



IBM Storwize V7000

Troubleshooting, Recovery, and Maintenance Guide

Version 6.1.0

GC27-2291-00





IBM Storwize V7000

Troubleshooting, Recovery, and Maintenance Guide

Version 6.1.0

GC27-2291-00

Note

Before using this information and the product it supports, read the general information in "Notices" on page 119, the information in the "Safety and environmental notices" on page ix, as well as the information in the *IBM Environmental Notices and User Guide* on the documentation CD.

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

This edition applies to GC27-2291-00.

© **Copyright IBM Corporation 2010.**

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

Contents

Figures v

Tables vii

Safety and environmental notices . . . ix

Sound pressure ix

About this guide xi

Who should use this guide xi

Emphasis xi

Storwize V7000 library and related publications . . . xi

How to order IBM publications xii

Sending your comments xii

Chapter 1. Storwize V7000 hardware components 1

Storwize V7000 components in the front of the enclosure 1

Drives 1

Drive indicators 2

Enclosure end cap indicators 3

Storwize V7000 components in the rear of the enclosure 4

Power supply unit and battery for the control enclosure 5

Power supply unit for the expansion enclosure . . . 6

Node canister ports and indicators 7

Expansion canister ports and indicators 13

Chapter 2. Best practices for troubleshooting 17

Record access information 17

Follow power management procedures 17

Set up event notifications 18

Set up inventory reporting 18

Back up your data 18

Manage your spare and failed drives 19

Resolve alerts in a timely manner 19

Keep your software up to date 19

Keep your records up to date 19

Subscribe to support notifications 19

Know your IBM warranty and maintenance agreement details 20

Chapter 3. Understanding the Storwize V7000 battery operation for the control enclosure 21

Maintenance discharge cycles 22

Chapter 4. Storwize V7000 user interfaces for servicing your system . . 25

Management GUI interface 25

When to use the management GUI 25

Accessing the management GUI 26

Service assistant interface 26

When to use the service assistant 26

Accessing the service assistant 27

Cluster command-line interface 27

When to use the cluster CLI 28

Accessing the cluster CLI 28

Service command-line interface 28

When to use the service CLI 28

Accessing the service CLI 28

USB key and Initialization tool interface 28

When to use the USB key 28

Using a USB key 29

Using the initialization tool 29

satask.txt commands 30

Chapter 5. Resolving a problem 35

Start here: Use the management GUI recommended actions 35

Problem: Management IP address unknown . . . 36

Problem: Unable to connect to the management GUI 36

Problem: Unable to log on to the management GUI 37

Problem: Cannot create a cluster 37

Problem: Unknown service address of a node canister 38

Problem: Cannot connect to the service assistant . . 38

Problem: Management GUI or service assistant does not display correctly 39

Problem: Node canister location error 39

Problem: SAS cabling not valid 39

Problem: New expansion enclosure not detected . . 40

Problem: Mirrored volume copies no longer identical 40

Procedure: Resetting superuser password 40

Procedure: Identifying which enclosure or canister to service 41

Procedure: Checking the status of your system . . 42

Procedure: Getting node canister and cluster information using the service assistant 42

Procedure: Getting node canister and cluster information using a USB key 42

Procedure: Understanding the system status using the LEDs 43

Procedure: Finding the status of the Ethernet connections 47

Procedure: Removing cluster data from a node canister 47

Procedure: Deleting a cluster completely 48

Procedure: Fixing node errors 48

Procedure: Changing the service IP address of a node 48

Procedure: Initializing a cluster with a USB key without using the initialization tool 49

Procedure: Initializing a cluster using the service assistant 50

Procedure: Accessing a canister using a directly attached Ethernet cable	51
Procedure: Reseating a node canister	51
Procedure: Powering off your system.	52
Procedure: Collecting information for support	52
SAN problem determination.	52
Fibre Channel link failures	53
Servicing storage systems.	53

Chapter 6. Recovery procedures. 55

Recover cluster procedure	55
When to run the recover cluster procedure	55
Fix hardware errors.	56
Removing cluster information for nodes with error code 550 or error code 578 using the service assistant	56
Performing cluster recovery using the service assistant	57
What to check after running cluster recovery	58
Backing up and restoring the cluster configuration	58
Backing up the cluster configuration using the CLI	59
Restoring the cluster configuration	61
Deleting backup configuration files using the CLI	63

Chapter 7. Removing and replacing parts 65

Preparing to remove and replace parts	65
Replacing a node canister.	65
Replacing an expansion canister	67
Replacing the SFP transceiver	68
Replacing a power supply unit for a control enclosure	70
Replacing a power supply unit for an expansion enclosure	73
Replacing a battery in a power supply unit.	76
Releasing the cable retention bracket	78
Replacing a 3.5" drive assembly or blank carrier	80
Replacing a 2.5" drive assembly or blank carrier	82
Replacing an enclosure end cap	83
Replacing a SAS cable.	83
Replacing a control enclosure chassis	84
Replacing an expansion enclosure chassis	87
Replacing the support rails	89
Storwize V7000 replaceable units	90

Chapter 8. Event reporting 93

Understanding events	93
--------------------------------	----

Viewing the event log	93
Managing the event log	93
Describing the fields in the event log.	94
Event notifications	94
Power-on self-test	95
Understanding the error codes	95
Event codes	95
Error codes	99
Node error code overview	107
Cluster code overview	108
Error code range	108
Node errors	108
Cluster recovery and states	116

Appendix. Accessibility 117

Notices 119

Trademarks	121
Electronic emission notices	121
Federal Communications Commission (FCC) statement.	121
Industry Canada compliance statement.	122
Avis de conformité à la réglementation d'Industrie Canada	122
Australia and New Zealand Class A Statement	122
European Union Electromagnetic Compatibility Directive	122
Germany Electromagnetic compatibility directive	122
Japan VCCI Council Class A statement	123
People's Republic of China Class A Electronic Emission Statement	123
International Electrotechnical Commission (IEC) statement.	124
United Kingdom telecommunications requirements	124
Korean Communications Commission (KCC) Class A Statement	124
Russia Electromagnetic Interference (EMI) Class A Statement	124
Taiwan Class A compliance statement	124
European Contact Information.	124
Taiwan Contact Information	125

Index 127

Figures

1.	12 drives	2	21.	Opening the handle on a canister	66
2.	24 drives	2	22.	Removing the canister from the enclosure	66
3.	LED indicators on a single 3.5" drive	2	23.	Opening the handle on a canister	68
4.	LED indicators on a single 2.5" drive	3	24.	Removing the canister from the enclosure	68
5.	12 drives and two end caps	4	25.	Fibre-channel SFP transceiver	69
6.	Left enclosure end cap	4	26.	Directions for lifting the handle on the power supply unit	71
7.	Rear view of a control enclosure	5	27.	Using the handle to remove a power supply unit	72
8.	Rear view of an expansion enclosure	5	28.	Directions for lifting the handle on the power supply unit	74
9.	LEDs on the power supply units of the control enclosure.	6	29.	Using the handle to remove a power supply unit	75
10.	LEDs on the power supply units of the expansion enclosure	7	30.	Removing the battery from the control enclosure power-supply unit.	77
11.	Fibre-channel ports on the node canisters	8	31.	Unlocking the cable retention bracket	79
12.	LEDs on the fibre-channel ports	8	32.	Pressing the lever to slide the cable retention bracket.	80
13.	USB ports on the node canisters.	10	33.	Unlocking the 3.5" drive	81
14.	Ethernet ports on the node canisters	11	34.	Removing the 3.5" drive	81
15.	SAS ports on the node canisters.	12	35.	Unlocking the 2.5" drive	82
16.	LEDs on the node canisters	12	36.	Removing the 2.5" drive	83
17.	SAS ports and LEDs in rear of expansion enclosure	14	37.	SAS cable	84
18.	LEDs on the expansion canisters	15	38.	Removing a rail assembly from a rack cabinet	90
19.	LEDs on the power supply units of the control enclosure	44			
20.	LEDs on the node canisters	45			

Tables

1. IBM documentation and related websites	xii	12. Expansion canister LEDs	15
2. Drive LEDs	3	13. Access information for your system	17
3. LED descriptions	4	14. Power-supply unit LEDs	44
4. Power supply unit LEDs in the rear of the control enclosure	6	15. Power LED status	45
5. Power supply unit LEDs in the rear of the expansion enclosure	7	16. Cluster status and fault LEDs	46
6. Fibre Channel port LED locations on canister 1	9	17. Status on the control enclosure batteries	46
7. Fibre Channel port LED status descriptions	9	18. Replaceable units	90
8. Ethernet port LEDs	11	19. Description of data fields for the event log	94
9. SAS port LEDs on the node canister	12	20. Notification types	95
10. Node canister LEDs	13	21. Informational event codes	96
11. SAS port LEDs on the expansion canister	14	22. Event IDs and codes	99
		23. Message classification number range	108

Safety and environmental notices

Review the multilingual safety notices for the IBM® Storwize® V7000 system before you install and use the product.

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

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

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

CAUTION:

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

DANGER

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

2. Locate *IBM Systems Safety Notices* with the user publications that were provided with the Storwize V7000 hardware.
3. Find the matching identification number in the *IBM Systems Safety Notices*. Then review the topics concerning the safety notices to ensure that you are in compliance.
4. Optionally, read the multilingual safety instructions on the Storwize V7000 website. Go to Support for Storwize V7000 website at www.ibm.com/storage/support/storwize/v7000, click the current product documentation link, and then click **Multi-language**.

Sound pressure

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

About this guide

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

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

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

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

Who should use this guide

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

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.

Storwize V7000 library and related publications

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

Storwize V7000 Information Center

The IBM Storwize V7000 Information Center contains all of the information that is required to install, configure, and manage the Storwize V7000. The information center is updated between Storwize V7000 product releases to provide the most current documentation. The information center is available at the following website:

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

Storwize V7000 library

Unless otherwise noted, the publications in the Storwize V7000 library are available in Adobe portable document format (PDF) from the following website:

Support for Storwize V7000 website at www.ibm.com/storage/support/storwize/v7000

IBM documentation and related websites

Table 1 lists websites that provide publications and other information about the Storwize V7000 or related products or technologies.

Table 1. IBM documentation and related websites

Website	Address
Support for Storwize V7000 (2076)	Support for Storwize V7000 website at www.ibm.com/storage/support/storwize/v7000
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 Storwize V7000 documentation:

- Go to the feedback page on the website for the Storwize V7000 Information Center at publib.boulder.ibm.com/infocenter/storwize/ic/index.jsp?topic=/com.ibm.storage.StorwizeV7000.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 Storwize V7000 Troubleshooting, Recovery, and Maintenance Guide*
 - Publication form number: GC27-2291-00
 - Page, table, or illustration numbers that you are commenting on
 - A detailed description of any information that should be changed

Chapter 1. Storwize V7000 hardware components

A Storwize V7000 system consists of one or more machine type 2076 rack-mounted enclosures.

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

- The number of drives that an enclosure can hold. Drives are located on the front of the enclosure. An enclosure can hold up to 12 3.5-inch drives or up to 24 2.5-inch drives.
- Whether the model is a control enclosure or an expansion enclosure.

Control enclosures contain the main processing units that control the whole system. They are where external systems, such as host application servers, other storage systems, and management workstations are connected through the Ethernet ports or Fibre Channel ports. Control enclosures can also be connected to expansion enclosures through the serial-attached SCSI (SAS) ports.

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

There are two control enclosure models.

- Machine type and model 2076-112, which can hold up to 12 3.5-inch drives
- Machine type and model 2076-124, which can hold up to 24 2.5-inch drives

There are two expansion enclosure models.

- Machine type and model 2076-212, which can hold up to 12 3.5-inch drives
- Machine type and model 2076-224, which can hold up to 24 2.5-inch drives

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

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

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

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

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

Storwize V7000 components in the front of the enclosure

This topic describes the components in the front of the enclosure.

Drives

An enclosure can hold up to 12 3.5-inch drives or up to 24 2.5-inch drives.

The drives are located in the front of the enclosure. The 12 drives are mounted in four columns with three rows.

The 24 drives are mounted vertically in one row.

Note: The drive slots cannot be empty. A drive assembly or blank carrier must be in each slot.

Figure 1 shows 12 drives, and Figure 2 shows 24 drives.



Figure 1. 12 drives



Figure 2. 24 drives

Drive indicators

The drives have two LED indicators each. They have no controls or connectors.

The LED color is the same for both drives. The LEDs for the 3.5-inch drives are placed vertically above and below each other. The LEDs for the 2.5-inch drives are placed next to each other at the bottom.




Figure 3. LED indicators on a single 3.5" drive



Figure 4. LED indicators on a single 2.5" drive

- 1** Fault LED
- 2** Activity LED

Table 2 shows the status descriptions for the two LEDs.

Table 2. Drive LEDs

Name	Description	Color
Activity	Indicates if the drive is ready or active. <ul style="list-style-type: none"> • If the LED is on, the drive is ready to be used. • If the LED is off, the drive is not ready. • If the LED is flashing, the drive is ready, and there is activity. 	Green
Fault	Indicates a fault or identifies a drive. <ul style="list-style-type: none"> • If the LED is on, a fault exists on the drive. • If the LED is off, no known fault exists on the drive. • If the LED is flashing, the drive is being identified. A fault might or might not exist. 	Amber

Enclosure end cap indicators

This topic describes the indicators on the enclosure end cap.

Figure 5 on page 4 shows where the end caps are located on the front of an enclosure with 12 drives. The end caps are located in the same position for an enclosure with 24 drives.

- **1** Left end cap
- **2** Drives
- **3** Right end cap

Figure 6 on page 4 shows the indicators on the front of the enclosure end cap.

The left enclosure end caps for both enclosures are identical and contain only indicators. The left enclosure end cap contains no controls or connectors. The right enclosure end cap for both enclosures has

no controls, indicators, or connectors.




Figure 5. 12 drives and two end caps




Figure 6. Left enclosure end cap

Table 3. LED descriptions

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

Storize V7000 components in the rear of the enclosure

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

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

Figure 7 shows the rear view of a control enclosure. Figure 8 shows the rear of an expansion enclosure.




Figure 7. Rear view of a control enclosure




Figure 8. Rear view of an expansion enclosure

- 1** Power supply unit 1
- 2** Power supply unit 2
- 3** Canister 1
- 4** Canister 2


Power supply unit and battery for the control enclosure

The control enclosure contains two power supply units, each with an integrated battery.

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

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

Figure 9 on page 6 shows the location of the LEDs **1** in the rear of the power supply unit.



svc00670

Figure 9. LEDs on the power supply units of the control enclosure

Table 4 identifies the LEDs in the rear of the control enclosure.

Table 4. Power supply unit LEDs in the rear of the control enclosure

Name	Color	Symbol
ac power failure	Amber	~
Power supply OK	Green	⏻
Fan failure	Amber	⚙️
dc power failure	Amber	⚡
Battery failure	Amber	🔋
Battery state	Green	🔋+

See “Procedure: Understanding the system status using the LEDs” on page 43 for help in diagnosing a particular failure.

Power supply unit for the expansion enclosure

The expansion enclosure contains two power supply units.

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

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

Figure 10 shows the locations of the LEDs **1** in the rear of the power supply unit.




Figure 10. LEDs on the power supply units of the expansion enclosure

Table 5 identifies the LEDs in the rear of the expansion enclosure.

Table 5. Power supply unit LEDs in the rear of the expansion enclosure

Name	Color	Symbol
ac power failure	Amber	~
Power supply OK	Green	Ⓜ
Fan failure	Amber	⊕
dc power failure	Amber	≡

See “Procedure: Understanding the system status using the LEDs” on page 43 for help in diagnosing a particular failure.

Node canister ports and indicators

The node canister has indicators and ports but no controls.

Fibre Channel ports and indicators

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

Each node canister has four Fibre Channel ports located on the left side of the canister as shown in Figure 11. The ports are in two rows of two ports. The ports are numbered 1 - 4 from left to right and top to bottom.

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




Figure 11. Fibre-channel ports on the node canisters

There are two green LEDs associated with each port: the speed LED and the link activity LED. These LEDs are in the shape of a triangle. The LEDs are located in between the two rows of the ports as shown in Figure 12. Figure 12 shows the LEDs for the Fibre Channel ports on canister 1. Each LED points to the associated port. The first and second LEDs in each set show the speed state, and the third and fourth LEDs show the link state.




Figure 12. LEDs on the fibre-channel ports

Table 6. Fibre Channel port LED locations on canister 1

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

Table 7 provides the status descriptions for the LEDs on the Fibre Channel ports.

Table 7. Fibre Channel port LED status descriptions

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

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) that is allocated to the Storwize V7000 node in which the ports are installed. The WWNN for each node is stored within the enclosure. When you replace a node canister, the WWPNs of the ports do not change.

The WWNN is in the form 50050768020XXXXX, where XXXXX is specific to a Storwize V7000 enclosure.

The WWPNs are in the form 50050768020QXXXXX, where XXXXX is as previously stated and Q is the port number.

USB ports

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

The USB ports are numbered 1 on the left and 2 on the right as shown in Figure 13 on page 10. One port is used during installation.

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




Figure 13. USB ports on the node canisters


The USB ports have no indicators.

Ethernet ports and indicators

Two Ethernet ports are located side by side on the node canister.

The Ethernet ports are numbered 1 on the left and 2 on the right as shown in Figure 14 on page 11. Port 1 must be connected; the use of port 2 is optional. Two LEDs are associated with each port.

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



svc00691

Figure 14. Ethernet ports on the node canisters

Table 8 provides a description of the two LEDs.

Table 8. Ethernet port LEDs

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

Node canister SAS ports and indicators

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

The SAS ports are numbered 1 on the left and 2 on the right as shown in Figure 15 on page 12. Port 1 is used if you add one expansion enclosure. Port 2 is used if you add a second expansion enclosure. Each port provides four data channels.

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




Figure 15. SAS ports on the node canisters.

SAS ports must be connected to Storwize V7000 enclosures only. See “Problem: SAS cabling not valid” on page 39 for help in attaching the SAS cables.

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

Table 9. SAS port LEDs on the node canister

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

Node canister LEDs

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

The three LEDs are located in a horizontal row near the upper right of the canister **1**. Figure 16 shows the rear view of the node canister LEDs.







Figure 16. LEDs on the node canisters

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

Table 10. Node canister LEDs

Name	Description	Color	Symbol
Cluster status	<p>Indicates the status of the node.</p> <ul style="list-style-type: none"> The on status indicates that the node is active, that is, it is an active member of a cluster. When the node is active, do not remove it. If the cluster status LED is on and the fault LED is off, the node canister is an active member of a cluster. If the cluster status LED is on and the fault LED is on, there is a problem establishing a cluster. The off state indicates there is no power to the canister or the canister is in standby mode. These conditions can cause the off state: <ul style="list-style-type: none"> The main processor is off and only the service processor is active. A power-on self-test (POST) is running on the canister. The operating system is loading. The flashing status indicates that the node is in candidate state or service state. It is not able to perform I/O in a cluster. When the node is in either of these states, it can be removed. Do not remove the canister unless directed by a service procedure. 	Green	
Fault	<p>Indicates if a fault is present and identifies which canister.</p> <ul style="list-style-type: none"> The on status indicates that the node is in service state or an error exists that might be stopping the software from starting. Do not assume that this status indicates a hardware error. Further investigation is required before replacing the node canister. The off status indicates that the node is a candidate or is active. This status does not mean that there is not a hardware error on the node. Any error that was detected is not severe enough to stop the node from participating in a cluster. The flashing status indicates that the canister is being identified. This status might or might not be a fault. 	Amber	
Power	<p>Indicates if power is available and the boot status of the canister.</p> <ul style="list-style-type: none"> The on status indicates that the canister is powered on and that the main processor or processors are running. The off status indicates that no power is available. The slow flashing (1 Hz) status indicates that power is available and that the canister is in standby mode. The main processor or processors are off and only the service processor is active. The fast flashing (2 Hz) indicates that the canister is running the power-on self-test (POST). 	Green	

Expansion canister ports and indicators

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

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

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

Expansion canister SAS ports and indicators

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

The SAS ports are numbered 1 on the left and 2 on the right as shown in Figure 17. Use of port 1 is required. Use of port 2 is optional. Each port connects four data channels.

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




Figure 17. SAS ports and LEDs in rear of expansion enclosure

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

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

Table 11. SAS port LEDs on the expansion canister

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

Expansion canister LEDs

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

The two LEDs are located in a vertical row on the left side of the canister. Figure 18 on page 15 shows the LEDs (**1**) in the rear of the expansion canister.




Figure 18. LEDs on the expansion canisters

Table 12. Expansion canister LEDs

Name	Description	Color	Symbol
Status	<p>Indicates if the canister is active.</p> <ul style="list-style-type: none"> • If the LED is on, the canister is active. • If the LED is off, the canister is not active. • If the LED is flashing, there is a vital product data (VPD) error. 	Green	⚡
Fault	<p>Indicates if a fault is present and identifies the canister.</p> <ul style="list-style-type: none"> • If the LED is on, a fault exists. • If the LED is off, no fault exists. • If the LED is flashing, the canister is being identified. This status might or might not be a fault. 	Amber	!

Chapter 2. Best practices for troubleshooting

Troubleshooting is made easier by taking advantage of certain configuration options and ensuring that you have recorded vital information that is required to access your system.

Record access information

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

Record the following information and ensure that authorized people know how to access the information:

- The management IP addresses. This address connects to the system using the management GUI or starts a session that runs the command-line interface (CLI) commands. The system has two Ethernet ports. Each port can have either an IPv4 address or an IPv6 address or both. Record these addresses and any limitations regarding where they can be accessed from within your Ethernet network.
- The service IP address of the node canisters. The service IP address connects to a node canister in the control enclosure. Access to the address is sometimes required if the canister has a fault that stops it from becoming an active member of the system. Each of the two node canisters can have a service IP address that is specified for Ethernet port 1. Each address can have either an IPv4 address or an IPv6 address or both. Ensure that the address specified for each node canister is different.
- The Storwize V7000 system password for user superuser. The password is required to access the system through the service IP address. The authentication of superuser is always local; therefore, the user ID can be used when a remote authentication server that is used for other users is not available.

Table 13. Access information for your system

Item	Value	Notes
Management IP address: Ethernet port 1		
Management IP address: Ethernet port 2		
Service IP address: node canister 1		
Service IP address: node canister 2		
Superuser password		

Follow power management procedures

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

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

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

Remove a node canister only when directed to do so by a service action. Physically removing an active node canister means that it is unable to write any configuration data or volume data that it has cached to

its internal disk and the data is lost. If both node canisters in a control enclosure are removed in quick succession, you will have to run recovery actions, which might include restoring your volume data from a backup.

Set up event notifications

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

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

- Email. An event notification can be sent to one or more email addresses. This mechanism notifies individuals of problems. Individuals can receive notifications wherever they have email access which includes mobile devices.
- Simple Network Management Protocol (SNMP). An SNMP trap report can be sent to a data-center management system, such as IBM Systems Director, that consolidates SNMP reports from multiple systems. Using this mechanism, you can monitor your data center from a single workstation.
- Syslog. A syslog report can be sent to a data-center management system that consolidates syslog reports from multiple systems. Using this mechanism, you can monitor your data center from a single workstation.

If your system is within warranty, or you have a hardware maintenance agreement, configure your system to send email reports to IBM if an issue that requires hardware replacement is detected. This mechanism is called a Call Home email. When this email is received, IBM automatically opens a problem report, and if appropriate, contacts you to verify if replacement parts are required.

If you set up Call Home to IBM, ensure that the contact details that you configure are correct and kept up to date as personnel change.

Set up inventory reporting

Inventory reporting is an extension to the Call Home email.

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

Back up your data

Back up your system configuration data and volume data.

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

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

Manage your spare and failed drives

Your RAID arrays that are created from drives in the Storwize V7000 system consist of drives that are active members and drives that are spares.

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

Resolve alerts in a timely manner

Storwize V7000 reports an alert when there is an issue or a potential issue that requires user attention. The management GUI helps resolve these problems through the Recommended Actions panel.

Perform the recommended actions as quickly as possible after the problem is reported. Storwize V7000 is designed to be resilient to most single hardware failures. However, if you operate for any period of time with a hardware failure, the possibility increases that a second hardware failure can result in some volume data that is unavailable.

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

Keep your software up to date

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

Check the IBM support website to see if new code releases are available:

Support for Storwize V7000 website at www.ibm.com/storage/support/storwize/v7000

The release notes provide information about new function in a release plus any issues that have been resolved. Update your code regularly if the release notes indicate an issue that you might be exposed to.

Keep your records up to date

Record the location information for your enclosures

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

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

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

Subscribe to support notifications

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

Subscribe to support notifications by visiting the Storwize V7000 support page on the IBM website. By subscribing, you are informed of new and updated support site information, such as publications, hints and tips, technical notes, product flashes (alerts), and downloads.

Know your IBM warranty and maintenance agreement details

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

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

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

Chapter 3. Understanding the Storwize V7000 battery operation for the control enclosure

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

If the power fails, the cache and state data is written to a local solid-state drive (SSD) that is held within the canister. The batteries within the control enclosure provide the power to write the cache and state data to a local drive.

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

There are two power supply units in the control enclosure. Each one contains an integrated battery. Both power supply units and batteries provide power to both control canisters. Each battery has a sufficient charge to power both node canisters for the duration of saving critical data to the local drive. In a fully redundant system with two batteries and two canisters, there is enough charge in the batteries to support saving critical data from both canisters to a local drive twice. In a system with a failed battery, there is enough charge in the remaining battery to support saving critical data from both canisters to a local drive once.

If the ac power to a control enclosure is lost, the canisters do not start saving critical data to a local drive until approximately 10 seconds after the loss of ac power is first detected. If the power is restored within this period, the system continues to operate. This loss in power is called a *brown out*. As soon as the saving of the critical data starts, the system stops handling I/O requests from the host applications, and Metro Mirror and Global Mirror relationships go offline. The system powers off when the saving of the critical data completes.

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

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

When the ac power is restored following a power outage that causes both canisters to save their critical data, the system restarts only when the batteries have sufficient charge to power both canisters for the duration of saving the critical data again. In a fully redundant system with two batteries, this condition means that after one ac power outage and a saving of critical data, the system can restart as soon as the power is restored. If a second ac power outage occurs before the batteries have completed charging, then the system starts in service state and does not permit I/O operations to be restarted until the batteries are half charged. The recharging takes approximately 30 minutes.

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

A battery is considered failed for the following conditions:

- When the system can communicate with it and it reports an error.
- When the system is unable to communicate with the battery. Failed communication exists because the power supply, which contains the battery, has been removed or because the power supply has failed in a manner that makes communication with the battery impossible.

There are conditions other than loss of ac power that can cause critical data to be saved and the nodes to go into service state and not permit I/O operations. The node canisters save critical data if they detect there is no longer sufficient battery charge to support a saving of critical data. This situation happens when, for example, both batteries have two-thirds of a charge. The total charge that is available in the enclosure is sufficient to support a saving of critical data once; therefore, both canisters are in active state and I/O operations are permitted. If one battery fails though, the remaining battery has only two-thirds of a charge, and the total charge that is available in the enclosure is now insufficient to perform a saving of the critical data if the ac power fails. Data protection cannot be guaranteed in this case. The nodes save the critical data using the ac power and enter service state. The nodes do not handle I/O operations until the remaining battery has sufficient charge to support the saving of the critical data. When the battery has sufficient charge, the system automatically restarts.

Important: Although Storwize V7000 is resilient to power failures and brown outs, always install Storwize V7000 in an environment where there is reliable and consistent ac power that meets the Storwize V7000 requirements. Consider uninterruptible power supply units to avoid extended interruptions to data access.

Maintenance discharge cycles

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

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

Maintenance discharges are scheduled for the following situations:

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

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

Under the following conditions, a battery is not considered when calculating whether there is sufficient charge to protect the system. This condition persists until a maintenance discharge cycle is completed.

- A battery is performing a maintenance discharge.
- A battery has provided protection for saving critical data at least four times without any intervening maintenance discharge.
- A battery has provided protection for at least 20 brown outs, which lasted up to 10 seconds each.

- A battery must restart a maintenance discharge because the previous maintenance cycle was disrupted by an ac power outage.

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

If one of the batteries in a system fails and is not replaced, it prevents the other battery from performing a maintenance discharge. Not only does this condition reduce the lifetime of the remaining battery, but it also prevents a maintenance discharge cycle from occurring after the battery has provided protection for at least 2 critical saves or 10 brown outs. Preventing this maintenance cycle from occurring increases the risk that the system accumulates a sufficient number of power outages to cause the remaining battery to be discounted when calculating whether there is sufficient charge to protect the system. This condition results in the system entering service state while the one remaining battery performs a maintenance discharge. I/O operations are not permitted during this process. This activity takes approximately 10 hours.

Chapter 4. Storwize V7000 user interfaces for servicing your system

Storwize V7000 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 initialization tool to do the initial setup of your system. Use the management GUI to monitor and maintain the configuration of storage that is associated with your 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 Storwize V7000 website at www.ibm.com/storage/support/storwize/v7000

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.

If you encounter problems logging on the management GUI or connecting to the management GUI, see “Problem: Unable to log on to the management GUI” on page 37 or “Problem: Unable to connect to the management GUI” on page 36.

Service assistant interface

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

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

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 or if both node canisters in an enclosure are in service state.

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

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.
- Configure an enclosure chassis after replacement.
- 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.

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

Accessing the service assistant

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

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

Support for Storwize V7000 website at www.ibm.com/storage/support/storwize/v7000

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. If you are unable to connect to the service assistant, see “Problem: Cannot connect to the service assistant” on page 38.

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

If you are accessing a new node canister, the default password is `passwd`. If the node canister is a member of a cluster or has been a member of a cluster, use the password for the cluster user superuser.

If you do not know the current superuser password, reset the password. Go to “Procedure: Resetting superuser password” on page 40.

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

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

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 Storwize V7000 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 Storwize V7000 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 Storwize V7000 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 Storwize V7000 Information Center to initialize and use a service CLI session.

USB key and Initialization tool interface

Use a USB key to retrieve status from a system and to manage a system.

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

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

When to use the USB key

The USB key is normally used to create a cluster after installing a new system.

Using the USB key is required in the following situations:

- When you cannot connect to a node canister using the service assistant and you want to see the status of the node.
- When you do not know, or cannot use, the service IP address for the node and must set the address.
- When you have forgotten the superuser password and must reset the password.

Using a USB key

Use any USB key that is formatted with FAT32, EXT2, or EXT3 file systems on its first partition.

When a USB key is plugged into a node, the node software searches for a text file named `satask.txt` in the root directory. If the software finds the file, it attempts to run a command that is specified in the file. When the command completes, a file called `satask_result.html` is written to the root directory of the USB key. If this file does not exist, it is created. If it exists, the data is inserted at the start of the file. The file contains the details and results of the command that was run and the status and the configuration information from the node. The status and configuration information matches the detail that is shown on the service assistant home page panels.

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

The fault LED on the node canister flashes when the USB service action is being performed. When the fault LED stops flashing, it is safe to remove the USB key.

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

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

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

Using the initialization tool

The initialization tool is a graphical user interface (GUI) application. You must have Microsoft Windows XP Professional or higher to run the application.

The initialization tool is available on the USB key that is shipped with the control enclosures. The name of the application file is `InitTool.exe`. If you cannot locate the USB key, you can download the application from the support website:

Support for Storwize V7000 website at www.ibm.com/storage/support/storwize/v7000

If you download the initialization tool, you must copy the file onto the USB key that you are going to use.

To start the initialization tool, insert the USB key that contains the program into a USB slot on a suitable personal computer. Run the `InitTool.exe` program from the USB drive.

The initialization tool is used to create the `satask.txt` file on a USB key. After the `satask.txt` file is created, follow the instructions in "Using a USB key" to run the commands on the node.

The initialization tool prompts you for the task that you want to perform and for the parameters that are relevant to that task. It prompts you when to put it in the node canister. When the commands have run, return the USB key to your personal computer and start the tool again to see the results.

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

- Create a new cluster.
- Reset the superuser password.
- Set or reset the service IP address.

For any other tasks, you must create the `satask.txt` file using a text editor.

satask.txt commands

This topic identifies the commands that can be run from a USB key.

If you are creating the `satask.txt` command file using a text editor, the file must contain a single command on a single line in the file. The commands that you use are the same as the service CLI commands except where noted. Not all service CLI commands can be run from the USB key. The `satask.txt` commands always run on the node that the USB key is plugged into.

Reset service IP address and superuser password command

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

Syntax

```
▶▶ satask -- chserviceip -- --serviceip--ipv4-- [--gw--ipv4] [--mask--ipv4] [--resetpassword]
▶▶ satask -- chserviceip -- --serviceip_6--ipv6-- [--gw_6--ipv6] [--prefix_6--int] [--resetpassword]
▶▶ satask -- chserviceip -- --default-- [--resetpassword]
```

Parameters

-serviceip

(Required) The IPv4 address for the service assistant.

-gw

(Optional) The IPv4 gateway for the service assistant.

-mask

(Optional) The IPv4 subnet for the service assistant.

-serviceip_6

(Required) The IPv6 address for the service assistant.

-gw_6

(Optional) The IPv6 gateway for the service assistant.

-prefix_6

(Optional) Resets to the default IPv4 address.

-default

(Required) The IPv6 prefix for the service assistant.

-resetpassword

(Optional) Sets the service assistant password to default.

Description

This command resets the service assistant IP address to the default value `passwd0rd` or to the address that is specified. If the node is active in a cluster, the cluster superuser password is reset; otherwise, the superuser password is reset on the node.

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

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

Reset service assistant password command

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

Syntax

```
▶▶ satask — resetpassword —————▶▶
```

Parameters

None.

Description

This command resets the service assistant password to the default value `passwd0rd`. If the node is active in a cluster, the cluster superuser password is reset; otherwise, the superuser password is reset on the node.

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

This command calls the `satask resetpassword` command.

Snap command

Use this command to collect diagnostic information from the node and to write the output to a USB key.

Syntax

```
▶▶ satask — snap — --options————▶▶
```

Parameters

-options

(Optional) Specifies which diagnostic information to collect.

Description

This command moves a snap file to a USB key.

This command calls the `satask snap` command.

Apply software command

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

Syntax

```
▶▶ satask — installsoftware — — -file filename — [ -ignore ] —————▶▶
```

Parameters

-file

(Required) The file name of software installation package.

-ignore

(Optional) Overrides prerequisite checking and forces installation of the software.

Description

This command copies the file from the USB key to the upgrade directory on the node.

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

Create Cluster command

Use this command to create a cluster.

Syntax

```
▶▶ satask — mkcluster — — -clusterip ipv4 — [ -gw ipv4 ] [ -mask ipv4 ] [ -name cluster_name ] —————▶▶
```

```
▶▶ satask — mkcluster — — -clusterip_6 ipv6 — [ -gw_6 ipv6 ] [ -prefix_6 int ] [ -name cluster_name ] —————▶▶
```

Parameters

-clusterip

(Required) The IPv4 address for cluster Ethernet port 1.

-gw

(Required) The IPv4 gateway for cluster Ethernet port 1.

-mask

(Required) The IPv4 subnet for cluster Ethernet port 1.

-clusterip_6

(Required) The IPv6 address for cluster Ethernet port 1.

gw_6

(Required) The IPv6 gateway for cluster Ethernet port 1.

prefix_6

(Required) The IPv6 prefix for cluster Ethernet port 1.

name

(Optional) The name of the new cluster.

Description

This command creates a cluster

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

Query status command

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

Syntax

```
▶▶— sainfo — getstatus — —————▶▶
```

Parameters

None.

Description

This command writes the output from each node to the USB key.

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

Chapter 5. Resolving a problem

This topic describes the procedures that you follow to resolve fault conditions that exist on your system. This topic assumes that you have a basic understanding of the Storwize V7000 system concepts.

The following procedures are often used to find and resolve problems:

- Procedures that involve data collection and system configuration
- Procedures that are used for hardware replacement.

Always use the recommended actions of the management GUI as the starting point to diagnose and resolve a problem. The topics that follow describe the type of problem that you might experience that are not resolved by using the management GUI. In those situations, review the symptoms and follow the actions that are provided here.

Unless you are unable to detect a newly installed enclosure, problems on expansion enclosures are resolved using the recommended actions in the management GUI. The "Start here" topic gives the starting point for any service action. The situations covered in this section are the cases where you cannot start the management GUI or the node canisters in the control enclosure are unable to run the system software.

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

Start here: Use the management GUI recommended actions

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

You can connect to and manage a Storwize V7000 system as soon as you have created a cluster. If you cannot create a cluster, see the problem that contains information about what to do if you cannot create a cluster. Go to "Problem: Cannot create a cluster" on page 37.

To run the management GUI, start a supported web browser and point it to the management IP address of your system. Up to four addresses can be configured for your use. There are two addresses for IPv4 access, and two addresses for IPv6 access. If you do not know the system management IP address, go to "Problem: Management IP address unknown" on page 36. After the connection is successful, you see a login panel. If you are unable to access the login panel, go to "Problem: Unable to connect to the management GUI" on page 36.

Log on using your user name and password. If you are unable to log on, go to "Problem: Unable to log on to the management GUI" on page 37.

When you have logged on, select **Troubleshooting** and then **Recommended Actions**. Any unfixed alerts are listed, and you see the alert that requires fixing first.

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

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

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

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

After all the alerts are fixed, go to “Procedure: Checking the status of your system” on page 42.

Problem: Management IP address unknown

This topic helps you if you are not able to run the management GUI because you do not know the IP address. This address is also known as the cluster IP address.

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

If you do not know the cluster address, it is part of the data that is shown in the service assistant home panel or the data that is returned by the USB key. If you know the service address of a node canister, go to “Procedure: Getting node canister and cluster information using the service assistant” on page 42; otherwise, go to “Procedure: Getting node canister and cluster information using a USB key” on page 42.

Problem: Unable to connect to the management GUI

This topic helps you if you are unable to connect to the management GUI from your web browser. You might get a Page not found or a similar error from the browser.

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

- You cannot connect if the system is not operational with at least one node online. If you know the service address of a node canister, go to “Procedure: Getting node canister and cluster information using the service assistant” on page 42; otherwise, go to “Procedure: Getting node canister and cluster information using a USB key” on page 42 and obtain the state of each of the node canisters from the data that is returned. If there is not a node canister with a state of active, resolve the reason why it is not in active state. If the state of both node canisters is candidate, then there is not a cluster to connect to. If the node state is service, go to “Procedure: Fixing node errors” on page 48.
- Ensure that you are using the correct cluster management IP address. If you know the service address of a node canister, go to “Procedure: Getting node canister and cluster information using the service assistant” on page 42; otherwise, go to “Procedure: Getting node canister and cluster information using a USB key” on page 42 and obtain the cluster IP address from the data that is returned.
- Ensure that you are using the correct cluster management IP address. If you know the service address of a node canister, use the service assistant to gather information about the node canister and cluster. You can also use the USB key to obtain the cluster IP address from the data that is returned.
- Ensure that all node canisters have an Ethernet cable that is connected to port 1 and that the port is working. To understand the port status, go to “Procedure: Finding the status of the Ethernet connections” on page 47.
- Ping the cluster address to see if the Ethernet network permits the connection. If the ping fails, check the Ethernet network configuration to see if there is a routing or a firewall issue. Ensure that the Ethernet network configuration is compatible with the gateway and subnet or prefix settings. Ensure

that you have not used the Ethernet address of another device as the cluster address. If necessary, modify your network settings to establish a connection.

- If you have just created the cluster using the USB key or service assistant, and entered a cluster IP address that is not accessible, you can safely delete the cluster and perform the cluster creation again using a suitable address. It is safe to delete the cluster in this situation because you have not configured any volumes or loaded any data onto the system.

Do not delete the cluster if you have created any volumes on your system because any data on those volumes will be lost. In this case, you must gain access to the cluster IP address. If you just created the cluster and cannot access the cluster IP address, go to “Procedure: Deleting a cluster completely” on page 48 to reset the enclosure. After that procedure, repeat the procedures to create a new cluster using an accessible IP address.

Problem: Unable to log on to the management GUI

This topic assists you when you can see the management GUI login screen but cannot log on.

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

- If you are not logging on as superuser, contact your system administrator who can verify your user name and reset your account password.
- If the user name that you are using is authenticated through a remote authentication server, verify that the server is available. If the authentication server is unavailable, you can log on as user name superuser. This user is always authenticated locally.
- If you do not know the password for superuser, go to “Procedure: Resetting superuser password” on page 40.

Problem: Cannot create a cluster

This topic helps if your attempt to create a cluster has failed.

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

- USB key
- Service assistant
- Service command line

The create cluster function protects the system from loss of volume data. If you create a cluster on a control enclosure that was previously used, you lose all of the volumes that you previously had. To protect against accidental creation of a cluster on a system that has required volume data, you have to follow the procedure to delete a cluster before you create a new one. For more information, go to “Procedure: Deleting a cluster completely” on page 48.

Use data that is returned by “Procedure: Getting node canister and cluster information using the service assistant” on page 42 or “Procedure: Getting node canister and cluster information using a USB key” on page 42 to check the following items:

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

If the node canister that you are attempting to create the cluster on is in service state, fix the reported node errors. For more information, go to “Procedure: Fixing node errors” on page 48. After the node error is corrected, attempt to create a cluster again.

If the reason the create cluster function failed is because this node or the other node in the cluster has cluster data or the enclosure cluster ID is not 0, fix the existing cluster versus re-creating a new cluster. If you want to create a cluster and do not want to use any data from the volumes used in the previous cluster, go to “Procedure: Deleting a cluster completely” on page 48, and then run the create cluster function again.

Problem: Unknown service address of a node canister

This topic describes the methods that you can use to determine the service address of a node canister.

- If you can access the management GUI, start the service assistant on the cluster configuration node by pointing your browser to the address: *cluster address*/service. For example, if your cluster address is 11.22.33.44, point your browser to 11.22.33.44/service.

The service assistant home page lists the node canister within the cluster that it can communicate with. If the service address of the node canister that you are looking for is listed in the **Change Node** window, make the node the current node. Its service address is listed under the **Access** tab of the node details. Use the service address to connect to the node or continue to manage the node using this session. If the service address of the node that you want is not listed, go to “Procedure: Getting node canister and cluster information using a USB key” on page 42 to get the service address.

- If you know the service address of any node canister in the system, follow a similar procedure to the one described previously. Rather than using *cluster address*/service to start the service assistant, use the service address that you know.
- Use a USB key to find the service address of a node. For more information, go to “Procedure: Getting node canister and cluster information using a USB key” on page 42.

Problem: Cannot connect to the service assistant

This topic provides assistance if you are unable to display the service assistant on your browser.

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

- Check that you have entered the “/service” path after the service IP address. Point your web browser to *<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.
- Check that you are using the correct service address for the node canister. To find the IPv4 and IPv6 addresses that are configured on the node, go to “Problem: Unknown service address of a node canister.” Try accessing the service assistant through these addresses. Verify that the IP address, subnet, and gateway are specified correctly for IPv4 addresses. Verify that the IP address, prefix, and gateway are specified for the IPv6 addresses. If any of the values are incorrect, see “Procedure: Changing the service IP address of a node” on page 48.
- You cannot connect to the service assistant if the node canister is not able to start the Storwize V7000 code. To verify that the LEDs indicate that the Storwize V7000 code is active, see “Procedure: Understanding the system status using the LEDs” on page 43.
- The service assistant is configured on Ethernet port 1 of a node. Verify that an Ethernet cable is connected to this port and to an active port on your Ethernet network. See “Procedure: Finding the status of the Ethernet connections” on page 47 for details.
- Ping the cluster address to see if the Ethernet network permits the connection. If the ping fails, check the Ethernet network configuration to see if there is a routing or a firewall issue. Check that the Ethernet network configuration is compatible with the gateway and subnet or prefix settings. Check that you have not used an address that is used by another device on your Ethernet network. If necessary, change the network configuration or see “Procedure: Changing the service IP address of a node” on page 48 to change the service IP address of a node.
- A default service address is initially assigned to each node canister. The service IP address 192.168.70.121 is preconfigured on Ethernet port 1 of the upper canister, canister 1. The service IP address 192.168.70.122 is preconfigured on Ethernet port 2 of the lower canister, canister 2.

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

- These addresses are the same as the addresses that are used by other devices on the network.
- These addresses cannot be accessed on your network.
- There are other reasons why they are not suitable for use on your network.

If the previous conditions apply, see “Procedure: Changing the service IP address of a node” on page 48 to change the service IP address to one that works in your environment.

If you are unable to change the service address, for example, because you cannot use a USB key in the environment, see “Procedure: Accessing a canister using a directly attached Ethernet cable” on page 51.

Problem: Management GUI or service assistant does not display correctly

This topic provides assistance if the management GUI or the service assistant does not display correctly.

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

Support for Storwize V7000 website at www.ibm.com/storage/support/storwize/v7000

Switch to using a supported browser. If the problem continues, contact IBM Support.

Problem: Node canister location error

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

To find out how to resolve the node error, go to “Procedure: Fixing node errors” on page 48.

Be aware of the following items:

- Each control enclosure must have two node canisters installed.
- Node canisters and expansion canisters are not interchangeable. A node canister cannot operate in an expansion enclosure. An expansion canister cannot operate in a control enclosure.
- After the node has been used in a cluster, the node canister has saved information that can identify whether the canister has been moved to a different enclosure or a different slot in the same enclosure from where it was previously used. Moving a node canister might compromise its access to storage or access to volumes by a host application. Do not move the canister from its original location unless directed to do so by a service action.

Problem: SAS cabling not valid

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

Check the following items:

- No more than five expansion enclosures can be chained to port 1 (below the control enclosure). The connecting sequence from port 1 of the node canister is called chain 1.
- No more than four expansion enclosures can be chained to port 2 (above the control enclosure). The connecting sequence from port 2 of the node canister is called chain 2.
- Do not connect a SAS cable between a port on an upper canister and a port on a lower canister.
- In any enclosure, the same ports must be used on both canisters.
- No SAS cable can be connected between ports in the same enclosure.

- For any enclosure, the cables that are connected to SAS port 1 on each canister must attach to the same enclosure. Similarly, for any enclosure, the cables that are connected to SAS port 2 on each canister must attach to the same enclosure. Cable attachments for SAS port 1 and cable attachments for SAS port 2 do not go to the same enclosure.
- For cables connected between expansion enclosures, one end is connected to port 1 while the other end is connected to port 2.
- For cables that are connected between a control enclosure and expansion enclosures, port 1 must be used on the expansion enclosures.
- The last enclosure in a chain must not have cables in port 2 of canister 1 and port 2 of canister 2.
- Ensure that each SAS cable is fully inserted.

Problem: New expansion enclosure not detected

This topic helps you resolve why a newly installed expansion enclosure was not detected by the system.

When installing a new expansion enclosure, follow the management GUI Add Enclosure wizard, which is available from the **Manage Devices Actions** menu.

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

- Verify the status of the LEDs at the back of the expansion enclosure. At least one power supply unit must be on with no faults shown. At least one canister must be active, with no fault LED on, and all the SAS port 1 LEDs must be on. For details about the LED status, see “Procedure: Understanding the system status using the LEDs” on page 43.
- Verify that the SAS cabling to the expansion enclosure has been correctly installed. To review the requirements, see “Problem: SAS cabling not valid” on page 39.

Problem: Mirrored volume copies no longer identical

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

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

You have the following choices:

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

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

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

Procedure: Resetting superuser password

You can reset the superuser password to the default password of `passw0rd` by using a USB key command action.

You can use this procedure to reset the superuser password if you have forgotten the password. This command runs differently depending on whether you run it on a node that is active in a cluster.

- If the node is active in a cluster, the password for superuser is changed on the cluster.

- If the node is not in active state, the superuser password for the node is changed. If the node joins a cluster later, the superuser password is reset to that of the cluster.

Note: If a node is not in active state, the superuser password is still required to log on to the service assistant.

It is possible to configure your system so that resetting the cluster superuser password with the USB key command action is not permitted. If your system is configured this way, there is no workaround. Contact the person who knows the password.

To use a USB key to reset the superuser password, see “USB key and Initialization tool interface” on page 28.

Procedure: Identifying which enclosure or canister to service

Use this procedure to identify which enclosure or canister must be serviced.

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

- The model type that is shown on the labels. Model types 112 and 124 are control enclosures. Model types 212 and 224 are expansion enclosures.
- The model description that is shown on the left end cap. The description shows either Control or Expansion.
- The number of ports at the rear of the enclosure. Control enclosures have Ethernet ports, Fibre Channel ports, and USB ports on the canisters. Expansion enclosures do not have any of these ports.
- The number of LEDs on the power supply units. The power supply units in the control enclosure have six LEDs. The power supply units in the expansion enclosure have four LEDs.

Identify the enclosure. An enclosure is identified by its ID and serial number.

- The ID is shown on the LCD panel on the front left of the enclosure. The serial number is also found on the front left end cap of the enclosure and is repeated on the rear left flange of the enclosure. The enclosure ID is unique within a Storwize V7000 system. However, if you have more than one Storwize V7000 system, the same ID can be used within more than one system. The serial number is always unique.

Note: Use the **Manage Device** options from the management GUI to change the ID of an enclosure. Use this option to set a unique ID on all your enclosures.

- Within an enclosure, a canister is identified by its slot location. Slot 1 is the upper canister. Slot 2 is the lower canister. A canister is uniquely identified by the enclosure that it is in and the slot location. The ID can be shown as E-C or E|C where *E* is the enclosure ID and *C* is the canister location. On the service assistant, the ID is known as the *Panel*.

Note: When a node canister is added to a cluster as a node, it is given a node name and a node ID. The default node name is *nodeN*, where *N* is an integer number. This number does not represent the slot location of the node. Similarly, the node ID does not indicate the slot location. The **Manage Device** → **Canister** panel from the management GUI shows both the node name and the canister location. The service assistant home page also shows both the node name and the canister location. If you have only the node name, use these panels to determine the node canister location.

- Use the service assistant to identify a node canister by turning on the identify LED of the containing enclosure. This option is at the upper left of the service assistant page. It is a good practice to identify a node in this way before performing any service action. Performing a service action on the wrong canister can lead to loss of access to data or loss of data.

Procedure: Checking the status of your system

Use this procedure to verify the status of objects in your system. If the status of the object is not online, view the alerts and run the recommended fix procedures.

Volumes normally show offline because another object is offline. A volume is offline if one of the MDisks that make up the storage pool that it is in is offline. You do not see an alert that relates to the volume. Instead, the alert relates to the MDisk. Performing the fix procedures for the MDisk enables the volume to go online.

An overview of the status is displayed under **Connectivity** in the lower-left corner of the management GUI window.

Use the following management GUI functions to find a more detailed status:

- **Home** → **System Status**
- **Home** → **Manage Device**
- **Home** → **Physical Storage MDisks and Pools**
- **Volumes** → **All Volumes**
- **Troubleshooting** → **Recommended Actions**

Procedure: Getting node canister and cluster information using the service assistant

This procedure explains how to view information about the node canister and cluster using the service assistant.

To obtain the information, connect to and log on to the service assistant using the starting service assistant procedure. For more information, go to “Accessing the service assistant” on page 27.

1. Log on to the service assistant.
2. View the information about the node canister that you connected to or the other node canister in the same enclosure.

Note: If the node that you want to see information about is not the current node, change it to the current node from the home page.

3. Examine the data shown for the current node.

The home page shows a table of node errors that exist on the node and a table of node details for the current node. The node errors are shown in priority order.

The node details are divided into several sections. Each section has a tab. Examine the data that is reported in each tab for the information that you want. The **Node** tab shows general information about the node that includes the node state and whether it is a configuration node. The **Hardware** tab shows information about the hardware. The **Access** tab shows the cluster IP address and the service address for this node. The **Location** tab identifies the enclosure in which the node canister is located.

Procedure: Getting node canister and cluster information using a USB key

This procedure explains how to view information about the node canister and cluster using a USB key.

Use any USB key with a FAT32 file system, a EXT2 file system, or a EXT3 file system on its first partition.

1. Ensure that the USB key does not contain a file named `satask.txt` in the root directory.

If `satask.txt` does exist in the directory, the node attempts to run the command that is specified in the file. The information that is returned is appended to the `satask_result.html` file. Delete this file if you no longer want the previous output.

2. Insert the USB key in one of the USB ports of the node canister from which you want to collect data.
3. The node canister fault LED flashes. It continues to flash while the information is collected and written to the USB key.
4. Wait until the LED stops flashing before removing the USB key.
Because the LED is a fault indicator, it might remain permanently on or off.
5. View the results in a web browser.

Procedure: Understanding the system status using the LEDs

This procedure helps you determine the system status using the LED indicators on the system.

The LEDs provide a general idea of the system status. You can obtain more detail from the management GUI and the service assistant. Examine the LEDs when you are not able to access the management GUI or the service assistant, or when the system is not showing any information about a device. For information about the LEDs, go to “Power supply unit and battery for the control enclosure” on page 5, “Power supply unit for the expansion enclosure” on page 6, “Fibre Channel ports and indicators” on page 7, “Ethernet ports and indicators” on page 10, “Node canister SAS ports and indicators” on page 11, “Node canister LEDs” on page 12, “Expansion canister SAS ports and indicators” on page 14, and “Expansion canister LEDs” on page 14.

The procedure shows the status for the enclosure chassis, power supply units and batteries, and canisters. It does not show the status for the drives.

The first step is to determine the state of the control enclosure, which includes its power supply units, batteries, and node canisters. Your control enclosure is operational if you can manage the system using the management GUI. You might also want to view the status of the individual power supply units, batteries, or node canisters.

Find the control enclosure for the system that you are troubleshooting. There is one control enclosure in a system. If you are unsure which one is the control enclosure, go to “Procedure: Identifying which enclosure or canister to service” on page 41.

1. Use the state of the ac power failure, power supply OK, fan failure, and dc power failure LEDs on each power supply unit in the enclosure to determine if there is power to the system, or if there are power problems. Figure 19 on page 44 shows the LEDs on the power supply unit.




Figure 19. LEDs on the power supply units of the control enclosure

Table 14. Power-supply unit LEDs

Power supply OK	ac failure	Fan failure	dc failure	Status	Action
On	On	On	On	Communication failure between the power supply unit and the enclosure chassis	Replace the power supply unit. If failure is still present, replace the enclosure chassis.
Off	Off	Off	Off	No ac power to the enclosure.	Turn on power.
Off	Off	Off	On	The ac power is on but power supply unit is not seated correctly in the enclosure.	Seat the power supply unit correctly in the enclosure.
Off	On	Off	On	No ac power to this power supply	<ol style="list-style-type: none"> 1. Check that the switch on the power supply unit is on. 2. Check that the ac power is on. 3. Reseat and replace the power cable.
On	Off	Off	Off	Power supply is on and operational.	No actions
Off	Off	On	Off	Fan failure	Replace the power supply unit.

Table 14. Power-supply unit LEDs (continued)

Power supply OK	ac failure	Fan failure	dc failure	Status	Action
Off	On	On	On	Communication failure and power supply problem	Replace the power supply unit. If replacing the power supply unit does not fix the problem, replace the enclosure chassis.
Flashing	X	X	X	No canister is operational.	Both canisters are either off or not seated correctly. Turn off the switch on both power supply units and then turn on both switches. If this action does not resolve the problem, remove both canisters slightly and then push the canisters back in.
Off	Flashing	Flashing	Flashing	Firmware is downloading.	No actions. Do not remove ac power.

- At least one power supply in the enclosure must indicate Power supply OK or Power supply firmware downloading for the node canisters to operate. For this situation, review the three canister status LEDs on each of the node canisters. Start with the power LED.

Table 15. Power LED status

Power LED status	Description
Off	There is no power to the canister. Try reseating the canister. Go to “Procedure: Reseating a node canister” on page 51. If the state persists, follow the hardware replacement procedures for the parts in the following order: node canister, enclosure chassis.
Slow flashing (1 Hz)	Power is available, but the canister is in standby mode. Try to start the node canister by reseating it. Go to “Procedure: Reseating a node canister” on page 51.
Fast flashing (2 Hz)	The canister is running its power-on self-test (POST). Wait for the test to complete. If the canister remains in this state for more than 10 minutes, try reseating the canister. Go to “Procedure: Reseating a node canister” on page 51. If the state persists, follow the hardware replacement procedure for the node canister.

Figure 20 shows the LEDs on the node canister.




Figure 20. LEDs on the node canisters

- If the power LED is on, consider the states of the cluster status and fault LEDs.

Table 16. Cluster status and fault LEDs

Cluster status LED	Fault LED	Status	Action
Off	Off	Code is not active.	<ul style="list-style-type: none"> Follow procedures for reviewing power LEDs. If the power LEDs show green, reseal the node canister. See "Procedure: Reseating a node canister" on page 51. If the LED status does not change, see "Replacing a node canister" on page 65.
Off	On	Code is not active. The BIOS or the service processor has detected a hardware fault.	Follow the hardware replacement procedures for the node canister.
On	Off	Code is active. Node state is active.	No action. The node canister is part of a cluster and can be managed by the management GUI.
On	On	Code is active and is in starting state. However, it does not have enough resources to form the cluster.	The node canister cannot become active in a cluster. There are no detected problems on the node canister itself. However, it cannot connect to enough resources to safely form a cluster. Follow the procedure to fix the node errors. Go to "Procedure: Fixing node errors" on page 48
Flashing	Off	Code is active. Node state is candidate.	Create a cluster on the node canister, or add the node canister to the cluster. If the other node canister in the enclosure is in active state, it automatically adds this node canister into the cluster. A node canister in this state can be managed using the service assistant.
Flashing	On	Code is active. Node state is service.	The node canister cannot become active in a cluster. Several problems can exist: hardware problem, a problem with the environment or its location, or problems with the code or data on the canister. Follow the procedure to fix the node errors. Go to "Procedure: Fixing node errors" on page 48.
Any	Flashing	The node canister is being identified so that you can locate it.	The fix procedures in the management GUI might have identified the component because it requires servicing. Continue to follow the fix procedures. The service assistant has a function to identify node canisters. If the identification LED is on in error, use the service assistant node actions to turn off the LED.

To review the status of the control enclosure batteries, see Table 17.

Table 17. Status on the control enclosure batteries

Battery Good	Battery Fault	Description	Action
On	Off	Battery is good and fully charged.	None
Flashing	off	Battery is good but not fully charged. The battery is either charging or a maintenance discharge is being performed.	None

Table 17. Status on the control enclosure batteries (continued)

Battery Good	Battery Fault	Description	Action
Off	On	Nonrecoverable battery fault.	Replace the battery. If replacing the battery does not fix the issue, replace the power supply unit.
Off	Flashing	Recoverable battery fault.	None

Procedure: Finding the status of the Ethernet connections

This procedure explains how to find the status of the Ethernet connections when you cannot connect.

Ensure that the software is active on the node before you begin this procedure. Go to “Procedure: Understanding the system status using the LEDs” on page 43. Ethernet port 1 must be connected to an active port on your Ethernet network.

Determine the state of the Ethernet LEDs by using one of the following methods:

- Use the USB key to obtain the most comprehensive information for the node status. Go to “Procedure: Getting node canister and cluster information using a USB key” on page 42.
The status, speed, and MAC address are returned for each port. Information is returned that identifies whether the node is the configuration node and if any node errors were reported.
- Examine the two LEDs of the Ethernet ports.
The activity LED flashes when there is activity on the connection. The link state LED must be permanently on. If it is off, the link is not connected.

Until the Ethernet link is connected, perform the following actions to check the port status each time:

1. Verify that each end of the cable is securely connected.
2. Verify that the port on the Ethernet switch or hub is configured correctly.
3. Connect the cable to a different port on your Ethernet network.
4. If the status is obtained using the USB key, review all the node errors that are reported.
5. Replace the Ethernet cable.
6. Follow the hardware replacement procedure for a node canister. See “Replacing a node canister” on page 65.

Procedure: Removing cluster data from a node canister

This procedure guides you through the process to remove cluster information from a node canister. The information that is removed includes configuration data, cache data, and location data.

Attention: If the enclosure reaches a point where the cluster data is not available on any node canister in the system, you will have to perform a cluster recovery. This recovery is an extended service action and might not be able to recover all your volumes. Do not perform this action to remove the cluster data from a node unless there is a node canister with saved cluster information in the enclosure. Do not remove the cluster data from a node unless instructed to do so by a service procedure.

1. Start the service assistant on the node canister.
2. Use the service assistant node action to hold the node in service state.
3. Use the **Manage Cluster** option to remove the cluster data from the node.

The node restarts in service state. When you want the node canister to be active again, use the service assistant home page action to leave service state.

Procedure: Deleting a cluster completely

This procedure guides you through the process to completely remove all cluster information from a system. When the procedure is finished, the system performs like a new installation.

Attention: This procedure makes all the volume data that you have on your system inaccessible. You cannot recover the data. This procedure affects all volumes that are managed by your system, which includes the drives in the control enclosure, in the expansion enclosures, and on the managed disks on external storage systems. The only volumes not affected are image mode volumes on external storage systems.

Do not continue unless you are certain that you want to remove all the volume data and configuration data from your system. This procedure is not used as part of any recovery action.

There are two stages to this procedure. First, the node canisters are reset. Second, the enclosure data is reset.

1. Start the service assistant on one of the node canisters.
2. Use the service assistant node action to hold the node in service state.
3. Use the **Manage Cluster** option to remove the cluster data from the node.
4. Perform the previous steps on the second node canister in the enclosure.

Note: Instead of connecting to the other node, you can make the second node the current node.

5. Open the service assistant **Configure Enclosure** and select **Reset Cluster ID** option.

This action causes the system to reset.

6. Power the system off and on before creating a cluster.

Procedure: Fixing node errors

This procedure describes how to fix a node error that is detected on one of the node canisters in your system.

Node errors are reported when there is an error that is detected that affects a specific node canister.

1. Use the service assistant to view the current node errors on any node.
2. If available, use the management GUI to run the recommended action for the alert.
3. Follow the fix procedure instructions.
4. If the recommended action does not provide sufficient information to determine the service action, review the node error descriptions and service actions. Go to “Error code range” on page 108.

See the node error descriptions if you cannot access the management GUI or if the management GUI is not reporting an alert because it cannot connect to the node. When you cannot connect to the management GUI, follow the procedure for getting node canister and cluster information using the service assistant. Go to “Procedure: Getting node canister and cluster information using the service assistant” on page 42. Start with the node that displays an error. The home page shows the node errors on the current node in the priority that you must service them. Start with the node error with the highest priority.

5. Select a different node in the system to see the node errors on that node.
6. Attempt to service the node errors in the priority order that they are listed.
7. Use the error number as an index when reviewing the node error descriptions. The service actions for each error are listed with the error code. Go to “Error code range” on page 108.

Procedure: Changing the service IP address of a node

This procedure identifies many methods that you can use to change the service IP address of a node.

When you change an IPv4 address, you change the IP address, the subnet, mask, and gateway. When you change an IPv6 address, you change the IP address, prefix, and gateway.

Which method to use depends on the status of the system and the other node canisters in the system. Follow the methods in the order shown until you are successful in setting the IP address to the required value.

You can set an IPv4 address, an IPv6 address, or both, as the service address of a node. Enter the required address correctly. If you set the address to 0.0.0.0 or 0000:0000:0000:0000:0000:0000, you disable the access to the port on that protocol.

Change the service IP address.

- Use the management GUI when the management GUI and the system are operating and the system is able to connect to the node with the service IP address that you want to change.
 1. Select **Configuration** → **Network** from the navigation.
 2. Select **Service IP Addresses**.
 3. Complete the panel. Be sure to select the correct node to configure.
- Use the service assistant when you can connect to the service assistant on either the node canister that you want to configure or on a node canister that can connect to the node canister that you want to configure:
 1. Make the node canister that you want to configure the current node.
 2. Select **Change Service IP** from the menu.
 3. Complete the panel.
- Use one of the following procedures if you cannot connect to the node canister from another node:
 - Use the initialization tool to write the correct command file to the USB key. Go to “Using the initialization tool” on page 29.
 - Use a text editor to create the command file on the USB key. Go to “Using a USB key” on page 29.

Procedure: Initializing a cluster with a USB key without using the initialization tool

Use this procedure to initialize a cluster using a USB key when you do not have a Microsoft Windows workstation to run the initialization tool or you do not have a copy of the tool.

In these situations, you must manually create an `satask.txt` file on a USB key to initialize your cluster. Use the USB key that was supplied with your system or any USB key that is formatted with a FAT32, an EXT2, or an EXT3 file system on its first partition.

1. Open a file editor that can create ASCII text files.
2. Create a file called `satask.txt`.
3. Add a single line of command text to the file.

If you are creating a cluster with an IPv4 address, the command line is like the following string:

```
satask mkcluster -clusterip aaa.aaa.aaa.aaa -gw ggg.ggg.ggg.ggg -mask mmm.mmm.mmm.mmm
```

where you must replace `aaa.aaa.aaa.aaa` with the management IP address, `ggg.ggg.ggg.ggg` with the network gateway address, and `mmm.mmm.mmm.mmm` with the subnet mask address.

If you are creating a cluster with an IPv6 address, the command line is like the following string:

```
satask mkcluster -clusterip_6 aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa  
-gw_6 gggg:gggg:gggg:gggg:gggg:gggg:gggg:gggg -prefix_6 pp
```


where you must replace `aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa` with the management IPv6 address, `gggg:gggg:gggg:gggg:gggg:gggg:gggg:gggg` with the network gateway IPv6 address, and `pp` with the prefix value.

4. Save the file to a USB key.
5. Plug the USB key into a USB port on a control canister.
6. The system detects the USB key, reads the `satask.txt` file, runs the command, and writes the results to the USB key. The `satask.txt` file is deleted after the command is run.
7. Wait for the fault LED on the node canister to stop flashing before removing the USB key.
8. Remove the USB key and insert it into your workstation to view the results.
9. Use a web browser to view the results file, `satask_result.html`.

Check that there were no errors returned by the command. If there is insufficient battery charge to protect the system, the cluster creates successfully, but the cluster does not start immediately. In the results, look for the `time_to_charge` field for the battery. The results provide an estimate of the time, in minutes, before the cluster can start. If the time is not 0, wait for the required time. Check that the node canister that you inserted the USB key into has its cluster state LED on permanently. For additional information, see “Procedure: Understanding the system status using the LEDs” on page 43.

10. If the cluster initialization was successful and the batteries are sufficiently charged, point a supported browser to the management IP address that you specified to start the management GUI. You see the management GUI logon panel.
11. Log on as superuser. Use `passwd` for the password.
12. Follow the on-screen instructions.

For more information about using the USB key, see “USB key and Initialization tool interface” on page 28.

Procedure: Initializing a cluster using the service assistant

Use this procedure to initialize a cluster using the service assistant rather than the USB key.

Note: The service assistant gives you the option to create a cluster only if the node state is candidate.

To initialize a cluster using the service assistant, perform the following steps:

1. Point your web browser to the service address of the upper node canister in your control enclosure.
2. Log on with the superuser password. The default password is `passwd`.

If you are unable to connect, see “Problem: Cannot connect to the service assistant” on page 38.

If the default service assistant address cannot be used in your network environment, connect using a direct Ethernet connection. To make this connection, see “Procedure: Accessing a canister using a directly attached Ethernet cable” on page 51.

3. Select **Manage Cluster**.
4. Enter the system name and the management IP address.
5. Click **Create cluster**.

Attention: Without a USB key to service the system, it is not possible to reset the superuser password or to change the service IP addresses in the event of a fault that prevents access to the management interface. It is essential that you take steps to record this information for use in the event of a failure.

Procedure: Accessing a canister using a directly attached Ethernet cable

Use this procedure if you need to use a direct Ethernet connection to attach a personal computer to a node canister to run the service assistant or to use the service CLI.

Perform this procedure if you are not authorized to use a USB key in your data center and when the service address of your nodes cannot be accessed over your Ethernet network. This situation might occur for a new installation where the default service IP addresses 192.168.70.121 and 190.168.70.122 cannot be accessed on your network.

Note: Do not attempt to use a directly attached Ethernet cable to a canister that is active in a cluster. You might disrupt access from host applications or the management GUI. If the node is active, use the management GUI network configuration options to set the service IP to an address that is accessible on the network.

Perform the following steps to access a canister using a directly attached Ethernet cable:

1. Connect one end of an Ethernet cable to Ethernet port 1 of the upper node canister.

Note: A cross-over Ethernet cable is not required.

2. Connect the other end of the Ethernet cable directly to the Ethernet port on a personal computer that has a web browser installed.
3. Use the operating system tools on the computer to set the IP address of the Ethernet port that is used in the previous step to 192.168.70.10.
4. Point your web browser to the service address.
 - If you have connected to node canister 1, the upper canister, point your web browser from your personal computer to <https://192.168.70.121>.
 - If you have connected to node canister 2, the lower canister, point your web browser from your personal computer to <https://192.168.70.122>.
5. Log on with the superuser password. The default password is `passwd`.
6. After the action completes, disconnect your personal computer and reconnect the node canister to the Ethernet network.
7. Set the service address of the canister to one that can be accessed on the network as soon as possible.

Procedure: Reseating a node canister

Use this procedure to reseat a canister that is in service state or because a service action has directed you.

Verify that you are reseating the correct node canister and that you use the correct canister handle for the node that you are reseating. Handles for the node canisters are located next to each other. The handle on the right operates the upper canister. The handle on the left operates the lower canister.

1. Verify the cluster status LED on the node canister. If it is permanently on, the node is active. If the node is active, no reseating is required.
2. Verify that you have selected the correct node canister and verify why you are reseating it. Go to “Procedure: Identifying which enclosure or canister to service” on page 41.

If you reseat a node that is active, it cannot store its state data and cannot restart without other service actions.

If the other node canister in the enclosure is not active, reseating the node canister while it is active results in loss of the data on your volumes and the system is unavailable to hosts.

3. Grasp the handle between the thumb and forefinger.
4. Squeeze them together to release the handle.

5. Pull out the handle to its full extension.
6. Grasp the canister and pull it out 2 or 3 inches.
7. Push the canister back into the slot until the handle starts to move.
8. Finish inserting the canister by closing the handle until the locking catch clicks into place.
9. Verify that the cables were not displaced.
10. Verify that the LEDs are on.

Procedure: Powering off your system

Use this procedure to power off your system when it must be serviced or to permit other maintenance actions in your data center.

To power off your system, use the following steps:

1. Stop hosts.
2. Shut down the cluster using the management GUI. Click **Home** → **Manage Device**. From the **Actions** menu, select **Shut down Cluster**.
3. Using the power switches, power off the control enclosure.
4. Using the power switches, power off the expansion enclosures.
5. (Optional) Shut down external storage systems.
6. (Optional) Shut down fibre-channel switches.

Procedure: Collecting information for support

IBM support might ask you to collect trace files and dump files from your system to help them resolve a problem.

The management GUI and the service assistant have features to assist you in collecting the required information. The management GUI collects information from all the components in the system. The service assistant collects information from a single node canister. When the information that is collected is packaged together in a single file, the file is called a *snap*.

Special tools that are only available to the support teams are required to interpret the contents of the support package. The files are not designed for customer use.

Always follow the instructions that are given by the support team to determine whether to collect the package using the management GUI or the service assistant. Instruction is also given for which package content option is required.

- If you are collecting the package using the management GUI, select **Troubleshooting** → **Support**. Select the option to **Download Support Package**. Follow the instructions.
- If you are collecting the package using the service assistant, ensure that the node that you want to collect logs from is the current node. Select the **Collect Logs** option from the navigation. You can collect a support package or copy an individual file from the node canister. Follow the instructions to collect the information.

SAN problem determination

The procedures that are provided here help you solve problems on the Storwize V7000 system and its connection to the storage area network (SAN).

SAN failures might cause Storwize V7000 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 the error codes:

1. Verify that the power is turned on to all switches and storage controllers that the Storwize V7000 system uses, and that they are not reporting any hardware failures. If problems are found, resolve those problems before proceeding further.
2. Verify that the Fibre Channel cables that connect the systems to the switches are securely connected.
3. If you have a SAN management tool, use that tool to view the SAN topology and isolate the failing component.

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 status LEDs at the rear of the Storwize V7000 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 Storwize V7000 node.

Note: Storwize V7000 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. Contact IBM Support for assistance in replacing the node canister.

Servicing storage systems

Storage systems that are supported for attachment to the Storwize V7000 system are designed with redundant components and access paths to enable concurrent maintenance. Hosts have continuous access to their data during component failure and replacement.

The following guidelines apply to all storage systems that are attached to the Storwize V7000 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

Chapter 6. Recovery procedures

This topic describes two recovery procedures: recover a cluster and back up and restore a cluster configuration.

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 node canisters 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 node canisters are in candidate status.
- Run the procedure on one cluster in a fabric at a time. Do not perform the procedure on different node canisters 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.

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 canister that is part of this system. Resolve all hardware errors except node error 578 or node error 550.
 - All nodes must be reporting either a node error 578 or a node error 550. These error codes indicate that the 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 550 node error. The 550 error can also indicate that a node is not able to join a cluster.
 - If any nodes show a node error 550, record the error data that is associated with the 550 error from the service assistant.
 - In addition to the node error 550, the report can show data that is separated by spaces in one of the following forms:

- Node identifiers in the format: <enclosure_serial>-<canister slot ID><7 characters, hyphen, 1 number), for example, 01234A6-2
 - Quorum drive identifiers in the format: <enclosure_serial>:<drive slot ID>[<drive 11S serial number>] (7 characters, colon, 1 or 2 numbers, open square bracket, 22 characters, close square bracket), for example, 01234A9:21[11S1234567890123456789]
 - Quorum MDisk identifier in the format: WWPN/LUN (16 hexadecimal digits followed by a forward slash and a decimal number), for example, 1234567890123456/12
- If the error data contains a node identifier, ensure that the node that is referred to by the ID is showing node error 578. If the node is showing a node error 550, ensure that the two nodes can communicate with each other. Verify the SAN connectivity and restart one of the two nodes by clicking **Restart Node** from the service assistant.
 - If the error data contains a quorum drive identifier, locate the enclosure with the reported serial number. Verify that it is powered on and that the node canister that is reporting the fault has SAS connectivity to the listed enclosure. Also verify that the drive in the reported slot is powered on and functioning. After verifying these things, restart the node by clicking **Restart Node** from the service assistant.
 - If the error data contains a quorum MDisk identifier, verify the SAN connectivity between this node and that WWPN. Check the storage controller to ensure that the LUN referred to is online. After verifying these things, restart the node by clicking **Restart Node** from the service assistant.

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. You can also call IBM Support for further assistance.

- For any nodes that are reporting a node error 550, ensure that all the missing hardware that is identified by these errors is powered on and connected without faults. If you cannot contact the service assistant from any node, isolate the problems using the LED indicators.
 - 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. For example, both power supplies have been removed.

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 node canisters are in candidate status. If there are any node canisters that display error code 550 or error code 578, you must remove their cluster data.

To remove cluster information from a node canister 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 using the service assistant

Start recovery when all node canisters that were members of the cluster are online and are in candidate status. Do not run the recovery procedure on different node canisters 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 node canisters in the cluster. The node canisters 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 canister that was in I/O group 0.

1. Point your browser to the service IP address of one of the node canisters.
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.
- 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 Storwize V7000 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 Storwize V7000 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, Storwize V7000, 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 Storwize V7000.

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 60.

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 55.

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. Use the USB key to initialize the system with the IP address.

3. In a supported browser, enter the IP address that you used to initialize the system, the superuser name, and the password. On the Configure Storage panel, deselect **Yes, automatically configure internal storage now**. Any internal storage configuration is recovered after the system is restored.
4. 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.
5. 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.

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.
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:

```
svconfig 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:

```
svconfig 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 `svconfig 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:

```
svconfig 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_2076:your_cluster_name:admin>
```

17. After the cluster configuration is restored, verify that the quorum disks are restored to the MDisk that you want by using the `svcinfolsqorum` command. To restore the quorum disks to the correct MDisk, 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
```

Chapter 7. Removing and replacing parts

You can remove and replace field-replaceable units (FRUs) from the control enclosure or the expansion enclosure.

Attention: If your system is powered on and performing I/O operations, go to the management GUI and follow the fix procedures. Performing the replacement actions without the assistance of the fix procedures can result in loss of data or access to data.

Even though many of these procedures are hot-swappable, these procedures are intended to be used only when your system is not up and running and performing I/O operations. Unless your system is offline, go to the management GUI and follow the fix procedures.

Each replaceable unit has its own removal procedure. Sometimes you can find that a step within a procedure might refer you to a different remove and replace procedure. You might want to complete the new procedure before you continue with the first procedure that you started.

Remove or replace parts only when you are directed to do so.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Preparing to remove and replace parts

Before you remove and replace parts, you must be aware of all safety issues.

First, read the safety precautions in the *IBM Systems Safety Notices*. These guidelines help you safely work with the Storwize V7000.

Replacing a node canister

This topic describes how to replace a node canister.

Attention: If your system is powered on and performing I/O operations, go to the management GUI and follow the fix procedures. Performing the replacement actions without the assistance of the fix procedures can result in loss of data or access to data.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.


Be aware of the following canister LED states:

- If both the power LED and cluster status LED are on, do not remove a node canister unless directed to do so by a service procedure.
- If the cluster status is off, it is acceptable to remove a node canister. However, do not remove a node canister unless directed to do so by a service procedure.
- If the power LED is flashing or off, it is safe to remove a node canister. However, do not remove a node canister unless directed to do so by a service procedure.

Attention: Even if a node canister is powered off, it is still possible to lose data. Do not remove a node canister unless directed to do so by a service procedure.

To replace the node canister, perform the following steps:


1. Read the safety information to which “Preparing to remove and replace parts” on page 65 refers.
2. Confirm that you know which canister to replace. Go to “Procedure: Identifying which enclosure or canister to service” on page 41.
3. Record which data cables are plugged into the specific ports of the node canister. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
4. Disconnect the data cables for each canister.
5. Grasp the handle between the thumb and forefinger.
 Handles for the node canisters are located next to each other. The handle on the right operates the upper canister. The handle on the left operates the lower canister.



svc00649

Figure 21. Opening the handle on a canister

6. Squeeze them together to release the handle.



svc00631

Figure 22. Removing the canister from the enclosure

7. Pull out the handle to its full extension.
8. Grasp canister and pull it out.

9. Insert the new canister into the slot with the handle pointing towards the center of the slot. Insert the unit in the same orientation as the one that you removed.
10. Push the canister back into the slot until the handle starts to move.
11. Finish inserting the canister by closing the handle until the locking catch clicks into place.
If the enclosure is powered on, the canister starts automatically.
12. Reattach the data cables.

Replacing an expansion canister

This topic describes how to replace an expansion canister.

Attention: If your system is powered on and performing I/O operations, go to the management GUI and follow the fix procedures. Performing the replacement actions without the assistance of the fix procedures can result in loss of data or access to data.

Even though many of these procedures are hot-swappable, these procedures are intended to be used only when your system is not up and running and performing I/O operations. Unless your system is offline, go to the management GUI and follow the fix procedures.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Be aware of the following canister LED states:


- If the power LED is on, do not remove an expansion canister unless directed to do so by a service procedure.
- If the power LED is flashing or off, it is safe to remove an expansion canister. However, do not remove an expansion canister unless directed to do so by a service procedure.

Attention: Even if an expansion canister is powered off, it is still possible to lose data. Do not remove an expansion canister unless directed to do so by a service procedure.

To replace an expansion canister, perform the following steps:

1. Read the safety information to which “Preparing to remove and replace parts” on page 65 refers.
2. Record which SAS cables are plugged into the specific ports of the expansion canister. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
3. Disconnect the SAS cables for each canister.
4. Grasp the handle between the thumb and forefinger.


Handles for the expansion canisters are located above and below each other. The upper handle operates the upper canister. The lower handle operates the lower canister.



svc00649

Figure 23. Opening the handle on a canister

5. Squeeze them together to release the handle.



svc00631

Figure 24. Removing the canister from the enclosure

6. Pull out the handle to its full extension.
7. Grasp canister and pull it out.
8. Insert the new canister into the slot with the handle pointing towards the center of the slot. Insert the unit in the same orientation as the one that you removed.
9. Push the canister back into the slot until the handle starts to move.
10. Finish inserting the canister by closing the handle until the locking catch clicks into place.
11. Reattach the SAS cables.

Replacing the SFP transceiver

When a failure occurs on a single Fibre Channel link, the SFP transceiver might need to be replaced.

Even though many of these procedures are hot-swappable, these procedures are intended to be used only when your system is not up and running and performing I/O operations. Unless your system is offline, go to the management GUI and follow the fix procedures.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

CAUTION:

Some laser products contain an embedded Class 3A or Class 3B laser diode. Note the following information: laser radiation when open. Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam. (C030)

Perform the following steps to remove and then replace the SFP transceiver:

1. Carefully determine the failing physical port connection.

Important: The enclosures are supported with both longwave SFP transceivers and shortwave SFP transceivers. A longwave SFP transceiver has some blue components that are visible even when the SFP transceiver is plugged in. You must replace an SFP transceiver with the same type of SFP transceiver that you are replacing. If the SFP transceiver to replace is a longwave SFP transceiver, for example, you must replace with another longwave SFP transceiver. Removing the wrong SFP transceiver might result in loss of data access.

2. Remove the Fibre Channel cable by pressing the release tab and pulling the cable out. Be careful to exert pressure only on the connector and do not pull on the Fibre Channel cables.
3. Remove the SFP transceiver. There are a number of different handling or locking mechanisms that are used on the SFP transceivers. Some SFP transceivers might have a plastic tag. If so, pull the tag to remove the SFP transceiver.

Important: Always check that the SFP transceiver that you replace matches the SFP transceiver that you remove.

4. Push the new SFP transceiver into the aperture and ensure that it is securely pushed home. The SFP transceiver usually locks into place without having to swing the release handle until it locks flush with the SFP transceiver. Figure 25 illustrates an SFP transceiver and its release handle.



Figure 25. Fibre-channel SFP transceiver

5. Reconnect the Fibre Channel cable.
6. Confirm that the error is now fixed. Either mark the error as fixed or restart the node depending on the failure indication that you originally noted.

Replacing a power supply unit for a control enclosure

You can replace either of the two 764 watt hot-swap redundant power supplies in the control enclosure. These redundant power supplies operate in parallel, one continuing to power the canister if the other fails.

DANGER

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

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

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

To disconnect:

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

To connect:

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

(D005)

Attention: If your system is powered on and performing I/O operations, go to the management GUI and follow the fix procedures. Performing the replacement actions without the assistance of the fix procedures can result in loss of data or access to data.

Attention: A powered-on enclosure must not have a power supply removed for more than five minutes because the cooling does not function correctly with an empty slot. Ensure that you have read and understood all these instructions and have the replacement available, and unpacked, before you remove the existing power supply.

Even though many of these procedures are hot-swappable, these procedures are intended to be used only when your system is not up and running and performing I/O operations. Unless your system is offline, go to the management GUI and follow the fix procedures.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.


Attention: In some instances, it might not be advisable to remove a power supply unit when a system is performing I/O. For example, the charge in the backup battery might not be sufficient enough within the partner power-supply unit to continue operations without causing a loss of access to the data. Wait until the partner battery is 100% charged before replacing the power supply unit.

Ensure that you are aware of the procedures for handling static-sensitive devices before you replace the power supply.

This task assumes that you are not replacing a battery in the power supply unit of the control enclosure. To replace a battery, go to “Replacing a battery in a power supply unit” on page 76.

To replace the power supply, perform the following steps:

1. Read the safety information to which “Preparing to remove and replace parts” on page 65 refers.
2. Examine the Identify LED that is lit on the front of the enclosure to identify the correct enclosure.
3. Turn off the power to the power supply units using the switches at the back of the units.
4. Disconnect the cable retention bracket and the power cord from the power supply that you are replacing.
5. Remove the power supply unit. Record the orientation of the power supply unit. Power supply unit 1 is top side up, and power supply unit 2 is inverted.
 - a. Depress the black locking catch from the side with the colored sticker as shown in Figure 26.



svc00633

Figure 26. Directions for lifting the handle on the power supply unit

- b. Grip the handle to pull the power supply out of the enclosure as shown in Figure 27 on page 72.




Figure 27. Using the handle to remove a power supply unit

6. Insert the replacement power supply unit into the enclosure with the handle pointing towards the center of the enclosure. Insert the unit in the same orientation as the one that you removed.
7. Push the power supply unit back into the enclosure until the handle starts to move.
8. Finish inserting the power supply unit into the enclosure by closing the handle until the locking catch clicks into place.
9. Reattach the power cable and cable retention bracket.
10. Turn on the power switch to the power supply unit.

If required, return the power supply. Follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Replacing a power supply unit for an expansion enclosure

You can replace either of the two 580 watt hot-swap redundant power supplies in the expansion enclosure. These redundant power supplies operate in parallel, one continuing to power the canister if the other fails.

DANGER

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

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

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

To disconnect:

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

To connect:

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

(D005)

Attention: If your system is powered on and performing I/O operations, go to the management GUI and follow the fix procedures. Performing the replacement actions without the assistance of the fix procedures can result in loss of data or access to data.

Attention: A powered-on enclosure must not have a power supply removed for more than five minutes because the cooling does not function correctly with an empty slot. Ensure that you have read and understood all these instructions and have the replacement available, and unpacked, before you remove the existing power supply.


Even though many of these procedures are hot-swappable, these procedures are intended to be used only when your system is not up and running and performing I/O operations. Unless your system is offline, go to the management GUI and follow the fix procedures.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Ensure that you are aware of the procedures for handling static-sensitive devices before you replace the power supply.

To replace the power supply unit in an expansion enclosure, perform the following steps:


1. Read the safety information to which “Preparing to remove and replace parts” on page 65 refers.
2. Examine the Identify LED that is lit on the front of the enclosure to identify the correct enclosure.
3. Turn off the power to the power supply units using the switches at the back of the units.
4. Disconnect the cable retention bracket and the power cord from the power supply that you are replacing.
5. Remove the power supply unit. Record the orientation of the power supply unit. Power supply unit 1 is top side up, and power supply unit 2 is inverted.
 - a. Depress the black locking catch from the side with the colored sticker as shown in Figure 28.



svc00633

Figure 28. Directions for lifting the handle on the power supply unit

- b. Grip the handle to pull the power supply out of the enclosure as shown in Figure 29 on page 75.



svc00632

Figure 29. Using the handle to remove a power supply unit

6. Insert the replacement power supply unit into the enclosure with the handle pointing towards the center of the enclosure. Insert the unit in the same orientation as the one that you removed.
7. Push the power supply unit back into the enclosure until the handle starts to move.
8. Finish inserting the power supply unit in the enclosure by closing the handle until the locking catch clicks into place.
9. Reattach the power cable and cable retention bracket.
10. Turn on the power switch to the power supply unit.

If required, return the power supply. Follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Replacing a battery in a power supply unit

This topic describes how to replace the battery in the control enclosure power-supply unit.

DANGER

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

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

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

To disconnect:

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

To connect:

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

(D005)

CAUTION:

The battery is a lithium ion battery. To avoid possible explosion, do not burn. Exchange only with the IBM-approved part. Recycle or discard the battery as instructed by local regulations. In the United States, IBM has a process for the collection of this battery. For information, call 1-800-426-4333. Have the IBM part number for the battery unit available when you call. (C007)

Attention: If your system is powered on and performing I/O operations, go to the management GUI and follow the fix procedures. Performing the replacement actions without the assistance of the fix procedures can result in loss of data or access to data.

Even though many of these procedures are hot-swappable, these procedures are intended to be used only when your system is not up and running and performing I/O operations. Unless your system is offline, go to the management GUI and follow the fix procedures.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Each power supply unit in a control enclosure contains an integrated battery that is used during temporary short-term power outages. You must replace the battery with the exact same model.

To replace the battery in the power supply unit of the control enclosure, perform the following steps:

1. Read the safety information to which “Preparing to remove and replace parts” on page 65 refers.
2. Follow the removing steps of the replacing a power-supply unit procedure. Go to “Replacing a power supply unit for a control enclosure” on page 70.
3. Remove the battery, as shown in Figure 30.




Figure 30. Removing the battery from the control enclosure power-supply unit

- a. Press the catch to release the handle **1**.

svc00611

- b. Lift the handle on the battery **2**
 - c. Lift the battery out of the power supply unit **3**.
4. Install the replacement battery.

Attention: The replacement battery has protective end caps that must be removed prior to use.

 - a. Remove the battery from the packaging.
 - b. Remove the end caps.
 - c. Attach the end caps to both ends of the battery that you removed and place the battery in the original packaging.
 - d. Place the replacement battery in the opening on top of the power supply in its proper orientation.
 - e. Press the battery to seat the connector.
 - f. Place the handle in its downward location
 5. Push the power supply unit back into the enclosure until the handle starts to move.
 6. Finish inserting the power supply unit into the enclosure by closing the handle until the locking catch clicks into place.
 7. Reattach the power cable and cable retention bracket.
 8. Turn on the power switch to the power supply unit.

If required, return the battery. Follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Releasing the cable retention bracket

This topic provides instructions for releasing the cable retention bracket when removing the power cords from the power supply unit.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Each cable retention bracket comes attached to the back of the power supply unit by the power cord plug-in.

To release a cable retention bracket, perform these steps:

1. Unlock the cable retention bracket that is around the end of the power cord as shown in Figure 31 on page 79.



svc00686

Figure 31. Unlocking the cable retention bracket

2. Pull the lever next to the black plastic loop towards you as shown in Figure 32 on page 80.



Figure 32. Pressing the lever to slide the cable retention bracket.

3. Continue to pull the lever towards you as you slide the cable retention bracket away from the end of the cable.

Replacing a 3.5" drive assembly or blank carrier

This topic describes how to replace a 3.5" drive assembly or blank carrier.


Attention: If your drive is configured for use, go to the management GUI and follow the fix procedures. Performing the replacement actions without the assistance of the fix procedures results in loss of data or access to data.

Attention: Do not leave a drive slot empty. Do not remove a drive or drive assembly before you have a replacement available.

The drives can be distinguished from the blank carriers by the color-coded striping on the drive. The drives are marked with an orange striping. The blank carriers are marked with a blue striping.

To replace the drive assembly or blank carrier, perform the following steps:


1. Read the safety information to which "Preparing to remove and replace parts" on page 65 refers.
2. Unlock the assembly by squeezing together the tabs on the side.



svc00612

Figure 33. Unlocking the 3.5" drive

3. Open the handle to the full extension.



svc00613

Figure 34. Removing the 3.5" drive

4. Pull out the drive.
5. Push the new drive back into the slot until the handle starts to move.
6. Finish inserting the drive by closing the handle until the locking catch clicks into place.

Replacing a 2.5" drive assembly or blank carrier


This topic describes how to remove a 2.5" drive assembly or blank carrier.

Attention: If your drive is configured for use, go to the management GUI and follow the fix procedures. Performing the replacement actions without the assistance of the fix procedures results in loss of data or access to data.

Attention: Do not leave a drive slot empty. Do not remove a drive or drive assembly before you have a replacement available.

To replace the drive assembly or blank carrier, perform the following steps:

1. Read the safety information to which "Preparing to remove and replace parts" on page 65 refers.
2. Unlock the module by squeezing together the tabs at the top.



svc00614

Figure 35. Unlocking the 2.5" drive

3. Open the handle to the full extension.




Figure 36. Removing the 2.5" drive

4. Pull out the drive.
5. Push the new drive back into the slot until the handle starts to move.
6. Finish inserting the drive by closing the handle until the locking catch clicks into place.

Replacing an enclosure end cap

This topic describes how to replace an enclosure end cap.

To replace the enclosure end cap, perform the following steps:

1. On either side of the drive assemblies, remove the enclosure end caps by squeezing the middle of the cap and pulling it away from the front of the rack.
2. Reattach the end cap by relocating it on either side of the drive assemblies and gently pushing it on.

Replacing a SAS cable

This topic describes how to replace a SAS cable.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

To replace a SAS cable, perform the following steps:

1. Record which SAS cable is plugged into the specific port of the expansion canister. The cable must be inserted back into the same port after the replacement is complete; otherwise, the system cannot function properly.

Note: If you are replacing a single cable, this step is not necessary.

2. Pull the tab with the arrow away from the connector.



Figure 37. SAS cable

3. Plug the replacement cable into the specific port.
4. Ensure that the SAS cable is fully inserted. A click is heard when the cable is successfully inserted.

Replacing a control enclosure chassis

This topic describes how to replace a control enclosure chassis.

Note: Ensure that you know the type of enclosure chassis that you are replacing. The procedures for replacing a control enclosure chassis are different from those procedures for replacing an expansion enclosure chassis. For information about replacing an expansion enclosure chassis, see “Replacing an expansion enclosure chassis” on page 87.

DANGER

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

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

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

To disconnect:

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

To connect:

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

(D005)

Attention: Perform this procedure only if instructed to do so by a service action or the IBM support center. This procedure requires that you shut down your Storwize V7000 system to replace the control enclosure. If the system is still performing I/O requests, schedule the replacement during a maintenance period or other time when the I/O can be stopped.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Ensure that you are aware of the procedures for handling static-sensitive devices before you remove the enclosure.

To replace a control enclosure chassis, perform the following steps:

1. If you are able to access either of the node canisters with the service assistant, record the machine type and model of the enclosure, the serial number of the enclosure, and the two WWNNs for the enclosure.
 - From the service assistant home page, open the location data for the node. Record the machine type and model (MTM), the serial number, WWNN 1 and WWNN 2 from the enclosure column.

- If you are replacing the enclosure because neither node canister can start, retrieve this information after you have completed the replacement.
 - a. Start the service assistant on one of the canisters.
 - b. Go to the node location data on the home page.
 - c. Record the machine type and model, the serial number, WWNN 1 and WWNN 2 from the node copy column.

The machine type and model and the serial number are also shown on the labels at the front and back of the enclosure.

2. If the enclosure is still active, shut down the host I/O and the Metro Mirror and Global Mirror activity to the enclosure.
3. If the cluster is still online, shut it down.
 - a. From the management GUI, go to **Home** → **Manage Device**.
 - b. Select **Shut Down Cluster** from the **Actions** menu.
 - c. Wait for the shutdown to complete.

4. Verify that it is safe to remove the power from the enclosure.

For each of the canisters, verify the status of the cluster status LED. If the LED is lit on either of the canisters, do not continue because the cluster is still online. Determine why the cluster did not shut down in step 3.

Note: If you continue while the cluster is still active, you risk losing the cluster configuration and volume cache data that is stored in the canister.

5. Turn off the power to the enclosure using the switches on the power supply units.
6. Record which data cables are plugged into the specific ports. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
7. Disconnect the cable retention brackets and the power cords from the power supply units.
8. Disconnect the data cables for each canister.
9. Remove the power supply units from the enclosure.
10. Remove the canisters from the enclosure. Record the location of each canister. They must be inserted back into the same location in the new enclosure.
11. Remove all the drives and blank drive assemblies from the enclosure. Record the location for each drive. They must be inserted back into the same location in the new enclosure.
12. Remove both enclosure end caps from the enclosure. Keep the left end cap because it is used again.
13. Remove the clamping screws that attached the enclosure to the rack cabinet.
14. Remove the enclosure chassis from the front of the rack cabinet and take the chassis to a work area.
15. Install the new enclosure chassis in the rack cabinet.
16. Remove the end caps from the new enclosure and install the clamping screws that attach the enclosure to the rack cabinet.
17. Replace the end caps. Use the new right end cap and use the left end cap that you removed in step 12.

Using the left end cap that you removed preserves the model and serial number identification.

18. Reinstall the drives in the new enclosure. The drives must be inserted back into the same location from which they were removed on the old enclosure.
19. Reinstall the canisters in the enclosure. The canisters must be inserted back into the same location from which they were removed on the old enclosure.
20. Install the power supply units.
21. Reattach the data cables to each canister using the information that you recorded previously.

Note: The cables must be inserted back into the same ports from which they were removed on the old enclosure; otherwise, the system cannot function properly.

22. Attach the power cords and the cable retention brackets to the power supply units.
23. Write the old enclosure machine type and model (MTM) and serial number on the repair identification (RID) tag that is supplied. Attach the tag to the left flange at the back of the enclosure.
24. Turn on the power to the enclosure using the switches on the power supply units.

The node canisters boot up. The fault LEDs are on because the new enclosure has not been set with the identity of the old enclosure. The node canisters report that they are in the wrong location.

- a. Connect to the service assistant on one of the node canisters to configure the machine type and model, serial number, and WWNNs that are stored in the enclosure. If you have replaced a node canister, connect to the canister that has not been replaced.
- b. Use the **Configure enclosure** panel.
- c. Select the options to **Update WWNN 1**, **Update WWNN 2**, **Update the machine type and model**, and **Update the serial number**. Do not update the cluster ID. Use the node copy data for each of the values. Check that these values match the values that you recorded in 1 on page 85.
If you were not able to record the values, use the node copy values only if none of them have all zeroes as their value. If any of the node copy values are all zeroes, connect the service assistant to the other node canister and configure the enclosure there. If you still do not have a full set of values, contact IBM support.
- d. After you modify the configuration, the nodes restart and the cluster starts.

Note: The cluster starts and can handle I/O requests from the host systems. The configuration changes that are described in the following steps must be performed to ensure that the system is operating correctly. If you do not perform these steps, the system is unable to report certain errors.

25. Start the management GUI and select **Home** → **Manage Devices**. You see an additional enclosure in the system list because the system has detected the replacement control enclosure. The original control enclosure is still listed in its configuration. The original enclosure is listed with its original enclosure ID. It is offline and managed. The new enclosure has a new enclosure ID. It is online and unmanaged.
26. Select the original enclosure in the tree view.
Verify that it is offline and managed and that the serial number is correct.
27. From the **Actions** menu, select **Remove enclosure** and confirm the action. The physical hardware has already been removed. You can ignore the messages about removing the hardware. Verify that the original enclosure is no longer listed in the tree view.
28. Add the new enclosure to the cluster.
 - a. Select the cluster name from the tree view.
 - b. From the **Actions** menu, select **Add enclosure**.
 - c. Because you have already added the hardware, select **Next** on the first panel that asks you to install the hardware. The next panel shows the unmanaged new enclosure.
 - d. Select the enclosure and add it to the system.
29. Select the new enclosure in the tree view and verify that it is now online and managed.
30. Change the enclosure ID of the replaced enclosure to that of the original enclosure. From the **Enclosure ID** field, select the ID value of the original enclosure.
31. Check the status of all volumes and physical storage to ensure everything is online.
32. Restart the host application and any FlashCopy activities, Global Mirror activities, or Metro Mirror activities that were stopped.

Replacing an expansion enclosure chassis

This topic describes how to replace an expansion enclosure chassis.

Note: Ensure that you know the type of enclosure chassis that you are replacing. The procedures for replacing an expansion enclosure chassis are different from those procedures for replacing a control enclosure chassis. For information about replacing a control enclosure chassis, see “Replacing a control enclosure chassis” on page 84.

DANGER

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

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

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

To disconnect:

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

To connect:

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

(D005)

Attention: If your system is powered on and performing I/O operations, go to the management GUI and follow the fix procedures. Performing the replacement actions without the assistance of the fix procedures can result in loss of data or access to data.

Even though many of these procedures are hot-swappable, these procedures are intended to be used only when your system is not up and running and performing I/O operations. Unless your system is offline, go to the management GUI and follow the fix procedures.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Ensure that you are aware of the procedures for handling static-sensitive devices before you remove the enclosure.

Note: If your system is online, replacing an expansion enclosure can cause one or more of your volumes to go offline or your quorum disks to be inaccessible. Before you proceed with these procedures, verify which volumes might go offline. From the management GUI, go to **Home** → **Manage Devices**. Select the enclosure that you want to replace. Then select **Show Dependent Volumes** in the **Actions** menu.

To replace an expansion enclosure chassis, perform the following steps:

1. Shut down the I/O activity to the enclosure, which includes host access, FlashCopy, Metro Mirror and Global Mirror access.
2. Turn off the power to the enclosure using the switches on the power supply units.
3. Record which data cables are plugged into the specific ports. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
4. Disconnect the cable retention brackets and the power cords from the power supply units.
5. Disconnect the data cables for each canister.
6. Remove the power supply units from the enclosure.
7. Remove the canisters from the enclosure.
8. Remove all the drives and blank drive assemblies from the enclosure. Record the location for each drive. They must be inserted back into the same location in the new enclosure.
9. Remove both enclosure end caps from the enclosure. Keep the left end cap because it is used again.
10. Remove the clamping screws that attached the enclosure to the rack cabinet.
11. Remove the enclosure chassis from the front of the rack cabinet and take the chassis to a work area.
12. Install the new enclosure chassis in the rack cabinet.
13. Remove the end caps from the new enclosure and install the clamping screws that attach the enclosure to the rack cabinet.
14. Replace the end caps. Use the new right end cap and use the left end cap that you removed in step 9.
Using the left end cap that you removed preserves the model and serial number identification.
15. Reinstall the drives in the new enclosure. The drives must be inserted back into the same location from which they were removed on the old enclosure.
16. Reinstall the canisters in the enclosure.
17. Install the power supply units.
18. Reattach the data cables to each canister using the information that you recorded previously.

Note: The cables must be inserted back into the same ports from which they were removed on the old enclosure; otherwise, the system cannot function properly.

19. Attach the power cords and the cable retention brackets to the power supply units.
20. Write the old enclosure machine type and model (MTM) and serial number on the repair identification (RID) tag that is supplied. Attach the tag to the left flange at the back of the enclosure.
21. Turn on the power to the enclosure using the switches on the power supply units.

The system records an error that indicates that an enclosure FRU replacement was detected. Go to the management GUI to use the fix procedure to change the machine type and model and serial number in the expansion enclosure.

Replacing the support rails

This topic describes how to replace the support rails.

Perform the following steps to replace the support rails:

1. Remove the enclosure.
2. Record the location of the rail assembly in the rack cabinet.

- Working from the back of the rack cabinet, remove the clamping screw **1** from the rail assembly on both sides of the rack cabinet.




Figure 38. Removing a rail assembly from a rack cabinet

- Working from the front of the rack cabinet, remove the clamping screw from the rail assembly on both sides of the rack cabinet.
- From one side of the rack cabinet, grip the rail and slide the rail pieces together to shorten the rail.
- Disengage the rail location pins **2**.
- From the other side the rack cabinet, grip the rail and slide the rail pieces together to shorten the rail.
- Disengage the rail location pins **2**.
- Starting from the location of the previous rail assembly, align the bottom of the rail with the bottom of the two rack units. Insert the rail location pins through the holes in the rack cabinet.
- Insert a clamping screw into the upper mounting hole between the rail location pins.
- Tighten the screw to secure the rail to the rack.
- Working from the rear of the rack cabinet, extend the rail that you secured to the front to align the bottom of the rail with the bottom of the two rack units.

Note: Ensure that the rail is level between the front and the back.

- Insert the rail location pins through the holes in the rack cabinet.
- Insert a clamping screw into the upper mounting hole between the rail location pins.
- Tighten the screw to secure the rail to the rack from the back side.
- Repeat the steps to secure the opposite rail to the rack cabinet.

Storwize V7000 replaceable units

The Storwize V7000 consists of several replaceable units. Generic replaceable units are cables, SFP transceivers, canisters, power supply units, battery assemblies, and enclosure chassis.

Table 18 provides a brief description of each replaceable unit.

Table 18. Replaceable units

Part	Part number	Applicable models	FRU or customer replaced
2U24 enclosure chassis (empty chassis)	85Y5897	124, 224	FRU

Table 18. Replaceable units (continued)

Part	Part number	Applicable models	FRU or customer replaced
2U14 enclosure chassis (empty chassis)	85Y5896	112, 212	FRU
Node canister	85Y5899	112, 124	Customer replaced
Expansion canister	85Y5850	212, 224	Customer replaced
764 W power supply unit	85Y5847	112, 124	Customer replaced
580 W power supply unit	85Y5846	212, 224	Customer replaced
Battery backup unit	85Y5898	112, 124	Customer replaced
1 m SAS cable	44V4041	212, 224	Customer replaced
3 m SAS cable	44V4163	212, 224	Customer replaced
6 m SAS cable	44V4164	212, 224	Customer replaced
1 m fibre-channel cable	39M5699	112, 124	Customer replaced
5 m fibre-channel cable	39M5700	112, 124	Customer replaced
25 m fibre-channel cable	39M5701	112, 124	Customer replaced
1.8 m power cord (Chicago)	39M5080	All	Customer replaced
2.8 m power cord (EMEA)	39M5151	All	Customer replaced
2.8 m power cord (Australia)	39M5102	All	Customer replaced
2.8 m power cord (Africa)	39M5123	All	Customer replaced
2.8 m power cord (Denmark)	39M5130	All	Customer replaced
2.8 m power cord (South Africa)	39M5144	All	Customer replaced
2.8 m power cord (Switzerland)	39M5158	All	Customer replaced
2.8 m power cord (Chile)	39M5165	All	Customer replaced
2.8 m power cord (Israel)	39M5172	All	Customer replaced
2.8 m power cord (Group 1 including the United States)	39M5081	All	Customer replaced
2.8 m power cord (Argentina)	39M5068	All	Customer replaced
2.8 m power cord (China)	39M5206	All	Customer replaced
2.8 m power cord (Taiwan)	39M5247	All	Customer replaced
2.8 m power cord (Brazil)	39M5233	All	Customer replaced
2.80 m jumper cable	39M5376	All	Customer replaced
2.8 m power cord (India)	39M5226	All	Customer replaced
4.3 m power cord (Japan)	39M5200	All	Customer replaced
2.8 m power cord (Korea)	39M5219	All	Customer replaced
2.5" SSD, - 300 GB, in carrier assembly	85Y5861	124, 224	Customer replaced
2.5" 10 K, 300 GB, in carrier assembly	85Y5862	124, 224	Customer replaced
2.5" 10 K, 450 GB, in carrier assembly	85Y5863	124, 224	Customer replaced
2.5" 10 K, 600 GB, in carrier assembly	85Y5864	124, 224	Customer replaced
3.5" 7.2 K nearline SAS - 2 TB in carrier assembly	85Y5869	112, 212	Customer replaced
Blank 2.5" carrier	85Y5893	124, 224	Customer replaced
Blank 3.5" carrier	85Y5894	112, 212	Customer replaced

Table 18. Replaceable units (continued)

Part	Part number	Applicable models	FRU or customer replaced
Shortwave small form-factor pluggable (SFP)	85Y5958	112, 124	Customer replaced
Longwave small form-factor pluggable (SFP)	85Y5957	112, 124	Customer replaced
Rail kit	85Y5852	All	Customer replaced
Left enclosure cap including RID tag but no black MTM label	85Y5901	All	Customer replaced
Right enclosure cap (2U12)	85Y5903	112, 212	Customer replaced
Right enclosure cap (2U24)	85Y5904	124, 224	Customer replaced

Chapter 8. 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:

- If you enabled Simple Network Management Protocol (SNMP), an SNMP trap is sent to an SNMP manager that is configured by the customer.
- If enabled, log messages can be forwarded 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.

Understanding events

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

Error data

Events are classified as either alerts or messages:

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

Viewing the event log

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

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.

Managing the event log

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

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

Describing the fields in the event log

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

Table 19 describes some of the fields that are available to assist you in diagnosing problems.

Table 19. 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

Storwize V7000 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 Storwize V7000 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 20 on page 95 describes the types of event notifications.

Table 20. 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 Storwize V7000. 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 Storwize V7000. 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>

Power-on self-test

When you turn on the Storwize V7000 system, the node canisters perform self-tests.

A series of tests is performed to check the operation of components and some of the options that have been installed when the units are first turned on. This series of tests is called the power-on self-test (POST).

If a critical failure is detected during the POST, the software is not loaded and the fault LED is illuminated. To determine if there is a POST error on a canister, go to “Procedure: Understanding the system status using the LEDs” on page 43.

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

Understanding the error codes

Error codes are generated for the Storwize V7000 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.

Event codes

The Storwize V7000 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 21 provides a list of informational event codes, the notification type, and the meaning of the event code.

Table 21. 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.
980301	I	Degraded or offline managed disk is now online.
980310	I	A degraded or offline storage pool is now online.
980320	I	Offline volume is now online.
980321	W	Volume is offline because of degraded or offline storage pool.
980330	I	All nodes can see the port.
980340	I	All ports in this host are now logged in.
980341	W	One or more ports in this host is now degraded.
980342	W	One or more ports in this host is now offline.
980343	W	All ports in this host are now offline.
980349	I	A node has been successfully added to the cluster.
980350	I	The node is now a functional member of the cluster.
980351	I	A noncritical hardware error occurred.
980352	I	Attempt to automatically recover offline node starting.
980370	I	Both nodes in the I/O group are available.
980371	I	One node in the I/O group is unavailable.
980372	W	Both nodes in the I/O group are unavailable.
980380	I	Maintenance mode was started.
980381	I	Maintenance mode has ended.
980392	I	Cluster recovery completed.
980435	W	Failed to obtain directory listing from remote node.
980440	W	Failed to transfer file from remote node.
980445	I	The migration is complete.
980446	I	The secure delete is complete.
980501	W	The virtualization amount is close to the limit that is licensed.
980502	W	The FlashCopy feature is close to the limit that is licensed.
980503	W	The Metro Mirror or Global Mirror feature is close to the limit that is licensed.
980504	I	The limit was reached for the external virtualization feature.

Table 21. Informational event codes (continued)

Event code	Notification type	Description
981002	I	Fibre-channel discovery occurred; configuration changes are required.
981003	I	Fibre-channel discovery occurred; configuration changes are complete.
981004	I	Fibre-channel discovery occurred; no configuration changes were detected.
981007	W	The managed disk is not on the preferred path.
981009	W	The initialization for the managed disk failed.
981014	W	The LUN discovery has failed. The cluster has a connection to a device through this node but this node cannot discover the unmanaged or managed disk that is associated with this LUN.
981015	W	The LUN capacity equals or exceeds the maximum. Only part of the disk can be accessed.
981020	W	The managed disk error count warning threshold has been met.
981022	I	Managed disk view smoothing start
981025	I	Drive firmware download started
981026	I	Drive FPGA download started
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.
984501	I	The firmware level of an enclosure component is being updated.
984502	I	The firmware level updated has completed.

Table 21. Informational event codes (continued)

Event code	Notification type	Description
984503	I	The battery conditioning completed.
984504	I	The battery conditioning started.
984505	I	The statesave information for the enclosure was collected.
984506	I	The debug from an IERR was extracted to disk.
984507	I	An attempt was made to power on the slots.
984508	I	All the expanders on the strand were reset.
984509	I	The component firmware update paused to allow the battery charging to finish.
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 Storwize V7000 software to the required version or delete the volume.
986011	I	The thin-provisioned volume copy import is successful.
986020	W	A thin-provisioned volume copy space warning has occurred.
986030	I	A thin-provisioned volume copy repair has started.
986031	I	A thin-provisioned volume copy repair is successful.
986032	I	A thin-provisioned volume copy validation is started.
986033	I	A thin-provisioned volume copy validation is successful.
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 Storwize V7000 cluster. If the problem persists, contact your IBM service representative for assistance.

Table 21. Informational event codes (continued)

Event code	Notification type	Description
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.

Error codes

Error codes describe a service procedure that must be followed. Each event ID that requires service has an associated error code.

Table 22 lists the event IDs and corresponding error codes.

Table 22. 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

Table 22. Event IDs and codes (continued)

Event ID	Notification type	Condition	Error code
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
010010	E	The managed disk is excluded because of excessive errors.	1310
010011	E	The remote port is excluded for a managed disk and node.	1220
010012	E	The local port is excluded.	1210
010013	E	The login is excluded.	1230
010014	E	The local port is excluded.	1211
010015	None	A timeout has occurred because of a nonresponsive device.	1340
010016	None	A timeout has occurred because of a lost command.	1340
010017	E	A timeout has occurred as a result of excessive processing time.	1340
010018	E	An error recovery procedure has occurred.	1370
010019	E	A managed disk 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

Table 22. Event IDs and codes (continued)

Event ID	Notification type	Condition	Error code
010031	E	One or more MDisks on a controller are degraded.	1623
010032	W	The controller configuration limits failover.	1625
010033	E	The controller configuration uses the RDAC mode; this is not supported.	1624
010034	E	Persistent unsupported controller configuration.	1695
010035	W	A quorum disk is configured on the controller that is set to not allow quorum.	1570
010040	E	The controller system device is only connected to the node through a single initiator port.	1627
010041	E	The controller system device is only connected to the node through a single target port.	1627
010042	E	The controller system device is only connected to the cluster nodes through a single target port.	1627
010043	E	The controller system device is only connected to the cluster nodes through half of the expected target ports.	1627
010044	E	The controller system device has disconnected all target ports to the cluster nodes.	1627
010050	W	A solid-state drive (SSD) failed. A rebuild is required.	1201
010051	E	A solid-state drive (SSD) is missing.	1202
010052	E	A solid-state drive (SSD) is offline as a result of a drive hardware error.	1205
010053	E	A solid-state drive (SSD) is reporting a predictive failure analysis (PFA).	1215
010054	E	A solid-state drive (SSD) is reporting too many errors.	1215
010055	W	An unrecognized SAS device.	1665
010056	E	SAS error counts exceeded the warning thresholds.	1216
010057	E	SAS errors exceeded critical thresholds.	1216
010058	E	A solid-state drive (SSD) initialization failed.	1661
010059	E	A solid-state drive (SSD) is offline due to excessive errors.	1311
010060	E	A solid-state drive (SSD) exceeded the warning temperature threshold.	1217
010061	E	A solid-state drive (SSD) exceeded the offline temperature threshold.	1218
010062	E	A drive exceeded the warning temperature threshold.	1217
010063	W	Drive medium error.	1321
010066	W	Controller indicates that it does not support descriptor sense for LUNs that are greater than 2 TBs.	1625
010067	W	Too many enclosures were presented to a cluster.	1200
010068	E	The solid-state drive (SSD) format was corrupted.	1204
010069	E	The block size for the solid-state drive (SSD) was incorrect.	1204
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
010072	E	The drive is offline as a result of a drive hardware error.	1680
010073	E	The drive is reporting predictive failure analysis (PFA) errors.	1680
010080	E	The drive is reporting too many errors.	1680
010081	E	The drive format is corrupted.	1206

Table 22. Event IDs and codes (continued)

Event ID	Notification type	Condition	Error code
010082	E	The block size for the drive was incorrect.	1206
010083	E	The drive is offline due to excessive errors.	1680
010084	E	The error counts for the SAS drive exceeded the warning thresholds.	1285
010085	W	The SAS device was not recognized.	1666
010086	W	The SAS enclosure was not recognized.	1666
010087	W	The SAS device was not able to be identified.	1666
010088	E	There were excessive medium errors on the drive.	1680
010089	E	There were excessive overall timeout errors on the drive.	1680
010090	E	There were excessive times when the drive stopped.	1680
010091	E	A drive failed validation testing.	1680
010092	E	There were excessive medium errors on the solid-state drive (SSD).	1215
010093	E	There were excessive overall timeout errors on the solid-state drive (SSD).	1204
010094	E	Login excluded.	1231
010096	E	Drive initialization failed.	1680
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
02005	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
045001	E	One or more power supply unit fans have failed.	1124
045002	E	A fan is operating outside the expected range.	1126
045003	E	There was a fan status communications failure.	1126
045004	E	The power supply unit is not installed.	1128
045005	W	The power supply unit has indicated an input power failure.	1138
045006	E	The power supply unit has indicated a dc failure.	1126
045007	E	The power supply unit has failed.	1124
045008	E	There is no communication with the power supply unit.	1148
045009	E	The model type for this enclosure is not valid.	1124
045010	E	The power supply unit type is unknown to this product.	1124
045011	E	The power supply unit serial number is not valid.	1124
045012	W	The canister temperature is at the warning level.	1098
045013	W	The canister temperature is at the critical level.	1095

Table 22. Event IDs and codes (continued)

Event ID	Notification type	Condition	Error code
045014	E	The SAS cable was excluded because of a missing device.	1260
045015	E	A SAS cable was excluded because too many change events were caused.	1260
045016	E	A SAS cable was excluded.	1255
045017	E	A SAS cable is operating at a reduced speed.	1260
045018	E	A SAS cable was excluded because frames were dropped.	1260
045019	E	A SAS cable was excluded because the enclosure discovery timed out.	1260
045020	W	A SAS cable is not present.	1265
045021	E	A canister was removed from the system.	1036
045022	E	A canister has been in a degraded state for too long and cannot be recovered.	1034
045023	E	A canister is encountering communication problems.	1038
045024	E	The canister VPD is not valid.	1032
045025	E	The canister has experienced too many resets.	1032
045026	E	The drive slot is causing the network to be unstable.	1686
045027	E	The drive slot is not running at 6 Gbps	1686
045028	E	The drive slot is dropping frames.	1686
045029	E	The drive is visible through only one SAS port.	1686
045031	E	The drive power control is not functional.	1008
045033	E	The drive slot contains a device that is not responding to queries.	1685
045034	E	The managed enclosure is not visible from any node canisters.	1042
045035	E	The electronics in the enclosure has failed.	1694
045036	E	The electronics in the enclosure has experienced a critical failure.	1008
045037	E	The SAS network has too many errors.	1048
045038	E	The SAS network has too many errors.	1048
045040	E	The firmware update for the enclosure component has failed.	3015
045041	E	More than one initiator ports were detected on the same strand.	1005
045042	W	The order of the enclosures is different on each strand.	1005
045044	W	Multiple canisters are connected to a single canister port.	1005
045045	W	Canister 1 is connected to canister 2.	1005
045046	W	An enclosure is connected to more than one I/O group.	1005
045047	W	A managed enclosure is connected to the wrong I/O group.	1005
045048	W	An enclosure is connected to more than one chain.	1005
045049	W	Too many canisters are connected to a strand.	1005
045050	W	The canister is connected to the wrong port.	1005
045051	E	A SAS cable is excluded because of single port active drives.	1260
045052	W	More than one canister was detected at the same hop count.	1005
045053	E	The node location is not able to be detected.	1031
045054	E	An enclosure display cannot be updated.	1694
045055	E	There is an enclosure battery fault.	1118

Table 22. Event IDs and codes (continued)

Event ID	Notification type	Condition	Error code
045056	E	An enclosure battery is missing.	1112
045057	E	An enclosure battery is nearing end of life.	1114
045058	E	An enclosure battery is at end of life.	1113
045062	W	An enclosure battery conditioning is not required but possible.	1131
045063	E	There was an enclosure battery communications error.	1116
045064	W	A SAS port is active, but no enclosures can be detected.	1005
045065	E	There is a connectivity problem between a canister and an enclosure.	1036
045066	E	The FRU identity of the enclosure is not valid.	1008
045067	W	A new enclosure FRU was detected and needs to be configured.	1041
045068	E	The internal device on a node canister was excluded because of too many change events.	1034
045069	E	The internal connector on the node canister was excluded as the cause of single ported drives.	1034
045070	W	The canister temperature sensor cannot be read.	1034
045071	W	The enclosure contains both a node canister and an expansion canister.	1037
045072	E	The discovery failed to complete.	1048
045073	E	The VPD for the enclosure cannot be read.	1048
045080	E	There are too many self-initiated resets on the enclosure.	1048
045082	E	The slots are powered off..	1048
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
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

Table 22. Event IDs and codes (continued)

Event ID	Notification type	Condition	Error code
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
070690	W	Node held in service state.	1189
071500	W	Incorrect enclosure.	1021
071501	E	Incorrect canister position.	1192
071502	E	No enclosure identity; cannot get status from partner.	1192
071503	E	Incorrect enclosure type.	1192
071504	E	No enclosure identity and partner does match.	1192
071505	E	No enclosure identity and partner does not match.	1192
071506	E	No enclosure identity and no state on partner.	1192
071507	E	No enclosure identity and no node state.	1192
071508	W	Cluster identity different on enclosure and node.	1023
071509	E	Cannot read enclosure identity.	1036
071510	E	Detected memory size does not match the expected memory size.	1032
071523	W	The internal disk file system of the node is damaged.	1187
071524	E	Unable to update the BIOS settings.	1034
071525	E	Unable to update the service processor firmware of the system board.	1032
071535	E	The internal PCIe switch of the node canister failed.	1192
071550	E	Cannot form cluster due to lack of cluster resources.	1192
071556	E	Duplicate WWNN detected on the SAN.	1133
071562	E	The node hardware does not meet the minimum requirements.	1036
071564	W	Too many software failures.	1188
071565	E	The internal drive of the node is failing.	1032

Table 22. Event IDs and codes (continued)

Event ID	Notification type	Condition	Error code
071573	E	The node software is inconsistent.	1187
071574	E	The cluster data cannot be read.	1032
071578	E	The cluster data was not saved when power was lost.	1194
071671	E	There is not enough battery available to start the node. Two batteries are charging.	1192
071672	E	There is not enough battery available to start the node. One battery is charging.	1192
071673	E	There is not enough battery available to start the node. No batteries are charging.	1192
071690	W	Node held in service state.	1189
072900	E	There was a PCIe link failure between canisters.	1006
072901	E	The PCIe link is degraded between canisters.	1052
072911	E	The PCIe link for the CPU is degraded.	1034
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 Storwize V7000 port.	1800
073007	W	There are fewer fibre-channel ports operational than are configured.	1061
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
075033	None	There was a service controller read failure.	None
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
076511	E	A high-speed SAS controller is missing.	1032
076512	E	Degraded PCIe lanes on a high-speed SAS adapter.	1032
076513	E	A PCI bus error occurred on a high-speed SAS adapter.	1032

Table 22. Event IDs and codes (continued)

Event ID	Notification type	Condition	Error code
076514	E	A high-speed SAS adapter requires a PCI bus reset.	1034
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
081002	E	An Ethernet port failure has occurred.	1401
082001	E	A server error has occurred.	2100
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

Node error code overview

Node error codes describe failure that relate to a specific node canister.

Because node errors are specific to a node, for example, memory has failed, the errors are only reported on that node. However, some of the conditions that the node detects relate to the shared components of the enclosure. In these cases both node canisters in the enclosure report the error.

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

Critical errors

A critical error means that the node is not able to participate in a 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 it is possible to communicate with the canister with a node error, an alert that describes the error is logged in the event log. If the system cannot communicate with the node canister, a Node missing alert is reported. If a node has a critical node error, it is in service state, and the fault LED on the node is on. The exception is when the node cannot connect to enough resources to form a cluster. It shows a critical node error but is in the starting state. The range of errors that are reserved for critical errors are 500 - 699.

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 node error is shown to make it clear which of the node canisters the alert refers to. The range of errors that are reserved for noncritical errors are 800 - 899.

Cluster code overview

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

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.

Error code range

This topic shows the number range for each message classification.

Table 23 lists the number range for each message classification.

Table 23. Message classification number range

Message classification	Range	
Node errors	Critical node errors	500-699
	Noncritical node errors	800-899
Cluster recovery	920, 990	

Node errors

500 Incorrect enclosure

Explanation: The node canister has saved cluster information, which indicates that the canister is now located in a different enclosure chassis from where it was previously used. Using the node canister in this state might corrupt the data held on the enclosure drives.

User response: Follow troubleshooting procedures to relocate the nodes to the correct location.

1. Review the saved location information of the node canister and the status of the other node canister in the enclosure (the partner canister). Determine whether the enclosure is part of an active system with volumes that contain required data.
2. If you have unintentionally moved the canister into this enclosure, move the canister back to its original location, and put the original canister back in this enclosure. Follow the hardware remove and replace canister procedures.
3. If you have intentionally moved the node canister into this enclosure, check whether it is safe to continue or whether a loss of data can result. Do not continue if any of the following conditions apply, contact IBM technical support instead:
 - a. You require the volume data on the system from which the node canister was removed, and that system is not running with two online nodes.
 - b. You require the volume data on this system, and the partner node is not online.

4. If you have determined that the node canister can be used in this location, follow procedures to remove cluster data from the node canister.

Possible Cause-FRUs or other:

- None

501 Incorrect slot

Explanation: The node canister has saved cluster information, which indicates that the canister is now located in the expected enclosure chassis, but in a different slot from where it was previously used. Using the node canister in this state might mean that hosts are not able to connect correctly.

User response: Follow troubleshooting procedures to relocate the nodes to the correct location.

1. Review the saved location information of the node canister and the status of the other node canister in the enclosure (the partner canister). If the node canister has been inadvertently swapped, the other node canister will have the same error.
2. If the canisters have been swapped, use the hardware remove and replace canister procedure to swap the canisters. The system should start.
3. If the partner canister is in candidate state, use the hardware remove and replace canister procedure to swap the canisters. The system should start.
4. If the partner canister is in active state, it is running the cluster on this enclosure and has replaced the original use of this canister. You must follow the

procedure to remove cluster data from this node canister. The node canister will then become active in the cluster in its current slot.

- If the partner canister is in service state, review its node error to determine the correct action. Generally, you will fix the errors reported on the partner node in priority order, and review the situation again after each change. If you have to replace the partner canister with a new one you should move this canister back to the correct location at the same time.

Possible Cause-FRUs or other:

- None

502 No enclosure identity exists and a status from the partner node could not be obtained.

Explanation: The enclosure has been replaced and communication with the other node canister (partner node) in the enclosure is not possible. The partner node could be missing, powered off or unable to boot, or an internode communication failure may exist.

User response: Follow troubleshooting procedures to fix the hardware:

- Follow the service procedures to get the partner node started. An error will still exist because the enclosure has no identity. If the error has changed, follow the service procedure for that error.
- If the partner has started and is showing a location error (probably this one), then the PCI link is probably broken. Since the enclosure chassis was recently replaced, this is likely the problem. Obtain a replacement enclosure chassis, and restart the hardware remove and replace control enclosure chassis procedure.
- If this action does not resolve the issue, contact IBM technical support. They will work with you to ensure that the cluster state data is not lost while resolving the problem.

Possible Cause-FRUs or other:

- Enclosure chassis (100%)

503 Incorrect enclosure type

Explanation: The node canister has been moved to an expansion enclosure. A node canister will not operate in this environment.

User response: Follow troubleshooting procedures to relocate the nodes to the correct location.

- Review the saved location information of the node canister to determine which control enclosure the node canister should be in.
- Use the hardware remove and replace canister procedures to move the node canister to the correct

location and move the expansion canister, that is probably in that location, to here. If there is a node canister that is in active state where this node canister should be, do not replace that node canister with this one.

504 No enclosure identity and partner node matches.

Explanation: The enclosure vital product data indicates that the enclosure chassis has been replaced. This node canister and the other node canister in the enclosure were previously operating in the same enclosure chassis.

User response: Follow troubleshooting procedures to configure the enclosure.

- This is an expected situation during the hardware remove and replace procedure for a control enclosure chassis. Continue following the remove and replace procedure and configure the new enclosure.

Possible Cause-FRUs or other:

- None

505 No enclosure identity and partner has cluster data that does not match.

Explanation: The enclosure vital product data indicates that the enclosure chassis has been replaced. This node canister and the other node canister in the enclosure do not come from the same original enclosure.

User response: Follow troubleshooting procedures to relocate nodes to the correct location.

- Review the saved location information of the node canister and the saved location information of the other node canister in the enclosure (the partner canister). The correct node canister is the one that comes from the enclosure chassis that is being replaced. The drives now in this enclosure should also come from that enclosure.
- Decide what to do with the node canister that did not come from the enclosure that is being replaced.
 - If the other node canister from the enclosure being replaced is available, use the hardware remove and replace canister procedures to remove the incorrect canister and replace it with the second node canister from the enclosure being replaced. The two node canister should show node error 504 and the actions for that error should be followed.
 - If the other node canister from the enclosure being replaced is not available, check the enclosure of the node canister that did not come from the replaced enclosure. Do not use this canister in this enclosure if you require the volume data on the system from which the node

canister was removed, and that system is not running with two online nodes. You should return the canister to its original enclosure and use a different canister in this enclosure.

- c. When you have checked it is not required elsewhere, follow the procedure to remove cluster data from the node canister that did not come from the enclosure that is being replaced. Restart both nodes. Expect node error 506 to now be reported and follow the service procedures for that error.

Possible Cause-FRUs or other:

- None

506 No enclosure identity and no node state on partner

Explanation: The enclosure vital product data indicates that the enclosure chassis has been replaced. There is no cluster state information on the other node canister in the enclosure (the partner canister), so both node canisters from the original enclosure have not been moved to this one.

User response: Follow troubleshooting procedures to relocate nodes to the correct location:

1. Review the saved location information of the node canister and why the second node canister from the original enclosure was not moved into this enclosure.
2. If you are sure that this node canister came from the enclosure that is being replaced, and the original partner canister is available, use the hardware remove and replace node canister procedure to install the second node canister in this enclosure. Restart the node canister. The two node canister should show node error 504 and the actions for that error should be followed.
3. If you are sure this node canister came from the enclosure that is being replaced, and that the original partner canister has failed, continue following the remove and replace procedure for an enclosure chassis and configure the new enclosure.

Possible Cause-FRUs or other:

- None

507 No enclosure identity and no node state

Explanation: The node canister has been placed in a replacement enclosure chassis. The node canister is also a replacement, or has had all cluster state removed from it.

User response: Follow troubleshooting procedures to relocate the nodes to the correct location.

1. Check the status of the other node in the enclosure. It should show node error 506. Unless it also shows error 507, use these procedures to resolve the errors on the other node.
2. If the other node in the enclosure is also reporting 507, the enclosure and both node canisters have no state information. You should contact IBM technical support. They will assist you in setting the enclosure vital product data and running cluster recovery.

Possible Cause-FRUs or other:

- None

508 Cluster identifier is different between enclosure and node

Explanation: The node canister location information shows it is in the correct enclosure, however the enclosure has had a new cluster created on it since the node was last shut down. Therefore, the cluster state data stored on the node is not valid.

User response: Follow troubleshooting procedures to relocate the nodes to the correct location.

1. Check whether a new cluster has been created on this enclosure while this canister was not operating or whether the node canister was recently installed in the enclosure.
2. If this node canister is the one to be used in this enclosure, you should follow the procedure to remove cluster data from the node canister. It will then join the cluster.
3. If this is not the node canister that you intended to use, follow the hardware remove and replace canister procedure to replace the node canister with the intended one.

Note: If both nodes are reporting this node error, the cause may be a damaged enclosure.

Possible Cause-FRUs or other:

- Service procedure error (90%)
- Enclosure chassis (10%)

509 The enclosure identity cannot be read.

Explanation: The canister was unable to read vital product data (VPD) from the enclosure. The canister requires this data to be able to initialize correctly.

User response: Follow troubleshooting procedures to fix the hardware:

1. Check errors reported on the other node canister in this enclosure (the partner canister).
2. If it is reporting the same error follow the hardware remove and replace procedure to replace the enclosure chassis.

- If the partner canister is not reporting this error, follow the hardware remove and replace procedure to replace this canister.

Note: If a newly installed system has this error on both node canister, the data that needs to be written to the enclosure will not be available on the canisters, you should contact IBM technical support for the WWNNs to use.

Possible Cause-FRUs or other:

- Node canister (50%)
- Enclosure chassis (50%)

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

Explanation: The amount of memory detected in the node canister is less than the expected memory. The error code date shows the detected memory, in MB, followed by the expected memory, in MB.

User response: Follow troubleshooting procedures to fix the hardware:

- Use the hardware remove and replace node canister procedure to install a new node canister.

Possible Cause-FRUs or other:

- Node canister (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.

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

Possible Cause-FRUs or other:

- Node canister (100%)

525 Unable to update system board service processor firmware.

Explanation: The node startup procedures have been unable to update the firmware configuration of the node canister.

User response: Follow troubleshooting procedures to fix the hardware:

- Follow the hardware remove and replace procedures for the node canister.

Possible Cause-FRUs or other:

- Node canister (100%)

535 Canister internal PCIe switch failed

Explanation: The PCI Express switch has failed or cannot be detected. In this situation, the only connectivity to the node canister is through the Ethernet ports.

User response: Follow troubleshooting procedures to fix the hardware:

- Follow the procedures to reseat the node canister.
- If reseating the canister does not resolve the situation, follow the hardware remove and replace node canister procedures to replace the canister.

Possible Cause-FRUs or other:

- Node canister (100%)

550 A cluster cannot be formed because of a lack of cluster resources.

Explanation: The node cannot become active in a cluster because it is unable to connect to enough cluster resources. The cluster resources are the node canisters in the system and the active quorum disk or drive. The node needs to be able to connect to a majority of the resources before that group will form an online cluster. This prevents the cluster splitting into two or more active parts, with both parts independently performing I/O.

The error data lists the missing resources. This will include a list of node canisters and optionally a drive that is operating as the quorum drive or a LUN on an external storage system that is operating as the quorum disk.

If a drive in one of the 2076 enclosures is the missing quorum disk, it is listed as enclosure:slot[part identification] where enclosure:slot is the location of the drive when the node shut down, enclosure is the seven digit product serial number of the enclosure, slot is a number between 1 and 24. The part identification is the 22 character string starting "11S" found on a label on a drive. The part identification cannot be seen until the drive is removed from the enclosure.

If a LUN on an external storage system is the missing quorum disk, it is listed it is listed as WWWWWWWWWWWWWWWWW/LL, where WWWWWWWWWWWWWWWWW is a worldwide port name (WWPN) on the storage system that contains the missing quorum disk and LL is the Logical Unit Number (LUN).

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

- Check the status of other node canisters in the system, resolve any faults on them.

2. Check all enclosures in the system are powered on and that the SAS cabling between the enclosures has not been disturbed. If any wiring changes have been made check all cables are securely connected and that the cabling rules have been followed.
3. If a quorum drive in a system enclosure is shown as missing, find the drive and check that it is working. The drive may have been moved from the location shown, in that case find the drive and ensure it is installed and working. If the drive is not located in the control enclosure, try moving it to the control enclosure, a problem in SAS connectivity may be the issue.

Note: If you are able to reestablish the systems operation you will be able to use the extra diagnostics the system provides to diagnose problems on SAS cables and expansion enclosures.

4. If a quorum disk on an external storage system is shown as missing, find the storage control and confirm that the LUN is available, check the fibre channel connections between the storage controller and the 2076 are working and that any changes made to the SAN configuration and zoning have not effected the connectivity. Check the status of the fibre channel ports on the node and resolve any issues.
5. If all nodes have either node error 578 or 550, attempt to reestablish a cluster by following the service procedures for the nodes showing node error 578. If this is not successful, follow the cluster recovery procedures.

556 **A duplicate WWNN has been detected.**

Explanation: The node canister 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 Storwize V7000, the first 11 digits are always 50050768020. The last 5 digits of the WWNN are given in the additional data of the error. The fibre-channel ports of the node canister are disabled to prevent disruption of the fibre-channel network. One or both node canisters with the same WWNN can show the error. Because of the way WWNNs are allocated, a device with a duplicate WWNN is normally another Storwize V7000 node canister.

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

1. Find the Storwize V7000 node canister with the same WWNN as the node canister reporting the error. The WWNN for a Storwize V7000 node canister can be found from the node Vital Product Data (VPD) or from the node canister details shown by the service assistant. The node with the duplicate WWNN need not be part of the same cluster as the node reporting the error; it could be remote from the node reporting the error on a part of the fabric

connected through an inter-switch link. The two node canisters within a control enclosure must have different WWNNs. The WWNN of the node canister is stored within the enclosure chassis, so the duplication is most likely caused by the replacement of a control enclosure chassis.

2. If a Storwize V7000 node canister with a duplicate WWNN is found, determine whether it, or the node reporting the error, has the incorrect WWNN. Generally, it is the node canister that has had its enclosure chassis 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 enclosure chassis has been replaced as part of a service action, the WWNN for the node canister should have been written down. If the correct WWNN cannot be determined contact your support center for assistance.
4. Use the service assistant to modify the incorrect WWNN. If it is the node showing the error that should be modified, this can safely be done immediately. If it is an active node that should be modified, use caution because the node will restart when the WWNN is changed. If this node is the only operational node in an enclosure, access to the volumes that it is managing will be lost. You should ensure that the host systems are in the correct state before you change the WWNN.
5. If the node showing the error had the correct WWNN, it can be restarted, using the service assistant, after the node with the duplicate WWNN is updated.
6. If you are unable to find a Storwize V7000 node canister with the same WWNN as the node canister showing the error, use the SAN monitoring tools to determine whether there is another device on the SAN with the same WWNN. This device should not be using a WWNN assigned to a Storwize V7000, so you should follow the service procedures for the device to change its WWNN. Once the duplicate has been removed, restart the node canister.

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. It is not possible to service parts within the node canister. Reseat the existing node canister to see

whether the problem fixes. If it does not, use the hardware node canister remove and replace procedures to change the node canister.

Possible Cause-FRUs or other:

- Node canister (100%)

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. The drive of the node canister cannot be replaced individually. Follow the hardware remove and replace instruction to change the node canister.

Possible Cause-FRUs or other:

- Node canister (100%)

573 The node software is inconsistent.

Explanation: Parts of the node software package are receiving unexpected results; there may be an inconsistent set of subpackages installed, or one subpackage may be damaged.

User response: Follow troubleshooting procedures to reload the software.

1. Follow the procedure to run a node rescue.
2. If the error occurs again, contact IBM technical support.

Possible Cause-FRUs or other:

- None

574 The node 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: Follow troubleshooting procedures to fix the hardware:

1. The drive of the node canister cannot be replaced individually. Follow the hardware remove and replace instructions to change the node canister.

Possible Cause-FRUs or other:

- None

578 The state data was not saved following a power loss.

Explanation: On startup, the node was unable to read its state data. When this happens, it expects to be automatically added back into a 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

671 The available battery charge is not enough to allow the node canister to start. Two batteries are charging.

Explanation: The battery charge within the enclosure is not sufficient for the node to safely become active in a cluster. The node will not start until sufficient charge exists to store the state and configuration data held in the node canister memory if power were to fail. Two batteries are within the enclosure, one in each of the power supplies. Neither of the batteries indicate an error—both are charging.

The node will start automatically when sufficient charge is available. The batteries do not have to be fully charged before the nodes can become active.

Both nodes within the enclosure share the battery charge, so both node canisters report this error.

The service assistant shows the estimated start time in the node canister hardware details.

Possible Cause-FRUs or other:

- None

User response: Wait for the node to automatically fix the error when sufficient charge becomes available.

Possible Cause-FRUs or other:

- None

672 The available battery charge is not enough to allow the node canister to start. One battery is charging.

Explanation: The battery charge within the enclosure is not sufficient for the node to safely become active in a cluster. The node will not start until sufficient charge exists to store the state and configuration data held in

the node canister memory if power were to fail. Two batteries are within the enclosure, one in each of the power supplies. Only one of the batteries is charging, so the time to reach sufficient charge will be extended.

The node will start automatically when sufficient charge is available. The batteries do not have to be fully charged before the nodes can become active.

Both nodes within the enclosure share the battery charge, so both node canisters report this error.

The service assistant shows the estimated start time, and the battery status, in the node canister hardware details.

Possible Cause-FRUs or other:

- None

User response:

1. Wait for the node to automatically fix the error when sufficient charge becomes available.
2. If possible, determine why one battery is not charging. Use the battery status shown in the node canister hardware details and the indicator LEDs on the PSUs in the enclosure to diagnose the problem. If the issue cannot be resolved, wait until the cluster is operational and use the troubleshooting options in the management GUI to assist in resolving the issue.

Possible Cause-FRUs or other:

- Battery (33%)
- Control power supply (33%)
- Power cord (33%)

673 The available battery charge is not enough to allow the node canister to start. No batteries are charging.

Explanation: A node cannot be in active state if it does not have sufficient battery power to store configuration and cache data from memory to internal disk after a power failure. The system has determined that both batteries have failed or are missing. The problem with the batteries must be resolved to allow the system to start.

User response: Follow troubleshooting procedures to fix hardware:

1. Resolve problems in both batteries by following the procedure to determine status using the LEDs.
2. If the LEDs do not show a fault on the power supplies or batteries, power off both power supplies in the enclosure and remove the power cords. Wait 20 seconds, then replace the power cords and restore power to both power supplies. If both node canisters continue to report this error replace the enclosure chassis.

Possible Cause-FRUs or other:

- Battery (33%)
- Power supply (33%)
- Power cord (33%)
- Enclosure chassis (1%)

690 The node is held in the service state.

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

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

1. Use the service assistant action to release the service state.

Possible Cause-FRUs or other:

- None

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.

1. If possible, use the troubleshooting fix procedures in the management GUI to correct the associated cluster error.
2. If no cluster is operational, check whether all fibre-channel cables are fully inserted into the node canister, and that the SAN switch is powered on and not showing errors. This node error will not prevent the node from becoming active in a cluster. If the specified checks do not resolve the problem, continue with cluster creation and then use the management GUI to resolve the problem.

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.

1. If possible, use the troubleshooting fix procedures in the management GUI to correct the associated cluster error. If this error has occurred in the configuration node, you may not be able to run the management GUI.
2. If you are not able to run the management GUI, check whether all Ethernet cables are fully inserted into the node canister, and that the Ethernet switch is powered on and not showing errors. This node error will not prevent the node from becoming active in a cluster. If the specified checks do not resolve the problem, and you have not created the cluster yet, continue with cluster creation on the other node canister in the enclosure and then use the management GUI to resolve the problem.
3. Use a USB key to get the status of the node and check whether it is the configuration node. If the node showing this error is the configuration node, use a USB key or the service assistant to get the status of the other node in the enclosure. If the other node is active, run the service assistant on the other node, but make this node the active node. Check whether the active node is reporting node error 805. Use the service assistant to hold the node in service state, which will cause the other node in the enclosure to become the configuration node. Once you have set the node canister into service state, you should immediately release the service state so that the node becomes active in the cluster. Now use the troubleshooting fix procedures in the management GUI to correct the associated cluster error.

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.

835 Unable to communicate with partner node over PCIe link.

Explanation: The partner node is not in a state that allows communication because the node is powered off, a boot failure has occurred, or the PCIe link is broken.

User response: Follow troubleshooting procedures to fix the hardware:

1. Determine status of other node.

- Restart or replace the node if it has failed (should be node error on partner).

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:

- View hardware WWNN information

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.

User response: None.

990 **Cluster recovery has failed.**

Explanation: Cluster recovery has failed.

User response: Contact IBM technical support.

- Reconfigure your SAN zoning so that only the Storwize V7000 nodes, the host system ports, and the storage system ports to which you want to connect are visible to the node that is reporting the error. Ensure that no more than 1024 exist.

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.

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).

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 Storwize V7000 information at the following website:

publib.boulder.ibm.com/infocenter/storwize/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。

taieimi

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

- 2.5" drive assembly
 - replacing 82
- 3.5" drive assembly
 - replacing 80

A

- about this document
 - sending comments xii
- accessibility
 - keyboard 117
 - repeat rate of up and down buttons 117
 - shortcut keys 117
- accessing
 - canisters
 - Ethernet cable 51
 - cluster CLI 28
 - management GUI 26
 - service assistant 27, 52
 - service CLI 28
- actions
 - reset service IP address 30
 - reset superuser password 30
- alerts
 - best practices 19
 - apply software command 32

B

- backing up
 - best practices 18
- backing up cluster configuration files
 - creating 59
- backup configuration files
 - deleting
 - using the CLI 63
 - restoring 61
- battery
 - maintenance 21, 22
 - removing 76
- best practices
 - alerts 19
 - backing up data 18
 - drive characteristics 19
 - inventory reporting 18
 - IP address 17
 - notifications 18
 - passwords 17
 - power management 17
 - RAID 19
 - record
 - location information 19
 - subscribe
 - notifications 20
 - troubleshooting 17
 - warranty agreement
 - maintenance agreement 20

- blank carrier
 - replacing 80, 82
- browsers
 - supported 39

C

- cable retention bracket
 - releasing 78
- Canadian electronic emission notice 122
- canister
 - expansion 67
 - identification 41
 - node 65
 - replacing 65, 67
- changing
 - service IP address 49
- CLI
 - cluster commands 28
 - service commands 28
- cluster
 - deleting 48
 - error codes 108
 - failure to create 37
 - initializing
 - with service assistant 50
 - with USB key 49
 - restore 55, 58
 - T3 recovery 55
- cluster CLI
 - accessing 28
 - when to use 28
- cluster commands
 - CLI 28
- cluster data
 - removing 47
- clusters
 - backing up configuration file using the CLI 59
 - error codes 108
 - recovery codes 108
 - restoring backup configuration files 61
- clusterT3 recovery
 - restore 55
- codes
 - event 95
 - informational event 96
 - node error
 - critical 107
 - noncritical 107
- commands
 - apply software 32
 - create cluster 32
 - query status 33
 - reset service assistant password 31
 - snap 31
 - svconfig backup 59
 - svconfig restore 61
- comments, sending xii

- components
 - enclosure 1
 - illustration 5
 - end cap
 - indicators 3
 - hardware 1
- contact information
 - European 124
 - Taiwan 125
- control enclosure
 - power supply unit 5
 - replacing 84
- create cluster command 32
- critical
 - node errors 107
- Customer replaced 90

D

- deleting
 - backup configuration files
 - using the CLI 63
 - cluster 48
- detection error
 - expansion location 40
- determining
 - SAN problem 52
- Deutschsprachiger EU Hinweis 122
- documentation
 - improvement xii
- drive characteristics
 - best practices 19
- drives
 - 2.5-inch drives 2
 - 3.5-inch drives 2
 - LED indicator 2

E

- electronic emission notices
 - Avis de conformité à la réglementation d'Industrie Canada 122
 - Deutschsprachiger EU Hinweis 122
 - European Union (EU) 122
 - Federal Communications Commission (FCC) 121
 - French Canadian 122
 - Germany 122
 - Industry Canada 122
 - International Electrotechnical Commission (IEC) 124
 - Japanese Voluntary Control Council for Interference (VCCI) 123
 - Korean 124
 - New Zealand 122
 - People's Republic of China 123
 - Taiwan 124
 - United Kingdom 124

- EMC statement, People's Republic of China 123
- enclosure
 - components 5
 - identification 41
- enclosure end cap
 - replacing 83
- end cap
 - indicators 3
- environmental notices ix
- error
 - expansion enclosure 40
 - node canister 39
 - not detected 40
 - SAS cabling 39
- error codes 99
 - understanding 95
- error events 93
- error ID 99
- errors
 - logs
 - describing the fields 94
 - error events 93
 - managing 93
 - understanding 93
 - viewing 93
 - node 107
- Ethernet
 - accessing
 - canister 51
 - ports 10
 - status 47
- European contact information 124
- European Union (EU), EMC Directive conformance statement 122
- event
 - codes 95
 - informational 96
- event notification 94
- events
 - reporting 93
- expansion canister
 - LEDs 14
- expansion enclosure
 - detection error 40
 - power supply unit 6
 - replacing 88

F

- failure
 - cluster create 37
- FCC (Federal Communications Commission) electronic emission notice 121
- Federal Communications Commission (FCC) electronic emission notice 121
- Fibre Channel
 - link failures 53
 - port numbers 9
 - SFP transceiver 53, 69
- Fibre Channel ports
 - rear-panel indicators 8
- field replaceable units (FRUs) 90
- fields
 - event log 94

- finding
 - Ethernet
 - status 47
- fix
 - errors 56
- fixing
 - node errors 48
- French Canadian electronic emission notice 122

G

- Germany electronic emission compliance statement 122
- GUI connectivity issues
 - troubleshooting procedure 36, 38

H

- hardware components 1

I

- identifying
 - canister 41
 - enclosure 41
 - status 42
- IEC (International Electrotechnical Commission) electronic emission notice 124
- indicators
 - end cap 3
- information
 - center xi
- informational event 93
- initialization tool
 - interface 28
 - using 29
- initializing
 - clusters
 - with service assistant 50
 - with USB key 49
- International Electrotechnical Commission (IEC) electronic emission notice 124
- inventory information 94
- inventory reporting
 - best practices 18
- IP address
 - best practices 17

J

- Japanese electronic emission notice 123

K

- keyboard 117
- Korean electronic emission statement 124

L

- LEDs
 - expansion canister 14

- LEDs (*continued*)
 - Fibre Channel ports 8
 - node canister 12
 - rear-panel indicators 8
 - system status 43
- legal notices 119
- link failures
 - Fibre Channel 53
- location information
 - record
 - best practices 19
- log files
 - viewing 93

M

- maintenance
 - battery 21, 22
- maintenance agreement
 - best practices 20
- management GUI
 - accessing 26
 - cannot log on 37
 - supported browsers 39
 - troubleshooting procedure
 - start here 35
- management GUI interface
 - when to use 25
- management IP address
 - troubleshooting procedure 36
- managing
 - event log 93
 - message classification 108
- mirrored volumes
 - not identical 40

N

- New Zealand electronic emission statement 122
- node canister
 - LEDs 12
 - location error 39
 - reseating 51
 - unknown service address 38
 - USB ports 9
- node errors
 - fixing 48
- noncritical
 - node errors 107
- notifications
 - best practices 18
 - sending 94
 - subscribe
 - best practices 20
- number range 108

P

- panel
 - rear
 - Fibre Channel ports 8
- parts
 - removing
 - overview 65
 - preparing 65

- parts (*continued*)
 - replacing
 - overview 65
 - preparing 65
- passwords
 - best practices 17
- People's Republic of China, electronic emission statement 123
- ports
 - Ethernet 10
 - port names, worldwide 9
 - port numbers, Fibre Channel 9
 - SAS 11, 14
- POST (power-on self-test) 95
- power management
 - best practices 17
- power supply
 - replacing 70, 73
- power supply unit
 - control enclosure 5
 - expansion enclosure 6
 - with battery 5
- power-on self-test 95
- powering off
 - system 52
- problem
 - mirrored volumes
 - not identical 40

Q

- query status 33
 - reset service assistant password 33

R

- RAID
 - best practices 19
- reader feedback, sending xii
- rear-panel indicators
 - Fibre Channel ports 8
- recovery
 - cluster
 - starting 57
 - when to run 55
- related information xi
- releasing
 - cable retention bracket 78
- removing 69
 - 550 errors 57
 - 578 errors 57
 - cluster 48
 - cluster data 47
 - fibre-channel SFP transceiver 69
 - parts
 - overview 65
 - preparing 65
 - SFP transceiver 69
- replaceable units 90
- replacing
 - 2.5" drive assembly 82
 - 3.5" drive assembly 80
 - battery 76
 - blank carrier 80, 82
 - control enclosure 84
 - enclosure end cap 83

- replacing (*continued*)
 - expansion canister 67
 - expansion enclosure 88
 - fibre-channel SFP transceiver 69
 - node canister 65
 - parts
 - overview 65
 - preparing 65
 - power supply
 - control enclosure 70
 - expansion enclosure 73
 - SAS cable 83
 - SFP transceiver 69
 - support rails 89
- reporting
 - events 93
- reseating
 - node canister 51
- reset service assistant password 31
- reset service IP address 30
- reset superuser password 30
- resetting
 - superuser password 40
- restore
 - cluster 55, 58

S

- safety notices ix
 - sound pressure ix
- SAN (storage area network)
 - problem determination 52
- SAS
 - ports 11, 14
- SAS cable
 - replacing 83
- SAS cabling
 - location error 39
- satask.txt commands 30
- sending
 - comments xii
- service address
 - unknown 38
- service assistant
 - accessing 27, 52
 - interface 26
 - supported browsers 39
 - when to use 26
- service CLI
 - accessing 28
- service commands
 - apply software 32
 - CLI 28
 - create cluster 32
 - reset service assistant password 31
 - reset service IP address 30
 - reset superuser password 30
 - snap 31
- service IP address
 - changing 49
- SFP transceiver 69
 - replacing 69
- shortcut keys 117
- snap command 31
- sound pressure
 - safety notices ix

- starting
 - cluster recovery 57
- status
 - cluster 42
 - Ethernet 47
 - identifying 42
 - node canister 42
- storage area network (SAN)
 - problem determination 52
- storage systems
 - servicing 53
- Storwize V7000 library
 - related publications xi
- superuser
 - password
 - resetting 40
- supported browsers 39
- system status
 - LEDs 43

T

- T3 recovery
 - removing
 - 550 errors 57
 - 578 errors 57
 - when to run 55
- Taiwan
 - contact information 125
 - electronic emission notice 124
- trademarks 121
- troubleshooting
 - best practices 17
 - event notification email 94
 - node errors 48
 - SAN failures 52
- troubleshooting procedure
 - GUI connectivity issues
 - main GUI 36
 - service assistant 38
 - management IP address 36
 - start here
 - management GUI 35

U

- understanding
 - cluster recovery codes 108
 - error codes 95
 - event log 93
- United Kingdom electronic emission notice 124
- USB key
 - using 29
- USB ports 9
- using
 - GUI interfaces 25
 - initialization tool 29
 - initialization tool interface 28
 - management GUI 25
 - service assistant 26
 - USB key 29

V

- viewing
 - cluster
 - status 42
 - event log 93
 - node canister
 - status 42

W

- warranty agreement
 - best practices 20
- when to use
 - cluster CLI 28
 - management GUI interface 25
 - service assistant 26
- worldwide port names (WWPNs)
 - description 9



Part Number: 85Y5975

Printed in USA

GC27-2291-00



(1P) P/N: 85Y5975

