 191 |
| Performing cluster recovery procedure using the front panel | 191 |
| Performing cluster recovery using the service assistant | 193 |
| What to check after running cluster recovery | 193 |
| Backing up and restoring the cluster configuration | 194 |
| Backing up the cluster configuration using the CLI. | 195 |
| Restoring the cluster configuration | 197 |
| Deleting backup configuration files using the CLI. | 199 |
| Performing the node rescue when the node boots | 199 |

Chapter 9. Using the maintenance analysis procedures 201

| | |
|---------------------------|-----|
| MAP 5000: Start | 201 |
|---------------------------|-----|

| | |
|--|-----|
| MAP 5050: Power 2145-CF8, 2145-8G4, 2145-8F4, and 2145-8F2 | 207 |
| MAP 5060: Power 2145-8A4 | 213 |
| MAP 5150: 2145 UPS-1U. | 216 |
| MAP 5250: 2145 UPS-1U repair verification | 221 |
| MAP 5320: Redundant ac power | 222 |
| MAP 5340: Redundant ac power verification | 223 |
| MAP 5350: Powering off a SAN Volume Controller node | 225 |
| Using the management GUI to power off a cluster. | 226 |
| Using the SAN Volume Controller CLI to power off a node | 227 |
| Using the SAN Volume Controller Power control button | 228 |
| MAP 5400: Front panel | 229 |
| MAP 5500: Ethernet | 231 |
| Defining an alternate configuration node | 233 |
| MAP 5600: Fibre-channel | 234 |
| MAP 5700: Repair verification | 239 |
| MAP 5800: Light path | 240 |
| Light path for SAN Volume Controller 2145-CF8 | 240 |
| Light path for SAN Volume Controller 2145-8A4 | 245 |
| Light path for SAN Volume Controller 2145-8G4 | 248 |
| Light path for SAN Volume Controller 2145-8F2 and SAN Volume Controller 2145-8F4 | 252 |
| MAP 5900: Hardware boot | 256 |

Appendix. Accessibility 261

Notices 263

| | |
|--|-----|
| Trademarks | 265 |
| Electronic emission notices | 265 |
| Federal Communications Commission (FCC) statement. | 265 |
| Industry Canada compliance statement. | 266 |
| Avis de conformité à la réglementation d'Industrie Canada | 266 |
| Australia and New Zealand Class A Statement | 266 |
| European Union Electromagnetic Compatibility Directive | 266 |
| Germany Electromagnetic compatibility directive | 266 |
| Japan VCCI Council Class A statement | 267 |
| People's Republic of China Class A Electronic Emission Statement | 267 |
| International Electrotechnical Commission (IEC) statement. | 268 |
| United Kingdom telecommunications requirements | 268 |
| Korean Communications Commission (KCC) Class A Statement | 268 |
| Russia Electromagnetic Interference (EMI) Class A Statement | 268 |
| Taiwan Class A compliance statement | 268 |
| European Contact Information. | 268 |
| Taiwan Contact Information | 268 |

Index 271

Figures

| | | | | | |
|-----|--|----|-----|---|-----|
| 1. | SAN Volume Controller system in a fabric | 2 | 32. | SAN Volume Controller 2145-8F2 external connectors | 25 |
| 2. | Data flow in a SAN Volume Controller system | 3 | 33. | Power connector | 26 |
| 3. | Configuration node | 4 | 34. | SAN Volume Controller 2145-CF8 ac, dc, and power-error LEDs | 28 |
| 4. | SAN Volume Controller 2145-CF8 front panel | 9 | 35. | SAN Volume Controller 2145-8G4 ac and dc LEDs. | 29 |
| 5. | SAN Volume Controller 2145-8A4 front-panel assembly | 10 | 36. | SAN Volume Controller 2145-8F4 and SAN Volume Controller 2145-8F2 ac and dc LEDs | 29 |
| 6. | SAN Volume Controller 2145-8G4 front-panel assembly | 10 | 37. | The physical port numbers for the SAN Volume Controller 2145-8A4 | 30 |
| 7. | SAN Volume Controller 2145-8F2 and SAN Volume Controller 2145-8F4 front-panel assembly | 11 | 38. | The physical port numbers for the SAN Volume Controller 2145-8G4 | 30 |
| 8. | SAN Volume Controller 2145-CF8 operator-information panel | 13 | 39. | The physical port numbers for the SAN Volume Controller 2145-8F4 | 31 |
| 9. | SAN Volume Controller 2145-8A4 operator-information panel | 14 | 40. | The physical port numbers for the SAN Volume Controller 2145-8F2 | 31 |
| 10. | SAN Volume Controller 2145-8G4 operator-information panel | 14 | 41. | Photo of the redundant ac-power switch | 40 |
| 11. | SAN Volume Controller 2145-8F2 and SAN Volume Controller 2145-8F4 operator-information panel | 15 | 42. | A four-node SAN Volume Controller cluster with the redundant ac-power switch feature | 42 |
| 12. | SAN Volume Controller 2145-CF8 rear-panel indicators | 17 | 43. | 2145 UPS-1U front-panel assembly | 44 |
| 13. | Connectors on the rear of the SAN Volume Controller 2145-CF8 | 17 | 44. | 2145 UPS-1U connectors and switches. | 46 |
| 14. | Power connector | 18 | 45. | 2145 UPS-1U dip switches | 46 |
| 15. | Service ports of the SAN Volume Controller 2145-CF8 | 18 | 46. | Ports not used by the 2145 UPS-1U | 47 |
| 16. | SAN Volume Controller 2145-CF8 port not used | 19 | 47. | Power connector | 47 |
| 17. | SAN Volume Controller 2145-8A4 rear-panel indicators | 19 | 48. | SAN Volume Controller front-panel assembly | 81 |
| 18. | SAN Volume Controller 2145-8A4 external connectors | 19 | 49. | Example of a boot progress display | 81 |
| 19. | Power connector | 20 | 50. | Example of a cluster error code | 82 |
| 20. | Service ports of the SAN Volume Controller 2145-8A4 | 20 | 51. | Example of a node error code | 82 |
| 21. | SAN Volume Controller 2145-8G4 rear-panel indicators | 21 | 52. | Node rescue display | 83 |
| 22. | SAN Volume Controller 2145-8G4 external connectors | 21 | 53. | Validate WWNN? navigation. | 85 |
| 23. | Power connector | 22 | 54. | SAN Volume Controller options on the front-panel display | 86 |
| 24. | Service ports of the SAN Volume Controller 2145-8G4 | 22 | 55. | Viewing the IPv6 address on the front-panel display | 88 |
| 25. | SAN Volume Controller 2145-8F4 rear-panel indicators | 23 | 56. | Upper options of the actions menu on the front panel. | 92 |
| 26. | SAN Volume Controller 2145-8F4 external connectors | 23 | 57. | Middle options of the actions menu on the front panel. | 93 |
| 27. | Power connector | 24 | 58. | Lower options of the actions menu on the front panel. | 94 |
| 28. | Service ports of the SAN Volume Controller 2145-8F4 | 24 | 59. | Language? navigation. | 103 |
| 29. | Ports not used during normal operation by the SAN Volume Controller 2145-8F4 | 24 | 60. | Example of a boot error code | 131 |
| 30. | Ports not used on the front panel of the SAN Volume Controller 2145-8F4 | 25 | 61. | Example of a boot progress display | 132 |
| 31. | SAN Volume Controller 2145-8F2 rear-panel indicators | 25 | 62. | Example of a displayed node error code | 132 |
| | | | 63. | Example of a node-rescue error code | 133 |
| | | | 64. | Example of a create-cluster error code | 133 |
| | | | 65. | Example of a cluster-recovery error code | 133 |
| | | | 66. | Example of a cluster error code | 134 |
| | | | 67. | Node rescue display | 200 |
| | | | 68. | SAN Volume Controller service controller error light. | 203 |
| | | | 69. | Error LED on the SAN Volume Controller models 2145-CF8, 2145-8A4, 2145-8G4, and 2145-8F4 or 2145-8F2 operator-information panel | 203 |

| | | | | | |
|-----|---|-----|-----|--|-----|
| 70. | Hardware boot display | 204 | 84. | SAN Volume Controller 2145-CF8 light path diagnostics panel | 241 |
| 71. | Power LED on the SAN Volume Controller models 2145-CF8, 2145-8G4, and 2145-8F4 or 2145-8F2 operator-information panel | 208 | 85. | SAN Volume Controller 2145-CF8 system board LEDs diagnostics panel | 242 |
| 72. | Power LED on the SAN Volume Controller models 2145-8G4, 2145-8F4, and 2145-8F2 rear panel | 209 | 86. | SAN Volume Controller 2145-8A4 operator-information panel | 246 |
| 73. | Power LED indicator on the rear panel of the SAN Volume Controller 2145-CF8. | 210 | 87. | SAN Volume Controller 2145-8A4 system board LEDs | 247 |
| 74. | SAN Volume Controller models 2145-8G4 and 2145-8F4 or 2145-8F2 ac and dc LED indicators on the rear panel. | 211 | 88. | SAN Volume Controller 2145-8G4 operator-information panel | 248 |
| 75. | Power LED indicator and ac and dc indicators on the rear panel of the SAN Volume Controller 2145-CF8 | 211 | 89. | SAN Volume Controller 2145-8G4 light path diagnostics panel | 249 |
| 76. | Power LED on the SAN Volume Controller 2145-8A4 operator-information panel. | 213 | 90. | SAN Volume Controller 2145-8G4 system board LEDs | 250 |
| 77. | SAN Volume Controller 2145-8A4 system board LEDs | 215 | 91. | SAN Volume Controller 2145-8F4 operator-information panel | 252 |
| 78. | 2145 UPS-1U front-panel assembly | 217 | 92. | SAN Volume Controller 2145-8F2 and SAN Volume Controller 2145-8F4 light path diagnostics panel | 253 |
| 79. | SAN Volume Controller models 2145-CF8, 2145-8A4, 2145-8G4, and 2145-8F4 or 2145-8F2 power control button | 228 | 93. | SAN Volume Controller 2145-8F2 and SAN Volume Controller 2145-8F4 system board LEDs | 254 |
| 80. | SAN Volume Controller service controller error light. | 229 | 94. | Hardware boot display | 256 |
| 81. | Front-panel display when push buttons are pressed | 230 | 95. | Node rescue display | 257 |
| 82. | Port 2 Ethernet link LED on the SAN Volume Controller rear panel | 232 | 96. | Keyboard and monitor ports on the SAN Volume Controller models 2145-8G4, 2145-8A4, 2145-8F4 and 2145-8F2 | 258 |
| 83. | SAN Volume Controller 2145-CF8 operator-information panel | 241 | 97. | Keyboard and monitor ports on the SAN Volume Controller 2145-CF8 | 258 |

Tables

| | | | | | |
|-----|---|-----|-----|---|-----|
| 1. | Terminology mapping table | x | 25. | Fields that are provided for the uninterruptible power supply assembly that is powering the node | 77 |
| 2. | Other IBM publications | xii | 26. | Fields that are provided for the SAS host bus adapter (HBA) | 77 |
| 3. | IBM documentation and related websites | xii | 27. | Fields that are provided for the small form factor pluggable (SFP) transceiver | 77 |
| 4. | SAN Volume Controller communications types | 3 | 28. | Fields that are provided for the cluster properties | 78 |
| 5. | Link state and activity for the bottom fibre-channel LED | 26 | 29. | Description of data fields for the event log | 109 |
| 6. | Link speed for the top fibre-channel LED | 26 | 30. | Notification types | 110 |
| 7. | Actual link speeds | 27 | 31. | SAN Volume Controller notification types and corresponding syslog level codes | 111 |
| 8. | Actual link speeds | 27 | 32. | SAN Volume Controller values of user-defined message origin identifiers and syslog facility codes | 111 |
| 9. | SAN Volume Controller 2145-CF8 FRU descriptions | 49 | 33. | Informational event codes | 113 |
| 10. | 2145 UPS-1U FRU descriptions | 50 | 34. | Configuration event codes | 116 |
| 11. | SAN Volume Controller 2145-8A4 FRU descriptions | 50 | 35. | SCSI status | 120 |
| 12. | SAN Volume Controller 2145-8G4 FRU descriptions | 51 | 36. | SCSI sense keys, codes, and qualifiers | 120 |
| 13. | SAN Volume Controller 2145-8F4 FRU descriptions | 52 | 37. | Reason codes | 121 |
| 14. | SAN Volume Controller 2145-8F2 FRU descriptions | 53 | 38. | Object types | 122 |
| 15. | Fields for the system board | 74 | 39. | Event IDs and codes | 123 |
| 16. | Fields for the processors | 74 | 40. | Message classification number range | 134 |
| 17. | Fields for the fans | 74 | 41. | 2145 UPS-1U error indicators | 217 |
| 18. | Fields that are repeated for each installed memory module | 75 | 42. | SAN Volume Controller fibre-channel adapter assemblies | 238 |
| 19. | Fields that are repeated for each adapter card that is installed | 75 | 43. | SAN Volume Controller fibre-channel adapter connection hardware | 238 |
| 20. | Fields that are repeated for each SCSI, IDE, SATA, and SAS device that is installed | 75 | 44. | Diagnostics panel LED prescribed actions | 243 |
| 21. | Fields that are specific to the node software | 76 | 45. | SAN Volume Controller 2145-8A4 diagnostics panel LED prescribed actions | 248 |
| 22. | Fields that are provided for the front panel assembly | 76 | 46. | Diagnostics panel LED prescribed actions | 251 |
| 23. | Fields that are provided for the Ethernet port | 76 | 47. | Diagnostics panel LED prescribed actions | 254 |
| 24. | Fields that are provided for the power supplies in the node | 76 | | | |

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 or with the management GUI.

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

The maintenance analysis procedures (MAPs) can help you analyze failures that occur in a SAN Volume Controller. With the MAPs, you can isolate the field-replaceable units (FRUs) of the SAN Volume Controller that fail. Begin all problem determination and repair procedures from “MAP 5000: Start” on page 201.

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, and the uninterruptible power supply.

Summary of changes for GC27-2284-00 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, GG27-2227-02. The following sections summarize the changes that have since been implemented from the previous version.

This version includes the following new information:

- New error codes
- New event codes
- New information codes
- Management GUI replaces the SAN Volume Controller Console. The management GUI provides extensive facilities to help troubleshoot and correct problems. Use it to also manage and service your system.
- The service assistant interface and **satask** and **sainfo** commands provide methods for servicing your system.
- USB key commands

To coincide with new and existing IBM products and functions, several common terms have changed and are incorporated in the SAN Volume Controller information. Certain SAN Volume Controller information, particularly command-line interface (CLI) documentation, remains primarily unchanged.

The following table shows the current and previous usage of the changed common terms.

Table 1. Terminology mapping table

| 6.1.0 SAN Volume Controller term | Previous SAN Volume Controller term | Description |
|---|-------------------------------------|--|
| event | error | An occurrence of significance to a task or system. Events can include completion or failure of an operation, a user action, or the change in state of a process. |
| host mapping | VDisk-to-host mapping | The process of controlling which hosts have access to specific volumes within a cluster. |
| storage pool | managed disk (MDisk) group | A collection of storage capacity that provides the capacity requirements for a volume. |
| thin provisioning (or thin-provisioned) | space-efficient | The ability to define a storage unit (full system, storage pool, volume) with a logical capacity size that is larger than the physical capacity assigned to that storage unit. |
| volume | virtual disk (VDisk) | A discrete unit of storage on disk, tape, or other data recording medium that supports some form of identifier and parameter list, such as a volume label or input/output control. |

Changed information

This version includes the following changed information:

- The front panel menu has been restructured into two paths. You can display objects and take action on objects.

Removed information

This version has had the following information removed:

- Service mode topics. The front panel menu has been restructured into two paths. You can access all of the service actions through the CLI interface for the service assistant commands, which include **satask** and **sainfo** or through the service assistant interface.

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:

- SAN Volume Controller 2145-4F2. With the exception of the message content, references to the SAN Volume Controller 2145-4F2 have been removed. The SAN Volume Controller 2145-4F2 node is supported only in SAN Volume Controller 4.3.1 and earlier releases.

For information about replacing a SAN Volume Controller 2145-4F2 node with a SAN Volume Controller 2145-CF8 node, see *Procedures for Replacing SAN Volume Controller 2145-4F2 Nodes with SAN Volume Controller 2145-CF8 Nodes* on the following website:

Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145

The link to the procedures is located under **Related Reading** on the page that contains the V5.1.x *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*.

- 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:

| | |
|-----------------|---|
| Boldface | Text in boldface 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 websites 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 website:

publib.boulder.ibm.com/infocenter/svc/ic/index.jsp

SAN Volume Controller library

Unless otherwise noted, the publications in the SAN Volume Controller library are available in Adobe portable document format (PDF) from the following website:

Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145

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 |

IBM documentation and related websites

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

Table 3. IBM documentation and related websites

| Website | Address |
|--|--|
| Support for SAN Volume Controller (2145) | Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145 |
| Support for IBM System Storage and IBM TotalStorage products | www.ibm.com/storage/support/ |
| IBM Publications Center | www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss |
| IBM Redbooks® publications | www.redbooks.ibm.com/ |

Related accessibility information

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

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 website:

www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss

Sending your comments

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

To submit any comments about this book or any other SAN Volume Controller documentation:

- Go to the feedback page on the website for the SAN Volume Controller Information Center at publib.boulder.ibm.com/infocenter/svc/ic/index.jsp?topic=/com.ibm.storage.svc.console.doc/feedback.htm. There you can use the feedback page to enter and submit comments or browse to the topic and use the feedback link in the running footer of that page to identify the topic for which you have a comment.
- Send your comments by email to starpubs@us.ibm.com. Include the following information for this publication or use suitable replacements for the publication title and form number for the publication on which you are commenting:
 - Publication title: *IBM System Storage SAN Volume Controller 2145 Troubleshooting Guide*
 - Publication form number: GC27-2284-00
 - Page, table, or illustration numbers that you are commenting on
 - A detailed description of any information that should be changed

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 volumes 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) function, including thin-provisioned FlashCopy to make multiple targets affordable
 - Metro Mirror (synchronous copy)
 - Global Mirror (asynchronous copy)
 - Data migration
- Space management
 - IBM System Storage Easy Tier™ to migrate the most frequently used data to higher performing storage
 - Metering of service quality when combined with IBM Tivoli® Storage Productivity Center
 - Thin-provisioned logical volumes

| Figure 1 on page 2 shows hosts, SAN Volume Controller nodes, and RAID storage systems connected to a
| SAN fabric. The redundant SAN fabric comprises a fault-tolerant arrangement of two or more
| counterpart SANs that provide alternate paths for each SAN-attached device.
|

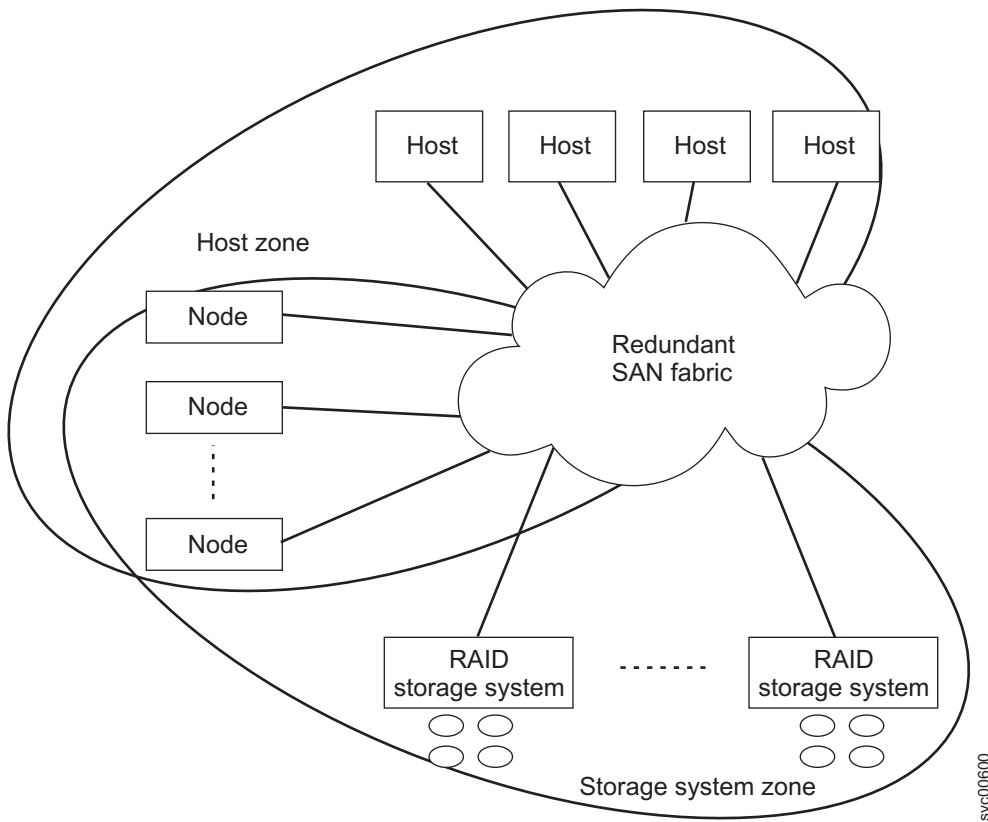


Figure 1. SAN Volume Controller system in a fabric

Volumes

A cluster of SAN Volume Controller nodes present volumes to the hosts. Most of the advanced functions that SAN Volume Controller provides are defined on volumes. These volumes are created from managed disks (MDisks) that are presented by the RAID storage systems. All data transfer occurs through the SAN Volume Controller nodes, which is described as symmetric virtualization.

Figure 2 shows the data flow across the fabric.

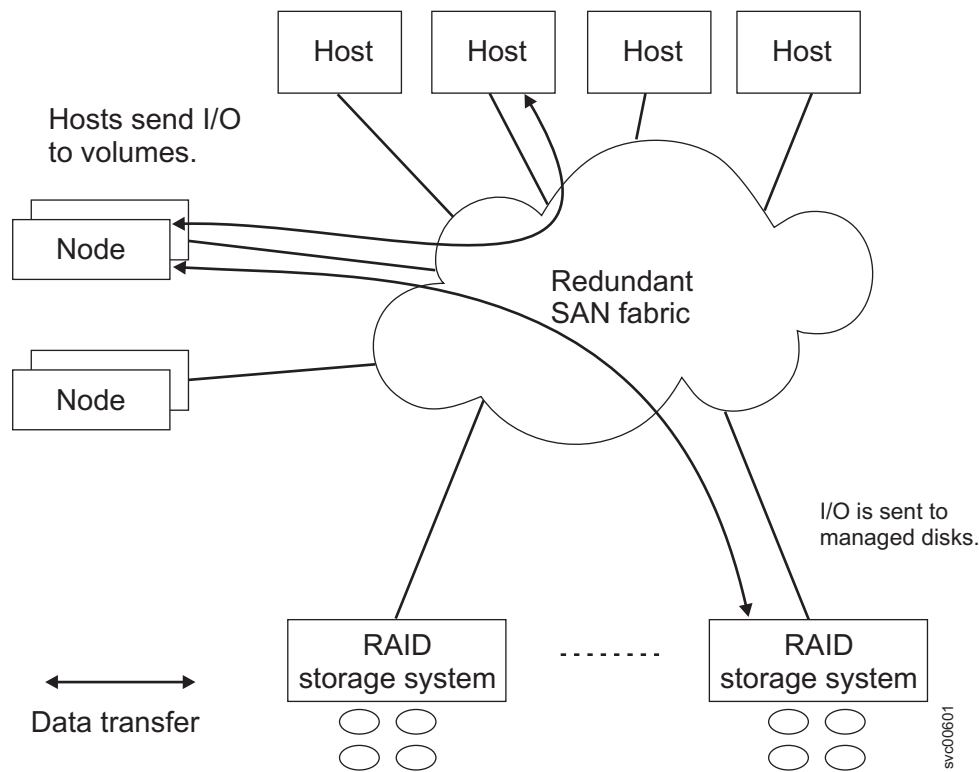


Figure 2. Data flow in a SAN Volume Controller system

The nodes in a cluster are arranged into pairs known as *I/O groups*. A single pair is responsible for serving I/O on a given volume. Because a volume is served by two nodes, there is no loss of availability if one node fails or is taken offline.

System management

The SAN Volume Controller nodes in a cluster operate as a single system and present a single point of control for system management and service. System management and error reporting are provided through an Ethernet interface to one of the nodes in the cluster, which is called the *configuration node*. The configuration node runs a web server and provides a command-line interface (CLI). The configuration node is a role that any node can take. If the current configuration node fails, a new configuration node is selected from the remaining nodes. Each node also provides a command-line interface and web interface for performing hardware service actions.

Fabric types

I/O operations between hosts and SAN Volume Controller nodes and between SAN Volume Controller nodes and RAID storage systems are performed by using the SCSI standard. The SAN Volume Controller nodes communicate with each other by using private SCSI commands.

SAN Volume Controller uses the SCSI commands over the fibre-channel SAN and 1 Gbps Ethernet. Table 4 shows the fabric type that can be used for communicating between hosts, nodes, and RAID storage systems. These fabric types can be used at the same time.

Table 4. SAN Volume Controller communications types

| Communications type | Host to SAN Volume Controller | SAN Volume Controller to storage system | SAN Volume Controller to SAN Volume Controller |
|---------------------|-------------------------------|---|--|
| Fibre Channel SAN | Yes | Yes | Yes |

Table 4. SAN Volume Controller communications types (continued)

| Communications type | Host to SAN Volume Controller | SAN Volume Controller to storage system | SAN Volume Controller to SAN Volume Controller |
|-------------------------|-------------------------------|---|--|
| iSCSI (1 Gbps Ethernet) | Yes | No | No |

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.

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 (volumes) that are used by applications on the hosts. A node is in only one I/O group and provides access to the volumes in that I/O group.

Clusters

All your configuration, monitoring, and service tasks are performed at the cluster (system) 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. Each interface accesses the cluster remotely through the Ethernet cluster-management address.

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 management GUI.

Figure 3 shows an example cluster that contains four nodes. Node 1 has been designated the configuration node. User requests (1) are handled by node 1.




Figure 3. 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 management GUI application or the CLI. This node is known as the *configuration node*.

If the configuration node is stopped or fails, the remaining nodes in the 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 subnet. 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 management IP addresses.

Routing considerations for event notification

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

- Email
- 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 external storage system. If you are using the Metro Mirror and Global Mirror feature, create a zone with at least one port from each node in each cluster; up to four clusters are supported.

Note: Some operating systems cannot tolerate other operating systems in the same host zone, although you might have more than one host type in the SAN fabric. For example, you can have a SAN that contains one host that runs on an IBM AIX® operating system and another host that runs on a Microsoft Windows operating system.

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

Chapter 2. Introducing the SAN Volume Controller hardware components

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

SAN Volume Controller nodes

SAN Volume Controller supports several different node types.

The following nodes are supported:

- The 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.
- The SAN Volume Controller 2145-8F4 node is no longer available for purchase, but remains supported.
- The SAN Volume Controller 2145-8F2 node is no longer available for purchase, but remains supported.

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




Figure 4. 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 5 shows the controls and indicators on the front panel of the SAN Volume Controller 2145-8A4.




Figure 5. 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 6 shows the controls and indicators on the front panel of the SAN Volume Controller 2145-8G4.




Figure 6. 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 7 shows the controls and indicators on the front panel of the SAN Volume Controller 2145-8F4 and SAN Volume Controller 2145-8F2.




Figure 7. 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?

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




Figure 8. 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 9 shows the operator-information panel for the SAN Volume Controller 2145-8A4.




Figure 9. 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 10 shows the operator information panel for the SAN Volume Controller 2145-8G4.




Figure 10. 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 11 shows the operator-information panel that is used by the SAN Volume Controller 2145-8F4 and the SAN Volume Controller 2145-8F2 models.




Figure 11. 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.

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.

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

Locator LED

The SAN Volume Controller does not use the locator LED.

Ethernet activity LED

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




Figure 12. 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 13 shows the external connectors on the SAN Volume Controller 2145-CF8 back panel assembly.




Figure 13. 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 14 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 14. 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 15 shows ports that are used only during service procedures.




Figure 15. 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 16 on page 19 shows the one port that is not used during service procedures or normal use.




Figure 16. 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.

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 17 shows the rear-panel indicators on the SAN Volume Controller 2145-8A4 back-panel assembly.




Figure 17. 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 18 shows the external connectors on the SAN Volume Controller 2145-8A4 back-panel assembly.




Figure 18. 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 19 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 19. 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 20.




Figure 20. 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 21 on page 21 shows the rear-panel indicators on the SAN Volume Controller 2145-8G4 back-panel assembly.




Figure 21. 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 22 shows the external connectors on the SAN Volume Controller 2145-8G4 back panel assembly.




Figure 22. 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 23 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 2145-8G4 to the power source from the uninterruptible power supply.




Figure 23. 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 24.




Figure 24. 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 25 on page 23 shows the rear-panel indicators on the SAN Volume Controller 2145-8F4 back-panel assembly.




Figure 25. 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 26 shows the external connectors on the SAN Volume Controller 2145-8F4 back panel assembly.




Figure 26. 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 27 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 2145-8F4 to the power source from the

uninterruptible power supply.




Figure 27. 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 28 provides the locations of the service ports.




Figure 28. 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 29 and Figure 30 on page 25 show the ports that are not used by the SAN Volume Controller.




Figure 29. 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 30. 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 31 shows the rear-panel indicators on the SAN Volume Controller 2145-8F2 back-panel assembly.




Figure 31. 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 32 shows the external connectors on the SAN Volume Controller 2145-8F2 back panel assembly.




Figure 32. 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 33 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 33. 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 5. Link state and activity for the bottom fibre-channel LED

| LED state | Link state and activity indicated |
|-----------|-----------------------------------|
| 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 6. Link speed for the top fibre-channel LED

| LED state | Link speed indicated |
|-----------|----------------------|
| OFF | SLOW |
| ON | FAST |
| BLINKING | MEDIUM |

Table 7 on page 27 shows the actual link speeds for the SAN Volume Controller models 2145-8A4, 2145-8G4, and 2145-8F4.

Table 7. Actual link speeds

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

Table 8 shows the actual link speeds for the SAN Volume Controller 2145-CF8.

Table 8. Actual link speeds

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

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 34 shows the location of the SAN Volume Controller 2145-CF8 ac, dc, and power-supply error LEDs.




Figure 34. 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 35 on page 29 shows the location of the ac and dc LEDs.




Figure 35. 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 36 shows the location of the ac and dc LEDs.




Figure 36. 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 ports are 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.

The WWPNS are in the form 50050768010QXXXXX, 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 |

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




Figure 37. 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 38 provides a view of the rear of the SAN Volume Controller 2145-8G4.




Figure 38. 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 39 on page 31 provides a view of the rear of the SAN Volume Controller 2145-8F4.




Figure 39. 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 40 provides a view of the rear of the SAN Volume Controller 2145-8F2.




Figure 40. 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.

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-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) |
| 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 |
| 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 41 shows a redundant ac-power switch.



Figure 41. 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 42 on page 42 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




Figure 42. 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.

2145 UPS-1U controls and indicators

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




Figure 43. 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

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.

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.

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




Figure 44. 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 45 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 45. 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 46 shows the 2145 UPS-1U ports that are not used.




Figure 46. 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 47 shows the power connector for the 2145 UPS-1U.




Figure 47. 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) |

| Height | Width | Depth | Maximum weight |
|---|-------|-------|----------------|
| Note: 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) |

2145 UPS environment

Ensure that your physical site meets the installation requirements of the 2145 UPS.

Uninterruptible power supply specifications

The following tables describe the physical characteristics of the 2145 UPS.

Attention: Ensure that you comply with the following requirements for the 2145 UPS units:

- 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 must also have input voltage capture that has a slew rate of no more than 3 Hz per second.

2145 UPS dimensions and weight

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

| Height | Width | Depth | Maximum weight |
|-----------------|-----------------|-------------------|----------------|
| 89 mm (3.5 in.) | 483 mm (19 in.) | 622 mm (24.5 in.) | 37 kg (84 lb) |

Heat output

The 2145 UPS unit produces the following approximate heat output.

| Model | Heat output during normal operation | Heat output during battery operation |
|--|-------------------------------------|--------------------------------------|
| 2145 UPS powering two SAN Volume Controller 2145-4F2 nodes | 140 W (480 Btu per hour) | 250 W (850 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. |
| Fibre-channel small form-factor pluggable (SFP) transceiver | 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. Note: It is possible that small form-factor pluggable (SFP) transceivers 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 SFP transceivers. 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 SFP transceivers, 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 to which the service controller cable connects. |
| 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. |
| 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. |

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

| FRU | Description |
|----------------|--|
| Alcohol wipe | A cleaning wipe. |
| Thermal grease | Grease that is used to provide a thermal seal between a processor and a heat sink. |

2145 UPS-1U FRUs

Table 10 provides a brief description of each 2145 UPS-1U FRU.

Table 10. 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 11 provides a brief description of each SAN Volume Controller 2145-8A4 FRU.

Table 11. 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) transceiver | 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. |

Table 11. SAN Volume Controller 2145-8A4 FRU descriptions (continued)

| FRU | Description |
|-------------------------------------|---|
| 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 12 provides a brief description of each SAN Volume Controller 2145-8G4 FRU.

Table 12. 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) transceiver | 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. |

Table 12. SAN Volume Controller 2145-8G4 FRU descriptions (continued)

| FRU | Description |
|----------------------------|--|
| 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 13 provides a brief description of each SAN Volume Controller 2145-8F4 FRU.

Table 13. 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) transceiver | 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. |

Table 13. SAN Volume Controller 2145-8F4 FRU descriptions (continued)

| FRU | Description |
|----------------------------|--|
| 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. |
| 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 14 provides a brief description of each SAN Volume Controller 2145-8F2 FRU.

Table 14. 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) transceiver | 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 DDR2 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 |

Table 14. SAN Volume Controller 2145-8F2 FRU descriptions (continued)

| FRU | Description |
|----------------------------|--|
| 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 |
| 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. |

Chapter 3. SAN Volume Controller user interfaces for servicing your system

SAN Volume Controller provides a number of user interfaces to troubleshoot, recover, or maintain your system. The interfaces provide various sets of facilities to resolve any situation that you encounter.

Use the management GUI to monitor and maintain the configuration of storage that is associated with your clusters. Perform service procedures from the service assistant. Use the command-line interface (CLI) to manage your cluster.

Management GUI interface

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

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

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

When to use the management GUI

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

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

Use the views that are available in the management GUI to verify the status of the system, the hardware devices, the physical storage, and the available volumes. The **Troubleshooting** panel provides access to problems that must be fixed and to maintenance procedures that step you through the process of correcting the problem.

The **Recommended Actions** panel provides an analysis of the event log and identifies the most important alert to resolve. If there is a service error code for the alert, you can run a fix procedure that assists you in resolving the problem. These fix procedures analyze the system and provide more information about the problem. They suggest actions to take and step you through the actions that automatically manage the system where necessary. Finally, they check that the problem is resolved.

If there is an error that is reported, always use the fix procedures within the management GUI to resolve the problem. Always use the fix procedures for both software configuration problems and hardware failures. The fix procedures analyze the system to ensure that the required changes do not cause volumes to be inaccessible to the hosts. The fix procedures automatically perform configuration changes that are required to return the system to its optimum state.

Accessing the management GUI

This procedure describes how to access the management GUI.

- | You must use a supported browser. Verify that you are using a supported browser from the following website:

Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145

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

1. Start a supported web browser and point the browser to the management IP address of your system. The management IP address is set when the system cluster is created. Up to four addresses can be configured for your use. There are two addresses for IPv4 access and two addresses for IPv6 access.
2. When the connection is successful, you see a login panel.
3. Log on by using your user name and password.
4. When you have logged on, select **Troubleshooting** and then **Recommended Actions**.
5. Select the recommended alert and run the fix procedure.
6. Continue to work through the alerts in the order suggested, if possible.

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

Deleting a node from a cluster using the management GUI

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 volumes within the I/O group are degraded), the system ensures that data loss does not occur as a result of deleting the only node with the cache data. If a failure occurs on the other node in the I/O group, the cache is flushed before the node is removed to prevent data loss.

Before 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. If the node is re-added to the cluster at a later time, recording this node information can avoid data corruption.

Attention:

- If you are removing a single node and the remaining node in the I/O group is online, the data on the remaining node goes into write-through mode. This data can be exposed to a single point of failure if the remaining node fails.
- If the volumes are already degraded before you remove a node, redundancy to the volumes is degraded. Removing a node might result in a loss of access to data and data loss.
- Removing the last node in the cluster destroys the cluster. Before you remove the last node in the cluster, ensure that you want to destroy the cluster.
- When you remove a node, you remove all redundancy from the I/O group. As a result, new or existing failures can cause I/O errors on the hosts. The following failures can occur:
 - Host configuration errors
 - Zoning errors
 - Multipathing-software configuration errors
- If you are deleting the last node in an I/O group and there are volumes that are assigned to the I/O group, you cannot remove the node from the cluster if the node is online. You must back up or migrate all data that you want to save before you remove the node. If the node is offline, you can remove the node.
- When you remove the configuration node, the configuration function moves to a different node within the cluster. This process can take a short time, typically less than a minute. The management GUI reattaches to the new configuration node transparently.
- If you turn 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. 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.

This task assumes that you have already accessed the management GUI.

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

1. From Home, select **System Status**.
2. Find the node that you want to remove.

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

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

If any volumes are listed, determine why and if access to the volumes is required while the node is removed from the cluster. If the volumes are assigned from MDisk groups that contain solid-state drives (SSDs) that are located in the node, check why the volume mirror, if it is configured, is not synchronized. There can also be dependent volumes because the partner node in the I/O group is offline. Fabric issues can also prevent the volume from communicating with the storage systems. Resolve these problems before continuing with the node removal.
4. Click **Remove Node**.
5. Click **OK** to remove the node. Before a node is removed, the SAN Volume Controller checks to determine if there are any volumes that depend on that node. If the node that you selected contains volumes within the following situations, the volumes go offline and become unavailable if the node is removed:
 - The node contains solid-state drives (SSD) and also contains the only synchronized copy of a mirrored volume
 - The other node in the I/O group is offline

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

Adding nodes to a cluster

This topic provides instructions for adding a node to a cluster. It also contains information about adding a node to a cluster if the node previously failed and is being replaced with a new node or if a repair action has caused the node to be unrecognizable by the 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.

Considerations 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 action might cause the hosts to access the wrong volumes.

- If the new node requires a level of software that is higher than the software level that is available on the cluster, the entire cluster must be upgraded before the new node can be added to the cluster.
- 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.
- 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 an I/O group in the cluster and are adding a new node, then there are no special procedures because this node was never added to a cluster and its WWNN did not exist.
- If you are creating an I/O group in the cluster and are adding a new node, but this node has been added to a cluster before, then the host system might still be configured to the node WWPNs and the node might still be zoned in the fabric. Because you cannot change the WWNN for the node, you must ensure that other components in your fabric are configured correctly. Verify that any host that was previously configured to use the node has been correctly updated.
- If the node that you are adding was previously replaced, either for a node repair or upgrade, you might have used the WWNN of that node for the replacement node. Ensure that the WWNN of this node was updated so that you do not have two nodes with the same WWNN attached to your fabric. Also ensure that the WWNN of the node that you are adding is not 00000. If it is 00000, contact your support representative.

Considerations 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 volume. This association uses an identifier (UID) which is unique to the volume and is never reused. The UID allows multipathing device drivers to directly associate vpaths with volumes.
- Multipathing device drivers operate within a protocol stack that contains disk and fibre-channel device drivers that 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 volume with the vpath on every I/O operation that it performs.

Adding nodes to a cluster using the management GUI

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 action 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. 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.
3. 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.
4. 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 website for the latest supported software levels:
Support for SAN Volume Controller (2145) website at 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 accessed the management GUI.

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

1. From Home, select **System Status**.
2. From the rack image, click an empty slot that is associated with the I/O group that you want to add the node.
3. Select the candidate node that you want to add.
If the node that you want to add is unavailable in the candidate list, the node is in service state. Actions are required to release the node from service state before it can be added to the cluster.
4. Select **Add node**. You are shown a warning.
5. Click **OK**.
6. 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 must remove and re-add the node to the cluster.

If a node shows node error 578 or node error 690, the node is in service state. Perform the following steps from the front panel to exit service state:

1. Press and release the up or down button until Actions? is displayed.
2. Press the select button.
3. Press and release the up or down button until Exit Service? is displayed.
4. Press the select button.
5. Press and release the left or right button until Confirm Exit? is displayed.
6. Press the select button.

For any other node errors, follow the appropriate service procedures to fix the errors. After the errors are resolved and the node is in candidate state, you can try to add the node to the cluster again.

Service assistant interface

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

You connect to the service assistant through the service IP address.

When to use the service assistant

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

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

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

Use the service assistant when you cannot access the management GUI to run the recommended actions or because the recommended action directs you to use the service assistant.

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

The service assistant provides detailed status and error summaries. You can also perform the following service-related actions:

- Collect logs to create and download a package of files to send to support personnel.
- Remove the cluster data from a node.
- Recover a cluster if it fails.
- Install a software package from the support site or rescue the software from another node.
- Upgrade software on nodes manually versus performing a standard upgrade procedure.
- Change the service IP address that is assigned to Ethernet port 1 for the current node.
- Install a temporary SSH key if a key is not installed and CLI access is required.
- Restart the services used by the system.

Accessing the service assistant

The service assistant is a web application that helps troubleshoot and resolve problems on a node canister.

- | You must use a supported browser. Verify that you are using a supported and an appropriately configured browser from the following website:

Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145

To start the application, perform the following steps:

1. Start a supported web browser and point your web browser to `<serviceaddress>/service` for the node that you want to work on.
For example, if you set a service address of 11.22.33.44 for a node canister, point your browser to 11.22.33.44/service.
2. Log on to the service assistant using the superuser password.
If you do not know the current superuser password, reset the password.
- | 3. Ensure that you perform the service assistant actions on the correct node.

Adding a node to a cluster

This topic provides information about adding a node to a cluster.

A node that has been removed from a cluster is automatically re-added when the node is in candidate state. No separate add action is required.

If a node is in service state, follow the appropriate service procedures to resolve a node error 578 or a node error 690. For all other node errors, follow the appropriate service procedures. Complete the following steps to release a node from service state:

1. From the **Service Assistant** Home page, select the node that is in service state.
2. From the available actions, click **Exit service state**.
The node is added as soon as it is available.

To validate that the node was added successfully, perform the following steps:

1. From the management GUI Home page, select **Manage Device**.
2. Select the canister to which the node was added.
3. Verify that the node was added.

Cluster command-line interface

Use the cluster command-line interface (CLI) to manage a cluster using the `svctask` commands and `svcinfo` commands.

For a full description of the commands and how to start an SSH command-line session, see the "Command-line interface" topic in the SAN Volume Controller Information Center.

When to use the cluster CLI

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

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

You might also find it useful to create command scripts using the CLI commands to monitor for certain conditions or to automate configuration changes that you make on a regular basis.

Accessing the cluster CLI

Follow the steps that are described in the "Command-line interface" topic in the SAN Volume Controller Information Center to initialize and use a cluster CLI session.

Service command-line interface

Use the service command-line interface (CLI) to manage a node using the satask commands and svcinfo commands.

For a full description of the commands and how to start an SSH command-line session, see the "Command-line interface" topic in the SAN Volume Controller Information Center.

When to use the service CLI

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

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

Accessing the service CLI

Follow the steps that are described in the "Command-line interface" topic in the SAN Volume Controller Information Center to initialize and use a service CLI session.

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

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

Validating and repairing mirrored VDisk copies using the CLI

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

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

When you issue the **repairvdiskcopy** command, you must use only one of the **-validate**, **-medium**, or **-resync** parameters. You must also specify the name or ID of the 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:

```
svcinfolrepairvdiskcopyprogress -delim :
```

The following example shows how the command output is displayed:

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

Repairing a 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 (volume) offline during the repair, but does not prevent the disk from being moved between I/O groups.

If a repair operation completes successfully and the volume was previously offline because of corrupted metadata, the command brings the volume back online. The only limit on the number of concurrent repair operations is the number of 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 (thin-provisioned volume) 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 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 (volumes).

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

Data loss scenario 1

One node in an I/O group 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 IBM FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the offline volumes.
3. Run the **recovervdisk**, **recovervdiskbyiogrp** or **recovervdiskbycluster** command.
4. Re-create all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the volumes.

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 volume 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 volumes.
2. Run the **recovervdisk**, **recovervdiskbyiogrp** or **recovervdiskbycluster** command.
3. Recreate all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the volumes.

Replacing nodes nondisruptively

These procedures describe how to replace most SAN Volume Controller nodes nondisruptively.

The procedures do not include replacing a SAN Volume Controller 2145-4F2 node with a SAN Volume Controller 2145-CF8 node.

For information about replacing a SAN Volume Controller 2145-4F2 node with a SAN Volume Controller 2145-CF8 node, see *Procedures for Replacing SAN Volume Controller 2145-4F2 Nodes with SAN Volume Controller 2145-CF8 Nodes* on the following website:

Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145

The link to the procedures is located under **Related Reading** on the page that contains the V5.1.x *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*.

| These procedures are nondisruptive because changes to your SAN environment are not required. The
| replacement (new) node uses the same worldwide node name (WWNN) as the node that you are
| replacing. An alternative to this procedure is to replace nodes disruptively either by moving volumes to a
| new I/O group or by rezoning the SAN. The disruptive procedures, however, require additional work on
| the hosts.

This task assumes that the following conditions have been met:

- The existing cluster software must be at a version that supports the new node. If a node is being replaced by a SAN Volume Controller 2145-CF8 node, the cluster software version must be 5.1.0 or later. If a node is being replaced by a SAN Volume Controller 2145-8A4 node, the cluster software version must be 4.3.1 or later.
- All nodes that are configured in the cluster are present and online.
- | • All errors in the cluster event log are addressed and marked as fixed.
- | • There are no volumes, managed disks (MDisks), or external storage systems with a status of degraded
| or offline.
- The replacement node is not powered on.
- The replacement node is not connected to the SAN.
- You have a 2145 UPS-1U unit (feature code 8115) for each new SAN Volume Controller 2145-CF8 or 2145-8A4 node.
- You have backed up the cluster configuration and saved the `svc.config.backup.xml` file.
- The replacement node must be able to operate at the fibre-channel or Ethernet connection speed of the node it is replacing.
- | • If the node being replaced contains solid-state drives (SSDs), all SSDs and SAS adapters should be
| transferred to the new node if it supports the drives. If the new node does not support the existing
| SSDs, you must transfer the data off of the SSDs before replacing the node to avoid losing access to the
| data.

Important:

1. Do not continue this task if any of the conditions listed above are not met unless you are instructed to do so by the IBM Support Center.
2. Review all of the steps listed below before you perform this task.
3. Do not perform this task if you are not familiar with SAN Volume Controller environments or the procedures described in this task.
4. If you plan to reuse the node that you are replacing, ensure that the WWNN of the node is set to a unique number on your SAN. If you do not ensure that the WWNN is unique, the WWNN and WWPNN are duplicated in the SAN environment and can cause issues.

Tip: You can change the WWNN of the node you are replacing to the factory default WWNN of the replacement node to ensure that the number is unique.

5. The node ID and possibly the node name change during this task. After the cluster assigns the node ID, the ID cannot be changed. However, you can change the node name after this task is complete.

Perform the following steps to replace active nodes in a cluster:

1. (If the cluster software version is at 5.1 or later, complete this step.)

Confirm that no hosts have dependencies on the node.

When shutting down a node that is part of a cluster or when deleting the node from a cluster, you can use either the management GUI or a command-line interface (CLI) command. In the management GUI, click **System Status**. Click the correct node and click **Manage** to display all the volumes that are dependent on a node. You can also use the **node** parameter with the **svcinfoldependentvdisks** CLI command to view dependent volumes.

If dependent volumes exist, determine if the volumes are being used. If the volumes are being used, either restore the redundant configuration or suspend the host application. If a dependent quorum disk is reported, repair the access to the quorum disk or modify the quorum disk configuration.

2. Perform the following steps to determine the cluster configuration node, and the ID, name, I/O group ID, and I/O group name for the node that you want to replace. If you already know the physical location of the node that you want to replace, you can skip this step and proceed to step 3.

Tip: If one of the nodes that you want to replace is the cluster configuration node, replace it last.

- a. Issue the following command from the command-line interface (CLI):

```
svcinfolnode -delim :
```

The following is an example of the output that is displayed for this command:

```
id:name:UPS_serial_number:WWNN:status:IO_group_id:IO_group_name:
config_node:UPS_unique_id:hardware:iscsi_name:iscsi_alias
3:dvt113294:100089J137:5005076801005A07:online:0:io_grp0:yes:
20400002096810C7:8A4:iqn.1986-03.com.ibm:2145.1dcluster-80.dvt113294:
14:des113004:10006BR010:5005076801004F0F:online:0:io_grp0:no:
2040000192880040:8G4:iqn.1986-03.com.ibm:2145.1dcluster-80.des113004:
```

- b. In the `config_node` column, find the value `yes` and record the values in the `id` and `name` columns.
 - c. Record the values in the `id` and the `name` columns for each node in the cluster.
 - d. Record the values in the `IO_group_id` and the `IO_group_name` columns for each node in the cluster.
 - e. Issue the following command from the CLI for each node in the cluster to determine the front panel ID:

```
svcinfolnodevpd node_name or node_id
```

where *node_name* or *node_id* is the name or ID of the node for which you want to determine the front panel ID.
 - f. Record the value in the `front_panel_id` column. The front panel ID is displayed on the front of each node. You can use this ID to determine the physical location of the node that matches the node ID or node name that you want replace.
3. Perform the following steps to record the WWNN or iSCSI name of the node that you want to replace:
 - a. Issue the following command from the CLI:

```
svcinfolnode -delim : node_name or node_id
```

where *node_name* or *node_id* is the name or ID of the node for which you want to determine the WWNN or iSCSI name.
 - b. Record the WWNN or iSCSI name of the node that you want to replace. Also record the order of the fibre-channel and Ethernet ports.

4. Issue the following command from the CLI to power off the node:

```
svctask stopcluster -node node_name
```

Important:

- a. Record and mark the order of the fibre-channel or Ethernet cables with the node port number (port 1 to 4 for fibre-channel, or port 1 to 2 for Ethernet) before you remove the cables from the back of the node. The fibre-channel ports on the back of the node are numbered 1 to 4 from left to right. You must reconnect the cables in the exact order on the replacement node to avoid issues when the replacement node is added to the cluster. If the cables are not connected in the same order, the port IDs can change, which impacts the ability of the host to access volumes. See the hardware documentation specific to your model to determine how the ports are numbered.
 - b. Do not connect the replacement node to different ports on the switch or director. The SAN Volume Controller can have 4 or 8 Gbps HBAs; however, do not move them to faster switch or director ports at this time to avoid issues when the replacement node is added to the cluster.
 - c. Do not move the fibre-channel cables of the node to faster or different ports on the switch or director at this time. This is a separate task that must be planned independently of replacing nodes in a cluster.
5. Issue the following CLI command to delete this node from the cluster and I/O group:

```
svctask rmnode node_name or node_id
```

Where *node_name* or *node_id* is the name or ID of the node that you want to delete. You can use the CLI to verify that the deletion process has completed.
 6. Issue the following CLI command to ensure that the node is no longer a member of the cluster:

```
svcinfolsnode
```

A list of nodes is displayed. Wait until the removed node is not listed in the command output.

7. Perform the following steps to change the WWNN or iSCSI name of the node that you just deleted from the cluster to FFFFF:

For SAN Volume Controller V6.1.0 or later:

- a. With the Cluster panel displayed, press the up or down button until the Actions option is displayed.
- b. Press and release the select button.
- c. Press the up or down button until Change WWNN? is displayed.
- d. Press and release the select button to display the current WWNN.
- e. Press and release the select button to switch into edit mode. The Edit WWNN? panel is displayed.
- f. Change the WWNN to FFFFF.
- g. Press and release the select button to exit edit mode.
- h. Press the right button to confirm your selection. The Confirm WWNN? panel is displayed.
- i. Press and release the select button to confirm.

For SAN Volume Controller V4.3 through 5.1:

- a. Power on the node.
- b. 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.
- c. Change the displayed number to FFFFF. 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 save your changes and apply FFFFF as the new WWNN for the node.

8. Install the replacement node and the uninterruptible power supply in the rack and connect the uninterruptible power supply cables. See the *IBM System Storage SAN Volume Controller Model 2145-XXX Hardware Installation Guide* to determine how to connect the node and the uninterruptible power supply.

Important: Do not connect the fibre-channel or Ethernet cables during this step.

9. Power on the replacement node.
10. Record the WWNN or iSCSI name of the replacement node. You can use this name if you plan to reuse the node that you are replacing.
11. Perform the following steps to change the WWNN or iSCSI name of the replacement node to match the name that you recorded in step 3 on page 67:

For SAN Volume Controller V6.1.0 or later:

- a. With the Cluster panel displayed, press the up or down button until the Actions option is displayed.
- b. Press and release the select button.
- c. Press the up or down button until Change WWNN? is displayed.
- d. Press and release the select button to display the current WWNN.
- e. Press the select button to switch into edit mode. The Edit WWNN? panel is displayed.
- f. Change the WWNN to the numbers that you recorded in step 3 on page 67.
- g. Press and release the select button to exit edit mode.
- h. Press the right button to confirm your selection. The Confirm WWNN? panel is displayed.
- i. Press the select button to confirm.

For SAN Volume Controller V4.3 through 5.1:

- 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 on page 67. 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 apply the numbers as the new WWNN for the node.

Wait one minute. If Cluster: is displayed on the front panel, this indicates that the node is ready to be added to the cluster. If Cluster: is not displayed, see the troubleshooting information to determine how to address this problem or contact the IBM Support Center before you continue with the next step.

12. Connect the fibre-channel or Ethernet cables to the same port numbers that you recorded for the original node in step 4 on page 68.
13. Issue the following CLI command to verify that the last five characters of the WWNN are correct:
`svcinfo lsnodecandidate`

Important: If the WWNN is not what you recorded in step 3 on page 67, you must repeat step 11.

14. Issue the following CLI command to add the node to the cluster and ensure that the node has the same name as the original node and is in the same I/O group as the original node. See the **svctask addnode** CLI command documentation for more information.

```
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.

Important:

- a. Both nodes in the I/O group cache data; however, the cache sizes are asymmetric. The replacement node is limited by the cache size of the partner node in the I/O group. Therefore, it is possible that the replacement node does not utilize the full cache size until you replace the other node in the I/O group.
 - b. You do not have to reconfigure the host multipathing device drivers because the replacement node uses the same WWNN and WWPN as the previous node. The multipathing device drivers should detect the recovery of paths that are available to the replacement node.
 - c. The host multipathing device drivers take approximately 30 minutes to recover the paths. Do not upgrade the other node in the I/O group until for at least 30 minutes after you have successfully upgraded the first node in the I/O group. If you have other nodes in different I/O groups to upgrade, you can perform those upgrades while you wait.
15. See the documentation that is provided with your multipathing device driver for information on how to query paths to ensure that all paths have been recovered before proceeding to the next step. If you are using the IBM System Storage Multipath Subsystem Device Driver (SDD), the command to query paths is **datapath query device**.
16. Repair the faulty node.
- If you want to use the repaired node as a spare node, perform the following steps.
- For SAN Volume Controller V6.1.0 or later:
- a. With the Cluster panel displayed, press the up or down button until the **Actions** option is displayed.
 - b. Press and release the select button.
 - c. Press the up or down button until **Change WWNN?** is displayed.
 - d. Press and release the select button to display the current WWNN.
 - e. Press and release the select button to switch into edit mode. The **Edit WWNN?** panel is displayed.
 - f. Change the WWNN to 00000.
 - g. Press and release the select button to exit edit mode.
 - h. Press the right button to confirm your selection. The **Confirm WWNN?** panel is displayed.
 - i. Press and release the select button to confirm.
- This node can now be used as a spare node.
- For SAN Volume Controller V4.3 through V5.1:
- 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.
17. Repeat steps 3 on page 67 to 16 for each node that you want to replace.

Chapter 5. Viewing the vital product data

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

The VPD lists the following types of information:

- Cluster-related values such as the software version, space in storage pools, and space allocated to volumes.
- Node-related values that include the specific hardware that is installed in each node. Examples include the FRU part number for the system board and the level of BIOS firmware that is installed. The node VPD is held by the 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 management GUI.

Viewing the vital product data using the management GUI

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

Perform the following steps to view the vital product data for a node:

1. From **Home**, click **System Status**.
2. Select the node for which you want to display the details.
3. Click **VPD** to view the data.

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:

```
svcinfolsservicenodes
svcinfolnodevpd nodename
svcinfolcluster clustername
svcinfolclusterip
svcinfoldrive
```

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."

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:WWNN:status:IO_group_id:IO_group_name:config_node:UPS_unique_id:hardware:iscsi_name:iscsi_alias:
panel_name:enclosure_id:canister_id:enclosure_serial_number
1:node1:UPS_Fake_SN:50050768010050B1:online:0:io_grp0:yes:10000000000050B1:8G4:iqn.1986-03.com.ibm:2145.cluster0.node1:000368:::
```

2. Issue the **svcinfolnode** 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:

```
svcinfolnode -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
WWNN: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.ndihill.node2
iscsi_alias
failover_active:no
failover_name:node1
failover_iscsi_name:iqn.1986-03.com.ibm:2145.ndihill.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 **svcinfolcluster** command to display the properties for a cluster.

The following is an example of the command you can issue:

```
svcinfolcluster -delim : build1
```

where *build1* is the name of the cluster.


```

id:000002007A00A0FE
name:build1
location:local
partnership:
bandwidth:
total_mdisk_capacity:90.7GB
space_in_mdisk_grps:90.7GB
space_allocated_to_vdisks:14.99GB
total_free_space:75.7GB
statistics_status:on
statistics_frequency:15
required_memory:0
cluster_locale:en_US
time_zone:522 UTC
code_level:6.1.0.0 (build 47.3.1009031000)
FC_port_speed:2Gb
console_IP:9.71.46.186:443
id_alias:000002007A00A0FE
gm_link_tolerance:300
gm_inter_cluster_delay_simulation:0
gm_intra_cluster_delay_simulation:0
email_reply:
email_contact:
email_contact_primary:
email_contact_alterate:
email_contact_location:
email_state:stopped
inventory_mail_interval:0
total_vdiskcopy_capacity:15.71GB
total_used_capacity:13.78GB
total_overallocation:17
total_vdisk_capacity:11.72GB
cluster_ntp_IP_address:
cluster_isns_IP_address:
iscsi_auth_method:none
iscsi_chap_secret:
auth_service_configured:no
auth_service_enabled:no
auth_service_url:
auth_service_user_name:
auth_service_pwd_set:no
auth_service_cert_set:no
relationship_bandwidth_limit:25
gm_max_host_delay:5
tier:generic_ssd
tier_capacity:0.00MB
tier_free_capacity:0.00MB
tier:generic_hdd
tier_capacity:90.67GB
tier_free_capacity:75.34GB
email_contact2:
email_contact2_primary:
email_contact2_alterate:
total_allocated_extnt_capacity:16.12GB

```

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, and SAS host bus adapter (HBA).

Table 15 on page 74 shows the fields you see for the system board.

Table 15. 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 Note: The service controller is a device. |
| | Number of power supplies |
| | Number of high-speed SAS adapters |
| | BIOS manufacturer |
| | BIOS version |
| | BIOS release date |
| | System manufacturer |
| | System product |
| | Planar manufacturer |
| | Power supply part number |
| | CMOS battery part number |
| Power cable assembly part number | |
| Service processor firmware | |
| SAS controller part number | |

Table 16 shows the fields you see for each processor that is installed.

Table 16. Fields for the processors

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

Table 17 shows the fields that you see for each fan that is installed.

Table 17. Fields for the fans

| Item | Field name |
|------|-------------|
| Fan | Part number |
| | Location |

Table 18 shows the fields that are repeated for each installed memory module.

Table 18. Fields that are repeated for each installed memory module

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

Table 19 shows the fields that are repeated for each installed adapter card.

Table 19. Fields that are repeated for each adapter card that is installed

| Item | Field name |
|--------------|----------------------|
| Adapter card | Card type |
| | Part number |
| | Port numbers |
| | Location |
| | Device serial number |
| | Manufacturer |
| | Device |
| | Card revision |
| | Chip revision |

Table 20 shows the fields that are repeated for each device that is installed.

Table 20. Fields that are repeated for each SCSI, IDE, SATA, and SAS device that is installed

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

Table 21 on page 76 shows the fields that are specific to the node software.

Table 21. Fields that are specific to the node software

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

Table 22 shows the fields that are provided for the front panel assembly.

Table 22. Fields that are provided for the front panel assembly

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

Table 23 shows the fields that are provided for the Ethernet port.

Table 23. Fields that are provided for the Ethernet port

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

Table 24 shows the fields that are provided for the power supplies in the node.

Table 24. Fields that are provided for the power supplies in the node

| Item | Field name |
|----------------|-------------|
| Power supplies | Part number |
| | Location |

Table 25 on page 77 shows the fields that are provided for the uninterruptible power supply assembly that is powering the node.

Table 25. 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 26 shows the fields that are provided for the SAS host bus adapter (HBA).

Table 26. 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 27 shows the fields that are provided for the small form factor pluggable (SFP) transceiver.

Table 27. Fields that are provided for the small form factor pluggable (SFP) transceiver

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

Fields for the 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 28 shows the fields that are provided for the cluster properties as shown by the management GUI.

Table 28. Fields that are provided for the cluster properties

| Item | Field name |
|---------------------------|---|
| General | ID Note: This is the unique identifier for the cluster. |
| | Name |
| | Location |
| | Time Zone |
| | Required Memory |
| | Licensed Code Version |
| | Channel Port Speed Note: 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. |
| IP Addresses ¹ | Cluster Ethernet Port 1 (attributes for both IPv4 and IPv6) <ul style="list-style-type: none"> • IP Address • Service IP Address • Subnet Mask • Prefix • Default Gateway |
| | Cluster Ethernet Port 2 (attributes for both IPv4 and IPv6) <ul style="list-style-type: none"> • IP Address • Service IP Address • Subnet Mask • Prefix • Default Gateway |
| Remote Authentication | Remote Authentication |
| | Web Address |
| | User Name |
| | Password |
| | SSL Certificate |
| Space | Total MDisk Capacity |
| | Space in Storage Pools |
| | Space Allocated to Volumes |
| | Total Free Space |
| | Total Used Capacity |
| | Total Allocation |
| | Total Volume Copy Capacity |
| | Total Volume Capacity |

Table 28. Fields that are provided for the cluster properties (continued)

| Item | Field name |
|---|---|
| Statistics | Statistics Status |
| | Statistics Frequency |
| Metro and Global Mirror | Link Tolerance |
| | Inter-Cluster Delay Simulation |
| | Intra-Cluster Delay Simulation |
| | Partnership |
| | Bandwidth |
| Email | SMTP Email Server |
| | Email Server Port |
| | Reply Email Address |
| | Contact Person Name |
| | Primary Contact Phone Number |
| | Alternate Contact Phone Number |
| | Physical Location of the System Reporting Error |
| | Email Status |
| | Inventory Email Interval |
| iSCSI | iSNS Server Address |
| | Supported Authentication Methods |
| | CHAP Secret |
| ¹ You can also use the <code>svcinfo lsclusterip</code> CLI command to view this data. | |

Chapter 6. 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 48 shows where the front-panel display **1** is located on the SAN Volume Controller node.




Figure 48. SAN Volume Controller front-panel assembly

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 49. 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, boot code 120 is displayed.



See the "Error code reference" topic where you can find a description of the failure and the appropriate steps that you must perform to correct the failure.

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 50 and Figure 51 show how error codes are displayed on the front panel.



Figure 50. Example of a cluster error code



Figure 51. 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.

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.

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 52 on page 83, 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 52. 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 53) to choose which WWNN that you want to use.




Figure 53. 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 nor disk is suitable, you must wait until
- | the node restarts before you can change it. After the node restarts, select **Change WWNN** to change the
- | WWNN to the value that you want.

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 and operations that you use to service the node.

Figure 54 on page 86 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.

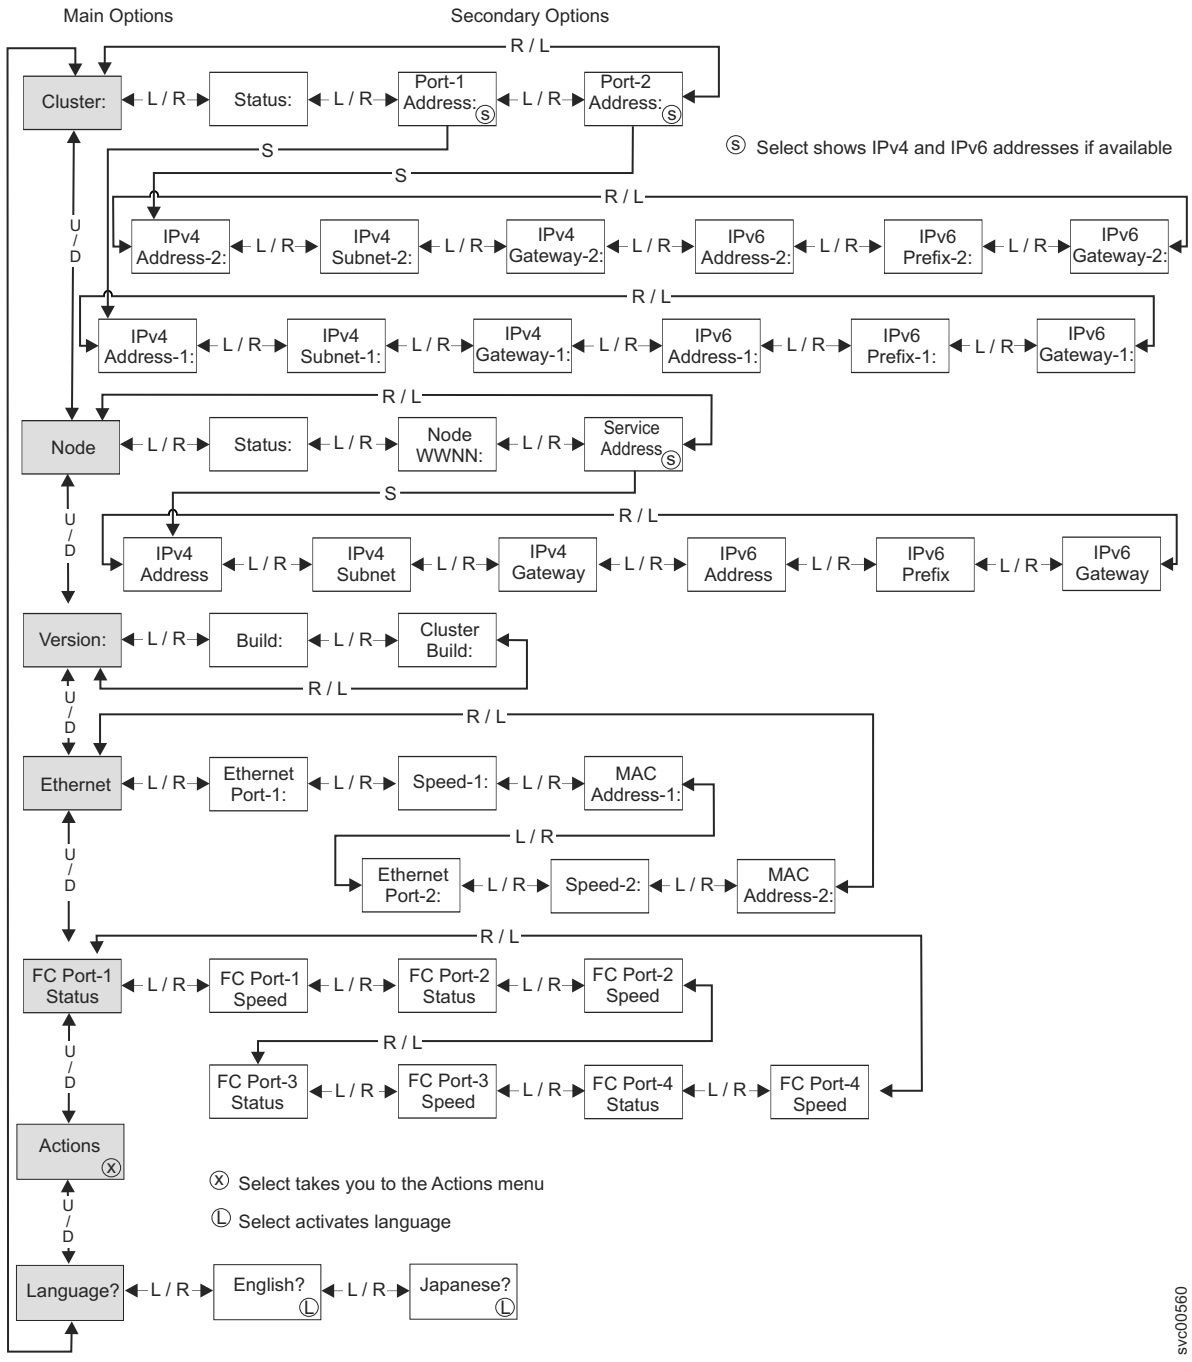


Figure 54. SAN Volume Controller options on the front-panel display

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 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 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
- | • Actions
- | • Language

Cluster options

The main cluster option from the menu can display the cluster name or the field 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 node is not assigned to a cluster, the field is blank.

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 node is a member of a cluster, the field indicates the operational status of the cluster, as follows:

Active

Indicates that this node is an active member of the cluster.

Inactive

Indicates that the node is a member of a cluster, but is not now operational. It is not operational because the other nodes that are in the cluster cannot be accessed or because this node was excluded from the cluster.

Degraded

Indicates that the cluster is operational, but one or more of the member 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 management GUI.

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.

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.

IPv6 Address

You can use the IPv6 addresses to access the cluster from the command line tools or the management GUI.

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.

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.

Displaying an IPv6 address

After you have set the IPv6 address, you can display the IPv6 addresses and the IPv6 gateway addresses.

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 55. 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 55. Viewing the IPv6 address on the front-panel display

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.

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, assigned to a cluster, and ready to perform I/O.
- | **Service**
 - | There is an error that is preventing the node from operating as part of a cluster. It is safe to shut down the node in this state.

| **Candidate**

| The SAN Volume Controller node is not assigned to a cluster and is not in service. It is safe to shut down the node in this state.

| **Starting**

| SAN Volume Controller node is part of a cluster and is attempting to join the cluster. It cannot perform I/O.

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 node. The first 11 digits are always 50050768010.

| **Service Address**

| Pressing select on the Service Address panel displays the IP address that is configured for access to the service assistant and the service CLI.

| **Version options**

| 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, 6.1.0.0.

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 options

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 options

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.

1. To display the current fibre-channel port speed, do the following:
 1. Press and release the up or down button until Actions is displayed in the menu panel.
 2. Press the select button.
 3. Press and release the up or down button until Set FC Speed? appears in the menu panel.
 4. Press the select button.
 5. Press the right button until Edit Speed? displays in the menu panel. The current fibre-channel port speed is displayed.
 6. Press the right button until Confirm Speed? appears in the menu panel.
 7. Press the select button.

For the SAN Volume Controller 2145-8F2, you can use the Set FC Speed action option to change the fibre-channel port speed of a node that is not participating in a cluster.

| **Actions options**

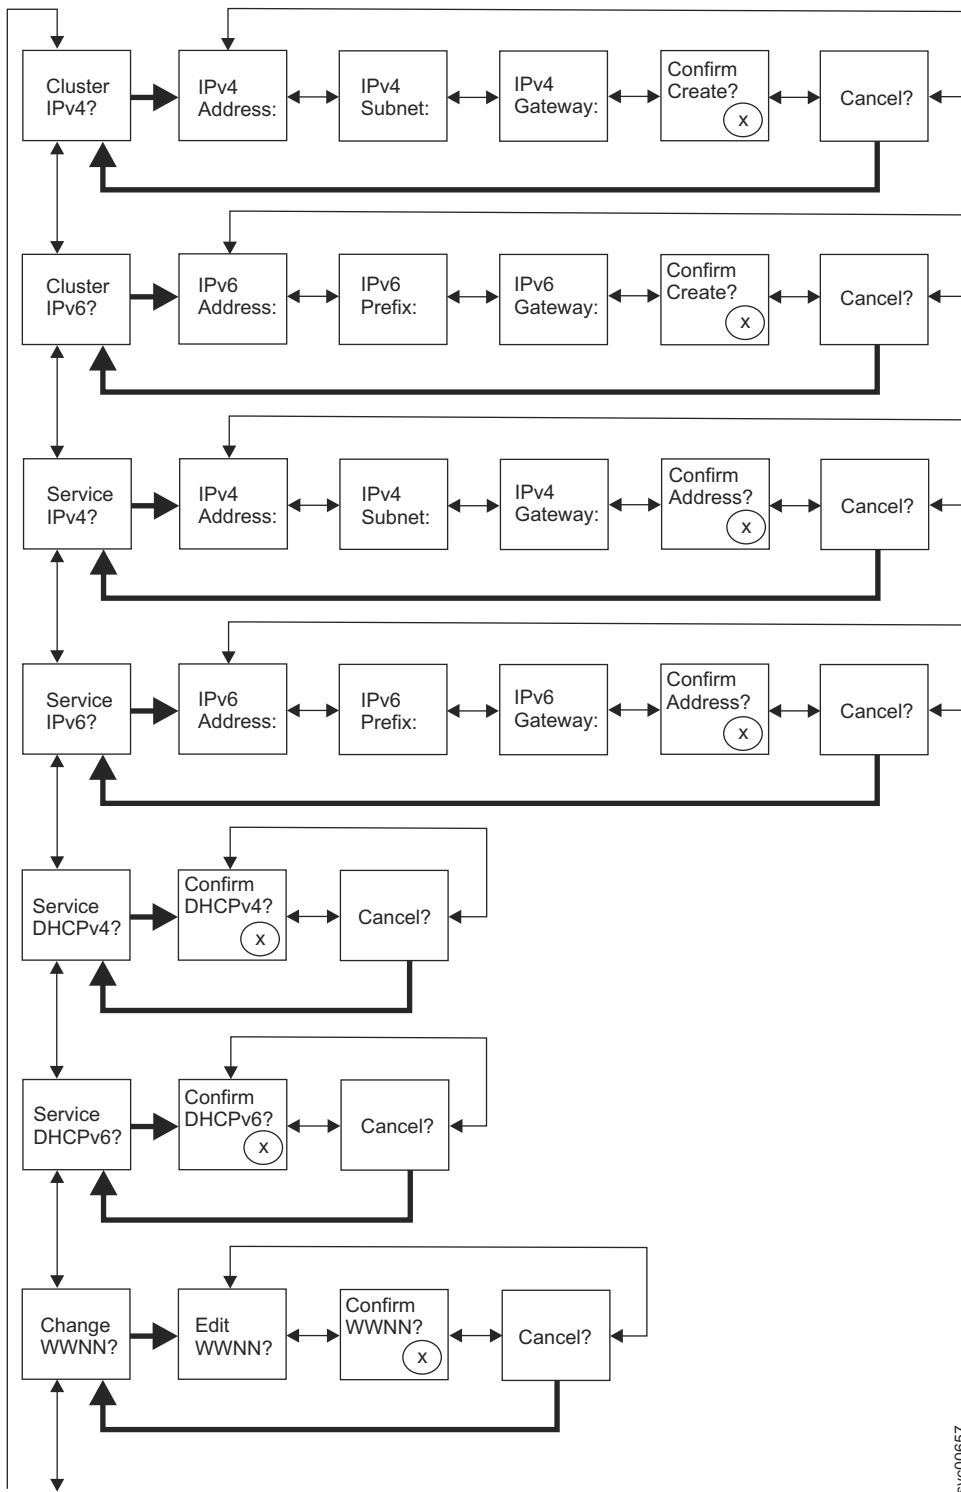
- | During normal operations, action menu options are available on the front panel display of the SAN Volume Controller node. Only use the front panel actions when directed to do so by a service procedure.
- | Inappropriate use can lead to loss of access to data or loss of data.

- | Figure 56 on page 92, Figure 57 on page 93, and Figure 58 on page 94 show the sequence of the actions options. In the figures, bold lines indicate that the select button was pressed. The lighter lines indicate the navigational path (up or down and left or right). The circled X indicates that if the select button is pressed, an action occurs using the data entered.

- | Only one action menu option at a time is displayed on the front-panel display.

- | **Note:** Options only display in the menu if they are valid for the current state of the node.

|



svc00657

Figure 56. Upper options of the actions menu on the front panel

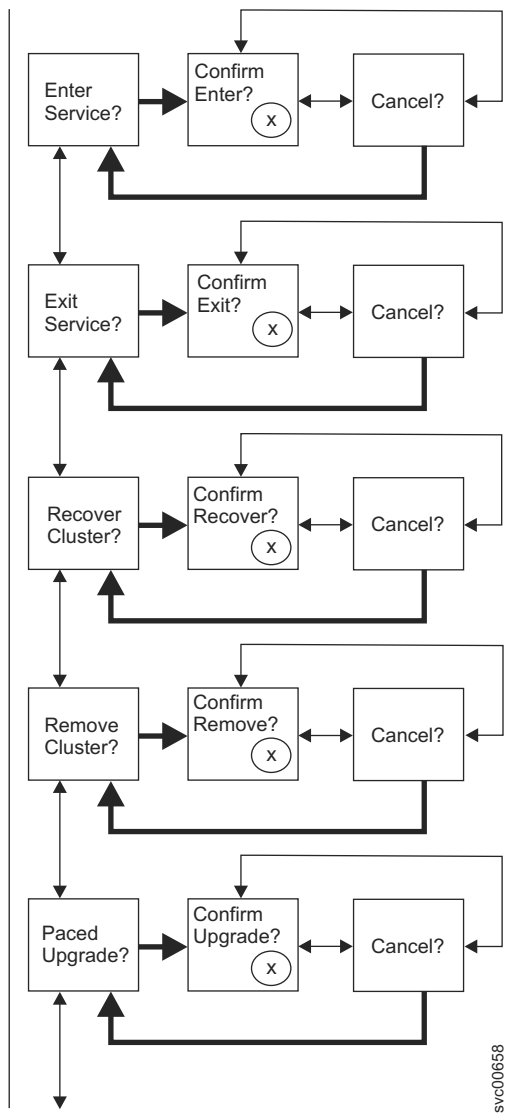
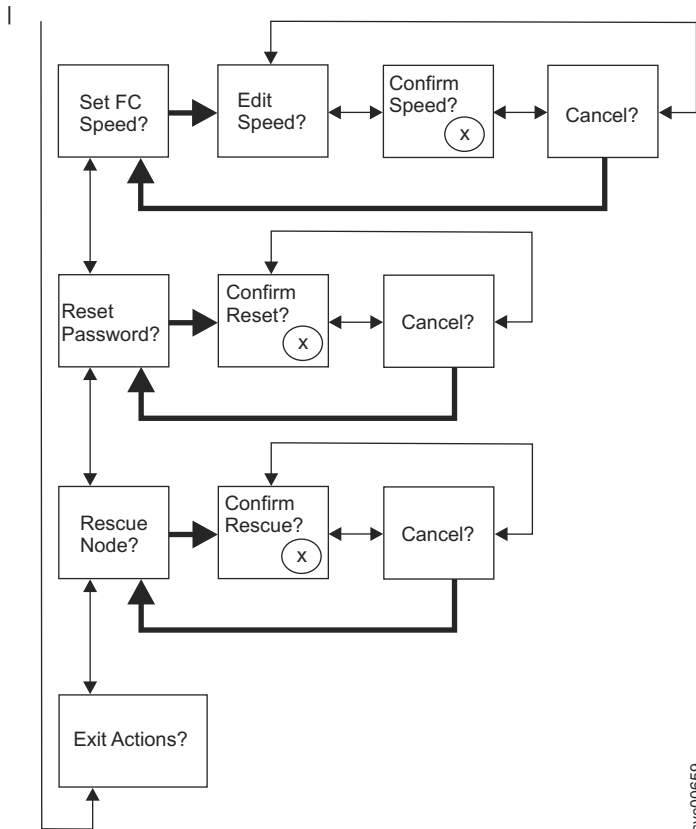


Figure 57. Middle options of the actions menu on the front panel



svc00659

Figure 58. Lower options of the actions menu on the front panel

To perform an action, navigate to the Actions option and press the select button. The action is initiated. Available parameters for the action are displayed. Use the left or right buttons to move between the parameters. The current setting is displayed on the second display line.

To set or change a parameter value, press the select button when the parameter is displayed. The value changes to edit mode. Use the left or right buttons to move between subfields, and use the up or down buttons to change the value of a subfield. When the value is correct, press select to leave edit mode.

Each action also has a Confirm? and a Cancel? panel. Pressing select on the Confirm? panel initiates the action using the current parameter value setting. Pressing select on the Cancel? panel returns to the Action option panel without changing the node.

Note: Messages might not display fully on the screen. You might see a right angle bracket (>) on the right 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 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 options are available from the Actions menu:

- Create a cluster with an IPv4 management address (Cluster IPv4)
- Create a cluster with an IPv6 management address (Cluster IPv6)
- Set the IPv4 service address of the node (Service IPv4)
- Set the IPv6 service address of the node (Service IPv6)

- | • Set a DHCP IPv4 service address (Service DHCPv4)
- | • Set a DHCP IPv6 service address (Service DHCPv6)
- | • Change the WWNN of the node (Change WWNN)
- | • Enter service state (Enter Service)
- | • Leave service state if possible (Exit Service)
- | • Recover cluster configuration (Recover Cluster)
- | • Remove cluster state (Remove Cluster)
- | • Perform user-paced CCU (Paced Upgrade)
- | • Set fibre-channel speed (Set FC Speed)
- | • Reset password (Reset Password)
- | • Rescue node software (Rescue Node)
- | • Exit actions menu (Exit Actions)

| **Cluster IPv4 or Cluster IPv6 options**

| You can create a cluster from the Cluster IPv4 or Cluster IPv6 action options.

| The Cluster IPv4 or Cluster IPv6 option allows you to create a SAN Volume Controller cluster.

| 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 management GUI or the CLI.

| Press the up and down buttons to navigate through the parameters that are associated with the Cluster option. When you have navigated to the desired parameter, press the select button.

| The parameters that are available include:

- | • IPv4 Address
- | • IPv4 Subnet
- | • IPv4 Gateway
- | • IPv4 Confirm Create?
- | • IPv6 Address
- | • IPv6 Subnet
- | • IPv6 Gateway
- | • IPv6 Confirm 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 Actions? is displayed. Press and release the select button.
- | 2. Press and release the up or down button until Cluster IPv4? is displayed. Press and release the select button.
- | 3. Edit the IPv4 address, the IPv4 subnet, and the IPv4 gateway.
- | 4. Press and release the left or right button until IPv4 Confirm Create? is displayed.
- | 5. 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 Actions? is displayed. Press and release the select button.
- | 2. Press and release the left or right button until Cluster Ipv6? is displayed. Press and release the select button.

- | 3. Edit the IPv6 address, the IPv6 prefix, and the IPv6 gateway.
- | 4. Press and release the left or right button until IPv6 Confirm Create? is displayed.
- | 5. 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 management GUI.

| 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 management GUI.

| 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 management GUI.

| 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 Confirm Create?

| Using this option, you can start an operation to create a cluster with an IPv4 address.

- | 1. Press and release the left or right button until IPv4 Confirm Create? is displayed.
- | 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 management GUI.

| **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 management GUI.

| 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 management GUI.

| Perform the following steps to set the IPv6 prefix:

| **Note:** This option is restricted to a value 0 - 127.

- | 1. Press and release the left or right button until IPv6 Prefix is displayed.
- | 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 management GUI.

- | Perform the following steps to set the IPv6 gateway address:
- | 1. Press and release the left or right button until IPv6 Gateway is displayed.
 - | 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 Confirm Create?**

| Using this option, you can start an operation to create a cluster with an IPv6 address.

- | 1. Press and release the left or right button until IPv6 Confirm Create? is displayed.
- | 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 management GUI.

| **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.

| **Service IPv4 or Service IPv6**

| You can use the front panel to change a service IPv4 address or a service IPv6 address.

| **IPv4 Address**

| The IPv4 Address panels show one of the following items for the selected Ethernet port:

- | • The active service address if the cluster has an IPv4 address. This address 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 address changes 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. 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 2 minutes, check that the selected address is valid on the subnetwork and that the Ethernet switch is working correctly.

IPv6 Address

The IPv6 Address panels show one of the following conditions for the selected Ethernet port:

- The active service address if the cluster has an IPv6 address. This address 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. 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 (four-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 2 minutes, check that the selected address is valid on the subnetwork and that the Ethernet switch is working correctly.

Service DHCPv4 or DHCPv6 options

The active service address for a cluster can be either a configured or fixed address, or it can be an address obtained through DHCP.

If a service IP address does not exist, you must assign a service IP address or use DHCP with this action.

To set the service IPv4 address to use DHCP, perform the following steps:

1. Press and release the up or down button until Service DHCPv4? is displayed.
2. Press and release the down button. Confirm DHCPv4? is displayed.
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.

To set the service IPv6 address to use DHCP, perform the following steps:

1. Press and release the up or down button until Service DHCPv6? is displayed.
2. Press and release the down button. Confirm DHCPv6? is displayed.
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.

Change WWNN?

The Change WWNN? 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 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 release the up or down button until Actions is displayed.
2. Press and release the select button.
3. Press and release the up or down button until Change 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.
4. 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.
5. 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.

Enter Service?

You can enter service state from the Enter Service? option. Service state can be used to remove a node from a candidate list or to prevent it from being readed to a cluster.

While in service state, the node cannot join or run as part of a cluster. If the node is active, entering service state can cause disruption to hosts if other faults exist in the system. Before entering service state for an active node, ensure that the following conditions are met:

- The other node in the I/O group is active and connected to all the hosts and storage systems.
- Verify that no volumes become unavailable if the action is performed.

To exit service state, ensure that all errors are resolved. You can exit service state by using the Exit Service? option or by restarting the node.

| **Exit Service?**

| You can exit service state from the Exit Service? option. This action releases the node from the service state.

| If there are no noncritical errors, the node enters candidate state. If possible, the node then becomes active in a cluster.

| To exit service state, ensure that all errors are resolved. You can exit service state by using this option or by restarting the node.

| **Recover Cluster?**

| You can recover an entire cluster if cluster data has been lost from all nodes by using the Recover Cluster? option.

| Perform service actions on nodes only when directed by the service procedures. If used inappropriately, service actions can cause loss of access to data or data loss

| **Remove Cluster?**

| The Remove Cluster? option deletes the cluster state data from the node.

| 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 management 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 management GUI.

| From the Remove 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 shows error 1195, Node missing, and the node is displayed in the list of nodes in the cluster. Remove the node by using the management GUI or CLI.

| **Paced Upgrade?**

| Use this option to control the time when individual nodes are upgraded within a concurrent code upgrade.

| **Note:** This action can be used only when the following conditions exist for the node:

- | • The node is in service state.
- | • The node has no errors.
- | • The node has been removed from the cluster.

| For additional information, see the "Upgrading the software manually" topic in the information center.

| **Set FC Speed?**

| You can change the speed of the fibre-channel ports on a SAN Volume Controller by using the Set FC Speed? option

| **Note:** This option is available only on SAN Volume Controller 2145-8F2 nodes.

| **Reset Password?**

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

| Use the Reset password? option 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.

| 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.

| If the node is in active state when the password is reset, the reset applies to all nodes in the cluster. If the node is in candidate or service state when the password is reset, the reset applies only to the single node.

| **Rescue Node?**

| You can start the automatic software recovery for this node by using the Rescue Node? option.

| **Note:** Another way to rescue a node is to force a node rescue when the node boots. It is the preferred method. Forcing a node rescue when a node boots 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.

| **Exit Actions?**

| Return to the main menu by selecting the Exit Actions? option.

| **Language? option**

You can change the language that displays on the front panel.

The Language? option allows you to change the language that is displayed on the menu. Figure 59 shows the Language? option sequence.




Figure 59. 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. Press and release the up or down button until Language? is displayed.
2. Press and release the select button.
3. Use the left and right buttons to move to the language that you want. 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 unavailable 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.

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. *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 225.

If the SAN Volume Controller software is running and you request it to power off from the management GUI, 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.

Chapter 7. Diagnosing problems

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

Event logs

- | By understanding the event log, you can do the following tasks:
- | • Manage the event log
- | • View the event log
- | • Describe the fields in the event log

Error codes

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

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

If the node is showing a boot message, failure message, or node error message, and you determined that the problem was caused by a software or firmware failure, you can restart the 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 225.
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.

Event reporting

- | Events that are detected are saved in an event log. As soon as an entry is made in this event log, the condition is analyzed. If any service activity is required, a notification is sent.

Event reporting process

The following methods are used to notify you and the IBM Support Center of a new event:

- The most serious cluster error 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 Systems 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, event notifications can be forwarded from a sender to a receiver through Call Home email.
- If you enabled Call Home, critical faults are reported directly to the IBM Support Center by email.

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 failure occurs, use “MAP 5000: Start” on page 201 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.

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 reports the failure condition, the SAN Volume Controller displays critical failure information about the front-panel display or sends noncritical failure information to the cluster event 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 events

When a significant change in status is detected, an event is logged in the event log.

Error data

Events are classified as either alerts or messages:

- An alert is logged when the event requires some action. Some alerts have an associated error code that defines the service action that is required. The service actions are automated through the fix procedures. If the alert does not have an error code, the alert represents an unexpected change in state. This situation must be investigated to see if it is expected or represents a failure. Investigate an alert and resolve it as soon as it is reported.
- A message is logged when a change that is expected is reported, for instance, an IBM FlashCopy operation completes.

Managing the event log

The event log has a limited size. After it is full, newer entries replace entries that are no longer required.

To avoid having a repeated event that fills the event log, some records in the event log refer to multiple occurrences of the same event. When event log entries are coalesced in this way, the time stamp of the first occurrence and the last occurrence of the problem is saved in the log entry. A count of the number of times that the error condition has occurred is also saved in the log entry. Other data refers to the last occurrence of the event.

Viewing the event log

You can view the event log by using the management GUI or the command-line interface (CLI).

You can view the event log by using the **Troubleshooting** options in the management GUI. The event log contains many entries. You can, however, select only the type of information that you need.

You can also view the event log by using the command-line interface (**svcinfo lseventlog**). See the "Command-line interface" topic for the command details.

Describing the fields in the event log

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

Table 29 describes some of the fields that are available to assist you in diagnosing problems.

Table 29. Description of data fields for the event log

| Data field | Description |
|----------------------|--|
| Event ID | This number precisely identifies why the event was logged. |
| Error code | This number describes the service action that should be followed to resolve an error condition. Not all events have error codes that are associated with them. Many event IDs can have the same error code because the service action is the same for all the events. |
| Sequence number | A number that identifies the event. |
| Event count | The number of events coalesced into this event log record. |
| Object type | The object type to which the event log relates. |
| Object ID | A number that uniquely identifies the instance of the object. |
| Fixed | When an alert is shown for an error condition, it indicates if the reason for the event was resolved. In many cases, the system automatically marks the events fixed when appropriate. There are some events that must be manually marked as fixed. If the event is a message, this field indicates that you have read and performed the action. The message must be marked as read. |
| First time | The time when this error event was reported. If events of a similar type are being coalesced together, so that one event log record represents more than one event, this field is the time the first error event was logged. |
| Last time | The time when the last instance of this error event was recorded in the log. |
| Root sequence number | If set, this number is the sequence number of an event that represents an error that probably caused this event to be reported. Resolve the root event first. |
| Sense data | Additional data that gives the details of the condition that caused the event to be logged. |

Event notifications

SAN Volume Controller can use Simple Network Management Protocol (SNMP) traps, syslog messages, and Call Home email to notify you and the IBM Support Center when significant events are detected. Any combination of these notification methods can be used simultaneously. Notifications are normally sent immediately after an event is raised. However, there are some events that might occur because of service actions that are being performed. If a recommended service action is active, these events are notified only if they are still unfixed when the service action completes.

Each event that SAN Volume Controller detects is assigned a notification type of Error, Warning, or Information. When you configure notifications, you specify where the notifications should be sent and which notification types are sent to that recipient.

Table 30 on page 110 describes the types of event notifications.

Table 30. 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 email 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. The allocation of notification type Warning does not imply that the event is less serious than one that has notification type Error.</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> |

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

SNMP traps

Simple Network Management Protocol (SNMP) is a standard protocol for managing networks and exchanging messages. The system can send SNMP messages that notify personnel about an event. You can use an SNMP manager to view the SNMP messages that SAN Volume Controller sends. You can use the management GUI or the 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 system. This file can be used with SNMP messages from all versions of the software. More information about the MIB file for SNMP is available at this website:

Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145

- | Search for **MIB**. Go to the downloads results to find **Management Information Base (MIB) file for SNMP**. Click this link to find download options.

Syslog messages

The syslog protocol is a standard protocol for forwarding log messages from a sender to a receiver on an IP network. The IP network can be either IPv4 or IPv6. The system can send syslog messages that notify personnel about an event. The system can transmit syslog messages in either expanded or concise format. You can use a syslog manager to view the syslog messages that the system sends. The system uses the User Datagram Protocol (UDP) to transmit the syslog message. You can use the management GUI or the SAN Volume Controller command-line interface to configure and modify your syslog settings.

- | Table 31 on page 111 shows how SAN Volume Controller notification codes map to syslog security-level codes.

Table 31. SAN Volume Controller notification types and corresponding syslog level codes

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

Table 32 shows how SAN Volume Controller values of user-defined message origin identifiers map to syslog facility codes.

Table 32. SAN Volume Controller values of user-defined message origin identifiers and syslog facility codes

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

Call Home email

The Call Home feature transmits operational and event-related data to you and IBM through a Simple Mail Transfer Protocol (SMTP) server connection in the form of an event notification email. When configured, this function alerts IBM service personnel about hardware failures and potentially serious configuration or environmental issues.

To send email, 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 email from the SAN Volume Controller cluster IP address. You can then use the management GUI or the SAN Volume Controller command-line interface to configure the email settings, including contact information and email recipients. Set the reply address to a valid email address. Send a test email to check that all connections and infrastructure are set up correctly. You can disable the Call Home function at any time using the management GUI or the SAN Volume Controller command-line interface.

Data that is sent with notifications

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

- Record type
- Machine type

- | • Machine serial number
- | • Error ID
- | • Error code
- | • Software version
- | • FRU part number
- | • Cluster name
- | • Node ID
- | • Error sequence number
- | • Time stamp
- | • Object type
- | • Object ID
- | • Problem data

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

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

| To send data and notifications to IBM service personnel, use one of the following email 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`

Inventory information email

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

Because inventory information is sent using the Call Home email function, you must meet the Call Home function requirements and enable the Call Home email function before you can attempt to send inventory information email. You can adjust the contact information, adjust the frequency of inventory email, or manually send an inventory email using the management GUI or the SAN Volume Controller command-line interface.

Inventory information that is sent to IBM includes the following information about the cluster on which the Call Home function is enabled. Sensitive information such as IP addresses is not included.

- | • Licensing information
- | • Details about the following objects and functions:
 - | Drives
 - | External storage systems
 - | Hosts
 - | MDisks
 - | Volumes
 - | RAID types
 - | Easy Tier
 - | FlashCopy

For detailed information about what is included in the Call Home inventory information, configure the system to send an inventory email to yourself.

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 can 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.

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 event log, and depending on the configuration, can be notified through email, 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 email, SNMP, or syslog.

Informational event codes

The informational event codes provide information on the status of an operation.

Informational events are recorded in the event log and, depending on the configuration, can be notified through email, 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 33 provides a list of informational event codes, the notification type, and the meaning of the event code.

Table 33. 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. |

Table 33. Informational event codes (continued)

| Event code | Notification type | Description |
|------------|-------------------|--|
| I 980301 | I | Degraded or offline managed disk is now online. |
| 980310 | I | A degraded or offline storage pool is now online. |
| I 980320 | I | Offline volume is now online. |
| I 980321 | W | Volume is offline because of degraded or offline storage pool. |
| I 980330 | I | All nodes can see the port. |
| I 980340 | I | All ports in this host are now logged in. |
| I 980341 | W | One or more ports in this host is now degraded. |
| I 980342 | W | One or more ports in this host is now offline. |
| I 980343 | W | All ports in this host are now offline. |
| 980349 | I | A node has been successfully added to the cluster. |
| I 980350 | I | The node is now a functional member of the cluster. |
| I 980351 | I | A noncritical hardware error occurred. |
| 980352 | I | Attempt to automatically recover offline node starting. |
| I 980370 | I | Both nodes in the I/O group are available. |
| 980371 | I | One node in the I/O group is unavailable. |
| 980372 | W | Both nodes in the I/O group are unavailable. |
| I 980392 | I | Cluster recovery completed. |
| 980435 | W | Failed to obtain directory listing from remote node. |
| 980440 | W | Failed to transfer file from remote node. |
| I 980445 | I | The migration is complete. |
| 980446 | I | The secure delete is complete. |
| 980501 | W | The virtualization amount is close to the limit that is licensed. |
| 980502 | W | The FlashCopy feature is close to the limit that is licensed. |
| 980503 | W | The Metro Mirror or Global Mirror feature is close to the limit that is licensed. |
| I 981002 | I | Fibre-channel discovery occurred; configuration changes are required. |
| I 981003 | I | Fibre-channel discovery occurred; configuration changes are complete. |
| I 981004 | I | Fibre-channel discovery occurred; no configuration changes were detected. |
| 981007 | W | The managed disk is not on the preferred path. |
| I 981009 | W | The initialization for the managed disk failed. |
| 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. |
| I 981015 | W | The LUN capacity equals or exceeds the maximum. Only part of the disk can be accessed. |
| 981020 | W | The managed disk error count warning threshold has been met. |
| 981022 | I | Managed disk view smoothing start |
| 981025 | I | Drive firmware download started |
| 981026 | I | Drive FPGA download started |

Table 33. Informational event codes (continued)

| Event code | Notification type | Description |
|------------|-------------------|---|
| 981101 | I | SAS discovery occurred; no configuration changes are required. |
| 981102 | I | SAS discovery occurred; configuration changes are required. |
| 981103 | I | SAS discovery occurred; configuration changes are complete. |
| 981104 | W | The LUN capacity equals or exceeds the maximum capacity. Only the first 1 PB of disk will be accessed. |
| 981105 | I | The drive format has started. |
| 981106 | I | The drive recovery was started. |
| 982003 | W | Insufficient virtual extents. |
| 982004 | W | The migration suspended because of insufficient virtual extents or too many media errors on the source managed disk. |
| 982007 | W | Migration has stopped. |
| 982009 | I | Migration is complete. |
| 982010 | W | Copied disk I/O medium error. |
| 983001 | I | The FlashCopy operation is prepared. |
| 983002 | I | The FlashCopy operation is complete. |
| 983003 | W | The FlashCopy operation has stopped. |
| 984001 | W | First customer data being pinned in a virtual disk working set. |
| 984002 | I | All customer data in a virtual disk working set is now unpinned. |
| 984003 | W | The volume working set cache mode is in the process of changing to synchronous destage because the volume working set has too much pinned data. |
| 984004 | I | Volume working set cache mode updated to allow asynchronous destage because enough customer data has been unpinned for the volume working set. |
| 984506 | I | The debug from an IERR was extracted to disk. |
| 984507 | I | An attempt was made to power on the slots. |
| 984508 | I | All the expanders on the strand were reset. |
| 984509 | I | The component firmware update paused to allow the battery charging to finish. |
| 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 thin-provisioned volume copy data in a node is pinned. |
| 986002 | I | All thin-provisioned volume copy data in a node is unpinned. |
| 986010 | I | The thin-provisioned volume copy import has failed and the new volume is offline; either upgrade the SAN Volume Controller software to the required version or delete the volume. |
| 986011 | I | The thin-provisioned volume copy import is successful. |
| 986020 | W | A thin-provisioned volume copy space warning has occurred. |
| 986030 | I | A thin-provisioned volume copy repair has started. |
| 986031 | I | A thin-provisioned volume copy repair is successful. |
| 986032 | I | A thin-provisioned volume copy validation is started. |

Table 33. Informational event codes (continued)

| Event code | Notification type | Description |
|------------|-------------------|---|
| 986033 | I | A thin-provisioned volume 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 marked as medium errors. |
| 986207 | I | The mirror disk repair has been started. |
| 986208 | W | A mirror copy repair, using the set medium error option, cannot complete. |
| 986209 | W | A mirror copy repair, using the resync option, cannot complete. |
| 987102 | W | Node coldstarted. |
| 987103 | W | A node power-off has been requested from the power switch. |
| 987104 | I | Additional fibre-channel ports were connected. |
| 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. |
| 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. |
| 988300 | W | An array MDisk is offline because it has too many missing members. |
| 988301 | I | The rebuild for an array MDisk was started. |
| 988302 | I | The rebuild for an array MDisk has finished. |
| 988304 | I | A RAID array has started exchanging an array member. |
| 988305 | I | A RAID array has completed exchanging an array member. |
| 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 emails. Their error fixed flags are ignored. Table 34 provides a list of the configuration event codes and their meanings.

Table 34. Configuration event codes

| Event code | Description |
|------------|---|
| 990101 | Modify cluster (attributes in the <code>svctask chcluster</code> command) |
| 990102 | The email test completed successfully |
| 990103 | The email test failed |
| 990105 | Delete node from cluster (attributes in the <code>svctask rmnode</code> command) |
| 990106 | Create host (attributes in the <code>svctask mkhost</code> command) |
| 990112 | Cluster configuration dumped to file (attributes from the <code>svcluster -x dumpconfig</code> command) |

Table 34. Configuration event codes (continued)

| Event code | Description |
|------------|---|
| 990117 | Create cluster (attributes in the svctask mkcluster command) |
| 990118 | Modify node (attributes in the svctask chnode command) |
| 990119 | Configure set controller name |
| 990120 | Shut down node (attributes in the svctask stopcluster command) |
| 990128 | Modify host (attributes in the svctask chhost command) |
| 990129 | Delete node (attributes in the svctask rmnode command) |
| 990138 | Volume modify (attributes in the svctask chvdisk command) |
| 990140 | Volume delete (attributes in the svctask rmvdisk command) |
| 990144 | Modify storage pool (attributes in the svctask chmdiskgrp command) |
| 990145 | Delete storage pool (attributes in the svctask rmdiskgrp command) |
| 990148 | Create storage pool (attributes in the svctask mkmdiskgrp command) |
| 990149 | Modify managed disk (attributes in the svctask chmdisk command) |
| 990150 | Modify managed disk |
| 990158 | Managed disk included |
| 990159 | Quorum created |
| 990160 | Quorum destroy |
| 990168 | Modify the I/O group a volume is assigned to |
| 990169 | Create a new volume (attributes in the svctask mkvdisk command) |
| 990173 | Add a managed disk to storage pool (attributes in the svctask addmdisk command) |
| 990174 | Delete a managed disk from storage pool (attributes in the svctask rmdisk command) |
| 990178 | Add a port to a Host (attributes in the svctask addhostport command) |
| 990179 | Delete a port from a Host (attributes in the svctask rmhostport command) |
| 990182 | Create a host mapping (attributes in the svctask mkvdiskhostmap command) |
| 990183 | Delete a host mapping (attributes in the svctask rmdiskhostmap command) |
| 990184 | Create a FlashCopy mapping (attributes in the svctask mkfcmap command) |
| 990185 | Modify a FlashCopy mapping (attributes in the svctask chfcmap command) |
| 990186 | Delete a FlashCopy mapping (attributes in the svctask rmfcmap command) |
| 990187 | Prepare a FlashCopy mapping (attributes in the svctask prestartfcmap command) |
| 990188 | Prepare a FlashCopy consistency group (attributes in the svctask prestartfcconsistgrp command) |
| 990189 | Trigger a FlashCopy mapping (attributes in the svctask startfcmap command) |
| 990190 | Trigger a FlashCopy consistency group (attributes in the svctask startfcconsistgrp command) |
| 990191 | Stop a FlashCopy mapping (attributes in the svctask stopfcmap command) |
| 990192 | Stop a FlashCopy consistency group (attributes in the svctask stopfcconsistgrp command) |
| 990193 | FlashCopy set name |
| 990194 | Delete a list of ports from a Host (attributes in the svctask rmhostport command) |
| 990196 | Shrink a volume. |
| 990197 | Expand a volume (attributes in the svctask expandvdisksize command) |

Table 34. Configuration event codes (continued)

| Event code | Description |
|------------|--|
| 990198 | Volume expanded by a single extent. |
| 990199 | Modify the I/O governing rate for a volume |
| 990203 | Initiate manual managed disk discovery (attributes in the svctask detectmdisk command) |
| 990204 | Create FlashCopy consistency group (attributes in the svctask mkfconsistgrp command) |
| 990205 | Modify FlashCopy consistency group (attributes in the svctask chfconsistgrp command) |
| 990206 | Delete FlashCopy consistency group (attributes in the svctask rmfconsistgrp command) |
| 990207 | Delete a list of Hosts (attributes in the svctask rmhost command) |
| 990213 | Change the I/O group a node belongs to (attributes in the svctask chiogrp command) |
| 990216 | Apply software upgrade (attributes in the satask installsoftware command) |
| 990219 | Analyze event log (attributes in the svctask finderr command) |
| 990220 | Dump event log (attributes in the satask snap command) |
| 990222 | Fix event log entry (attributes in the svctask cherrstate command) |
| 990223 | Migrate a single extent (attributes in the svctask migrateexts command) |
| 990224 | Migrate a number of extents |
| 990225 | Create a Metro Mirror or Global Mirror or Global Mirror relationship (attributes in the svctask mkcrrelationship command) |
| 990226 | Modify a Metro Mirror or Global Mirror relationship (attributes in the svctask chrrelationship command) |
| 990227 | Delete a Metro Mirror or Global Mirror relationship (attributes in the svctask rmcrrelationship command) |
| 990229 | Start a Metro Mirror or Global Mirror relationship (attributes in the svctask startcrrelationship command) |
| 990230 | Stop a Metro Mirror or Global Mirror relationship (attributes in the svctask stopcrrelationship command) |
| 990231 | Switch a Metro Mirror or Global Mirror relationship (attributes in the svctask switchcrrelationship command) |
| 990232 | Start a Metro Mirror or Global Mirror consistency group (attributes in the svctask startcrconsistgrp command) |
| 990233 | Stop a Metro Mirror or Global Mirror consistency group (attributes in the svctask stopcrconsistgrp command) |
| 990234 | Switch a Metro Mirror or Global Mirror consistency group (attributes in the svctask switchcrconsistgrp command) |
| 990235 | Managed disk migrated to a storage pool |
| 990236 | Volume migrated to a new managed disk |
| 990237 | Create partnership with remote cluster (attributes in the svctask mkpartnership command) |
| 990238 | Modify partnership with remote cluster (attributes in the svctask chpartnership command) |
| 990239 | Delete partnership with remote cluster (attributes in the svctask rmpartnership command) |

Table 34. Configuration event codes (continued)

| Event code | Description |
|------------|---|
| 990240 | Create a Metro Mirror or Global Mirror consistency group (attributes in the svctask mkrconsistgrp command) |
| 990241 | Modify a Metro Mirror or Global Mirror consistency group (attributes in svctask chrconsistgrp) |
| 990242 | Delete a Metro Mirror or Global Mirror consistency group (attributes in the svctask rmrconsistgrp command) |
| 990245 | Node shutdown imminent |
| 990246 | Node remove |
| 990247 | Node unpend |
| 990380 | Time zone changed (attributes in the svctask settimezone command) |
| 990383 | Change cluster time (attributes in the svctask setclustertime command) |
| 990385 | System time changed |
| 990386 | SSH key added (attributes in the svctask addsshkey command) |
| 990387 | SSH key removed (attributes in the svctask rmsshkey command) |
| 990388 | All SSH keys removed (attributes in the svctask rmallsshkeys 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 volume has been modified. |
| 991026 | The primary copy of the volume has been modified. |
| 991027 | The volume synchronization rate has been modified. |
| 991028 | The thin-provisioned volume warning capacity has been modified. |
| 991029 | A mirrored copy has been added to a volume. |
| 991030 | A repair of mirrored volume copies has started. |
| 991031 | A volume copy has been split from a mirrored volume. |
| 991032 | A volume copy has been removed from a mirrored volume. |

SCSI event reporting

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

SCSI status

Some events are part of the SCSI architecture and are handled by the host application or device drivers without reporting an event. Some events, such as read and write I/O events and events that are

associated with the loss of nodes or loss of access to backend devices, cause application I/O to fail. To help troubleshoot these events, SCSI commands are returned with the Check Condition status and a 32-bit event identifier is included with the sense information. The identifier relates to a specific event in the cluster event log.

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

Table 35 describes the SCSI status and codes that are returned by the nodes.

Table 35. 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 SAM-2 where a reserve or persistent reserve condition exists. |
| Task set full | 28h | The initiator has at least one task queued for that LUN on this port. |
| ACA active | 30h | This code is reported as specified in SAM-2. |
| Task aborted | 40h | This code is returned if TAS is set in the control mode page 0Ch. The node has a default setting of TAS=0, which cannot be changed; therefore, the node does not report this status. |

SCSI Sense

Nodes notify the hosts of events on SCSI commands. Table 36 defines the SCSI sense keys, codes and qualifiers that are returned by the nodes.

Table 36. 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. |
| 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"> The node lost sight of the cluster and cannot perform I/O operations. The additional sense does not have additional information. 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 volumes. |

Table 36. SCSI sense keys, codes, and qualifiers (continued)

| Key | Code | Qualifier | Definition | Description |
|-----|------|-----------|---|---|
| 3h | 00h | 00h | Medium event | This is only returned for read or write I/Os. The I/O suffered an event at a specific LBA within its scope. The location of the event is reported within the sense data. The additional sense also includes a reason code that relates the event to the corresponding event log entry. For example, a RAID controller event or a migrated medium event. |
| 4h | 08h | 00h | Hardware event. A command to logical unit communication failure has occurred. | The I/O suffered an event that is associated with an I/O event that is returned by a RAID controller. The additional sense includes a reason code that points to the sense data that is returned by the controller. This is only returned for I/O type commands. This event is also returned from FlashCopy target volumes in the prepared and preparing state. |
| 5h | 25h | 00h | Illegal request. The logical unit is not supported. | The logical unit does not exist or is not mapped to the sender of the command. |

Reason codes

The reason code appears in bytes 20-23 of the sense data. The reason code provides the node with a specific log entry. The field is a 32-bit unsigned number that is presented with the most significant byte first. Table 37 lists the reason codes and their definitions.

If the reason code is not listed in Table 37, the code refers to a specific event in the cluster event log that corresponds to the sequence number of the relevant event log entry.

Table 37. 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 volume. |
| 81 | Wait for the validation to complete, or delete the volume. |
| 82 | An offline thin-provisioned volume has caused data to be pinned in the directory cache. Adequate performance cannot be achieved for other thin-provisioned volumes, so they have been taken offline. |
| 85 | The volume has been taken offline because checkpointing to the quorum disk failed. |

Table 37. Reason codes (continued)

| Reason code (decimal) | Description |
|-----------------------|--|
| 86 | The <code>svctask repairvdiskcopy -medium</code> 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 38 lists the object codes and corresponding object types.

Table 38. 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 |
| 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 describe a service procedure that must be followed. Each event ID that requires service has an associated error code.

- | Table 39 on page 123 lists the event IDs and corresponding error codes.

Table 39. Event IDs and codes

| Event ID | Notification type | Condition | Error code |
|----------|-------------------|--|------------|
| 009002 | None | One or more paths to a managed disk have been lost. | None |
| 009003 | None | All paths to a managed disk have been lost. | None |
| 009011 | None | One or more managed disks in this group are offline. | None |
| 009012 | None | All managed disks in this group are offline. | None |
| 009020 | E | An automatic cluster recovery has started. All configuration commands are blocked. | 1001 |
| 009031 | None | One or more nodes cannot see this port. | None |
| 009032 | None | No node can see this port. | None |
| 009040 | E | The error event log is full. | 1002 |
| 009052 | W | The following causes are possible: <ul style="list-style-type: none"> • The node is missing • The node is no longer a functional member of the cluster | 1196 |
| 009053 | E | A node has been missing for 30 minutes. | 1195 |
| 009100 | W | The software install process has failed. | 2010 |
| 009101 | W | The software upgrade package delivery has failed. | 2010 |
| 009110 | None | Software installation process stalled due to lack of redundancy. | 2010 |
| 009115 | None | A software downgrade has failed. | 2008 |
| 009150 | W | Unable to connect to the SMTP (email) server | 2600 |
| 009151 | W | Unable to send mail through the SMTP (email) server | 2601 |
| 009170 | W | 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 |
| 009175 | W | The usage for the thin-provisioned volume is not licensed. | 3033 |
| 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 email to any of the configured email servers. | 3081 |
| 009182 | W | The external virtualization feature license limit was exceeded. | 3032 |
| 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 |
| 010004 | None | Device excluded because of excessive errors on all managed disks. | 1640 |
| 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 |

Table 39. Event IDs and codes (continued)

| Event ID | Notification type | Condition | Error code |
|----------|-------------------|--|------------|
| 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 |
| 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 | W | There are too many devices presented to the cluster. | 1200 |
| 010022 | W | There are too many managed disks presented to the cluster. | 1200 |
| 010023 | W | There are too many LUNs presented to a node. | 1200 |
| 010024 | W | There are too many drives presented to a cluster. | 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"> • Sense • Key • Code • Qualifier | 1370 |
| 010031 | E | One or more MDisks on a controller are degraded. | 1623 |
| 010032 | W | The controller configuration limits failover. | 1625 |
| 010033 | E | The controller configuration uses the RDAC mode; this is not supported. | 1624 |
| 010034 | 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 |
| 010055 | W | An unrecognized SAS device. | 1665 |
| 010056 | E | SAS error counts exceeded the warning thresholds. | 1216 |
| 010057 | E | SAS errors exceeded critical thresholds. | 1216 |

Table 39. Event IDs and codes (continued)

| Event ID | Notification type | Condition | Error code |
|----------|-------------------|--|------------|
| 010063 | W | Drive medium error. | 1321 |
| 010066 | W | Controller indicates that it does not support descriptor sense for LUNs that are greater than 2 TBs. | 1625 |
| 010070 | W | Too many controller target ports were presented to the cluster. | 1200 |
| 010071 | W | Too many target ports were presented to the cluster from a single controller. | 1200 |
| 010073 | E | The drive is reporting predictive failure analysis (PFA) errors. | 1680 |
| 010088 | E | There were excessive medium errors on the drive. | 1680 |
| 010089 | E | There were excessive overall timeout errors on the drive. | 1680 |
| 010090 | E | There were excessive times when the drive stopped. | 1680 |
| 010091 | E | A drive failed validation testing. | 1680 |
| 010095 | E | Drive failed. | 1687 |
| 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 |
| 020005 | None | Migration was suspended. | None |
| 029001 | W | The managed disk has bad blocks. | 1840 |
| 029002 | E | The system failed to create a bad block because MDisk already has the maximum number of allowed bad blocks. | 1226 |
| 029003 | E | The system failed to create a bad block because cluster already has the maximum number of allowed bad blocks. | 1225 |
| 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 thin-provisioned volume copy is offline because there is insufficient space. | 1865 |
| 060002 | W | The thin-provisioned volume copy is offline because the metadata is corrupt. | 1862 |
| 060003 | W | The thin-provisioned volume copy is offline because the repair has failed. | 1860 |
| 062001 | W | Unable to mirror medium error during volume copy synchronization | 1950 |
| 062002 | W | The mirrored volume is offline because the data cannot be synchronized. | 1870 |

Table 39. Event IDs and codes (continued)

| Event ID | Notification type | Condition | Error code |
|----------|-------------------|--|------------|
| 062003 | W | The repair process for the mirrored disk has stopped because there is a difference between the copies. | 1600 |
| 070000 | E | Unrecognized node error. | 1083 |
| 070510 | E | Detected memory size does not match the expected memory size. | 1022 |
| 070517 | E | The WWNN that is stored on the service controller and the WWNN that is stored on the drive do not match. | 1192 |
| 070521 | E | Unable to detect any Fibre Channel adapter. | 1016 |
| 070522 | E | The system board processor has failed. | 1020 |
| 070523 | W | The internal disk file system of the node is damaged. | 1187 |
| 070524 | E | Unable to update BIOS settings. | 1027 |
| 070525 | E | Unable to update the service processor firmware for the system board. | 1020 |
| 070528 | W | The ambient temperature is too high while the system is starting. | 1182 |
| 070550 | E | Cannot form cluster due to lack of cluster resources. | 1192 |
| 070556 | E | Duplicate WWNN detected on the SAN. | 1192 |
| 070558 | E | A node is unable to communicate with other nodes. | 1192 |
| 070562 | E | The node hardware does not meet minimum requirements. | 1183 |
| 070564 | E | Too many software failures. | 1188 |
| 070565 | E | The internal drive of the node is failing. | 1030 |
| 070574 | E | The node software is damaged. | 1187 |
| 070576 | E | The cluster data cannot be read. | 1030 |
| 070578 | E | The cluster data was not saved when power was lost. | 1194 |
| 070580 | E | Unable to read the service controller ID. | 1044 |
| 070581 | E | 2145 UPS-1U serial link error. | 1181 |
| 070582 | E | 2145 UPS-1U battery error. | 1181 |
| 070583 | E | 2145 UPS-1U electronics error. | 1171 |
| 070584 | E | 2145 UPS-1U overloaded. | 1166 |
| 070585 | E | 2145 UPS-1U failure | 1171 |
| 070586 | E | Power supply to 2145 UPS-1U does not meet requirements. | 1141 |
| 070587 | E | Incorrect type of uninterruptible power supply detected. | 1152 |
| 070588 | E | 2145 UPS-1U is not cabled correctly. | 1151 |
| 070589 | E | The ambient temperature limit for the 2145 UPS-1U was exceeded. | 1136 |
| 070590 | E | Repeated node restarts because of 2145 UPS-1U errors. | 1186 |
| 070670 | W | Insufficient uninterruptible power supply charge to allow node to start. | 1193 |
| 070690 | W | Node held in service state. | 1189 |
| 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 |

Table 39. Event IDs and codes (continued)

| Event ID | Notification type | Condition | Error code |
|----------|-------------------|---|------------|
| 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 | Cluster path failure. | 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 |
| 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 |

Table 39. Event IDs and codes (continued)

| Event ID | Notification type | Condition | Error code |
|----------|-------------------|--|------------|
| 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 |
| 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 |
| 075033 | None | There was a service controller read failure. | None |
| 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 |

Table 39. Event IDs and codes (continued)

| Event ID | Notification type | Condition | Error code |
|----------|-------------------|---|------------|
| 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 |
| 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 |

Table 39. Event IDs and codes (continued)

| Event ID | Notification type | Condition | Error code |
|----------|-------------------|---|------------|
| 077174 | E | System board - any voltage low. This error applies to only the SAN Volume Controller 2145-8G4 model. | 1106 |
| 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 | None | Data has not been recovered on volumes. | None |
| 079500 | W | The limit on the number of cluster secure shell (SSH) sessions has been reached. | 2500 |
| 079501 | I | 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 |
| 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 |

Table 39. Event IDs and codes (continued)

| Event ID | Notification type | Condition | Error code |
|----------|-------------------|--|------------|
| 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 |
| 084000 | W | An array MDisk has deconfigured members and has lost redundancy. | 1689 |
| 084100 | W | An array MDisk is corrupt because of lost metadata. | 1240 |
| 084200 | W | An array MDisk has taken a spare member that is not an exact match to the array goals. | 1692 |
| 084300 | W | An array MDisk is no longer protected by an appropriate number of suitable spares. | 1690 |
| 084500 | W | An array MDisk is offline. The metadata for the inflight writes is on a missing node. | 1243 |
| 084600 | W | An array MDisk is offline. Metadata on the missing node contains needed state information. | 1243 |

Determining a hardware boot failure

During the hardware boot, you see progress messages. If the boot detects a situation where it cannot continue, it fails. The cause might be that the software on the hard disk drive is missing or damaged. If possible, the boot sequence loads and starts the SAN Volume Controller software. Any faults that are detected are reported as a node error.

Line 1 of the front panel displays the message `Boot ing` that is followed by the boot code. Line 2 of the display shows a boot progress indicator. If the boot code detects an error that makes it impossible to continue, `Failed` is displayed. You can use the code to isolate the fault.

The following figure shows an example of a hardware boot display.



Figure 60. Example of a boot error code

Perform the following steps to determine a 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.

Boot code reference

Boot codes are displayed on the screen when a node is booting.

The codes indicate the progress of the boot operation. Line 1 of the front panel displays the message `Boot ing` that is followed by the boot code. Line 2 of the display shows a boot progress indicator. Figure 61 on page 132 provides a view of the boot progress display.



Figure 61. Example of a boot progress display

Node error code overview

Node error codes describe failure that relate to a specific node. Node rescue codes are displayed on the menu screen during node rescue.

Because node errors are specific to a node, for example, memory has failed, the errors are only reported on that node.

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 62 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.



Figure 62. 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.

| There are two types of node errors: critical node errors and noncritical node errors.

| Critical errors

A critical error means that the node is not able to participate in a cluster until the issue that is preventing it from joining a cluster is resolved. This error occurs because part of the hardware has failed or the system detects that the software is corrupt. If a node has a critical node error, it is in service state, and the fault LED on the node is on. The exception is when the node cannot connect to enough resources to form a cluster. It shows a critical node error but is in the starting state. Resolve the errors in priority order. The range of errors that are reserved for critical errors are 500 - 699.

| Noncritical errors

| A noncritical error code is logged when there is a hardware or software failure that is related to just one
| specific node. These errors do not stop the node from entering active state and joining a cluster. If the
| node is part of a cluster, there is also an alert that describes the error condition. The range of errors that
| are reserved for noncritical errors are 800 - 899.

Node rescue codes

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 63 shows an example of a displayed node rescue code.



Figure 63. Example of a node-rescue error code

The three-digit code that is shown in Figure 63 represents a node rescue code.

Note: The 2145 UPS-1U will not power off following a node rescue failure.

Cluster code overview

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. Cluster recovery codes indicate that a critical software error has occurred that might corrupt your cluster. Cluster error codes describe errors other than creation and recovery errors. Each error-code topic includes an error code number, a description, action, and possible field-replaceable units (FRUs).

Create cluster error codes

Figure 64 provides an example of a create-cluster error code.



Figure 64. 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.

Cluster recovery codes

You must perform software problem analysis before you can perform further cluster operations to avoid the possibility of corrupting your cluster configuration.

Figure 65 provides an example of a cluster-recovery error code.



Figure 65. Example of a cluster-recovery error code

Cluster error codes

Cluster error codes describe errors other than recovery errors.

Figure 66 provides an example of a cluster error code.




Figure 66. Example of a cluster error code

Error code range

This topic shows the number range for each message classification.

Table 40 lists the number range for each message classification.

Table 40. Message classification number range

| Message classification | Range | |
|------------------------|-------------------------|---------|
| Booting codes | 100-299 | |
| Node errors | Node rescue errors | 300-399 |
| | Critical node errors | 500-699 |
| | Noncritical node errors | 800-899 |
| Create cluster error | 700, 710 | |
| Cluster recovery | 920, 990 | |
| Cluster error codes | 1001-3081 | |

Booting codes

100 Boot is running

Explanation: The SAN Volume Controller node has started. It is running diagnostics and loading the runtime code.

User response: Go to the hardware boot MAP to resolve the problem.

Possible Cause-FRUs or other:

2145-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%)

120 Disk drive hardware error

Explanation: The internal disk drive of the node has reported an error. The node is unable to start.

User response: 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%)

130 Checking the internal disk file system

Explanation: The file system on the internal disk drive of the node is being checked for inconsistencies.

User response: If the progress bar has been stopped for at least five minutes, power off the node and then power on the node. If the boot process stops again at this point, run the node rescue procedure.

Possible Cause-FRUs or other:

- None.

132 Updating BIOS settings of the node

Explanation: The system has found that changes are required to the BIOS settings of the node. These changes are being made. The node will restart once the changes are complete.

User response: If the progress bar has stopped for more than 10 minutes, or if the display has shown codes 100 and 132 three times or more, go to MAP 5900: Hardware boot to resolve the problem.

135 Verifying the software

Explanation: The software packages of the node are being checked for integrity.

User response: Allow the verification process to complete.

137 Updating system board service processor firmware

Explanation: The service processor firmware of the node is being updated to a new level. This process can take 90 minutes. Do not restart the node while this is in progress.

User response: Allow the updating process to complete.

150 Loading cluster code

Explanation: The SAN Volume Controller code is being loaded.

User response: If the progress bar has been stopped for at least 90 seconds, power off the node and then power on the node. If the boot process stops again at this point, run the node rescue procedure.

Possible Cause-FRUs or other:

- None.

155 Loading cluster data

Explanation: The saved cluster state and cache data is being loaded.

User response: If the progress bar has been stopped for at least 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.

160 Updating the service controller

Explanation: The firmware on the service controller is being updated. This can take 30 minutes.

User response: When a node rescue is occurring, if the progress bar has been stopped for at least 30 minutes, exchange the FRU for a new FRU. When a node rescue is not occurring, if the progress bar has been stopped for at least 15 minutes, exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

2145-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.

User response: Exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

2145-CF8

- Service controller (95%)
- Service controller cable (5%)

All previous 2145 models

- Service controller (100%)

182 Checking uninterruptible power supply

Explanation: The node is checking whether the uninterruptible power supply is operating correctly.

User response: Allow the checking process to complete.

232 Checking uninterruptible power supply connections

Explanation: The node is checking whether the power and signal cable connections to the uninterruptible power supply are correct.

User response: Allow the checking process to complete.

Create cluster errors

870 **The cluster cannot be created because the counter maximum has been reached.**

Explanation: Each time a node creates a new cluster, a unique ID is generated by the service controller of the node. Once 255 clusters have been created, the service controller must be replaced.

User response: Use a different node to create the cluster.

871 **The cluster cannot be created because**

Node errors

300 **The 2145 is running node rescue.**

Explanation: The 2145 is running node rescue.

User response: If the progress bar has been stopped for at least two minutes, exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

2145-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.

User response: If the progress bar has been stopped for two minutes, exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

2145-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%)

the counter increment failed.

Explanation: When a new cluster ID is requested from the service controller, the service controller must increase the ID counter. The new ID is returned for verification. If the ID counter has not been increased, this error code is displayed. This error has occurred because the service controller failed.

User response: Exchange the FRU for a new FRU.

- Disk drive cables (5%)

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

Explanation: A 2145 format operation has failed.

User response: Exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

2145-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.

User response: If the progress bar has been stopped for two minutes, exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

2145-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.

User response: If the progress bar has been stopped for more than two minutes, exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

- Fibre-channel adapter (100%)

345 The 2145 is searching for a donor node from which to copy the software.

Explanation: The node is searching at 1 Gb/s for a donor node.

User response: If the progress bar has stopped for more than two minutes, exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

- Fibre-channel adapter (100%)

350 The 2145 cannot find a donor node.

Explanation: The 2145 cannot find a donor node.

User response: If the progress bar has stopped for more than two minutes, perform the following steps:

1. Ensure that all of the fibre-channel cables are connected correctly and securely to the cluster.
2. Ensure that at least one other node is operational, is connected to the same fibre-channel network, and is a donor node candidate. A node is a donor node candidate if the version of software that is installed on that node supports the model type of the node that is being rescued.

3. Ensure that the fibre-channel zoning allows a connection between the node that is being rescued and the donor node candidate.
4. Perform the problem determination procedures for the network.

Possible Cause-FRUs or other:

- None

Other:

- Fibre-channel network problem

360 The 2145 is loading software from the donor.

Explanation: The 2145 is loading software from the donor.

User response: If the progress bar has been stopped for at least two minutes, restart the node rescue procedure.

Possible Cause-FRUs or other:

- None

365 Cannot load SW from donor

Explanation: None.

User response: None.

370 Installing software

Explanation: The 2145 is installing software.

User response:

1. If this code is displayed and the progress bar has been stopped for at least ten minutes, the software install process has failed with an unexpected software error.
2. Power off the 2145 and wait for 60 seconds.
3. Power on the 2145. The software upgrade operation continues.
4. Report this problem immediately to your Software Support Center.

Possible Cause-FRUs or other:

- None

510 The detected memory size does not match the expected memory size.

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.

User response: 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%)

511 Memory bank 1 of the 2145 is failing.

Explanation: Memory bank 1 of the 2145 is failing.

User response: 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%)

513 Memory bank 2 of the 2145 is failing.

Explanation: Memory bank 2 of the 2145 is failing.

User response: 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%)

514 Memory bank 3 of the 2145 is failing.

Explanation: This cannot occur on the 2145-4F2.

User response: 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%)

515 Memory bank 4 of the 2145 is failing.

Explanation: This cannot occur on the 2145-4F2.

User response: 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%)

517 The WWNNs of the service controller and the disk do not match.

Explanation: The node is unable to determine the WWNN that it should use. This is because of the service controller or the nodes internal drive being replaced.

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

1. Continue to follow the hardware remove and replace procedure for the service controller or disk these explain the service actions.
2. If you have not followed the hardware remove and replace procedures, you should determine the correct WWNN. If you do not have this information recorded, examine your fibre channel switch configuration to see whether it is listed there. Follow the procedures to change the WWNN of a node.

Possible Cause-FRUs or other:

- None

521 Unable to detect a fibre-channel adapter

Explanation: The 2145 cannot detect any fibre-channel adapter cards.

Explanation: The 2145 cannot detect any fibre-channel adapter cards.

User response: 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%)

522 The system board service processor has failed.

Explanation: The service processor on the system board has failed.

User response: 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%)

523 The internal disk file system is damaged.

Explanation: The node startup procedures have found problems with the file system on the internal disk of the node.

User response: Follow troubleshooting procedures to reload the software.

1. Follow the procedures to rescue the software of a node from another node.
2. If the rescue node does not succeed, use the hardware remove and replace procedures.

Possible Cause-FRUs or other:

- Disk drive (100%)

524 Unable to update BIOS settings.

Explanation: Unable to update BIOS settings.

User response: Power off node, wait 30 seconds, and then power on again. If the error code is still reported, replace the system board.

Possible Cause-FRUs or other:

- System board (100%)

525 Unable to update system board service processor firmware.

Explanation: The process of updating the system board service processor firmware might take up to 90 minutes.

User response: 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%)

528 Ambient temperature is too high during system startup.

Explanation: The ambient temperature read during the node startup procedures is too high for the node to continue. The startup procedure will continue when the temperature is within range.

User response: Reduce the temperature around the system.

1. Resolve the issue with the ambient temperature, by checking and correcting:
 - a. Room temperature and air conditioning
 - b. Ventilation around the rack
 - c. Airflow within the rack

Possible Cause-FRUs or other:

- Environment issue (100%)

530 A problem with one of the node's power supplies has been detected.

Explanation: The 530 error code is followed by two numbers. The first number is either 1 or 2 to indicate which power supply has the problem.

The second number is either 1, 2 or 3 to indicate the reason. 1 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.

User response:

1. Ensure that the power supply is seated correctly and that the power cable is attached correctly to both the node and to 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%)

550 A cluster cannot be formed because of 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

WWWWWWWWWWWWWWWWWW/LL, where WWWWWWWWWWWWWWWWWWW 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.

User response: Follow troubleshooting procedures to correct connectivity issues between the cluster nodes and the quorum devices.

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.

555 Power Domain error

Explanation: Both 2145s in an I/O group are being powered by the same uninterruptible power supply. The ID of the other 2145 is displayed with the node error code on the front panel.

User response: Ensure that the configuration is correct and that each 2145 in an I/O group is connected from a separate uninterruptible power supply.

556 A duplicate WWNN has been detected.

Explanation: The node has detected another device that has the same World Wide Node Name (WWNN) on the fibre-channel network. A WWNN is 16 hexadecimal digits long. For a cluster, the first 11 digits are always 50050768010. The last 5 digits of the WWNN are given in the additional data of the error and appear on the front panel displays. The fibre-channel ports of the node are disabled to prevent disruption of the fibre-channel network. One or both nodes with the same WWNN can show the error. Because of the way WWNNs are allocated, a device with a duplicate WWNN is normally another cluster node.

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

1. Find the cluster node with the same WWNN as the node reporting the error. The WWNN for a cluster node can be found from the node Vital Product Data (VPD) or from the Node 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 WWNN of the node is stored within the service controller, so the duplication is most likely caused by the replacement of a service controller.

2. If a cluster node with a duplicate WWNN is found, determine whether it, or the node reporting the error, has the incorrect WWNN. Generally, it is the node that has had its service controller that was recently replaced or had its WWNN changed incorrectly. Also consider how the SAN is zoned when making your decision.
3. Determine the correct WWNN for the node with the incorrect WWNN. If the 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.
4. Use the 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 enclosure, access to the volumes that it is managing will be lost. You should ensure that the host systems are in the correct state before you change the WWNN.
5. If the node showing the error had the correct WWNN, it can be restarted, using the the front panel power control button, after the node with the duplicate WWNN is updated.
6. If you are unable to find a cluster node with the same WWNN as the node showing the error, use the SAN monitoring tools to determine whether there is another device on the SAN with the same WWNN. This device should not be using a WWNN assigned to a cluster, so you should follow the service procedures for the device to change its WWNN. Once the duplicate has been removed, restart the node canister.

Possible Cause-FRUs or other:

- None

558 The node is unable to communicate with other nodes.

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.

User response: 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

562 The nodes hardware configuration does not meet the minimum requirements.

Explanation: The node hardware is not at the minimum specification for the node to become active in a cluster. This may be because of hardware failure, but is also possible after a service action has used an incorrect replacement part.

User response: Follow troubleshooting procedures to fix the hardware:

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

564 Too many software crashes have occurred.

Explanation: The node has been determined to be unstable because of multiple resets. The cause of the resets can be that the system encountered an unexpected state or has executed instructions that were not valid. The node has entered the service state so that diagnostic data can be recovered.

The node error does not persist across restarts of the node software and operating system.

User response: Follow troubleshooting procedures to reload the software:

1. Get a support package (snap), including dumps, from the node using the management GUI or the service assistant.
2. If more than one node is reporting this error, contact IBM technical support for assistance. The support package from each node will be required.
3. Check the support site to see whether the issue is known and whether a software upgrade exists to resolve the issue. Update the cluster software if a resolution is available. Use the manual upgrade process on the node that reported the error first.
4. If the problem remains unresolved, contact IBM technical support and send them the support package.

Possible Cause-FRUs or other:

- None

565 The internal drive of the node is failing.

Explanation: The internal drive within the node is reporting too many errors. It is no longer safe to rely on the integrity of the drive. Replacement is recommended.

User response: Follow troubleshooting procedures to fix the hardware:

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

Possible Cause-FRUs or other:

- 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%)

574 The node software is damaged.

Explanation: A checksum failure has indicated that the node software is damaged and needs to be reinstalled.

User response: If the other node canister is operational, run node rescue. Otherwise, install new software using the service assistant. Node rescue failures or the repeated return of this node error after reinstallation is symptomatic of a hardware fault with the node canister.

Possible Cause-FRUs or other:

- None

576 The cluster state and configuration data cannot be read.

Explanation: The node has been unable to read the saved cluster state and configuration data from its internal drive because of a read or medium error.

User response: 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 The state data was not saved following a power loss.

Explanation: On startup, the node was unable to read its state data. When this happens, it expects to be automatically added back into a cluster. However, if it has not joined a cluster in 60 sec, it raises this node error. This is a critical node error and user action is required before the node can become a candidate to join a cluster.

User response: Follow troubleshooting procedures to correct connectivity issues between the cluster nodes and the quorum devices.

1. Manual intervention is required once the node reports this error.
2. Attempt to reestablish the cluster using other nodes. This may involve fixing hardware issues on other nodes or fixing connectivity issues between nodes.
3. If you are able to reestablish the cluster, remove the cluster data from the node showing 578 so it goes to candidate state, it will then be automatically added back to the cluster. If the node does not automatically add back to the cluster, note the name and I/O group of the node, then delete the node from the cluster configuration (if this has not already happened) and then add the node back to the cluster using the same name and I/O group.
4. If all nodes have either node error 578 or 550, follow the cluster recovery procedures.
5. Attempt to determine what caused the nodes to shut down.

Possible Cause-FRUs or other:

- None

580 The service controller ID could not be read.

Explanation: The 2145 cannot read the unique ID from the service controller, so the fibre-channel adapters cannot be started.

User response: In the sequence shown, exchange the following FRUs for new FRUs.

Possible Cause-FRUs or other:

2145-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

581 A serial link error in the 2145 UPS-1U has occurred.

Explanation: There is a fault in the communications cable, the serial interface in the uninterruptible power supply 2145 UPS-1U, or 2145.

User response: Check that the communications cable is correctly plugged 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%)

582 A battery error in the 2145 UPS-1U has occurred.

Explanation: A problem has occurred with the uninterruptible power supply 2145 UPS-1U battery.

User response: Exchange the FRU for a new FRU. After replacing the battery assembly, if the 2145 UPS-1U service indicator is on, press and hold the 2145

UPS-1U Test button for three seconds to start the self-test and verify the repair.

Possible Cause-FRUs or other:

- UPS-1U battery assembly (50%)
- UPS-1U assembly (50%)

583 An electronics error in the 2145 UPS-1U has occurred.

Explanation: A problem has occurred with the 2145 UPS-1U electronics.

User response: Exchange the FRU for a new FRU.

Possible Cause-FRUs or other:

- 2145 UPS-1U assembly

584 The 2145 UPS-1U is overloaded.

Explanation: A problem with output overload has been reported by the uninterruptible power supply 2145 UPS-1U. The Overload Indicator on the 2145 UPS-1U front panel is illuminated red.

User response:

1. Ensure that only one 2145 is receiving power from the 2145 UPS-1U. Also ensure that no other devices are connected to the 2145 UPS-1U.
2. Disconnect the 2145 from the 2145 UPS-1U. If the Overload Indicator is still illuminated, on the disconnected 2145 replace the 2145 UPS-1U.
3. If the Overload Indicator is now off, and the node is a 2145-8F2, 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.
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 whether the error is reported. Then shut down the node and connect the other power cable into the left hand power supply and start the node and see whether the error is repeated. Then repeat the two tests for the right hand power supply. If the error is repeated for both cables on one power supply but not the other, replace the power supply that showed the error; otherwise, replace the power cable assembly.

Possible Cause-FRUs or other:

- Power cable assembly (45%)
- Power supply assembly (45%)
- UPS-1U assembly (10%)

586 **The power supply to the 2145 UPS-1U does not meet requirements.**

Explanation: None.

User response: Follow troubleshooting procedures to fix the hardware.

587 **An incorrect type of uninterruptible power supply has been detected.**

Explanation: An incorrect type of 2145 UPS-1U was installed.

User response: Exchange the 2145 UPS-1U for one of the correct type.

Possible Cause-FRUs or other:

- 2145 UPS-1U (100%)

588 **The 2145 UPS-1U is not cabled correctly.**

Explanation: The signal cable or the 2145 power cables are probably not connected correctly. The power cable and signal cable might be connected to different 2145 UPS-1U assemblies.

User response:

1. Connect the cables correctly.
2. Restart the node.

Possible Cause-FRUs or other:

- None.

Other:

- Cabling error (100%)

589 **The 2145 UPS-1U ambient temperature limit has been exceeded.**

Explanation: The ambient temperature threshold for the 2145 UPS-1U has been exceeded.

User response: Reduce the temperature around the system:

1. Turn off the 2145 UPS-1U and unplug it from the power source.
2. Clear the vents and remove any heat sources.
3. Ensure that the air flow around the 2145 UPS-1U is not restricted.
4. Wait at least five minutes, and then restart the 2145 UPS-1U. If the problem remains, exchange 2145 UPS-1U assembly.

590 **Repetitive node restarts have occurred because of errors from the 2145 UPS-1U.**

Explanation: Multiple node restarts have occurred because of 2145 UPS-1U errors.

User response: Follow troubleshooting procedures to fix the hardware:

1. Verify that the room temperature is within specified limits and that the input power is stable.
2. Verify that the 2145 UPS-1U signal cable is fastened securely at both ends.

Note: The condition will be reset by powering off the node from the node front panel.

670 **The UPS battery charge is not enough to allow the node to start.**

Explanation: The uninterruptible power supply connected to the node does not have sufficient battery charge for the node to safely become active in a cluster. The node will not start until a sufficient charge exists to store the state and configuration data held in the node memory if power were to fail. The front panel of the node will show "charging".

User response: Wait for sufficient battery charge for enclosure to start:

1. Wait for the node to automatically fix the error when there is sufficient charge.
2. Ensure that no error conditions are indicated on the uninterruptible power supply.

690 **The node is held in the service state.**

Explanation: The node is in service state and has been instructed to remain in service state. While in service state, the node will not run as part of a cluster. A node must not be in service state for longer than necessary while the cluster is online because a loss of redundancy will result. A node can be set to remain in service state either because of a service assistant user action or because the node was deleted from the cluster.

User response: When it is no longer necessary to hold the node in the service state, exit the service state to allow the node to run:

1. Use the service assistant action or use the front panel **Exit Service** action to release the service state.

Possible Cause-FRUs or other:

- None

801 **Memory reduced.**

Explanation: Memory is reduced but sufficient memory exists to run I/O operations.

User response: Follow troubleshooting procedures to fix the hardware.

803 **One or more fibre-channel ports are not operational.**

Explanation: One or more fibre-channel ports are not operational.

User response: Follow troubleshooting procedures to fix the hardware.

805 **One or more configured Ethernet ports are not operational.**

Explanation: One or more configured Ethernet ports are not operational.

User response: Follow troubleshooting procedures to fix the hardware.

815 **Cannot determine the VPD for a component.**

Explanation: An FRU in the system has been changed, and the VPD is unreadable or unrecognized.

User response:

1. Check whether the replacement part that you have installed is the correct part.
 2. See whether there is an updated software package that correctly supports the part that was used. If an updated software package exists, upgrade to that software version. Otherwise, obtain the correct replacement part for the enclosure model and software version that you are operating.
-

818 **Unable to recover the service controller flash disk.**

Explanation: Unable to recover the service controller flash disk.

User response: Follow troubleshooting procedures to fix the hardware.

860 **The fibre-channel network fabric is too large.**

Explanation: This is a non-critical node error. The node will continue to operate but only the first 1024 fibre-channel logins will be used. Connectivity problems to the controllers, hosts, or other nodes could exist.

User response: Fix the fibre-channel network configuration:

1. View hardware WWNN information.
2. Reconfigure your SAN zoning.

878 **Attempting recovery after loss of state data**

Explanation: During startup, the node was unable to read its state data. It expects to be added back into a cluster, and reports this error while it is waiting.

User response: Allow time for recovery. No further action is required.

Cluster recovery and states

920 **Unable to perform cluster recovery**

because of a lack of cluster resources.

Explanation: The node is looking for a quorum of resources which also require cluster recovery.

User response: Contact IBM technical support.

950 Special upgrade mode.

Explanation: Special upgrade mode.

Cluster error codes

1001 Automatic cluster recovery has run.

Explanation: All cluster configuration commands are blocked.

User response: Call your software support center.

Caution: You can unblock the configuration commands 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 Event log full.

Explanation: Event log full.

User response: To fix the errors in the event log, go to the start MAP.

Possible Cause-FRUs or other:

- Unfixed errors in the log.

1011 Fibre-channel adapter (4 port) in slot 1 is missing.

Explanation: Fibre-channel adapter (4 port) in slot 1 is missing.

User response:

1. In the sequence shown, exchange the FRUs for new FRUs.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired as "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem 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%)

User response: None.

990 Cluster recovery has failed.

Explanation: Cluster recovery has failed.

User response: Contact IBM technical support.

- System board (2%)

2145-8F4

N/A

2145-8F2

N/A

2145-4F2

N/A

1013 Fibre-channel adapter (4-port) in slot 1 PCI fault.

Explanation: Fibre-channel adapter (4-port) in slot 1 PCI fault.

User response:

1. In the sequence shown, exchange the FRUs for new FRUs.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired as "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem 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

1014 Fibre-channel adapter in slot 1 is missing.

Explanation: Fibre-channel adapter in slot 1 is missing.

User response:

1. In the sequence shown, exchange the FRUs for new FRUs.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "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

1015 Fibre-channel adapter in slot 2 is missing.

Explanation: Fibre-channel adapter in slot 2 is missing.

User response:

1. In the sequence shown, exchange the FRUs for new FRUs.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "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

1016 Fibre-channel adapter (4 port) in slot 2 is missing.

Explanation: Fibre-channel adapter (4 port) in slot 2 is missing.

User response:

1. In the sequence shown, exchange the FRUs for new FRUs.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "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

1017 Fibre-channel adapter in slot 1 PCI bus error.

Explanation: Fibre-channel adapter in slot 1 PCI bus error.

User response:

1. In the sequence shown, exchange the FRUs for new FRUs.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "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

1018 Fibre-channel adapter in slot 2 PCI fault.

Explanation: Fibre-channel adapter in slot 2 PCI fault.

User response:

1. In the sequence shown, exchange the FRUs for new FRUs.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "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

1019 Fibre-channel adapter (4-port) in slot 2 PCI fault.

Explanation: Fibre-channel adapter (4-port) in slot 2 PCI fault.

User response:

1. In the sequence shown, exchange the FRUs for new FRUs.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "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

1020 The system board service processor has failed.

Explanation: The cluster is reporting that a node is not operational because of critical node error 522. See the details of node error 522 for more information.

User response: See node error 522.

1021 Incorrect enclosure

Explanation: The cluster is reporting that a node is not operational because of critical node error 500. See the details of node error 500 for more information.

User response: See node error 500.

1022 The detected memory size does not match the expected memory size.

Explanation: The cluster is reporting that a node is not operational because of critical node error 510. See the details of node error 510 for more information.

User response: See node error 510.

1025 The 2145 system assembly is failing.

Explanation: The 2145 system assembly is failing.

User response:

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

1027 Unable to update BIOS settings.

Explanation: The cluster is reporting that a node is not operational because of critical node error 524. See the details of node error 524 for more information.

User response: See node error 524.

1030 The internal disk of a node has failed.

Explanation: An error has occurred while attempting to read or write data to the internal disk of one of the nodes in the cluster. The disk has failed.

User response: Determine which node's internal disk has failed using the node information in the error. Replace the FRUs in the order shown. Mark the error as fixed.

Possible Cause-FRUs or other:

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%)
-

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.

User response:

1. Replace the FRUs below in the order listed
2. Check node status. If all nodes show a status of Online, mark the error that you have just repaired “fixed”. If any nodes do not show a status of Online, go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145.

3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-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%)

1044 A service controller read failure occurred.

Explanation: A service controller read failure occurred.

User response:

1. Replace the FRUs below in the order listed.
2. Check node status. If all nodes show a status of Online, mark the error that you have just repaired "fixed". If any nodes do not show a status of Online, go to start MAP. If you return to this step, contact your support center to resolve the problem 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%)

1052 Incorrect type of uninterruptible power supply detected

Explanation: The cluster is reporting that a node is not operational because of critical node error 587. See the details of node error 587 for more information.

User response: See node error 587.

1054 Fibre-channel adapter in slot 1 adapter present but failed.

Explanation: Fibre-channel adapter in slot 1 adapter present but failed.

User response:

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

1055 Fibre-channel adapter (4 port) in slot 1 adapter present but failed.

Explanation: Fibre-channel adapter (4 port) in slot 1 adapter present but failed.

User response:

1. Exchange the FRU for new FRU.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem 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.

User response:

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

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.

User response:

1. Exchange the FRU for new FRU.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem 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

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.

User response:

1. Go to MAP 5600: Fibre-channel to isolate and repair the problem.
2. Go to the repair verification MAP.

Possible Cause-FRUs or other:

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%)

1065 One or more fibre channel ports are running at lower than the previously saved speed.

Explanation: The fibre channel ports will normally operate at the highest speed permitted by the fibre-channel switch, but this speed might be reduced if the signal quality on the fibre-channel connection is

poor. The fibre-channel switch could have been set to operate at a lower speed by the user, or the quality of the fibre-channel signal has deteriorated.

User response:

- Go to MAP 5600: Fibre Channel to resolve the problem.

Possible Cause-FRUs or other:

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 connector, or GBIC (25%)

1083 Unrecognized node error

Explanation: The cluster is reporting that a node is not operational because of critical node error 562. See the details of node error 562 for more information.

User response: See node error 562.

1089 One or more fans are failing.

Explanation: One or more fans are failing.

User response:

1. Determine the failing fan(s) from the fan indicator on the system board or from the text of the error data in the log. The reported fan for the 2145-8A4 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.
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

1090 One or more fans (40x40x28) are failing.

Explanation: One or more fans (40x40x28) are failing.

User response:

1. Determine the failing fan(s) from the fan indicator on the system board or from the text of the error data in the log.
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.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F2 or 2145-8F4

- Fan 40x40x28 (98%)
- Fan power cable assembly (2%)

2145-4F2 or 2145-8G4

N/A

1091 One or more fans (40x40x56) are failing.

Explanation: One or more fans (40x40x56) are failing.

User response:

1. Determine the failing fan(s) from the fan indicator on the system board or from the text of the error data in the log.
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.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145-8F2 or 2145-8F4

- Fan 40x40x56 (98%)
- Fan power cable assembly (2%)

2145-4F2 or 2145-8G4

N/A

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.

User response:

1. Ensure that the operating environment meets specifications.
2. Ensure that the airflow is not obstructed.
3. Ensure that the fans are operational.
4. Go to the light path diagnostic MAP and perform the light path diagnostic procedures.
5. Check node status. If all nodes show a status of "online", mark the error that you have just repaired as "fixed". If any nodes do not show a status of "online", go to the start MAP. If you return to this step, contact your support center to resolve the problem 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%)

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%)

1094 The ambient temperature threshold has been exceeded.

Explanation: The ambient temperature threshold has been exceeded.

User response:

1. Check that the room temperature is within the limits allowed.
2. Check for obstructions in the air flow.
3. Mark the errors as fixed.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

None

Other:

System environment (100%)

1096 A Power Supply Unit is missing or has failed.

Explanation: One of the two power supply units in the node is either missing or has failed.

NOTE: This error is reported when a hot-swap power supply is removed from an active node, so it might be reported when a faulty power supply is removed for replacement. Both the missing and faulty conditions report this error code.

User response: Error code 1096 is reported when the power supply either cannot be detected or reports an error.

1. Ensure that the power supply is seated correctly and that the power cable is attached correctly to both the node and to the 2145 UPS-1U.
2. If the error has not been automatically marked fixed after two minutes, note the status of the three LEDs on the back of the power supply. For the 2145-CF8,

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.

User response:

1. Ensure that the internal airflow of the node has not been obstructed.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired 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.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

the AC LED is the top green LED, the DC LED is the middle green LED and the error LED is the bottom amber LED.

3. If the power supply error LED is off and the AC and DC power LEDs are both on, this is the normal condition. If the error has not been automatically fixed after two minutes, replace the system board.
4. Follow the action specified for the LED states noted in the table below.
5. If the error has not been automatically fixed after two minutes, contact support.
6. Go to repair verification MAP.

Error,AC,DC:Action

ON,ON or OFF,ON or OFF:The power supply has a fault. Replace the power supply.

OFF,OFF,OFF:There is no power detected. Ensure that the power cable is connected at the node and 2145 UPS-1U. If the AC LED does not light, check the status of the 2145 UPS-1U to which the power supply is connected. Follow MAP 5150 2145 UPS-1U if the UPS-1U is showing no power or an error; otherwise, replace the power cable. If the AC LED still does not light, replace the power supply.

OFF,OFF,ON:The power supply has a fault. Replace the power supply.

OFF,ON,OFF:Ensure that the power supply is installed correctly. If the DC LED does not light, replace the power supply.

Possible Cause-FRUs or other:

Failed PSU:

- Power supply (90%)
- Power cable assembly (5%)
- System board (5%)

Missing PSU:

- Power supply (19%)
- System board (1%)
- Other: Power supply not correctly installed (80%)

1097 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.

User response:

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%)

1100 One of the voltages that is monitored on the system board is over the set threshold.

Explanation: One of the voltages that is monitored on the system board is over the set threshold.

User response:

1. See the light path diagnostic MAP.
2. If the light path diagnostic MAP does not resolve the issue, exchange the frame assembly.
3. Check node status. If all nodes show a status of "online", mark the error that you have just repaired as "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem 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

1101 **One of the voltages that is monitored on the system board is over the set threshold.**

Explanation: One of the voltages that is monitored on the system board is over the set threshold.

User response:

1. See the light path diagnostic MAP.
2. If the light path diagnostic MAP does not resolve the issue, exchange the system board assembly.
3. Check node status. If all nodes show a status of "online", mark the error that you have just repaired as "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem 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%)

1105 **One of the voltages that is monitored on the system board is under the set threshold.**

Explanation: One of the voltages that is monitored on the system board is under the set threshold.

User response:

1. Check the cable connections.
2. See the light path diagnostic MAP.
3. If the light path diagnostic MAP does not resolve the issue, exchange the frame assembly.
4. Check node status. If all nodes show a status of "online", mark the error that you have just repaired as "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem 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

1106 **One of the voltages that is monitored on the system board is under the set threshold.**

Explanation: One of the voltages that is monitored on the system board is under the set threshold.

User response:

1. Check the cable connections.
2. See the light path diagnostic MAP.
3. If the light path diagnostic MAP does not resolve the issue, exchange the system board assembly.
4. Check node status. If all nodes show a status of "online", mark the error that you have just repaired as "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem 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%)

1110 **The power management board detected a voltage that is outside of the set thresholds.**

Explanation: The power management board detected a voltage that is outside of the set thresholds.

User response:

1. In the sequence shown, exchange the FRUs for new FRUs.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired as "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem 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

1120 A high speed SAS adapter is missing.

Explanation: This node has detected that a high speed SAS adapter that was previously installed is no longer present.

User response: If the high speed SAS adapter was deliberately removed, mark the error “fixed.”

Otherwise, the high speed SAS adapter has failed and must be replaced. In the sequence shown, exchange the FRUs for new FRUs.

Go to the repair verification MAP.

Possible Cause-FRUs or other:

1. High speed SAS adapter (90%)
2. System board (10%)

1121 A high speed SAS adapter has failed.

Explanation: A fault has been detected on a high speed SAS adapter.

User response: In the sequence shown, exchange the FRUs for new FRUs.

Go to the repair verification MAP.

Possible Cause-FRUs or other:

1. High speed SAS adapter (90%)
2. System board (10%)

1122 A high speed SAS adapter error has occurred.

Explanation: The high speed SAS adapter has detected a PCI bus error and requires service before it can be restarted. The high speed SAS adapter failure has caused all of the solid-state drives that were being accessed through this adapter to go Offline.

User response: If this is the first time that this error has occurred on this node, complete the following steps:

1. Power off the node.
2. Reseat the high speed SAS adapter card.
3. Power on the node.

4. Submit the **svcinfolsmdisk** 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, complete the following steps:

1. In the sequence shown, exchange the FRUs for new FRUs.
2. Submit the **svcinfolsmdisk** 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%)

1133 A duplicate WWNN has been detected.

Explanation: The cluster is reporting that a node is not operational because of critical node error 556. See the details of node error 556 for more information.

User response: See node error 556.

1135 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.

User response:

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.
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%)

1136 The 2145 UPS-1U has reported an ambient over temperature.

Explanation: The 2145 UPS-1U has reported an ambient over temperature.

User response:

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.
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%)

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.

User response:

1. Check the input AC power, whether it is missing or out of specification. Correct if necessary. Otherwise, exchange the FRU for a new FRU.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- 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%)

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.

User response:

1. Check the input AC power, whether it is missing or out of specification. Correct if necessary. Otherwise, exchange the FRU for a new FRU.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- 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%)

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.

User response:

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.
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%)

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.

User response:

1. Exchange the FRUs for new FRUs in the sequence shown.
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%)

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.

User response:

1. Connect the cables correctly. See your product's installation guide.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of

"online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.

3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Configuration error

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.

User response:

1. Connect the cables correctly. See your product's installation guide.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Configuration error

1152 Incorrect type of uninterruptible power supply detected.

Explanation: The cluster is reporting that a node is not operational because of critical node error 587. See the details of node error 587 for more information.

User response: See node error 587.

1155 A power domain error has occurred.

Explanation: Both 2145s of a pair are powered by the same uninterruptible power supply.

User response:

1. List the 2145s of the cluster and check that 2145s in the same I/O group are connected to a different uninterruptible power supply.
2. Connect one of the 2145s as identified in step 1 to a different uninterruptible power supply.
3. Mark the error that you have just repaired, "fixed".
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Configuration error

1160 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.

User response:

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.
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%)

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).

User response:

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. 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%)

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.

User response:

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

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.

User response:

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%)

1170 2145 UPS electronics fault (reported by the 2145 UPS alarm bits).

Explanation: 2145 UPS electronics fault (reported by the 2145 UPS alarm bits).

User response:

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%)

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).

User response:

1. Replace the uninterruptible power supply assembly.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145 UPS-1U.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145 UPS-1U assembly (100%)

1175 A problem has occurred with the uninterruptible power supply frame fault (reported by uninterruptible power supply alarm bits).

Explanation: A problem has occurred with the uninterruptible power supply frame fault (reported by the uninterruptible power supply alarm bits).

User response:

1. Replace the uninterruptible power supply assembly.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired

"fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.

3. Go to repair verification MAP.

Possible Cause-FRUs or other:

Uninterruptible power supply assembly (100%)

1180 2145 UPS battery fault (reported by 2145 UPS alarm bits).

Explanation: 2145 UPS battery fault (reported by 2145 UPS alarm bits).

User response:

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%)

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).

User response:

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%)

1182 Ambient temperature is too high during system startup.

Explanation: The cluster is reporting that a node is not operational because of critical node error 528. See the details of node error 528 for more information.

User response: See node error 528.

1183 The nodes hardware configuration does not meet the minimum requirements.

Explanation: The cluster is reporting that a node is not operational because of critical node error 562. See the details of node error 562 for more information.

User response: See node error 562.

1185 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).

User response:

1. In the sequence shown, exchange the FRU for a new FRU.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem 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%)
-

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).

User response:

1. In the sequence shown, exchange the FRU for a new FRU.
2. Check node status. If all nodes show a status of "online", mark the error that you have just repaired "fixed". If any nodes do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the uninterruptible power supply.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

2145 UPS-1U assembly (100%)

1187 Node software is inconsistent or damaged

Explanation: The cluster is reporting that a node is not operational because of critical node error 523,573-574. See the details of node error 523,573-574 for more information.

User response: See node error 523,573-574.

1188 Too many software crashes have occurred.

Explanation: The cluster is reporting that a node is not operational because of critical node error 564. See the details of node error 564 for more information.

User response: See node error 564.

1189 The node is held in the service state.

Explanation: The cluster is reporting that a node is not operational because of critical node error 690. See the details of node error 690 for more information.

User response: See node error 690.

1190 The 2145 UPS battery has reached its end of life.

Explanation: The 2145 UPS battery has reached its end of life.

User response:

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%)

1191 The 2145 UPS-1U battery has reached its end of life.

Explanation: The 2145 UPS-1U battery has reached its end of life.

User response:

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%)

1192 Unexpected node error

Explanation: A node is missing from the cluster. The error that it is reporting is not recognized by the system.

User response: Find the node that is in service state and use the service assistant to determine why it is not active.

1193 The UPS battery charge is not enough to allow the node to start.

Explanation: The cluster is reporting that a node is not operational because of critical node error 587. See the details of node error 587 for more information.

User response:

1194 Automatic recovery of offline node has failed.

Explanation: The cluster has an offline node and has determined that one of the candidate nodes matches the characteristics of the offline node. The cluster has attempted but failed to add the node back into the cluster. The cluster has stopped attempting to automatically add the node back into the cluster.

If a node has incomplete state data, it remains offline after it starts. This occurs if the node has had a loss of power or a hardware failure that prevented it from completing the writing of all of the state data to disk. The node reports a node error 578 when it is in this state.

If three attempts to automatically add a matching candidate node to a cluster have been made, but the node has not returned online for 24 hours, the cluster stops automatic attempts to add the node and logs error code 1194 "Automatic recovery of offline node failed".

Two possible scenarios when this error event is logged are:

1. The node has failed without saving all of its state data. The node has restarted, possibly after a repair, and shows node error 578 and is a candidate node for joining the cluster. The cluster attempts to add the node into the cluster but does not succeed. After 15 minutes, the cluster makes a second attempt to add the node into the cluster and again does not succeed. After another 15 minutes, the cluster makes a third attempt to add the node into the cluster and again does not succeed. After another 15

minutes, the cluster logs error code 1194. The node never came online during the attempts to add it to the cluster.

2. The node has failed without saving all of its state data. The node has restarted, possibly after a repair, and shows node error 578 and is a candidate node for joining the cluster. The cluster attempts to add the node into the cluster and succeeds and the node becomes online. Within 24 hours the node fails again without saving its state data. The node restarts and shows node error 578 and is a candidate node for joining the cluster. The cluster again attempts to add the node into the cluster, succeeds, and the node becomes online; however, the node again fails within 24 hours. The cluster attempts a third time to add the node into the cluster, succeeds, and the node becomes online; however, the node again fails within 24 hours. After another 15 minutes, the cluster logs error code 1194.

A combination of these scenarios is also possible.

Note: If the node is manually removed from the cluster, the count of automatic recovery attempts is reset to zero.

User response:

1. If the node has been continuously online in the cluster for more than 24 hours, mark the error as fixed and go to the Repair Verification MAP.
2. Determine the history of events for this node by locating events for this node name in the event log. Note that the node ID will change, so match on the WWNN and node name. Also, check the service records. Specifically, note entries indicating one of three events: 1) the node is missing from the cluster (cluster error 1195 event 009052), 2) an attempt to automatically recover the offline node is starting (event 980352), 3) the node has been added to the cluster (event 980349).
3. If the node has not been added to the cluster since the recovery process started, there is probably a hardware problem. The node's internal disk might be failing in a manner that it is unable to modify its software level to match the software level of the cluster. If you have not yet determined the root cause of the problem, you can attempt to manually remove the node from the cluster and add the node back into the cluster. Continuously monitor the status of the nodes in the cluster while the cluster is attempting to add the node. Note: If the node type is not supported by the software version of the cluster, the node will not appear as a candidate node. Therefore, incompatible hardware is not a potential root cause of this error.
4. If the node was added to the cluster but failed again before it has been online for 24 hours, investigate the root cause of the failure. If no events

in the event log indicate the reason for the node failure, collect dumps and contact IBM technical support for assistance.

5. When you have fixed the problem with the node, you must use either the cluster console or the command line interface to manually remove the node from the cluster and add the node into the cluster.
6. Mark the error as fixed and go to the verification MAP.

Possible Cause-FRUs or other:

None, although investigation might indicate a hardware failure.

1195 A 2145 is missing from the cluster.

Explanation: You can resolve this problem by repairing the failure on the missing 2145.

User response:

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

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.

User response:

1. Remove unwanted devices from the fibre-channel network fabric.
2. Start a cluster discovery operation to find devices/disks by rescanning the fibre-channel network.
3. List all connected managed disks. Check with the customer that the configuration is as expected. Mark the error that you have just repaired fixed.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

Fibre-channel network fabric fault (100%)

1201 A 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.

User response: The following actions delete all the volumes and host maps for the storage pool that contains the offline managed disk. Therefore, it is strongly recommended that you back up the cluster 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 volume copies, including unmirrored volumes, that use storage from the storage pool that contains the offline managed disk. Record all of the properties of the storage pool and the list of all of the managed disks in that storage pool.
2. Delete the storage pool 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 storage pool with the same properties as the deleted storage pool and add the drive back into the storage pool.
5. Adding the disk into the storage pool 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 storage pool.

7. Add all of the other managed disks that were in the deleted storage pool into the new storage pool.
8. Add a volume copy for all of the volumes that were mirrored in the original storage pool.
9. Create all of the standard volumes that were members of the original storage pool 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 volume.
2. Delete the image mode volume 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, storage pool and add the MDisk into this new storage pool.
5. Adding the disk into the storage pool 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 storage pool.
7. Delete the temporary storage pool that was created in step 4.
8. Either re-create the image mode volume or add the image mode volume copy, depending on the original configuration of the volume.
9. NOTE: If the original configuration of the volume was an image mode volume, ALL of the data has been lost and must be re-created.

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 cluster event log will identify a managed disk ID. This can be determined by running the maintenance procedure on the 1202 error.

User response: Determine whether the managed disk

referred to in the error is currently offline using the graphical user interface (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 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%)

1203 A duplicate fibre channel frame has been received.

Explanation: A duplicate fibre channel frame should never be detected. Receiving a duplicate fibre channel frame indicates that there is a problem with the fibre channel fabric. Other errors related to the fibre channel fabric might be generated.

User response:

1. Use the transmitting and receiving WWPNs indicated in the error data to determine the section of the fibre channel fabric that has generated the duplicate frame. Search for the cause of the problem by using fabric monitoring tools. The duplicate frame might be caused by a design error in the topology of the fabric, by a configuration error, or by a software or hardware fault in one of the components of the fibre channel fabric, including inter-switch links.
2. When you are satisfied that the problem has been corrected, mark the error that you have just repaired “fixed”.
3. Go to MAP 5700: Repair verification.

Possible Cause-FRUs or other:

- Fibre-channel cable assembly (1%)
- Fibre-channel adapter (1%)

Other:

- Fibre-channel network fabric fault (98%)

1210 A local fibre-channel port has been excluded.

Explanation: A local fibre-channel port has been excluded.

User response:

1. Repair faults in the order shown.
2. Check the status of the disk controllers. If all disk controllers show a “good” status, mark the error that you just repaired as “fixed”.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- Fibre-channel cable assembly (75%)
- Small Form-factor Pluggable (SFP) connector (10%)
- Fibre-channel adapter (5%)

Other:

- Fibre-channel network fabric fault (10%)

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 cluster event log will identify a managed disk ID for the solid-state drive that caused the error.

User response: 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 rmmddisk -force (mdisk name/id)**, where (mdisk name/id) is the name or ID of the managed disk identified in the event log. This command migrates all of the data from the failing managed disk into the free extents in the rest of the storage pool. If the command fails with a message that indicates that there are not enough free extents, create more free extents in the storage pool 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 volume copies that exist in this storage pool. The second option is to migrate some of the volume copies into other storage pools. The third option is to temporarily add more managed disks to the storage pool.
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 event 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 storage pool.
8. Reverse any procedures that might have been performed in step 1 to create free extents in the storage pool.
9. This step is optional. Run the 'balance.pl' script that is available from the alphaworks package called SVCTools to redistribute the volume extents equally across all of the managed disks in the storage pool.

Possible Cause-FRUs or other:

- solid-state drive (100%)

1216 SAS errors have exceeded thresholds.

Explanation: The cluster has experienced a large number of SAS communication errors, which indicates a faulty SAS component that must be replaced.

User response: In the sequence shown, exchange the FRUs for new FRUs.

Go to the repair verification MAP.

Possible Cause-FRUs or other:

1. SAS Cable (70%)
2. High speed SAS adapter (20%)
3. SAS drive backplane (5%)
4. solid-state drive (5%)

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.

User response: Take steps to reduce the temperature of the drive.

1. Determine the temperature of the room, and reduce the room temperature if this action is appropriate.
2. Replace any failed fans.
3. Ensure that there are no obstructions to air flow for the node.
4. Mark the error as fixed. If the error recurs, contact hardware support for further investigation.

Possible Cause-FRUs or other:

- 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.

User response:

1. View the event log. Note the MDisk ID associated with the error code.
2. From the MDisk, determine the failing disk controller ID.
3. Refer to the service documentation for the disk controller and the fibre-channel network to resolve the reported problem.
4. After the disk drive is repaired, start a cluster discovery operation to recover the excluded fibre-channel port by rescanning the fibre-channel network.
5. To restore MDisk online status, include the managed disk that you noted in step 1.
6. Check the status of the disk controller. If all disk controllers show a "good" status, mark the error that you have just repaired, "fixed".
7. If all disk controllers do not show a good status, contact your support center to resolve the problem with the disk controller.
8. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Enclosure/controller fault (50%)
- Fibre-channel network fabric (50%)

1230 A login has been excluded.

Explanation: A port to port fabric connection, or login, between the cluster node and either a controller or another cluster has had excessive errors. The login has therefore been excluded, and will not be used for I/O operations.

User response: Determine the remote system, which might be either a controller or a SAN Volume Controller cluster. Check the event log for other 1230 errors. Ensure that all higher priority errors are fixed.

This error event is usually caused by a fabric problem. If possible, use the fabric switch or other fabric diagnostic tools to determine which link or port is reporting the errors. If there are error events for links from this node to a number of different controllers or clusters, then it is probably the node to switch link that is causing the errors. Unless there are other contrary

indications, first replace the cable between the switch and the remote system.

1. From the fabric analysis, determine the FRU that is most likely causing the error. If this FRU has recently been replaced while resolving a 1230 error, choose the next most likely FRU that has not been replaced recently. Exchange the FRU for a new FRU.
2. Mark the error as fixed. If the FRU replacement has not fixed the problem, the error will be logged again; however, depending on the severity of the problem, the error might not be logged again immediately.
3. Start a cluster discovery operation to recover the login by re-scanning the fibre-channel network.
4. Check the status of the disk controller or remote cluster. If the status is not “good”, go to the Start MAP.
5. Go to repair verification MAP.

Possible Cause-FRUs or other:

- Fibre-channel cable, switch to remote port, (30%)
- Switch or remote device SFP connector or adapter, (30%)
- Fibre-channel cable, local port to switch, (30%)
- Cluster SFP connector, (9%)
- Cluster fibre channel adapter, (1%)

Note: The first two FRUs are not cluster FRUs.

1310 A managed disk is reporting excessive errors.

Explanation: A managed disk is reporting excessive errors.

User response:

1. Repair the enclosure/controller fault.
2. Check the managed disk status. If all managed disks show a status of “online”, mark the error that you have just repaired as “fixed”. If any managed disks show a status of “excluded”, include the excluded managed disks and then mark the error as “fixed”.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

Enclosure/controller fault (100%)

1311 A solid-state drive is reporting excessive errors.

Explanation: The drive that is reporting excessive errors has been taken offline.

User response: 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 event 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 event 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.

User response:

1. Check whether the volume the error is reported against is mirrored. If it is, check if there is a “1870 Mirrored volume offline because a hardware read error has occurred” error relating to this volume in the event log. Also check if one of the mirror copies is synchronizing. If all these tests are true then you must delete the volume copy that is not synchronized from the volume. Check that the volume is online before continuing with the following actions. Wait until the medium error is corrected before trying to re-create the volume mirror.
2. If the medium error was detected by a read from a host, ask the customer to rewrite the incorrect data to the block logical block address (LBA) that is reported in the host systems SCSI sense data. If an individual block cannot be recovered it will be

necessary to restore the volume from backup. (If this error has occurred during a migration, the host system does not notice the error until the target device is accessed.)

3. If the medium error was detected during a mirrored volume synchronization, the block might not be being used for host data. The medium error must still be corrected before the mirror can be established. It may be possible to fix the block that is in error using the disk controller or host tools. Otherwise, it will be necessary to use the host tools to copy the volume content that is being used to a new volume. Depending on the circumstances, this new volume can be kept and mirrored, or the original volume can be repaired and the data copied back again.
4. Check managed disk status. If all managed disks show a status of "online", mark the error that you have just repaired as "fixed". If any managed disks do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the disk controller.
5. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

Enclosure/controller fault (100%)

1330 A suitable managed disk (MDisk) or drive for use as a quorum disk was not found.

Explanation: A quorum disk is needed to enable a tie-break when some cluster members are missing. Three quorum disks are usually defined. By default, the cluster automatically allocates quorum disks when managed disks are created; however, the option exists to manually assign quorum disks. This error is reported when there are managed disks or image mode disks but no quorum disks.

To become a quorum disk:

- The MDisk must be accessible by all nodes in the cluster.
- The MDisk must be managed; that is, it must be a member of a storage pool.
- The MDisk must have free extents.
- The MDisk must be associated with a controller that is enabled for quorum support. If the controller has multiple WWNNs, all of the controller components must be enabled for quorum support.

A quorum disk might not be available because of a fibre-channel network failure or because of a fibre-channel switch zoning problem.

User response:

1. Resolve any known fibre-channel network problems.
2. Ask the customer to confirm that MDisks have been added to storage pools and that those MDisks have free extents and are on a controller that is enabled for use as a provider of quorum disks. Ensure that any controller with multiple WWNNs has all of its components enabled to provide quorum disks. Either create a suitable MDisk or if possible enable quorum support on controllers with which existing MDisks are associated. If at least one managed disk shows a mode of managed and has a non-zero quorum index, mark the error that you have just repaired as "fixed".
3. If the customer is unable to make the appropriate changes, ask your software support center for assistance.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

Configuration error (100%)

1335 Quorum disk not available.

Explanation: Quorum disk not available.

User response:

1. View the event log entry to identify the managed disk (MDisk) being used as a quorum disk, that is no longer available.
2. Perform the disk controller problem determination and repair procedures for the MDisk identified in step 1.
3. Include the MDisks into the cluster.
4. Check the managed disk status. If the managed disk identified in step 1 shows a status of "online", mark the error that you have just repaired as "fixed". If the managed disk does not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the disk controller.
5. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

Enclosure/controller fault (100%)

1340 A managed disk has timed out.

Explanation: This error was reported because a large number of disk timeout conditions have been detected. The problem is probably caused by a failure of some other component on the SAN.

User response:

1. Repair problems on all enclosures/controllers and switches on the same SAN as this 2145 cluster.
2. If problems are found, mark this error as “fixed”.
3. If no switch or disk controller failures can be found, take an event log dump and call your hardware support center.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Enclosure/controller fault
 - Fibre-channel switch
-

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.

User response:

1. View the event log entry to determine the node that logged the problem. Determine the 2145 node or controller that the problem was logged against.
2. Perform fibre-channel 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 event 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
-

1370 A managed disk error recovery procedure (ERP) has occurred.

Explanation: This error was reported because a large number of disk error recovery procedures have been performed by the disk controller. The problem is probably caused by a failure of some other component on the SAN.

User response:

1. View the event log entry and determine the managed disk that was being accessed when the problem was detected.
2. Perform the disk controller problem determination and repair procedures for the MDisk determined in step 1.
3. Perform problem determination and repair procedures for the fibre channel 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 event log dump. Call your hardware support center.
6. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Enclosure/controller fault
 - Fibre-channel switch
-

1400 The 2145 cannot detect an Ethernet connection.

Explanation: The 2145 cannot detect an Ethernet connection.

User response:

1. Go to the Ethernet MAP.
2. Go to the repair verification MAP.

Possible Cause-FRUs or other:

2145-8G4 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%)

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.

User response:

1. Check for incorrect switch zoning.
2. Repair the fault in the fibre-channel network fabric.
3. Check the status of the node ports. 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%)

1570 Quorum disk configured on controller that has quorum disabled

Explanation: This error can occur with a storage controller that can be accessed through multiple WWNNs and have a default setting of not allowing quorum disks. When these controllers are detected by a cluster, although multiple component controller definitions are created, the cluster recognizes that all of the component controllers belong to the same storage system. To enable the creation of a quorum disk on this storage system, all of the controller components must be configured to allow quorum.

A configuration change to the SAN, or to a storage system with multiple WWNNs, might result in the cluster discovering new component controllers for the storage system. These components will take the default setting for allowing quorum. This error is reported if there is a quorum disk associated with the controller and the default setting is not to allow quorum.

User response:

- Determine if there should be a quorum disk on this storage system. Ensure that the controller supports quorum before you allow quorum disks on any disk controller. You can check the support website 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 volume two copy disks were found to contain different data for the same logical block address (LBA). The validate option was used, so the repair process has halted.

Read operations to the LBAs that differ might return the data of either volume copy. Therefore it is important not to use the volume unless you are sure that the host applications will not read the LBAs that differ or can manage the different data that potentially can be returned.

User response: Perform one of the following actions:

- Continue the repair starting with the next LBA after the difference to see how many differences there are for the whole mirrored volume. This can help you decide which of the following actions to take.
- Choose a primary disk and run repair resynchronizing differences.
- Run a repair and create medium errors for differences.
- Restore all or part of the volume from a backup.
- Decide which disk has correct data, then delete the copy that is different and re-create it allowing it to be synchronized.

Then mark the error as “fixed”.

Possible Cause-FRUs or other:

- None

1610 There are too many copied media errors on a managed disk.

Explanation: The cluster maintains a virtual medium error table for each MDisk. This table is a list of logical block addresses on the managed disk that contain data that is not valid and cannot be read. The virtual medium error table has a fixed length. This error event indicates that the system has attempted to add an entry to the table, but the attempt has failed because the table is already full.

There are two circumstances that will cause an entry to be added to the virtual medium error table:

1. FlashCopy, data migration and mirrored volume synchronization operations copy data from one managed disk extent to another. If the source extent contains either a virtual medium error or the RAID controller reports a real medium error, the system creates a matching virtual medium error on the target extent.
2. The mirrored volume validate and repair process has the option to create virtual medium errors on sectors that do not match on all volume copies. Normally zero, or very few, differences are expected; however, if the copies have been marked as synchronized inappropriately, then a large number of virtual medium errors could be created.

User response: Ensure that all higher priority errors are fixed before you attempt to resolve this error.

Determine whether the excessive number of virtual medium errors occurred because of a mirrored disk validate and repair operation that created errors for differences, or whether the errors were created because of a copy operation. Follow the corresponding option shown below.

1. If the virtual medium errors occurred because of a mirrored disk validate and repair operation that created medium errors for differences, then also ensure that the volume copies had been fully synchronized prior to starting the operation. If the copies had been synchronized, there should be only a few virtual medium errors created by the validate and repair operation. In this case, it might be possible to rewrite only the data that was not consistent on the copies using the local data recovery process. If the copies had not been synchronized, it is likely that there are now a large number of medium errors on all of the volume copies. Even if the virtual medium errors are expected to be only for blocks that have never been written, it is important to clear the virtual medium errors to avoid inhibition of other operations. To recover the data for all of these virtual medium errors it is likely that the volume will have to be recovered from a backup using a process that rewrites all sectors of the volume.
2. If the virtual medium errors have been created by a copy operation, it is best practice to correct any medium errors on the source volume and to not propagate the medium errors to copies of the volume. Fixing higher priority errors in the event log would have corrected the medium error on the source volume. Once the medium errors have been fixed, you must run the copy operation again to clear the virtual medium errors from the target volume. It might be necessary to repeat a sequence of copy operations if copies have been made of already copied medium errors.

An alternative that does not address the root cause is to delete volumes on the target managed disk that have the virtual medium errors. This volume deletion

reduces the number of virtual medium error entries in the MDisk table. Migrating the volume to a different managed disk will also delete entries in the MDisk table, but will create more entries on the MDisk table of the MDisk to which the volume is migrated.

Possible Cause-FRUs or other:

- None

1620 **A storage pool is offline.**

Explanation: A storage pool is offline.

User response:

1. Repair the faults in the order shown.
2. Start a cluster discovery operation by rescanning the fibre-channel network.
3. Check managed disk (MDisk) status. If all MDisks show a status of "online", mark the error that you have just repaired as "fixed". If any MDisks do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the disk controller.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Fibre-channel network fabric fault (50%)
- Enclosure/controller fault (50%)

1623 **One or more MDisks on a controller are degraded.**

Explanation: At least one MDisk on a controller is degraded because the MDisk is not available through one or more nodes. The MDisk is available through at least one node. Access to data might be lost if another failure occurs.

In a correctly configured system, each node accesses all of the MDisks on a controller through all of the controller's ports.

This error is only logged once per controller. There might be more than one MDisk on this controller that has been configured incorrectly, but the error is only logged for one MDisk.

To prevent this error from being logged because of short-term fabric maintenance activities, this error condition must have existed for one hour before the error is logged.

User response:

1. Determine which MDisks are degraded. Look for MDisks with a path count lower than the number of nodes. Do not use only the MDisk status, since other errors can also cause degraded MDisks.

2. Ensure that the controller is zoned correctly with all of the nodes.
3. Ensure that the logical unit is mapped to all of the nodes.
4. Ensure that the logical unit is mapped to all of the nodes using the same LUN.
5. Run the console or CLI command to discover MDisks and ensure that the command completes.
6. Mark the error that you have just repaired as "fixed". When you mark the error as "fixed", the controller's MDisk availability is tested and the error will be logged again immediately if the error persists for any MDisks. It is possible that the new error will report a different MDisk.
7. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Fibre-channel network fabric fault (50%)
- Enclosure/controller fault (50%)

1624 Controller configuration has unsupported RDAC mode.

Explanation: The cluster has detected that an IBM DS series disk controller's configuration is not supported by the cluster. The disk controller is operating in RDAC mode. The disk controller might appear to be operating with the cluster; however, the configuration is unsupported because it is known to not work with the cluster.

User response:

1. Using the IBM DS series console, ensure that the host type is set to 'IBM TS SAN VCE' and that the AVT option is enabled. (The AVT and RDAC options are mutually exclusive).
2. Mark the error that you have just repaired as "fixed". If the problem has not been fixed it will be logged again; this could take a few minutes.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Enclosure/controller fault

1625 Incorrect disk controller configuration.

Explanation: While running an MDisk discovery, the cluster has detected that a disk controller's configuration is not supported by the cluster. The disk controller might appear to be operating with the cluster; however, the configuration detected can

potentially cause issues and should not be used. The unsupported configuration is shown in the event data.

User response:

1. Use the event data to determine changes required on the disk controller and reconfigure the disk controller to use a supported configuration.
2. Mark the error that you have just repaired as "fixed". If the problem has not been fixed it will be logged again by the managed disk discovery that automatically runs at this time; this could take a few minutes.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Enclosure/controller fault

1627 The cluster has insufficient redundancy in its controller connectivity.

Explanation: The cluster has detected that it does not have sufficient redundancy in its connections to the disk controllers. This means that another failure in the SAN could result in loss of access to the application data. The cluster SAN environment should have redundant connections to every disk controller. This redundancy allows for continued operation when there is a failure in one of the SAN components.

To provide recommended redundancy, a cluster should be configured so that:

- each node can access each disk controller through two or more different initiator ports on the node.
- each node can access each disk controller through two or more different controller target ports. **Note:** Some disk controllers only provide a single target port.
- each node can access each disk controller target port through at least one initiator port on the node.

If there are no higher-priority errors being reported, this error usually indicates a problem with the SAN design, a problem with the SAN zoning or a problem with the disk controller.

If there are unfixed higher-priority errors that relate to the SAN or to disk controllers, those errors should be fixed before resolving this error because they might indicate the reason for the lack of redundancy. Error codes that must be fixed first are:

- 1210 Local FC port excluded
- 1230 Login has been excluded

Note: This error can be reported if the required action, to rescan the fibre-channel network for new MDisks,

has not been performed after a deliberate reconfiguration of a disk controller or after SAN rezoning.

The 1627 error code is reported for a number of different error IDs. The error ID indicates the area where there is a lack of redundancy. The data reported in an event log entry indicates where the condition was found.

The meaning of the error IDs is shown below. For each error ID the most likely reason for the condition is given. If the problem is not found in the suggested areas, check the configuration and state of all of the SAN components (switches, controllers, disks, cables and cluster) to determine where there is a single point of failure.

010040 A disk controller is only accessible from a single node port.

- A node has detected that it only has a connection to the disk controller through exactly one initiator port, and more than one initiator port is operational.
- The error data indicates the device WWNN and the WWPN of the connected port.
- A zoning issue or a fibre-channel connection hardware fault might cause this condition.

010041 A disk controller is only accessible from a single port on the controller.

- A node has detected that it is only connected to exactly one target port on a disk controller, and more than one target port connection is expected.
- The error data indicates the WWPN of the disk controller port that is connected.
- A zoning issue or a fibre-channel connection hardware fault might cause this condition.

010042 Only a single port on a disk controller is accessible from every node in the cluster.

- Only a single port on a disk controller is accessible to every node when there are multiple ports on the controller that could be connected.
- The error data indicates the WWPN of the disk controller port that is connected.
- A zoning issue or a fibre-channel connection hardware fault might cause this condition.

010043 A disk controller is accessible through only half, or less, of the previously configured controller ports.

- Although there might still be multiple ports that are accessible on the disk controller, a hardware component of the controller might have failed or one of the SAN fabrics has failed such that the operational system configuration has been reduced to a single point of failure.

- The error data indicates a port on the disk controller that is still connected, and also lists controller ports that are expected but that are not connected.
- A disk controller issue, switch hardware issue, zoning issue or cable fault might cause this condition.

010044 A disk controller is not accessible from a node.

- A node has detected that it has no access to a disk controller. The controller is still accessible from the partner node in the I/O group, so its data is still accessible to the host applications.
- The error data indicates the WWPN of the missing disk controller.
- A zoning issue or a cabling error might cause this condition.

User response:

1. Check the error ID and data for a more detailed description of the error.
2. Determine if there has been an intentional change to the SAN zoning or to a disk controller configuration that reduces the cluster's access to the indicated disk controller. If either action has occurred, continue with step 8.
3. Use the GUI or the CLI command `svcinfolsfabric` 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

1630 The number of device logins was reduced.

Explanation: The number of port to port fabric connections, or logins, between the node and a storage controller has decreased. This might be caused by a problem on the SAN or by a deliberate reconfiguration of the SAN.

User response:

1. Check the error in the cluster event log to identify the object ID associated with the error.
2. Check the availability of the failing device using the following command line: **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 event log and continue with the repair verification MAP.
 - If "no" or if the command lists details of the failing controller, continue with the next step.
3. Check whether the device has regained connectivity. If it has not, check the cable connection to the remote-device port.
4. If all attempts to log in to a remote-device port have failed and you cannot solve the problem by changing cables, check the condition of the remote-device port and the condition of the remote device.
5. Start a cluster discovery operation by rescanning the fibre-channel network.
6. Check the status of the disk controller. If all disk controllers show a "good" status, mark the error that you have just repaired as "fixed". If any disk controllers do not show "good" status, go to 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%)

1660 The initialization of the managed disk has failed.

Explanation: The initialization of the managed disk has failed.

User response:

1. View the event log entry to identify the managed disk (MDisk) that was being accessed when the problem was detected.
2. Perform the disk controller problem determination and repair procedures for the MDisk identified in step 1.
3. Include the MDisk into the cluster.
4. Check the managed disk status. If all managed disks show a status of "online", mark the error that you have just repaired as "fixed". If any managed disks do not show a status of "online", go to the start MAP. If you return to this step, contact your support center to resolve the problem with the disk controller.
5. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

Enclosure/controller fault (100%)

1661 The solid-state drive initialization process has failed.

Explanation: The drive has not been initialized.

User response: If there was data on the drive, contact IBM technical support for assistance.

If the drive is currently in a storage pool 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 storage pool 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%)

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.

User response: 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 cluster, 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.

User response:

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%)

1695 Persistent unsupported disk controller configuration.

Explanation: A disk controller configuration that might prevent failover for the cluster has persisted for more than four hours. The problem was originally logged through a 010032 event, service error code 1625.

User response:

1. Fix any higher priority error. In particular, follow the service actions to fix the 1625 error indicated by this error's root event. This error will be marked as “fixed” when the root event is marked as “fixed”.
2. If the root event cannot be found, or is marked as “fixed”, perform an MDisk discovery and mark this error as “fixed”.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Enclosure/controller fault

1700 Unrecovered 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.

User response: To fix this error it is necessary to delete all of the relationships that could not be recovered and then re-create the relationships.

1. 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 volume in this I/O group. Use the volume

view to determine which volumes 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 re-created.
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, re-create 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 re-create 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 cluster prevents you from creating a new Metro Mirror and Global Mirror cluster partnership if a resulting partnership set would exceed the maximum of four clusters. However, if you restore a broken link between two clusters that have a partnership, the number of clusters in the set might exceed four. If this occurs, Metro Mirror and Global Mirror cluster partnerships are excluded from the set until only four clusters remain in the set. A cluster partnership that is excluded from a set has all of its Metro Mirror and Global Mirror cluster partnerships excluded.

Event ID 0x050030 is reported if the cluster is retained in the partnership set. Event ID 0x050031 is reported if the cluster is excluded from the partnership set. All clusters that were in the partnership set report error 1710.

All inter-cluster Metro Mirror or Global Mirror relationships that involve an excluded cluster will lose connectivity. If any of these relationships are in the consistent_synchronized state and they receive a write I/O, they will stop with error code 1720.

User response: To fix this error it is necessary to delete all of the relationships that could not be recovered and then re-create the relationships.

1. Determine which clusters are still connected and members of the partnership set, and which clusters have been excluded.
2. Determine the Metro Mirror and Global Mirror relationships that exist on those clusters.
3. Determine which of the Metro Mirror and Global Mirror relationships you want to maintain, which determines which cluster partnerships you want to maintain. Ensure that the partnership set or sets that would result from configuring the cluster partnerships that you want contain no more than four clusters in each set. NOTE: The reduced partnership set created by the cluster might not contain the clusters that you want in the set.
4. Remove all of the Metro Mirror and Global Mirror relationships that you do not want to retain.
5. Remove all of the Metro Mirror and Global Mirror cluster partnerships that you do not want to retain.
6. Restart all relationships and consistency groups that were stopped.
7. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

1720 **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.

User response:

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

1800 **The SAN has been zoned incorrectly.**

Explanation: This has resulted in more than 512 other ports on the SAN logging into one port of a 2145 node.

User response:

1. Ask the user to reconfigure the SAN.
2. Mark the error as “fixed”.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

- Fibre-channel switch configuration error
- Fibre-channel switch

1850 **A cluster recovery operation was performed but data on one or more volumes has not been recovered.**

Explanation: A cluster recovery operation was performed but data on one or more volumes has not been recovered.

User response:

1. The support center will direct the user to restore the data on the affected volumes.
2. When the volume data has been restored or the user has chosen not to restore the data, mark the error as “fixed”.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

1860 Thin-provisioned volume copy offline because of failed repair.

Explanation: The attempt to repair the metadata of a thin-provisioned volume that describes the disk contents has failed because of problems with the automatically maintained backup copy of this data. The error event data describes the problem.

User response: Delete the thin-provisioned volume and reconstruct a new one from a backup or mirror copy. Mark the error as “fixed”. Also mark the original 1862 error as “fixed”.

Possible Cause-FRUs or other:

- None

1862 Thin-provisioned volume copy offline because of corrupt metadata.

Explanation: A thin-provisioned volume has been taken offline because there is an inconsistency in the cluster metadata that describes the disk contents. This might occur because of corruption of data on the physical disk (e.g., medium error or data miscompare), the loss of cached metadata (because of a cluster recovery) or because of a software error. The event data gives information on the reason.

The cluster maintains backup copies of the metadata and it might be possible to repair the thin-provisioned volume using this data.

User response: The cluster is able to repair the inconsistency in some circumstances. Run the repair volume option to start the repair process. This repair process, however, can take some time. In some situations it might be more appropriate to delete the thin-provisioned volume and reconstruct a new one from a backup or mirror copy.

If you run the repair procedure and it completes, this error is automatically marked as “fixed”; otherwise, another error event (error code 1860) is logged to indicate that the repair action has failed.

Possible Cause-FRUs or other:

- None

1865 Thin-provisioned volume copy offline because of insufficient space.

Explanation: A thin-provisioned volume has been taken offline because there is insufficient allocated real capacity available on the volume for the used space to increase further. If the thin-provisioned volume is auto-expand enabled, then the storage pool it is in also has no free space.

User response: The service action differs depending on whether the thin-provisioned volume copy is auto-expand enabled or not. Whether the disk is auto-expand enabled or not is indicated in the error event data.

If the volume copy is auto-expand enabled, perform one or more of the following actions. When you have performed all of the actions that you intend to perform, mark the error as “fixed”; the volume copy will then return online.

- Determine why the storage pool free space has been depleted. Any of the thin-provisioned volume copies, with auto-expand enabled, in this storage pool might have expanded at an unexpected rate; this could indicate an application error. New volume copies might have been created in, or migrated to, the storage pool.
- Increase the capacity of the storage pool that is associated with the thin-provisioned volume copy by adding more MDisk to the group.
- Provide some free capacity in the storage pool by reducing the used space. Volume copies that are no longer required can be deleted, the size of volume copies can be reduced or volume copies can be migrated to a different storage pool.
- Migrate the thin-provisioned volume copy to a storage pool that has sufficient unused capacity.
- Consider reducing the value of the storage pool warning threshold to give more time to allocate extra space.

If the volume copy is not auto-expand enabled, perform one or more of the following actions. In this case the error will automatically be marked as “fixed”, and the volume copy will return online when space is available.

- Determine why the thin-provisioned volume copy used space has grown at the rate that it has. There might be an application error.
- Increase the real capacity of the volume copy.
- Enable auto-expand for the thin-provisioned volume copy.
- Consider reducing the value of the thin-provisioned volume copy warning threshold to give more time to allocate more real space.

Possible Cause-FRUs or other:

- None

1870 Mirrored volume offline because a hardware read error has occurred.

Explanation: While attempting to maintain the volume mirror, a hardware read error occurred on all of the synchronized volume copies.

The volume copies might be inconsistent, so the volume is now offline.

User response:

- Fix all higher priority errors. In particular, fix any read errors that are listed in the sense data. This error event will automatically be fixed when the root event is marked as “fixed”.

- If you cannot fix the root error, but the read errors on some of the volume copies have been fixed, mark this error as “fixed” to run without the mirror. You can then delete the volume copy that cannot read data and re-create it on different MDisks.

Possible Cause-FRUs or other:

- None

1895 Unrecovered FlashCopy mappings

Explanation: This error might be reported after the recovery action for a cluster failure or a complete I/O group failure. The error is reported because some FlashCopies, whose control data is stored by the I/O group, were active at the time of the failure and the current state of the mapping could not be recovered.

User response: To fix this error it is necessary to delete all of the FlashCopy mappings on the I/O group that failed.

1. Note the I/O group index against which the error is logged.
2. List all of the FlashCopy mappings that are using this I/O group for their bitmaps. You should get the detailed view of every possible FlashCopy ID. Note the IDs of the mappings whose IO_group_id matches the ID of the I/O group against which this error is logged.
3. Note the details of the FlashCopy mappings that are listed so that they can be re-created.
4. Delete all of the FlashCopy mappings that are listed. Note: The error will automatically be marked as “fixed” once the last mapping on the I/O group is deleted. New mappings cannot be created until the error is fixed.
5. Using the details noted in step 3, re-create all of the FlashCopy mappings that you just deleted.

Possible Cause-FRUs or other:

- None

1900 A FlashCopy, Trigger Prepare command has failed because a cache flush has failed.

Explanation: A FlashCopy, Trigger Prepare command has failed because a cache flush has failed.

User response:

1. Correct higher priority errors, and then try the Trigger Prepare command again.
2. Mark the error that you have just repaired as “fixed”.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

Cache flush error (100%)

1910 A FlashCopy mapping task was stopped because of the error that is indicated in the sense data.

Explanation: A stopped FlashCopy might affect the status of other volumes in the same I/O group. Preparing the stopped FlashCopy operations as soon as possible is advised.

User response:

1. Correct higher priority errors, and then prepare and start the FlashCopy task again.
2. Mark the error that you have just repaired as “fixed”.
3. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

1920 Global and Metro Mirror persistent error.

Explanation: This error might be caused by a problem on the primary 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 because of 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 information in the software installation and configuration documentation while diagnosing this error.

User response:

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 because of 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 planning documentation 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 software installation and configuration documentation 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 volume data on the primary and secondary clusters. During this period the data on the Metro Mirror or Global Mirror auxiliary volumes on the secondary cluster is inconsistent and the volumes 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 volumes 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%)

1930 Migration suspended.

Explanation: Migration suspended.

User response:

1. Ensure that all error codes of a higher priority have already been fixed.
2. Ask the customer to ensure that all storage pools that are the destination of suspended migrate operations have available free extents.
3. Mark this error as “fixed”. This causes the migrate operation to be restarted. If the restart fails, a new error is logged.
4. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

1950 Unable to mirror medium error.

Explanation: During the synchronization of a mirrored volume copy it was necessary to duplicate the record of a medium error onto the volume copy, creating a virtual medium error. Each managed disk has a table of virtual medium errors. The virtual medium error could not be created because the table is full. The volume copy is in an inconsistent state and has been taken offline.

User response: Three different approaches can be taken to resolving this problem: 1) the source volume copy can be fixed so that it does not contain medium

errors, 2) the number of virtual medium errors on the target managed disk can be reduced or 3) the target volume copy can be moved to a managed disk with more free virtual medium error entries.

The managed disk with a full medium error table can be determined from the data of the root event.

Approach 1) - This is the preferred procedure because it restores the source volume copy to a state where all of the data can be read. Use the normal service procedures for fixing a medium error (rewrite block or volume from backup or regenerate the data using local procedures).

Approach 2) - This method can be used if the majority of the virtual medium errors on the target managed disk do not relate to the volume copy. Determine where the virtual medium errors are using the event log events and re-write the block or volume from backup.

Approach 3) - Delete the offline volume copy and create a new one either forcing the use of different MDisks in the storage pool or using a completely different storage pool.

Follow your selection option(s) and then mark the error as "fixed".

Possible Cause-FRUs or other:

- None

2008 A software downgrade has failed.

Explanation: Cluster configuration changes are restricted until the downgrade is completed. The cluster downgrade process waits for user intervention when this error is logged.

User response: The action required to recover from a stalled downgrade depends on the current state of the cluster being downgraded. Call IBM Support for an action plan to resolve this problem.

Possible Cause-FRUs or other:

- None

Other:

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.

User response: The action required to recover from a stalled upgrade depends on the current state of the cluster being upgraded. Call IBM technical support for an action plan to resolve this problem.

Possible Cause-FRUs or other:

- None

Other:

2145 software (100%)

2030 Software error.

Explanation: The 2145 software has restarted because of a problem in the cluster, on a disk system or on the fibre-channel fabric.

User response:

1. Collect the software dump file(s) generated at the time the error was logged on the cluster.
2. Contact your product support center to investigate and resolve the problem.
3. Ensure that the software is at the latest level on the cluster and on the disk systems.
4. Use the available SAN monitoring tools to check for any problems on the fabric.
5. Mark the error that you have just repaired as "fixed".
6. Go to repair verification Map.

Possible Cause-FRUs or other:

- Your support center might indicate a FRU based on their problem analysis (2%)

Other:

- 2145 software (48%)
- Enclosure/controller software (25%)
- Fibre-channel switch or switch configuration (25%)

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.

User response:

1. If a FRU has been replaced, ensure that the correct replacement part was used. The node VPD indicates which part is not recognized.
2. Ensure that the cluster software is at the latest level.
3. Save dump data with configuration dump and logged data dump.
4. Contact your product support center to resolve the problem.
5. Mark the error that you have just repaired as "fixed".
6. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

2145 software (100%)

2100 A software error has occurred.

Explanation: One of the 2145 server software components (sshd, crond, or httpd) has failed and reported an error.

User response:

1. Ensure that the software is at the latest level on the cluster.
2. Save dump data with configuration dump and logged data dump.
3. Contact your product support center to resolve the problem.
4. Mark the error that you have just repaired as “fixed”.
5. Go to repair verification MAP.

Possible Cause-FRUs or other:

- None

Other:

2145 software (100%)

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”.

User response:

- Because this error indicates a problem with the number of sessions that are attempting external access to the cluster, determine the reason that so many SSH sessions have been opened.
- Select Launch Maintenance Procedures for the cluster on the Viewing Clusters panel of the cluster 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”.

2600 The cluster was unable to send an email.

Explanation: The cluster has attempted to send an email in response to an event, but there was no acknowledgement that it was successfully received by the SMTP mail server. It might have failed because the cluster was unable to connect to the configured SMTP server, the email might have been rejected by the server, or a timeout might have occurred. The SMTP server might not be running or might not be correctly configured, or the cluster might not be correctly configured. This error is not logged by the test email function because it responds immediately with a result code.

User response:

- Ensure that the SMTP email server is active.
- Ensure that the SMTP server TCP/IP address and port are correctly configured in the cluster email configuration.
- Send a test email and validate that the change has corrected the issue.
- Mark the error that you have just repaired as fixed.
- Go to MAP 5700: Repair verification.

Possible Cause-FRUs or other:

- None

2601 Error detected while sending an email.

Explanation: An error has occurred while the cluster was attempting to send an email in response to an event. The cluster is unable to determine if the email has been sent and will attempt to resend it. The problem might be with the SMTP server or with the cluster email configuration. The problem might also be caused by a failover of the configuration node. This error is not logged by the test email function because it responds immediately with a result code.

User response:

- If there are higher-priority unfixed errors in the log, fix those errors first.
- Ensure that the SMTP email server is active.
- Ensure that the SMTP server TCP/IP address and port are correctly configured in the cluster email configuration.
- Send a test email and validate that the change has corrected the issue.
- Mark the error that you have just repaired as fixed.
- Go to MAP 5700: Repair verification.

Possible Cause-FRUs or other:

- None

2700 Unable to access NTP network time server

Explanation: Cluster time cannot be synchronized with the NTP network time server that is configured.

User response: There are three main causes to examine:

- The cluster NTP network time server configuration is incorrect. Ensure that the configured IP address matches that of the NTP network time server.
- The NTP network time server is not operational. Check the status of the NTP network time server.
- The TCP/IP network is not configured correctly. Check the configuration of the routers, gateways and firewalls. Ensure that the cluster can access the NTP network time server and that the NTP protocol is permitted.

The error will automatically fix when the cluster is able to synchronize its time with the NTP network time server.

Possible Cause-FRUs or other:

- None

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.

User response:

1. Ensure that the room ambient temperature is within the permitted limits.
2. Ensure that the air vents at the front and back of the 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".

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.

User response:

1. Ensure that the room ambient temperature is within the permitted limits.
2. Ensure that the air vents at the front and back of the 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".

3010 Internal uninterruptible power supply software error detected.

Explanation: Some of the tests that are performed during node startup did not complete because some of the data reported by the uninterruptible power supply during node startup is inconsistent because of a software error in the uninterruptible power supply. The node has determined that the uninterruptible power supply is functioning sufficiently for the node to continue operations. The operation of the cluster is not affected by this error. This error is usually resolved by power cycling the uninterruptible power supply.

User response:

1. Power cycle the uninterruptible power supply at a convenient time. The one or two nodes attached to the uninterruptible power supply should be powered off before powering off the uninterruptible power supply. Once the nodes have powered 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%)

3025 A virtualization feature license is required.

Explanation: The cluster has no virtualization feature license registered. You should have either an Entry Edition Physical Disk virtualization feature license or a Capacity virtualization feature license that covers the cluster.

The cluster will continue to operate, but it might be violating the license conditions.

User response:

- If you do not have a virtualization feature license that is valid and sufficient for this cluster, contact your IBM sales representative, arrange a license and change the license settings for the cluster to register the license.
- The error will automatically fix when the situation is resolved.

Possible Cause-FRUs or other:

- None

3029 Virtualization feature capacity is not valid.

Explanation: The setting for the amount of space that can be virtualized is not valid. The value must be an integer number of terabytes.

This error event is created when a cluster is upgraded from a version prior to 4.3.0 to version 4.3.0 or later. Prior to version 4.3.0 the virtualization feature capacity value was in gigabytes and therefore could be set to a fraction of a terabyte. With version 4.3.0 and later the licensed capacity for the virtualization feature must be an integer number of terabytes.

User response:

- Review the license conditions for the virtualization feature. If you have one cluster, change the license settings for the cluster to match the capacity that is licensed. If your license covers more than one cluster, apportion an integer number of terabytes to each cluster. You might have to change the virtualization capacity that is set on the other clusters to ensure that the sum of the capacities for all of the clusters does not exceed the licensed capacity.
- You can view the event data or the feature log to ensure that the licensed capacity is sufficient for the space that is actually being used. Contact your IBM sales representative if you want to change the capacity of the license.
- This error will automatically be fixed when a valid configuration is entered.

Possible Cause-FRUs or other:

- None

3030 Global and Metro Mirror feature capacity not set.

Explanation: The Global and Metro Mirror feature is set to On for the cluster, but the capacity has not been set.

This error event is created when a cluster is upgraded from a version prior to 4.3.0 to version 4.3.0 or later. Prior to version 4.3.0 the 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.

User response: 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 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.

User response: 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 cluster 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 cluster will continue to operate, but it might be violating the license conditions.

User response:

- Determine which feature license limit has been exceeded. This might be:
- Virtualization 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 volumes. The used virtualization size is the sum of the capacities of all of the managed disks and image mode disks.
- To reduce the FlashCopy capacity delete some FlashCopy mappings. The used FlashCopy size is the sum of all of the volumes that are the source volume of a FlashCopy mapping.
- To reduce Global and Metro Mirror capacity delete some Global Mirror or Metro Mirror relationships. The used Global and Metro Mirror size is the sum of the capacities of all of the volumes that are in a Metro Mirror or Global Mirror relationship; both master and auxiliary volumes are counted.
- The error will automatically be fixed when the licensed capacity is greater than the capacity that is being used.

Possible Cause-FRUs or other:

- None

3035 Physical Disk FlashCopy feature license required

Explanation: The Entry Edition cluster has some FlashCopy mappings defined. There is, however, no Physical Disk FlashCopy license registered on the cluster. The cluster will continue to operate, but it might be violating the license conditions.

User response:

- Check whether you have an Entry Edition Physical Disk FlashCopy license for this cluster that you have not registered on the cluster. Update the cluster license configuration if you have a license.
- Decide whether you want to continue to use the FlashCopy feature or not.
- If you want to use the FlashCopy feature contact your IBM sales representative, arrange a license and change the license settings for the cluster to register the license.
- If you do not want to use the FlashCopy feature, you must delete all of the FlashCopy mappings.
- The error will automatically fix when the situation is resolved.

Possible Cause-FRUs or other:

- None

3036 Physical Disk Global and Metro Mirror feature license required

Explanation: The Entry Edition cluster has some Global Mirror or Metro Mirror relationships defined. There is, however, no Physical Disk Global and Metro Mirror license registered on the cluster. The cluster will continue to operate, but it might be violating the license conditions.

User response:

- Check if you have an Entry Edition Physical Disk Global and Metro Mirror license for this cluster that you have not registered on the cluster. Update the cluster license configuration if you have a license.
- Decide whether you want to continue to use the Global Mirror or Metro Mirror features or not.
- If you want to use either the Global Mirror or Metro Mirror feature contact your IBM sales representative, arrange a license and change the license settings for the cluster to register the license.
- If you do not want to use both the Global Mirror and Metro Mirror features, you must delete all of the Global Mirror and Metro Mirror relationships.
- The error will automatically fix when the situation is resolved.

Possible Cause-FRUs or other:

- None

3080 Global or Metro Mirror relationship or consistency group with deleted partnership

Explanation: A Global Mirror or Metro Mirror relationship or consistency group exists with a cluster whose partnership is deleted.

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.

User response: The issue can be resolved either by deleting all of the Global Mirror or Metro Mirror relationships or consistency groups that exist with a cluster whose partnership is deleted, or by recreating all of the partnerships that they were using.

The error will automatically fix when the situation is resolved.

- 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 cluster 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 email to any of the configured email servers.

Explanation: Either the system was not able to connect to any of the SMTP email servers, or the email transmission has failed. A maximum of six email servers can be configured. Error event 2600 or 2601 is raised when an individual email server is found to be not working. This error indicates that all of the email servers were found to be not working.

User response:

- Check the event log for all unresolved 2600 and 2601 errors and fix those problems.
- If this error has not already been automatically marked fixed, mark this error as fixed.
- Perform the check email function to test that an email server is operating properly.

Possible Cause-FRUs or other:

- None

SAN problem determination

The procedures that are provided here help you solve problems on the SAN Volume Controller system and its connection to the storage area network (SAN).

SAN failures might cause SAN Volume Controller drives 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) transceiver
- 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 storage controllers that the SAN Volume Controller system 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 systems 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.

Fibre Channel link failures

When a failure occurs on a single Fibre Channel link, the small form-factor pluggable (SFP) transceiver might need to be replaced.

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
- An error that indicates that 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 transceiver for the failing port on the SAN Volume Controller node.

Note: SAN Volume Controller nodes are supported with both longwave SFP transceivers and shortwave SFP transceivers. You must replace an SFP transceiver with the same type of SFP transceiver. If the SFP transceiver to replace is a longwave SFP transceiver, for example, you must provide a suitable replacement. Removing the wrong SFP transceiver could result in loss of data access.

4. Perform the Fibre Channel switch service procedures for a failing Fibre Channel link. This might involve replacing the SFP transceiver 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 system, or the host.

For network problems, you can attempt any of the following actions:

- Test your connectivity between the host and SAN Volume Controller ports.
- Try to ping the SAN Volume Controller system from the host.
- Ask the Ethernet network administrator to check the firewall and router settings.
- Check that the subnet mask and gateway are correct for the SAN Volume Controller host configuration.

For SAN Volume Controller problems, you can attempt any of the following actions:

- | • View the configured node port IP addresses by using the **svcinfo lspportip** CLI command.
- | • View the list of volumes that are mapped to a host by using the **svcinfo lshostvdiskmap** command to ensure that the volume host mappings are correct.
- | • Verify that the volume is online by using the **svcinfo lsvdisk** command.

For host problems, you can attempt any of the following actions:

- Verify that the host iSCSI qualified name (IQN) is correctly configured.
- Use operating system utilities (such as Windows device manager) to verify that the device driver is installed, loaded, and operating correctly.

Servicing storage systems

Storage systems that are supported for attachment to the SAN Volume Controller system are designed with redundant components and access paths to enable concurrent maintenance. Hosts have continuous access to their data during component failure and replacement.

The following guidelines apply to all storage systems that are attached to the SAN Volume Controller system:

- Always follow the service instructions that are provided in the documentation for your storage system.
- Ensure that there are no unfixed errors in the event log before you perform any service procedures.
- After you perform a service procedure, check the event 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 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.

Chapter 8. Recovery procedures

This topic describes two recovery procedures: recover a cluster and back up and restore a cluster configuration. This topic also contains information about the performing the node rescue.

Recover cluster procedure

The recover cluster procedure recovers the entire cluster if cluster data has been lost from all nodes. The procedure re-creates the cluster using saved configuration data. The recovery might not be able to restore all volume data. This procedure is also known as Tier 3 (T3) recovery.

Attention: Perform service actions on nodes only when directed by the service procedures. If used inappropriately, service actions can cause loss of access to data or even data loss. Before attempting to recover a cluster, investigate the cause of the cluster failure and attempt to resolve those issues by using other service procedures.

Attention: Do not attempt to run the recovery procedure unless you are familiar with the following conditions:

- All hardware errors are fixed.
- All nodes are in candidate status.
- Run the procedure on one cluster in a fabric at a time. Do not perform the procedure on different nodes in the same cluster. This restriction also applies to remote clusters.

When the procedure is successful, other actions must be performed to get your environment operational.

You can run the cluster recovery procedure using the front panel or the service assistant.

When to run the recover cluster procedure

A recover cluster procedure must be attempted only after a complete and thorough investigation of the cause of the cluster failure. Attempt to resolve those issues by using other service procedures.

Attention: If at any time while you are running the recover cluster procedure, you experience failures, call the IBM Support Center. Do not attempt to do further recovery actions because these actions might prevent IBM Support from restoring the system to an operational status.

Certain conditions must be met before you run the recovery procedure. Use the following items to help you determine when to run the recovery procedure:

- Check to see if any node in the cluster has a node status of active. This status means that the cluster is still available. In this case, cluster recovery is not necessary.
- Do not recover the cluster if the cluster IP address is available from another node. Ensure that all service procedures have been run.
- Check the node status of every node that is a member of the cluster. Resolve all errors.
 - All nodes must be reporting either a node error 578 or a cluster error. These error codes indicate that the cluster has lost its configuration data. If any nodes report anything other than these error codes, do not perform a cluster recovery. You can encounter situations where non-configuration nodes report other node errors, such as a node error 550. The 550 error can also indicate that a node is not able to join a cluster.

Note: If any of the buttons on the front panel have been pressed after these two error codes are reported, the report for the node returns to the 578 node error. The change in the report happens

after approximately 60 seconds. Also, if the node was rebooted or if hardware service actions were taken, the node might show only the Cluster: error.

- If any nodes show Node Error: 550, record the data from the second line of the display. If the last character on the second line of the display is >, use the right button to scroll the display to the right.
 - In addition to the Node Error: 550, the second line of the display can show a list of node front panel IDs (7 digits) that are separated by spaces. The list can also show the WWPN/LUN ID (16 hexadecimal digits followed by a forward slash and a decimal number).
 - If the error data contains any front panel IDs, ensure that the node referred to by that front panel ID is showing Node Error 578:. If it is not reporting node error 578, ensure that the two nodes can communicate with each other. Verify the SAN connectivity and restart one of the two nodes by pressing the front panel power button twice.
 - If the error data contains a WWPN/LUN ID, verify the SAN connectivity between this node and that WWPN. Check the storage system to ensure that the LUN referred to is online. After verifying these items, restart the node by pressing the front panel power button twice.

Note: If after resolving all these scenarios, half or greater than half of the nodes are reporting Node Error: 578, it is appropriate to run the cluster recovery procedure.

- For any nodes that are reporting a node error 550, ensure that all the missing hardware that is identified by these errors is powered on and connected without faults.
- If you have not been able to restart the cluster and if any node other than the current node is reporting node error 550 or 578, you must remove cluster data from those nodes. This action acknowledges the data loss and puts the nodes into the required candidate state.
- Do not attempt to recover the cluster if you have been able to restart the cluster.
- If back-end MDisks are removed from the configuration, those volumes that depended on that hardware cannot be recovered. All previously configured back-end hardware must be present for a successful recovery.
- Any nodes that were replaced must have the same WWNN as the nodes that they replaced.
- The configuration backup file must be up to date. If any configuration changes had been made since the backup was taken, the data is inconsistent and further investigation is needed. Manual changes are required after the cluster is recovered.
- Any data that was in the cache at the point of cluster failure is lost. The loss of data can result in data corruption on the affected volumes. If the volumes are corrupted, call the IBM Support Center.

Fix hardware errors

Before you can run a cluster recovery procedure, it is important that the root cause of the hardware issues be identified and fixed.

Obtain a basic understanding about the hardware failure. In most situations when there is no cluster, a power issue is the cause.

- The node has been powered off or the power cords were unplugged.
- A 2145 UPS-1U might have failed and shut down one or more nodes because of the failure. In general, this cause might not happen because of the redundancy provided by the second 2145 UPS-1U.

Removing cluster information for nodes with error code 550 or error code 578 using the front panel

The cluster recovery procedure works only when all nodes are in candidate status. If there are any nodes that display error code 550 or error code 578, you must remove their cluster data.

To remove cluster information from a node with an error 550 or 578, follow this procedure using the front panel:

1. Press and release the up or down button until the **Actions** menu option is displayed.

- | 2. Press and release the select button.
 - | 3. Press and release the up or down button until **Remove Cluster?** option is displayed.
 - | 4. Press and release the select button.
 - | 5. The node displays **Confirm Remove?**.
 - | 6. Press and release the select button.
 - | 7. The node displays **Cluster:**.
- | When all nodes show **Cluster:** on the top line and blank on the second line, the nodes are in candidate status. The 550 or 578 error has been removed. You can now run the cluster recovery procedure.

| **Removing cluster information for nodes with error code 550 or error code 578 using the service assistant**

| The cluster recovery procedure works only when all nodes are in candidate status. If there are any nodes that display error code 550 or error code 578, you must remove their cluster data.

| To remove cluster information from a node with an error 550 or 578, follow this procedure using the service assistant:

- | 1. Point your browser to the service IP address of one of the nodes.
| If you do not know the IP address or if it has not been configured, you must assign an IP address using the initialization tool.
- | 2. Log on to the service assistant.
- | 3. Select **Manage cluster**.
- | 4. Click **Remove cluster data**.
- | 5. Click **Confirm Remove Cluster**.
- | 6. Go to the home page to view the error condition and node status for the nodes in the cluster.
| All nodes for this cluster must be in candidate status. The error conditions must be **None**.
- | **Note:** A node that is powered off might not show up in this list of nodes for the cluster. Diagnose hardware problems directly on the node using the service assistant IP address and by physically verifying the LEDs for the hardware components.
- | 7. Resolve any hardware errors until the error condition for all nodes in the cluster is **None**.
- | 8. Ensure that all nodes in the cluster display a status of candidate.

| When all nodes display a status of candidate and all error conditions are **None**, you can run the cluster recovery procedure.

| **Performing cluster recovery procedure using the front panel**

| Start recovery when all nodes that were members of the cluster are online and are in candidate status. Do not run the recovery procedure on different nodes in the same cluster. This restriction includes remote clusters also.

| **Attention:** This service action has serious implications if not performed properly. If at any time during the procedure, you encounter an error, stop, and call IBM Support.

| You might see any one of the following categories of messages:

- | • T3 successful. If the volumes are back online, use the final checks to get your environment operational again.
- | • T3 incomplete. One or more of the volumes is offline because there was fast write data in the cache. Further actions are required to bring the volumes online again. Contact IBM Support for more details regarding how to bring the volumes online again.

- T3 failed. Call IBM Support. Do not attempt any further action.

The cluster recovery procedure can start from any node in the cluster. The node must not have participated in any other cluster. To receive optimal results in maintaining the I/O group ordering, run the recovery from a node that was in I/O group 0.

1. Press and release the up or down button until the **Actions** menu option is displayed.
2. Press and release the select button.
3. Press and release the up or down button until the **Recover Cluster?** option is displayed.
4. Press and release the select button.
5. The node displays **Confirm Recover?**.
6. Press and release the select button.
7. The node displays **Retrieving** on the top line.
After a short delay, the second line displays a sequence of progress messages that indicate the actions that are taking place, for example, finding qdisks. The backup files are scanned to find the most recent configuration backup data.
8. The node displays **T3 data:** on the top line after the file and quorum data retrieval is complete.
9. Verify the date and time on the second line of the display. The time stamp shown is the date and time of the last quorum update and must be less than 10 minutes before the cluster failure. The time stamp format is *YYYYMMDD hh:mm*, where *YYYY* is the year, *MM* is the month, *DD* is the day, *hh* is the hour, and *mm* is the minute.
Attention: If the time stamp is not less than 10 minutes before the cluster failure, call IBM support.
10. When you have verified that the time stamp is correct, perform the following steps:
 - a. Press and hold the up key.
 - b. Press and release select.
 - c. Release the up key.
11. The node displays **Backup file** on the top line.
12. Verify the date and time on the second line of the display. The time stamp shown is the date and time of the last configuration backup and must be less than 24 hours before the cluster failure. The time stamp format is *YYYYMMDD hh:mm*, where *YYYY* is the year, *MM* is the month, *DD* is the day, *hh* is the hour, and *mm* is the minute.
Attention: If the time stamp is not less than 24 hours before the cluster failure, call IBM support.
Changes made after the time of this configuration backup might not be restored.
13. When you have verified that the time stamp is correct, perform the following steps:
 - a. Press and hold the up key.
 - b. Press and release select.
 - c. Release the up key.
14. The node displays **Restoring**. After a short delay, the second line displays a sequence of progress messages that indicate the actions that are taking place.
15. The software on the node restarts.
16. The node displays **Cluster** on the top line and a cluster IP address on the second line.
17. After a few moments, the node displays **T3 Completing**.
Note: If there are any cluster errors logged at this time, they might temporarily overwrite the display. If you see **Cluster Error: 3025**, you can ignore the message. After a short delay, the second line displays a sequence of progress messages that indicate the actions that are taking place.
18. When each node is added to the cluster, the display shows **Cluster:** on the top line and the cluster name on the second line.
After the last node is added to the cluster, there is a short delay to allow the cluster to stabilize.

| **Attention:** Do not attempt to use the cluster. The recovery is still in progress.

| 19. The node displays **T3 Succeeded** on the top line.

| 20. Press select to return the node to its normal display.

| When the node displays **T3 Succeeded**, the recovery is complete. You can now perform the checks to get your environment operational.

| **Performing cluster recovery using the service assistant**

| Start recovery when all nodes that were members of the cluster are online and are in candidate status. Do not run the recovery procedure on different nodes in the same cluster. This restriction includes remote clusters also.

| **Attention:** This service action has serious implications if not performed properly. If at any time during the procedure, you encounter an error, stop and call IBM Support.

| **Note:** Your web browser must not block pop-up windows; otherwise, progress windows cannot open.

| You might see any one of the following categories of messages:

- | • T3 successful. The volumes are back online. Use the final checks to get your environment operational again.
- | • T3 incomplete. One or more of the volumes is offline because there was fast write data in the cache. Further actions are required to bring the volumes online again. Contact IBM Support for more details regarding how to bring the volumes online again.
- | • T3 failed. Call IBM Support. Do not attempt any further action.

| The recovery can be run from any nodes in the cluster. The nodes must not have participated in any other cluster. To receive optimal results in maintaining the I/O group ordering, run the recovery from a node that was in I/O group 0.

| 1. Point your browser to the service IP address of one of the nodes.

| If you do not know the IP address or if it has not been configured, you must assign an IP address using the initialization tool.

| 2. Log on to the service assistant.

| 3. Select **Recover cluster**.

| 4. Follow the online instructions to complete the recovery procedure.

| Verify the date and time of the last quorum time. The time stamp must be less than 10 minutes before the cluster failure. The time stamp format is *YYYYMMDD hh:mm*, where *YYYY* is the year, *MM* is the month, *DD* is the day, *hh* is the hour, and *mm* is the minute.

| **Attention:** If the time stamp is not less than 10 minutes before the cluster failure, call IBM support. If you do not meet all the conditions that were listed previously, call IBM Support.

| Verify the date and time of the last backup date. The time stamp must be less than 24 hours before the cluster failure. The time stamp format is *YYYYMMDD hh:mm*, where *YYYY* is the year, *MM* is the month, *DD* is the day, *hh* is the hour, and *mm* is the minute.

| **Attention:** If the time stamp is not less than 24 hours before the cluster failure, call IBM Support.

| Changes made after the time of this backup date might not be restored.

| After the recovery completes successfully, perform the checks to get your environment operational.

| **What to check after running cluster recovery**

| Several tasks must be performed before you use the volumes.

| Differences to be aware of regarding the recovered configuration:

- | • FlashCopy mappings are restored as "idle_or_copied" with 0% progress. Both volumes must have been restored to their original I/O groups.
- | • The cluster ID is different. Any scripts or associated programs that refer to the cluster ID of the cluster must be changed.
- | • Any FlashCopy mappings that were not in the "idle_or_copied" state with 100% progress at the point of disaster have inconsistent data on their target disks. These mappings must be restarted.
- | • Intercluster remote copy partnerships and relationships are not restored and must be re-created manually.
- | • Consistency groups are not restored and must be re-created manually.
- | • Intracluster remote copy relationships are restored if all dependencies were successfully restored to their original I/O groups.
- | • The cluster time zone might not have been restored.

| Before using the volumes, perform the following tasks:

- | • Start the host systems.
- | • Manual actions might be necessary on the hosts to trigger them to rescan for devices. You can perform this task by disconnecting and reconnecting the Fibre Channel cables to each host bus adapter (HBA) port.
- | • Verify that all mapped volumes can be accessed by the hosts.
- | • Run file system consistency checks.

| **Note:** Any data that was in the SAN Volume Controller write cache at the time of the failure is lost.

- | • Run the application consistency checks.

Backing up and restoring the cluster configuration

You can back up and restore the cluster configuration data after preliminary tasks are completed.

Cluster configuration data provides information about your cluster and the objects that are defined in it. The backup and restore functions of the **svconfig** command can only back up and restore your cluster configuration data. You must regularly back up your application data using the appropriate backup methods.

| You can maintain your cluster configuration data by completing the following tasks:

- | • Backing up the cluster configuration data
- | • Restoring the cluster configuration data
- | • Deleting unwanted backup configuration data files

Before you back up your cluster configuration data, the following prerequisites must be met:

- No independent operations that change the cluster configuration can be running while the backup command is running.
- No object name can begin with an underscore character (_).

Note:

- The default object names for controllers, I/O groups, and managed disks (MDisks) do not restore correctly if the ID of the object is different from what is recorded in the current cluster configuration data file.
- | • All other objects with default names are renamed during the restore process. The new names appear in the format *name_r* where *name* is the name of the object in your cluster.

Before you restore your cluster configuration data, the following prerequisites must be met:

- You have the Security Administrator role associated with your user name and password.
- You have a copy of your backup cluster configuration files on a server that is accessible to the cluster.
- You have a backup copy of your application data that is ready to load on your cluster after the restore configuration operation is complete.
- You know the current license settings for your cluster.
- You have not removed any hardware since the last backup of your cluster configuration. If you had to replace a faulty node, the new node must use the same worldwide node name (WWNN) as the faulty node that it replaced.

Note: You can add new hardware, but you must not remove any hardware because the removal can cause the restore process to fail.

- No zoning changes have been made on the Fibre Channel fabric which would prevent communication between the SAN Volume Controller and other nodes or storage controllers which are present in the configuration.

The restore must be performed to a single node cluster. You can restore the configuration using any node as the configuration node. However, if you do not use the node that was the configuration node when the cluster was first created, the unique identifier (UID) of the volumes that are within the I/O groups can change. This can affect IBM Tivoli Storage Productivity Center for Fabric, VERITAS Volume Manager, and any other programs that record this information.

The SAN Volume Controller analyzes the backup configuration data file and the cluster to verify that the required disk controller system nodes are available.

Before you begin, hardware recovery must be complete. The following hardware must be operational: hosts, SAN Volume Controller, disk controller systems, disks, the Ethernet network, and the SAN fabric.

Backing up the cluster configuration using the CLI

You can back up your cluster configuration data using the command-line interface (CLI).

Before you back up your cluster configuration data, the following prerequisites must be met:

- No independent operations that change the cluster configuration can be running while the backup command is running.
- No object name can begin with an underscore character (_).
- All objects must have non-default names, that is, names that are not assigned by the SAN Volume Controller.

Note:

- The default object names for controllers, I/O groups, and managed disks (MDisks) do not restore correctly if the ID of the object is different from what is recorded in the current cluster configuration data file.
- All other objects with default names are renamed during the restore process. The new names appear in the format *name_r*, where *name* is the name of the object in your cluster.

The backup feature of the **svconfig** CLI command is designed to back up information about your cluster configuration, such as volumes, local Metro Mirror information, local Global Mirror information, managed disk (MDisk) groups, and nodes. All other data that you have written to the volumes is *not* backed up. Any application that uses the volumes on the cluster as storage, must back up its application data using the appropriate backup methods.

You must regularly back up your cluster configuration data and your application data to avoid data loss. If a cluster is lost after a severe failure occurs, both cluster configuration and application data is lost. You must reinstate the cluster to the exact state it was in before the failure, and then recover the application data.

Perform the following steps to back up your cluster configuration data:

1. Back up all of the application data that you have stored on your volumes using your preferred backup method.
2. Open a command prompt.
3. Issue the following command to log on to the cluster:

```
ssh -l admin your_cluster_name -p 22
```

where *your_cluster_name* is the name of the cluster for which you want to back up cluster configuration data.

4. Issue the following CLI command to remove all of the existing cluster configuration backup and restore files that are located on your configuration node in the /tmp directory.

```
svconfig clear -all
```

5. Issue the following CLI command to back up your cluster configuration:

```
svconfig backup
```

The following output is an example of the messages that are displayed during the backup process:

```
CMMVC6112W io_grp io_grp1 has a default name
CMMVC6112W io_grp io_grp2 has a default name
CMMVC6112W mdisk mdisk14 ...
CMMVC6112W node node1 ...
CMMVC6112W node node2 ...
.....
```

The **svconfig backup** CLI command creates three files that provide information about the backup process and cluster configuration. These files are created in the /tmp directory of the configuration node.

The following table describes the three files that are created by the backup process:

| File name | Description |
|-----------------------|---|
| svc.config.backup.xml | This file contains your cluster configuration data. |
| svc.config.backup.sh | This file contains the names of the commands that were issued to create the backup of the cluster. |
| svc.config.backup.log | This file contains details about the backup, including any error information that might have been reported. |

6. Issue the following command to exit the cluster:

```
exit
```

7. Issue the following command to copy the backup files to a location that is not in your cluster:

```
scp -P 22 admin@your_cluster:/tmp/svc.config.backup.*
  /offclusterstorage/
```

where *your_cluster* is the name of your cluster and *offclusterstorage* is the location where you want to store the backup files.

You must copy these files to a location outside of your cluster because the /tmp directory on this node becomes inaccessible if the configuration node changes. The configuration node might change in response to an error recovery action or to a user maintenance activity.

Tip: To maintain controlled access to your cluster configuration data, copy the backup files to a location that is password protected.

8. Ensure that the copies of the backup files are stored in the location that you specified in step 7 on page 196.

You can rename the backup files to include the configuration node name either at the start or end of the file names so that you can easily identify these files when you are ready to restore your configuration.

Issue the following command to rename the backup files that are stored on a Linux or IBM AIX host:

```
mv /offclusterstorage/svc.config.backup.xml  
/offclusterstorage/svc.config.backup.xml_myconfignode
```

where *offclusterstorage* is the name of the directory where the backup files are stored and *myconfignode* is the name of your configuration node.

To rename the backup files that are stored on a Windows host, right-click the name of the file and select **Rename**.

Restoring the cluster configuration

| **Attention:** Use this procedure only if the recover cluster procedure has failed or if the data that is stored on the volumes is not required. For directions on the recover cluster procedure, see “Recover cluster procedure” on page 189.

| This configuration restore procedure is designed to restore information about your cluster configuration, such as volumes, local Metro Mirror information, local Global Mirror information, storage pools, and nodes. All the data that you have written to the volumes is not restored. To restore the data on the volumes, you must restore application data from any application that uses the volumes on the cluster as storage separately. Therefore, you must have a backup of this data before you follow the cluster configuration recovery process.

You must regularly back up your cluster configuration data and your application data to avoid data loss. If a cluster is lost after a severe failure occurs, both cluster configuration and application data is lost. You must reinstate the cluster to the exact state it was in before the failure, and then recover the application data.

| **Important:** There are two phases during the restore process: prepare and execute. You must not change the fabric or cluster between these two phases.

| If you do not understand the instructions to run the CLI commands, refer to the command-line interface reference information.

| To restore your cluster configuration data, follow these steps:

1. The node must be available as a candidate node before you can run this recovery procedure. You must remove error code 550 or 578 to put the node in candidate status.
2. Create a new cluster from the front panel. If possible, use the node that was originally in I/O group 0.
3. From the management GUI, click **User Management** → **Users** to set up your system and configure an SSH key for superuser. This allows access to the CLI.
4. Using the command-line interface, issue the following command to log on to the cluster:

```
plink -i ssh_private_key_file admin@cluster_ip
```

where *ssh_private_key_file* is the name of the SSH private key file for superuser and *cluster_ip* is the IP address or DNS name of the cluster for which you want to restore the cluster configuration.

| **Note:** Because the RSA host key has changed, a warning message might display when you connect to the cluster using SSH.

5. Issue the following CLI command to ensure that only the configuration node is online:

```
svcinfolsnode
```

The following output is an example of what is displayed:

```
id name status IO_group_id IO_group_name config_node
1 node1 online 0 io_grp0 yes
```

6. Identify the configuration backup file that you want to restore from.

Configuration data is automatically backed up daily at 01:00 system time.

Attention: You must copy the required backup file to another computer before you continue. To save a copy of the data, perform the following steps:

- a. From the management GUI, click **Troubleshooting** → **Support**.
 - b. Click **Show full log listing**.
 - c. Find the file name that begins with `svc.config.cron.xml`.
 - d. Double-click the file to download the file to your computer.
 - e. If a recent configuration file is not present on this node, configure service IP addresses for other nodes and connect to the service assistant to look for configuration files on other nodes. For details on how to do this, see the information regarding service IPv4 or service IPv6 at “Service IPv4 or Service IPv6” on page 99.
7. Issue the following CLI command to remove all of the existing backup and restore cluster configuration files that are located on your configuration node in the `/tmp` directory:

```
svcconfig clear -all
```

8. The XML files contain a date and time that can be used to identify the most recent backup. After you identify the backup XML file that is to be used when you restore the cluster, rename the file to `svc.config.backup.xml`. From your desktop, issue the following command to copy the file back on to the cluster.

```
pscp -i ssh_private_key_file
full_path_to_identified_svc.config.backup.xml
admin@cluster_ip:/tmp/
```

9. Issue the following CLI command to compare the current cluster configuration with the backup configuration data file:

```
svcconfig restore -prepare
```

This CLI command creates a log file in the `/tmp` directory of the configuration node. The name of the log file is `svc.config.restore.prepare.log`.

Note: It can take up to a minute for each 256-MDisk batch to be discovered. If you receive error message CMMVC6119E for an MDisk after you enter this command, all the managed disks (MDisks) might not have been discovered yet. Allow a suitable time to elapse and try the `svcconfig restore -prepare` command again.

10. Issue the following command to copy the log file to another server that is accessible to the cluster:

```
pscp -i ssh_private_key_file
admin@cluster_ip:/tmp/svc.config.restore.prepare.log
full_path_for_where_to_copy_log_files
```

11. Open the log file from the server where the copy is now stored.

12. Check the log file for errors.

- If there are errors, correct the condition that caused the errors and reissue the command. You must correct all errors before you can proceed to step 13.
- If you need assistance, contact the IBM Support Center.

13. Issue the following CLI command to restore the cluster configuration:

```
svcconfig restore -execute
```

Note: Issuing this CLI command on a single node cluster adds the other nodes to the cluster.

This CLI command creates a log file in the /tmp directory of the configuration node. The name of the log file is svc.config.restore.execute.log.

14. Issue the following command to copy the log file to another server that is accessible to the cluster:

```
| pscp -i ssh_private_key_file  
| admin@cluster_ip:/tmp/svc.config.restore.execute.log  
| full_path_for_where_to_copy_log_files
```

15. Open the log file from the server where the copy is now stored.
16. Check the log file to ensure that no errors or warnings have occurred.

Note: You might receive a warning stating that a licensed feature is not enabled. This message means that after the recovery process, the current license settings do not match the previous license settings. The recovery process continues normally and you can enter the correct license settings in the management GUI at a later time.

The following output is displayed after a successful cluster configuration restore operation:

```
IBM_2145:your_cluster_name:admin>
```

17. After the cluster configuration is restored, verify that the quorum disks are restored to the MDisks that you want by using the **svcinfolsqorum** command. To restore the quorum disks to the correct MDisks, issue the appropriate **svctask chquorum** CLI commands.

You can remove any unwanted configuration backup and restore files from the /tmp directory on your configuration by issuing the following CLI command:

```
svconfig clear -all
```

Deleting backup configuration files using the CLI

You can use the command-line interface (CLI) to delete backup configuration files.

Perform the following steps to delete backup configuration files:

1. Issue the following command to log on to the cluster:

```
ssh -l admin your_cluster_name -p 22
```

where *your_cluster_name* is the name of your cluster.

2. Issue the following CLI command to erase all of the files that are stored in the /tmp directory:

```
svconfig clear -all
```

Performing the node rescue when the node boots

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, 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.

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, run only 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 67).



Figure 67. 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.

Chapter 9. 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 225.

- 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 event log after the repair, but before you verify the repair.

Note: Start all problem determination procedures and repair procedures with “MAP 5000: Start.”

MAP 5000: Start

MAP 5000: Start is an entry point to the maintenance analysis procedures (MAPs) for the SAN Volume Controller.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 9, “Using the maintenance analysis procedures.”

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 fix procedures 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.
- 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 fix procedure?**

NO Go to step 2

YES Go to step 8

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 using the user ID and password that is provided by the user.

4. (from step 3)

Log into the management GUI using the user ID and password that is provided by the user and launch the management GUI for the cluster that you are repairing.

5. (from step 4)

Does the management GUI start?

NO Go to step 8.

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)

Start the fix procedures.

Did the fix procedures find an error that needs to be fixed?

NO Go to step 8.

YES Follow the fix procedures.

8. (from steps 1, 5, 6, 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.

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 104.

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 213, “MAP 5050: Power 2145-CF8, 2145-8G4, 2145-8F4, and 2145-8F2” on page 207.

9. (from step 8 on page 202)

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 68 shows the location of the service controller error light.




Figure 68. 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 for SAN Volume Controller models 2145-CF8, 2145-8A4, 2145-8G4, 2145-8F4, and 2145-8F2.

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 239.

10. (from step 9)

Is the operator-information panel error LED **1 that you see in Figure 69 illuminated or flashing?**




Figure 69. 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 11 on page 204.

YES Go to “MAP 5800: Light path” on page 240.

11. (from step 10 on page 203)

Is the hardware boot display that you see in Figure 70 displayed on any of the SAN Volume Controller nodes?



Figure 70. Hardware boot display

NO Go to step 13.

YES Go to step 12.

12. (from step 11)

Has the hardware boot display that you see in Figure 70 displayed for more than three minutes?

NO Go to step 13.

YES Perform the following:

a. Go to “MAP 5900: Hardware boot” on page 256.

b. Go to “MAP 5700: Repair verification” on page 239.

13. (from steps 11)

Is Failed displayed on the top line of the front-panel display of any SAN Volume Controller node?

NO Go to step 14.

YES Perform the following:

a. Note the failure code and go to boot code “Boot code reference” on page 131 to perform the repair actions.

b. Go to “MAP 5700: Repair verification” on page 239.

14. (from step 13)

Is Booting displayed on the top line of the front-panel display of any SAN Volume Controller node?

NO Go to step 16.

YES Go to step 15.

15. (from step 14)

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 16.

YES Perform the following:

a. Note the failure code and go to boot code “Boot code reference” on page 131 to perform the repair actions.

b. Go to “MAP 5700: Repair verification” on page 239.

16. (from steps 14 and step 15)

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 17.
- YES** Perform the following steps:
- Note the failure code and go to node error code “Node error code overview” on page 132 to perform the repair actions.
 - Go to “MAP 5700: Repair verification” on page 239.
17. (from step 16 on page 204)
- Is Cluster Error displayed on the top line of the front-panel display of any SAN Volume Controller node?**
- NO** Go to step 18.
- 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 fix procedures. Perform the following steps:
- Go to step 2 on page 202 to perform the fix procedure. If you return here, go to cluster error code “Cluster code overview” on page 133 to perform the repair actions.
 - Go to “MAP 5700: Repair verification” on page 239.
18. (from step 17)
- Is Powering Off, Restarting, Shutting Down, or Power Failure displayed in the top line of the front-panel display?**
- NO** Go to step 20.
- YES** The progress bar moves every few seconds. Wait for the operation to complete and then return to step 1 on page 202 in this MAP. If the progress bar does not move for three minutes, press the power button and go to step 19.
19. (from step 18)
- Did the SAN Volume Controller node power off?**
- NO** Perform the following steps:
- Remove the power cord from the rear of the box.
 - Wait 60 seconds.
 - Replace the power cord.
 - 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 202 in this MAP.
- YES** Perform the following steps:
- Wait 60 seconds.
 - Click the power button to turn on the SAN Volume Controller node and then return to step 1 on page 202 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.
20. (from step 19)
- Is Charging or Recovering displayed in the top line of the front-panel display of any SAN Volume Controller node?**
- NO** Go to step 21 on page 206.
- 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 216.

- 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 216.

21. (from step 20 on page 205)

Is Validate WWNN? displayed on the front-panel display of any SAN Volume Controller node?

NO Go to step 22.

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.

a. 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).

b. 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 cluster recover procedure to delete the node from the cluster and add the node back into the cluster.

22. (from step 21)

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 23.

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.

23. (from step 22 on page 206)

Is the front-panel display unreadable?

NO Go to step 24.

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 229.

24. (from step 23)

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 system" documentation for
| details about how to upgrade your entire SAN Volume Controller environment.

If the problem is still not fixed, collect diagnostic information and contact the IBM support center.

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 9, "Using the maintenance analysis procedures," on page 201.

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 211.

YES Go to step 2.

2. (from step 1)

Is the power LED on the operator-information panel continuously illuminated? Figure 71 shows the location of the power LED **1** on the operator-information panel.




Figure 71. 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 201 or go to “MAP 5700: Repair verification” on page 239 to verify the correct operation.

3. (from step 2 on page 207)

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.

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" to see how to remove the cable-retention brackets when removing the power cords from the node.
- 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.
- 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 239.

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 and step 4)

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 239.

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 239 to verify the correct operation.

6. (from step 3 on page 208 and step 4 on page 208)

Is the rear panel power LED on or flashing? Figure 72 shows the location of the power LED **1** on the rear panel of the 2145-8G4, 2145-8F4, or 2145-8F2 nodes. Figure 73 on page 210 shows the location of the power LED **1** on the 2145-CF8.





Figure 72. Power LED on the SAN Volume Controller models 2145-8G4, 2145-8F4, and 2145-8F2 rear panel



svc00574

Figure 73. Power LED indicator on the rear panel of the SAN Volume Controller 2145-CF8

NO Go to step 7.

YES The operator-information panel is failing.

Verify that the operator-information panel cable is seated on the system board.

If 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 209)

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 216.

YES Go to step 8.

8. (from step 7)

Are the ac LED indicators on the rear of the power supply assemblies illuminated? Figure 74 on page 211 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 75 on page 211 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 74. SAN Volume Controller models 2145-8G4 and 2145-8F4 or 2145-8F2 ac and dc LED indicators on the rear panel




Figure 75. 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 239.

YES Go to step 10.

9. (from step 8 on page 210)

Is the power supply error LED on the rear of the SAN Volume Controller 2145-CF8 power supply assemblies illuminated? Figure 73 on page 210 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 207, step 8 on page 210, 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 239.

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 239.

11. (from step 1 on page 207)

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 management GUI might cause a loss of data in the node cache. If you are performing concurrent maintenance, this node must be deleted from the 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 201 to resolve the problem.

YES Go to step 14.

14. (from step 13 on page 212)

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.

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 9, “Using the maintenance analysis procedures,” on page 201.

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 216.

YES Go to step 2.

2. (from step 1)

Is the power LED on the operator-information panel continuously illuminated? Figure 76 shows the location of the power LED **1** on the operator-information panel.




Figure 76. 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 201 or go to “MAP 5700: Repair verification” on page 239 to verify the correct operation.

3. (from step 2)

Is the power LED on the operator-information panel flashing?

NO Go to step 5 on page 214.

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:

- a. Operator-information panel
- b. Operator-information panel cable
- c. System board

Verify the repair by continuing with “MAP 5700: Repair verification” on page 239.

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 239 to verify the correct operation.

5. (from step 3 on page 213)

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 216.

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 239.

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 77 on page 215 shows where the diagnostics LEDs are located on the system board.

NO Go to step 7 on page 215.

YES Replace the SAN Volume Controller 2145-8A4 parts in the following sequence:

- a. Operator-information panel
- b. Operator-information panel cable




Figure 77. 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 214)

Is the voltage regulator LED that is on the system board illuminated?

NO Go to step 8 on page 216.

YES Replace the system board.

8. (from step 7 on page 215)

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 239 to verify the correct operation.

9. (from step 1 on page 213)

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 management GUI might cause a loss of data in the node cache. If you are performing concurrent maintenance, this node must be deleted from the 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 201 to resolve the problem.

YES Go to step 12.

12. (from step 11)

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 9, “Using the maintenance analysis procedures,” on page 201.

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 78 shows an illustration of the front of the panel for the 2145 UPS-1U.




Figure 78. 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 41 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 41. 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 |

Table 41. 2145 UPS-1U error indicators (continued)

| [1] Load2 | [2] Load1 | [3] Alarm | [4] Battery | [5] Overload | [6] Power-on | Buzzer | Error condition |
|------------------|------------------|--------------|----------------|--------------|----------------|---------------------------------------|---|
| 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 |
| | | Flashing red | Amber | | Green | Solid on | Over temperature |
| | | Flashing red | Amber | Red | Green | | Output short circuit |

Table 41. 2145 UPS-1U error indicators (continued)

| [1] Load2 | [2] Load1 | [3] Alarm | [4] Battery | [5] Overload | [6] Power-on | Buzzer | Error condition |
|---|-----------|-----------|-------------|--------------|--------------|--------|-----------------|
| 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. | | | | | | | |

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.

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 222. Otherwise, complete these steps:

a. Restore main power to installation.

b. Verify the repair by continuing with “MAP 5250: 2145 UPS-1U repair verification” on page 221.

3. (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 221.

4. (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 on page 220.

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 223. 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 221.

5. (from step 4 on page 219)

Is the 2145 UPS-1U overload indicator illuminated solid red?

NO Go to step 6.

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 221.

6. (from step 5)

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 221.

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 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."

8. (from step 7 on page 220)

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."

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."

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 9, "Using the maintenance analysis procedures," on page 201.

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 201.

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 201.

4. (from step 3 on page 221)

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 239.

YES Continue with “MAP 5000: Start” on page 201.

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 9, “Using the maintenance analysis procedures,” on page 201.

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 216 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 223.

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 225. 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 222)
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 9, “Using the maintenance analysis procedures,” on page 201.

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 management 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 management 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 225.

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 on page 225.

YES Go to step 6.

6. (from step 5)

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 on page 225.

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 239.

8. (from steps 4 on page 224, 5 on page 224, 6 on page 224, and 7 on page 224)

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 management GUI
2. Use the CLI command `svctask stopcluster -node name`

It is preferable to use either the management GUI 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 volumes 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 volumes 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.

Using the management GUI to power off a cluster

This topic describes how to power off a cluster using the management GUI.

Perform the following steps to use the management GUI to power off a cluster:

1. Sign on to the IBM System Storage Productivity Center as an administrator and then launch the management GUI for the cluster that you are servicing.
2. Find the cluster that you are about to shut down.
If the nodes that you want to power off are shown as `Offline`, then the nodes are not participating in the cluster. In these circumstances, you must use the power button on the nodes to power off the nodes.
If the nodes that you want to power off are shown as `Online`, powering off the nodes can result in the dependent volumes to also go offline. Verify whether or not the nodes have any dependent volumes.
3. Select the node and click **Show Dependent Volumes**.
4. Make sure that the status of each volume in the I/O group is `Online`. You might need to view more than one page.
If any volumes are shown as degraded, only one node in the I/O is processing I/O requests for that volume. If that node is powered off, it impacts all the hosts that are submitting I/O requests to the degraded volume.
If any volumes are degraded and you believe that 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 volumes online. All the volumes should be online within 30 minutes of the partner node being powered off.
Note: If, after waiting 30 minutes, you have a degraded volume and all of the associated nodes and MDisks are online, contact the IBM Support Center for assistance.
Ensure that all volumes that are being used by hosts are online before you continue.
5. If possible, check that all the hosts that access the volumes that are managed by this I/O group are able to fail over to use paths that are provided by the other node in the group.
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 the volume.
6. If you have decided it is okay to continue and power off the nodes, select the cluster that you want to power off, and then click **Shut Down Cluster**.

- | 7. Click **OK**. If you have selected a node that is the last remaining node that provides access to a volume for example, a node that contains solid-state drives (SSDs) with unmirrored volumes, the Shutting Down a Node-Force panel is displayed with a list of volumes that will go offline if this node is shut down.
- | 8. Check that no host applications are accessing the volumes that will go offline; only continue with the shut down if the loss of access to these volumes 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 cluster 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 volumes that are managed by the I/O group. Ensure that the other node in the I/O group is online before you continue.

2. Issue the **svcinfoldependentvdisks** CLI command to list the volumes that are dependent on the status of a specified node.

```
| svcinfoldependentvdisks group1node1
```

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

| If the node goes offline or is removed from the cluster, the dependent volumes 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 volumes.

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 volumes 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 volumes will be taken offline. Use the **force** parameter with caution; access to data on node-dependent volumes 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 management GUI or the CLI commands, described in the previous topics, to power off an active node.

If you must use this method, notice in Figure 79 that each SAN Volume Controller model type has a power control button **1** on the front.




Figure 79. 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 volumes 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 9, “Using the maintenance analysis procedures,” on page 201.

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 207 or “MAP 5060: Power 2145-8A4” on page 213.

YES Go to step 2.

2. (from step 1)

Is the service controller error light that you see in Figure 80 illuminated and showing a solid amber?




Figure 80. 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 on page 230.

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 239.

3. (from step 2 on page 229)

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 239.

YES Go to step 4.

4. (from step 3)

Figure 81 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.



23njjj

Figure 81. 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 81?

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 239.

YES Press and hold the select button for five seconds to exit the test. Go to step 5.

5. Is the front-panel display now showing **Cluster:?**

NO Continue with “MAP 5000: Start” on page 201.

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 on page 230)

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 239.

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 239.

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 9, “Using the maintenance analysis procedures,” on page 201.

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 233

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 232.

NO Go to step 2.

2. **Is the cluster reporting cluster error 1400 either on the front panel or in the event log?**

YES Go to step 4 on page 232.

NO Go to step 3.

3. **Are you experiencing Ethernet performance issues?**

YES Go to step 9 on page 233.

NO Go to step 10 on page 233.

4. (from step 2 on page 231) **On all nodes perform the following actions:**

- 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 2145 Troubleshooting Guide*, Release 4.3.1.
- b. Press right until the top line displays Ethernet port 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 Ethernet port 2 and repeat the previous step.
- e. 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 233.

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.

7. (from step 6) **Perform the following actions:**

- a. Plug the Ethernet cable from that node into the Ethernet port 2 from a different node.
- b. If the Ethernet link light is illuminated when the cable is plugged into Ethernet port 2 of the other node, replace the system board of the original node.




Figure 82. 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 239.

8. (from step 5 or step 6) **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 239.

9. (from step 3 on page 231) **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 2145 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 231)

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:
- a. Use the Ethernet hub problem determination procedure to check for and resolve an Ethernet network connection problem. If you resolve a problem continue with “MAP 5700: Repair verification” on page 239.
 - b. Determine if similar Ethernet connection problems have occurred recently on this node. If they have, replace the system board.
 - c. Verify the repair by continuing with “MAP 5700: Repair verification” on page 239.
- YES** Verify the repair by continuing with “MAP 5700: Repair verification” on page 239.

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 management GUI 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:

1. Press and release the power button on the node that is displaying Node Error 540.

2. When Powering off is displayed on the front panel display, press the power button again.
3. Restarting is displayed.

The cluster will select a new configuration node. The management GUI 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 9, "Using the maintenance analysis procedures," on page 201.

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 238.

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 6, "Using the front panel of the SAN Volume Controller," on page 81.

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 small form-factor pluggable (SFP) transceiver 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 236.
- **Failed:** The port is not operational because of a hardware failure. Make a note of port-1. Go to step 10 on page 237.
- **Not installed:** This port is not installed. Make a note of port-1. Go to step 11 on page 238.

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 small form-factor pluggable (SFP) transceiver 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 236.
- **Failed:** The port is not operational because of a hardware failure. Make a note of port-2. Go to step 10 on page 237.
- **Not installed:** This port is not installed. Make a note of port-2. Go to step 11 on page 238.

YES Press and release the right button to display fibre-channel port-3. Go to step 5.

5. (from step 4 on page 234)

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 small form-factor pluggable (SFP) transceiver 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 236.
- **Failed:** The port is not operational because of a hardware failure. Make a note of port-3. Go to step 10 on page 237.
- **Not installed:** This port is not installed. Make a note of port-3. Go to step 11 on page 238.

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 small form-factor pluggable (SFP) transceiver 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 236.
- **Failed:** The port is not operational because of a hardware failure. Make a note of port-4. Go to step 9 on page 237.
- **Not installed:** This port is not installed. Make a note of port-4. Go to step 11 on page 238.

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 239.
- 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 transceiver, unless it has already been replaced.

Note: SAN Volume Controller nodes are supported with both longwave SFP transceivers and shortwave SFP transceivers. You must replace an SFP transceiver with the same type of SFP transceiver. If the SFP transceiver to replace is a longwave SFP transceiver, for example, you must provide a suitable replacement. Removing the wrong SFP transceiver could result in loss of data access. Refer to the Removing and replacing the fibre-channel SFP transceiver on a SAN Volume Controller node documentation to find out how to replace an SFP transceiver.

- 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 host bus adapter (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 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 239.

YES Verify the repair by continuing with “MAP 5700: Repair verification” on page 239.

- 8. (from steps 3 on page 234, 4 on page 234, 5 on page 235, and 6 on page 235)

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 237.

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 2145 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 on page 237.

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 239. Otherwise, go to step 9.

9. (from step 8 on page 236)

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 transceiver 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 transceiver.

Note: SAN Volume Controller nodes are supported with both longwave SFPs and shortwave SFPs. You must replace an SFP with the same type of SFP transceiver that you are replacing. If the SFP transceiver to replace is a longwave SFP transceiver, for example, you must provide a suitable replacement. Removing the wrong SFP transceiver could result in loss of data access. Refer to the Removing and replacing the fibre-channel SFP transceiver on a SAN Volume Controller node documentation to find out how to replace an SFP transceiver.

- 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 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 239.

10. (from steps 3 on page 234, 4 on page 234, 5 on page 235, and 6 on page 235)

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 transceiver.

Note: SAN Volume Controller nodes are supported with both longwave SFP transceivers and shortwave SFP transceivers. You must replace an SFP transceiver with the same type of SFP transceiver. If the SFP transceiver to replace is a longwave SFP transceiver, for example, you must provide a suitable replacement. Removing the wrong SFP transceiver could result in loss of data access. See the "Removing and replacing the fibre-channel SFP transceiver on a SAN Volume Controller node" documentation to find out how to replace an SFP transceiver.

- 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 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 239.

11. (from steps 3 on page 234, 4 on page 234, 5 on page 235, and 6 on page 235)

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 42.

Table 42. 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 HBA - Low profile |

b. If the problem is not fixed, replace the fibre-channel connection hardware in the order that is shown in Table 43.

Table 43. 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 239.

YES Verify the repair by continuing with “MAP 5700: Repair verification” on page 239.

12. (from step 2 on page 234)

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 Status: is displayed.
- b. Press and release the select button.
- c. Press the left or right button until FC Port-1 Speed: is displayed.
- d. Press and release the select button.
- e. Press 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.”

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 transceiver
 - Fibre-channel switch gigabit interface converter (GBIC) or SFP transceiver
 - 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.”

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 9, “Using the maintenance analysis procedures,” on page 201.

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 201.

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. 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 201.

YES Go to step 3.

3. (from step 2 on page 239)

Using the SAN Volume Controller application for the cluster you have just repaired, check the status of all configured volumes. **Do all volumes 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 volumes are showing a status of offline. Volumes will be held offline if SAN Volume Controller cannot confirm the integrity of the data. The volumes might be the target of a copy that did not complete, or cache write data that was not written back to disk might have been lost. Determine why the volume is offline. If the volume was the target of a copy that did not complete, you can start the copy again. Otherwise, write data might not 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 volume online, you must move all the offline disks to the recovery I/O group and then move them back to an active I/O group.

Go to "MAP 5000: Start" on page 201.

5. (from step 3)

You have successfully repaired the SAN Volume Controller.

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 9, "Using the maintenance analysis procedures," on page 201.

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

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 83 on page 241, on the SAN Volume Controller 2145-CF8 operator-information panel on or flashing?**




Figure 83. 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 201.

YES Go to step 2.

2. (from step 1 on page 240)

Press the release latch and open the light path diagnostics panel, which is shown in Figure 84.

Are one or more LEDs on the light path diagnostics panel on or flashing?




Figure 84. 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 239.

YES See Table 44 on page 243 and perform the action specified for the specific light path diagnostics LEDs. Then go to step 3 on page 245. Some actions will require that you observe the state of LEDs on the system board. Figure 85 on page 242 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 225 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. . Refer to the Removing the SAN Volume Controller from a rack documentation to find out how to perform the procedure.
- c. Remove the top cover.

- d. See Table 44 on page 243 and perform the action specified for the specific light path diagnostics LEDs. Then go to step 3 on page 245.




Figure 85. 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 44. 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"> 1. Turn off the node, pull the node forward in the rack, and remove the cover. Do not disconnect power from the node. 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"> • LED A: fans, disk drive, any solid-state drives (SSDs), or disk backplane • LED B: fibre-channel adapter and riser, all memory • LED C: disk controller, all memory • LED D: go to f. • LED E: High-speed SAS adapter and riser, if installed • LED AUX: fibre-channel adapter and high-speed SAS adapter, if installed 3. Restart the node to see whether the problem remains. 4. Reinstall each device that you removed in substep b one at a time. Start the node each time to isolate the failing device. 5. Replace any failing device. 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. 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. 8. If no failing device is isolated, replace the system board. |
| LOG | An error occurred. Connect a keyboard and a monitor. Check the IMM system event log and the system event log for information about the error. Replace any components that are identified in the event logs. |
| LINK | This is not used on the SAN Volume Controller 2145-CF8. Replace the system board. |

Table 44. Diagnostics panel LED prescribed actions (continued)

| Diagnostics panel LED | Action |
|-----------------------|---|
| 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"> 1. Check the power supply that has a lit amber LED. 2. Make sure that the power supplies are seated correctly. 3. Remove one of the power supplies to isolate the failed power supply. 4. Replace the failed power supply. |
| 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"> 1. Identify the failing adapter by checking the LEDs on the PCI slots. 2. If PCI slot 1 shows an error, replace the 4-port fibre-channel adapter assembly. 3. If PCI slot 2 shows an error, replace the high-speed SAS adapter assembly. 4. If the error is not resolved, replace the system board. |
| SP | <p>A service processor error was detected. Perform the following actions to resolve the problem:</p> <ol style="list-style-type: none"> 1. Remove power from the node. Reconnect the server to the power, and restart the node. 2. If the problem remains, replace the system board. |
| 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"> 1. Reseat the failing fan, which is indicated by a lit LED near the fan connector on the system board. 2. If the problem remains, replace the failing fan. |
| 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"> 1. Make sure that the heat sink is seated correctly. 2. Determine whether a fan has failed. If it has, replace it. 3. Verify that the ambient temperature is within normal operating specifications. 4. Make sure that airflow in and around the SAN Volume Controller 2145-CF8 is not obstructed. |
| 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"> 1. Check that all the memory DIMMs are correctly installed. 2. If any memory error LEDs are lit, replace the indicated memory module. 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. |
| 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"> 1. Operator-information panel 2. Operator-information panel cable 3. System board |

Table 44. 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"> 1. Check the system board error LEDs. 2. If CPU 1 error LED is lit, check that the microprocessor is correctly installed. 3. If the error persists, replace the microprocessor. 4. If the error persists, replace the system board. |
| VRM | <p>This is not used on the SAN Volume Controller 2145-CF8.</p> |
| 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"> 1. The system disk drive. 2. The disk backplane. |
| RAID | <p>This is not used on the SAN Volume Controller 2145-CF8.</p> |
| BRD | <p>An error occurred on the system board. Perform the following actions to resolve the problem:</p> <ol style="list-style-type: none"> 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"> • Battery • Missing PCI riser-card assembly. There must be a riser card in PCI slot 2 even if the optional adapter is not present. • Failed voltage regulator 2. Replace any failed or missing replacement components, such as the battery or PCI riser-card assembly. 3. If a voltage regulator fails, replace the system board. |

3. Continue with “MAP 5700: Repair verification” on page 239 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 86 on page 246, on the SAN Volume Controller 2145-8A4 operator-information panel on or flashing?**




Figure 86. SAN Volume Controller 2145-8A4 operator-information panel

1 Error LED

NO Reassess your symptoms and return to “MAP 5000: Start” on page 201.

YES Go to step 2.

2. (from step 1 on page 245)

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 225 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 248.

YES Identify any diagnostic LEDs on the system board that are on. Figure 87 on page 247 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 87. 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 246)

Are any diagnostic LEDs other than **13**, **14**, and **15** on the system board illuminated?

NO Go to step 5 on page 248.

YES Refer to Table 45 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 45. SAN Volume Controller 2145-8A4 diagnostics panel LED prescribed actions

| Diagnostics 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 247)

Replace the top cover and place the node in the rack. Refer to the Removing the SAN Volume Controller from a rack documentation to find out how to perform the procedure. Then continue with “MAP 5700: Repair verification” on page 239 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 88, on the SAN Volume Controller 2145-8G4 operator-information panel illuminated or flashing?**



Figure 88. 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 201.

YES Go to step 2.

2. (from step 1)

Press the release latch and open the light path diagnostics panel, which is shown in Figure 89.




Figure 89. 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 239.

YES See Table 46 on page 251 and perform the action specified for the specific light path diagnostics LEDs. Then go to step 3 on page 252. Some actions will require that you observe the state of LEDs on the system board. Figure 90 on page 250 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 225 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. Refer to the Removing the SAN Volume Controller from a rack documentation to find out how to perform the procedure.
- c. Remove the top cover and open the fan doors.
- d. Press the light path diagnostics button (**7** in Figure 90 on page 250).

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 90. 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 46. Diagnostics panel LED prescribed actions

| Diagnostics panel LED | Action |
|-----------------------|--|
| OVER SPEC | Replace parts in the following sequence: <ol style="list-style-type: none"> 1. Power supply 2. Power backplane 3. System board |
| 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"> 1. Power supply 2. Power backplane |
| 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"> 1. Power backplane 2. Operator-information panel 3. System board |
| 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"> 1. Failing DIMM 2. System board <p>Note: 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. |
| 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"> 1. Operator-information panel 2. System board |
| 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. |

Table 46. Diagnostics panel LED prescribed actions (continued)

| Diagnostics panel LED | Action |
|-----------------------|--|
| 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. |

3. Continue with “MAP 5700: Repair verification” on page 239 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:

1. Is the Error LED, shown in Figure 91, on the SAN Volume Controller 2145-8F2 or the SAN Volume Controller 2145-8F4 operator-information panel illuminated or flashing?




Figure 91. 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 201.

YES Go to step 2.

2. (from step 1)

Press the release latch and open the light path diagnostics panel, which is shown in Figure 92 on page 253.




Figure 92. 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 239.

- YES** Refer to Table 47 on page 254 and perform the action specified for the specific light path diagnostics LEDs, then go to step 3 on page 256. 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 93 on page 254. 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 225 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. Refer to the Removing the SAN Volume Controller from a rack documentation to find out how to perform the procedure.
 - c. Remove the top cover and open the fan doors.
 - d. Press the light path diagnostics button **1**. See Figure 93 on page 254.

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 93. 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 47. Diagnostics panel LED prescribed actions

| Diagnostics panel LED | Action |
|-----------------------|--------------------------|
| OVER SPEC | Replace the power supply |

Table 47. Diagnostics panel LED prescribed actions (continued)

| Diagnostics panel LED | Action |
|-----------------------|---|
| 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"> 1. Power supply 2. Power backplane |
| 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"> 1. Power backplane 2. Operator-information panel 3. Frame assembly |
| 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"> 1. Microprocessor 2. Frame assembly |
| 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"> 1. VRM 2. Frame assembly |
| 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"> 1. Component adjacent to the illuminated LED 2. Frame assembly |
| 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"> 1. Failing DIMM 2. Frame assembly <p>Note: 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> |
| 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"> 1. Operator-information panel 2. Frame assembly |

Table 47. Diagnostics panel LED prescribed actions (continued)

| Diagnostics panel LED | Action |
|-----------------------|---|
| FAN | Observe the LEDs on the fan backplanes. The fan adjacent to the failing LED is failing. Replace parts in the following sequence: <ol style="list-style-type: none"> 1. Fan 2. Fan backplane |
| 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 2145-8F2 or SAN Volume Controller 2145-8F4 is not obstructed. Replace the frame assembly. |
| 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 frame assembly. |
| PCI A | 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: <ol style="list-style-type: none"> 1. Operator-information panel 2. Frame assembly |
| PCI B | 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: <ol style="list-style-type: none"> 1. Fibre-channel adapter card 2. Frame assembly |
| PCI C | Replace the frame assembly. |

3. Continue with “MAP 5700: Repair verification” on page 239 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 9, “Using the maintenance analysis procedures,” on page 201.

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 94, is displayed continuously.



Figure 94. Hardware boot display

- The node rescue display, shown in Figure 95, is displayed continuously.



Figure 95. Node rescue display

- The boot progress is hung and an error 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 the Error LED on the operator-information panel illuminated or flashing?**

NO Go to step 2.

YES Go to “MAP 5800: Light path” on page 240 to resolve the problem.

2. (From step 1)

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:

- Ensure that the correct power cable assembly from the 2145 UPS-1U to the node is installed. The correct power cable assembly has tape that binds the cables together.
- 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 the Removing the SAN Volume Controller from a rack documentation to find out how to perform the procedure.
- Remove the top cover. See the Removing the top cover documentation to find out how to perform the procedure.
- If you have just replaced a FRU, ensure that the FRU is correctly placed and that all connections to the FRU are secure.
- Ensure that all memory modules are correctly installed and that the latches are fully closed. See the Replacing the memory modules (DIMM) documentation to find out how to perform the procedure.
- Ensure that the fibre-channel adapter cards are correctly installed. See the Replacing the fibre-channel adapter assembly documentation to find out how to perform the procedure.
- Ensure that the disk drive and its connectors are correctly installed. See the Replacing the disk drive documentation to find out how to perform the procedure.
- Ensure that the service controller is correctly installed. See the Replacing the service controller documentation to find out how to perform the procedure.
- Replace the top cover. See the Replacing the top cover documentation to find out how to perform the procedure.
- Place the node in the rack. See the Replacing the SAN Volume Controller in a rack documentation to find out how to perform the procedure.
- Turn on the node.

Does the boot operation still hang?

NO Verify the repair by continuing with “MAP 5700: Repair verification” on page 239.

YES Go to step 3.

3. (from step 2)

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 225.
- b. Connect the keyboard **1** and the display **2**. Figure 96 shows the location of the keyboard and monitor ports. Figure 97 shows the location of the keyboard and monitor ports on the 2145-CF8.




Figure 96. Keyboard and monitor ports on the SAN Volume Controller models 2145-8G4, 2145-8A4, 2145-8F4 and 2145-8F2




Figure 97. 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 in the rack, 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 239.

YES Go to step 4.

4. (from step 3 on page 257)

- 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 225.
 - 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 the Removing the SAN Volume Controller from a rack documentation to find out how to perform the procedure.
 - c. Remove the top cover. See the Removing the top cover documentation to find out how to perform the procedure.
 - 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 the Replacing the top cover documentation to find out how to perform the procedure.
 - h. Place the node in the rack. See the Replacing the SAN Volume Controller in a rack documentation to find out how to perform the procedure.
 - i. Turn on the node.
5. **Does the boot operation still hang with the booting display or has the boot operation progressed?**

Note: With the FRUs removed, the boot will hang with a different boot failure code.

NO Go to step 6 to replace the FRUs, one-at-a-time, until the failing FRU is isolated.

YES Go to step 7

6. (From step 5)
- Remove all hardware except the hardware that is necessary to power up. Continue to add in the FRUs one at a time and power on each time until the original failure is introduced.

Does the boot operation still hang?

NO Verify the repair by continuing with “MAP 5700: Repair verification” on page 239.

YES Go to step 7.

7. (from steps 4 on page 258 and 6)
- 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 225 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 the Replacing the SAN Volume Controller in a rack documentation to find out how to perform the procedure.
 - c. Remove the top cover. See the Replacing the top cover documentation to find out how to perform the procedure.
 - 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 module in slots 3 and 6 with any of the removed memory modules from 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 the Replacing the top cover documentation to find out how to perform the procedure.
- g. Place the node in the rack. See the Replacing the SAN Volume Controller in a rack documentation to find out how to perform the procedure.
- h. Turn on the node.

Does the boot operation still hang with the booting display or does the display progress beyond the initial booting panel?

NO Exchange the failing memory modules for new FRUs and verify the repair by continuing with “MAP 5700: Repair verification” on page 239.

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 239.

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

This list includes the major accessibility features in the management GUI:

- | • 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: JAWS 11.
- | • Most of the GUI features are accessible by using the keyboard. For those features that are not accessible, equivalent function is available by using the command-line interface (CLI).
- | • When setting or changing an IP address on the SAN Volume Controller front panel, you can disable the fast increase function to reduce the address scrolling speed 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 management GUI and help system from the keyboard by using the following key combinations:

- | • To navigate between different GUI panels, select the Low-graphics mode option on the GUI login panel. You can use this option to navigate to all the panels without manually typing the web addresses.
- | • To go to the next frame, press Ctrl+Tab.
- | • To move to the previous frame, press Shift+Ctrl+Tab.
- | • To navigate to the next link, button, or topic within a panel, press Tab inside a frame (page).
- | • To move to the previous link, button, or topic within a panel, press Shift+Tab.
- | • To select GUI objects, press Enter.
- | • To print the current page or active frame, press Ctrl+P.
- | • To expand a tree node, press the Right Arrow key. To collapse a tree node, press the Left Arrow key.
- | • To scroll all the way up, press Home; to scroll all the way down, press End.
- | • To go back, press Alt+Left Arrow key.
- | • To go forward, press Alt+Right Arrow key.
- | • For actions menus:
 - | – Press Tab to navigate to the grid header.
 - | – Press the Left or Right Arrow keys to reach the drop-down field.
 - | – Press Enter to open the drop-down menu.
 - | – Press the Up or Down Arrow keys to select the menu items.
 - | – Press Enter to launch the action.
- | • For filter panes:
 - | – Press Tab to navigate to the filter panes.
 - | – Press the Up or Down Arrow keys to change the filter or navigation for nonselection.
 - | – Press Tab to navigate to the magnifying glass icon in the filter pane and press Enter.
 - | – Type the filter text.

- | – Press Tab to navigate to the red X icon and press Enter to reset the filter.
- | • For information areas:
 - | – Press Tab to navigate to information areas.
 - | – Press Tab to navigate to the fields that are available for editing.
 - | – Type your edit and press Enter to issue the change command.

Accessing the publications

- | You can find the HTML version of the IBM System Storage SAN Volume Controller information at the
- | following website:

`publib.boulder.ibm.com/infocenter/svc/ic/index.jsp`

- | You can access this information using screen-reader software and a digital speech synthesizer to hear
- | what is displayed on the screen. The information was tested using the following screen reader: JAWS
- | Version 10 or later.

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.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 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 PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you 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 are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Adobe 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

This explains the Federal Communications Commission's (FCC) statement.

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, might cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors, or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

Industry Canada compliance statement

This Class A digital apparatus complies with ICES-003.

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

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.

Australia and New Zealand Class A Statement

Attention: This is a Class A product. In a domestic environment this product might cause radio interference in which case the user might be required to take adequate measures.

European Union Electromagnetic Compatibility Directive

This product is in conformity with the protection requirements of European Union (EU) Council Directive 2004/108/EC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

Attention: This is an EN 55022 Class A product. In a domestic environment this product might cause radio interference in which case the user might be required to take adequate measures.

Responsible Manufacturer:

International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
914-499-1900

European community contact:

IBM Technical Regulations, Department M456
IBM-Allee 1, 71137 Ehningen, Germany
Tel: +49 7032 15-2937
E-mail: tjahn@de.ibm.com

Germany Electromagnetic compatibility directive

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2004/108/EG zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55022 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung der IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung der IBM gesteckt/eingebaut werden.

EN 55022 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:

"Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem "Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)." Dies ist die Umsetzung der EU-Richtlinie 2004/108/EG in der Bundesrepublik Deutschland.

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

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

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:

International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:

IBM Deutschland
Technical Regulations, Department M456
IBM-Allee 1, 71137 Ehningen, Germany
Tel: +49 7032 15-2937
e-mail: tjahn@de.ibm.com

Generelle Informationen: Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 Klasse A.

Japan VCCI Council Class A statement

この装置は、クラス A 情報技術装置です。この装置を家庭環境で使用する
と電波妨害を引き起こすことがあります。この場合には使用者が適切な対策
を講ずるよう要求されることがあります。 VCCI-A

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 Communications Commission (KCC) Class A Statement

이기는 업무용으로 전자파 적합등록을 받은 기기
이오니, 판매자 또는 사용자는 이점을 주의하시기
바라며, 만약 잘못 구입하셨을 때에는 구입한 곳에
서 비업무용으로 교환하시기 바랍니다.

Russia Electromagnetic Interference (EMI) Class A Statement

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать
радиопомехи, для снижения которых необходимы
дополнительные меры

rusemi

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。

taieami

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
Email: <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
 - environment 48
 - test and alarm-reset button 46
- 2145 UPS-1U
 - alarm 45
 - circuit breakers 46
 - connecting 44
 - connectors 46
 - controls and indicators on the front panel 44
 - description of parts 46
 - dip switches 46
 - environment 47
 - Load segment 1 indicator 45
 - Load segment 2 indicator 45
 - MAP
 - 5150: 2145 UPS-1U 217
 - 5250: repair verification 221
 - on-battery indicator 45
 - on/off button 45
 - operation 43
 - overload indicator 45
 - ports not used 46
 - power-on indicator 45
 - service indicator 45
 - test and alarm-reset button 46
 - unused ports 46
- 2145-8A4 node
 - SAN Volume Controller 2145-8A4 34
- 2145-8F2 node
 - SAN Volume Controller 38
- 2145-8F4 node
 - SAN Volume Controller 38
- 2145-8G4 node
 - SAN Volume Controller 36
- 2145-CF8 node
 - SAN Volume Controller 31

A

- about this document
 - sending comments xiii
- ac and dc LEDs 28
- ac power switch, cabling 41
- accessibility
 - keyboard 261
 - repeat rate of up and down buttons 103, 261
 - shortcut keys 261
- accessing
 - cluster CLI 62
 - management GUI 56
 - service assistant 61
 - service CLI 63
- action menu options
 - front panel display 91
 - sequence 91

- action options
 - node
 - create cluster 95
- active status 87
- adding
 - nodes 58, 61
- address
 - MAC 89
- Address Resolution Protocol (ARP) 5
- addressing
 - configuration node 5

B

- back panel assembly
 - SAN Volume Controller 2145-8A4
 - connectors 19
 - indicators 19
 - SAN Volume Controller 2145-8F2
 - connectors 25
 - SAN Volume Controller 2145-8F4
 - connectors 23
 - SAN Volume Controller 2145-8G4
 - connectors 21
 - SAN Volume Controller 2145-CF8
 - connectors 17
 - indicators 17
- back-panel assembly
 - SAN Volume Controller 2145-8F2
 - indicators 25
 - SAN Volume Controller 2145-8F4
 - indicators 22
 - SAN Volume Controller 2145-8G4
 - indicators 20
- backing up cluster configuration files
 - creating 195
- backup configuration files
 - deleting
 - using the CLI 199
 - restoring 197
- battery
 - Charging, front panel display 82
 - power 83
- boot
 - codes, understanding 131
 - failed 81
 - progress indicator 81
- buttons, navigation 12

C

- Call Home 109, 112
- Canadian electronic emission notice 266
- charging 82
- circuit breakers
 - 2145 UPS-1U 46
 - requirements
 - SAN Volume Controller 2145-8A4 34

- circuit breakers (*continued*)
 - requirements (*continued*)
 - SAN Volume Controller 2145-8G4 36
 - SAN Volume Controller 2145-CF8 32
- CLI
 - cluster commands 61
 - service commands 62
- CLI commands
 - svcinfo lscluster
 - displaying cluster properties 72
- cluster
 - error codes 133
 - restore 189, 193
 - T3 recovery 189
 - T3 Recovery 193
- cluster CLI
 - accessing 62
 - when to use 61, 62
- cluster commands
 - CLI 61
- clusters
 - adding nodes 58, 61
 - backing up configuration file using the CLI 195
 - Call Home email 109, 112
 - deleting nodes 56
 - diagnose failures 88
 - error codes 133
 - IP address
 - configuration node 5
 - IP failover 5
 - IPv4 address 87
 - IPv6 address 88
 - metadata, saving 83
 - options 87
 - overview 4
 - properties 72
 - recovery codes 133
 - removing nodes 56
 - restoring backup configuration files 197
- clusterT3 recovery
 - restore 189
- codes
 - configuration events 116
 - event 113
 - informational event 113
 - node error
 - critical 132
 - noncritical 132
 - node rescue 132
- commands
 - svcconfig backup 195
 - svcconfig restore 197
- comments, sending xiii
- configuration
 - event codes 116
 - node failover 5
 - configuration node 5

- connecting
 - 2145 UPS-1U 44
- connectors
 - 2145 UPS-1U 46
 - SAN Volume Controller 2145-8A4 19
 - SAN Volume Controller 2145-8F2 25
 - SAN Volume Controller 2145-8F4 23
 - SAN Volume Controller 2145-8G4 21
 - SAN Volume Controller 2145-CF8 17
- contact information
 - European 268
 - Taiwan 269
- controls and indicators on the front panel
 - 2145 UPS
 - test and alarm-reset button 46
 - 2145 UPS-1U
 - alarm 45
 - illustration 44
 - Load segment 1 indicator 45
 - Load segment 2 indicator 45
 - on-battery indicator 45
 - on/off button 45
 - overload indicator 45
 - power-on indicator 45
 - test and alarm-reset button 46
 - front-panel display 12
 - SAN Volume Controller
 - navigation buttons 12
 - node status LED 11
 - select button 12
 - SAN Volume Controller 2145-8A4
 - illustration 10
 - operator-information panel 14
 - SAN Volume Controller 2145-8F2
 - error LED 13
 - illustration 11
 - operator information panel 15
 - SAN Volume Controller 2145-8F4
 - illustration 11
 - operator information panel 15
 - SAN Volume Controller 2145-8G4
 - illustration 10
 - operator information panel 14
 - SAN Volume Controller 2145-CF8
 - illustration 9
 - operator-information panel 13
 - status indicators
 - action menu options 91
 - boot failed 81
 - boot progress 81
 - charging 82
 - error codes 82
 - hardware boot 82
 - menu options 85
 - node rescue request 82
 - power failure 83
 - powering off 83
 - recovering 83
 - restarting 84
 - shutting down 84
- create cluster
 - action option 95
 - error codes 133
- critical
 - node errors 132

D

- defining FRUs
 - for the redundant ac-power switch 54
 - for the SAN Volume Controller 49
- degraded status 87
- deleting
 - backup configuration files
 - using the CLI 199
 - nodes 56
- determining
 - hardware boot failure 131
 - SAN problem 186
- Deutschsprachiger EU Hinweis 266
- diagnosing problems
 - through event logs 107
 - through error codes 107
 - with SAN Volume Controller 107
- display on front panel
 - Change WWNN 101
 - Enter Service 101
 - Exit Actions 103
 - Exit Service 102
 - IPv6 address 88
 - Node rescue request 199
 - Node WWNN 89
 - overview 12
 - Paced Upgrade 102
 - Recover Cluster 102
 - Rescue Node 103
 - Service Address 89
 - Service DHCPv4 100
 - Service DHCPv6 100
 - status indicators
 - action menu options 91
 - boot failed 81
 - boot progress 81
 - charging 82
 - error codes 82
 - hardware boot 82
 - menu options 85
 - node rescue request 82
 - power failure 83
 - powering off 83
 - recovering 83
 - restarting 84
 - shutting down 84
 - validate WWNN? 84
 - version 89
- displaying
 - IPv6 address 88
- displaying vital product data 71
- documentation
 - improvement xiii

E

- electronic emission notices
 - Avis de conformité à la réglementation d'Industrie Canada 266
 - Deutschsprachiger EU Hinweis 266
 - European Union (EU) 266
 - Federal Communications Commission (FCC) 265
 - French Canadian 266

- electronic emission notices (*continued*)
 - Germany 266
 - Industry Canada 266
 - International Electrotechnical Commission (IEC) 268
 - Japanese Voluntary Control Council for Interference (VCCI) 267
 - Korean 268
 - New Zealand 266
 - People's Republic of China 267
 - Taiwan 268
 - United Kingdom 268
- email
 - Call Home 111, 112
 - inventory information 112
- EMC statement, People's Republic of China 267
- Enter Service
 - option 101
- error codes 122
 - front panel display 82
 - understanding 113
- error events 108
- error ID 122
- error LED 13
- errors
 - logs
 - describing the fields 109
 - error events 108
 - managing 108
 - understanding 108
 - viewing 109
 - node 132
- Ethernet
 - activity LED 27
 - link failures 5, 231
 - link LED 27
 - MAP 5500 231
 - port 89
 - Ethernet activity LED 17
 - European contact information 268
 - European Union (EU), EMC Directive conformance statement 266
- event
 - codes 113
 - informational 113
- event notification 109, 112
- events
 - codes
 - configuration 116
 - reporting 107
- examples
 - redundant ac power switch
 - cabling 41
 - SAN Volume Controller cluster in a SAN fabric 6
- Exit Actions
 - option 103
- Exit Service
 - option 102

F

- fabric, SAN 6
- failover, configuration node 5

- FCC (Federal Communications Commission) electronic emission notice 265
- Federal Communications Commission (FCC) electronic emission notice 265
- Fibre Channel
 - link failures 186
 - port numbers 29
 - SFP transceiver 186
- fibre-channel
 - LEDs 26
 - MAP 234
 - port menu option 90
- field replaceable units
 - redundant ac-power switch
 - describing 54
 - SAN Volume Controller
 - describing 49
 - disk drive assembly 49
 - disk drive cables 49
 - Ethernet cable 49
 - fan assembly 49
 - fibre-channel cable 49
 - fibre-channel SFP transceiver 49
 - frame assembly 49
 - front panel 49
 - operator-information panel 49
 - power cable assembly 49
 - service controller 49
 - system board assembly 49
- fields
 - cluster 78
 - description for the cluster vital product data 78
 - description for the node vital product data 73
 - device 73
 - event log 109
 - fibre-adapter card 73
 - front panel 73
 - memory module 73
 - processor 73
 - processor cache 73
 - software 73
 - system board 73
 - uninterruptible power supply 73
- fix
 - errors 190
- French Canadian electronic emission notice 266
- front panel
 - 2145 UPS-1U 44
 - action menu options 91
 - booting 104
 - buttons and indicators 81
 - charging 104
 - display 12
 - ID 13
 - menu options 85
 - Ethernet 89
 - fibre-channel port-1 through 4 90
 - IPv4 address 87
 - IPv6 address 88
 - Language? 103
 - node 88
 - version 89
 - power failure 104

- front panel (*continued*)
 - powering off the SAN Volume Controller 104
 - recovering 104
 - SAN Volume Controller 81

G

- gateway
 - menu option 88
 - node option 97, 98
- Germany electronic emission compliance statement 266

H

- hard-disk drive activity LED 15
- hardware
 - boot 82, 256
 - boot failure 131
 - components 9
 - failure 82
 - node 9

I

- I/O operations, stopped 83
- identification
 - label, node 13
 - name 88
 - number 88
- IEC (International Electrotechnical Commission) electronic emission notice 268
- inactive status 87
- indicators and controls on the front panel
 - 2145 UPS
 - test and alarm-reset button 46
 - 2145 UPS-1U
 - alarm 45
 - illustration 44
 - Load segment 1 indicator 45
 - Load segment 2 indicator 45
 - on-battery indicator 45
 - on/off button 45
 - overload indicator 45
 - power-on indicator 45
 - test and alarm-reset button 46
 - SAN Volume Controller
 - navigation buttons 12
 - node status LED 11
 - select button 12
 - SAN Volume Controller 2145-8A4
 - illustration 10
 - operator-information panel 14
 - SAN Volume Controller 2145-8F2
 - error LED 13
 - illustration 11
 - operator information panel 15
 - SAN Volume Controller 2145-8F4
 - illustration 11
 - operator information panel 15
 - SAN Volume Controller 2145-8G4
 - illustration 10
 - operator information panel 14

- indicators and controls on the front panel (*continued*)
 - SAN Volume Controller 2145-CF8
 - illustration 9
 - operator-information panel 13
 - status indicators
 - action menu options 91
 - boot failed 81
 - boot progress 81
 - charging 82
 - error codes 82
 - hardware boot 82
 - menu options 85
 - node rescue request 82
 - power failure 83
 - powering off 83
 - recovering 83
 - restarting 84
 - shutting down 84
 - indicators on the rear panel
 - ac and dc LEDs 28, 29
 - Ethernet
 - activity LED 27
 - link LED 27
 - fibre-channel LEDs 26
 - power-supply error LED 28
 - power, location, and system-error LEDs 27
 - SAN Volume Controller 2145-8A4
 - Ethernet activity LED 27
 - SAN Volume Controller 2145-8G4
 - Ethernet activity LED 27
 - information
 - center xi
 - information, system
 - LED 16
 - informational event 108
 - International Electrotechnical Commission (IEC) electronic emission notice 268
 - inventory information 109, 112
 - IP address
 - cluster 87, 88
 - IPv6 88
 - service 99
 - IPv4 address 87
 - IPv6
 - address 88
 - gateway menu option 88
 - prefix mask menu option 88
 - iSCSI
 - link problems 187

J

- Japanese electronic emission notice 267

K

- keyboard 261
- Korean electronic emission statement 268

L

- language menu selection options 103

- LEDs
 - ac and dc 28, 29
 - diagnostics 240
 - Ethernet
 - activity 27
 - link 27
 - Ethernet activity LED 17
 - fibre-channel 26
 - hard-disk drive activity 15
 - location 16, 27
 - power 16, 27
 - power-supply error 28
 - rear-panel indicators 17, 19, 20, 22, 25
 - SAN Volume Controller 2145-8A4 19
 - SAN Volume Controller 2145-8F2 25
 - SAN Volume Controller 2145-8F4 22
 - SAN Volume Controller 2145-8G4 20
 - SAN Volume Controller 2145-CF8 17
 - system information 16
 - system-error 15, 27
- legal notices 263
- light path MAP 240
- link failures
 - Fibre Channel 186
- link problems
 - iSCSI 187
- Load segment 1 indicator 45
- Load segment 2 indicator 45
- locator LED 16
- log files
 - viewing 109

M

- MAC address 89
- maintenance analysis procedures (MAPs)
 - 2145 UPS-1U 217
 - Ethernet 231
 - fibre-channel 234
 - front panel 229
 - hardware boot 256
 - light path 240
 - overview 201
 - power
 - SAN Volume Controller 2145-8A4 213
 - SAN Volume Controller 2145-8F2 207
 - SAN Volume Controller 2145-8F4 207
 - SAN Volume Controller 2145-8G4 207
 - repair verification 239
 - start 201
- management GUI
 - accessing 56
 - shut down a node 225
- management GUI interface
 - when to use 55
- managing
 - event log 108
- MAP
 - 5000: Start 201
 - 5050: Power 2145-CF8, 2145-8G4, 2145-8F4, and 2145-8F2 207
 - 5060: Power 2145-8A4 213

- MAP (*continued*)
 - 5150: 2145 UPS-1U 217
 - 5250: 2145 UPS-1U repair verification 221
 - 5320: Redundant ac power 222
 - 5340: Redundant ac power verification 223
 - 5400: Front panel 229
 - 5500: Ethernet 231
 - 5600: Fibre-channel 234
 - 5700: Repair verification 239
 - 5800: Light path 240
 - 5900: Hardware boot 256
 - power off SAN Volume Controller node 225
- MAPs (maintenance analysis procedures)
 - 2145 UPS-1U 217
 - 2145 UPS-1U repair verification 221
 - Ethernet 231
 - fibre-channel 234
 - front panel 229
 - hardware boot 256
 - light path 240
 - power
 - SAN Volume Controller 2145-8A4 213
 - SAN Volume Controller 2145-8F2 207
 - SAN Volume Controller 2145-8F4 207
 - SAN Volume Controller 2145-8G4 207
 - SAN Volume Controller 2145-CF8 207
 - power off 225
 - redundant ac power 222, 223
 - repair verification 239
 - start 201
 - using 201
- media access control (MAC) address 89
- menu options
 - clusters
 - gateway 88
 - IPv4 address 87
 - IPv4 gateway 88
 - IPv4 subnet 87
 - IPv6 address 88
 - IPv6 prefix 88
 - options 87
 - reset password 103
 - status 87, 88
 - Ethernet
 - MAC address 89
 - port 89
 - speed 89
 - fibre-channel port-1 through 4 90
 - front panel display 85
 - IPv4 gateway 88
 - IPv6 gateway 88
 - IPv6 prefix 88
 - Language? 103
 - node
 - options 88
 - status 88
 - SAN Volume Controller
 - active 87
 - degraded 87

- menu options (*continued*)
 - SAN Volume Controller (*continued*)
 - inactive 87
 - IPv4 subnet 87
 - sequence 85
- message classification 134

N

- navigation
 - buttons 12
 - create cluster 95
 - Language? 103
 - recover cluster 103
- New Zealand electronic emission statement 266
- node
 - create cluster 95
 - failover 5
 - options
 - create cluster? 95
 - gateway 98
 - IPv4 address 96
 - IPv4 confirm create? 97
 - IPv4 gateway 97
 - IPv4 subnet mask 96
 - IPv6 address 97
 - IPv6 Confirm Create? 99
 - IPv6 prefix 98
 - Remove Cluster? 102
 - status 88
 - subnet mask 96
 - rescue request 82
 - software failure 207, 213
- node canisters
 - configuration 4
- node status LED 11
- nodes
 - adding 58, 61
 - cache data, saving 83
 - configuration 4
 - addressing 5
 - failover 5
 - deleting 56
 - hard disk drive failure 82
 - identification label 13
 - options
 - main 88
 - removing 56
 - rescue
 - performing 199
 - rescue codes
 - understanding 132
 - viewing
 - general details 71
 - vital product data 71
- noncritical
 - node errors 132
- not used
 - 2145 UPS-1U ports 46
 - location LED 27
- notifications
 - Call Home information 112
 - inventory information 112
 - sending 109
- number range 134

O

- object classes and instances 122
 - object codes 122
 - object types 122
 - on/off button 45
 - operator information panel
 - locator LED 16
 - SAN Volume Controller 2145-8F2 15
 - SAN Volume Controller 2145-8F4 15
 - SAN Volume Controller 2145-8G4 14
 - system-information LED 16
 - operator-information panel
 - Ethernet activity LED 17
 - hard-disk drive activity LED 15
 - power button 16
 - power LED 16
 - release latch 16
 - reset button 15
 - SAN Volume Controller 2145-8A4 14
 - SAN Volume Controller 2145-CF8 13
 - system-error LED 15
 - overload indicator 45
 - overview
 - redundant ac-power switch 40
 - SAN fabric 6
 - SAN Volume Controller 1
 - vital product data 71
- ## P
- Paced Upgrade
 - option 102
 - panel
 - front 12
 - name 13
 - operator information
 - SAN Volume Controller 2145-8A4 14
 - SAN Volume Controller 2145-8F2 15
 - SAN Volume Controller 2145-8F4 15
 - SAN Volume Controller 2145-8G4 14
 - SAN Volume Controller 2145-CF8 13
 - rear
 - SAN Volume Controller 2145-8A4 19
 - SAN Volume Controller 2145-8F2 25
 - SAN Volume Controller 2145-8F4 22
 - SAN Volume Controller 2145-8G4 20
 - SAN Volume Controller 2145-CF8 17
 - passwords
 - resetting 103
 - People's Republic of China, electronic emission statement 267
 - physical characteristics
 - 2145 UPS 48
 - 2145 UPS-1U 47
 - redundant ac-power switch 40
 - physical characteristics (*continued*)
 - SAN Volume Controller 2145-8A4
 - connectors 19
 - SAN Volume Controller 2145-8F2
 - connectors 25
 - SAN Volume Controller 2145-8F4
 - connectors 23
 - SAN Volume Controller 2145-8G4
 - connectors 21
 - SAN Volume Controller 2145-CF8
 - connectors 17
 - service ports 18
 - unused ports 18
 - port speed
 - fibre-channel 90
 - ports
 - Ethernet 27
 - not used
 - 2145 UPS-1U 46
 - SAN Volume Controller 2145-8A4 19
 - SAN Volume Controller 2145-8F4 23
 - SAN Volume Controller 2145-8G4 21
 - port names, worldwide 29
 - port numbers, Fibre Channel 29
 - SAN Volume Controller 2145-CF8 17
 - POST (power-on self-test) 108
 - power
 - button 16
 - controls 104
 - failure 83
 - off
 - operation 83
 - requirements
 - SAN Volume Controller 2145-8A4 34
 - SAN Volume Controller 2145-8F2 38
 - SAN Volume Controller 2145-8F4 38
 - SAN Volume Controller 2145-8G4 36
 - SAN Volume Controller 2145-CF8 32
 - restored 83
 - switch, failure 207, 213
 - uninterruptible power supply 104
 - power LED 16
 - Power MAP 2145-8A4 213
 - Power MAP 2145-CF8, 2145-8G4, 2145-8F4, and 2145-8F2 207
 - power off
 - SAN Volume Controller 225
 - power-supply error LED 28
 - preparing
 - SAN Volume Controller environment 31
 - uninterruptible power supply environment 47

R

- reader feedback, sending xiii
- rear-panel indicators
 - SAN Volume Controller 2145-8A4 19

- rear-panel indicators (*continued*)
 - SAN Volume Controller 2145-8F2 25
 - SAN Volume Controller 2145-8F4 22
 - SAN Volume Controller 2145-8G4 20
 - SAN Volume Controller 2145-CF8 17
- Recover Cluster
 - option 102
- recovering
 - front panel display 83
 - offline virtual disks (volumes) using CLI 65
- recovery
 - cluster
 - starting 193
 - when to run 189
- redundant ac power switch
 - cabling 41
 - examples 41
- redundant ac-power switch
 - environment preparation 40
 - field replaceable units 54
 - MAP 222, 223
 - overview 40
 - problems 222
 - specifications 40
 - verifying 223
- related information xi
- release latch 16
- removing
 - 550 errors 190, 191
 - 578 errors 190, 191
 - node from a cluster 102
 - nodes 56
- Repair verification MAP 239
- repairing
 - space-efficient VDisk 64
- reporting
 - events 107
- requirements
 - 2145 UPS-1U 43
 - ac voltage 31, 33, 34, 35, 36, 37
 - circuit breakers 32, 34, 36
 - electrical 31, 32, 34, 36
 - power 32, 34, 36
 - SAN Volume Controller 2145-8A4 34
 - SAN Volume Controller 2145-8G4 36
 - SAN Volume Controller 2145-CF8 31
- rescue
 - node
 - codes, understanding 132
 - performing 199
- Rescue Node
 - option 103
- reset button 15
- reset password menu option 103
 - navigation 103
 - resetting the password 103
- resetting passwords 103
- restore
 - cluster 189, 193

S

- SAN (storage area network)
 - fabric overview 6
 - problem determination 186

- SAN Volume Controller 2145 UPS-1U 44
 - action options
 - create cluster 95
 - field replaceable units
 - 4-port fibre-channel adapter 49
 - 40×40×28 fan 49
 - 40×40×56 fan 49
 - alcohol wipe 49
 - CMOS battery 49
 - disk backplane 49
 - disk controller 49
 - disk drive assembly 49
 - disk drive cables 49
 - disk power cable 49
 - disk signal cable 49
 - Ethernet cable 49
 - fan assembly 49
 - fan power cable 49
 - fibre-channel adapter assembly 49
 - fibre-channel cable 49
 - fibre-channel HBA 49
 - frame assembly 49
 - front panel 49
 - memory module 49
 - microprocessor 49
 - operator-information panel 49
 - power backplane 49
 - power cable assembly 49
 - power supply assembly 49
 - riser card, PCI 49
 - riser card, PCI Express 49
 - service controller 49
 - service controller cable 49
 - system board 49
 - thermal grease 49
 - voltage regulator module 49
 - front-panel display 81
 - hardware 1
 - hardware components 9
 - menu options
 - active 87
 - degraded 87
 - Ethernet 89
 - fibre-channel port-1 through 4 90
 - inactive 87
 - IPv4 subnet 87
 - IPv6 prefix 88
 - Language? 103
 - node 88
 - node 9
 - overview 1
 - power control 104
 - power off 225
 - power-on self-test 108
 - preparing environment 31
 - properties 71
 - software
 - overview 1
- SAN Volume Controller 2145-8A4
 - additional space requirements 35
 - air temperature without redundant ac power 34
 - circuit breaker requirements 34
 - connectors 19

- SAN Volume Controller 2145-8A4
 - (continued)*
 - controls and indicators on the front panel 10
 - dimensions and weight 35
 - fibre-channel
 - port number 30
 - heat output of node 35
 - humidity with redundant ac power 35
 - humidity without redundant ac power 34
 - indicators and controls on the front panel 10
 - input-voltage requirements 34
 - light path MAP 245
 - MAP 5800: Light path 245
 - nodes
 - heat output 35
 - not used, service ports 19
 - operator-information panel 14
 - ports 19
 - power requirements for each node 34
 - product characteristics 34
 - rear-panel indicators 19
 - requirements 34
 - specifications 34
 - temperature with redundant ac power 35
 - weight and dimensions 35
- SAN Volume Controller 2145-8F2
 - air temperature 38
 - connectors 25
 - controls and indicators on the front panel 11
 - dimensions and weight 38
 - fibre-channel port number 31
 - heat output 38
 - humidity 38
 - indicators and controls on the front panel 11
 - light path MAP 252
 - MAP 5800: Light path 252
 - operator information panel 15
 - product characteristics 38
 - rear-panel indicators 25
 - specifications 38
 - weight and dimensions 38
- SAN Volume Controller 2145-8F4
 - air temperature 38
 - connectors 23
 - controls and indicators on the front panel 11
 - dimensions and weight 38
 - fibre-channel
 - port number 30
 - heat output 38
 - humidity 38
 - indicators and controls on the front panel 11
 - light path MAP 252
 - MAP 5800: Light path 252
 - operator information panel 15
 - product characteristics 38
 - rear-panel indicators 22
 - specifications 38

- SAN Volume Controller 2145-8F4
 - (continued)*
 - weight and dimensions 38
- SAN Volume Controller 2145-8G4
 - additional space requirements 37
 - air temperature without redundant ac power 36
 - circuit breaker requirements 36
 - connectors 21
 - controls and indicators on the front panel 10
 - dimensions and weight 37
 - fibre-channel
 - port number 30
 - heat output of node 37
 - humidity with redundant ac power 37
 - humidity without redundant ac power 36
 - indicators and controls on the front panel 10
 - input-voltage requirements 36
 - light path MAP 248
 - MAP 5800: Light path 248
 - nodes
 - heat output 37
 - not used, service ports 21
 - operator information panel 14
 - ports 21
 - power requirements for each node 36
 - product characteristics 36
 - rear-panel indicators 20
 - requirements 36
 - specifications 36
 - temperature with redundant ac power 37
 - weight and dimensions 37
- SAN Volume Controller 2145-CF8
 - additional space requirements 33
 - air temperature without redundant ac power 32
 - circuit breaker requirements 32
 - connectors 17
 - controls and indicators on the front panel 9
 - dimensions and weight 33
 - heat output of node 33
 - humidity with redundant ac power 33
 - humidity without redundant ac power 32
 - indicators and controls on the front panel 9
 - input-voltage requirements 31
 - light path MAP 240
 - MAP 5800: Light path 240
 - nodes
 - heat output 33
 - operator-information panel 13
 - ports 17
 - power requirements for each node 32
 - product characteristics 31
 - rear-panel indicators 17
 - requirements 31
 - service ports 18

- SAN Volume Controller 2145-CF8
 - (continued)
 - specifications 31
 - temperature with redundant ac power 33
 - unused ports 18
 - weight and dimensions 33
- SAN Volume Controller library
 - related publications xi
- self-test, power-on 108
- sending
 - comments xiii
- serial number 12
- service
 - actions, uninterruptible power supply 43
- service address
 - navigation 99
 - options 99
- Service Address
 - option 89
- service assistant
 - accessing 61
 - interface 60
 - when to use 60
- service CLI
 - accessing 63
- service commands
 - CLI 62
- service controller
 - replacing
 - validate WWNN 84
- Service DHCPv4
 - option 100
- Service DHCPv6
 - option 100
- service ports
 - SAN Volume Controller 2145-CF8 18
- shortcut keys 261
- shutting down
 - front panel display 84
- SNMP traps 109
- software
 - failure, MAP 5050 207
 - failure, MAP 5060 213
 - overview 1
 - version
 - display 89
- space requirements
 - SAN Volume Controller 2145-8A4 35
 - SAN Volume Controller 2145-8G4 37
 - SAN Volume Controller 2145-CF8 33
- specifications
 - redundant ac-power switch 40
- speed
 - fibre-channel port 90
- Start MAP 201
- starting
 - cluster recovery 191, 193
 - T3 recovery 191
- status
 - active 87
 - degraded 87
 - inactive 87
 - operational 87, 88
- storage area network (SAN)
 - fabric overview 6

- storage area network (SAN) (continued)
 - problem determination 186
- storage systems
 - servicing 187
- subnet
 - menu option 87
- subnet mask
 - node option 96
- summary of changes ix, x
- switches
 - 2145 UPS-1U 46
 - redundant ac power 40
- syslog
 - messages 110
- syslog messages 109
- system-error LED 15

T

- T3 recovery
 - removing
 - 550 errors 190, 191
 - 578 errors 190, 191
 - restore
 - cluster 189
 - starting 191 |
 - when to run 189
- T3 Recovery
 - what to check 193
- Taiwan
 - contact information 269
 - electronic emission notice 268
- test and alarm-reset button 46
- trademarks 265
- troubleshooting
 - event notification email 109, 112
 - SAN failures 186
 - using error logs 82
 - using the front panel 81

U

- understanding
 - event log 108
 - cluster error codes 133
 - cluster recovery codes 133
 - create cluster error codes 133
 - error codes 113
 - fields for the cluster vital product data 78
 - fields for the node vital product data 73
 - node rescue codes 132
- uninterruptible power supply
 - 2145 UPS
 - environment 48
 - 2145 UPS-1U
 - controls and indicators 44
 - environment 47
 - operation 43
 - overview 43
 - front panel MAP 229
 - operation 43
 - overview 43
 - preparing environment 47

- United Kingdom electronic emission notice 268
- unused ports
 - 2145 UPS-1U 46
 - SAN Volume Controller 2145-8A4 19
 - SAN Volume Controller 2145-8F4 23
 - SAN Volume Controller 2145-8G4 21
 - SAN Volume Controller 2145-CF8 18
- using
 - CLI 63
 - error code tables 113
 - GUI interfaces 55
 - management GUI 55
 - service assistant 60

V

- validating
 - VDisk copies 63
- VDisk copies
 - validating 63
- VDisks (volumes)
 - recovering from offline
 - using CLI 65
- viewing
 - event log 109
- vital product data (VPD)
 - displaying 71
 - overview 71
 - understanding the fields for the cluster 78
 - understanding the fields for the node 73
 - viewing
 - nodes 71
- volumes (VDisks)
 - recovering from offline
 - using CLI 65
- VPD (vital product data)
 - displaying 71
 - overview 71
 - understanding the fields for the cluster 78
 - understanding the fields for the node 73

W

- when to use
 - cluster CLI 61, 62
 - management GUI interface 55
 - service assistant 60
- worldwide node names
 - change 101
 - choose 84
 - display 89
 - node, front panel display 89, 101
 - validate, front panel display 84
- worldwide port names (WWPNs)
 - description 29



Printed in USA

GC27-2284-00

