

IBM Flex System V7000 Storage Node
Version 6.4.1

*Troubleshooting, Recovery, and
Maintenance Guide*



Note

Before using this information and the product it supports, read the general information in “Notices” on page 163, the information in the “Safety and environmental notices” on page xi, as well as the information in the *IBM Environmental Notices and User Guide* , which is provided on a DVD.

This edition applies to IBM Flex System V7000 Storage Node, Version 6.4.0, and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2011, 2012.**

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

Contents

Figures vii

Tables ix

Safety and environmental notices xi

Safety xi
Guidelines for trained service technicians xiii
 Inspecting for unsafe conditions xiii
 Guidelines for servicing electrical equipment xiv
Danger notices for Flex System V7000 Storage Node xv
Caution notices for the Flex System V7000 Storage Node xvii

About this guide xix

Who should use this guide xix
Accessibility xix
Emphasis xix
Flex System V7000 Storage Node library and related publications xx
How to order IBM publications xxii
Sending your comments. xxii

Chapter 1. Major components of the Flex System V7000 Storage Node 1

4939 Control enclosure 3
 4939 Node canister 4
 4939 Node canister indicators. 5
4939 Expansion enclosure 6
 Expansion canister 7
 4939 Expansion canister indicators 8
4939 drives 9
2076 Expansion enclosure. 10
 Components in the front of the enclosure 10
 Components in the rear of the expansion enclosure 13

Chapter 2. Best practices for troubleshooting 17

Record access information 17
Follow power management procedures 18
Set up event notifications. 18
Set up inventory reporting 19
Back up your data 19
Manage your spare and failed drives 20
Resolve alerts in a timely manner 20
Keep your software up to date 20
Keep your records up to date 20
Subscribe to support notifications 21
Know your IBM warranty and maintenance agreement details 21

Chapter 3. Understanding the battery operation for the node canister 23

Chapter 4. Understanding the medium errors and bad blocks 25

Chapter 5. Flex System V7000 Storage Node user interfaces for servicing your system. 27

Management GUI interface 27
 When to use the management GUI 28
 Accessing the management GUI 28
Service assistant interface. 29
 When to use the service assistant 29
 Accessing the service assistant 30
Cluster (system) command-line interface. 31
 When to use the cluster (system) CLI. 31
 Accessing the cluster (system) CLI. 31
Service command-line interface. 31
 When to use the service CLI. 31
 Accessing the service CLI. 32
USB flash drive interface 32
 When to use the USB flash drive 32
 Using a USB flash drive 32
 satask.txt commands 33

Chapter 6. Removing and replacing 4939 enclosure parts 35

Preparing to remove and replace parts 35
4939 Tier 1 CRUs 35
Removing and replacing CRUs 36
 Replacing a node canister. 36
 Replacing an expansion canister 38
 Replacing the node canister battery 38
 Replacing a host interface card 40
 Replacing a 2.5" drive assembly 42
 Replacing a SAS cable 44
 Replacing the CMOS battery in the node canister 46
 Replacing the canister bezel 48
 Removing a canister from an enclosure 48
 Installing a canister into an enclosure. 50
 Removing the node canister cover 51
 Installing the node canister cover 52
 Removing the RFID tag 54
 Installing the RFID tag 54
 Removing and replacing FRUs (trained service technician only) 55

Chapter 7. Resolving a problem 63

Start here: Use the management GUI recommended actions 63
Problem FSM unable to discover Flex V7000 64
Problem: Management IP address unknown 65
Problem: Unable to connect to the management GUI 65
Problem: Unable to log on to the management GUI 66
Problem: Cannot create a clustered storage system 66

Problem: CMM start management panel does not display create system panel	67
Problem: Node canister service IP address unknown	68
Problem: Cannot connect to the service assistant	69
Problem: Unable to obtain a DHCP service address	70
Problem: Management GUI or service assistant does not display correctly	70
Problem: A node canister has a location node error	70
Problem: SAS cabling not valid	70
Problem: New expansion enclosure not detected	71
Problem: Control enclosure not detected.	72
Problem: Mirrored volume copies no longer identical	72
Problem: Command file not processed from USB flash drive.	72
Procedure: Resetting superuser password	73
Procedure: Identifying which enclosure or canister to service	74
Procedure: Checking the status of your system	75
Procedure: Getting node canister and system information using the service assistant	75
Procedure: Getting node canister and system information using a USB flash drive	76
Procedure: Troubleshooting the node canister using the status LEDs	76
Procedure: Understanding the system state using the enclosure LEDs.	81
Procedure: Troubleshooting the expansion canister using the status LEDs	83
Procedure: Finding the status of the Ethernet connections	86
Procedure: Removing system data from a node canister.	86
Procedure: Deleting a system completely	87
Procedure: Fixing node errors	87
Procedure: Changing the service IP address of a node canister	88
Procedure: Initializing a clustered system using the service assistant	89
Procedure: Re-seating a node canister.	90
Procedure: Powering off your system.	90
Procedure: Collecting information for support.	91
Procedure: Rescuing node canister software from another node (node rescue)	91
Fibre Channel link failures	92
Ethernet iSCSI connectivity problems.	92
Servicing storage systems.	93
Procedure: Resolving PCI errors or HIC card faults	93

Chapter 8. Using FSM to perform troubleshooting tasks 95

Chapter 9. Using the CMM for troubleshooting tasks 97

Procedure: Viewing and understanding Flex V7000 information using the CMM.	97
Procedure: Powering off a node using the CMM	99
Procedure: Powering on a node using the CMM	99
Procedure: Starting the Flex System V7000 Storage Node GUI using the CMM	100

Procedure: Identifying a Flex System V7000 Storage Node canister using the CMM.	100
Procedure: Changing the node canister service IP address using the CMM.	101
Procedure: Viewing the node canister management GUI and service IP address using the CMM	101
Procedure: Finding the temperature of a canister using the CMM	102
Procedure: Finding the power usage of a canister using the CMM	102
Procedure: Re-seating a canister using the CMM	102
Procedure: Collecting System Management Server service data using the CMM	103
Procedure: Initializing a clustered system using the CMM	104

Chapter 10. Recovery procedures . . . 107

Recover system procedure	107
When to run the recover system procedure	108
Fix hardware errors	110
Removing system information for node canisters with error code 550 or error code 578 using the service assistant	110
Performing system recovery using the service assistant	110
Recovering from offline VDIs using the CLI	112
What to check after running the system recovery	113
Backing up and restoring the system configuration	114
Backing up the system configuration using the CLI.	115
Restoring the system configuration	117
Deleting backup configuration files using the CLI.	120

Chapter 11. Event reporting 121

Understanding events	121
Viewing the event log	121
Managing the event log	122
Describing the fields in the event log	122
Event notifications.	123
Power-on self-test	123
Understanding the error codes	124
Event IDs	124
Error event IDs and error codes	128
Node error code overview	139
Clustered-system code overview	140
Error code range	140

Chapter 12. Installing optional features 157

Installing an optional host interface card	157
Installing optional drives	159

Appendix. Accessibility features for IBM Flex System V7000 Storage Node . 161

Notices 163

Trademarks	165
Electronic emission notices	165

Federal Communications Commission (FCC) statement	165
Industry Canada compliance statement	166
Avis de conformité à la réglementation d'Industrie Canada	166
Australia and New Zealand Class A Statement	166
European Union Electromagnetic Compatibility Directive	166
Germany Electromagnetic compatibility directive	167
Japan VCCI Council Class A statement	168
People's Republic of China Class A Electronic Emission Statement	168
International Electrotechnical Commission (IEC) statement.	168

United Kingdom telecommunications requirements	168
Korean Communications Commission (KCC) Class A Statement	168
Russia Electromagnetic Interference (EMI) Class A Statement	169
Taiwan Class A compliance statement	169
European Contact Information.	169
Taiwan Contact Information	169

Index 171

Figures

1.	4939 enclosure	2	23.	Releasing the host interface card	41
2.	Machine type 2076 models 212 and 224	3	24.	Positioning the host interface card	42
3.	Control enclosure	4	25.	Unlocking the drive assembly	43
4.	Node canister	4	26.	Removing a drive assembly	43
5.	SAS port	5	27.	Disconnecting SAS cable	45
6.	Node canister LEDs	6	28.	Connecting a SAS cable	46
7.	Expansion enclosure	7	29.	Location of CMOS battery	48
8.	Expansion canister	8	30.	Removing a canister	50
9.	Expansion canister LEDs	8	31.	Installing the canister	51
10.	Drive indicators	9	32.	Removing the cover.	52
11.	2076-212	10	33.	Positioning the cover	53
12.	2076-224	10	34.	Cover installation	54
13.	LED indicators on a single 3.5" drive	11	35.	Two person lift	56
14.	LED indicators on a single 2.5" drive	11	36.	Removing the enclosure	57
15.	12 drives and two end caps	12	37.	Installing the enclosure.	58
16.	Left enclosure end cap	12	38.	Two person lift	60
17.	Rear view of 2076 expansion enclosure	13	39.	Removing the enclosure	61
18.	LEDs on the power supply units of the expansion enclosure.	14	40.	Installing the enclosure.	62
19.	SAS ports and LEDs in rear of expansion enclosure	15	41.	Node canister LEDs.	77
20.	LEDs on the expansion canisters	16	42.	Enclosure fault LEDs	82
21.	Removing the battery	39	43.	Expansion canister LEDs	84
22.	Installing the battery	40	44.	Releasing the host interface card	158
			45.	Positioning the adapter	158

Tables

1. Flex System V7000 Storage Node library	xx	21. Node canister system status LED	78
2. Other IBM publications	xxi	22. Node canister fault LEDs	79
3. IBM documentation and related websites	xxi	23. Node canister battery status	79
4. Node canister components	5	24. Node canister system status LEDs	80
5. SAS port	5	25. Enclosure fault LEDs	82
6. Node canister LEDs	6	26. LED state descriptions used in the IBM Flex System V7000 expansion canister	84
7. Expansion canister LEDs	8	27. Expansion canister power LED	85
8. Drive LEDs	9	28. Expansion canister status LEDs	85
9. Drive LEDs	11	29. IBM Flex System V7000 canister statuses in CMM	98
10. LED descriptions	13	30. Results of hovering over a storage enclosure component in CMM	98
11. Power supply unit LEDs in the rear of the expansion enclosure	14	31. Results displayed by clicking a storage enclosure component in CMM	98
12. SAS port LEDs on the expansion canister	15	32. Actions displayed in menu by right-clicking a storage enclosure component in CMM	99
13. Expansion canister LEDs	16	33. Description of data fields for the event log	122
14. Access information for your system	17	34. Notification levels	123
15. Bad block errors	26	35. Informational events	124
16. Replaceable units	35	36. Error event IDs and error codes	129
17. FRUs	55	37. Message classification number range	140
18.	71		
19. LED state descriptions used in IBM Flex System V7000 Storage Node	77		
20. Node canister power LED	78		

Safety and environmental notices

Review the multilingual safety notices for the IBM® Flex System V7000 Storage Node 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

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

2. Locate *Safety Information* with the user publications that were provided with the Flex System V7000 Storage Node hardware.
3. Find the matching identification number in the *Safety Information*. Then review the topics concerning the safety notices to ensure that you are in compliance.
4. Optionally, read the multilingual safety instructions on the Flex System V7000 Storage Node website. Go to www.ibm.com/support/entry/portal/overview/hardware/puresystems/pureflex_system/storage_node/flex_system_v7000 and click the documentation link.

Safety

Before installing this product, read the Safety Information.

قبل تركيب هذا المنتج، يجب قراءة الملاحظات الأمنية

Antes de instalar este produto, leia as Informações de Segurança.

在安装本产品之前，请仔细阅读 **Safety Information** (安全信息)。

安裝本產品之前，請先閱讀「安全資訊」。

Prije instalacije ovog produkta obavezno pročitajte Sigurnosne Upute.

Před instalací tohoto produktu si přečtěte příručku bezpečnostních instrukcí.

Læs sikkerhedsforskrifterne, før du installerer dette produkt.

Lees voordat u dit product installeert eerst de veiligheidsvoorschriften.

Ennen kuin asennat tämän tuotteen, lue turvaohjeet kohdasta Safety Information.

Avant d'installer ce produit, lisez les consignes de sécurité.

Vor der Installation dieses Produkts die Sicherheitshinweise lesen.

Πριν εγκαταστήσετε το προϊόν αυτό, διαβάστε τις πληροφορίες ασφάλειας (safety information).

לפני שתתקינו מוצר זה, קראו את הוראות הבטיחות.

A termék telepítése előtt olvassa el a Biztonsági előírásokat!

Prima di installare questo prodotto, leggere le Informazioni sulla Sicurezza.

製品の設置の前に、安全情報をお読みください。

본 제품을 설치하기 전에 안전 정보를 읽으십시오.

Пред да се инсталира овој продукт, прочитајте информацијата за безбедност.

பெரிய அளவு
பெரிய அளவு
பெரிய அளவு
பெரிய அளவு
பெரிய அளவு
பெரிய அளவு
பெரிய அளவு
பெரிய அளவு
பெரிய அளவு
பெரிய அளவு

Les sikkerhetsinformasjonen (Safety Information) før du installerer dette produktet.

Przed zainstalowaniem tego produktu, należy zapoznać się z książką "Informacje dotyczące bezpieczeństwa" (Safety Information).

Antes de instalar este produto, leia as Informações sobre Segurança.

Перед установкой продукта прочтите инструкции по технике безопасности.

Pred inštaláciou tohto zariadenia si pečítajte Bezpečnostné predpisy.

Pred namestitvijo tega proizvoda preberite Varnostne informacije.

Antes de instalar este producto, lea la información de seguridad.

Läs säkerhetsinformationen innan du installerar den här produkten.

ཐོན་རྒྱུ་ལ་འདི་བའི་སྐོར་ལ་བྱས་ཤིང་། རྒྱ་ལྷོ་ཡི་ཡི་གེ་གཟུགས་
བྱ་ལ་འདྲ་མིན་ཡིན་པའི་འོད་སྤེལ་བཟང་ལྟ་དགོས།

Bu ürünü kurmadan önce güvenlik bilgilerini okuyun.

مەزكۇر مەھسۇلاتنى ئورنىتىشتىن بۇرۇن بىخەتەرلىك ئۇچۇرلىرىنى ئوقۇپ چىقىڭ.

Youq mwngz yungh canjbinj neix gaxgonq, itdingh aeu doeg aen
canjbinj soengq cungj vahgangj ancien siusik.

Guidelines for trained service technicians

This section contains information for trained service technicians.

Inspecting for unsafe conditions

Use this information to help you identify potential unsafe conditions in an IBM product that you are working on.

Each IBM product, as it was designed and manufactured, has required safety items to protect users and service technicians from injury. The information in this section addresses only those items. Use good judgment to identify potential unsafe conditions that might be caused by non-IBM alterations or attachment of non-IBM features or optional devices that are not addressed in this section. If you identify an unsafe condition, you must determine how serious the hazard is and whether you must correct the problem before you work on the product.

Consider the following conditions and the safety hazards that they present:

- Electrical hazards, especially primary power. Primary voltage on the frame can cause serious or fatal electrical shock.
- Explosive hazards, such as a damaged CRT face or a bulging capacitor.
- Mechanical hazards, such as loose or missing hardware.

To inspect the product for potential unsafe conditions, complete the following steps:

1. Make sure that the power is off and the power cords are disconnected.
2. Make sure that the exterior cover is not damaged, loose, or broken, and observe any sharp edges.
3. Check the power cords:
 - Make sure that the third-wire ground connector is in good condition. Use a meter to measure third-wire ground continuity for 0.1 ohm or less between the external ground pin and the frame ground.
 - Make sure that the power cords are the correct type.
 - Make sure that the insulation is not frayed or worn.
4. Remove the cover.
5. Check for any obvious non-IBM alterations. Use good judgment as to the safety of any non-IBM alterations.

6. Check inside the system for any obvious unsafe conditions, such as metal filings, contamination, water or other liquid, or signs of fire or smoke damage.
7. Check for worn, frayed, or pinched cables.
8. Make sure that the power-supply cover fasteners (screws or rivets) have not been removed or tampered with.

Guidelines for servicing electrical equipment

Observe these guidelines when you service electrical equipment.

- Check the area for electrical hazards such as moist floors, nongrounded power extension cords, and missing safety grounds.
- Use only approved tools and test equipment. Some hand tools have handles that are covered with a soft material that does not provide insulation from live electrical current.
- Regularly inspect and maintain your electrical hand tools for safe operational condition. Do not use worn or broken tools or testers.
- Do not touch the reflective surface of a dental mirror to a live electrical circuit. The surface is conductive and can cause personal injury or equipment damage if it touches a live electrical circuit.
- Some rubber floor mats contain small conductive fibers to decrease electrostatic discharge. Do not use this type of mat to protect yourself from electrical shock.
- Do not work alone under hazardous conditions or near equipment that has hazardous voltages.
- Locate the emergency power-off (EPO) switch, disconnecting switch, or electrical outlet so that you can turn off the power quickly in the event of an electrical accident.
- Disconnect all power before you perform a mechanical inspection, work near power supplies, or remove or install main units.
- Before you work on the equipment, disconnect the power cord. If you cannot disconnect the power cord, have the customer power-off the wall box that supplies power to the equipment and lock the wall box in the off position.
- Never assume that power has been disconnected from a circuit. Check it to make sure that it has been disconnected.
- If you have to work on equipment that has exposed electrical circuits, observe the following precautions:
 - Make sure that another person who is familiar with the power-off controls is near you and is available to turn off the power if necessary.
 - When you work with powered-on electrical equipment, use only one hand. Keep the other hand in your pocket or behind your back to avoid creating a complete circuit that could cause an electrical shock.
 - When you use a tester, set the controls correctly and use the approved probe leads and accessories for that tester.
 - Stand on a suitable rubber mat to insulate you from grounds such as metal floor strips and equipment frames.
- Use extreme care when you measure high voltages.
- To ensure proper grounding of components such as power supplies, pumps, blowers, fans, and motor generators, do not service these components outside of their normal operating locations.
- If an electrical accident occurs, use caution, turn off the power, and send another person to get medical aid.

Danger notices for Flex System V7000 Storage Node

Ensure that you are familiar with the danger notices for Flex System V7000 Storage Node.

Use the reference numbers in parentheses at the end of each notice, such as (C003) for example, to find the matching translated notice in *Safety Information*.

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.
- Sharp edges, corners and joints may be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching.

(D005)

DANGER

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

DANGER

Observe the following precautions when working on or around your IT rack system:

- Heavy equipment—personal injury or equipment damage might result if mishandled.
- Always lower the leveling pads on the rack cabinet.
- Always install stabilizer brackets on the rack cabinet.
- To avoid hazardous conditions due to uneven mechanical loading, always install the heaviest devices in the bottom of the rack cabinet. Always install servers and optional devices starting from the bottom of the rack cabinet.
- Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices.



- Each rack cabinet might have more than one power cord. Be sure to disconnect all power cords in the rack cabinet when directed to disconnect power during servicing.
- Connect all devices installed in a rack cabinet to power devices installed in the same rack cabinet. Do not plug a power cord from a device installed in one rack cabinet into a power device installed in a different rack cabinet.
- An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock.

(R001 part 1 of 2)

DANGER

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

(R001 part 2 of 2)

Caution notices for the Flex System V7000 Storage Node

Ensure that you understand the caution notices for Flex System V7000 Storage Node.

Use the reference numbers in parentheses at the end of each notice, such as (C003) for example, to find the matching translated notice in *Safety Information*.

CAUTION:

The battery contains lithium. To avoid possible explosion, do not burn or charge the battery.

Do not:

- Throw or immerse into water
- Heat to more than 100°C (212°F)
- Repair or disassemble

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. (C003)

CAUTION:

Hazardous energy is present when the blade is connected to the power source. Always replace the blade cover before installing the blade. (21)


CAUTION:

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

CAUTION:

The battery is a lithium ion battery. To avoid possible explosions, do not burn. Exchange only with the approved part. Recycle or discard the battery as instructed by local regulations. (C007a)

Note: CAUTION: Use safe practices when lifting.



UL regulatory information

This device is for use only with Listed IBM Flex System™ Enterprise Chassis.

About this guide

This guide describes how to service, maintain, and troubleshoot the IBM Flex System V7000 Storage Node.

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

The troubleshooting procedures can help you analyze failures that occur in a Flex System V7000 Storage Node 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 Flex System V7000 Storage Node.

Accessibility

IBM has a long-standing commitment to people with disabilities. In keeping with that commitment to accessibility, IBM strongly supports the U.S. Federal government's use of accessibility as a criterion in the procurement of Electronic Information Technology (EIT).

IBM strives to provide products with usable access for everyone, regardless of age or ability.

This product uses standard Windows navigation keys.

For more information, see “Accessibility features for *IBM Flex System V7000 Storage Node*,” on page 161.

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.
Bold monospace	Text in bold monospace represents 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 system.

Monospace	Text in monospace identifies the data or commands that you type, samples of command output, examples of program code or messages from the system, or names of command flags, parameters, arguments, and name-value pairs.
-----------	---

Flex System V7000 Storage Node library and related publications

Product manuals, other publications, and websites contain information that relates to Flex System V7000 Storage Node.

Flex System V7000 Storage Node Information Center

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

ralfss21.raleigh.ibm.com:8099/help/index.jsp

Flex System V7000 Storage Node library

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

www.ibm.com/support/entry/portal/overview/hardware/puresystems/pureflex_system/storage_node/flex_system_v7000

Each of the PDF publications in Table 1 is available in this information center by clicking the number in the “Order number” column:

Table 1. Flex System V7000 Storage Node library

Title	Description	Order number
<i>IBM Flex System V7000 Storage Node Installation Guide</i>	This guide provides instructions for unpacking your shipping order and installing your system. The first of three chapters describes verifying your order, becoming familiar with the hardware components, and meeting environmental requirements. The second chapter describes installing the hardware and attaching data cables. The last chapter describes initially configuring your system.	GC27-4203

Table 1. Flex System V7000 Storage Node library (continued)

Title	Description	Order number
<i>IBM Storwize V7000 Expansion Enclosure Installation Guide, Machine type 2076</i>	This guide provides instructions for unpacking your shipping order and installing the 2076 expansion enclosure for the Flex System V7000 Storage Node.	GC27-4234
<i>IBM Flex System V7000 Storage Node Troubleshooting, Recovery, and Maintenance Guide</i>	This guide describes how to service, maintain, and troubleshoot the Flex System V7000 Storage Node.	GC27-4205
<i>IBM Statement of Limited Warranty (2145 and 2076)</i>	This multilingual document provides information about the IBM warranty for machine types 2145 and 2076.	Part number: 85Y5978

Other IBM publications

Table 2 lists IBM publications that contain information related to the Flex System V7000 Storage Node.

Table 2. Other IBM publications

Title	Description	Order number
<i>IBM Storage Management Pack for Microsoft System Center Operations Manager User Guide</i>	This guide describes how to install, configure, and use the IBM Storage Management Pack for Microsoft System Center Operations Manager (SCOM).	GC27-3909 publibfp.dhe.ibm.com/epubs/pdf/c2739092.pdf
<i>IBM Storage Management Console for VMware vCenter, version 3.0.0, User Guide</i>	This publication describes how to install, configure, and use the IBM Storage Management Console for VMware vCenter, which enables Flex System V7000 Storage Node and other IBM storage systems to be integrated in VMware vCenter environments.	GA32-0929 publibfp.dhe.ibm.com/epubs/pdf/a3209295.pdf

IBM documentation and related websites

Table 3 lists websites that provide publications and other information about the Flex System V7000 Storage Node or related products or technologies.

Table 3. IBM documentation and related websites

Website	Address
Support for Flex System V7000 Storage Node (4939)	www.ibm.com/support/entry/portal/overview/hardware/puresystems/pureflex_system/storage_node/flex_system_v7000
Support for IBM System Storage® and IBM TotalStorage products	www.ibm.com/storage/support/

Table 3. IBM documentation and related websites (continued)

Website	Address
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 Flex System V7000 Storage Node documentation:

- Go to the feedback page on the website for the Flex System V7000 Storage Node Information Center at ralfss21.raleigh.ibm.com:8099/help/index.jsp?topic=/com.ibm.flex.v7000.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 Flex System V7000 Storage Node Troubleshooting, Recovery, and Maintenance Guide*
 - Publication form number: GC27-4205-00
 - Page, table, or illustration numbers that you are commenting on
 - A detailed description of any information that should be changed

Chapter 1. Major components of the Flex System V7000 Storage Node

Use this information to locate the major components on the Flex System V7000 Storage Node. The major components of the Flex System V7000 Storage Node include field replaceable units (FRUs), customer replaceable units (CRUs), and optional devices.

The Flex System V7000 Storage Node consists of one or more machine type 4939 storage enclosures mounted in an IBM Flex System chassis. In addition, the Flex System V7000 Storage Node has the capability to attach to machine type 2076 expansion enclosures. The machine types and models that make up the system are:

- Control enclosure - machine type 4939 models A49, H49, and X49
- Expansion enclosure - machine type 4939 models A29, H29, and X29
- Expansion enclosure - machine type 2076 models 212 and 224


Each 4939 control enclosure contains two node canisters and can hold up to 24 2.5-inch drives. It can connect to an expansion enclosure through the SAS port.

Each 4939 expansion enclosure contains two expansion canisters and can hold up to 24 2.5-inch drives. It connects either to a control enclosure or to another expansion enclosure through the SAS ports.

See “2076 Expansion enclosure” on page 10 for more information about machine type 2076.

Figure 1 on page 2 shows the 4939 enclosure. Because of the differences between the 4939 canisters and the enclosures, you must be able to distinguish between them when you service the system. Be aware of the differences. The machine type, model number, and serial number are located on a pull-out tab **1** on the front of the 4939 enclosure.

Note: You must know the serial number when you contact IBM support.



ite00031

Figure 1. 4939 enclosure


Figure 2 on page 3 shows the machine type 2076 expansion enclosures. Because of the differences between the 2076 enclosures, you must be able to distinguish between the enclosures when you service the system. Be aware of the following:

The machine type and model (MTM) are shown on labels that are located on the front and the rear of each 2076 enclosure.

- A machine type and model (MTM) label is located on the left end cap on the front of the enclosure.
- A machine type and model (MTM) label is located at the rear on the left enclosure flange.
- The labels also indicate if the enclosure is a control enclosure or an expansion enclosure.



- 1** Left end cap
- 2** 12 or 24 drives
- 3** Right end cap



- 1** Left end cap
- 2** 12 drives
- 3** Right end cap

Figure 2. Machine type 2076 models 212 and 224


4939 Control enclosure

The 4939 control enclosure is installed in the IBM Flex System chassis.

The enclosure contains two node canisters and up to 24 2.5-inch drives. Both spinning hard disk drives and solid state drives are supported.

The drives are mounted vertically in one row located in the front of the enclosure.

Note: The drive slots cannot be empty. A drive assembly or blank carrier must be in each slot to ensure proper cooling of the enclosure.



ite00006


Figure 3. Control enclosure

4939 Node canister


The node canister contains the main processing unit, battery, and host interface cards that control the control enclosure.

The node canister is where external systems, such as host application servers, other storage systems, and management devices are connected through the Ethernet ports or Fibre Channel ports that are part of the Flex chassis. A Node canister can also be connected to expansion enclosures through the serial-attached SCSI (SAS) ports.

Figure 4. Node canister



ite00001



ite00016

Table 4. Node canister components

1 Canister release handles	6 Host interface card slot 2 (port 2)
2 Node canister	7 Battery
3 Host interface card slot 1 (port 1)	8 Battery alignment arrows
4 Host interface card retainers	9 Cover alignment arrow
5 Host interface card locator	10 CMOS battery




Figure 5. SAS port

Table 5. SAS port

1 SAS port status indicator
2 SAS HD cable socket
3 SAS port activity indicator

4939 Node canister indicators

The indicators on the node canisters are used to provide status of the various components within the canister as well as the control enclosure. There are two canisters in each control enclosure.




Figure 6. Node canister LEDs

The node canisters have several indicators or light-emitting diodes (LEDs) that provide status of the various components. See Table 6 for a description of the various LED states that you might observe.

Table 6. Node canister LEDs

LED	Icon	Description	Color	LED	Icon	Description	Color
1	N/A	SAS port status	Amber	8		Power on	Green
2	N/A	SAS port activity	Green	9		Canister status	Green
3		Canister fault	Amber	10		Canister activity	Green
4		Internal canister fault	Amber	11		Enclosure fault	Amber
5		Battery in use	Green	12		Check log	Amber
6		Battery fault	Amber	13		Canister or control enclosure identify	Blue
7		Battery status	Green				


4939 Expansion enclosure

This topic describes the components in the 4939 expansion enclosure.

The 4939 expansion enclosure contains two expansion canisters and up to 24 2.5-inch drives.

The drives are located in the front of the enclosure and the 24 drives are mounted vertically in one row and can be spinning or solid state drives.

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



ite00008

Figure 7. Expansion enclosure

Expansion canister

This topic describes the expansion canister.

The expansion canister contains the main processing units that control the expansion enclosure. An Expansion canister must be connected to another expansion enclosure or to a control enclosure through the serial-attached SCSI (SAS) ports. There are 2 mini-SAS ports. The left-to-right features are SAS Port 1 and SAS Port 2.




Figure 8. Expansion canister

4939 Expansion canister indicators

The indicators on the expansion canisters are used to provide status of the various components within the canister as well as the expansion enclosure. There are two canisters in each expansion enclosure.

The expansion canisters have several indicators or light-emitting diodes (LEDs) that provide status of the various components. See Table 7 for a description of the various LED states that you might observe.









Figure 9. Expansion canister LEDs

Table 7. Expansion canister LEDs

LED	Icon	Description	Color	LED	Icon	Description	Color
1	N/A	SAS port status	Amber	6		Expansion canister Internal fault	Amber
2	N/A	SAS port activity	Green	7		Power on	Green
3	N/A	SAS port status	Amber	8		Canister or expansion enclosure identify	Blue
4	N/A	SAS port activity	Green	9		Expansion enclosure fault	Amber
5		Expansion canister fault	Amber				

4939 drives

Each 4939 control enclosure and 4939 expansion enclosure can hold up to 24 2.5-inch drives.

The 24 drives are mounted vertically in one row in the enclosure and can be spinning drives or solid state drives.

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

Drive indicators

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




Figure 10. Drive indicators

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

Table 8. 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

2076 Expansion enclosure

The 2076 expansion enclosures contain additional storage capacity. 2076 expansion enclosures connect either to a 4939 control enclosure or to other expansion enclosures through the SAS ports.

These are the 2076 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 drives are located in the front of the enclosure.

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

Components in the front of the enclosure

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

Drives

A 2076 expansion enclosure can hold up to 12 3.5-inch drives or up to 24 2.5-inch drives depending on the model.

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.



Figure 11. 2076-212



Figure 12. 2076-224

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

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 13. LED indicators on a single 3.5" drive



Figure 14. LED indicators on a single 2.5" drive

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

Table 9 shows the status descriptions for the two LEDs.

Table 9. 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 15 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 16 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 15. 12 drives and two end caps




Figure 16. Left enclosure end cap

Table 10. 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

Components in the rear of the expansion enclosure

This topic describes the hardware components in the rear of the expansion 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 17. Rear view of 2076 expansion enclosure

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

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




Figure 18. LEDs on the power supply units of the expansion enclosure

Table 11 identifies the LEDs in the rear of the expansion enclosure.

Table 11. 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 2076 status using the LEDs for help in diagnosing a particular failure.

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 19. 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 19. 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 12. 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 20 on page 16 shows the LEDs (**1**) in the rear of the expansion canister.




Figure 20. LEDs on the expansion canisters

Table 13. 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

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

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 this address and any limitations regarding where it can be accessed from within your Ethernet network.
- The service IP address of the node canisters on the control enclosures is used only in certain circumstances. This address is used to access a node canister when the management IP address cannot be accessed. 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. These addresses are not set during the installation of a Flex System V7000 Storage Node system, but you can set these IP addresses later by using the **chserviceip** CLI command.
- The system password for user superuser. The password is required to access the system through the service IP address. The authentication of superuser is always local; therefore, the user ID can be used when a remote authentication server that is used for other users is not available.

Table 14. Access information for your system

Item	Value	Notes
The management IP address for the GUI and CLI		
The management user ID (the default is admin)		
The management user ID password (the default is admin)		
The control enclosure management IP address		
Control enclosure service IP address: node canister 1		
Control enclosure service IP address: node canister 2		

Table 14. Access information for your system (continued)

Item	Value	Notes
The control enclosure superuser password (the default is passw0rd)		

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 a canister within an enclosure unless instructed to do so. If you power off a canister within 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 a canister within an expansion enclosure can prevent the control enclosure from flushing all the data that it has cached to the RAID arrays.

Do not use the Flex management interfaces to power off Storage Nodes unless instructed to do so by service instructions. Use the Storage Node management interfaces.

Remove a node canister only when directed to do so by a service action. When a node canister is removed, the loss of power to the node canister is detected, battery power is enabled and the critical data and state information in volatile memory is saved to a solid-state drive (SSD) contained within the node canister.

Set up event notifications

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

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

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

If your system is within warranty, or you have a hardware maintenance agreement, configure your system to send email events to IBM if an issue that requires

hardware replacement is detected. This mechanism is called Call Home. When this event is received, IBM automatically opens a problem report, and if appropriate, contacts you to verify if replacement parts are required.

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

Configuring Call Home

You can configure Call Home differently, depending on whether IBM Flex System V7000 includes FSM or not.

If your IBM Flex System V7000 includes FSM:

- You configure Call Home on FSM. FSM automatically takes control of Call Home. See *Enabling a system for monitoring by the Service and Support Manager* for information on enabling system monitoring (including performing a system discovery, requesting access to a secured system, and collecting inventory). See *Event actions for general information on event actions, and Configure e-mail notification for Service and Support Manager using automation plans* for instructions on configuring e-mail notification for Service and Support Manager using automation plans.
- If you launch the IBM Flex System V7000 installation sequence from FSM, you will not see Call Home configuration prompts in the Setup wizard. You do not perform Call Home configuration tasks on the IBM Flex System V7000.
- FSM performs Call Home for problems that SVC and CMM identify. They appear in Service and Support Manager. See *Setting up Electronic Service Agent (ESA)* for additional information.

If your IBM Flex System V7000 does not include FSM:

- You configure Call Home on the IBM Flex System V7000 by SVC email.
- You also configure Call Home on CMM to cover chassis related problems, which are relevant to IBM Flex System V7000 operation.
- CMM performs Call Home for problems with IBM Flex System V7000 canisters, including if they are dead and cannot boot, or if they have a bad DIMM.

Set up inventory reporting

Inventory reporting is an extension to the Call Home email.

Note: If you are using FSM, you do not have to set up inventory reporting.

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.

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

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

Manage your spare and failed drives

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

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

Resolve alerts in a timely manner

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

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

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

Keep your software up to date

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

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

www.ibm.com/support/entry/portal/overview/hardware/puresystems/pureflex_system/storage_node/flex_system_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.

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

Subscribe to support notifications

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

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

www.ibm.com/support/entry/portal/overview/hardware/puresystems/pureflex_system/storage_node/flex_system_v7000

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 4939 or 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 battery operation for the node canister

Each node canister in the control enclosure caches critical data and holds state information in volatile memory that must be automatically saved in the event of a power loss.

When the system detects a loss of power to the node canister, battery power is enabled and the critical data and state information in volatile memory is saved to a solid-state drive (SSD) contained within the node canister.

Each node canister contains a battery that provides power sufficient to save the critical data and state information.

Note: An expansion canister within an expansion enclosure does not cache critical data or store state information in volatile memory. Therefore, it does not require battery power.

The battery is maintained in a fully charged state by the battery subsystem and can provide enough power for two back-to-back saves of the critical data and state information. If power to a control enclosure is lost, saving critical data starts immediately. The system stops handling I/O requests from the host applications, and Metro Mirror and Global Mirror relationships go offline. If power to only a node canister in a control enclosure is lost, saving critical data on that node canister will start immediately. The I/O requests and the Metro Mirror and Global Mirror relationships will be handled by the surviving node canister in the control enclosure. The saving of critical data runs to completion, even if power is restored during this time. The loss of power could be because the input power to the Flex chassis is lost (the power supplies for the Storage node are within the chassis and are shared with other Flex components), or because the node canister is removed from the enclosure. The Flex system management components are capable of controlling power to any Flex system bay. These capabilities should only be used to power on a storage node, they should not be used to power a node off.

When the 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. The system restarts only when the battery has sufficient charge to power the node canister for the duration of saving the critical data again. If a second power outage occurs before the battery has completed charging, then the system starts in service state and does not permit I/O operations to be restarted until the battery has sufficient charge.

A battery is considered failed for the following conditions:

- When the system can communicate with the battery, and it reports an error.
- When the system cannot communicate with the battery. Failed communication exists because there is a failure within the system that makes communication with the battery impossible.

There are conditions other than loss of power that can cause critical data to be saved and the node canister to go into service state and not permit I/O operations. The node canister saves critical data if it detects there is no longer sufficient battery charge to support a saving of critical data. Data protection cannot be guaranteed in

this case. The node canister saves the critical data by using the power supplied by the chassis and then enters service state. The node canister does not handle I/O operations until the battery has sufficient charge to support the saving of the critical data. When the battery has sufficient charge, the system automatically restarts.

The front panel on the node canister has three light-emitting diodes (LEDs) or indicators that indicate the state of the battery.

- In use LED
- Status LED
- Fault LED

See “Procedure: Troubleshooting the node canister using the status LEDs” on page 76 for a complete description of the battery LEDs

Important: Although the IBM Flex System chassis is resilient to power failures and brown outs, you should always install all the enclosures in the chassis in an environment where there is reliable and consistent ac power that meets the IBM Flex System chassis requirements. Consider uninterruptible power supply units to avoid extended interruptions to data access.

Chapter 4. Understanding the medium errors and bad blocks

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

The Flex System V7000 Storage Node provides volume virtualization. This means it can move or copy volumes from one area of physical disk to another. It must maintain an exact replica of the original data. If there is a medium error within the original volume, the system must remember this and return a medium error response when the moved or copied data is read. This happens even though there is no longer a problem reading a physical drive block. Because of this difference to non-virtualized systems, the Flex System V7000 Storage Node uses the term *bad blocks* rather than medium errors.

The Flex System V7000 Storage Node allocates volumes from the extents that are on the managed disks (MDisks). The MDisk can be a volume on an external storage controller or a RAID array that is created from internal drives. In either case, depending on the RAID level used, there is normally protection against a read error on a single drive. However, it is still possible to get a medium error on a read request if multiple drives have errors or if the drives are rebuilding or are offline due to other issues.

The Flex System V7000 Storage Node provides migration facilities to move a volume from one underlying set of physical storage to another or to replicate a volume that uses FlashCopy or Metro Mirror or Global Mirror. In all these cases, the migrated volume or the replicated volume returns a medium error to the host when the logical block address on the original volume is read. The system maintains tables of bad blocks to record where the logical block addresses that cannot be read are. These tables are associated with the MDisks that are providing storage for the volumes.

The **dumpdiskbadblocks** command and the **dumpalldiskbadblocks** command are available to query the location of bad blocks.

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

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

The system creates alerts in the event log for the following situations:

- When it detects medium errors and creates a bad block
- When the bad block tables fill up

The following errors are identified:

Table 15. Bad block errors

Error code	Description
1840	The managed disk has bad blocks. On an external controller, this can only be a copied medium error.
1226	The system has failed to create a bad block because the MDisk already has the maximum number of allowed bad blocks.
1225	The system has failed to create a bad block because the system already has the maximum number of allowed bad blocks.

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

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

Chapter 5. Flex System V7000 Storage Node user interfaces for servicing your system

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

The Flex System V7000 Storage Node interfaces are integrated into the Flex chassis management interfaces. The chassis management interface provides links to the appropriate parts of the Storage Node GUIs.

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

Management GUI interface

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

About this task

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

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

Recommended actions (default)

Shows only the alerts that require attention. Alerts are listed in priority order and should be fixed sequentially by using the available fix procedures. For each problem that is selected, you can:

- Run a fix procedure.
- View the properties.

Unfixed messages and alerts

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

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

Show all

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

- Run a fix procedure.
- Mark an event as fixed.
- Filter the entries to show them by specific minutes, hours, or dates.

- Reset the date filter.
- View the properties.

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

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

When to use the management GUI

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

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

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

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

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

Accessing the management GUI

This procedure describes how to access the management GUI.

About this task

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

www.ibm.com/support/entry/portal/overview/hardware/puresystems/pureflex_system/storage_node/flex_system_v7000

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

Procedure

1. The management GUI is launched from the CMM. In the CMM chassis map, select a storage node canister.
2. From the **Actions menu** select **Launch V7000 Storage Node Console**.
3. When the connection is successful, you see a login panel.
4. Log on by using your user name and password.

Results

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

If you encounter problems logging on the management GUI or connecting to the management GUI, see “Problem: Unable to log on to the management GUI” on page 66 or “Problem: Unable to connect to the management GUI” on page 65.

Service assistant interface

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

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

When to use the service assistant

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

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

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

Use the service assistant in the following situations:

- When you cannot access the system from the management GUI and you cannot access the storage Flex System V7000 Storage Node to run the recommended actions
- When the recommended action directs you to use the service assistant.

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

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

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

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 data for the system from a node.
- Recover a system if it fails.
- Install a code package from the support site or rescue the code from another node.
- Upgrade code on node canisters manually versus performing a standard upgrade procedure.
- Configure a control enclosure chassis after replacement.
- Change the service IP address that is assigned to Ethernet port 1 for the current node canister.
- Install a temporary SSH key if a key is not installed and CLI access is required.
- Restart the services used by the system.

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 canister when it restarts. If the current node canister on which the tasks are performed is also the node canister that the browser is connected to and you lose your connection, reconnect and log on to the service assistant again after running the tasks.

Accessing the service assistant

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

About this task

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

www.ibm.com/support/entry/portal/overview/hardware/puresystems/pureflex_system/storage_node/flex_system_v7000

To start the application, perform the following steps:

Procedure

1. Start a supported web browser and point your web browser to `<serviceaddress>/service` for the node canister that you want to work on.
For example, if you set a service address of 11.22.33.44 for a node canister, point your browser to 11.22.33.44/service. If you are unable to connect to the service assistant, see “Problem: Cannot connect to the service assistant” on page 69.
2. Log on to the service assistant using the superuser password.
If you are accessing a new node canister, the default password is `passw0rd`. If the node canister is a member of a system or has been a member of a system, use the password for the superuser password.
If you do not know the current superuser password, reset the password. Go to “Procedure: Resetting superuser password” on page 73.

Results

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

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

Cluster (system) command-line interface

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

For a full description of the commands and how to start an SSH command-line session, see the “Command-line interface” topic in the “Reference” section of the Flex System V7000 Storage Node Information Center.

When to use the cluster (system) CLI

The cluster (system) 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 (system) CLI

Follow the steps that are described in the “Command-line interface” topic in the “Reference” section of the Flex System V7000 Storage Node Information Center to initialize and use a CLI session.

Service command-line interface

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

For a full description of the commands and how to start an SSH command-line session, see the “Command-line interface” topic in the “Reference” section of the Flex System V7000 Storage Node 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 canister 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 “Reference” section of the Flex System V7000 Storage Node Information Center to initialize and use a CLI session.

USB flash drive interface

Use a USB flash drive to help service the node canisters in a control enclosure.

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

Note: If an oversized USB flash drive is used, remove the canister bezel to allow proper access to the USB port. Refer to “Replacing the canister bezel” on page 48 for information on how to remove the canister bezel.

When to use the USB flash drive

The USB flash drive can be used to help when servicing your system.

Using the USB flash drive is required in the following situations:

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

Using a USB flash drive

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

About this task

Note: The node canister LED will be lit while the USB command is processing. When the command completes, the node canister LED returns to its previous state.

When a USB flash drive is plugged into a node canister, the node canister 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 flash drive. If this file does not exist, it is created. If it exists, the data is inserted at the start of the file. The file contains the details and results of the command that was run and the status and the configuration information from the node canister. The status and configuration information matches the detail that is shown on the service assistant home page panels.

The `satask.txt` file can be created on any workstation by using a text editor.

Results

The USB flash drive 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 flash drive, the result file is still created, if necessary, and the status and configuration data is written to it.

satask.txt commands

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

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

To find details on running the following CLI commands from the USB flash drive, see the Service Task Commands section in the Flex System V7000 Storage Node information center.

Reset service IP address and superuser password

Use this command to obtain service assistant access to a node canister using a service IP address and superuser password. The physical access to the node canister is required and is used to authenticate the action.

This command resets the service assistant IP address to the default address value. A default address in the range of 192.168.70.131 through 192.168.70.144, depending on the nodes slot ID, is pre-configured and will be used. If the node canister is active in a system, the superuser password for the system is reset; otherwise, the superuser password is reset on the node canister.

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

Reset service assistant password

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

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

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

Snap

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

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

Apply software

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

This command copies the file from the USB flash drive to the upgrade directory on the node canister and then installs the software.

Create cluster

Use this command to create a storage system.

Query status

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

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

Chapter 6. Removing and replacing 4939 enclosure parts

You can remove and replace customer-replaceable units (CRUs) 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. Removing a node 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, they are intended to be used only when your system is not up and running and performing I/O operations. 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 loss of access to data.

Each replaceable unit has its own removal procedure. Sometimes you can find that a step within a procedure might refer you to a different remove and replace procedure. 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.

Preparing to remove and replace parts

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

Before you begin

First, read the safety precautions in the *Safety Information*. These guidelines help you safely work with the Flex System V7000 Storage Node.

4939 Tier 1 CRUs

TheFlex System V7000 Storage Node consists of several customer replaceable units (CRUs). Generic replaceable parts are cables, canisters, and battery assemblies.

The following provides a brief description of each CRU.

Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation.

Table 16. Replaceable units

CRU	Part number	Applicable models
Controller Module, SRC w/Hsg, card (Node canister)	90Y7691	4939-A49, 4939-H49, 4939-X49
Expansion Module, SDE w/Hsg, card (Expansion canister)	90Y7692	4939-A29, 4939-H29, 4939-X29
Backup battery	90Y7689	4939-A49, 4939-H49, 4939-X49

Table 16. Replaceable units (continued)

CRU	Part number	Applicable models
Bezel Kit, Controller Module	00Y4592	4939-A49, 4939-H49, 4939-X49
Bezel Kit, Expansion Module	00Y4593	4939-A29, 4939-H29, 4939-X29
Cable (MINI SAS TO MINI SAS) .328M	90Y7687	All
Cable (HD SAS TO MINI SAS) 6M	44X3322	All
10GB iSCSI Host Interface Card	90Y7693	4939-A49, 4939-H49, 4939-X49
8GB Fibre Channel Host Interface Card	90Y7694	4939-A49, 4939-H49, 4939-X49
Hard Disk Drive Filler	90Y7688	All
RFID label tag assembly (EMEA)	00E6323	All
RFID label tag assembly (US)	74Y8800	All
500GB 7.2K NL SAS Hard Disk Drive	90Y7641	All
1000GB 7.2K NL SAS Hard Disk Drive	90Y7699	All
300GB 10K SAS Hard Disk Drive	44X3207	All
600GB 10K SAS Hard Disk Drive	44X3212	All
900GB 10K SAS Hard Disk Drive	44X3217	All
146GB 15K SAS Hard Disk Drive	90Y7666	All
300GB 15K SAS Hard Disk Drive	90Y7671	All
200GB Solid State Disk Drive	90Y7700	All
400GB Solid State Disk Drive	90Y7680	All

Removing and replacing CRUs

Use this information for removing and replacing 4939 CRUs.

About this task

Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation.

Replacing a node canister

This topic describes how to replace a node canister.

About this task

Attention:

Even though many of these procedures are hot-swappable, they are intended to be used only when your system is not up and running and performing I/O operations. 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 loss of access to data.

Be aware of the following node canister LED states:

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

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

To replace a node canister, perform the following steps:

Procedure

1. Read the safety information to which “Preparing to remove and replace parts” on page 35 refers.
2. Record which SAS cable, if any, is attached to the SAS port of the canister. The cable must be inserted back into the same port after the replacement is complete; otherwise, the system cannot function properly.
3. If a SAS cable is attached, disconnect the cable.
4. Remove the node canister from the control enclosure. See “Removing a canister from an enclosure” on page 48.
5. Carefully lay the canister on a flat, static-protective surface, orienting the canister with the bezel pointing toward you.
6. Remove the cover from the failing node canister. See “Removing the node canister cover” on page 51. For additional information about removing the cover, see the service label attached to the canister cover.
7. Remove the cover from the new node canister in the same manner.
8. Remove the battery from the failed node canister and install into the new node canister. See “Replacing the node canister battery” on page 38. For additional information about removing the battery, see the service label attached to the canister cover.
9. Remove the host interface card(s) from the failed node canister. See “Replacing a host interface card” on page 40.

Note: There is either one or two host interface cards installed in the failed node canister. Note the location of each card so that it can be installed into the new node canister in the same location.

For additional information about removing and replacing the host interface cards, see the service label attached to the canister cover.

10. Install the host interface card(s) into the new node canister. See “Replacing a host interface card” on page 40.

Attention: Make sure to install the host interface card(s) into the same location as they were in the failed node canister.

11. Reinstall the cover on both canisters. See “Installing the node canister cover” on page 52.
12. Install the new node canister into the control enclosure. See “Installing a canister into an enclosure” on page 50.
The canister starts automatically.
13. If a SAS cable was disconnected, reattach the cable.

Replacing an expansion canister

This topic describes how to replace a 4939 expansion canister.

About this task

Attention:

Even though many of these procedures are hot-swappable, they are intended to be used only when your system is not up and running and performing I/O operations. 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 loss of access to data.

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.

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.

To replace the expansion canister, complete the following steps:

Procedure

1. Read the safety information to which “Preparing to remove and replace parts” on page 35 refers.
2. Record which SAS cables are plugged into the specific ports of the 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 cable(s) for the expansion canister.
4. Remove the failing expansion canister. See “Removing a canister from an enclosure” on page 48.
5. Install the new expansion canister into the expansion enclosure. See “Installing a canister into an enclosure” on page 50.
6. Reattach the SAS cable(s).

The canister starts automatically when inserted into the enclosure.

Replacing the node canister battery

This topic describes how to replace the node canister battery.

About this task

CAUTION:

The battery is a lithium ion battery. To avoid possible explosions, do not burn. Exchange only with the approved part. Recycle or discard the battery as instructed by local regulations. (C007a)

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

Be aware of the following canister LED states:

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

To replace the node canister battery, complete the following steps:

Procedure

1. Read the safety information to which “Preparing to remove and replace parts” on page 35 refers.
2. Record which SAS cable, if any, is plugged into the SAS port of the canister. The cable must be inserted back into the same port after the replacement is complete; otherwise, the system cannot function properly.
3. If a SAS cable is attached, disconnect the cable from the node canister.
4. Remove the node canister. See “Removing a canister from an enclosure” on page 48.
5. Remove the cover from the canister. See “Removing the node canister cover” on page 51. For additional information about removing and replacing the canister cover, see the service label attached to the cover.
6. Notice that the battery has two blue touch points **1** and **2**. This is where you will lift the battery to remove it from the canister. For additional information about removing and replacing the battery, see the service label attached to the canister cover.
7. Lift the battery straight up and out of the control canister.




Figure 21. Removing the battery

8. To position the battery for installation, align the alignment arrows **3** and **4** on the battery with those on the canister. See Figure 22 on page 40. The bottom of the battery has positioning pins that protrude through holes in the system board when it is fully seated.

Important: To avoid damaging the battery contacts, make sure to lower the battery straight down into the canister.




Figure 22. Installing the battery

9. Holding the battery at the blue touch points **1** and **2**, lower the new battery into the canister. Make sure that it is fully seated onto the system board.
10. Install the canister cover. See “Installing the node canister cover” on page 52.
11. Install the canister into the control enclosure. See “Installing a canister into an enclosure” on page 50.
12. If a SAS cable was disconnected, reattach the cable to the SAS port.
The canister starts automatically when inserted into the control enclosure.

Replacing a host interface card

This topic describes how to replace a node canister host interface card.

About this task

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

Be aware of the following canister LED states:

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

To replace a node canister host interface card, complete the following steps:

Procedure

1. Remove the node canister from the control enclosure. See “Removing a canister from an enclosure” on page 48.

2. Remove the cover from the node canister. See “Removing the node canister cover” on page 51. For additional information about replacing the host interface cards, see the service label attached to the canister cover.
3. Determine which host interface card is to be replaced.

Note:

- The location of the host interface card in slot 1 (port 1) is on the left side when facing the front of the canister
 - The location of the host interface card in slot 2 (port 2) is on the right side when facing the front of the canister
4. Push the card retainer **2** to the side to release the failed host interface card **1**.




Figure 23. Releasing the host interface card

5. Slide the host interface card toward the rear of the canister until the four locating pins are aligned with larger opening of the four keyholes in the card.
6. Lift and rotate the host interface card upward to gain access to the card cable connector on the system board.
7. Disconnect the cable connector from the system board. Do not disconnect the cable from the host interface card.
8. Connect the new host interface card and cable to the system board.
9. Push the card retainer **2** out of the way.
10. Rotate the host interface card **1** into position and align the four larger opening of the keyholes in the card with the four locating pins on the card carrier.




Figure 24. Positioning the host interface card

11. Slide the host interface card toward the front of the canister until the four locating pins are aligned in smaller openings of the four keyholes in the card. Then release the card retainer.
12. Install the cover on the node canister. See “Installing the node canister cover” on page 52.
13. Install the node canister into the control enclosure. See “Installing a canister into an enclosure” on page 50. Make sure to reattach the SAS cable to the canister if one was attached.
The canister starts up automatically.

Replacing a 2.5" drive assembly

This topic describes how to replace a 2.5" drive assembly in either a control enclosure or an expansion enclosure.

About this task


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 assembly slot empty. Do not remove a drive assembly before you have a replacement available.

To replace the drive assembly, complete the following steps:

Procedure


1. Unlock the drive assembly by pressing on the red touch point at the top of the assembly with your finger or thumb. See Figure 25 on page 43 The handle will pop outward.



ite00023

Figure 25. Unlocking the drive assembly

2. Rotate the handle **1** downward to the full extension. This will start moving the drive assembly outward.
3. Open the handle to the full extension.



ite00024

Figure 26. Removing a drive assembly

4. Pull the drive assembly out until it is free of the enclosure slot.
5. Push the new drive assembly back into the slot until the handle starts to move.

6. Finish inserting the drive assembly by closing the handle until the locking catch clicks into place.

Replacing a SAS cable

This topic describes how to replace a SAS cable.

About this task

Attention: Be careful when you are replacing a cable that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

To replace a SAS cable, complete the following steps:

Procedure

1. Record which SAS cable is plugged into the specific port of the canister. The cable must be inserted back into the same port after the replacement is complete; otherwise, the system cannot function properly. Be aware of the different types of SAS cables used when connecting:
 - A 4939 control enclosure to a 4939 expansion enclosure
 - A 4939 control enclosure to a 2076 expansion enclosure
 - A 4939 expansion enclosure to a 4939 expansion enclosure
 - A 4939 expansion enclosure to a 2076 expansion enclosure
 - A 2076 expansion enclosure to a 2076 expansion enclosure
2. Disconnect one end of existing cable. Pull the blue tab with the arrow away from the connector to release it from the port. Figure 27 on page 45 shows the SAS cable.




Figure 27. Disconnecting SAS cable

3. Connect the replacement cable into the specific port.

Attention: The SAS connector sockets are recessed into the canister bezel such that you cannot see the connection. The cable connector and socket are keyed and it is important that you have proper alignment of the keys as the cable is inserted. Ensure that each SAS cable is fully inserted into the socket. A click is heard or felt when the cable is successfully inserted and you should not be able to disconnect the cable without pulling on the blue tag at the top of the cable connector. Figure 28 on page 46 shows the SAS cable socket and connector.




Figure 28. Connecting a SAS cable

Ensure that the replacement cable is fully inserted. A click is heard or felt when the cable is successfully inserted.

4. Trace the cable being replaced to the other end and disconnect it from the port.
5. Connect the replacement cable into the specific port. Ensure that the replacement cable is fully inserted. A click is heard or felt when the cable is successfully inserted.
6. Ensure that when you are finished replacing the cable that all cables are left in a neat and orderly manner.

Replacing the CMOS battery in the node canister

This topic describes how to replace the CMOS battery.

Before you begin

Be aware of the following canister LED states:

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

About this task

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

The following notes describe information that you must consider when you are replacing the CMOS battery in the storage node:

- You must replace the CMOS battery with a lithium CMOS battery of the same type.

- To order replacement batteries, call 1-800-426-7378 within the United States, and 1-800-465-7999 or 1-800-465-6666 within Canada. Outside the U.S. and Canada, call your IBM marketing representative or authorized reseller.
- After you replace the CMOS battery, you must reconfigure the compute node and reset the system date and time.
- To avoid possible danger, read and follow the following safety statement.

Statement 2



CAUTION:

When replacing the lithium battery, use only an equivalent type battery recommended by the manufacturer. If your system has a module containing a lithium battery, replace it only with the same module type made by the same manufacturer. The battery contains lithium and can explode if not properly used, handled, or disposed of.

Do not:

- Throw or immerse into water
- Heat to more than 100° C (212° F)
- Repair or disassemble

Dispose of the battery as required by local ordinances or regulations.

Attention: Touching the CMOS battery on a metal surface, such as the side of the compute node, when you replace the CMOS battery can cause it to fail.

To replace the CMOS battery, complete the following steps:

Procedure

1. Read the safety information to which “Preparing to remove and replace parts” on page 35 refers.
2. Record which SAS cable, if any, is plugged into the SAS port of the canister. The cable must be inserted back into the same port after the replacement is complete; otherwise, the system cannot function properly.
3. If a SAS cable is attached, disconnect the cable from the node canister.
4. Remove the node canister. See “Removing a canister from an enclosure” on page 48.
5. Remove the cover from the canister. See “Removing the node canister cover” on page 51. For additional information about removing and replacing the canister cover, see the service label attached to the cover.
6. Follow any special handling and installation instructions that come with the CMOS battery.
7. Remove the defective CMOS battery. Figure 29 on page 48 shows the location of the CMOS battery **10**.




Figure 29. Location of CMOS battery

8. Install the new CMOS battery.
9. Install the canister cover. See “Installing the node canister cover” on page 52.
10. Install the canister into the control enclosure. See “Installing a canister into an enclosure” on page 50.
11. If a SAS cable was disconnected, reattach the cable to the SAS port.
The canister starts automatically when inserted into the control enclosure.

Replacing the canister bezel

This topic describes how to replace the bezel on the front of the canister.

About this task

If the canister came with a RFID tag attached to the bezel, you must obtain and install a replacement RFID tag.

- To obtain an RFID tag, see “4939 Tier 1 CRUs” on page 35.
- For installation instructions, see “Installing the RFID tag” on page 54.

To replace the bezel, complete the following steps.

Procedure

1. Read the safety information to which “Preparing to remove and replace parts” on page 35 refers.
2. Use a flat-blade screwdriver to carefully pull the bezel from the canister.
3. Lift the bezel from the canister. If you are instructed to return the bezel, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.
4. Position the replacement bezel on the front of the compute node.
5. Press the bezel securely against the canister.

Removing a canister from an enclosure

This topic describes how to remove a canister from either a control enclosure or an expansion enclosure.

About this task

Attention:

Even though many of these procedures are hot-swappable, they are intended to be used only when your system is not up and running and performing I/O operations. 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 loss of access to data.

Be aware of the following enclosure LED states:


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

To remove a canister from an enclosure, complete the following steps:

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

Procedure

1. Record which SAS cables, if any, are plugged into the specific ports of the canister. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
2. Disconnect the data cables connected to the canister.
3. Rotate the two release handles **1** outward to release the canister from the enclosure.



ite00021

Figure 30. Removing a canister

4. Grasp the canister and pull it completely out of the enclosure.

Installing a canister into an enclosure

This topic describes how to install a canister into either a control enclosure or an expansion enclosure.

About this task

To Install a canister into an enclosure, complete the following steps:

Procedure

1. Rotate the canister release handles **1** outward as shown.
2. Insert the canister into the enclosure in the same location as the one that you removed.




Figure 31. Installing the canister

3. Insert the canister into the enclosure until the release handles **1** start to move.
4. Rotate the release handles inward until the canister is locked into position.
When the canister is powered on, it starts automatically.
5. If any SAS cables were disconnected, reattach the cables.

Removing the node canister cover


This topic describes how to remove the cover from the node canister.

About this task

To remove the cover from the node canister, complete the following steps:

Procedure

1. Carefully lay the canister on a flat, static-protective surface, orienting the canister with the bezel pointing toward you.
2. Press on the blue release button **1** with a finger on your left hand and at the same time push at the blue push point **2** with a finger on your right hand. Slide the cover toward the rear of the canister.
3. Lift the cover away from the canister.



ite00019

Figure 32. Removing the cover

4. Lay the cover in a safe place for future use.

Installing the node canister cover


This topic describes how to install the cover on the node canister.

About this task

To install the cover on the node canister, complete the following steps:

Procedure

1. Carefully lay the canister on a flat, static-protective surface, orienting the canister with the bezel pointing toward you.
2. Position the cover so that the arrow on the cover **1** is aligned with the alignment mark on the canister. Be aware of the 'T' shaped alignment pins **2** on the underside of the cover that must engage the slots in the canister chassis.



ite00018

Figure 33. Positioning the cover

Note: Before you close the cover, make sure that all components are installed and seated correctly and that you have not left tools or loose parts inside the canister.

3. While lightly pushing downward at the rear of the enclosure cover, slide the cover forward to the closed position, until it clicks into place. See Figure 34 on page 54.




Figure 34. Cover installation

Removing the RFID tag

This topic describes how to remove the RFID tag on the front of the canister bezel.

About this task

To remove the RFID tag, complete the following steps.

Procedure

1. Read the safety information to which “Preparing to remove and replace parts” on page 35 refers.
2. If the RFID tag is attached to the hinge, open the RFID tag.
3. Use pliers to carefully grip the hinge and gently pull and twist the base from the bezel. If you are instructed to return the RFID tag, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.
- 4.

Installing the RFID tag

This topic describes how to install the RFID tag on the front of the canister bezel.

About this task

To install the RFID tag, complete the following steps.

Procedure

1. Read the safety information to which “Preparing to remove and replace parts” on page 35 refers.
2. Make sure that the surface you are installing the tag on is thoroughly clean and dry.
3. Remove the backing from the base to expose the adhesive.

4. Position the RFID tag over the ID label on the front of the bezel.
5. Press the RFID tag securely against the bezel and hold firmly in place for 30 seconds.

Note:

- Allow 30 minutes before pivoting the RFID tag.
- Allow 24 hours to achieve near maximum adhesion.

Removing and replacing FRUs (trained service technician only)

Use this information to remove and replace field replaceable units (FRUs).

About this task

The FRUs that are described in this section must be installed only by trained service technicians.

4939 FRUs

TheFlex System V7000 Storage Node consists of the enclosure chassis FRU.

The following provides a brief description of the enclosure chassis FRU.

Table 17. FRUs

Part	Part number	Applicable models	FRU or customer replaced
Enclosure (empty chassis)	90Y7690	All	FRU (Trained Service Personnel)

Replacing a 4939 control enclosure

This topic describes how to replace a 4939 Control enclosure.

Before you begin

Note: Ensure that you know the type of enclosure that you are replacing. The procedures for replacing a control enclosure are different from those procedures for replacing an expansion enclosure. For information about replacing an expansion enclosure chassis, see “Replacing a 4939 expansion enclosure” on page 60.

About this task

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

CAUTION:

To lift and install the enclosure into the Flex chassis requires at least two people.




Figure 35. Two person lift

Attention: Perform this procedure only if instructed to do so by a service action or the IBM support center. If you have a single control enclosure, this procedure requires that you shut down your system to replace the control enclosure. If you have more than one control enclosure, you can keep part of the system running, but you lose access to the volumes that are on the affected I/O group and any volumes that are in other I/O groups that depend on the drives that are in the affected I/O group. If the system is still performing I/O requests in all the I/O groups, schedule the replacement during a maintenance period or other time when the I/O can be stopped.

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

To replace a 4939 control enclosure, complete the following steps:

Procedure

1. Read the safety information to which “Preparing to remove and replace parts” on page 35 refers.
2. Confirm that you know which enclosure to replace. Go to “Procedure: Identifying which enclosure or canister to service” on page 74.
3. If you are able to access either of the node canisters with the service assistant, record the machine type and model of the enclosure, the serial number of the enclosure, and the two worldwide node numbers (WWNN)s 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 a label tab at the front enclosure as well as on a label on the side.
4. If the enclosure is still active, shut down the host I/O and the Metro Mirror and Global Mirror activity to all the volumes that depend on the affected enclosure.

This statement applies to all volumes in the I/O group that are managed by this enclosure plus any volumes in other I/O groups that depend on the drives in the affected I/O group.

5. If your system contains a single I/O group and if the clustered system is still online, shut the system down by using the management GUI.
 - a. From the management GUI, go to **Monitoring > Manage Device**.
 - b. Select **Shut Down System** from the **Actions** menu.
 - c. Wait for the shutdown to complete.
6. If your system contains more than one I/O group and if this I/O group is still online, shut down the I/O group by using the CLI.
 - a. Identify the two nodes in the I/O group.
 - b. To shut down each node, issue the following CLI command once for each of the two node canisters:

```
stopssystem -force -node <node ID>
```
 - c. Wait for the shutdown to complete.
7. Rotate the two handles outward to release the enclosure from the chassis.





Figure 36. Removing the enclosure

8. Grasp the failed enclosure and pull it completely out of the chassis.
9. Note the position of each canister and remove the canisters from the enclosure. See “Removing a canister from an enclosure” on page 48.

Attention: The canisters must be replaced to the same position from which they were removed.
10. Note the position of each drive carrier and remove the drives from the enclosure. See “Replacing a 2.5” drive assembly” on page 42.

Attention: The drives must be replaced to the same position from which they were removed.

11. Install the drive modules into the new enclosure in the same position from which they were removed. See “Replacing a 2.5” drive assembly” on page 42.
12. Install the canisters into the new enclosure in the same position from which they were removed. See “Installing a canister into an enclosure” on page 50.
13. Write the failed enclosure machine type and model (MTM) and serial number on the repair identification (RID) label that is supplied. Attach the label to the top of the enclosure next to the agency label.
14. Install the new enclosure by sliding it into the chassis until the handles begin to move. Then seat the enclosure completely into the chassis by rotating the release handles inward until it is locked into position.



ite00026

Figure 37. Installing the enclosure

15. 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.
 You can connect using the previous service address. However, it is not always possible to maintain this address. If you cannot connect through the original service address, attempt to connect using the default service address. If you still cannot access the system, see Problem: Cannot connect to the service assistant.
 - b. Use the **Configure enclosure** panel.
 - c. Select the options to **Update WWNN 1**, **Update WWNN 2**, **Update the machine type and model**, and **Update the serial number**. Do not update the system ID. Use the node copy data for each of the values. Check that these values match the values that you recorded in step 3 on page 56.

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.

After you modify the configuration, the node attempts to restart.

Note: There are situations where the canisters restart and report critical node error 508. If the node canisters fail to become active after they restart when the enclosure is updated, check their status by using the service assistant. If both node canisters show critical node error 508, use the service assistant to restart the nodes. For any other node error, see “Procedure: Fixing node errors” on page 87. To restart a node from the service assistant, perform the following steps:

- 1) Log on to the service assistant.
 - 2) From the home page, select the node that you want to restart from the **Changed Node List**.
 - 3) Select **Actions > Restart**.
- d. The system starts and can handle I/O requests from the host systems.

Note: The configuration changes that are described in the following steps must be performed to ensure that the system is operating correctly. If you do not perform these steps, the system is unable to report certain errors.

16. Start the management GUI and select **Monitoring > System Details**. 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.
17. Select the original enclosure in the tree view.
Verify that it is offline and managed and that the serial number is correct.
18. 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.
19. Add the new enclosure to the system.
 - a. Select the enclosure from the tree view.
 - b. From the **Actions** menu, select **Add Control and Expansion Enclosures**.
 - c. Because you 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. Follow the steps in the wizard. The wizard changes the control enclosure to Managed.
 - e. Select the enclosure and add it to the system.
20. Select the new enclosure in the tree view and verify that it is now online and managed.
21. 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.
22. Check the status of all volumes and physical storage to ensure everything is online.

- Restart the host application and any FlashCopy[®] activities, Global Mirror activities, or Metro Mirror activities that were stopped.

Replacing a 4939 expansion enclosure

This topic describes how to replace a 4939 Expansion enclosure.

Before you begin

Note: Ensure that you know the type of enclosure chassis that you are replacing. The procedures for replacing an expansion enclosure chassis are different from those procedures for replacing a control enclosure chassis. For information about replacing a control enclosure chassis, see “Replacing a 4939 control enclosure” on page 55.

Attention: If your system is powered on and performing I/O operations, go 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.

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

About this task

CAUTION:

To lift and install the enclosure into the Flex chassis requires at least two people.




Figure 38. Two person lift

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 a 4939 expansion enclosure, complete the following steps:

Procedure

1. Read the safety information to which “Preparing to remove and replace parts” on page 35 refers.
2. Shut down the I/O activity to the enclosure, which includes host access, FlashCopy, Metro Mirror and Global Mirror access.
3. Confirm that you know which enclosure to replace. Go to “Procedure: Identifying which enclosure or canister to service” on page 74.
4. Record which SAS 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.
5. Disconnect the SAS cables for each canister.
6. Rotate the two handles outward to release the enclosure from the chassis.




Figure 39. Removing the enclosure

7. Grasp the failed enclosure and pull it completely out of the chassis.
8. Note the position of each canister and remove the canisters from the enclosure. See “Removing a canister from an enclosure” on page 48.
Attention: The canisters must be replaced to the same position from which they were removed.
9. Note the position of each hard disk drive carrier and remove the hard disk drives from the enclosure. See “Replacing a 2.5” drive assembly” on page 42.
Attention: The hard disk drives must be replaced to the same position from which they were removed.
10. Install the hard disk drive modules into the new enclosure in the same position from which they were removed. See “Replacing a 2.5” drive assembly” on page 42.

11. Install the canisters into the new enclosure in the same position from which they were removed. See “Installing a canister into an enclosure” on page 50.
12. Write the failed enclosure machine type and model (MTM) and serial number on the repair identification (RID) label that is supplied. Attach the label to the top of the enclosure next to the agency label.
13. Install the new enclosure by sliding it into the chassis until the handles begin to move. Then seat the enclosure completely into the chassis by rotating the release handles inward until it is locked into position.




Figure 40. Installing the enclosure

14. Reattach the SAS cables to each canister by using the information that you recorded previously.

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

15. 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.

Chapter 7. Resolving a problem

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

Note: For information on 2076 expansion enclosures, see Table 3 on page xxi.

Always use the recommended actions on the Events panel of the management GUI of the Flex System V7000 Storage Node as the starting point to diagnose and resolve a problem. The *Start here: Use the management GUI to run fix procedures* topic explains how to run the recommended action for any unfixed alert, and what to do if you cannot run the management GUI. Symptoms of a problem can include hosts being unable to access data, the system sent an e-mail notification, an LED indicates an issue, or the CMM or FSM event logs indicate a problem. In all these cases the management GUI recommended action is the correct starting point.

The following topics describe a type of problem that you might experience, that is not resolved by using the management GUI. In those situations, review the symptoms and follow the actions that are provided here. You may be directed to the IBM Flex System V7000 service assistant GUI or to Flex System chassis management devices.

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

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

Start here: Use the management GUI recommended actions

The Flex System V7000 Storage Node management GUI provides extensive facilities to help you troubleshoot and correct problems on your system.

You can connect to and manage a Flex System V7000 Storage Node system using the management GUI as soon as you have created a clustered system. If you cannot create a clustered system, see the problem that contains information about what to do if you cannot create one. Go to *“Problem: Cannot create a clustered storage system”* on page 66.

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. 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 65.

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 66.

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

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

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

The fix procedure will check that alerts are addressed in the correct order. So, for example, if a battery in a node canister is nearing end of life, that node canister should not be removed if the partner node canister is offline for some reason. The fix procedure will direct you to resolve the more serious issue before removing a component that is on the last path to data

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 75.

Problem FSM unable to discover Flex V7000

This topic provides information to be aware of should FSM or the **manageV7000** CLI command cannot discover the IBM Flex System V7000.

In the scenario where either the FSM **manageV7000** CLI command or the FSM Discover Storage dialog reports errors when discovering a IBM Flex System V7000 there may be 2 possible causes.

- There is a non-default password configured in the device.
- More than 1 FSM is trying to discover or manage the device.

To resolve the problem, use the GUI to determine if an SSH key file exists for the administrative userid you are using. If the key exists, specify that key when you issue the **manageV7000** CLI command or use the Discover Storage dialog.

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 management IP address.

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

If you do not know the management IP 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 flash drive. If you know the service address of a node canister, go to “Procedure: Getting node canister and system information using the service assistant” on page 75; otherwise, go to “Procedure: Getting node canister and system information using a USB flash drive” on page 76.

The management IP address for any IBM Flex System V7000 canister can be found using the CMM for the chassis containing the canister. The management IP address is the same for all canisters in a system, whether they are node canisters or expansion canisters. See “Procedure: Viewing and understanding Flex V7000 information using the CMM” on page 97 for details.

Problem: Unable to connect to the management GUI

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

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

- You cannot connect if the system is not operational with at least one node online. If you know the service address of a node canister, go to “Procedure: Getting node canister and system information using the service assistant” on page 75; otherwise, go to Procedure: Getting node canister and system information using a USB flash drive 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 clustered system to connect to. If the node state is service, go to “Procedure: Fixing node errors” on page 87.
- Ensure that you are using the correct system IP address. If you know the service address of a node canister, go to “Procedure: Getting node canister and system information using the service assistant” on page 75; otherwise, go to Procedure: Getting node canister and system information using a USB flash drive and obtain the management IP address from the data that is returned.
- Ping the management address to see if the Ethernet network permits the connection. If the ping fails, check the Ethernet network configuration to see if there is a routing or a firewall issue. Ensure that the Ethernet network configuration is compatible with the gateway and subnet or prefix settings. Ensure that you have not used the Ethernet address of another device as the management address. If necessary, modify your network settings to establish a connection.
- If the system IP address settings are wrong, you can use the **chsystemip** CLI command to correct the system IP address settings by using ssh and the service IP of the main configuration node. See the "Clustered system commands" topic

in the "Command-line interface" section of the IBM Flex System V7000 Storage Node Information Center for the syntax and parameters for this command.

- You can use the USB flash drive to obtain the service IP address, of the main configuration node, from the data that is returned. You can use the USB flash drive or CMM to set the service IP addresses if needed. See "USB flash drive interface" on page 32 and "Procedure: Changing the node canister service IP address using the CMM" on page 101.
- If you continue to have problems connecting to the management GUI using the system IP address, an alternative is to use the service address of the configuration node. You should use the service assistant to determine which canister is the configuration node. Then direct your browser to the service IP address of that canister; do not include the \service in the address as you do when accessing the service assistant.

Note: Access to the management GUI through this address will only persist while the canister remains the configuration node.

- Access the primary system address through the Flex Chassis internal service Ethernet. The connection to this is through the CMMs in the chassis. If there is not a functioning CMM, you cannot connect to the primary service address. You should troubleshoot the CMM.
- Access the alternate system address through the first port on the optional 10 Gbps Ethernet host interface card in the node canister. You cannot access this address if the 10Gbps Ethernet switch in the Flex chassis is not operational.

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 73.

Problem: Cannot create a clustered storage system

This topic helps if your attempt to create a clustered storage system has failed.

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

- FSM
- CMM
- USB flash drive
- Service assistant
- Service command line

The create clustered-system function protects the system from loss of volume data. If you create a clustered system on a control enclosure that was previously used, you lose all of the volumes that you previously had. To determine if there is an existing system, use data that is returned by “Procedure: Getting node canister and system information using the service assistant” on page 75 or Procedure: Getting node canister and system information using a USB flash drive.

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

If the create function failed because there is an existing system, fix the existing clustered system; do not re-create a new clustered system. If you want to create a clustered system and do not want to use any data from the volumes used in the previous clustered system, go to “Procedure: Deleting a system completely” on page 87, and then run the create function again.

You might not be able to create a cluster if the node canister (the one on which you are attempting to create the clustered system) is in service state. Check whether the node canister is in service state by using the data returned by “Procedure: Getting node canister and system information using the service assistant” on page 75 or Procedure: Getting node canister and system information using a USB flash drive. If the node is in service state, fix the reported node errors. For more information, go to “Procedure: Fixing node errors” on page 87. After the node error is corrected, attempt to create a clustered storage system again.

Problem: CMM start management panel does not display create system panel

Usually, the initialization process displays a Create system panel to initialize a management interface. If the panel does not display, use this procedure to correct the problem.

Procedure

Identify the error symptom and take the appropriate action.

Symptom	Cause and action
A “Page not found” or similar browser warning	<p>Either the service address for the node canister is inaccessible or the node has gone offline. In either case, these steps will correct the problem.</p> <ol style="list-style-type: none"> 1. Check the address in the CMM. Correct the address if it is not accessible. 2. Determine if an incorrect static IP address is in use. This commonly occurs because it was not possible to obtain an IP address through DHCP. To correct do either of the following. <ul style="list-style-type: none"> • Correct the problem that prevented DHCP address allocation. • Set an accessible unique static IP address.

Symptom	Cause and action
The management GUI is displayed	There is already an operational management system and there is no need to create a new one.
The Service assistant is displayed	The Service assistant is displayed if there is an error on the node canister that does not allow the create system panel to run. Use the Service assistance to determine if the node state is Service. If so, correct the problem using the procedures to correct critical node errors. See "Procedure: Fixing node errors" on page 87

Note: It is possible that the enclosure may have the history of an established cluster. If this is the case, the enclosure VPD cluster ID is not 0. Confirm you do not want to keep the previous cluster and reset the cluster ID using the Service assistant.

Problem: Node canister service IP address unknown

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

Note: You can also use CMM to view or set the service IP address of either canister.

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

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

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

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

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

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

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 `<control enclosure management IP address>/service` for the node that you want to work on. For example, if you set a service address of 11.22.33.44 for a node canister, point your browser to 11.22.33.44/service.
- Check that you are using the correct service address for the node canister. To find the IPv4 and IPv6 addresses that are configured on the node, go to “Problem: Node canister service IP address unknown” on page 68. Try accessing the service assistant through these addresses. Verify that the IP address, subnet, and gateway are specified correctly for IPv4 addresses. Verify that the IP address, prefix, and gateway are specified for the IPv6 addresses. If any of the values are incorrect, see “Procedure: Changing the service IP address of a node canister” on page 88.
- You cannot connect to the service assistant if the node canister is not able to start the Flex System V7000 Storage Node code. To verify that the LEDs indicate that the code has started, see “Procedure: Troubleshooting the node canister using the status LEDs” on page 76.
- Ping the service address to see if the Ethernet network permits the connection. If the ping fails, check the Ethernet network configuration to see if there is a routing or a firewall issue. Check that the Ethernet network configuration is compatible with the gateway and subnet or prefix settings. Check that you have not used an address that is used by another device on your Ethernet network. If necessary, change the network configuration or see “Procedure: Changing the service IP address of a node canister” on page 88 to change the service IP address of a node.
- A default service address is assigned to each node canister via DHCP if possible. Otherwise a default address in the range of 192.168.70.131 through 192.168.70.144, depending on the node bay number, is pre-configured and will be used. This can be changed to a static address using the CMM Component IP Configuration Page.

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

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

If the previous conditions apply, see “Procedure: Changing the service IP address of a node canister” on page 88 to change the service IP address to one that works in your environment.

Problem: Unable to obtain a DHCP service address

This topic provides assistance if you are unable to obtain a DHCP service address.

1. Identify the DHCP server from which the IBM Flex System V7000 would obtain an IP address.
2. Examine that server's logs to determine the cause of the problem and correct, if possible. If the issue cannot be corrected and you need to assign a static IP address, proceed as directed in "Procedure: Changing the service IP address of a node canister" on page 88.

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 web browser. Verify that you are using a supported web browser from the following website:

www.ibm.com/support/entry/portal/overview/hardware/puresystems/pureflex_system/storage_node/flex_system_v7000

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

Problem: A node canister has a location node error

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

The node error, that is listed on the service assistant home page, can be classified as a location error. This indicates that the node canister, or the enclosure midplane has been moved or changed. This is normally due to a service action not being completed or not being implemented correctly.

There are a number of different conditions that report as a location error. Each condition is indicated by different node error. To find out how to resolve the node error, go to "Procedure: Fixing node errors" on page 87.

Be aware that after a node canister has been used in a system, the node canister must not be moved to a different location, either within the same enclosure or in a different enclosure because this might compromise its access to storage, or a host application's access to volumes. Do not move the canister from its original location unless directed to do so by a service action.

Problem: SAS cabling not valid

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 nine expansion enclosures can be chained to the control enclosure.
- 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 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.
- Port 2 of a 2076 expansion enclosure canister must not be connected to 4939 enclosure.
- Ensure that each SAS cable is fully inserted.
- Cables connected between 4939 enclosures must not go to an enclosure in a different Flex chassis.
- Ensure that the connections in the following cases are correct.

Table 18.

Connection between	Correct connection
Two Flex enclosures	Cable should connect the two left hand canisters together or the two right hand canisters together. Do not connect a left hand canister to a right hand canister.
A Flex enclosure to a Storwize® enclosure	Cable should connect the left hand canister of the Flex enclosure to the top canister of the Storwize enclosure or the right hand canister of the Flex enclosure to the bottom canister of the Storwize enclosure.
Two Storwize enclosures	Cable should connect the two top canisters together or the two bottom canisters together. Do not connect the top canister of one enclosure to the bottom canister of the other enclosure.

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:

- For a 4939 expansion enclosure, verify the status of the LEDs on the canisters. At least one canister must be active, with no fault LED on, and all the serial-attached SCSI (SAS) port 1 LEDs must be on. For details about the LED status, see “Procedure: Troubleshooting the node canister using the status LEDs” on page 76.
- For a 2076 expansion enclosure, 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 serial-attached SCSI (SAS) port 1 LEDs must be on. For details about the LED status, see Procedure: Understanding the 2076 status using the LEDs.
- Verify that the SAS cabling to the expansion enclosure is correctly installed. To review the requirements, see “Problem: SAS cabling not valid” on page 70.

Problem: Control enclosure not detected

This topic helps you resolve why a control enclosure was not detected by the system.

When installing a new control enclosure, follow the management GUI Add Control and Expansion Enclosures wizard, which is available from the **Monitoring > System Details** menu. After selecting the control enclosure from the navigation tree, click the **Actions** menu, and then select **Add Enclosures > Control and Expansions**.

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

- The enclosure is powered on.
- The enclosure is not part of another system.
- At least one node is in candidate state.
- If the control enclosures are in different Flex chassis, ensure there is a SAN connection between the chassis FC switches, and that zoning is correctly configured.
- Zoning is set up according to the zoning rules defined in the ""Configuring"" topic of the information center. There must be a zone that includes all ports from all node canisters.

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.

Problem: Command file not processed from USB flash drive

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

You might encounter this problem when running commands using your USB flash drive.

If you encounter this situation, verify the following items:

- That an `satask_result.html` file is in the root directory on the USB flash drive. If the file does not exist, then the following problems are possible:
 - The USB flash drive is not formatted with the correct file system type. Use any USB flash drive that is formatted with FAT32, EXT2, or EXT3 file system on its first partition; for example, NTFS is not a supported type. Reformat the key or use a different key.
 - The USB port is not working. Try the key in the other USB port.
 - The node is not operational. Check the node status using the LEDs. See “Procedure: Troubleshooting the node canister using the status LEDs” on page 76.
- If there is a `satask_result.html` file, check the first entry in the file. If there is no entry that matches the time the USB flash drive was used, it is possible that the USB port is not working or the node is not operational. Check the node status using the LEDs. See “Procedure: Troubleshooting the node canister using the status LEDs” on page 76.
- If there is a status output for the time the USB flash drive was used, then the `satask.txt` file was not found. Check that the file was named correctly. The `satask.txt` file is automatically deleted after it has been processed.

Procedure: Resetting superuser password

You can reset the superuser password to the default password of `passwd` by using a USB flash drive command action.

About this task

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

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

It is possible to configure your system so that resetting the superuser password with the USB flash drive command action is not permitted. If your system is configured this way, there is no work-around. Contact the person who knows the password.

To use a USB flash drive to reset the superuser password, see “USB flash drive interface” on page 32.

See also “Problem: Unable to log on to the management GUI” on page 66.

Results

If the node canister is active in a clustered system, the password for superuser is changed on the clustered system. If the node canister is not in active state, the superuser password for the node canister is changed. If the node canister joins a clustered system later, the superuser password is reset to that of the clustered system.

Procedure: Identifying which enclosure or canister to service

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

About this task

Procedure

Use the following options to identify an enclosure. An enclosure is identified by its ID and serial number.

- For 4939 enclosures, the enclosure ID is unique within a Flex System V7000 Storage Node system. However, if you have more than one Flex System V7000 Storage Node 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.

- For 2076 expansion enclosures, 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 Flex System V7000 Storage Node system. However, if you have more than one Flex System V7000 Storage Node 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 a 4939 enclosure, a canister is identified by its slot location. Slot 1 is the canister in the left and Slot 2 is the canister in the right slot. 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.

Note: When a node canister is added to a clustered system as a node, it is given a node name and a node ID. The default node name is node*N*, where *N* is an integer number. This number does not represent the slot location of the node. Similarly, the node ID does not indicate the slot location. 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.

- Within a 2076 expansion 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.

Note: When a node canister is added to a clustered system as a node, it is given a node name and a node ID. The default node name is node*N*, where *N* is an integer number. This number does not represent the slot location of the node. Similarly, the node ID does not indicate the slot location. 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.
- Use the CMM to help identify which enclosure or canister to service. See “Procedure: Viewing and understanding Flex V7000 information using the CMM” on page 97.

Procedure: Checking the status of your system

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

About this task

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

Procedure

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

- **Monitoring > System Details**
- **Pools > MDisks by Pools**
- **Volumes > Volumes**
- **Monitoring > Events**, and then use the filtering options to display alerts, messages, or event types.

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

This procedure explains how to view information about the node canisters and system using the service assistant.

About this task

To obtain the information:

1. Log on to the service assistant, as described in “Accessing the service assistant” on page 30
2. View the information about the node canister to which you connected or the other node canister in the enclosure. To change which node's information is shown, select the node in the **Change Node** table of the Home page.

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

The node details are divided into several sections. Each section has a tab. Examine the data that is reported in each tab for the information that you want.

- The Node tab shows general information about the node canister that includes the node state and whether it is a configuration node.
- The Hardware tab shows information about the hardware.
- The Access tab shows the management IP addresses and the service addresses for this node.
- The Location tab identifies the enclosure in which the node canister is located.
- The Ports tab shows information about the I/O ports.

Procedure: Getting node canister and system information using a USB flash drive

This procedure explains how to view information about the node canister and system using a USB flash drive.

About this task

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

Procedure

1. Ensure that the USB flash drive 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.
2. The information that is returned is added to the beginning of the `satask_result.html` file. Delete this file if you no longer want the previous output.
3. Insert the USB flash drive in one of the USB ports of the node canister from which you want to collect data. There is no indication when the command has completed.
4. View the results in a web browser.

Results

The file contains the details and results of the command that was run and the status and the configuration information from the node canister.

Procedure: Troubleshooting the node canister using the status LEDs

This procedure helps you determine the node canister status using the LED indicators on the node canisters and indicates the action you should take for each state.

About this task

The LEDs provide a general idea of the node canister status. You can obtain more detail from the management GUI and the service assistant. Examine the LEDs when you are not able to access the management GUI or the service assistant, or when the system is not showing any information about a device.

The CMM also shows status information about your system, which you can use when the management GUI or the service assistant cannot be accessed. The CMM

indicates the current status of many of the LEDs on your IBM Flex System V7000 system. Also see “Procedure: Viewing and understanding Flex V7000 information using the CMM” on page 97.

The procedure shows the status for the control enclosure and node canisters. It does not show the status for the drives.

The first step is to determine the state of the control enclosure, which includes the node canisters with a battery in each canister. Your control enclosure is operational if you can manage the system using the management GUI.

Find all of the control enclosures for the system that you are troubleshooting. If you are unsure which one is the control enclosure, go to “Procedure: Identifying which enclosure or canister to service” on page 74.

Also see the following procedures that might help in troubleshooting:

- “Procedure: Understanding the system state using the enclosure LEDs” on page 81 for system state LEDs **11** and **12**.
- “Procedure: Troubleshooting the expansion canister using the status LEDs” on page 83 for expansion enclosure LEDs.

Table 19 defines the various LED states for the node canisters in the IBM Flex System V7000 Storage Node.

Table 19. LED state descriptions used in IBM Flex System V7000 Storage Node

State	Description
Off	The LED is continuously not lit or off.
Blinking slowly	The LED turns on and off at a frequency of 1 Hz: It is on for 500 ms, then off for 500 ms, then repeats.
Blinking	The LED turns on and off at a frequency of 2 Hz: It is on for 250 ms, then off for 250 ms, then repeats.
Blinking fast	The LED turns on and off at a frequency of 4 Hz: It is on for 125 ms, then off for 125 ms, then repeats.
On	The LED is continuously lit or on.
Flashing	The LED is lit to indicate some activity, then turns off. The rate and duration that the LED is lit depends on the rate and duration of the activity.

Procedure

1. Start by considering the state of the power LED **8**. Continue to troubleshoot the canister until the power LED is ON.






Figure 41. Node canister LEDs

Table 20. Node canister power LED

LED Name	Color	States
8 Power 	green	<ul style="list-style-type: none"> • OFF – There is no power to the canister. Make sure that the CMM has powered on the storage node. See Chapter 9, “Using the CMM for troubleshooting tasks,” on page 97. The CMM event log might indicate why the node canister has no power. Otherwise try re-seating the canister. Go to “Procedure: Re-seating a node canister” on page 90. If the state persists, follow the hardware replacement procedures for the parts in the following order: node canister, control enclosure. • ON SOLID – The canister is powered on. This is the normal state. Continue to step 2. • BLINK – The canister is in a standby state after being powered down. Use the CMM to power on the canister. See “Procedure: Powering on a node using the CMM” on page 99. • FAST BLINK – The node canister is in the process of communicating with the CMM during initial insertion of the canister. If the canister remains in this state for more than 10 minutes, Try re-seating the canister. Go to “Procedure: Re-seating a node canister” on page 90. If the state persists, follow the hardware replacement procedure for the node canister.



2. If the power LED is on solid, consider the state of the status LED **9** for each node canister.

Table 21. Node canister system status LED

LED Name	Color	States
9 Canister status 	green	<ul style="list-style-type: none"> • OFF – The IBM Flex System V7000 system code is not running. If the Power LED is ON, the operating system has not started the IBM Flex System V7000 system code yet. This state is normal just after the canister is powered on. If it persists for more than 10 minutes, check the canister status using the CMM. See “Procedure: Viewing and understanding Flex V7000 information using the CMM” on page 97. Unless it indicates otherwise, you should replace the node canister. • ON SOLID – The canister is active and able to perform I/O as part of a IBM Flex System V7000. The IBM Flex System V7000 management GUI is available if any node canister in the system has its status LED ON. Do not power off, or remove, a node canister whose status LED is on solid. You might lose access to data or corrupt volume data. Follow the procedures to shut down a node so that access to data is not compromised. See “Procedure: Powering off your system” on page 90. • BLINK – The canister is in candidate or service state and is not participating in I/O. If possible, use the recommended action in the event log event log shown by the management GUI to service the node canister. If this is not possible, use the service assistant to determine the state of the node canister and resolve any critical node errors that stop the canister from becoming active. See “Procedure: Getting node canister and system information using the service assistant” on page 75. If you are unable to access either GUI, continue by examining the state of the node canister fault LED.

3. If the power LED is on solid, consider the states of the fault LEDs **3** and **4** for each node canister.

Table 22. Node canister fault LEDs

LED Name	Color	States
3 Canister fault 	amber	<ul style="list-style-type: none"> • OFF – The node canister is operating correctly, or it has not been possible to isolate an issue to the node canister. • ON SOLID – The node canister has failed and should be replaced. If it is available, use the IBM Flex System V7000 management GUI event log fix procedures to guide the replacement. Otherwise, go to Chapter 7, “Resolving a problem,” on page 63 and Chapter 6, “Removing and replacing 4939 enclosure parts,” on page 35 for further information.
4 Internal fault 	amber	<p>There is a fault on one or more of the replaceable parts in the canister, it has been determined that the part needs replacing. Note: This LED only applies to the host interface cards. (There may be one or two of them.) Review the IBM Flex System V7000 event logs or the CMM to identify the host interface card that failed.</p> <ul style="list-style-type: none"> • OFF – There are no failures isolated to internal components of the canister • ON SOLID – Replace the failing host interface card

4. If the canister fault LED and the internal fault LED do not indicate a problem, consider the state of the battery LEDs **5**, **6**, and **7** on each node canister. The canister will not run as part of the IBM Flex System V7000 unless it is sure it can save its state and cache data in the event of power loss to the canister. See Chapter 3, “Understanding the battery operation for the node canister,” on page 23.

Table 23. Node canister battery status




LED Name	Color	States
5 Battery status 	green	<ul style="list-style-type: none"> • OFF – Indicates the battery is not in a state where it can support a save of cache and system state data. This is an error condition. Refer to the battery fault LED and the node status provided by the service assistant for further information. • ON SOLID – Indicates the battery is fully charged and can support a save of cache and system state data. This is the normal state. • BLINK – Indicates the battery is charging and can support at least one save of cache and system state data. No action is required; the node canister can function as an active member of the system in this state. • FAST BLINK – Indicates the battery is charging, but cannot yet support a save of cache and system state data. The node canister cannot operate as an active member of the system in this state. No action is required; when there is sufficient charge to save the cache and system state data, the node canister automatically becomes active in the system.

Table 23. Node canister battery status (continued)



LED Name	Color	States
6 Battery fault 	amber	<ul style="list-style-type: none"> • OFF – No faults have detected with the battery • ON SOLID – A fault has been detected with the battery. The battery needs to be replaced. See “Replacing the node canister battery” on page 38. • BLINK – The battery has been identified by user action.
7 Battery in use 	green	<ul style="list-style-type: none"> • OFF – The battery is not providing power to the node canister. This is the normal state • FAST BLINK – The system is currently in the process of saving cache and system state data to its internal storage device using battery power. Do not remove, or perform any other service action on the canister until the LED stops blinking. If the LED continues to fast blink for more than fifteen minutes, perform the following actions: <ol style="list-style-type: none"> 1. Re-seat the node canister. See “Procedure: Re-seating a node canister” on page 90. 2. Re-seat the battery in the node canister. See “Replacing the node canister battery” on page 38. 3. Replace the battery. See “Replacing the node canister battery” on page 38 4. Replace the node canister. See “Replacing a node canister” on page 36.

5. The following table describes the SAS port status and activity LEDs, canister activity LED, and identify LEDs on the canisters.

Table 24. Node canister system status LEDs

LED Name	Color	States
1 SAS port status	amber	<ul style="list-style-type: none"> • OFF – There are no reported faults on the SAS link. This is the normal state. • ON SOLID – There is no connectivity on the link • SLOW BLINK – The port has been disabled and will not service SAS traffic • BLINK – One or more of the narrow ports of the wide SAS port link have failed, the port is not operating at full capacity. <p>If the LED is not OFF, use the IBM Flex System V7000 management GUI, if available, and use event log fix procedures to troubleshoot. These procedures verify that the SAS cable is seated correctly at each end, it is working, or if one of the canisters it connects to needs to be replaced.</p>

Table 24. Node canister system status LEDs (continued)

LED Name	Color	States
2 SAS port activity	green	<ul style="list-style-type: none"> • OFF – Power is not present or there is no SAS link connectivity established. Check that the SAS cable is correctly connected at each end, and that the canisters it connects to are both powered on. The cable may need replacing. • ON SOLID – There is at least 1 active SAS link in the wide port established and there is no external port activity. This is a normal state and no action is required. • FLASHING – The expansion port activity LED should flash at a rate proportional to the level of SAS port interface activity as determined by the canister. The port also blinks when routing updates or configuration changes are being performed on the port. This is a normal state and no action is required.
10 Canister activity 	green	<ul style="list-style-type: none"> • OFF – There is no host I/O activity. • FLASHING – The canister is actively processing input/output traffic (host I/O activity).
13 Canister or control enclosure identify 	blue	<ul style="list-style-type: none"> • OFF – There has been no request to identify the canister or enclosure. • ON SOLID – The canister or enclosure has been identified in request to user action. • BLINK – This occurs during power on and power-on self-test (POST) activities and when a command is being run from a USB flash drive.

Procedure: Understanding the system state using the enclosure LEDs

This procedure helps you determine the system state using the enclosure LED indicators on the canisters.

About this task

The Fault LED on both the expansion and control canister and the check log LED on the control canister represent enclosure or system conditions, and not necessarily conditions that exist on just the single canister. These LEDs are therefore called enclosure LEDs. If possible, only a single enclosure LED will be lit to indicate an issue. Where possible, the LED on the right canister will be lit. However, if it is not possible to light the right hand canister's LED, the left hand canister LED will be lit. The system will err on the side of ensuring there is at least one LED indication of an issue, so there are situations where both the right canister's and left canister's enclosure LED are lit at the same time.





Figure 42. Enclosure fault LEDs

Table 25. Enclosure fault LEDs

LED Name	Color	States
11 Enclosure fault !	amber	<ul style="list-style-type: none"> • OFF – There are no isolated hardware failures on the storage enclosure. This does not mean there is no problem, An issue may have been detected that requires further diagnosis before determining that a FRU replacement is the correct action. • ON SOLID – There is a problem on the enclosure, and a single FRU part has been isolated as needing replacement. A FRU fault LED is typically lit on the enclosure, indicating the FRU that needs replacing. If the enclosure fault LED is lit, and there is not a FRU fault LED lit on a canister or drive within the enclosure, the indication is that the enclosure chassis needs replacing. As with other service actions, check the system logs for further information before replacing the enclosure chassis.

Table 25. Enclosure fault LEDs (continued)

LED Name	Color	States
12 Check log 	amber	<ul style="list-style-type: none"> • OFF – There are no conditions that require the user to login to the management interface and review the error logs • ON SOLID – A problem is detected, but it is not possible to determine that exchanging a single FRU part will fix the issue. There are multiple reasons that the Check Log LED could be illuminated. For example, it may be that there is an inconsistency in the system configuration, one of the system tasks failed, a system environment issue exists, or a hardware issue exists but it cannot be determined exactly which part has failed. <p>Note: Node canisters have Check Log LEDs. Expansion canisters do not. Normally only one canister's Check Log LED will be lit on a IBM Flex System V7000 system. If there are multiple control enclosures, check all control enclosures before determining that there are no problems being reported. The system attempts to always light the same Check Log LED. However in some circumstances, a different LED is lit. There is an exception to the Check Log LED not being lit when a failure has been isolated to a single FRU part. This is when Storwize V7000 expansion enclosures have been included in the system. The Check Log LED for the system is lit if there is any problem on a Storwize V7000 enclosure even if it has been isolated to a single FRU part such as a drive assembly. If a node canister has a problem which prevents it from operating as part of the active system, then either the Fault LED or Check Log LED is lit, depending on whether the problem has been isolated to a single replaceable part or not. If you notice a Check Log LED lit on a IBM Flex System V7000, use the IBM Flex System V7000 management GUI event log fix procedures to guide the troubleshooting. If you cannot access the management GUI, use the service assistance to view the node canister information and information in the Troubleshooting section of the Information Center to guide the service action. If you cannot access the management GUI or service assistant, use either "Procedure: Troubleshooting the node canister using the status LEDs" on page 76 or "Procedure: Troubleshooting the expansion canister using the status LEDs" depending on th type of the enclosure.</p>

Procedure: Troubleshooting the expansion canister using the status LEDs

This procedure helps you determine the expansion enclosure status using the LED indicators on the expansion canisters and indicates the action you should take for each state.

About this task

The LEDs provide a general idea of the expansion canister status. You can obtain more detail from the management GUI and the service assistant. Examine the LEDs when you are not able to access the management GUI or the service assistant, or when the system is not showing any information about a device.

The CMM also shows status information about your expansion canisters, which you can use when the management GUI or the service assistant cannot be accessed. The CMM indicates the current status of many of the LEDs on your IBM Flex System V7000.

The procedure shows the status for the expansion canisters. It does not show the status for the drives.

Table 26 defines the various LED states for the canisters in the IBM Flex System V7000 expansion canister.

Table 26. LED state descriptions used in the IBM Flex System V7000 expansion canister

State	Description
Off	The LED is continuously not lit or off.
Blinking slowly	The LED turns on and off at a frequency of 1 Hz: It is on for 500 ms, then off for 500 ms, then repeats.
Blinking	The LED turns on and off at a frequency of 2 Hz: It is on for 250 ms, then off for 250 ms, then repeats.
Blinking fast	The LED turns on and off at a frequency of 4 Hz: It is on for 125 ms, then off for 125 ms, then repeats.
On	The LED is continuously lit or on.
Flashing	The LED is lit to indicate some activity, then turns off. The rate and duration that the LED is lit depends on the rate and duration of the activity.

Procedure

1. Start by considering the states of the power LED **7** of each 4939 expansion canister.





Figure 43. Expansion canister LEDs

Table 27. Expansion canister power LED

LED Name	Color	States
7 Power 	green	<ul style="list-style-type: none"> • OFF – There is no power to the expansion canister. Make sure that the CMM has powered on the expansion canister. See Chapter 9, “Using the CMM for troubleshooting tasks,” on page 97. The CMM event log might indicate why the expansion canister has no power. Try reseating it. Go to “Procedure: Re-seating a node canister” on page 90. If the state persists, follow the hardware replacement procedures for the parts in the following order: <ol style="list-style-type: none"> 1. Expansion canister. See “Replacing an expansion canister” on page 38. 2. Expansion enclosure. See Replacing an enclosure. • ON SOLID – The expansion canister is powered on. This is the normal state. • BLINK – The expansion canister is in a powered down state. Use the CMM to power on the expansion canister. • FAST BLINK – The management controller is in the process of communicating with the Chassis Management Module (CMM) during initial insertion of the expansion canister. If the expansion canister remains in this state for more than 10 minutes, try reseating it. Go to “Procedure: Re-seating a node canister” on page 90. If the state persists, follow the hardware replacement procedure for the expansion canister.

2. If the power LED is on, consider the states of the status and fault LEDs for each expansion canister.

Note: Although it is useful to understand the meaning of the expansion canister LEDs, service actions should always start by examining the IBM Flex System V7000 management GUI event log and running the Fix Procedures for the reported errors in the order suggested.

Table 28. Expansion canister status LEDs




LED Name	Color	States
5 Expansion canister fault 	amber	<ul style="list-style-type: none"> • OFF – The expansion canister is operating correctly, or it has not been possible to isolate the issue to make it the expansion canister. • ON SOLID – The expansion canister has failed and should be replaced. If it is available, use the IBM Flex System V7000 management GUI event log fix procedures to guide the replacement. Otherwise, go to Chapter 7, “Resolving a problem,” on page 63 and Chapter 6, “Removing and replacing 4939 enclosure parts,” on page 35 for further information.
8 Expansion canister identify 	blue	<ul style="list-style-type: none"> • OFF – The expansion canister is not in identify state by the controller management system • ON SOLID – The expansion canister has been identified in response to the controller management system • BLINK – This occurs during power on and power-on self-test (POST) activities
9 Expansion enclosure fault 	amber	<p>This LED is not used and it is always off.</p>

Table 28. Expansion canister status LEDs (continued)

LED Name	Color	States
1 and 3 SAS port status	amber	<ul style="list-style-type: none"> • OFF – There have been no faults or conditions detected by the expansion canister on the SAS port or down stream device connected to the port • ON SOLID – There has been a fault condition isolated by the expansion canister on the external SAS port • SLOW BLINK – The port has been disabled and will not service SAS traffic • BLINK – One or more of the narrow ports of the SAS links on the wide SAS port link have failed, the port is not operating as a full wide port
2 and 4 SAS port activity	green	<ul style="list-style-type: none"> • OFF – Power is not present or there is no SAS link connectivity established • ON SOLID – There is at least 1 active SAS link in the wide port established and there is no external port activity • FLASHING – The expansion port activity LED should flash at a rate proportional to the level of SAS port interface activity as determined by the expansion canister. The port shall also blink when routing updates or configuration changes are being performed on the port

Procedure: Finding the status of the Ethernet connections

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

About this task

Ensure that the software is active on the node before you begin this procedure. Ethernet port 1 must be connected to an active port on your Ethernet network.

- Use the CMM to look at the status of the Flex chassis Ethernet switches.
- Use the USB flash drive to obtain the most comprehensive information for the node status. Go to Procedure: Getting node canister and system information using a USB flash drive.

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.

Procedure: Removing system data from a node canister

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

About this task

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

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 System** option to remove the system data from the node.

Results

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 system completely

This procedure guides you through the process of completely removing all system information. When the procedure is finished, the system performs like a new installation.

About this task

Attention: This procedure makes all the volume data that you have on your system inaccessible. You cannot recover the data. This procedure affects all volumes that are managed by your system.

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

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

Procedure

1. Start the service assistant on one of the node canisters.
2. Use the service assistant node action to hold the node in service state.
3. Use the **Manage System** option to remove the system data from the node.
4. Perform the previous steps on the second node canister in the enclosure.
5. On one node, open the service assistant **Configure Enclosure** and select the **Reset System ID** option. This action causes the system to reset.

Procedure: Fixing node errors

To fix a node error that is detected on one of the node canisters in your system, use this procedure.

About this task

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

Procedure

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 140.

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 clustered-system information using the service assistant. Go to “Procedure: Getting node canister and system information using the service assistant” on page 75. 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 140.

Procedure: Changing the service IP address of a node canister

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

About this task

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

Which method to use depends on the status of the system and the other node canisters in the system. Follow the methods in the order 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.

Note: You have the option of being allocated an address from a DHCP server.

Procedure

Change the service IP address.

- Use the control enclosure management GUI when the system is operating and the system is able to connect to the node with the service IP address that you want to change.
 1. Select **Settings > Network** from the navigation.
 2. Select **Service IP Addresses**.
 3. Complete the panel. Be sure to select the correct node to configure.
- Use the service assistant when you can connect to the service assistant on either the node canister that you want to configure or on a node canister that can connect to the node canister that you want to configure:
 1. Make the node canister that you want to configure the current node.
 2. Select **Change Service IP** from the menu.
 3. Complete the panel.
- Use the CMM to change the service IP address of a node canister if, for example, the canister has failed to obtain a DHCP address and a default Service IP address in the range between 192.168.70.131 and 192.168.70.144 is displayed, the canister has failed to obtain a DHCP address.
 1. Open the CMM and navigate to the Component IP Configuration page.
 2. In the list of Storage Nodes, click on the device name of the node canister.
 3. In the IP Address configuration panel, click the IPv4 tab.
 4. Select Node Service as the Network Interface and select the Configuration Method.

Tip: If the selected Configuration Method is Use Static IP Address, type the Static IP address, Subnet Mask, and Gateway address for the New Static IP Configuration.
 5. Click Apply.

Procedure: Initializing a clustered system using the service assistant

To initialize a clustered system using the service assistant rather than the CMM, use this procedure.

About this task

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

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

Procedure

1. Point your web browser to the service address of the left node canister in your control enclosure: *192.168.70.121* subnet mask: *255.255.255.0*.
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 69.
3. Select **Manage System**.
4. Enter the system name and the management IP address.
5. Click **Create System**.

Results

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

Procedure: Re-seating a node canister

A service action has directed you to re-seat a node canister. There is no need to physically remove a canister to do a re-seat, a IBM Flex System V7000 canister can be "virtually re-seated" using the CMM.

About this task

Verify that you are re-seating the correct node canister.

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

Attention: Do not physically remove the canister hardware from the enclosure. This procedure directs you to a virtual re-seat using the CMM.

Be aware of the following canister LED states:

- If the system status LED is on, do not re-seat a canister unless directed to do so by a service procedure.
- If the system status is blinking or off, it is acceptable to re-seat a canister.
- If the power LED is blinking or off, it is safe to re-seat a canister.

Procedure

1. Verify that you have selected the correct node canister and verify why you are re-seating it.

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

2. To do a virtual re-seat of the canister, go to "Procedure: Re-seating a canister using the CMM" on page 102.

Procedure: Powering off your system

Use this procedure to power off your Flex System V7000 Storage Node system when it must be serviced or to permit other maintenance actions in your data center.

About this task

To power off your Flex System V7000 Storage Node system, use the following steps:

Procedure

1. Stop hosts.
2. Shut down the system by using the management GUI. Click **Monitoring > >System Details**. From the **Actions** menu, select **Shut Down System**.

3. Wait for the power LED on both node canisters in all control enclosures to start flashing, which indicates that the shutdown operation has completed.
4. Once the power LEDs in all the control enclosures in the system are flashing, use the CMM GUI to power off any 4939 expansion enclosures. If there are 2076 expansion enclosures within the system, power them off using the switches on the power supply units in the enclosures.

Procedure: Collecting information for support

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

About this task

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

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

Procedure

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

- If you are collecting the package by using the management GUI, select **Settings > Support**. Click **Download Support Package**. Follow the instructions to download the appropriate log files.
- If you are collecting the package by using the service assistant, ensure that the node that you want to collect logs from is the current node. Select the **Collect Logs** option from the navigation. You can collect a support package or copy an individual file from the node canister. Follow the instructions to collect the information.
- If you are collecting System Management Server service data using the CMM, click **Service and Support > Download Service Data > Storage Node Service Data**. Select the storage node, dump file type, and dump file action.

Procedure: Rescuing node canister software from another node (node rescue)

Use this procedure to perform a node rescue.

About this task

A failure has indicated that the node software is damaged and must be reinstalled.

Procedure

1. Ensure that the node you want to reinstall the code on is the current node. Go to “Accessing the service assistant” on page 30.
2. Select **Reinstall Machine Code** from the navigation.
3. Select **Rescue from another node**.

Results

Fibre Channel link failures

The Fibre Channel Host interface channel connects with the chassis fibre channel switch through the backplane. Use Fibre Channel Switch diagnostics to see if there is any additional information about the failure to determine the appropriate action.

Before you begin

The diagnostics may indicate problems with:

- A component in the switch
- The host interface in the IBM Flex System V7000
- The IBM Flex System V7000 chassis

Contact IBM support with this diagnostic information.

Ethernet iSCSI connectivity problems

If you are having problems attaching to the Ethernet iSCSI hosts, your problem might be related to the network, the Flex System V7000 Storage Node system, or the host.

Procedure

Determine the problem area and take the indicated steps.

Problem area	Actions to take
Network	<ol style="list-style-type: none">1. Test your connectivity between the host and Flex System V7000 Storage Node ports.2. Try to ping the Flex System V7000 Storage Node system from the host.3. Ask the Ethernet network administrator to check the firewall and router settings. Make sure the firewall settings do not block the connection.4. Check that the subnet mask and gateway are correct for the Flex System V7000 Storage Node host configuration.
Flex System V7000 Storage Node	<ol style="list-style-type: none">1. View the configured node port IP addresses by using the lspport ip CLI command.2. View the list of volumes that are mapped to a host by using the lshostvdiskmap command to ensure that the volume host mappings are correct.3. Verify that the volume is online by using the lsvdisk command.

Problem area	Actions to take
Flex System V7000 Storage Node	<ol style="list-style-type: none"> 1. Verify that the host iSCSI qualified name (IQN) is correctly configured. 2. Use operating system utilities (such as Windows device manager) to verify that the device driver is installed, loaded, and operating correctly.

Servicing storage systems

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

Procedure: Resolving PCI errors or HIC card faults

This procedure describes the steps to take if there is a CMM event indicating a bus error.

Before you begin

Attention: Make sure that you can safely remove the canister before proceeding.

Procedure

1. Remove the canister
2. Open top cover.
3. Re-seat the reported mezzanine cards and mezzanine flex cables securely on both ends.
4. Close top cover and re-insert the canister.
5. Power on the cluster and I/O devices.
6. Check to determine whether problem is resolved.

Problem resolved?	Action
Yes	You have completed the procedure.
No	Take these steps. <ol style="list-style-type: none"> 1. Remove the canister. 2. Open the cover and replace the mezzanine card indicated by the error message. 3. If the problem continues to persist, replace the canister.

Chapter 8. Using FSM to perform troubleshooting tasks

You can use FSM to service the Flex System V7000 Storage Node with the rest of the chassis.

FSM can provide information to help you pinpoint the source of an error. These topics can help you determine the failure.

- Managing a chassis
- Managing compute nodes
- Managing storage resources
- Managing virtualized resources
- Managing network resources
- Solving problems

Chapter 9. Using the CMM for troubleshooting tasks

This topic describes CMM procedures that can help to resolve fault conditions that exist on your system. This topic assumes that you have a basic understanding of the Flex Storage Management node's operation.

The CMM procedures can often be used to help resolve problems:

- Procedures that are used to power on and off the node canisters and expansion canisters.
- Procedures that involve data collection and system configuration
- Procedures that are used to launch the management GUI

Always use the recommended actions of the Flex System V7000 Storage Node management GUI as the starting point to diagnose and resolve a problem. The topics that follow describe how to use the CMM in helping to resolve a problem that you might experience that cannot be resolved by using the management GUI.

Procedure: Viewing and understanding Flex V7000 information using the CMM

Use this procedure to view and understand Flex System V7000 Storage Node information using the CMM.

About this task

The CMM manages individual nodes within the Flex chassis. For IBM Flex System V7000 the CMM manages each canister within an enclosure individually. If you use the CMM to manage the IBM Flex System V7000 you should be aware of the inter-dependencies between all canisters in all the enclosures in the system, and not perform operations that might cause the system as a whole to stop providing access to data. It is always preferable to use the IBM Flex System V7000 management GUI to manage the system that is available. You should only use the CMM when directed to do so by service actions. To view Flex System V7000 Storage Node information using the CMM, use the following steps:

Procedure

1. Open the CMM and navigate to the chassis map. The chassis map provides a front and rear graphical view of the Flex chassis components. The front view shows compute nodes and storage control enclosures, including canisters and drives. The rear view shows power modules, cooling devices, fan-mux, IO module, and Management module.

The front view of the chassis map also shows the bay numbering scheme. Generally, bay numbers increment across each row (left to right) from the bottom left to the top. A storage control enclosure takes up four bays. Each canister within a storage control enclosure takes up a bay. Canisters have a slot designation that appears when you hover over them. For example, starting at the bottom left, drives in the first enclosure take up bays 1 through 4. The left canister is in bay 1. The right canister is in bay 2. Moving up, the next enclosure takes up bays 5-8. The left canister is in bay 5, and the right canister is in bay 6. The bay numbering scheme continues in the same manner as you move up.

2. Look for colors as an indication of a IBM Flex System V7000 canister's status.

Table 29. IBM Flex System V7000 canister statuses in CMM

Color	Status
Gray	Normal
Amber	Attention
Red	Error

3. Hover the mouse over a component to view information about it.

Table 30. Results of hovering over a storage enclosure component in CMM

Component type	Information displayed
Storage canister	<ul style="list-style-type: none"> • Name • Bay • Status • Model • Serial number
Drives	<ul style="list-style-type: none"> • Name • Bay • FRU number

4. Click a component to view detailed information, which appears in tabs in the Details panel.

Table 31. Results displayed by clicking a storage enclosure component in CMM

Component type	Information displayed
Storage canister	<ul style="list-style-type: none"> • Events • General • Hardware • Firmware • Power • Environmentals • I/O connectivity • LEDs
Drives	<ul style="list-style-type: none"> • Device name • Number of bays • Number of installed devices • Bay location • Bay width • Module description

5. Right-click a component to access an action pop-up window. Right-clicking:

Note: Attention: The use of the actions below are described on other pages. Do not use them without referring to those pages.

Table 32. Actions displayed in menu by right-clicking a storage enclosure component in CMM

Component	Action pop-up displays
Storage canister	<ul style="list-style-type: none"> • Power on • Power off (Controller will shutdown OS) • Restart System Mgmt Processor • Launch Storage Node Console • Manage identify LED

Procedure: Powering off a node using the CMM

Use this procedure when it is necessary to power off the Flex System V7000 Storage Node using the CMM.

About this task

Attention: This operation can result in permanent loss of data on storage volumes or loss of access to the volumes. Use the storage management user interface provided by the storage node to power off the storage system or individual storage controllers. Only continue if the service procedures have directed the use of this CMM power control operation.

Notes:

- Power off a Flex System V7000 Storage Node canister using the Flex System V7000 Storage Node service GUI.
- Make sure you select the correct node canister before proceeding.

To power off your Flex System V7000 Storage Node system, use the following steps:

Procedure

1. Open the CMM and navigate to the chassis map. The chassis map displayed by the CMM provides a graphical view of the Flex System chassis components.
2. Using the mouse, locate the node that you want to power off.
3. Click the node to display the **Actions** menu.
4. In the **Actions** menu click Power Off.

Procedure: Powering on a node using the CMM

Use this procedure when it is necessary to power on the Flex System V7000 Storage Node using the CMM.

About this task

To power on your Flex System V7000 Storage Node system using the CMM, use the following steps:

Procedure

1. Open the CMM and navigate to the chassis map. The chassis map displayed by the CMM provides a graphical view of the Flex chassis components.
2. Using the mouse, locate the node that you want to power on.

3. Click the node to display the **Actions** menu.
4. In the **Actions** menu click Power On.

Procedure: Starting the Flex System V7000 Storage Node GUI using the CMM

Use this procedure when it is necessary to start the Flex System V7000 Storage Node GUI using the CMM.

About this task

CMM allows you to start the most appropriate GUI for a node canister. If the node canister is an active member of a cluster, the system management GUI is started. If the node canister has an error condition that stops it operating in a cluster, the service assistant is started.

To start the Flex System V7000 Storage Node system management GUI, use the following steps:

Procedure

1. Open the CMM and navigate to the chassis map. The chassis map displayed by the CMM provides a graphical view of the Flex chassis components.
2. Using the mouse, do one of the following choices.
 - Locate any node canister in the system for which you wish to start the management GUI.
 - Select the node canister that has a critical node error, if you need to the service assistant.
3. Click on the storage node to display the **Actions** menu.
4. Click Storage Node Console.

Procedure: Identifying a Flex System V7000 Storage Node canister using the CMM

Use this procedure when it is necessary to identify a Flex System V7000 Storage Node canister using the CMM.

About this task

To identify a Flex System V7000 Storage Node canister, use the following steps:

Procedure

1. Open the CMM and navigate to the chassis map. The chassis map displayed by the CMM provides a graphical view of the Flex System chassis components.
2. Using the mouse, locate the canister that you want to identify.
3. Click the node to display the **Actions** menu.
4. In the **Actions** menu click Manage Identify LED. The Node Identify LED status (on or off) is displayed and can be turned on or off by clicking on the appropriate action.

Results

When the canister is identified the blue identity LED on the canister will turn on. See

- “4939 Node canister indicators” on page 5
- “4939 Expansion canister indicators” on page 8
- “Enclosure end cap indicators” on page 12
- “Expansion canister LEDs” on page 15

Procedure: Changing the node canister service IP address using the CMM

This procedure provides instructions on how to change the service IP address of a node canister using the CMM.

About this task

You can change the service IP address that currently applies to a canister.

Procedure

1. Open the CMM and navigate to the chassis map. The chassis map displayed by the CMM provides a graphical view of the Flex System chassis components.
2. In the menu bar at the top, click **Chassis Management**.
3. Navigate to the **Component IP Configuration** page.
4. In the list of Storage Nodes, double-click on the device name of the node canister whose service address you wish to change. A pop-up window appears.

Tip: If you select Use Static IP Address or Try DHCP server, if it fails use the static IP configuration. You must enter the Static IP address, Subnet Mask, and Gateway address for the New Static IP Configuration.

5. The IP address Configuration panel shows the IP setting for the selected node. Select to view the IPv4 or IPv6 address, by clicking the correct tab on the top of the panel. Then select the management interface you need details about in the **Network Interface** field. The currently assigned address is shown on the panel.
6. Click Apply.

Procedure: Viewing the node canister management GUI and service IP address using the CMM

This procedure provides instructions on how to view the management GUI and service IP address of a node canister using the CMM.

About this task

You can view the management GUI and service IP address that currently applies to a canister.

Procedure

1. Open the CMM and navigate to the chassis map. The chassis map displayed by the CMM provides a graphical view of the Flex System chassis components.
2. In the menu bar at the top, click **Chassis Management**.

3. Navigate to the **Component IP Configuration** page.
4. In the list of Storage Nodes, double-click on the device name of the node canister whose service address you wish to view. A pop-up window appears.
5. The IP address Configuration panel shows the IP setting for the selected node. Select to view the IPv4 or IPv6 address, by clicking the correct tab on the top of the panel. Then select the management interface you need details about in the **Network Interface** field. The currently assigned address is shown on the panel.

Procedure: Finding the temperature of a canister using the CMM

Use this procedure when it is necessary to find the temperature of a canister using the CMM.

About this task

To find the temperature of a canister using the CMM, use the following steps:

Procedure

1. Open the CMM and navigate to the chassis map. The chassis map displayed by the CMM provides a graphical view of the Flex System chassis components.
2. Using the mouse, locate the storage node canister for which you are looking for temperature.
3. Click on the canister to display a series of tabs below the graphical view. These tabs provide several choices for details related to the storage node.
4. Click the **Environmentals** tab. The temperature and voltage values are displayed for the selected canister.

Procedure: Finding the power usage of a canister using the CMM

Use this procedure when it is necessary to find the power usage of a canister using the CMM.

About this task

To find the power usage of a canister using the CMM, use the following steps:

Procedure

1. Open the CMM and navigate to the chassis map. The chassis map displayed by the CMM provides a graphical view of the Flex System chassis components.
2. Using the mouse, locate the storage node canister for which you are looking for power usage values.
3. Click on the canister to display a series of tabs below the graphical view. These tabs provide several choices for details related to the storage node.
4. Click the **Power** tab. The power usage values are displayed for the selected canister.

Procedure: Re-seating a canister using the CMM

The CMM provides a "virtual re-seat" option. This removes and then restores power to the canister. With the IBM Flex System V7000 this should be used when the service actions require a re-seat of the canister.

About this task

You must only perform this action when directed to do so by a IBM Flex System V7000 service action.

To do the equivalent of a physical canister re-seat using the CMM, use the following steps:

Procedure

1. Log in to the CMM and navigate to the chassis map. The chassis map displayed by the CMM provides a graphical view of the Flex System chassis components.
2. Carefully locate the canister to re-seat and note its bay number and CMM name.
3. Select **Service and Support** from the CMM top menu and choose the **advanced** option.
4. Select the **Service Reset Tab**.
5. Click the radio button next to the canister you want to reset. Use the bay number and CMM name noted in Step 2.
6. Click the **Reset Button** and choose the **Virtual re-seat** option. The canister will restart as if it had been physically removed from the enclosure. The restart will take a number of minutes.

Procedure: Collecting System Management Server service data using the CMM

Use this procedure when it is necessary to collect System Management Server service data using the CMM.

About this task

To collect System Management Server service data using the CMM, use the following steps:

Procedure

1. Open the CMM. In the menu bar at the top, click **Service and Support**.
2. In the Service and Support menu, click **Download Service Data**, and then click **Storage Node Service Data**.
3. Select the storage node from the **Storage Node to download data from:** pull-down menu.
4. Select the dump file type from the **Dump type:** pull-down menu. IBM service requires a service processor dump.
5. Perform the desired dump file action. The Storage Node creates a new dump file on the Storage Node, which overwrites any previous dump files. The CMM transfers the dump file from the storage node to the CMM file system. After the file is transferred, view the `/tftpboot/service` directory using the File Management page. Then, obtain the file using FTP.

Note: If the file transfer fails, you can view the error details in `Not_available.txt`, located in `/tftpboot/service` directory.

6. Click **OK**.

Procedure: Initializing a clustered system using the CMM

Use this procedure to initialize a clustered system using the CMM.

Before you begin

- Check that your IBM Flex System V7000 Storage Nodes have access to the DHCP server and there are at least two DHCP addresses available on your subnet. These are the addresses that will be used when initializing.
- The Chassis Management Module (CMM) automatically recognizes the new hardware when a new control enclosure is detected in the IBM Flex System chassis.

About this task

The CMM gives you the option to create a clustered system only if the node state is candidate.

To initialize a clustered system using the CMM, perform the following steps:

Procedure

1. Make sure that you know the address you want to assign to the new system before continuing with this procedure.
2. Open the CMM and navigate to the chassis map. In the chassis map displayed by the CMM, you will see the new **Flex System V7000 Storage Node**.
3. Select the **Actions menu**. The Action menu can be found by left clicking on the left canister in the graphical representation of the storage node.
4. From the **Actions for Canister [node nn]** select **Launch Storage Node Controller Console**. A browser window displays the IBM Flex System V7000 system startup screen. The Service IP address displayed will be one that has been allocated by your DHCP server.
 - If an IP address of 192.168.70.131 or 192.168.70.144 is displayed, the canister has failed to obtain a DHCP address and is set to the default service IP address, see “Procedure: Changing the service IP address of a node canister” on page 88.
 - If you have problems connecting to the service address, see “Problem: Cannot connect to the service assistant” on page 69.
5. When asked **Do you want to create a new system or add to an existing system?**, select **Create a new system**.
6. The **Create a new system** screen is displayed.
7. Select whether it is an **IPv4** or **IPv6** address and type in the IP address. The subnet mask and gateway will already have defaults listed. You can edit these if you wish.
8. Click **Finish** to set the management IP address for the system. System initialization begins and might take several minutes to complete.
9. When system initialization is complete, System Setup is launched automatically. The setup wizard will take you through the steps to configure basic system settings such as time and date, system name, and hardware detection and verification.
10. Log into the system using superuser as the ID and passw0rd for the password. The License agreement is displayed.
11. Review the International Program License Agreement (IPLA). You must accept the agreement before you can continue to use the product.

12. On the next screen, enter the system name and a new superuser password. At this point you have the option to use the setup wizard to configure more settings such as notifications and storage configuration. If you are not ready to complete more configuration steps now, use the configuration tasks in the management GUI to complete the configuration at a later time.

Chapter 10. Recovery procedures

This topic describes these recovery procedures: recover a system and back up and restore a system configuration.

Recover system procedure

The recover system procedure recovers the entire storage system if the data has been lost from all control enclosure node canisters. The procedure re-creates the storage system by 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 only when directed by the fix procedures. If used inappropriately, service actions can cause loss of access to data or even data loss. Before attempting to recover a storage system, investigate the cause of the failure and attempt to resolve those issues by using other fix procedures. Read and understand all of the instructions before performing any action.

Attention: Do not attempt the recovery procedure unless the following conditions are met:

- All hardware errors are fixed.
- All node canisters have candidate status.
- All node canisters must be at the same level of code that the storage system had before the system failure. If any node canisters were modified or replaced, use the service assistant to verify the levels of code, and where necessary, to upgrade or downgrade the level of code.

The system recovery procedure is one of several tasks that must be performed. The following list is an overview of the tasks and the order in which they must be performed:

1. Preparing for system recovery
 - a. Review the information regarding when to run the recover system procedure
 - b. Fix your hardware errors
 - c. Remove the system information for node canisters with error code 550 or error code 578 by using the service assistant.
2. Performing the system recovery. After you prepared the system for recovery and met all the pre-conditions, run the system recovery.

Note: Run the procedure on one system in a fabric at a time. Do not perform the procedure on different node canisters in the same system. This restriction also applies to remote systems.

3. Performing actions to get your environment operational
 - Recovering from offline VDisks (volumes) by using the CLI
 - Checking your system, for example, to ensure that all mapped volumes can access the host.

When to run the recover system procedure

A recover procedure must be attempted only after a complete and thorough investigation of the cause of the system failure. Attempt to resolve those issues by using other service procedures.

Attention: If you experience failures at any time while you are running the recover system procedure, 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:

Note: It is important that you know the number of control enclosures in the system, and when the instructions indicate that every node is checked, you must check the status of both nodes in every control enclosure. For some system problems or Fibre Channel network problems, you must run the service assistant directly on the node to get its status.

- Check to see if any node in the system has a node status of active. This status means that the system is still available. In this case, recovery is not necessary.
- Do not recover the system if the management 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 system has lost its configuration data. If any nodes report anything other than these error codes, do not perform a recovery. You can encounter situations where non-configuration nodes report other node errors, such as a 550 node error. The 550 error can also indicate that a node is not able to join a system.
 - If any nodes show a node error 550, record the error data that is associated with the 550 error from the service assistant.
 - In addition to the node error 550, the report can show data that is separated by spaces in one of the following forms:
 - Node identifiers in the format: *<enclosure_serial>-<canister slot ID>*(7 characters, hyphen, 1 number), for example, 01234A6-2
 - Quorum drive identifiers in the format: *<enclosure_serial>:<drive slot ID>[<drive 11S serial number>]* (7 characters, colon, 1 or 2 numbers, open square bracket, 22 characters, close square bracket), for example, 01234A9:21[11S1234567890123456789]
 - Quorum MDisk identifier in the format: *WWPN/LUN* (16 hexadecimal digits followed by a forward slash and a decimal number), for example, 1234567890123456/12
 - If the error data contains a node identifier, ensure that the node that is referred to by the ID is showing node error 578. If the node is showing a node error 550, ensure that the two nodes can communicate with each other. Verify the SAN connectivity, and if the 550 error is still present, restart one of the two nodes 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 the enclosure is powered on and that the drive in the reported slot is powered on and functioning. If the node canister that is reporting the fault is in the I/O group of the listed

enclosure, ensure that it has SAS connectivity to the listed enclosure. If the node canister that is reporting the fault is in a different I/O group from the listed enclosure, ensure that the listed enclosure has SAS connectivity to both node canisters in the control enclosure in its I/O group. After 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, if the 550 error is still present, restart the node by clicking **Restart Node** from the service assistant.
- If there is no error data, the error is because there are insufficient connections between nodes over the Fibre Channel network. Each node must have at least two independent Fibre Channel logical connections, or logins, to every node that is not in the same enclosure. An independent connection is one where both physical ports are different. In this case, there is a connection between the nodes, but there is not a redundant connection. If there is no error data, wait 3 minutes for the SAN to initialize. Next check that the following items:
 - That there are at least two Fibre Channel ports that are operational and connected on every node.
 - That the SAN zoning allows every port to connect to every port on every other node
 - If redundant SANs are being used, that all of them are operational.After verifying these things, if the 550 error is still present, 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 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 by using the LED indicators.
- If you have not been able to restart the system and if any node other than the current node is reporting node error 550 or 578, you must remove system data from those nodes. This action acknowledges the data loss and puts the nodes into the required candidate state.
- Do not attempt to recover the system if you have been able to restart it.
- If back-end MDisk 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.
- If any of the node canisters were replaced, they must not have participated in any other system. You can resolve this issue by performing a node rescue on the affected canister by using the service assistant. Do not perform this action on any of the other node canisters.
- 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 system is recovered.

- Any data that was in the cache at the point of 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 system 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 clustered system, a power issue is the cause. Use the CMM to verify that the enclosure has been powered on.

Removing system information for node canisters with error code 550 or error code 578 using the service assistant

The system recovery procedure works only when all node canisters are in candidate status. If there are any node canisters that display error code 550 or error code 578, you must remove their data.

About this task

Before performing this task, ensure that you have read the introductory information in the overall recover system procedure.

To remove system information from a node canister with an error 550 or 578, follow this procedure using the service assistant:

Procedure

1. Point your browser to the service IP address of one of the nodes, for example, https://node_service_ip_address/service/.
If you do not know the IP address or if it has not been configured, you must assign an IP address using the CMM. See “Procedure: Changing the node canister service IP address using the CMM” on page 101.
2. Log on to the service assistant.
3. Select **Manage System**.
4. Click **Remove System Data**.
5. Confirm that you want to remove the system data when prompted.
6. Remove the system data for the other nodes that display a 550 or a 578 error.
All nodes previously in this system must have a node status of Candidate and have no errors listed against them.
7. Resolve any hardware errors until the error condition for all nodes in the system is **None**.
8. Ensure that all nodes in the system display a status of candidate.

Results

When all nodes display a status of candidate and all error conditions are **None**, you can run the recovery procedure.

Performing system recovery using the service assistant

Start recovery when all node canisters that were members of the system are online and have candidate status. If any nodes display error code 550 or 578, remove their

system data to place them into candidate status. Do not run the recovery procedure on different node canisters in the same system.

About this task

All node canisters must be at the original level of code, prior to the system failure. If any node canisters were modified or replaced, use the service assistant to verify the levels of code, and where necessary, to upgrade or downgrade the level of code.

Attention: This service action has serious implications if not performed properly. If at any time an error is encountered not covered by this procedure, stop and call IBM Support.

Note: The web browser must not block pop-up windows, otherwise progress windows cannot open.

Any one of the following categories of messages may be displayed:

- T3 successful

The volumes are back online. Use the final checks to get your environment operational again.

- T3 failed

T3 recovery completed with errors: One or more of the volumes are offline because there was fast write data in the cache. To bring the volumes online, see *Recovering from offline volumes using the CLI* for details.

- T3 failed

Call IBM Support. Do not attempt any further action.

Run the recovery from any node canisters in the system; the node canisters must not have participated in any other system.

Note: Each individual stage of the recovery procedure might take significant time to complete, dependant upon the specific configuration.

Before performing this procedure, read the recover system procedure introductory information; see “Recover system procedure” on page 107.

Procedure

1. Point your browser to the service IP address of one of the node canisters.

If the IP address is unknown or has not been configured, assign an IP address using the initialization tool; see “Procedure: Changing the service IP address of a node canister” on page 88.

2. Log on to the service assistant.

3. Check that all node canisters that were members of the system are online and have candidate status.

If any nodes display error code 550 or 578, remove their system data to place them into candidate status; see “Procedure: Removing system data from a node canister” on page 86.

4. Select **Recover System** from the navigation.

5. Follow the online instructions to complete the recovery procedure.

- a. Verify the date and time of the last quorum time. The time stamp must be less than 30 minutes before the failure. The time stamp format is *YYYYMMDD hh:mm*, where *YYYY* is the year, *MM* is the month, *DD* is the day, *hh* is the hour, and *mm* is the minute.

Attention: If the time stamp is not less than 30 minutes before the failure, call IBM Support.

- a. Verify the date and time of the last backup date. The time stamp must be less than 24 hours before the failure. The time stamp format is *YYYYMMDD hh:mm*, where *YYYY* is the year, *MM* is the month, *DD* is the day, *hh* is the hour, and *mm* is the minute.

Attention: If the time stamp is not less than 24 hours before the failure, call IBM Support.

Changes made after the time of this backup date might not be restored.

Results

Verify the environment is operational by performing the checks provided in “What to check after running the system recovery” on page 113.

If any errors are logged in the error log after the system recovery procedure completes, use the fix procedures to resolve these errors, especially the errors related to offline arrays.

If the recovery completes with offline volumes, go to “Recovering from offline VDisks using the CLI.”

Recovering from offline VDisks using the CLI

If a recovery procedure (T3 procedure) completes with offline volumes, you can use the command-line interface (CLI) to access the volumes.

About this task

If you have performed the recovery procedure, and it has completed successfully but there are offline volumes, you can perform the following steps to bring the volumes back online. Any volumes that are offline and are not thin-provisioned volumes are offline because of the loss of write-cache data during the event that led both nodes to lose their hardened data. These volumes might need additional recovery steps after the volume is brought back online.

Note: If you encounter errors in the error log after running the recovery procedure that are related to offline arrays, use the fix procedures to resolve the offline array errors before fixing the offline volume (VDisk) errors.

Example

Perform the following steps to recover an offline volume after the recovery procedure has completed:

1. Delete all IBM FlashCopy function mappings and Metro Mirror or Global Mirror relationships that use the offline volumes.
2. Run the **recovervdisk** or **recovervdiskbysystem** command.

You can recover individual volumes by using the **recovervdisk** command. You can recover all the volumes in a clustered system by using the **recovervdiskbysystem** command.

3. Recreate all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the volumes.

What to check after running the system recovery

Several tasks must be performed before you use the system.

The recovery procedure performs a recreation of the old system from the quorum data. However, some things cannot be restored, such as cached data or system data managing in-flight I/O. This latter loss of state affects RAID arrays managing internal storage. The detailed map about where data is out of synchronization has been lost, meaning that all parity information must be restored, and mirrored pairs must be brought back into synchronization. Normally this results in either old or stale data being used, so only writes in flight are affected. However, if the array had lost redundancy (such as syncing, or degraded or critical RAID status) prior to the error requiring system recovery, then the situation is more severe. Under this situation you need to check the internal storage:

- Parity arrays will likely be syncing to restore parity; they do not have redundancy when this operation proceeds.
- Because there is no redundancy in this process, bad blocks may have been created where data is not accessible.
- Parity arrays could be marked as corrupt. This indicates that the extent of lost data is wider than in-flight IO, and in order to bring the array online, the data loss must be acknowledged.
- Raid-6 arrays that were actually degraded prior the system recovery may require a full restore from backup. For this reason, it is important to have at least a capacity match spare available.

Be aware of the following differences regarding the recovered configuration:

- FlashCopy mappings are restored as “idle_or_copied” with 0% progress. Both volumes must have been restored to their original I/O groups.
- The management ID is different. Any scripts or associated programs that refer to the system-management ID of the clustered system must be changed.
- Any FlashCopy mappings that were not in the “idle_or_copied” state with 100% progress at the point of disaster have inconsistent data on their target disks. These mappings must be restarted.
- Intersystem 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.
- Intrasystem remote copy relationships are restored if all dependencies were successfully restored to their original I/O groups.
- The system 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 system configuration

You can back up and restore the configuration data for the system after preliminary tasks are completed.

Configuration data for the system provides information about your system and the objects that are defined in it. The backup and restore functions of the **svconfig** command can back up and restore only your configuration data for the Flex System V7000 Storage Node system. You must regularly back up your application data by using the appropriate backup methods.

You can maintain your configuration data for the system by completing the following tasks:

- Backing up the configuration data
- Restoring the configuration data
- Deleting unwanted backup configuration data files

Before you back up your configuration data, the following prerequisites must be met:

- No independent operations that change the configuration for the system can be running while the backup command is running.
- No object name can begin with an underscore character (_).

Note:

- The default object names for controllers, I/O groups, and managed disks (MDisks) do not restore correctly if the ID of the object is different from what is recorded in the current configuration data file.
- All other objects with default names are renamed during the restore process. The new names appear in the format *name_r* where *name* is the name of the object in your system.

Before you restore your 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 configuration files on a server that is accessible to the system.
- You have a backup copy of your application data that is ready to load on your system after the restore configuration operation is complete.
- You know the current license settings for your system.
- You did not remove any hardware since the last backup of your configuration.
- No zoning changes were made on the Fibre Channel fabric which would prevent communication between the Flex System V7000 Storage Node and any storage controllers which are present in the configuration.
- For configurations with more than one I/O group, if a new system is created on which the configuration data is to be restored, the I/O groups for the other control enclosures must be added.

Use the following steps to determine how to achieve an ideal T4 recovery:

- Open the appropriate `svc.config.backup.xml` (or `svc.config.cron.xml`) file with a suitable text editor or browser and navigate to the **node section** of the file.
- For each node entry, make a note of the value of following properties; `IO_group_id`, `canister_id`, `enclosure_serial_number`.
- Use the CLI `sainfo lsservicenodes` command and the `adata` to determine which node canisters previously belonged in each IO group.

Restoring the system configuration should be performed via one of the nodes previously in IO group zero. For example, **property name="IO_group_id" value="0"** . The remaining enclosures should be added, as required, in the appropriate order based on the previous **IO_group_id** of its node canisters.

Note: It is not currently possible to determine which canister within the identified enclosure was previously used for cluster creation. Typically the restoration should be performed via canister 1.

The Flex System V7000 Storage Node analyzes the backup configuration data file and the system to verify that the required disk controller system nodes are available.

Before you begin, hardware recovery must be complete. The following hardware must be operational: hosts, Flex System V7000 Storage Node, drives, the Ethernet network, and the SAN fabric.

Backing up the system configuration using the CLI

You can back up your configuration data using the command-line interface (CLI).

Before you begin

Before you back up your configuration data, the following prerequisites must be met:

- No independent operations that change the configuration can be running while the backup command is running.
- No object name can begin with an underscore character (`_`).

About this task

The backup feature of the `svconfig` CLI command is designed to back up information about your system configuration, such as volumes, local Metro Mirror information, local Global Mirror information, managed disk (MDisk) groups, and nodes. All other data that you wrote to the volumes is *not* backed up. Any application that uses the volumes on the system as storage, must back up its application data using the appropriate backup methods.

You must regularly back up your configuration data and your application data to avoid data loss. It is recommended that this is performed after any significant changes in configuration have been made to the system. Note that the system automatically creates a backup of the configuration data each day at 1AM. This is known as a **cron** backup and is written to `/dumps/svc.config.cron.xml_<serial#>` on the configuration node. A manual backup can be generated at any time using the instructions in this task. If a severe failure occurs, both the configuration of the system and application data may be lost. The backup of the configuration data can be used to restore the system configuration to the exact state it was in before the failure. In some cases it may be possible to automatically recover the application data. This can be attempted via the <Recover System Procedure> also known as a

Tier 3 (T3) procedure. Restoring the system configuration without attempting to recover the application data is performed via the <Restoring the System Configuration> procedure also known as a Tier 4 (T4) recovery. Both of these procedures require a recent backup of the configuration data.

Perform the following steps to back up your configuration data:

Procedure

1. Back up all of the application data that you stored on your volumes using your preferred backup method.

2.

```
svconfig clear -all
```

3. Issue the following CLI command to back up your configuration:

```
svconfig backup
```

The following output is an example of the messages that may be displayed during the backup process:

```
CMMVC6112W io_grp io_grp1 has a default name
CMMVC6112W io_grp io_grp2 has a default name
CMMVC6112W mdisk mdisk14 ...
CMMVC6112W node node1 ...
CMMVC6112W node node2 ...
.....
```

The **svconfig backup** CLI command creates three files that provide information about the backup process and the configuration. These files are created in the /dumps directory of the configuration node canister.

The following table describes the three files that are created by the backup process:

File name	Description
svc.config.backup.xml_<serial#>	This file that contains your configuration data.
svc.config.backup.sh_<serial#>	This file that contains the names of the commands that were issued to create the backup of the system.
svc.config.backup.log_<serial#>	This file contains details about the backup, including any reported errors or warnings.

4. Check that the **svconfig backup** command completes successfully, and examine the command output for any warnings or errors. The following output is an example of the message that is displayed when the backup process is successful:

```
CMMVC6155I SVCONFIG processing completed successfully.
```

If the process fails, resolve the errors, and run the command again.

5. It is recommended to keep backup copies of the files above outside the system to protect them against a system hardware failure. Copy the backup files off the system to a secure location using either the management GUI or scp command line. For example:

```
pscp superuser@cluster_ip:/dumps/svc.config.backup.*
/offclusterstorage/
```

The `cluster_ip` is the IP address or DNS name of the system and `offclusterstorage` is the location where you want to store the backup files.

Tip: To maintain controlled access to your configuration data, copy the backup files to a location that is password-protected.

Restoring the system configuration

Use this procedure in the following situations: only if the recover procedure has failed or if the data that is stored on the volumes is not required. For directions on the recover procedure, see “Recover system procedure” on page 107.

Before you begin

This configuration restore procedure is designed to restore information about your configuration, such as volumes, local Metro Mirror information, local Global Mirror information, storage pools, and nodes. 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 clustered system as storage separately. Therefore, you must have a backup of this data before you follow the configuration recovery process.

About this task

You must regularly back up your configuration data and your application data to avoid data loss. If a system is lost after a severe failure occurs, both configuration for the system and application data is lost. You must reinstate the system 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 system between these two phases.

If you do not understand the instructions to run the CLI commands, see the command-line interface reference information.

To restore your configuration data, follow these steps:

Procedure

1. Verify that all nodes are available as candidate nodes before you run this recovery procedure. You must remove errors 550 or 578 to put the node in candidate state. For all nodes that display these errors, perform the following steps:
 - a. Point your browser to the service IP address of one of the nodes, for example, `https://node_service_ip_address/service/`.
 - b. Log on to the service assistant.
 - c. Select **Home**.
 - d. If a node is not in service state already, select it from the table of nodes.
 - e. Select **Enter service state** from the Actions list, then select **Go** to put the node in to service state.
 - f. Select **Manage System**.
 - g. Click **Remove System Data**.
 - h. Confirm that you want to remove the system data when prompted.
 - i. Exit service state from the **Home** page. The 550 or 578 errors are removed and the node appears as a candidate node.

- j. Remove the system data for the other nodes that display a 550 or a 578 error. All nodes previously in this system must have a node status of Candidate and have no errors listed against them.

Note: A node that is powered off might not show up in this list of nodes for the system. Diagnose hardware problems directly on the node using the service assistant IP address and by physically verifying the LEDs for the hardware components.

2. Verify that all nodes are available as candidate nodes with blank system fields. Perform the following steps on one node in each control enclosure:
 - a. Connect to the service assistant on either of the nodes in the control enclosure.
 - b. Select **Configure Enclosure**.
 - c. Select the **Reset the system ID** option. Do not make any other changes on the panel.
 - d. Click **Modify** to make the changes.
3. Open the CMM and navigate to the chassis map. In the chassis map displayed by the CMM, you will see **Flex System V7000 Storage Node**.
 - a. Select the **Actions menu**. The Action menu can be found by left clicking on the left canister in the graphical representation of the storage node.
 - b. From the **Actions for Canister [node nn]** select **Launch Storage Node Controller Console**. A browser window displays the IBM Flex System V7000 system startup screen. The Service IP address displayed will be one that has been allocated by your DHCP server.
 - c. When asked **Do you want to create a new system or add to an existing system?**, select **Create a new system**. The **Create a new system** screen is displayed.
 - d. Select whether it is an **IPv4** or **IPv6** address and type in the IP address. The subnet mask and gateway will already have defaults listed. You can edit these if you wish.
 - e. Click **Finish** to set the management IP address for the system. System initialization begins and might take several minutes to complete.
When system initialization is complete, System Setup is launched automatically. The setup wizard will take you through the steps to configure basic system settings.
4. In a supported browser, enter the IP address that you used to initialize the system and the default superuser password (passwd).
5. At this point the setup wizard is shown. Be aware of the following items:
 - a. Accept the license agreements.
 - b. Set the values for the system name, date and time settings, and the system licensing. The original settings are restored during the configuration restore process.
 - c. Verify the hardware. Only the control enclosure on which the clustered system was created and directly attached expansion enclosures are displayed. Any other control enclosures and expansion enclosures in other I/O groups will be added to the system.
 - d. On the **Configure Storage** panel, deselect **Yes, automatically configure internal storage now**. Any internal storage configuration is recovered after the system is restored.
6. Add the rest of the control enclosures into the clustered system.
 - a. From the management GUI, select **Monitoring > System Details**.

- b. Select the system name in the tree.
- c. Go to **Actions > Add Enclosures > Control and Expansions**
- d. Continue to follow the on-screen instructions to add the control enclosures. Decline the offer to configure storage for the new enclosures when asked if you want to do so.

7. Identify the configuration backup file that you want to restore from.

The file can be either a local copy of the configuration backup XML file that you saved when backing up the configuration or an up-to-date file on one of the nodes.

Configuration data is automatically backed up daily at 01:00 system time on the configuration node.

Download and check the configuration backup files on all nodes that were previously in the system to identify the one containing the most recent complete backup

For each node in the system:

- a. From the management GUI, click **Settings > Support**.
- b. Click **Show full log listing**.
- c. Select the node to operate on from the selection box at the top of the table.
- d. Find the file name that begins with `svc.config.cron.xml`.
- e. Double-click the file to download the file to your computer.

The XML files contain a date and time that can be used to identify the most recent backup. After you identify the backup XML file that is to be used when you restore the system, rename the file to `svc.config.backup.xml`.

8. Issue the following CLI command to remove all of the existing backup and restore configuration files that are located on your configuration node in the `/tmp` directory: **svconfig clear -all**
9. Copy the XML backup file that you wish to restore from back onto the system.

```
pscp full_path_to_identified_svc.config.backup.xml
superuser@cluster_ip:/tmp/
```

10. Issue the following CLI command to compare the current 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 CMMVC6200W 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.

11. Issue the following command to copy the log file to another server that is accessible to the system:

```
pscp superuser@cluster_ip:/tmp/svc.config.restore.prepare.log
full_path_for_where_to_copy_log_files
```

12. Open the log file from the server where the copy is now stored.

13. 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 14 on page 120.

- If you need assistance, contact the IBM Support Center.
14. Issue the following CLI command to restore the configuration:

```
svconfig restore -execute
```

This CLI command creates a log file in the /tmp directory of the configuration node. The name of the log file is `svc.config.restore.execute.log`.
 15. Issue the following command to copy the log file to another server that is accessible to the system:

```
pscp superuser@cluster_ip:/tmp/svc.config.restore.execute.log  
full_path_for_where_to_copy_log_files
```
 16. Open the log file from the server where the copy is now stored.
 17. Check the log file 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.

When you log into the CLI again over SSH, you see this output:

```
IBM_2076:your_cluster_name:superuser>
```

What to do next

You can remove any unwanted configuration backup and restore files from the /tmp directory on your configuration by issuing the following CLI command:

```
svconfig clear -all
```

Deleting backup configuration files using the CLI

You can use the command-line interface (CLI) to delete backup configuration files.

About this task

Perform the following steps to delete backup configuration files:

Procedure

1. Issue the following command to log on to the system:

```
plink -i ssh_private_key_file superuser@cluster_ip
```

where `ssh_private_key_file` is the name of the SSH private key file for the superuser and `cluster_ip` is the IP address or DNS name of the clustered system from which you want to delete the configuration.
2. Issue the following CLI command to erase all of the files that are stored in the /tmp directory:

```
svconfig clear -all
```

Chapter 11. 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 Call Home is enabled, critical faults generate a problem management record (PMR) that is sent directly to the appropriate IBM Support Center.

Understanding events

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

Error data

Events are classified as either alerts or messages:

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

Viewing the event log

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

About this task

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

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

Managing the event log

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

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

Describing the fields in the event log

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

Table 33 describes some of the fields that are available to assist you in diagnosing problems.

Table 33. 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

The SAN Volume Controller product can use Simple Network Management Protocol (SNMP) traps, syslog messages, emails and Call Homes notify you and the IBM(r) Remote Technical 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.

Only events recorded in the event log can be notified. Most CLI messages in response to some CLI commands are not recorded in the event log so do not cause an event notification.

Table 34 describes the levels of event notifications.

Table 34. Notification levels

Notification level	Description
Critical	<p>Critical 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 Flex System V7000 Storage Node. For example, the event that is being reported could indicate a loss of redundancy in the system, and it is possible that another failure could result in loss of access to data. The most typical reason that this type of notification is sent is because of a hardware failure, but some configuration errors or fabric errors also are included in this notification level. Critical 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 Flex System V7000 Storage Node. 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 level Critical.</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 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: Troubleshooting the node canister using the status LEDs” on page 76.

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

Understanding the error codes

Error codes are generated by the event-log analysis and system configuration code.

Error codes help you to identify the cause of a problem, the failing field-replaceable units (FRUs), and the service actions that might be needed to solve the problem.

Event IDs

The Flex System V7000 Storage Node software generates events, such as informational events and error events. An event ID or number is associated with the event and indicates the reason for the event.

Informational events provide information about the status of an operation. Informational events are recorded in the event log, and depending on the configuration, can be notified through email, SNMP, or syslog.

Error events are generated when a service action is required. An error event maps to an alert with an associated error code. Depending on the configuration, error events can be notified through email, SNMP, or syslog.

Informational events

The informational events provide information about the status of an operation.

Informational events are recorded in the event log and, depending on the configuration, can be notified through email, SNMP, or syslog.

Informational events can be either notification type I (information) or notification type W (warning). An informational event report of type (W) might require user attention. Table 35 provides a list of informational events, the notification type, and the reason for the event.

Table 35. Informational events

Event ID	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.

Table 35. Informational events (continued)

Event ID	Notification type	Description
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 (system).
980350	I	The node is now a functional member of the cluster (system).
980351	I	A noncritical hardware error occurred.
980352	I	Attempt to automatically recover offline node starting.
980370	I	Both nodes in the I/O group are available.
980371	I	One node in the I/O group is unavailable.
980372	W	Both nodes in the I/O group are unavailable.
980380	I	Maintenance mode was started.
980381	I	Maintenance mode has ended.
980392	I	Cluster (system) recovery completed.
980435	W	Failed to obtain directory listing from remote node.
980440	W	Failed to transfer file from remote node.
980445	I	The migration is complete.
980446	I	The secure delete is complete.
980501	W	The virtualization amount is close to the limit that is licensed.
980502	W	The FlashCopy feature is close to the limit that is licensed.
980503	W	The Metro Mirror or Global Mirror feature is close to the limit that is licensed.
980504	I	The limit was reached for the external virtualization feature.
980505	I	The limit was reached for the compression feature license.
981002	I	Fibre Channel discovery occurred; configuration changes are pending.
981003	I	Fibre Channel discovery occurred; configuration changes are complete.
981004	I	Fibre Channel discovery occurred; no configuration changes were detected.
981007	W	The managed disk is not on the preferred path.
981009	W	The initialization for the managed disk failed.
981014	W	The LUN discovery has failed. The cluster (system) has a connection to a device through this node but this node cannot discover the unmanaged or managed disk that is associated with this LUN.

Table 35. Informational events (continued)

Event ID	Notification type	Description
981015	W	The LUN capacity equals or exceeds the maximum. Only part of the disk can be accessed.
981020	W	The managed disk error count warning threshold has been met.
981022	I	Managed disk offline imminent, offline prevention started
981025	I	Drive firmware download started
981026	I	Drive FPGA download started
981101	I	SAS discovery occurred; no configuration changes were detected.
981102	I	SAS discovery occurred; configuration changes are pending.
981103	I	SAS discovery occurred; configuration changes are complete.
981104	W	The LUN capacity equals or exceeds the maximum capacity. Only the first 1 PB of disk will be accessed.
981105	I	The drive format has started.
981106	I	The drive recovery was started.
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.
984503	I	The battery conditioning completed.
984504	I	The battery conditioning started.

Table 35. Informational events (continued)

Event ID	Notification type	Description
984505	I	The statesave information for the enclosure was collected.
984506	I	The debug from an IERR was extracted to disk.
984507	I	An attempt was made to power on the slots.
984508	I	All the expanders on the strand were reset.
984509	I	The component firmware update paused to allow the battery charging to finish.
984511	I	The update for the component firmware paused because the system was put into maintenance mode.
984512	I	A component firmware update is needed but is prevented from running.
985001	I	The Metro Mirror or Global Mirror background copy is complete.
985002	I	The Metro Mirror or Global Mirror is ready to restart.
985003	W	Unable to find path to disk in the remote cluster (system) within the timeout period.
986001	W	The thin-provisioned volume copy data in a node is pinned.
986002	I	All thin-provisioned volume copy data in a node is unpinned.
986010	I	The thin-provisioned volume copy import has failed and the new volume is offline; either upgrade the Flex System V7000 Storage Node software to the required version or delete the volume.
986011	I	The thin-provisioned volume copy import is successful.
986020	W	A thin-provisioned volume copy space warning has occurred.
986030	I	A thin-provisioned volume copy repair has started.
986031	I	A thin-provisioned volume copy repair is successful.
986032	I	A thin-provisioned volume copy validation is started.
986033	I	A thin-provisioned volume copy validation is successful.
986034	I	The import of the compressed-virtual volume copy was successful.
986035	W	A compressed-virtual volume copy space warning has occurred.
986036	I	A compressed-virtual volume copy repair has started.
986037	I	A compressed-virtual volume copy repair is successful.
986038	I	A compressed-virtual volume copy has too many bad blocks.
986201	I	A medium error has been repaired for the mirrored copy.

Table 35. Informational events (continued)

Event ID	Notification type	Description
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 (system) has been lost.
987400	W	The node unexpectedly lost power but has now been restored to the cluster (system).
988100	W	An overnight maintenance procedure has failed to complete. Resolve any hardware and configuration problems that you are experiencing on the cluster (system). If the problem persists, contact your IBM service representative for assistance.
988300	W	An array MDisk is offline because it has too many missing members.
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.
988306	I	A RAID array needs resynchronization.
989001	W	A managed disk group space warning has occurred.

Error event IDs and error codes

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

Error codes can be either notification type E (error) or notification type W (warning). Table 36 on page 129 lists the event IDs and corresponding error codes, the notification type, and the condition of the event.

Table 36. Error event IDs and error codes

Event ID	Notification type	Condition	Error code
009020	E	An automatic system recovery has started. All configuration commands are blocked.	1001
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 system. 	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
009150	W	Unable to connect to the SMTP (email) server	2600
009151	W	Unable to send mail through the SMTP (email) server	2601
009170	W	The Metro Mirror or Global Mirror feature capacity is not set.	3030
009171	W	The FlashCopy feature capacity is not set.	3031
009172	W	The Virtualization feature has exceeded the amount that is licensed.	3032
009173	W	The FlashCopy feature has exceeded the amount that is licensed.	3032
009174	W	The Metro Mirror or Global Mirror feature has exceeded the amount that is licensed.	3032
009175	W	The usage for the thin-provisioned volume is not licensed.	3033
009176	W	The value set for the virtualization feature capacity is not valid.	3029
009177	E	A physical disk FlashCopy feature license is required.	3035
009178	E	A physical disk Metro Mirror and Global Mirror feature license is required.	3036
009179	E	A virtualization feature license is required.	3025
009180	E	Automatic recovery of offline node failed.	1194
009181	W	Unable to send email to any of the configured email servers.	3081
009182	W	The external virtualization feature license limit was exceeded.	3032
009183	W	Unable to connect to LDAP server.	2251
009184	W	The LDAP configuration is not valid.	2250
009185	E	The limit for the compression feature license was exceeded.	3032
009186	E	The limit for the compression feature license was exceeded.	3032
010002	E	The node ran out of base event sources. As a result, the node has stopped and exited the system.	2030
010003	W	The number of device logins has reduced.	1630
010006	E	A software error has occurred.	2030

Table 36. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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
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 (system).	1200
010022	W	There are too many managed disks presented to the cluster (system).	1200
010023	W	There are too many LUNs presented to a node.	1200
010024	W	There are too many drives presented to a cluster (system).	1200
010025	W	A disk I/O medium error has occurred.	1320
010026	W	A suitable MDisk or drive for use as a quorum disk was not found.	1330
010027	W	The quorum disk is not available.	1335
010028	W	A controller configuration is not supported.	1625
010029	E	A login transport fault has occurred.	1360
010030	E	A managed disk error recovery procedure (ERP) has occurred. The node or controller reported the following: <ul style="list-style-type: none"> • Sense • Key • Code • Qualifier 	1370
010031	E	One or more MDisks on a controller are degraded.	1623
010032	W	The controller configuration limits failover.	1625
010033	E	The controller configuration uses the RDAC mode; this is not supported.	1624
010034	E	Persistent unsupported controller configuration.	1695
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 (system) nodes through a single target port.	1627

Table 36. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
010043	E	The controller system device is only connected to the cluster (system) nodes through half of the expected target ports.	1627
010044	E	The controller system device has disconnected all target ports to the cluster (system) 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	The drive initialization failed because of an unknown block size or a block size that is not valid; an unknown capacity or a capacity that is not valid; or was not able to set the required mode pages.	1661
010059	E	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 (system).	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 (system).	1200
010071	W	Too many target ports were presented to the cluster (system) from a single controller.	1200
010072	E	The drive is offline as a result of a drive hardware error.	1680
010073	E	The drive is reporting predictive failure analysis (PFA) errors.	1680
010080	E	The drive is reporting too many errors.	1680
010081	E	The drive format is corrupted.	1206
010082	E	The block size for the drive was incorrect.	1206

Table 36. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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
010095	E	Drive failed.	1687
010096	E	The drive initialization failed because of an unknown block size or a block size that is not valid; an unknown capacity or a capacity that is not valid; or was not able to set the required mode pages.	1680
010097	E	A drive is reporting excessive errors.	1685
010098	W	There are too many drives presented to a cluster (system).	1200
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
029001	W	The managed disk has bad blocks. On an external controller, this can only be a copied medium error.	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 the clustered system 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 clustered system 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

Table 36. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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
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	W	The firmware update for the enclosure component has failed.	3015

Table 36. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
045041	W	More than one initiator port was detected on the same strand.	1005
045042	W	The order of the enclosures is different on each strand.	1005
045044	W	Multiple canisters are connected to a single canister port.	1005
045045	W	Canister 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
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 required but not possible.	1131
045063	E	There was an enclosure battery communications error.	1116
045064	W	A SAS port is active, but no enclosures can be detected.	1005
045065	E	There is a connectivity problem between a canister and an enclosure.	1036
045066	E	The FRU identity of the enclosure is not valid.	1008
045067	W	A new enclosure FRU was detected and needs to be configured.	1041
045068	E	The internal device on a node canister was excluded because of too many change events.	1034
045069	E	The internal connector on the node canister was excluded as the cause of single ported drives.	1034
045070	W	The canister temperature sensor cannot be read.	1034
045071	W	The enclosure contains both a node canister and an expansion canister.	1037
045072	E	The discovery failed to complete.	1048
045073	E	The VPD for the enclosure cannot be read.	1048
045080	E	There are too many self-initiated resets on the enclosure.	1048
045082	E	The slots are powered off.	1048

Table 36. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
045087	W	The expansion or control enclosure temperature is at the critical level.	
050001	W	The relationship is stopped because of a clustered system 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 clustered system, but its partnership has been deleted.	3080
050010	W	A Global Mirror relationship has stopped because of a persistent I/O error.	1920
050011	W	A remote copy has stopped because of a persistent I/O error.	1915
050020	W	Remote copy has stopped.	1720
050030	W	There are too many cluster (system) partnerships. The number of partnerships has been reduced.	1710
050031	W	There are too many cluster (system) partnerships. The system has been excluded.	1710
050040	W	Background copy process for the Remote Copy was blocked.	1960
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
060004	W	The compressed volume copy is offline because there is insufficient space.	1865
060005	W	The compressed volume copy is offline because the metadata is corrupt.	1862
060006	W	The compressed volume copy is offline because the repair has failed.	1860
060007	W	The compressed volume copy has bad blocks.	1850
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
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

Table 36. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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 (system) due to lack of 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
070574	E	The node software is damaged.	1187
070576	E	The cluster (system) data cannot be read.	1030
070578	E	The cluster (system) data was not saved when power was lost.	1194
070580	E	Unable to read the service controller ID.	1044
070690	W	Node held in service state.	1189
071820	W	Node canister has the incorrect model for the enclosure.	3020
071840	W	Detected hardware is not a valid configuration.	1198
071841	W	Detected hardware needs activation.	1199
072528	W	The control canister temperature is at the critical level.	528
072653	W	The control canister battery is too cold, and at the critical level.	653
072654	W	The control canister battery is too hot, and at the critical level.	654
072768	W	The ambient temperature of the canister is close to the point where it stops performing I/O and enters a service state. The canister is currently continuing to operate.	768
072769	E	The temperature of the CPU in the canister is close to the point where the canister stops performing I/O and enters service state.	769
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 (system) 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 Flex System V7000 Storage Node port.	1800
073007	W	There are fewer Fibre Channel ports operational than are configured.	1061
073305	W	One or more Fibre Channel ports are running at a speed that is lower than the last saved speed.	1065

Table 36. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
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
073500	W	Incorrect enclosure.	1021
073501	E	Incorrect canister position.	1192
073502	E	No enclosure identity; cannot get status from partner.	1192
073503	E	Incorrect enclosure type.	1192
073504	E	No enclosure identity and partner does match.	1192
073505	E	No enclosure identity and partner does not match.	1192
073506	E	No enclosure identity and no state on partner.	1192
073507	E	No enclosure identity and no node state.	1192
073508	W	Cluster (system) identity is different on the enclosure and the node.	1023
073509	E	Cannot read enclosure identity.	1036
073510	E	Detected memory size does not match the expected memory size.	1032
073512	E	Enclosure VPD is inconsistent	1008
073522	E	The system board service processor has failed.	1034
073523	W	The internal disk file system of the node is damaged.	1187
073525	E	Unable to update the service processor firmware of the system board.	1034
073528	E	Ambient temperature is too high during system startup.	1098
073535	E	The internal PCIe switch of the node canister failed.	1034
073550	E	Cannot form cluster (system) due to lack of resources.	1192
073556	E	Duplicate WWNN detected on the SAN.	1133
073562	E	The node hardware does not meet the minimum requirements.	1034
073564	W	Too many software failures.	1188
073565	E	The internal drive of the node is failing.	1032
073573	E	The node software is inconsistent.	1187
073574	E	The cluster (system) data cannot be read.	1187
073578	E	The cluster (system) data was not saved when power was lost.	1194
073651	E	The canister battery is missing.	1153
073652	E	The canister battery has failed.	1154
073653	E	The canister battery's temperature is too low.	1156
073654	E	The canister battery's temperature is too high	1157
073655	E	The canister battery communications fault.	1158
073656	E	The canister battery has insufficient charge.	1184
073690	W	Node held in service state.	1189
073700	E	FC adapter missing.	1045

Table 36. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
073701	E	FC adapter failed.	1046
073702	E	FC adapter PCI error.	1046
073703	E	FC adapter degraded.	1045
073704	W	Fewer Fibre Channel ports operational.	1061
073705	W	Fewer Fibre Channel IO ports operational.	1450
073710	E	SAS adapter missing.	1045
073711	E	SAS adapter failed.	1046
073712	E	SAS adapter PCI error.	1046
073713	E	SAS adapter degraded.	1046
073720	E	Ethernet adapter missing.	1045
073721	E	Ethernet adapter failed.	1046
073722	E	Ethernet adapter PCI error.	1046
073723	E	Ethernet adapter degraded.	1046
073724	W	Fewer Ethernet ports operational.	1401
073730	E	Bus adapter missing.	1032
073731	E	Bus adapter failed.	1032
073732	E	Bus adapter PCI error.	1032
073733	E	Bus adapter degraded.	1032
073734	W	Inter-canister PCIe link failure.	1006
073768	W	Ambient temperature warning.	1094
073769	W	CPU temperature warning.	1093
073840	E	Detected hardware is not a valid configuration.	1198
073841	E	Detected hardware needs activation.	1199
073860	W	Fabric too large.	1800
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 (system) 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
074003	W	A connection to a configured remote system has been lost because of a connectivity problem.	1715
074004	W	A connection to a configured remote system has been lost because of too many minor errors.	1716
076001	E	The internal disk for a node has failed.	1030
076002	E	The hard disk is full and cannot capture any more output.	2030
076401	E	One of the two power supply units in the node has failed.	1096
076402	E	One of the two power supply units in the node cannot be detected.	1096

Table 36. Error event IDs and error codes (continued)

Event ID	Notification type	Condition	Error code
076403	E	One of the two power supply units in the node is without power.	1097
076502	E	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
076514	E	A high-speed SAS adapter requires a PCI bus reset.	1034
079500	W	The limit on the number of cluster (system) secure shell (SSH) sessions has been reached.	2500
079501	W	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
084201	W	An array has members that are located in a different I/O group.	1688
084300	W	An array MDisk is no longer protected by an appropriate number of suitable spares.	1690
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 clustered system until the issue that is preventing it from joining a clustered system is resolved. This

error occurs because part of the hardware has failed or the system detects that the code 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 clustered system. It shows a critical node error but is in the starting state. The range of errors that are reserved for critical errors are 500 - 699.

Noncritical errors

A noncritical error code is logged when there is a hardware or code failure that is related to just one specific node. These errors do not stop the node from entering active state and joining a clustered system. If the node is part of a clustered system, 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.

Clustered-system code overview

Recovery codes for clustered systems indicate that a critical software error has occurred that might corrupt your system. Each error-code topic includes an error code number, a description, action, and possible field-replaceable units (FRUs).

Error codes for recovering a clustered system

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

Error code range

This topic shows the number range for each message classification.

Table 37 lists the number range for each message classification.

Table 37. Message classification number range

Message classification	Range	
Node errors	Critical node errors	500-699
	Noncritical node errors	800-899
Error codes when recovering a clustered system	920, 990	

500 Incorrect enclosure

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

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

1. Follow the procedure: Getting node canister and system information and review the node canister's saved location information and the status of the other node canister in the enclosure (the partner

canister). Determine if the enclosure is part of an active system with volumes that contain required data. See "Procedure: Getting node canister and system information using the service assistant" on page 75.

2. If you have unintentionally moved the canister into this enclosure, move the canister back to its original location, and put the original canister back in this enclosure. Follow the Replacing a node canister procedure.
3. If you have intentionally moved the node canister into this enclosure you should check it is safe to continue or whether you will lose data on the

enclosure you removed it from. Do not continue if the system the node canister was removed from is offline, rather return the node canister to that system.

4. If you have determined it is alright to continue, follow the procedure to remove cluster data from node canister. See “Procedure: Removing system data from a node canister” on page 86.
5. If the partner node in this enclosure is not online, or is not present, you will have to perform a system recovery. Do not create a new system, you will lose all the volume data.

Possible Cause-FRUs or other cause:

- None

501 Incorrect slot

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

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

1. Follow the procedure: Getting node canister and system information and 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. See “Procedure: Getting node canister and system information using the service assistant” on page 75.
2. If the canisters have been swapped, use the Replacing a node canister procedure to swap the canisters. The system should start.
3. If the partner canister is in candidate state, use the hardware remove and replace canister procedure to swap the canisters. The system should start.
4. If the partner canister is in active state, it is running the cluster on this enclosure and has replaced the original use of this canister. 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. See “Procedure: Removing system data from a node canister” on page 86.
5. If the partner canister is in service state, review its node error to determine the correct action. Generally, you will fix the errors reported on the partner node in priority order, and review the situation again after each change. If you have to replace the partner canister with a new one you should move this canister back to the correct location at the same time.

Possible Cause-FRUs or other:

- None

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

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

User response: Follow troubleshooting procedures to configure the enclosure:

1. Follow the procedures to resolve a problem to get the partner node started. An error will still exist because the enclosure has no identity. If the error has changed, follow the service procedure for that error.
2. If the partner has started and is showing a location error (probably this one), then the PCI link is probably broken. Since the enclosure midplane was recently replaced, this is likely the problem. Obtain a replacement enclosure midplane, and replace it following the remove and replace control enclosure midplane procedure.
3. 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. Also see Chapter 7, “Resolving a problem,” on page 63.

Possible Cause—FRUs or other:

- Enclosure midplane (100%)

504 No enclosure identity and partner node matches.

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

User response: Follow troubleshooting procedures to configure the enclosure.

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

Possible Cause—FRUs or other:

- None
-

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

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

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

1. Follow the procedure: Getting node canister and system information and review the node canister's saved location information and the status of the other node canister in the enclosure (the partner canister). Determine if the enclosure is part of an active system with volumes that contain required data. See "Procedure: Getting node canister and system information using the service assistant" on page 75.
2. Decide what to do with the node canister that did not come from the enclosure that is being replaced.
 - a. If the other node canister from the enclosure being replaced is available, use the hardware remove and replace canister procedures to remove the incorrect canister and replace it with the second node canister from the enclosure being replaced. Restart both canisters. The two node canister should show node error 504 and the actions for that error should be followed.
 - b. If the other node canister from the enclosure being replaced is not available, check the enclosure of the node canister that did not come from the replaced enclosure. Do not use this canister in this enclosure if you require the volume data on the system from which the node canister was removed, and that system is not running with two online nodes. You should return the canister to its original enclosure and use a different canister in this enclosure.
 - c. When you have checked that it is not required elsewhere, follow the procedure to remove cluster data from the node canister that did not come from the enclosure that is being replaced. See "Procedure: Removing system data from a node canister" on page 86. 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 midplane has been replaced. There is no cluster state information on the other node canister in the enclosure (the partner canister), so both node canisters from the original

enclosure have not been moved to this one.

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

1. Follow the procedure: Getting node canister and system information and review the saved location information of the node canister and determine why the second node canister from the original enclosure was not moved into this enclosure. See "Procedure: Getting node canister and system information using the service assistant" on page 75.
2. If you are sure that this node canister came from the enclosure that is being replaced, and the original partner canister is available, use the Replacing a node canister procedure to install the second node canister in this enclosure. Restart the node canister. The two node canisters should show node error 504, and the actions for that error should be followed.
3. If you are sure this node canister came from the enclosure that is being replaced, and that the original partner canister has failed, continue following the remove and replace procedure for an enclosure midplane and configure the new enclosure.

Possible Cause—FRUs or other:

- None

507 No enclosure identity and no node state

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

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

1. Check the status of the other node in the enclosure. It should show node error 506. Unless it also shows error 507, check the errors on the other node and follow the corresponding procedures to resolve the errors.
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 correctly relocate the nodes.

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. Follow the procedure: Get node canister and system information using the service assistant, and check the partner node canister to see if it is also reporting node error 508; if it is, check that the saved system information on this and the partner node match. See “Procedure: Getting node canister and system information using the service assistant” on page 75. If the system information on both nodes matches, follow the Replacing the control enclosure midplane procedure to change the enclosure midplane.
3. If this node canister is the one to be used in this enclosure, follow “Procedure: Removing system data from a node canister” on page 86 to remove cluster data from the node canister. It will then join the cluster.
4. If this is not the node canister that you intended to use, follow the Replacing a node canister procedure to replace the node canister with the one intended for use.

Possible Cause—FRUs or other:

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

509 The enclosure identity cannot be read.

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

User response: Follow troubleshooting procedures to fix the hardware:

1. Check errors reported on the other node canister in this enclosure (the partner canister).
2. If it is reporting the same error follow the hardware remove and replace procedure to replace the enclosure midplane.
3. If the partner canister is not reporting this error, follow the hardware remove and replace procedure to replace this canister.

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

Possible Cause—FRUs or other:

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

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

Explanation: The amount of memory detected in the node canister is less than the amount required for the canister to operate as an active member of a system. The error code data shows the detected memory, in MB, followed by the minimum required memory, in MB, there is then a series of values indicating the amount of memory, in GB, detected in each memory slot.

Data:

- Detected memory on MB
- Minimum required memory in MB
- Memory in slot 1 in GB
- Memory in slot 2 in GB
- ... etc.

User response: Follow troubleshooting procedures to fix the hardware:

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

Possible Cause—FRUs or other:

- Node canister (100%)

522 The system board service processor has failed.

Explanation: The service processor (IMM) in the canister has failed or is taking a long time resetting. If the IMM recovers the error will clear itself and, if possible, the canister will become active.

User response:

1. Wait 30 minutes to see if the canister becomes active.
2. Reseat the node canister.
3. If the error persists, use the remove and replace procedures to replace the node canister.

Possible Cause—FRUs or other:

- node canister

523 The internal disk file system is damaged.

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

User response: Follow troubleshooting procedures to reload the software.

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

Possible Cause—FRUs or other:

- Node canister (100%)

525 Unable to update system board service processor firmware.

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

User response: Follow troubleshooting procedures to fix the hardware:

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

Possible Cause—FRUs or other:

- Node canister (100%)

528 Ambient temperature is too high during system startup.

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

User response: Reduce the temperature around the system.

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

Possible Cause—FRUs or other:

- Environment issue (100%)

535 Canister internal PCIe switch failed

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

User response: Follow troubleshooting procedures to fix the hardware:

1. Follow Procedure: Reseat the node canister. See Procedure: Reseating a node canister.
2. If reseating the canister does not resolve the situation, follow the Replacing a node canister procedure to replace the canister.

Possible Cause—FRUs or other:

- Node canister (100%)

541 Multiple, undetermined, hardware errors

Explanation: Multiple hardware failures have been reported on the data paths within the node canister, and the threshold of the number of acceptable errors within a given time frame has been reached. It has not been possible to isolate the errors to a single component.

After this node error has been raised, all ports on the node will be deactivated. The reason for this is that the node canister is considered unstable, and has the potential to corrupt data.

User response:

1. Follow the procedure for collecting information for support, and contact your support organization.
2. A software [code] upgrade may resolve the issue.
3. Replace the node canister.

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

Explanation: The node canister 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 system enclosures is the missing quorum disk, it is listed as enclosure:slot[part identification] where enclosure:slot is the location of the drive when the node 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 WWWWWWWWWWWWWWWW/LL, where WWWWWWWWWWWWWWWW is a worldwide port name (WWPN) on the storage system that contains the missing quorum disk and LL is the Logical Unit Number (LUN).

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

1. Check for any node errors that indicate issues with bus or Fibre Channel connectivity. Resolve any issues.
2. Check the status of other node canisters in the system, resolve any faults on them.
3. 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.
4. 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, because 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.

5. 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.
6. 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 system 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 canister and resolve any issues.
7. If all canisters 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 Flex System V7000 Storage Node, 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 Flex System V7000 Storage Node node canister.

User response:

1. Find the Flex System V7000 Storage Node node canister with the same WWNN as the node canister reporting the error. The WWNN for a Flex System V7000 Storage Node 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 Flex System V7000 Storage Node 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 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 Flex System V7000 Storage Node 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 Flex System V7000 Storage Node, so you should follow the service procedures for the device to change its WWNN. Once the duplicate has been removed, restart the node canister.

562 The nodes hardware configuration does not meet the minimum requirements

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

User response: Follow troubleshooting procedures to fix the hardware:

1. View node VPD information, to see whether anything looks inconsistent. Compare the failing node VPD with the VPD of a working node canister of the same type. Pay particular attention to the number and type of CPUs, memory, and adapters.
2. If the canister has just had parts replaced or added, check the parts are the correct ones for use in the system.
3. Replace any incorrect parts.
4. A code upgrade may be required to support the new hardware.

Possible Cause—FRUs or other:

- Node canister (100%)
-

564 Too many machine code crashes have occurred.

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

The node error does not persist across restarts of the node machine code.

User response: Follow troubleshooting procedures to reload the machine code:

1. Get a support package (snap), including dumps, from the node, using the management GUI or the service assistant.
2. If more than one node is reporting this error, contact IBM technical support for assistance. The support package from each node will be required.
3. Check the support site to see whether the issue is known and whether a machine code upgrade exists to resolve the issue. Update the cluster machine code 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 instructions to change the node canister.

Possible Cause—FRUs or other:

- Node canister (100%)
-

573 The node machine code is inconsistent.

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

User response: Follow troubleshooting procedures to reload the machine code.

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

Possible Cause—FRUs or other:

- None
-

574 The node machine code is damaged.

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

User response:

1. If the other nodes are operational, run node rescue; otherwise, install new machine code using the service assistant. Node rescue failures, as well as the repeated return of this node error after reinstallation, are symptomatic of a hardware fault with the node.

Possible Cause—FRUs or other:

- None
-

576 The cluster state and configuration data cannot be read.

Explanation: The node 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.
 - a. To remove the cluster data from the node, either go to the service assistant, select the radio button for the node with a 578, click **Manage System**, then choose **Remove System Data**.
 - b. Or use the CLI to **satask leavecluster**.

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

650 The canister battery is not supported

Explanation: The canister battery shows product data that indicates it cannot be used with the code version of the canister.

User response: This is resolved by either obtaining a battery which is supported by the system's code level, or the canister's code level is updated to a level which supports the battery.

1. Remove the canister and its lid and check the FRU part number of the new battery matches that of the replaced battery. Obtain the correct FRU part if it does not.
2. If the canister has just been replaced, check the code level of the partner node canister and use the service assistant to upgrade this canister's code level to the same level.

Possible cause—FRUs or other cause

- canister battery

651 The canister battery is missing

Explanation: The canister battery cannot be detected.

User response:

1. Use the remove and replace procedures to remove the node canister and its lid.
2. Use the remove and replace procedures to install a battery.
3. If there is a battery present ensure it is fully inserted. Replace the canister.
4. If this error persists, use the remove and replace procedures to replace the battery.

Possible cause—FRUs or other cause

- canister battery

652 The canister battery has failed

Explanation: The canister battery has failed. The battery may be showing an error state, it may have reached the end of life, or it may have failed to charge.

Data

Number indicators with failure reasons

- 1—battery reports a failure
- 2—end of life
- 3—failure to charge

User response:

1. Use the remove and replace procedures to replace the battery.

Possible cause—FRUs or other cause

- canister battery

653 The canister battery's temperature is too low

Explanation: The canister battery's temperature is below its minimum operating temperature.

User response:

- Wait for the battery to warm up, the error will clear when its minimum working temperature is reached.

- If the error persists for more than an hour when the ambient temperature is normal, use the remove and replace procedures to replace the battery.

Possible cause—FRUs or other cause

- canister battery

654 The canister battery's temperature is too high

Explanation: The canister battery's temperature is above its safe operating temperature.

User response:

- If necessary, reduce the ambient temperature.
- Wait for the battery to cool down, the error will clear when normal working temperature is reached. Keep checking the reported error as the system may determine the battery has failed.
- If the node error persists for more than two hours after the ambient temperature returns to the normal operating range, use the remove and replace procedures to replace the battery.

Possible cause—FRUs or other cause

- canister battery

655 The canister battery communications fault

Explanation: The canister cannot communicate with the battery.

User response:

- Use the remove and replace procedures to replace the battery.
- If the node error persists, use the remove and replace procedures to replace the node canister.

Possible cause—FRUs or other cause

- canister battery
- node canister

656 The canister battery has insufficient charge

Explanation: The canister battery has insufficient charge to save the canister's state and cache data to the internal drive if power were to fail.

User response:

- Wait for the battery to charge, the battery does not need to be fully charged for the error to automatically clear.

Possible cause—FRUs or other cause

- none

690 The node is held in the service state.

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

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

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

Possible Cause—FRUs or other:

- none

700 The Fibre Channel adapter that was previously present has not been detected.

Explanation: A Fibre Channel adapter that was previously present has not been detected. For Flex System V7000 Storage Node, the adapter is located on the node canister system board.

This node error does not, in itself, stop the node canister from becoming active in the system; however, the Fibre Channel network might be being used to communicate between the node canisters in a clustered system. It is possible that this node error indicates why the critical node error 550 A cluster cannot be formed because of a lack of cluster resources is reported on the node canister.

Data:

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

User response:

1. If possible, this noncritical node error should be serviced using the management GUI and running the recommended actions for the service error code.
- 2.

There are a number of possibilities.

- a. If you have deliberately removed the adapter (possibly replacing it with a different adapter type), you will need to follow the management GUI recommended actions to mark the hardware change as intentional.
- b. Use the remove and replace procedures to remove and open the canister and check the adapter is properly installed, particularly check the ribbon cable which connects the adapter to the system board is properly connected.

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

Possible Cause—FRUs or other cause:

- Fibre Channel adapter
- Node canister

701 A Fibre Channel adapter has failed.

Explanation: A Fibre Channel adapter has failed.

This node error does not, in itself, stop the node canister becoming active in the system. However, the Fibre Channel network might be being used to communicate between the node canisters in a clustered system. Therefore, it is possible that this node error indicates the reason why the critical node error 550 A cluster cannot be formed because of a lack of cluster resources is reported on the node canister.

Data:

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

User response:

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

Possible Cause-FRUs or other cause:

- Fibre Channel adapter
- Node canister

702 A Fibre Channel adapter has a PCI error.

Explanation: A Fibre Channel adapter has a PCI error.

This node error does not, in itself, stop the node canister becoming active in the system. However, the Fibre Channel network might be being used to communicate between the node canisters in a clustered system. Therefore, it is possible that this node error indicates the reason why the critical node error 550 A cluster cannot be formed because of a lack of cluster resources is reported on the node canister.

Data:

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

User response:

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

Possible Cause-FRUs or other cause:

- Fibre Channel adapter
- Node canister

703 A Fibre Channel adapter is degraded.

Explanation: A Fibre Channel adapter is degraded.

This node error does not, in itself, stop the node canister becoming active in the system. However, the Fibre Channel network might be being used to communicate between the node canisters in a clustered system. Therefore, it is possible that this node error indicates the reason why the critical node error 550 A cluster cannot be formed because of a lack of cluster resources is reported on the node canister.

Data:

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

User response:

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

Possible Cause-FRUs or other cause:

- Fibre Channel adapter
- Node canister

704 Fewer Fibre Channel ports operational.

Explanation: A Fibre Channel port that was previously operational is no longer operational. The physical link is down. The physical connection to the Fibre Channel switch is part of the IBM Flex System chassis backplane.

This node error does not, in itself, stop the node canister becoming active in the system. However, the Fibre Channel network might be being used to communicate between the node canisters in a clustered system. Therefore, it is possible that this node error indicates the reason why the critical node error 550 A cluster cannot be formed because of a lack of cluster resources is reported on the node canister.

Data:

Three numeric values are listed:

- The ID of the first unexpected inactive port. This ID is a decimal number.
- The ports that are expected to be active, which is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is expected to be active.
- The ports that are actually active, which is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is active.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Possibilities:
 - If the port has been intentionally disconnected, use the management GUI recommended action for the service error code and acknowledge the intended change.
 - Check the Fibre Channel switch the port is connected to is powered and enabled in a compatible mode. Rectify any issue. The switch service interface might indicate the issue.
 - Use the remove and replace procedures to replace the adapter.
 - Consider replacing the IBM Flex System chassis backplane.

Possible Cause-FRUs or other cause:

- Fibre Channel adapter

705 Fewer Fibre Channel I/O ports operational.

Explanation: One or more Fibre Channel I/O ports that have previously been active are now inactive. This situation has continued for one minute.

A Fibre Channel I/O port might be established on either a Fibre Channel platform port or an Ethernet platform port using FCoE. This error is expected if the associated Fibre Channel or Ethernet port is not operational.

Data:

Three numeric values are listed:

- The ID of the first unexpected inactive port. This ID is a decimal number.
- The ports that are expected to be active, which is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is expected to be active.
- The ports that are actually active, which is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is active.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Follow the procedure for mapping I/O ports to platform ports to determine which platform port is providing this I/O port.
3. Check for any 704 (Fibre channel platform port not operational) or 724 (Ethernet platform port not operational) node errors reported for the platform port.
4. Possibilities:
 - If the port has been intentionally disconnected, use the management GUI recommended action for the service error code and acknowledge the intended change.
 - Resolve the 704 or 724 error.
 - If this is an FCoE connection, use the information the view gives about the Fibre Channel forwarder (FCF) to troubleshoot the connection between the port and the FCF.

Possible Cause-FRUs or other cause:

- None

706 Fibre Channel clustered system path failure.

Explanation: This situation has continued for 5 minutes.

A Fibre Channel I/O port might be established on either a FC platform port or an Ethernet platform port using Fiber Channel over Ethernet (FCoE).

Data:

Three numeric values are listed:

- The ID of the first FC I/O port that does not have connectivity. This is a decimal number.
- The ports that are expected to have connections. This is a hexadecimal number, and each bit position represents a port - with the least significant bit representing port 1. The bit is 1 if the port is expected to have a connection to all online node canisters.
- The ports that actually have connections. This is a hexadecimal number, each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port has a connection to all online nodes.

User response:

1. If possible, this noncritical node canister error should be serviced using the management GUI and running the recommended actions for the service error code.

2. Follow the procedure: Mapping I/O ports to platform ports to determine which platform port does not have connectivity.
3. There are a number of possibilities.
 - If the port's connectivity has been intentionally reconfigured, use the management GUI recommended action for the service error code and acknowledge the intended change. You must have at least two I/O ports with connections to all other node canisters, except the node canisters in the same enclosure.
 - Resolve other node errors relating to this platform port or I/O port.
 - Check that the SAN zoning is correct.

710 The SAS adapter that was previously present has not been detected.

Explanation: A SAS adapter that was previously present has not been detected. The adapter is located on the node canister system board.

Data:

- A number indicating the adapter location. Location 0 indicates the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

- Node canister

711 A SAS adapter has failed.

Explanation: A SAS adapter has failed. The adapter is located on the node canister system board.

Data:

- A number indicating the adapter location. Location 0 indicates that the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

- Node canister

712 A SAS adapter has a PCI error.

Explanation: A SAS adapter has a PCI error.

Data:

- A number indicating the adapter location. Location 0 indicates the adapter that is integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Locate the adapter on the system board and replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

- Node canister

713 A SAS adapter is degraded.

Explanation: A SAS adapter is degraded. The adapter is located on the node canister system board.

Data:

- A number indicating the adapter location. Location 0 indicates that the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Use the remove and replace procedures to replace the node canister.

Possible Cause-FRUs or other cause:

- Node canister

720 Ethernet adapter that was previously present has not been detected.

Explanation: An Ethernet adapter that was previously present has not been detected. The adapter might have been removed deliberately, it might not be installed correctly, or it might have failed.

Data:

- A number indicating the adapter location. The location indicates an adapter slot. See the node canister description for the definition of the adapter slot locations. If the location is 0, the adapter integrated into the system board is being reported.

User response:

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

2. If the adapter location is 0, replace the node canister.
3. If the location is not 0, there are a number of possibilities:
 - a. If you deliberately removed the adapter (for example, replacing it with a different adapter type), use the management GUI recommended actions to mark the hardware change as intentional.
 - b. Use the remove and replace procedures to remove and open the canister and make sure the adapter is properly installed. In particular, make sure the ribbon cable that connects the adapter to the system board is properly connected.
 - c. If the previous steps have not isolated the problem, use the remove and replace procedures to replace the adapter. If this does not fix the problem, replace the canister.

Possible Cause—FRUs or other cause:

- Ethernet adapter
- Node canister

721 An Ethernet adapter has failed.

Explanation: An Ethernet adapter has failed.

Data:

•

A number indicating the adapter location. The location indicates an adapter slot. See the node canister description for the definition of the adapter slot locations. If the location is 0, the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. If the adapter location is 0, use the remove and replace procedures to replace the node canister.
3. If the location is not 0, use the remove and replace procedures to replace the adapter. If this does not fix the problem, replace the node canister.

Possible Cause—FRUs or other cause:

- Ethernet adapter
- Node canister

722 An Ethernet adapter has a PCI error.

Explanation: An Ethernet adapter has a PCI error.

Data:

•

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

slot locations. If the location is 0, the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. If the adapter location is 0, use the remove and replace procedures to replace the node canister.
3. If the location is not 0, use the remove and replace procedures to replace the adapter. If this does not fix the problem, replace the node canister.

Possible Cause—FRUs or other cause:

- Ethernet adapter
- Node canister

723 An Ethernet adapter is degraded.

Explanation: An Ethernet adapter is degraded.

Data:

•

A number indicating the adapter location. The location indicates an adapter slot. See the node canister description for the definition of the adapter slot locations. If the location is 0, the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. If the adapter location is 0, use the remove and replace procedures to replace the node canister.
3. If the location is not 0, use the remove and replace procedures to replace the adapter. If this does not fix the problem, replace the node canister.

Possible Cause—FRUs or other cause:

- Ethernet adapter
- Node canister

724 Fewer Ethernet ports active.

Explanation: An Ethernet port that was previously operational is no longer operational. The physical link is down. The physical connection to the Ethernet switch is part of the IBM Flex System chassis backplane.

Data:

Three numeric values are listed:

- The ID of the first unexpected inactive port. This is a decimal number.
- The ports that are expected to be active. This is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is expected to be active.

- The ports that are actually active. This is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is active.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Possibilities:
 - a. If the port has been intentionally disconnected, use the management GUI recommended action for the service error code and acknowledge the intended change.
 - b. Make sure the Ethernet switch the port is connected to is powered and enabled in a compatible mode. Rectify any issue. The switch service interface might indicate the issue.
 - c. Replace the adapter or node canister (depending on the port location) using the remove and replace procedures.
 - d. Consider replacing the Flex system chassis backplane.

Possible Cause—FRUs or other cause:

- Ethernet Adapter
- Node canister

730 The bus adapter has not been detected.

Explanation: The bus adapter that connects the canister to the enclosure midplane has not been detected.

This node error does not, in itself, stop the node canister becoming active in the system. However, the bus might be being used to communicate between the node canisters in a clustered system. Therefore, it is possible that this node error indicates the reason why the critical node error 550 A cluster cannot be formed because of a lack of cluster resources is reported on the node canister.

Data:

- A number indicating the adapter location. Location 0 indicates that the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

- Node canister

731 The bus adapter has failed.

Explanation: The bus adapter that connects the canister to the enclosure midplane has failed.

This node error does not, in itself, stop the node canister becoming active in the system. However, the bus might be being used to communicate between the node canisters in a clustered system. Therefore, it is possible that this node error indicates the reason why the critical node error 550 A cluster cannot be formed because of a lack of cluster resources is reported on the node canister.

Data:

- A number indicating the adapter location. Location 0 indicates that the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

- Node canister

732 The bus adapter has a PCI error.

Explanation: The bus adapter that connects the canister to the enclosure midplane has a PCI error.

This node error does not, in itself, stop the node canister becoming active in the system. However, the bus might be being used to communicate between the node canisters in a clustered system; therefore it is possible that this node error indicates the reason why the critical node error 550 A cluster cannot be formed because of a lack of cluster resources is reported on the node canister.

Data:

- A number indicating the adapter location. Location 0 indicates that the adapter integrated into the system board is being reported.

User response:

1. If possible, this noncritical node error should be serviced using the management GUI and running the recommended actions for the service error code.
2. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

- Node canister

733 The bus adapter degraded.

Explanation: The bus adapter that connects the canister to the enclosure midplane is degraded.

This node error does not, in itself, stop the node canister becoming active in the system. However, the bus might be being used to communicate between the node canisters in a clustered system. Therefore, it is possible that this node error indicates the reason why the critical node error 550 A cluster cannot be formed because of a lack of cluster resources is reported on the node canister.

Data:

- A number indicating the adapter location. Location 0 indicates that the adapter integrated into the system board is being reported.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. As the adapter is located on the system board, replace the node canister using the remove and replace procedures.

Possible Cause-FRUs or other cause:

- Node canister

734 Fewer bus ports.

Explanation: One or more PCI bus ports that have previously been active are now inactive. This condition has existed for over one minute. That is, the internode link has been down at the protocol level.

This could be a link issue but is more likely caused by the partner node unexpectedly failing to respond.

Data:

Three numeric values are listed:

- The ID of the first unexpected inactive port. This is a decimal number.
- The ports that are expected to be active. This is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is expected to be active.
- The ports that are actually active. This is a hexadecimal number. Each bit position represents a port, with the least significant bit representing port 1. The bit is 1 if the port is active.

User response:

1. If possible, this noncritical node error should be serviced using the management GUI and running the recommended actions for the service error code.
2. Follow the procedure for getting node canister and clustered-system information and determine the state of the partner node canister in the enclosure. Fix any errors reported on the partner node canister.

3. Use the remove and replace procedures to replace the enclosure.

Possible Cause-FRUs or other cause:

- Node canister
- Enclosure midplane

768 Ambient temperature warning.

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

Data:

- A text string identifying the thermal sensor reporting the warning level and the current temperature in degrees (Celsius).

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Check the temperature of the room and correct any air conditioning or ventilation problem.
3. Check the airflow around the system. Ensure no vents are blocked.

Possible Cause-FRUs or other cause:

- None

769 CPU temperature warning.

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

Data:

- A text string identifying the thermal sensor reporting the warning level and the current temperature in degrees (Celsius).

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Check the temperature of the room and correct any air conditioning or ventilation problems.
3. Check the airflow around the system. Ensure no vents are blocked.
4. If the error is still reported, replace the node canister.

Possible Cause—FRUs or other cause:

- CPU

840 Unsupported hardware change detected.

Explanation: A change has been detected to this node canister hardware configuration. The new configuration is not supported by the node canister software. User action is required to repair the hardware or update the software.

User response:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Follow the procedure for getting node canister and clustered-system information. A change to the hardware configuration is expected.
3. If the hardware configuration is unexpectedly reduced, make sure the component has not been unseated. Hardware replacement might be necessary.
4. If a new hardware component is shown as unsupported, check the software version required to support the hardware component. Upgrade the software to a version that supports the hardware.

If the hardware detected does not match the expected configuration, replace the hardware component that is reported incorrectly.

Possible Cause-FRUs or other cause:

- One of the optional hardware components might require replacement

841 Supported hardware change detected.

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

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

User response:

1. Use the management GUI to run the recommended actions for the associated service error code. Use the directed maintenance to accept or reject the new configuration.

850 The canister battery is reaching the end of its useful life.

Explanation: The canister battery is reaching the end of its useful life. It should be replaced within a week of the node error first being reported.

User response:

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

2. Replace the node canister battery by using the remove and replace procedures.

Possible Cause-FRUs or other cause:

- Canister battery

860 Fibre Channel network fabric is too big.

Explanation: The number of Fibre Channel (FC) logins made to the node canister exceeds the allowed limit. The node canister continues to operate, but only communicates with the logins made before the limit was reached. The order in which other devices log into the node canister cannot be determined, so the node canister's FC connectivity might vary after each restart. The connection might be with host systems, other storage systems, or with other node canisters.

This error might be the reason the node canister is unable to participate in a system.

The number of allowed logins per node is 1024.

Data:

- None

User response: This error indicates a problem with the Fibre Channel fabric configuration. It is resolved by reconfiguring the FC switch:

1. If possible, use the management GUI to run the recommended actions for the associated service error code.
2. Rezone the FC network so only the ports the node canister needs to connect to are visible to it.

Possible Cause-FRUs or other cause:

- None

878 Attempting recovery after loss of state data.

Explanation: During startup, the node canister cannot read its state data. It reports this error while waiting to be added back into a clustered system. If the node canister is not added back into a clustered system within a set time, node error 578 is reported.

User response:

1. Allow time for recovery. No further action is required.
2. Keep monitoring in case the error changes to error code 578.

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.

1021 Incorrect enclosure

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

User response: See node error 500.

1036 The enclosure identity cannot be read.

Explanation: The cluster is reporting that a node is not operational because of critical node error 509. See the details of node error 509 for more information.

User response: See node error 509.

1188 Too many software crashes have occurred.

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

User response: See node error 564.

1189 The node is held in the service state.

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

User response: See node error 690.

1202 A solid-state drive is missing from the configuration.

Explanation: The offline solid-state drive (SSD) identified by this error must be repaired.

User response: In the management GUI, click **Troubleshooting > Recommended Actions** to run the recommended action for this error. Otherwise, use MAP 6000 to replace the drive.

Chapter 12. Installing optional features

There are optional features that might need to be installed for the Flex System V7000 Storage Node.

Use these procedures to install optional features in the Flex System V7000 Storage Node

- Drives for the enclosures
- Adapters for the node canisters

Installing an optional host interface card

This topic describes how to install an optional host interface card into a node canister.

About this task

Note: You might have to remove an existing host interface card before installing the new optional host interface card.

To install an optional node canister host interface card, follow these steps:

Procedure

1. Use the management GUI to remove the node canister from the cluster.
Be aware of the following canister LED states:
 - If both the power LED and system status LED are on, do not remove a node canister.
 - If the system status is off, it is acceptable to remove a node canister.
 - If the power LED is flashing or off, it is safe to remove a node canister.
2. Open the service GUI. Select the node that is in service state with a node error of 690, and power it off.
3. Remove the node canister from the control enclosure. See “Removing a canister from an enclosure” on page 48.
4. Remove the cover from the node canister. See “Removing the node canister cover” on page 51. For additional information about replacing the host interface cards, see the service label attached to the canister cover.
5. Determine the location for the new host interface card and if you have to remove an existing host interface card first.

Note:

- The location of the host interface card in slot 1 (port 1) is on the left side when facing the front of the canister
 - The location of the host interface card in slot 2 (port 2) is on the right side when facing the front of the canister
6. If you are adding a new host interface card into an empty slot, continue at Step 11 on page 158. If you are exchanging host interface card types, continue to the next step.

7. Push the card retainer **2** to the side to release the host interface card being removed **1**.




Figure 44. Releasing the host interface card

8. Slide the host interface card toward the rear of the canister until the four locating pins are aligned with larger opening of the four keyholes in the card.
9. Lift and rotate the host interface card upward to gain access to the card cable connector on the system board.
10. Disconnect the cable connector from the system board. Do not disconnect the cable from the host interface card.
11. Connect the new adapter and cable to the system board.
12. Push the card retainer **2** out of the way.
13. Rotate the adapter **1** into position and align the four larger opening of the keyholes in the card with the four locating pins on the card carrier.




Figure 45. Positioning the adapter

14. Slide the adapter toward the front of the canister until the four locating pins are aligned in smaller openings of the four keyholes in the card. Then release the card retainer.
15. Install the cover on the node canister. See “Installing the node canister cover” on page 52.
16. Install the node canister into the control enclosure. See “Installing a canister into an enclosure” on page 50. Make sure to reattach the SAS cable to the canister if one was attached.

The canister starts up automatically.

17. After adding the new adapter there will be an alert in the Flex System V7000 event log indicating new hardware found, follow the fix procedure to accept the change and enable the adapter.

Installing optional drives


This topic describes how to install an optional 2.5-inch drive assembly.

About this task

To install an optional drive, complete the following steps:

Procedure

1. Rotate the handle **1** to the full extension. This will position the drive assembly for installation.



2. Push the new drive assembly back into the slot until the handle starts to move.
3. Finish inserting the drive assembly by closing the handle until the locking catch clicks into place.
4. After adding the new drive(s) there will be an alert indicating new hardware found, follow the fix procedure to accept the change and enable the drive(s).

Appendix. Accessibility features for *IBM Flex System V7000 Storage Node*

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

Accessibility features

These are the major accessibility features associated with the *Flex System V7000 Storage Node Information Center*:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. PDF documents have been tested using Adobe Reader version 7.0. HTML documents have been tested using JAWS version 13.0.
- This product uses standard Windows navigation keys.

Keyboard navigation

You can use keys or key combinations to perform operations and initiate menu actions that can also be done through mouse actions. You can navigate the *Flex System V7000 Storage Node Information Center* from the keyboard by using the shortcut keys for your browser or screen-reader software. See your browser or screen-reader software Help for a list of shortcut keys that it supports.

IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the 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 grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
Almaden Research
650 Harry Road
Bldg 80, D3-304, Department 277
San Jose, CA 95120-6099
U.S.A.*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements 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 should 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 trademarks and special non-IBM trademarks in this information are identified and attributed.

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.

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.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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 conform 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 Deutschland GmbH
Technical Regulations, Department M372
IBM-Allee 1, 71139 Ehningen, Germany
Tele: +49 7032 15 2941
e-mail: mailto:lugi@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 GmbH
Technical Regulations, Department M372

IBM-Allee 1, 71139 Ehningen, Germany
Tele: +49 7032 15 2941
e-mail: <mailto:lugi@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

이 기기는 업무용(A급)으로 전자파적합기기로서
판매자 또는 사용자는 이 점을 주의하시기
바라며, 가정외의 지역에서 사용하는 것을 목
적으로 합니다.

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。

tailemi

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](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-inch drives 9
- 2.5" drive assembly
 - replacing 42
- 2076 expansion enclosure
 - components 10
 - major components 1
- 4939
 - major components 1
 - 4939 control enclosure 3
 - 4939 drives 9
 - 4939 expansion canister
 - indicators 8
 - 4939 expansion enclosure 6
 - 4939 FRUs 55
 - 4939 node canister
 - components 4
 - indicators 6
 - 4939 parts
 - removing
 - overview 35
 - replacing
 - overview 35
 - 4939 tier 1 CRUs 35

A

- about this document
 - sending comments xxii
- accessibility
 - disability xix, 161
 - features 161
 - keyboard xix, 161
 - navigation 161
 - overview xix
 - repeat rate
 - up and down buttons 161
 - shortcut keys 161
- accessing
 - cluster (system) CLI 31
 - management GUI 28
 - publications 161
 - service assistant 30, 91
 - service CLI 32
- alerts
 - best practices 20

B

- backing up
 - best practices 20
 - system configuration files 115
- backup configuration files
 - deleting
 - using the CLI 120
 - restoring 117
- bad blocks 25
- battery
 - maintenance 23

- best practices
 - alerts 20
 - backing up data 20
 - drive characteristics 20
 - inventory reporting 19
 - IP address 17
 - notifications 18
 - passwords 17
 - power management 18
 - RAID 20
 - record
 - location information 21
 - subscribe
 - notifications 21
 - troubleshooting 17
 - warranty agreement
 - maintenance agreement 21
- bezel
 - replacing 48
- browsers
 - supported 70

C

- Canadian electronic emission notice 166
- canister
 - identification 74
- caution notices xvii
- changing
 - service IP address 88, 101
- CLI
 - cluster (system) commands 31
 - service commands 31
- cluster (system) CLI
 - accessing 31
 - when to use 31
- cluster (system) commands
 - CLI 31
- clustered storage system
 - failure to create 66
- clustered system
 - initializing
 - with service assistant 89
 - with the CMM 104
 - restore 113
- clustered systems
 - error codes 140
 - recovery codes 140
 - restore 108
 - T3 recovery 108
- CMM 102
 - collect System Management Server
 - service data 103
 - create system panel
 - not displaying 67
 - identifying an enclosure 100
 - start management panel
 - create system panel not displayed 67
 - starting the management GUI 100
 - using to power off a node 99

- CMM (*continued*)
 - using to power on a node 99
 - viewing and understanding Flex System V7000 information 97
 - virtual re-seat of a canister 103
- CMOS battery
 - replacing 46
- codes
 - node error
 - critical 139
 - noncritical 139
- commands
 - satask.txt 33
 - svcconfig backup 115
 - svcconfig restore 117
- comments, sending xxii
- components
 - 2076 expansion enclosure 10
 - 4939 node canister 4
 - control enclosure 3
 - enclosure 10
 - end cap
 - indicators 12
 - expansion enclosure 6
 - illustration 13
 - illustrated 1
- contact information
 - European 169
 - Taiwan 169
- control enclosure
 - components 3
 - detection error 72
 - replacing 55
- critical
 - node errors 139
- CRUs, Tier 1, removing and replacing 36
- customer replaceable units (CRUs) 35

D

- danger notices xv
- deleting
 - backup configuration files
 - using the CLI 120
 - system 87
- detection error
 - control location 72
 - expansion location 71
- Deutschsprachiger EU Hinweis 167
- DHCP service address
 - troubleshooting procedure 70
- disability
 - accessibility xix, 161
- discover
 - FSM cannot 64
 - manageV7000 cannot 64
- documentation
 - improvement xxii
- drive characteristics
 - best practices 20

- drive indicators 9
- drives 9
 - 2.5-inch drives 10
 - 3.5-inch drives 10
 - LED indicator 10

E

- electrical equipment, servicing xiv
- electronic emission notices
 - Avis de conformité à la réglementation d'Industrie Canada 166
 - Deutschsprachiger EU Hinweis 167
 - European Union (EU) 166
 - Federal Communications Commission (FCC) 165
 - French Canadian 166
 - Germany 167
 - Industry Canada 166
 - International Electrotechnical Commission (IEC) 168
 - Japanese Voluntary Control Council for Interference (VCCI) 168
 - Korean 168
 - New Zealand 166
 - People's Republic of China 168
 - Taiwan 169
 - United Kingdom 168
- EMC statement, People's Republic of China 168
- enclosure
 - identification 74
- end cap
 - indicators 12
- environmental notices xi
- error
 - control enclosure 72
 - expansion enclosure 71
 - node canister 70
 - not detected 71, 72
 - SAS cabling 70
 - USB flash drive 72
- error codes 128
- understanding 124
- error event IDs 128
- error events 121
- errors
 - logs
 - describing the fields 122
 - error events 121
 - managing 122
 - understanding 121
 - viewing 121
 - node 139
- Ethernet
 - status 86
- European contact information 169
- European Union (EU), EMC Directive conformance statement 166
- event IDs 124
- event notification 123
- events
 - reporting 121
- expansion canister 7
 - installing 50
 - LEDs 15

- expansion canister (*continued*)
 - removing 49
 - replacing 38
- expansion canister status
 - LEDs 83
- expansion enclosure
 - components 6, 13
 - detection error 71
 - power supply unit 13
 - replacing 60

F

- failure
 - storage system creation 66
- FCC (Federal Communications Commission) electronic emission notice 165
- features
 - accessibility 161
- Federal Communications Commission (FCC) electronic emission notice 165
- Fibre Channel
 - link failures 92
- field replaceable units (FRUs) 55
- fields
 - event log 122
- finding
 - Ethernet
 - status 86
- finding the power usage of canister 102
- finding the temperature of canister 102
- fix
 - errors 110
- fixing
 - node errors 88
- Flex System V7000 Storage Node
 - safety notices xv, xvii
- Flex System V7000 Storage Node library
 - related publications xx
- French Canadian electronic emission notice 166
- FRUs, removing and replacing 55
- FSM
 - cannot discover Flex V7000 64

G

- Germany electronic emission compliance statement 167
- GUI connectivity issues
 - troubleshooting procedure 65, 69
- guidelines
 - servicing electrical equipment xiv
 - trained service technicians xiii

H

- HIC card faults
 - resolving 93

I

- identifying
 - canister 74

- identifying (*continued*)
 - enclosure 74
 - status 75
- identifying an enclosure 100
- IEC (International Electrotechnical Commission) electronic emission notice 168
- indicators
 - 4939 expansion canister 8
 - 4939 node canister 6
 - end cap 12
- information center xx
- informational events 124
- initializing
 - clustered systems
 - with service assistant 89
 - with the CMM 104
- inspecting for unsafe conditions xiii
- installing
 - expansion canister 50
 - node canister 50
 - node canister cover 52
 - optional drive 157, 159
 - optional features 157
 - optional host interface card 157
 - RFID tag 54
- International Electrotechnical Commission (IEC) electronic emission notice 168
- inventory information 123
- inventory reporting
 - best practices 19
- IP address
 - best practices 17
- iSCSI
 - link problems 92

J

- Japanese electronic emission notice 168

K

- keyboard
 - accessibility xix, 161
- Korean electronic emission statement 168

L

- LEDs
 - expansion canister 15
 - expansion canister status 83
 - node canister status 76
 - system state 81
- legal notices
 - Notices 163
 - trademarks 165
- link failures
 - Fibre Channel 92
- link problems
 - iSCSI 92
- location information
 - record
 - best practices 21
- log files
 - viewing 121

M

- maintenance
 - battery 23
- maintenance agreement
 - best practices 21
- major components 1
- management GUI 100
 - accessing 28
 - cannot log on 66
 - supported browsers 70
 - troubleshooting procedure
 - start here 63
- management GUI interface
 - when to use 28
- management GUI IP address
 - viewing using CMM 101
- management IP address
 - troubleshooting procedure 65
- manageV7000 command
 - cannot discover Flex V7000 64
- managing
 - event log 122
- medium errors 25
- message classification 140
- mirrored volumes
 - not identical 72

N

- navigation
 - accessibility 161
- New Zealand electronic emission statement 166
- node canister
 - installing 50
 - location error 70
 - re-seating 90
 - removing 49
 - replacing 36
 - unknown service address 68
- node canister battery
 - replacing 38
- node canister cover
 - installing 52
 - removing 51
- node canister host interface card
 - replacing 40
- node canister status
 - LEDs 76
- node errors
 - fixing 88
- node rescue
 - performing 91
- noncritical
 - node errors 139
- notifications
 - best practices 18
 - sending 123
 - subscribe
 - best practices 21
- number range 140

O

- optional adapter
 - installing 157

- optional drive
 - installing 157, 159
- optional features
 - installing 157
- optional host interface card
 - installing 157

P

- parts
 - removing
 - preparing 35
 - replacing
 - preparing 35
- passwords
 - best practices 17
- PCI errors
 - resolving 93
- People's Republic of China, electronic emission statement 168
- performing
 - node rescue 91
- ports
 - SAS 15
- POST (power-on self-test) 123
- power management
 - best practices 18
- power supply unit
 - expansion enclosure 13
- power usage 102
- power-on self-test 123
- powering off
 - system 90
- powering off a node 99
- powering on a node 99
- problem
 - mirrored volumes
 - not identical 72
- publications
 - accessing 161

R

- RAID
 - best practices 20
- re-seating
 - node canister 90
- reader feedback, sending xxii
- recovering
 - offline virtual disks (volumes)
 - using CLI 112
- recovery
 - system
 - when to run 108
 - systems
 - starting 111
- related information xx
- removing
 - 4939 parts
 - overview 35
 - 550 errors 110
 - 578 errors 110
 - expansion canister 49
 - node canister 49
 - node canister cover 51

- removing (*continued*)
 - parts
 - preparing 35
 - RFID tag 54
 - system 87
 - system data 87
- removing and replacing
 - 4939 parts 36
 - FRUs 55
 - Tier 1 CRUs 36
- replacing
 - 2.5" drive assembly 42
 - 4939 parts
 - overview 35
 - a node canister host interface card 40
 - bezel 48
 - CMOS battery 46
 - control enclosure 55
 - expansion canister 38
 - expansion enclosure 60
 - node canister 36
 - node canister battery 38
 - parts
 - preparing 35
 - SAS cable 44
- reporting
 - events 121
- rescue
 - performing
 - for a node 91
- resetting
 - superuser password 73
- resolving
 - HIC card faults 93
 - PCI errors 93
- restore
 - system 107, 113
- RFID tag
 - installing 54
 - removing 54

S

- safety xi
 - caution notices xvii
 - danger notices xv
- safety notices xi
- safety statements xi
- SAS
 - ports 15
- SAS cable
 - replacing 44
- SAS cabling
 - location error 70
- satask.txt
 - commands 33
- sending
 - comments xxii
- service address
 - unknown 68
- service assistant
 - accessing 30, 91
 - interface 29
 - supported browsers 70
 - when to use 29
- service CLI
 - accessing 32

- service CLI (*continued*)
 - when to use 32
- service commands
 - CLI 31
- service IP address
 - changing 88
 - changing using CMM 101
 - viewing using CMM 101
- servicing electrical equipment xiv
- shortcut keys
 - accessibility 161
 - keyboard 161
- starting
 - system recovery 111
- status
 - Ethernet 86
 - identifying 75
 - node canister 75, 76
 - system 75, 76
- storage systems
 - restore 107
 - servicing 93
- superuser
 - password
 - resetting 73
- supported browsers 70
- system
 - backing up configuration file using
 - the CLI 115
 - deleting 87
 - restoring backup configuration
 - files 117
- system data
 - removing 87
- system state
 - LEDs 81

T

- T3 recovery
 - removing
 - 550 errors 110
 - 578 errors 110
 - when to run 108
- Taiwan
 - contact information 169
 - electronic emission notice 169
- temperature 102
- trademarks 165
- trained service technicians,
 - guidelines xiii
- troubleshooting
 - best practices 17
 - event notification email 123
 - node errors 88
- troubleshooting procedure
 - DHCP service address
 - service assistant 70
 - GUI connectivity issues
 - main GUI 65
 - service assistant 69
 - management IP address 65
 - start here
 - management GUI 63

U

- understanding
 - clustered-system recovery codes 140
 - error codes 124
 - event log 121
- United Kingdom electronic emission
 - notice 168
- unsafe conditions, inspecting for xiii
- usb flash drive
 - satask.txt commands 33
- USB flash drive
 - detection error 72
 - using 32
 - when to use 32
- using
 - GUI interfaces 27
 - management GUI 27
 - service assistant 29
 - USB flash drive 32

V

- VDisks (volumes)
 - recovering from offline
 - using CLI 112
- viewing
 - event log 121
 - management GUI IP address 101
 - node canister
 - status 75, 76
 - service IP address 101
 - system
 - status 75, 76
- volumes (VDisks)
 - recovering from offline
 - using CLI 112

W

- warranty agreement
 - best practices 21
- when to use
 - cluster (system) CLI 31
 - management GUI interface 28
 - service assistant 29
 - service CLI 32
 - USB flash drive 32



Printed in USA

GC27-4205-00

