

IBM Storage Virtualize for SAN Volume
Controller and FlashSystem Family

Troubleshooting Guide



Note

Before using this information and the product it supports, read the information in [Chapter 1, “Notices ,”](#) on [page 1](#).

This edition applies to version 8, release 7, modification x, and to all subsequent modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2024.**

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

Contents

Figures.....	v
Tables.....	vii
Chapter 1. Notices	1
Trademarks.....	2
Chapter 2. Best practices for troubleshooting.....	5
How to get information, help, and technical assistance.....	6
Chapter 3. User interfaces for servicing your system.....	9
Management GUI interface.....	9
Deleting a node from a clustered system by using the management GUI.....	10
Using the GUI for troubleshooting.....	11
Service assistant interface.....	13
When to use the service assistant.....	14
Accessing the service assistant.....	14
Adding a node to a system using the service assistant.....	15
Command-line interface.....	16
Service command-line interface.....	16
USB flash drive interface.....	16
Technician port for node access.....	23
Chapter 4. Event reporting.....	25
Understanding events.....	25
Event notifications.....	27
Power-on self-test.....	27
Chapter 5. Resolving a problem.....	29
Problem: Management IP address unknown.....	29
Problem: Unable to connect to the management GUI.....	29
Problem: Unable to log on to the management GUI.....	30
Problem: Cannot initialize or create a system.....	30
Problem: Node canister service IP address unknown.....	31
Problem: Cannot connect to the service assistant.....	31
Problem: Management GUI or service assistant does not display correctly.....	32
Problem: Node canister has a location node error.....	32
Problem: SAS cabling not valid.....	32
Problem: New expansion enclosure not detected.....	33
Problem: Control enclosure is not detected.....	33
Problem: Mirrored volume copies no longer identical.....	33
Problem: Command file not processed from USB flash drive.....	34
Problem: A storage pool has arrays that are not encrypted.....	34
Procedure: Resetting the superuser password by using the management GUI or CLI.....	34
Procedure: Identifying which enclosure or canister to service.....	34
Procedure: Checking the status of your system.....	35
Procedure: Getting node canister and system information using the service assistant.....	36
Procedure: Getting node canister and system information by using a USB flash drive.....	36
Procedure: Understanding the system status from the LEDs	37

Procedure: Finding the status of Ethernet connections.....	37
Procedure: Finding the status of SAS connections.....	38
Procedure: Removing system data from a node canister.....	38
Procedure: Removing a control enclosure from a system.....	39
Procedure: Fixing node errors.....	39
Procedure: Changing the service IP address of a node canister.....	40
Procedure: Accessing the service assistant from the technician port.....	40
Procedure: Powering on the system.....	41
Procedure: Powering on a node canister.....	41
Procedure: Collecting information for support.....	42
Procedure: Rescuing the software from another node canister (node rescue).....	42
Procedure: SAN problem determination.....	43
Replacing a expansion enclosure attachment SAS cable.....	43
Resolving a problem with dual boot drives.....	44
Resolving a problem with offloaded data transfer.....	46
Resolving a problem with SSL certificates.....	47
Resolving a problem with SSL/TLS clients.....	48
Procedure: Making drives support protection information.....	48
Debugging and performance-monitoring statistics for offloaded data transfer.....	49
iSCSI performance analysis and tuning.....	51
Problem: iSCSI backend MDisks are in a degraded or offline state.....	52
Problem: iSCSI backend controllers are in a degraded or offline state.....	53
Problem: Ethernet host connectivity issues.....	53
Fibre Channel link failures.....	54
Validating Ethernet partnerships configuration.....	55
Chapter 6. Recover system procedure.....	57
When to run the recover system procedure.....	59
Fix hardware errors.....	59
Removing system information for nodes with error code 550 or error code 578 using the service assistant.....	59
Running system recovery by using the service assistant.....	60
Recovering from offline volumes by using the CLI.....	63
What to check after running the system recovery.....	64
Verifying migration volumes after a system recovery.....	66
Chapter 7. Backing up and restoring the system configuration.....	69
Backing up the system configuration using the CLI.....	69
Restoring the system configuration.....	71
Deleting backup configuration files by using the CLI.....	78
Chapter 8. Enabling concurrent maintenance.....	79
Chapter 9. Removing a control enclosure and its expansion enclosures.....	81
Chapter 10. Deleting a node from a system by using the CLI.....	83
Chapter 11. Moving expansion enclosures.....	85
Chapter 12. Shutting down a system by using the CLI.....	87

Figures

1. Proper orientation for SAS cable connector..... 44

2. Location of the power LED on a node canister 87

Tables

1. IBM websites for help, services, and information..... 6

2. Description of data fields for the event log.....26

3. Default service IP addresses..... 31

4. ODX debug global parameters..... 50

5. ODX debug local parameters..... 51

6. Volume migration after a system recovery..... 66

7. Files created by the backup process..... 70

Chapter 1. Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux® and the Linux logo is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.
Other product and service names might be trademarks of IBM or other companies.

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.

Follow proper power management procedures

You might lose access to volume data if you incorrectly power off all or part of the system.

Always use the management GUI function to power off the system. Only power off or remove a node canister if instructed to do so in a service action.

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

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.

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 system helps resolve these problems through the **Recommended Actions** option from the **Monitoring > Events** page in the management GUI.

Complete 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.

This can be done using the management GUI, or by checking the IBM support website to see if new code releases are available: www.ibm.com/support

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 a potential issue.

Keep your system records up to date

Keep good records of the system location, names, and management addresses. 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.

Record the MT-M and serial number of each enclosure. The information can be found on the IBM Standard Asset Tag attached to the left enclosure bezel, which includes a machine-readable data matrix to ISO/IEC 15434.

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: <http://www.ibm.com/support>

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 warranty and maintenance agreement details

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

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

How to get information, help, and technical assistance

If you need help, service, technical assistance, or want more information about IBM products, you can find a wide variety of sources available from IBM to assist you.

Information

IBM maintains pages on the web where you can get information about IBM products and fee services, product implementation and usage assistance, break and fix service support, and the latest technical information. For more information, refer to this table.

Table 1. IBM websites for help, services, and information	
Website	Address
IBM home page	http://www.ibm.com
Directory of worldwide contacts	http://www.ibm.com/planetwide
Support for products	http://www.ibm.com/support

Note: Available services, telephone numbers, and web links are subject to change without notice.

Help and service

When you call from the US or Canada, choose the **storage** option.

Before you call

Make sure that you take steps to try to solve the problem yourself before you call. Some suggestions for resolving the problem before you call IBM Support include:

- Check all cables to make sure that they are connected.
- Use the troubleshooting information in your system documentation. The troubleshooting section of IBM Documentation for your product contains procedures to help you diagnose problems.

- Write down the machine type and model (MTM) and the serial number of the machine with the problem. IBM uses this information to engage the appropriate support team.

Using the documentation

Information about your IBM storage system is available in the information that is available on this IBM Documentation and on the IBM support website.

That information includes online documents, readme files, and help files. See the troubleshooting information for diagnostic instructions. The troubleshooting procedure might require you to download updated software. IBM maintains pages on the web where you can get the latest technical information and download updates. To access this information, go to <http://www.ibm.com/support> and follow the instructions. Also, some documents are available through the IBM Publications Center.

Sign up for the Support Line Offering

If you have questions about how to use and configure the machine, sign up for the IBM Support Line offering to get a professional answer.

The maintenance that is supplied with the system provides support when there is a problem with a hardware component or a fault in the system machine code. At times, you might need expert advice about using a function that is provided by the system or about how to configure the system. Purchasing the IBM Support Line offering gives you access to this professional advice for your system, and in the future.

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

Chapter 3. User interfaces for servicing your system

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

The interfaces for servicing your system connect through the 1 Gbps Ethernet ports that are accessible from port 1 of each canister.

- Use the initialization tool to do the initial setup of your system. (The USB flash drive is not required to initialize the system configuration. However, it can be used for other functions.)
- Use the management GUI to monitor and maintain the configuration of storage that is associated with your clustered systems.
- Use the service assistant to complete service procedures.
- Use the command line interface (CLI) to manage your system.

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

Management GUI interface

The management GUI is a browser-based interface 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** page can be filtered four ways:

Recommended action (default)

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

- Run a fix procedure.
- View the properties.

Unfixed alerts

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

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

Unfixed messages and alerts

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

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

Show all

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

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

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

You can also sort events by time or error code. When you sort by error code, the most serious events, those with the lowest numbers, are displayed first. You can select any event that is listed and select

Actions > Properties to view details about the event.

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

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

Before you begin

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

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



Attention:

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

- If you turn on the power to the node that is removed and it is still connected to the same fabric or zone, it attempts to rejoin the system. The system tells the node to remove itself from the system and the node becomes a candidate for addition to this system or another system.
- If you are adding this node into the system, ensure that you add it to the same I/O group that it was previously a member of. Failure to do so can result in data corruption.

This task assumes that you access the management GUI.

About this task

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

Procedure

1. Select **Monitoring > System**.
2. On the **System -- Overview** page, use the directional arrow next to the node.
3. On the **Node -- Details** page, select **Node Actions > Remove**.

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

If the node that you want to remove is shown as **Online**, deleting the node can result in the dependent volumes to also go offline. Verify whether the node has any dependent volumes. To verify dependent volumes, select **Node Actions > Dependent Volumes**.

Using the GUI for troubleshooting

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

About this task

You can connect to and manage a system that uses the management GUI as soon as you have created a clustered system.

To run the management GUI, start a supported web browser and point it to any of the management IP addresses of your system. After the connection is successful, you see a login panel.

Log in using your username and password.

When you have logged in, select **Monitoring > Events**. Depending upon how you choose to filter alerts, you might see the following options:

- Alerts that require attention.
- Alerts and messages that are not fixed.
- 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.

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

Monitoring events

The management GUI is a browser-based interface 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** page can be filtered four ways:

Recommended action (default)

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

- Run a fix procedure.
- View the properties.

Unfixed alerts

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

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

Unfixed messages and alerts

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

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

Show all

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

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

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

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

Diagnosing and resolving problems with fix procedures

You can use fix procedures to diagnose and resolve problems with the system.

About this task

For example, to repair a system, you might complete the following tasks:

- Analyze the event log (if it is available, or view node errors).
- Replace failed components.
- Verify the status of a repaired device.
- Restore a device to an operational state in the system.
- Mark the error as fixed in the event log.

Fix procedures help simplify this process by automating as many of the tasks as possible. The following example uses the management GUI to repair a system.

Procedure

Complete the following steps to start the fix procedure.

1. Click **Monitoring > Events** and ensure that you are filtering the event log to display **Recommended actions**.

The list might indicate that many errors must be repaired. If the list contains several errors, the error at the top of the list has the highest priority and must always be fixed first. If you do not fix the higher priority errors first, you might not be able to fix the lower priority errors.

2. Select the error at the top of the list or select the **Next recommended action**.
3. Click **Run Fix Procedure**.

The pane displays the error code and provides a description of the condition.

4. Click **Next** to go forward or **Cancel** to return to the previous pane.

One or more panes might be displayed with instructions for you to replace parts or complete other repair activity.

5. If you are not able to complete the actions now, click **Cancel** until you return to the previous pane. Click **Cancel** until you are returned to the **Next recommended action** pane. When you return to the fix procedures, the repair can be restarted from step [“1” on page 13](#). After you complete all the instructions, click **OK**. When the last repair action is completed, the procedures might attempt to restore failed devices to the system.
6. After you complete the fix, you see the statement **Click OK to mark the error as fixed**. Click **OK**. This action marks the error as fixed in the event log and prevents this instance of the error from being listed again.
7. When you see the statement **The repair has been completed.**, click **Exit**. If other errors must be fixed, those errors are displayed and the fix procedures continue.

If no errors remain, the following statement is displayed: **There are no unfixed errors in the event log**. The statement indicates that no further repair procedures are necessary.

Service assistant interface

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

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. On the system, you can also access the service assistant interface by using the technician port.

When to use the service assistant

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



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

Use the service assistant in the following situations:

- When you cannot access the system from the management GUI and you cannot access the system to run the recommended actions
- When the recommended action directs you to use the service assistant.

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

The service assistant 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 (WWNN) for each node.

You can also complete the following service-related actions:

- Collect logs to create and download a package of files to send to support personnel.
- Remove the data for the system from a node.
- Recover a system if it fails.
- Install a software package from the support site or rescue the software from another node.
- Update software on nodes manually versus completing a standard update procedure.
- Configure a control enclosure chassis after replacement.
- Change the service IP address that is assigned to Ethernet port 1 for the current node or 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.

The service assistant completes a number of tasks that 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 completed 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 or node canister. The service assistant can be accessed through a service IP address, or by using the technician port.

About this task

To access the service assistant by a service IP address, you must use a supported web browser. For a list of supported browsers, refer to the topic "Web browser requirements to access the management GUI".

Procedure

To access the service assistant, complete the following steps.

1. Start a supported web browser and point your web browser to *serviceaddress/service* for the node or 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.

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 was a member of a system, use the password for the superuser.

If you do not know the current superuser password, try to find out. If you cannot find out what the password is, reset the password.

Results

Complete the service assistant actions on the correct node or 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.

Adding a node to a system using the service assistant

The service assistant is a web application that can be used to troubleshoot nodes on your system. If a node goes offline, you can use the service assistant to add the node back into the system after any repair actions are completed.

Before you begin

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

About this task

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

Procedure

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

The node is added as soon as it is available.

What to do next

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

1. In the management GUI, select **Monitor > System**.
2. Select the canister to which the node was added.
3. Verify that the node was added.

Command-line interface

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

When to use the CLI

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

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

You might also find it useful to create command scripts that use CLI commands to monitor certain conditions or to automate configuration changes that you make regularly.

Accessing the system CLI

Follow the steps that are described in the command-line interface section to initialize and use a CLI session.

Service command-line interface

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

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

To initialize and use a CLI session, review the command-line interface topic of this product information.

To initialize and use a CLI session, refer to "Using the CLI" in the "Command-line interface" topic of this product information to learn how to log in to a node by using Secure Shell (SSH). To run service CLI commands on a node, you log in to the **service IP** address of the node that requires servicing. Alternatively, you can log in to the **management IP** address of the system to run service commands on the configuration node of the system.

When you are logged in to a node, you can generally issue service commands to any node in the system by specifying the panel ID of the node. For more information, see the CLI help for each command. Some error conditions, such as network problems, might prevent this function from working, so it is better to directly log in to the service IP address of the node that requires servicing.

USB flash drive interface

Use a USB flash drive to help service a node.

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

When to use the USB flash drive

The USB flash drive can be used for service functions.

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

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

Using a USB flash drive

Use any USB flash drive that is formatted with a FAT32 file system on its first partition.

About this task

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

The fault light-emitting diode (LED) on the node canister flashes when the USB service action is being completed. When the fault LED stops flashing, it is safe to remove the USB flash drive.

The fault light-emitting diode (LED) on the node flashes when the USB service action is being completed. When the fault LED stops flashing, it is safe to remove the USB flash drive.

Result

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

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

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

satask.txt commands

If you are creating the **satask.txt** command file by using a text editor, the file must contain a single command on a single line in the file.


The commands that you use are the same as the service CLI commands except where noted. Not all service CLI commands can be run from the USB flash drive. The **satask.txt** commands always run on the node that the USB flash drive is plugged into.

Reset service IP address and superuser password command

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

Syntax

```
➔ satask — chserviceip — -serviceip — ipv4 — -gw — ipv4 — -mask — ipv4 —  
— -resetpassword — <cleartext_password> ➔
```



Parameters

-serviceip *ipv4*

The IPv4 address for the service assistant.

-gw *ipv4*

The IPv4 gateway for the service assistant.

-mask *ipv4*

The IPv4 subnet for the service assistant.

-serviceip_6 *ipv6*

The IPv6 address for the service assistant.

-gw_6 *ipv6*

The IPv6 gateway for the service assistant.

-prefix_6 *int*

The IPv6 prefix for the service assistant.

-resetpassword <cleartext_password>

(Optional) Sets the service assistant password to **cleartext_password**. The password must be in the range of 6 - 64 characters.

Description

This command can be used to reset the service assistant IP address to a new value. If the **-default** parameter is provided, it resets the service assistant IP address to the default value.

If the **-resetpassword** parameter is provided, the superuser password for the system is reset to the value provided. The superuser password cannot be reset back to the default password. If the node or canister is active in a system, the superuser password for the system is reset; otherwise, the superuser password is reset on the individual node canister. If the node canister later becomes active in a system, the superuser password is reset to that of the system. You can configure the system to disable resetting the superuser password. If you disable that function, this action fails.

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

Reset service assistant password command

Use this command when you are unable to log on to the system because you forget the superuser password, and you wish to reset it.

Syntax

➤ **satask** — **resetpassword** — *-password <cleartext_password>* ➤

Parameters

None.

An invocation example

```
satask snap
```

The resulting output:

```
No feedback
```

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

An invocation example

```
satask snap -noimm
```

The resulting output:

```
No feedback
```

An invocation example

```
satask snap -dump 111584
```

The resulting output:

```
No feedback
```

An invocation example

```
satask snap -upload -pmr 63058,075,724
```

The resulting output:

```
No feedback
```

An invocation example

```
satask snap -upload -pmr TS000008577
```

The resulting output:

```
No feedback
```

Install software command

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

Syntax

```
➤ satask — installsoftware — -file — filename — { -ignore | -pacedccu } ➤
```

Parameters

-file *filename*

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

-ignore | -pacedccu

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

Description


This command copies the file from the USB flash drive to the update directory on the node canister, and then installs the update package.

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

Create system command

Use this command to create a storage system.

Syntax



Parameters

-clusterip ipv4

(Optional) The IPv4 address for Ethernet port 1 on the system.

-gw ipv4

(Optional) The IPv4 gateway for Ethernet port 1 on the system.

-mask ipv4

(Optional) The IPv4 subnet for Ethernet port 1 on the system.

-clusterip_6 ipv6

(Optional) The IPv6 address for Ethernet port 1 on the system.

-gw_6 ipv6

(Optional) The IPv6 gateway for Ethernet port 1 on the system.

-prefix_6 int

(Optional) The IPv6 prefix for Ethernet port 1 on the system.

-name cluster_name

(Optional) The name of the new system.

Description

This command creates a storage system.


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

Change system IP address

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

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

Syntax



Parameters

-systemip

The IPv4 address for Ethernet port 1 on the system.

-gw

The IPv4 gateway for Ethernet port 1 on the system.

-mask

The IPv4 subnet for Ethernet port 1 on the system.

-consoleip

The management IPv4 address of the system.

Description

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

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

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

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

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

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

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

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

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

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

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

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

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

Query status command

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

Syntax

➤ sainfo — getstatus ➤

Parameters

None.

Description

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

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

Technician port for node access

A technician port is available on the rear of the enclosure to use a direct connection for installing and servicing support.

The technician port provides direct access to the service assistant GUI and command-line interface (CLI). The technician port can be used by directly connecting a computer that has web browsing software and is configured for Dynamic Host Configuration Protocol (DHCP) through a standard 1 Gbps Ethernet cable.

If the node has candidate status when you open the web browser, the initialization tool is displayed. Otherwise, the service assistant interface is displayed.

Important: Do not use the initialization tool on a node if any other node in the system is already active. For example, a node status LED is solid on any node of the system.

After the Ethernet port of the personal computer is connected, open a supported browser and browse to address `https://service`. (If you do not have DHCP, open a supported browser and go to the following static IP address `192.168.0.1`.)

See the following information for how to access the CLI through the technician port. For other ways to access the CLI, see [Chapter 3, “User interfaces for servicing your system,” on page 9](#).

To use the technician port, you must be next to the node. The technician port does not work if it is connected to an Ethernet switch.

If you have Secure Shell (SSH) software on the computer that is directly connected to the technician port, then you can use it to access the CLI as superuser at `192.168.0.1`. The default superuser password is `passwd`.

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

Chapter 4. Event reporting

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

Event reporting process

The following methods are used to identify a new event:

- If you enabled Simple Network Management Protocol (SNMP), an SNMP trap is sent to an SNMP manager that is configured by the customer.
- If enabled, log messages can be forwarded on an IP network by using the syslog protocol.
- If enabled, event notifications can be forwarded by email by using Simple Mail Transfer Protocol (SMTP).
- Call Home can be enabled so that critical faults generate a problem management record (PMR) that is sent in an email to the appropriate support center.

Understanding events

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

Error data

Events are classified as either alerts or messages:

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

Understanding event codes

The system generates informational event codes when significant changes to the state of the system are detected.

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

Understanding the error codes

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

Error codes help you to identify the cause of a problem, a failing component, and the service actions that might be needed to solve the problem.

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

Viewing the event log

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

You can view the event log by using the **Monitoring > Events** options in the management GUI. The event log contains many entries. However, you can, 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 2 on page 26 describes some of the fields that are available to assist you in diagnosing problems.

Table 2. Description of data fields for the event log	
Data field	Description
Event ID	This number precisely identifies why the event was logged.
Description	A short description of the event.
Status	Indicates whether the event requires some attention. Alert: if a red icon with a cross is shown, follow the fix procedure or service action to resolve the event and turn the status green. Monitoring: the event is not yet of concern. Expired: the event no longer represents a concern. Message: provide useful information about system activity.
Error code	Indicates that the event represents an error in the system that can be fixed by following the fix procedure or service action that is identified by the error code. Not all events have an error code. Different events have the same error code if the same service action is required for each.
Sequence number	Identifies the event within the system.
Event count	The number of events that are coalesced into this event log record.
Object type	The object type to which the event relates.
Object ID	Uniquely identifies the object within the system to which the event relates.
Object name	The name of the object in the system to which the event relates.
Copy ID	If the object is a volume and the event refers to a specific copy of the volume, this field is the number of the copy to which the event relates.
Reporting node ID	Typically identifies the node responsible for the object to which the event relates. For events that relate to nodes, it identifies the node that logged the event, which can be different from the node that is identified by the object ID.
Reporting node name	Typically identifies the node that contains the object to which the event relates. For events that relate to nodes, it identifies the node that logged the event, which can be different from the node that is identified by the object name.

Table 2. Description of data fields for the event log (continued)

Data field	Description
Fixed	Where an alert is shown for an error or warning condition, it indicates that the user marked the event as fixed, completed the fix procedure, or that the condition was resolved automatically. For a message event, this field can be used to acknowledge the message.
First time stamp	The time when this error event was reported. If events of a similar type are being coalesced together, so that one event log record represents more than one event, this field is the time the first error event was logged.
Last time stamp	The time when the last instance of this error event was recorded into this event log record.
Root sequence number	If set, it is the sequence number of an event that represents an error that probably caused this event to be reported. Resolve the root event first.
Sense data	Extra data that gives the details of the condition that caused the event to be logged.

Event notifications

The system uses Simple Network Management Protocol (SNMP) traps, syslog messages, emails, and the Call Home function to notify you 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.

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

Power-on self-test

When you turn on the system, the node canisters complete self-tests.

The power-on self-test is completed to check the operation of components and some of the options that have been installed when the system is first turned on.

If a critical failure is detected during the POST, the software is not loaded and the fault LED is illuminated.

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

Chapter 5. Resolving a problem

Described here are some procedures to help resolve fault conditions that might exist on your system. A basic understanding of the system concepts is required.

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

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

Always use the recommended actions on the **Events** panel of the management GUI as the starting point to diagnose and resolve a problem.

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.

Note: After you have created your 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.

Problem: Management IP address unknown

If you are not able to run the management GUI because you do not know the management IP address, you can retrieve 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.

The management IP address 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.

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.

If you are unable to connect to the management GUI from your web browser and received a `Certificate expired` or similar error, see [“Resolving a problem with SSL certificates”](#) on page 47.

If you are unable to connect to the management GUI from your web browser and received a `cipher error`, `SSL error`, `TLS error`, or `handshake error` or similar error, see [“Resolving a problem with SSL/TLS clients”](#) on page 48.

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, either use the service assistant to verify that the state of at least one node canister is active, or if the node canister is not active, use the LEDs to see if any node canister state is active.

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 all node canisters is candidate, then there is not a clustered system to connect to. If the node state is service, see the topic that contains information about fixing node errors.

- Ensure that you are using the correct system IP address. If you can access the service assistant using the service address or the technician port of a node canister, log in to find the node and system addresses on the **Access** tab of the Node Detail table.
- Ensure that all node canisters have an Ethernet cable that is connected to port 1 and that the port is working. Use the Ethernet LEDs to understand the port status.

- 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 did not use 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 incorrect for your environment, take these steps:
 1. You can determine this if you can access the service assistant on any node canister. Access the service assistant using the technician port on the rear of a node canister if it cannot be accessed over your network. Alternatively use the summary data returned, when a USB flash drive is plugged into a node canister.
 2. You can temporarily run the management GUI on the service address of the configuration node. Point your browser to service address/gui. For example, if the service address of the configuration node is 11.22.33.44, point your browser to 11.22.33.44/gui.
 3. In the Management GUI, use the options in the **Settings > Network** to change the management IP settings.
 4. As an alternative to using the management GUI, you can use the **chsystemip** CLI command to correct the system IP address settings by using ssh to the service IP of the configuration node.

Problem: Unable to log on to the management GUI

If you can see the management GUI login screen but cannot log on, you have several options for correcting the problem.

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

- If you are not logging on as `superuser`, contact your system administrator to 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.
-

Problem: Cannot initialize or create a system

Use this information if your attempt to create a clustered system failed.

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

- System initialization wizard (using the technician port of a node canister)
- USB flash drive
- Management console
- Service assistant
- Service command line

To prevent accidental loss of volumes, a system cannot be initialized on an enclosure that is already configured in an existing system. Check the following details using the service assistant to confirm that the enclosure is not already configured in a system:

- The node canister that you are attempting to create a clustered system on must be in candidate state.
- The partner node canister in the control enclosure must not be in active state.
- The latest system ID of the control enclosure must be 0. If the partner node is in active state. If the partner code is not in the active state and the node canister on which you are attempting to create the system is in service state.

Problem: Node canister service IP address unknown

You can use several methods to determine the service address of a node canister.

A default service address is initially assigned to each node canister, as shown in Table 3 on page 31. Try using these addresses if the node has not been reconfigured, and the addresses are valid on your network.

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 in to 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 in to 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 in to 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 that 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 cannot find the service address of the node using the management GUI or service assistant of a different node, two options remain:

- You can connect directly to the service assistant of a node using the technician port of the node, as described in the Initializing the system with the technician port page of the relevant Hardware Guide..
- You can also use a USB flash drive to find the service address of the node. For more information, see “Procedure: Getting node canister and system information by using a USB flash drive” on page 36.

Table 3. Default service IP addresses		
Canister and port	IPv4 address	IPv4 subnet mask
Canister 1 (left) port 1 (left)	192.168.70.121	255.255.255.0
Canister 2 (right) port 1 (left)	192.168.70.122	255.255.255.0

Problem: Cannot connect to the service assistant

Use this information if you are unable to display the service assistant on your browser via the service IP address.

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 `service 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. 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.

- You cannot connect to the service assistant if the node canister is not able to start the code.
- 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.
- A default service address is initially assigned to each node canister. The service IP address 192.168.70.121 subnet mask 255.255.255.0 is preconfigured on Ethernet port 1. The service IP address 192.168.70.122 subnet mask 255.255.255.0 is preconfigured on Ethernet port 1.

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.

Problem: Management GUI or service assistant does not display correctly

If the Management GUI or the service assistant does not display correctly, verify that you are using a supported web browser.

For a list of supported browsers, see Management GUI.

Problem: 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.

A location error means that the node canister or the enclosure midplane was moved or changed. This is normally due to a service action not being completed or not being implemented correctly.

A number of different conditions are reported as location errors. 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 39](#).

After a node canister is 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. The used node canister 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:

- Do not connect a SAS cable between a port on a top canister and a port on a bottom canister.
- In any enclosure, the same ports must be used on both canisters.
- Do not connect a SAS cable between ports in the same enclosure.
- In each enclosure, where a cable connects from SAS port 1 of the top canister to another enclosure, a cable must also connect from SAS port 1 of the bottom canister to the other enclosure. Similarly, where a cable connects from SAS port 2 of the top canister to another enclosure, a cable must also connect from SAS port 2 of the bottom canister to the other 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 the top canister, nor port 2 of the bottom canister.
- Ensure that each SAS cable is fully inserted.

Problem: New expansion enclosure not detected

Determine why a newly installed expansion enclosure was not detected by the system.

When you install a new expansion enclosure, follow the Add Enclosure wizard in the management GUI. Select **Monitoring > System** and select **Add Enclosures**. If **Add Enclosures** is not available on the **System -- Overview** page, there might be a cabling issue that needs to be verified.

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

- Verify the status of the LEDs at the back of the expansion enclosure. At least one power supply unit must be on with no faults shown. At least one canister must be active, with no fault LED on.
- Verify that the SAS cabling to the expansion enclosure is correctly installed.

Problem: Control enclosure is not detected

If a control enclosure is not detected by the system, this procedure can help you resolve the problem.

When you install a new control enclosure, use the **Add Enclosures** wizard in the management GUI. To access this wizard, select **Monitoring > System**. If the management GUI detects a new control enclosure, click **Add Enclosures** and complete the wizard.

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 your system uses Fibre Channel, ensure that all cables are connected and zoning is set up according to the zoning rules. 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 the primary volume copy is correct and make the other volume copy match the primary volume 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 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 enclosure is not operational. Check the status using the light-emitting diode (LED) indicators.
- 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.
- 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.

Problem: A storage pool has arrays that are not encrypted

The management GUI provides a way to determine whether each of the arrays in a storage pool is encrypted.

Any array that is made after encryption is enabled is encrypted by default unless you select to not encrypt the array at the time it is created. However, the arrays that existed before encryption was enabled are not encrypted. There is no way to change an array that is not encrypted into one that is encrypted.

If you require that all of the arrays in the storage pool are encrypted, then you can add new encrypted arrays to the storage pool and then remove any non-encrypted arrays from the storage pool. The data on an array that you remove from a storage pool is automatically moved to other arrays or managed disks in the pool.

When an array is removed from a storage pool, the member drives return to candidate status. A new array that is encrypted can be made with these drives.

Procedure: Resetting the superuser password by using the management GUI or CLI

The primary method for resetting the superuser password is to change the password as you log in, with the link on the log-in page. You can also access the service assistant from the technician port to change the password.

If the password reset function is enabled, the log-in page displays a link for resetting the password. You can also use the technician port to access the service assistant.

If the password reset function is not enabled, there is no work-around. You must contact the person who knows the password. The USB flash drive interface also supports resetting the password.

Procedure: Identifying which enclosure or canister to service

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

To prevent loss of access to data or loss of data when you service your system, it is important to be able to identify the correct enclosure or canister when you complete a service action. Each enclosure is

identified by its model type and serial number. Model type and serial number are indicated by labels on the enclosure front and rear.

Depending on the system model, the rear of the enclosure might have different features. But, in general, looking at the rear of a rack:

- Control enclosures contain node canisters that have slots for networking adapters, Ethernet ports, SAS ports, and USB ports. LED indicators are also visible.
- Expansion enclosures have SAS ports and LED indicators and a service port. The model type is shown on a label.

A canister is uniquely identified by the enclosure that it is in and its slot location. This ID is shown as E-C or E|C, where E is the enclosure ID and C is the canister location. On the service assistant, the ID is known as the Panel.

Notes:

- When a node canister is added to a system as a node, it is given a node name and a node ID. The default node name is nodeN, where N is the node ID. The node ID does not represent the slot location of the node. To display node canister details in the management GUI, complete these steps:
 1. Select **Monitoring > System**.
 2. On the **System - Overview** page, select the directional arrow next to the enclosure that contains the node canister that you want to display.
 3. On the **Enclosure Details** page, select the node canister in the graphic to display details for the node canister.
- Use this procedure to identify which enclosure or canister must be serviced, as completing a service action on the wrong canister can lead to loss of access to data or loss of data.

To control the identify LED of an enclosure or online canister in the management GUI, complete these steps:

1. Select **Monitoring > System**.
2. On the **System - Overview** page, select the directional arrow next to the enclosure that contains the node canister that you want to identify.
3. On the **Enclosure Details** page, right-click the node canister in the graphic and select **Identify** from the **Actions** menu.

Alternatively, if a node canister is not online to the system, use the service assistant to control the identify LED.

1. Log in to the service assistant of the node canister to be identified.
2. Click **Identify** at the upper left of the page to control the identify LEDs.

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**

- **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 canister and system using the service assistant.

To obtain the information:

1. Log on to the service assistant.
2. View the information about the node canister to which you connected or another node canister in the enclosure. To change the node for which 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 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 by using a USB flash drive

You can view information about the node canister and system by using a USB flash drive.

About this task

Use any USB flash drive with a FAT32 file system on its first partition.

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

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

Procedure

1. Insert the USB flash drive in the USB port of the node from which you want to collect data.

The node fault light-emitting diode (LED) flashes while information is collected and written to the USB flash drive.

2. Wait until the LED stops flashing before you remove the USB flash drive.

Because the LED is a fault indicator, it might remain permanently on or off.

3. View the results in file `satask_result.html` in a web browser. 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: Understanding the system status from the LEDs

To determine the system status by using the LED indicators on a control enclosure or node, use this procedure.

The control enclosure or the node LEDs, which are visible on the front end cap. These are described in the Control enclosure indicators page in the relevant Hardware Guide..

A detailed view of the system state is provided in the **Monitoring** sections of the management GUI and by the service assistant. If neither the management GUI nor the service assistant is accessible, use this procedure to determine the system status by using the LED indicators on the control enclosure or node.

Complete the following steps to understand when a node canister or node is not participating in the system, and what remedial action to take.

1. Identify the control enclosure or node for the system you are troubleshooting.
2. See Drive indicators and Power supply indicators pages in the relevant Hardware Guide.
3. If at least one power supply unit is providing power to a control enclosure or node, see the Node canister indicators page in the relevant Hardware Guide to check the status of the node canister.
4. Repeat these steps for each node in the system.

Procedure: Finding the status of Ethernet connections

Use this procedure to find the status of Ethernet connections in the system when you cannot connect.

About this task

Ethernet port 1 of each node canister must be connected to an active port on your Ethernet network. You can determine the status of the Ethernet ports by using any of the following methods:

- Connect a personal computer directly to the node by following “Procedure: Accessing the service assistant from the technician port” on page 40. In the service assistant, the status, speed, and MAC address for each port are shown in the **Ports** tab of the **Node Details** table. Any node errors are shown in the **Node Errors** table.
- Use a USB flash drive to obtain node configuration information (see “Procedure: Getting node canister and system information by using a USB flash drive” on page 36). The results file contains the status, speed, and MAC address for each port; whether the node is the configuration node and if any node errors are being reported.
- Examine the LEDs of the Ethernet ports. The Ethernet ports on the left end of the rear of each node canister are 1 Gbps Ethernet ports. For these ports, the link state LED is ON if the link is connected. If optional 10 Gbps Ethernet adapters are fitted, then the green link LED is ON if the link is connected. If optional 2-port 25 Gbps Ethernet adapters are present, the port status is green if the link is up and flashes green with link activity.

Procedure

1. Verify that each end of the cable is securely connected.
2. Verify that the port on the Ethernet switch or hub is configured correctly.
3. Connect the cable to a different port on your Ethernet network.
4. Replace the Ethernet cable.
5. Review any node errors that are reported in the service assistant or on the USB flash drive. Follow fixing node errors for each node error that is reported.
6. Follow the hardware replacement procedure for a node canister.

Procedure: Finding the status of SAS connections

Find the status of the SAS connections between canisters in different enclosures.

About this task

Ensure that the software is active on the node before you begin this procedure.

Procedure

Determine the state of the SAS ports by using one of the following methods:

- In the management GUI, select to **Monitoring > System**.
- On the **System - Overview** page, select **SAS Chain View** to display the SAS ports. SAS ports that are not connected are displayed as offline.
 - It is normal for port 2 of the canisters in the expansion enclosure at the end of a SAS chain to be offline.
 - If no expansion enclosures are connected to a system, it is normal for port 4 of each canister in the control enclosure to be offline.



Attention: The system can identify some SAS cabling errors and log an event to alert you to the error. Go to the **Monitoring > Events** page of the management GUI and identify alerts concerning hardware errors and SAS cabling errors. Run fix procedures in the recommended order.

- To connect a link that is not connected, complete the following actions while checking the link status after each step until the link is connected:
 - a) Ensure that both ends of each SAS cable are correctly inserted into their correct ports, as described in [“Problem: SAS cabling not valid”](#) on page 32.
 - b) Replace the expansion enclosure at one end of the connection.
 - c) Replace the node canister at the other end of the connection.

Procedure: Removing system data from a node canister

You can safely remove system information from a node canister, if you follow the proper guidelines and procedure. The information that is removed includes configuration data, cache data, and location data.

About this task



Attention: Do not remove the system data from a node canister unless instructed to do so by a service procedure. Do not use this procedure to remove the system data from the only online node canister in a system. If the system data is removed or lost from all node canisters in the system, the system is effectively deleted. Attempting a system recovery procedure to restore a deleted system is not guaranteed to recover all of your volumes.

Procedure

1. Log in to the service assistant of the node canister.
2. Use the service assistant node action to hold the node canister in service state.
3. Click **Manage System**, then click **Remove system data** to remove the system data from the node canister.

What to do next

When you want the node canister to be active again, use the service assistant to leave service state. The node canister moves to candidate state, and can be added to the system. If the partner node canister is already active, the candidate node canister is added automatically.

Procedure: Removing a control enclosure from a system

If needed, you can completely remove a control enclosure from a system. When the procedure is finished, no data is retained on the control enclosure. The control enclosure operates like a new installation and can then be added again to another system.

About this task

This procedure is not used as part of any recovery action.



Attention:

- If the system contains only one control enclosure, the system is destroyed when the control enclosure is removed. Do not start this procedure unless you are certain that you want to remove all the volume and configuration data from this system.
- If the system contains several control enclosures (I/O groups), you can continue to operate the system by using the remaining I/O groups while the one control enclosure (I/O group) is removed. However, you must first reconfigure your storage so that Mdisks are not configured in the I/O group that is being removed. Otherwise, volumes that depend on those Mdisks go offline and cannot be restored after the control enclosure is removed.

Procedure

1. Required: If you are removing a control enclosure (I/O group) from an active system, you must safely remove storage in the I/O group from the system.
 - a. Use the management GUI to identify storage pools that contain managed disks (MDisks) that have members in the I/O group to be removed.
 - b. Remove the identified Mdisks from their pool and ensure that sufficient space remains available in the pools. You might need to add Mdisks from other I/O groups or migrate volumes to other pools.
 - c. If the I/O group that you are removing contains spare drives, ensure that sufficient spare drives are defined in other I/O groups to protect array Mdisks in the system. You might need to move or add spare drives to other I/O groups.
 - d. After the Mdisk removal tasks complete, set the state of each drive in the I/O group that is being removed to be *unused*.
 - e. If the I/O group contains expansion enclosures, use the management GUI to remove each expansion enclosure.
2. Start the service assistant on one of the node canisters.

Tip: If a node in the control enclosure is the configuration node for the system, it is quicker to start the service assistant on the non-configuration node first.
3. Use the service assistant to hold the node in *service* state.
4. Use the **Manage System** option to remove the system data from the node.
5. Repeat steps “2” on page 39 through “4” on page 39 on the second node canister in the enclosure.
6. 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 node errors that are detected by node canisters in your system, use this procedure.

About this task

Node errors are reported in the service assistant when a node detects erroneous conditions in a node canister.

Procedure

1. If possible, log into the management GUI and use the monitoring page to run the recommended fix procedure.
 - a) Follow the fix procedure instructions to completion.
 - b) Repeat this step for each subsequent recommended fix procedure.
2. If it is not possible to access the management GUI, or no recommended actions are listed, follow the identified user response for each reported node error.

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 that is shown until you are successful in setting the IP address to the required value.

You can either set an IPv4 address or an IPv6 address as the service IP address of a node. Enter the required address correctly to avoid disabling the node canister access. For more information, see [Configuring VLAN](#).

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.
 - a. Select **Settings > Network** from the navigation.
 - b. Select **Service IP Addresses**.
 - c. 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:
 - a. Make the node canister that you want to configure the current node.
 - b. Select **Change Service IP** from the menu.
 - c. Complete the panel.

Procedure: Accessing the service assistant from the technician port

If a node canister is inaccessible through your administrative network, use this procedure to connect a personal computer directly to the node canister to access the service assistant.

About this task

This procedure starts the initialization tool if the enclosure is not part of a system because the following conditions are true:

- The node canister is in candidate state.

- No system details are configured.
- The partner node is not in active state.

Otherwise, this procedure starts the service assistant.

Procedure

To connect a personal computer directly to the node canister, complete the following steps:

1. Configure DHCP on the Ethernet port of the personal computer to connect to the node canister.

If the personal computer cannot support DHCP, configure static IPv4 address 192.168.0.2 on the laptop port.

2. Connect an Ethernet cable between the port on the personal computer and the technician port.
3. After the notebook computer is connected, if you have DHCP, open a supported browser and browse to address <https://service>.

If you do not have DHCP, open a supported browser and go to static IP address 192.168.0.1.)

The browser is automatically directed to the initialization tool.

Note: If the system cannot be initialized, you are directed to the service assistant interface.

4. Complete the appropriate procedure.
 - If the initialization tool opens, complete the initialization as described in the installation procedure.
 - If the service assistant dialog opens, use it to service the node canister.
5. Log out of the service assistant.
6. Disconnect the Ethernet cable from the technician port.

Procedure: Powering on the system

After installing all hardware components, you must power on the system and check its status.

About this task



Attention: Do not power on the system with any open bays or slots. Open bays or slots disrupt the internal air flow, causing the drives to receive insufficient cooling.

- Every unused drive bay must be occupied by a filler panel.
- Filler plates must be installed in all empty host interface adapter slots.

Procedure

Power on the control enclosure by connecting the power supply units of the enclosure to their power sources, using the supplied power cables.

If the power sources have circuit breakers or switches, ensure that they are turned on. The enclosure does not have power switches.

Note: To provide power failure redundancy, connect the power cords to separate power circuits.

Procedure: Powering on a node canister

To power on a node canister in a control enclosure, you can use the fix procedure running in the management GUI to do this for you or you can reseal the node canisters.

About this task

If a node canister in a control enclosure has been powered off with the GUI or CLI, you can use the fix procedure running in the management GUI to power it back on, or you can reseating the node canisters.

Note: If enclosure input power was removed from the four power supply units, the input power LED of the four PSUs is off. To power on the enclosure in this case, refer to [“Procedure: Powering on the system”](#) on page 41.

Procedure

1. Refer to [“Procedure: Understanding the system status from the LEDs ”](#) on page 37.
2. Locate the ac power (input) and dc power (output) LEDs on each of the power supply units in the enclosure.
If these LEDs are not lit green, refer to [“Procedure: Powering on the system”](#) on page 41.
3. Locate the Power LED on the rear of the node canister.
If it is lit green and slowly flashing, you can use the fix procedure running in the management GUI to power it back on, or reseal the node canister.
The power LED of the canister flashes fast for a few seconds while it starts.
4. Wait a few minutes until the green status LED of the node canister is lit.

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.

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

Use this procedure to rescue a node.

About this task

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

Procedure

1. Ensure that the node you want to reinstall the software on is the current node.

2. Select **Re-install Software**.
3. Select **Rescue from another node**.

Procedure: SAN problem determination

You can solve problems on the system and its connection to the storage area network (SAN).

About this task

SAN failures might cause system volumes to be inaccessible to host systems. Failures can be caused by SAN configuration changes or by hardware failures in SAN components.

The following list identifies some of the hardware that might cause failures:

- Power, fan, or cooling
- Application-specific integrated circuits
- Installed small form-factor pluggable (SFP) transceiver
- Fiber-optic cables

If either the maintenance analysis procedures or the error codes sent you here, complete the following steps:

Procedure

1. If the SAN configuration was changed by changing the Fibre Channel cable connections or switch zoning, verify that the changes are correct and, if necessary, reverse those changes.
2. Verify that the power is turned on to all switches and storage controllers that the system uses, and that they are not reporting any hardware failures. If problems are found, resolve those problems before you proceed further.
3. Verify that the Fibre Channel cables that connect the systems to the switches are securely connected.
4. If you have a SAN management tool, use that tool to view the SAN topology and isolate the failing component.

Replacing a expansion enclosure attachment SAS cable

To replace a faulty expansion enclosure attachment SAS cable with a new one received from CRU / FRU stock, use this procedure.

About this task

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



Attention:

If you need to replace more than one cable, record which two ports, canisters, and enclosures each cable connects, so you can match the connections with the replacement cables. The system cannot operate if the expansion enclosure attachment SAS cabling is incorrect.

Expansion enclosure attachment SAS cables are connected only between SAS port 3 or 4 of a node canister and SAS port 1 of an expansion canister, or between SAS ports 1 and 2 of different expansion canisters.

More information about correct expansion enclosure attachment SAS cabling can be found in the troubleshooting description of a problem with SAS cabling.

Procedure

To replace a SAS cable, complete the following steps.

1. Locate the connector at one end of the SAS cable that is to be removed.
2. Grasp the connector by its blue tag. Pull the tag.
The connector is released and slides out of the port.
3. Repeat steps 2 and 3 on the other end of the SAS cable.
4. To connect the replacement expansion enclosure attachment SAS cable, connect each end to the vacated ports.



Attention: When you insert a SAS connector into a SAS port, ensure that the orientation of the connector matches the orientation of the port before you push the connector into the port.

- The cable connector and socket are keyed and it is important that you have proper alignment of the keys when the cable is inserted.
- Before you insert the connector into the port, ensure that the connector is rotated such that the blue tag is the lowest part.
- Figure 1 on page 44 shows the correct orientation. The blue tab is always below the port for expansion enclosure attachment SAS cables.




Figure 1. Proper orientation for SAS cable connector

- 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.
- When both ends of a SAS cable are correctly connected, the green link LED next to the connected SAS ports are lit.

For more information, see the troubleshooting procedure for finding the status of SAS connections.

Resolving a problem with dual boot drives

Complete the following steps to resolve most problems with dual boot drives.

Before you begin

When the IBM Storage Virtualize software starts, it reads identity information from the system board and compares it with the information stored on the boot drives in order to identify whether a system has had a fault or has had a part replaced.

Specific node errors are produced under the following conditions:

- Node error 543: This error indicates that none of the identification information data stored in the drives or the system board match. IBM Storage Virtualize requires that at least two devices contain the same identity information.
- Node error 545: This error indicates that the identity information on each boot drive match each other but are not the same as the identity information from the system board. If the node

canister was replaced, use the command-line interface (CLI) command **satask chbootdrive -replacecanister** to update the data stored on the boot drives. If the canister was not replaced this error indicates that the node canister and boot drives are from different systems and need to be returned to their original configuration.

- Node error 743: This error indicates that the identity information cannot be read from one of the two boot drives because that drive failed, is missing, or is out of sync with the other boot drive.
- Node error 744: This error indicates that the identity information from one of the boot drives belongs to a different node. If boot drives were swapped between drive slots 1 and 2, node error 744 is produced.
- Node error 745: This error indicates that a single boot drive system has booted with the boot device in slot 2 instead of slot 1.

About this task

An event is displayed in the **Monitoring > Events** pane of the management GUI if the problem produces node error 743, 744, or 745. Run the fix procedure for that event. Otherwise, connect to the technician port to see the boot drive slot information and use the serial number label on the node canister latch to determine the problem. Alternately, you can use the output of the **satask lsbootdrive** command.

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

Procedure

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

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

Note: If the node serial number does not match the node serial number on the system board, a drive slot has a status of `wrong_node`. If the serial number on the node canister latch matches the node serial number on the drive, you can ignore this status.

4. Move any drive that is in the wrong slot back to the correct slot.
5. Reseat the drive in any slot that has a status of `failed`. If the status remains `failed`, replace the drive with one from FRU stock.
6. If the drive slot has status `out of sync` and Yes in the **can_sync** column, then:
 - Use the service assistant GUI to synchronize boot drives, or
 - Use the command-line interface (CLI) command **satask chbootdrive -sync**.
 - If No is displayed in the **can_sync** column, you must resolve another boot drive problem first.

Replacing the node canister:

7. Replace the node canister. Replacing the node canister results in a node error 545.
8. Use the command-line interface (CLI) command **satask chbootdrive -replacecanister** to update the data stored on the boot drives to complete the procedure.

What to do next

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

Resolving a problem with offloaded data transfer

When Microsoft offloaded data transfer (ODX) is enabled on a system, it is possible to encounter problems. These procedures help you address some common issues that might arise.

The following issues might arise when ODX is enabled on a system:

1. ODX is not working for a volume or volumes, or ODX is not getting initiated from Microsoft Windows.
2. ODX performance is not as expected.
3. Existing read/write workload latency increases.

ODX is not working for a volume or volumes

Complete the following steps if ODX is not working for a volume or volumes.

1. Check whether a specific volume is NTFS formatted. Only NTFS volumes can initiate or enable ODX.
2. Verify that both the source and destination volumes that are involved are from the same system disk array. ODX can be initiated across any two volumes that are served by the same system.
3. Check if ODX is enabled, or if any filter does not support it.
 - a. Check the registry value to determine whether ODX is enabled. Run the command **Get-ItemProperty** `hklm:\system\currentcontrolset\control\filesystem -Name "FilterSupportedFeaturesMode"` to obtain the registry value.

For example, on a Windows PowerShell command prompt:

```
PS C:\Users\Administrator> Get-ItemProperty
hklm:\system\currentcontrolset\control\filesystem -Name "FilterSupportedFeaturesMode"
FilterSupportedFeaturesMode : 0
```

If the value is 0, ODX is enabled. If the value is 1, ODX is disabled.

- b. If ODX is disabled, enable it by running the following command:

```
Set-ItemProperty hklm:\system\currentcontrolset\control\filesystem -Name
"FilterSupportedFeaturesMode" -Value 0
```

4. Validate that the file system filter drivers that are attached to the volume support ODX. Certain Windows filters do not support ODX. If these filters are enabled for a specific volume or volumes, Windows does not initiate ODX.
 - a. To validate file system filter driver opt-in status, list all of the file system filter drivers that are attached to the volume on which you want to run ODX.

Open a Windows PowerShell session as an administrator, and then type the following command where *volume* is the drive letter of the volume:

```
Fltmc instances -v volume
```

For example, to check whether the configuration is correct, run the following command where *volume* is the drive letter for the VDisk that is formatted to NTFS:

```
Fltmc instances -v F:
```

The following output is an example of the result:

```
Instances for F: volume:
Filter           Altitude  Instance Name           Frame  SprtFtrs
-----
TSFairShare      400010    TSFairShare Instance    0       00000000 -->
Shows ODX incapability. You need to disable this filter.
PROCMON23        385200    Process Monitor 23 Instance 0       00000003 -->
The "3" in the end means ODX read/write capability.
```

5. Ensure that Windows is initiating ODX. To check that ODX is indeed getting initiated from your Windows host:
 - a. Install Microsoft process monitor.
 - b. Start Microsoft process monitor and start capturing before you attempt ODX.
 - c. When Microsoft process monitor is ready, initiate the operation that initiates ODX.
 - d. Search the process monitor capture for the following commands:

```
FSCTL_OFFLOAD_WRITE  
FSCTL_OFFLOAD_READ
```

Note: Windows initiates ODX only for transfers greater than 256 K.

ODX performance is not as expected

ODX performance is dependent upon various parameters.

- Verify that ODX is getting initiated for the copy operation by using the steps that are outlined in the previous section.
- If the ODX is initiated, but performance does not appear optimal, then make sure that both the source and destination volumes that are involved in offloaded copy:
 - Are served by the same system disk array (across array ODX is not supported).
 - Have an NTFS cluster size greater than or equal to 32 K (that is, greater than or equal to a 32 K allocation unit)
- Offload performance depends upon various parameters such as:
 - Controller side utilization due to other workloads.
 - Whether Windows initiates the offload operations in parallel.

Existing read/write workload latency increases

Workloads other than copy can experience higher latencies if the copy workload is high. The latency occurs because the offload jobs tend to be bigger in size and tend to complete faster. Therefore, they require more controller resources in a specific time slice. If the preference is to have existing workloads not experience extra latency due to copy offload rather than the copy offload benefits, then consider the following options:

1. Revisit the capacity that is planned to accommodate offloaded copies.

Note: Host side CPU or network bandwidth is freed up due to ODX, but it might add to the latency - depending upon the amount of copy work that is offloaded to controller.

2. Disable system-wide ODX by using the CLI.

Resolving a problem with SSL certificates

If you are unable to connect to the management GUI from your web browser and received a Certificate expired or similar error, this information might help you resolve the issue.

Procedure

1. Access the CLI by using ssh.
2. Run the **lseventlog** CLI command.
3. Use the fix procedures for any error events, for example, 2258 - System SSL certificate has expired.
 - a) If the system SSL certificate is expired and no longer valid, connections to the GUI, and service assistant are likely to generate security exceptions.

- b) Check that the system time and date is correct. If the system time or date is incorrect, it can cause the certificate to be incorrectly marked as expired.
- c) Create a new internally signed certificate or install an externally signed certificate, signed by trusted third-party certificate authority.

Note: If you are using an internally signed certificate, you can enable auto-renewal to renew the certificate before it expires.

Resolving a problem with SSL/TLS clients

Changing the security level of the system might cause the web interface, CIM clients, and other SSL/TLS clients to stop working. If any clients stop working, complete the following procedure.

Procedure

1. Wait 5 minutes and try again. The clients might still need to wait for the services to restart.
2. Confirm that the SSL/TLS implementation of the client (for example, the web browser or CIM management tool) is up to date and supports the level of security that is being enforced. If necessary, revert to a weaker SSL/TLS security level in the system and see whether this action resolves the issue.
3. If the problem is a browser problem, check the exact error message that is reported by the browser.

If the error message is `cipher error`, `SSL error`, `TLS error`, or `handshake error`, then the error implies that there is a problem with the secure connection. In this case, ensure that the following conditions are met:

- The browser is up to date.
- All of the supported browsers (Internet Explorer, Firefox, Firefox ESR, and Chrome) support the TLS version that you configured. To determine the TLS version that is configured, use the **lssecurity** command.
- The services support the protocol versions and ciphers for the selected security level.

If there is only a blank screen, it is likely that either the web service needs to restart, or there is a problem unrelated to the security level.

Procedure: Making drives support protection information

You can use this procedure to migrate drives and arrays to pick up support for protection information.

About this task

Drives cannot start by using protection information for I/O requests on demand. They must be validated as having a correct format and general support for the function within the code. The system can validate the format and general support when the drive object is first discovered by the system. The requirement for system validation means that no drive that exists can use protection information on an update from version 730 regardless of use in the configuration. The system can reject a request to make a drive a candidate if the media is not formatted correctly for use with protection information. The process to use protection information on an existing drive is to use the system interface (GUI/CLI) and involves unmanaging and rediscovering the drive to allow the software to reacquire the drive characteristics.

The **lsdrive** view contains the `protection_enabled` field that shows whether a drive is using protection information. Drives and arrays that exist on an update to version 740 do not automatically pick up support for protection information. All newly discovered drives at this code level support protection information. If the system has spare capacity, then migration can proceed an MDisk at a time. Otherwise, the migration to using protection information on drives proceeds drive by drive.

To migrate a MDisk that is using spare storage capacity, complete the following procedure.

Procedure

1. Migrate data off the MDisk.

The data migration can be accomplished by MDisk migration as part of MDisk delete (**rmmdisk**, **lsmigrate**) within a storage pool. You can also use volume mirroring to create an in-sync mirrored copy of each volume in another pool (**addvdiskcopy**). When it is copied (**lsvdisksyncprogress**), delete the original volume copies (**rmvdiskcopy**), and then delete the MDisk (**rmmdisk**) that has no data.

2. Follow the instructions in step “5” on page 49 for all the drives that are now candidates when the MDisk is deleted (see **lsmigrate**).
3. Re-create the array by using the system interface when all old members adopt protection information.
4. If the drive is a member, complete the following steps to adopt protection information on an individual drive.
 - a) Run the **charraymember** command to eject the drive from the array (either immediately with redundancy loss or after an exchange).
 - b) When the drive is no longer a member, follow the instructions in step “5” on page 49 for candidates or spares.
 - c) Repeat for the next member.
5. If the drive is a spare or candidate, complete the following steps:
 - a) Use the management GUI to take the drive offline.
 - b) When the drive is offline, use the system interfaces to change the drive's use to unused.
 - c) The system reacquires the drive and brings it back online, possibly changing the drive ID.
 - d) Attempt to make the drive a candidate.

Depending on the drive, this step might generate error CMMVC6624E, The command cannot be initiated because the drive is not in an appropriate state to perform the task. This step is necessary to run the format command in the next step.

- e) Run the following format command.

```
svctask chdrive -task format drive_id
```

- f) Wait approximately 3 minutes until the drive is online again.
Use **lsdrive drive_id** to see the drive's online/offline status.
- g) Use the system interface to change the drive's use to candidate. If required, use the system interface to change the drive's use to spare.
- h) Enter **lsdrive drive_id** and check that the `protection_enabled` field is yes. This drive can now be used in an array.

Debugging and performance-monitoring statistics for offloaded data transfer

Offloaded data transfer (ODX) infrastructure captures debug and monitoring information for specific ODX modules and makes it available in global and local views.

Debug and monitoring information for specific ODX modules is available in global and local views. The facility to maintain historical data for certain time periods is also available. It currently represents the historical data in eight entries in the global and local views.

Global view

In this view, ODX infrastructure displays the information at the node level. ODX infrastructure provides the memory to represent the module information. It can also capture the historical data for a certain time period. An example of the information is open system interconnect (OSI) command-related information

(populate token, write using token, command failure, token revoked, and other commands). ODX infrastructure provides the following facilities to capture the data as part of the global view:

- MAX - maximum

MAX represents the largest/maximum value along with the time stamp in the base object. Whenever you update the object, it can maintain the history for the last *X* time period if needed. Through the history, information on how many times the MAX is reached on the last *X* time period becomes known.

- MIN - minimum

- CUMULATIVE - cumulative

CUMULATIVE represents the cumulative value along with time stamp in the base object. It can also maintain the history for the last *X* time period. Through the history, the frequency of failures on the last *X* time period becomes known.

Table 4 on page 50 lists the global (node level) parameters.

<i>Table 4. ODX debug global parameters</i>	
Command	Statistic
max pt cmd entry	Maximum simultaneous ODX PT commands arrived at OSI layer
max pt cmd exit	Maximum simultaneous ODX PT commands exited at OSI layer
max wut cmd entry	Maximum simultaneous ODX WUT commands arrived at OSI layer
max wut cmd exit	Maximum simultaneous ODX WUT commands exited at OSI layer
cmd failure cumulative history	Total number of ODX commands failure happened during ODX copy operation
token revoked cumulative history	Total number of ODX tokens revoked happened during ODX copy operation
send ops history by comm	Total number of communication messages that are sent by a node to all nodes as part of ODX copy operation
send kb data history by comm	Total amount (KBs) of cluster-wide data that is sent by a node as part of ODX copy operation
send & rcvd ops history by comm	Represents clustering operations that are generated by a node (Number of messages sent & received)
send & rcvd kb data history by comm	Represents clustering traffic that is generated by a node (Data sent & received in KBs)

Local view

In this view, ODX infrastructure displays the information at the VDisk level. ODX infrastructure provides the following facility to capture the data as part of the local view:

- COPY SPECIAL - cumulative

COPY SPECIAL is specific to the copy job manager. It maintains the history of ODX I/O (read and write) average latency for the last *X* time period. It has a local view only.

Table 5 on page 51 lists the local (VDisk level) parameters.

Table 5. ODX debug local parameters	
Command	Statistic
Average READ latency(ms)	Offload read (ODX) average latency in milliseconds (ms)
Average WRITE latency(ms)	Offload write (ODX) average latency in milliseconds (ms)

iSCSI performance analysis and tuning

This procedure provides a solution for Internet Small Computer Systems Interface (iSCSI) host performance problems while connected to a system and its connectivity to the network switch.

About this task

Some of the attributes and host parameters that might affect iSCSI performance:

- Transmission Control Protocol (TCP) Delayed ACK
- Ethernet jumbo frame
- Network bottleneck or oversubscription
- iSCSI session login balance
- Priority flow control (PFC) setting and bandwidth allocation for iSCSI on the network

Procedure

1. Disable the TCP delayed acknowledgment feature.

To disable this feature, refer to OS/platform documentation.

- VMware: <https://kb.vmware.com/s/article/1002598>.
- Windows: <http://support.microsoft.com/kb/823764>.
- Linux (RHEL): <https://access.redhat.com/solutions/407743>.
- Linux (Other): Use IP utility to enable quickack on host side. e.g ip route show; ip route change 192.168.140.0/24 dev eth7 proto kernel scope link src 192.168.140.55 quickack 1.

The primary signature of this issue: read performance is significantly lower than write performance. Transmission Control Protocol (TCP) delayed acknowledgment is a technique that is used by some implementations of the TCP to improve network performance. However, in this scenario where the number of outstanding I/O is 1, the technique can significantly reduce I/O performance.

In essence, several ACK responses can be combined into a single response, reducing protocol overhead. As described in RFC 1122, a host can delay sending an ACK response by up to 500 ms. Additionally, with a stream of full-sized incoming segments, ACK responses must be sent for every second segment.

2. Enable jumbo frame for iSCSI.

Jumbo frames are Ethernet frames with a size in excess of 1500 bytes. The maximum transmission unit (MTU) parameter is used to measure the size of jumbo frames.

The system supports 9000-bytes MTU. Refer to the **chportethernet** command to enable jumbo frame. This command is disruptive as the link flips and the I/O operation through that port pauses.

The network must support jumbo frames end-to-end to be effective. Send a ping packet to be delivered without fragmentation to verify that the network supports jumbo frames. For example:

- Windows:

```
ping -t <iscsi target ip> -S <iscsi initiator ip> -f -l <new mtu size - packet overhead (usually 36, might differ)>
```

The following command is an example of a command that is used to check whether a 9000-bytes MTU is set correctly on a Windows 7 system:

```
ping -t -S 192.168.1.117 192.168.1.217 -f -l 8964
```

The following output is an example of a successful reply:

```
192.168.1.217: bytes=8964 time=1ms TTL=52
```

- Linux:

```
ping -l <source iscsi initiator ip> -s <new mtu size> -M do <iscsi target ip>
```

- ESXi:

```
ping <iscsi target ip> -I <source iscsi initiator ip> -s <new mtu size - 28> -d
```

3. Verify the switch's port statistic where initiator/target ports are connected to make sure that packet drops are not high.

Review network architecture to avoid any bottlenecks and oversubscription. The network needs to be balanced to avoid any packet drop; packet drop significantly reduces storage performance. Involve networking support to fix any such issues.

4. Optimize and utilize all iSCSI ports.

To optimize system resource utilization, all iSCSI ports must be used.

- Each port is assigned to one CPU, and by balancing the login, one can maximize CPU utilization and achieve better performance. Ideally, configure subnets equal to the number of iSCSI ports on the system node. Configure each port of a node with an IP on a different subnet and keep it the same for other nodes. The following example displays an ideal configuration:

Node 1

Port 1: 192.168.1.11

Port 2: 192.168.2.21

Port 3: 192.168.3.31

Node 2:

Port 1: 192.168.1.12

Port 2: 192.168.2.22

Port 3: 192.168.3.33

- Avoid situations where 50 hosts are logged in to port 1 and only five hosts are logged in to port 2.
- Use proper subnetting to achieve a balance between the number of sessions and redundancy.

5. Ensure that proper bandwidth is given to iSCSI on the network.

You can divide the bandwidth among the various types of traffic. It is important to assign proper bandwidth for good performance. To assign bandwidth for iSCSI traffic, you need to first enable the priority flow control for iSCSI.

Problem: iSCSI backend MDisks are in a degraded or offline state

Complete the steps in the following procedures to debug iSCSI backend sessions.

Note: iSCSI backend sessions are not supported on 100g ports.

1. Run the command **lseventlog** to check the system event log. If the issue is resolved by following the directed maintenance procedures (DMPs), then no further action is required. If the issue is not resolved, complete the following steps to debug iSCSI sessions.

2. Run the command **lsmdisk mdisk_id | mdisk_name** for the MDisk that is in degraded/offline state. In the detailed view, note the values for the fields *preferred_iscsi_port_id* and *active_iscsi_port_id*.
3. Issue the **svcinfo lsiscsistorageport** command, map the values from step 2, and verify the status field.
4. If the status is not full, then obtain a detailed view for that particular session by using the command **svcinfo lsiscsistorageport sessions_row_id**.
5. In the detailed view of **lsiscsistorageport** from step 4, check the values of the field **connected**. It indicates which nodes, which *src_ip*, is having an issue. The *src_port_id* and *node_id* attributes are also available in the detailed view.
6. Using the command **svcinfo lspportethernet**, check the link status for the *src_port_id* and *node_id* tuple.
7. If link status is not active, check the physical connection of the cable/small form-factor pluggable (SFP) transceiver.
8. Check the port status of the particular port on the switch.

Problem: iSCSI backend controllers are in a degraded or offline state

Complete the steps in the following procedures to debug iSCSI backend sessions.

Note: iSCSI backend sessions are not supported on 100g ports.

1. Run the command **lseventlog** to check the system event log. If the issue is resolved by following the directed maintenance procedures (DMPs), then no further action is required. If the issue is not resolved, complete the following steps to debug iSCSI sessions.
2. Run the command **lscontroller controller_id | controller_name** for the MDisk that is in degraded/offline state. In the detailed view, note the values for the field *iscsi_port_id*.
3. Issue the command **svcinfo lsiscsistorageport**, map the values from step 2, and verify the status field.
4. If the status is not full, then obtain a detailed view for that particular session by using the command **svcinfo lsiscsistorageport sessions_row_id**.
5. In the detailed view of **lsiscsistorageport** from step 4, check the values of the field **connected**. It indicates which nodes, which *src_ip*, is having an issue. The *src_port_id* and *node_id* attributes are also available in the detailed view.
6. Using the command **svcinfo lspportethernet**, check the link status for the *src_port_id* and *node_id* tuple.
7. If link status is not active, check the physical connection of the cable/small form-factor pluggable (SFP) transceiver.
8. Check the port status of the particular port on the switch.

Problem: Ethernet host connectivity issues

If you are having problems attaching to the Ethernet hosts, your problem might be related to the network, the system, or the host. If you face a ping issue with the host after changing the MTU of ethernet ports, you need to run **rmip** command to remove the IP address from the port and then consecutively execute a **mkip** command to add the IP address in order to resolve the issue.

For network problems, you can attempt any of the following actions:

- Test your connectivity between the host and system ports.
- Try to ping the system from the host.
- Ask the Ethernet network administrator to check the firewall and router settings.

- Check that the subnet mask and gateway are correct for the system host configuration.

Using the management GUI for system problems, you can attempt any of the following actions:

- View the configured node port IP addresses.
- View the list of volumes that are mapped to a host to ensure that the volume host mappings are correct.
- Verify that the volume is online.

For problems with iSCSI-attached hosts, you can attempt any of the following actions:

- Verify that the host qualified name (IQN) is correctly configured.
- Use operating system utilities (such as Windows device manager) to verify that the device driver is installed, loaded, and operating correctly.
- If you configured the VLAN, check that its settings are correct. Ensure that Host Ethernet port, system Ethernet ports IP address, and Switch port are on the same VLAN ID. Ensure that on each VLAN, a different subnet is used. Configuring the same subnet on different VLAN IDs can cause network connectivity problems.
- For iSCSI discovery or login failure, verify that the target IP and host object are part of the same portset.
- For missing target IP in the discovery records, verify that the target IP and the host object are part of the same portset.

Fibre Channel link failures

When a failure occurs on a single Fibre Channel link, the small form-factor pluggable (SFP) transceiver might need to be replaced.

Before you begin

The following items can indicate that a single Fibre Channel link failed:

- The Fibre Channel status LEDs at the rear of the node canister
- An error that indicates a single port failed.
- You can use **lsportstats** to see statistics about the port, such as error counts and RX power. If **lsportstats** is not available on your system software level, you can get a utility package from Fix Central that installs a **utilportstats** command.

Procedure

Attempt each of these actions, in the following order, until the failure is fixed.

1. Ensure that the Fibre Channel cable is securely connected at each end.
2. Replace the Fibre Channel cable.
3. Replace the SFP transceiver for the failing port on the node.

Note: If the system is supported by both longwave SFP transceiver and shortwave SFP transceiver. You must replace an SFP transceiver with the same type of SFP transceiver. If the SFP transceiver to replace is a longwave SFP transceiver, for example, you must provide a suitable replacement. Removing the wrong SFP transceiver might result in loss of data access.

4. Perform the Fibre Channel switch service procedures for a failing Fibre Channel link. This procedure might involve replacing the SFP transceiver at the switch.
5. Replace the Fibre Channel adapter on the node.
6. If you are using **lsportstats**, run the command to check that the Rx power and Tx power are greater than the low alarm thresholds.

Run the command again after an hour or more to check that the error counts are no longer increasing.

7. Contact your support center for assistance in replacing the canister in the enclosure or node canister.

Validating Ethernet partnerships configuration

You can validate the Ethernet partnerships configuration (short-distance and long-distance).

Validating IP (short-distance and long-distance) partnerships configuration

For configuring IP (short-distance and long-distance) partnership, you need to create a portset, add IP address to that portset, and then, assign this portset to new IP partnership. Hence, to validate this new IP partnership (short-distance and long-distance) configuration is running accurately, certain attribute needs to be checked.

To ensure that new IP (short-distance and long-distance) partnership is functioning correctly, run **lspartnership** command on each system in the configuration and complete the following checks:

Note: IP replication support is disabled on 100g ports.

- Ensure that the partnership is in a Fully Configured state.
- Ensure that backup IP addresses if available, are configured correctly for all the portsets that are used for the partnership.

Resolving a not_present state for an IP (long-distance and short-distance) partnership

The **not_present** state indicates that the remote system is not visible. This state can be caused by a problem with the connectivity between the local and remote system, the remote system is unavailable, or if the partnership was not configured correctly. You can use **lspartnership** command to check the status of partnerships.

Possible configuration mistakes of IP partnerships and several error codes can result in a **not_present** state. The system supports up to 3 IP partnerships and it is important to determine which system in the configuration that has faced connectivity issues. You can use the **Events** page in the management GUI to run dedicated fix procedures for these errors provide specific details on troubleshooting steps. You can also use the **lseventlog** command.

Configuration error that causes the not_present state

The IP addresses of the corresponding portsets from both the system are not mutually reachable. All the configured IP addresses for the portsets used while creating an IP partnership on one system should be reachable to the corresponding IP address on other system.

Complete the following steps to troubleshoot short-distance partnerships using RDMA configuration:

1. Ensuring connectivity between links

Ensure that all IP addresses associated with the link1 and link2 portsets on the primary IBM FlashSystem are connected with all IP addresses associated with the link1 and link2 portsets on the secondary IBM FlashSystem. The same can be validated as follows:

- **By using the GUI:** Go to **Settings > Network > Ethernet Connectivity**. Select the node and click **Show Results**.

or

- **By using the CLI:** Run the CLI command **sainfo lsnodeipconnectivity** on all the nodes of the system.

2. Ensuring resilient configuration and maximum redundancy

For a more resilient configuration and to have maximum redundancy, ensure that IP addresses configured on the ports for both the links are from different nodes.

3. Further troubleshooting if partnership status is not fully_configured

In case the partnership status, is anything but **fully_configured**, further troubleshooting is required to understand why there is change in partnership. Here are some of the partnerships states and ways to troubleshoot:

Partnership is not_present state

An IP partnership can transition to **not_present** state due to multiple reasons and it means that the replication services have come to a halt. Check for alerts, warnings, or errors that are associated with partnership by using the **lseventlog** CLI command or select **Monitoring > Events** tab in the GUI to view the list of events pertaining to the changes that have occurred in the system.

Check whether there are issues with sessions established

Run the command **sainfo lsnodeipconnectivity** on all the nodes of the system to understand if there are any issues with the sessions established. The session status should ideally be **Connected** but apart from this state, other possible states include **Protocol mismatch**, **Degraded**, or **Unreachable**. The output is as follows:

```
CabaroN1:~ # sainfo lsnodeipconnectivity 01-1
status      local_port_id local_vlan local_rdma_type local_ip_addr remote_port_id
remote_vlan remote_rdma_type remote_ip_addr remote_wwnn  remote_panel_name
cluster_id   error_data
Connected:iWARP 3          iWARP      xx.xx.xx.xx 6          iWARP
xx.xx.xx.xx 500507681000110C 01-1      0000020435202218
CabaroN1:~ # sainfo lsnodeipconnectivity 01-2
status      local_port_id local_vlan local_rdma_type local_ip_addr remote_port_id
remote_vlan remote_rdma_type remote_ip_addr remote_wwnn  remote_panel_name
cluster_id   error_data
Connected:iWARP 3          iWARP      xx.xx.xx.xx 6          iWARP
xx.xx.xx.xx 5005076810001143 01-2      0000020435202218
```

4. Check reachability to system

While reference of event logs and DMPs from the GUI are the recommended way to resolve any issue pertaining to the IBM FlashSystem, an alternate way is to verify connectivity to the remote system by checking reachability to the system.

- Check the connectivity to the remote cluster by using the **svctask ping** command:

For IPv4: `svctask ping -srcip4 <source_ip> <destination_ip>`

For IPv6: `svctask ping6 -srcip6 <source_ip> <destination_ip>`

Resolving a mkippartnership command failure

The **mkippartnership** command defines a new partnership (short-distance and long-distance) that is created over intersite WAN links. The **mkippartnership** might fail because of empty portsets and correct IP address not selected.

To fix failure of **mkippartnership** command complete these checks:

- Empty portsets leads to failure of **mkippartnership** command, hence associated IP addresses must be configured and mapped to the corresponding portsets before using them with **mkippartnership** command.
- A system that is running on version 8.4.2.x or later can partner with only one of three systems from version less than 8.4.2.x. The **mkippartnership** command fails if any attempt is made to establish IP partnership with more than one system that is running on version less than 8.4.2.x. Hence, try to restrict the IP partnership with only one system running on version lesser than 8.4.2.x if multiple IP partnerships are intended.
- The **mkippartnership** command fails while creating an IP partnership (short-distance and long-distance), in case either of the systems participating in the partnership fail to specify link1 or link2 on both the sides. Both systems *must* specify either link1 or link2 on both sides to establish a partnership.)
- The **mkippartnership** command fails while creation of IP partnership with a dual portsets. This happens when each portset ID or portset name is specified twice while executing **mkippartnership** command on each system. Hence, use different available portset ID or portset name while executing the **mkippartnership** command on each system.

Even after following the troubleshooting procedure, if the issue still is not resolved, raise a support ticket with IBM for further assistance.

Chapter 6. Recover system procedure

The recover system procedure recovers the entire storage system if the system state is lost from all control enclosure node canisters. The procedure re-creates the storage system by using saved configuration data and is also known as Tier 3 (T3) recovery. The saved configuration data is in the active quorum disk and the latest JSON configuration backup file. For IBM Storage Virtualize for Public Cloud, the saved configuration data is in the system disk and the latest JSON configuration backup file. The recovery might not be able to restore all volume data.



CAUTION: If the system encounters a state where:

- No nodes are active

Do not attempt to initiate a node rescue, contact IBMSupport. If you start the system recovery system procedure while in this specific state, then loss of the JSON configuration backup files can result.

- No nodes are active
- One or more nodes have node errors that require a node rescue, node canister replacement, or node firmware re-installation.

Do not attempt system recovery. Contact IBM Remote Technical Support. If you start the system recovery system procedure while in this specific state, then loss of the JSON backup of the block volume storage configuration can result.



Attention:

- Run service actions only when directed by the fix procedures. If used inappropriately, service actions can cause loss of access to data or even data loss. Before you attempt to recover a system, investigate the cause of the failure and attempt to resolve those issues by using other fix procedures. Read and understand all of the instructions before you complete any action.
- The recovery procedure can take several hours if the system uses large-capacity devices as quorum devices.
- If there are offline arrays after you run the recovery procedure, contact IBM Support.

Do not attempt the recover system procedure unless the following conditions are met:

- All of the conditions are met in [“When to run the recover system procedure” on page 59](#).
- All hardware errors are fixed. See [“Fix hardware errors” on page 59](#).
- All node canisters have candidate status. Otherwise, see step [“1” on page 58](#).
- All nodes must be at the same level of code that the system had before the failure. If any nodes were modified or replaced, use the service assistant to verify the levels of code, and where necessary, to reinstall the level of code so that it matches the level that is running on the other nodes in the system. For more information, see [“Removing system information for nodes with error code 550 or error code 578 using the service assistant” on page 59](#).
- 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 reinstall the level of code so that it matches the level that is running on the other node canisters in the system.
- If the system was using IP quorum for T3 metadata, verify that all the IP quorum applications are running.
- If the system recovery occurs during a non-disruptive system migration, recovery of system data is dependent on the point in the migration process when the system recovery action occurred. For more information, see [“Verifying migration volumes after a system recovery” on page 66](#).

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

1. Preparing for system recovery:

- a. Review the information about when to run the recover system procedure.
- b. Fix your hardware errors and make sure that all nodes in the system are shown in service assistant or in the output from **sainfo lsservicenodes**.
- c. Remove the system information for node canisters with error code 550 or error code 578 by using the service assistant, but only if the recommended user responses for these node errors are followed. See [“Removing system information for nodes with error code 550 or error code 578 using the service assistant”](#) on page 59.
- d. Remove the system information for nodes with error code 550 or error code 578 by using the service assistant, but only if the recommended user responses for these node errors are followed.
- e. For Virtual Volumes (VVols), shut down the services for any instances of Spectrum Control Base that are connecting to the system. Use the Spectrum Control Base command **service ibm_spectrum_control stop**.
- f. Remove hot spare nodes from the system and set them into candidate mode before you start the recovery process. Run the following CLI command to remove the node from the system.

```
satask leavecluster -force spare-node-panel-name
```

Once the node returns in service mode, run the following CLI command to set it into candidate mode.

```
satask stopservice spare-node-panel-name
```

- g. For IBM Storage Virtualize for Public Cloud on Amazon Web Services (AWS), if you run the recovery from a non-configuration node, the system ID might change, and then it is not able to detect its originally managed Amazon Elastic Block Store (EBS). In this case, it is necessary to delete the EBS tags first, and then start the T3 recovery.
 - h. For IBM Storage Virtualize for Public Cloud on Microsoft Azure, if you run the recovery from a non-configuration node, the cluster ID might change, and then it is not able to detect its originally managed Azure disk. In this case, it is necessary to delete the Azure disk tags first, and then start the T3 recovery.
 - i. For IBM Storage Virtualize for Public Cloud on Amazon Web Services (AWS) before you run the **satask t3recovery -prepare** and **svconfig restore** CLI commands, delete the tag values of IBM-SV-cluster-id and IBM-SV-cluster-name and detach the EBS volumes that are managed by the system from the AWS console. Continue to perform the T3 procedure after this action is complete.
 - j. For IBM Storage Virtualize for Public Cloud on Microsoft Azure before you run the **satask t3recovery -prepare** and **svconfig restore** CLI commands, delete the Azure disk tag values of IBM-SV-cluster-id and IBM-SV-cluster-name and detach the Azure disk volumes that are managed by the cluster from the Microsoft Azure console. Continue to perform the T3 procedure after this action is complete.
- ## 2. Running the system recovery. After you prepared the system for recovery and met all the pre-conditions, run the system recovery.

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

3. Completing actions to get your environment operational.

- Recovering from offline volumes by using the CLI.
- Checking your system, for example, to ensure that all mapped volumes can access the host.
- Locking the superuser account again, if the superuser account was locked before the procedure. To lock the superuser account, enter the command:

```
chuser -lock superuser
```

- If two person integrity was enabled before the recovery, superuser was locked, so lock superuser again as described in [Locking superuser](#) step.
- Enabling multifactor authentication for the superuser account again, if the superuser account has multifactor authentication enabled before the procedure. To enable multifactor authentication for the superuser account, enter the CLI command:

```
chsecurity -superusermultifactor yes
```

When to run the recover system procedure

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



Attention: If you experience failures at any time while running the recover system procedure, call the IBM remote technical support. Do not attempt to do further recovery actions, because these actions might prevent support from restoring the system to an operational status.

Certain conditions must be met before you run the recovery procedure. Use the following items to help you determine when to run the recovery procedure:

1. All enclosures and external storage systems are powered up and can communicate with each other.
2. Check that all nodes in the system are shown in the service assistant tool or using the service command: **sainfo lsservicenodes**. Investigate any missing nodes.
3. Check that no node in the system is active and that the management IP is not accessible. If any node has active status, it is not necessary to recover the system.
4. Resolve all hardware errors in nodes so that only node errors 578 or 550 are present. If this is not the case, go to “Fix hardware errors” on page 59.
5. Ensure that all backend storage is administered by the system is present before you run the recover system procedure.
6. If any nodes have been replaced, ensure that the WWNN of the replacement node matches that of the replaced node, and that no prior system data remains on this node.

Fix hardware errors

Before running a system recovery procedure, it is important to identify and fix the root cause of the hardware issues.

Identifying and fixing the root cause can help recover a system, if these are the faults that are causing the system to fail. The following are common issues that can be easily resolved:

- The node is powered off or the power cords were unplugged.
- – For any nodes that are reporting a node error 550, ensure that all the missing hardware that is identified by these errors is powered on and connected without faults.
- If you are not able to restart the system, and if any node other than the current node is reporting node error 550 or 578, you must remove system data from those nodes. This action acknowledges the data loss and puts the nodes into the required candidate state.

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

The system recovery procedure works only when all node canisters are in candidate status. Ensure that the service assistant displays all of the node canisters with the 550 error code. The 550 error code is the expected node error when more than half of the nodes in the system are missing or when the active quorum disk cannot be found. If the service assistant displays any node canisters with error codes 550 or 578 and all the recommended actions have been completed on these nodes, you must remove their system data.

About this task

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

Having used the service assistant to identify the system status and specific error, you will continue to use the service assistant to complete this procedure.

Selecting Change Node in the service assistant tool lists all of the Storage Virtualize nodes that have logged in to the node that is running the tool. Follow these guidelines when performing the recovery procedure:

- The system column of the node table identifies any nodes that are **not** in the system of nodes that must be recovered. Do not remove the system data for these nodes.
- Do not remove system information from any node that has online status, unless directed to do so by remote technical support.
- Do not remove the system data from the first node until you ensure that the following conditions are met:
 - All nodes in the system of nodes are listed in the Change Node part of the service assistant and are in service status with error 550 or 578
 - You have checked the extra node error data for each node to ensure that no other communication or hardware problem is causing the node error.

Procedure

1. In the service assistant tool, select the node with status service and error 550 or 578.
2. Select **Manage System**.
3. Click **Remove System Data**.

Note: Spare nodes do not go into the 878/578 state that active nodes do. As such, the **Manage System** screen does not have the **Remove System Data** button for spare nodes. To remove system data on spare nodes, ssh onto any spare node and run the following commands.

```
satask leavecluster -force
```

```
satask stopservice
```

Failure to remove the cluster state from the spare nodes results in the T3 failing, as the new cluster is unable to find the spare nodes as available candidates.

4. Confirm that you want to remove the system data when prompted.
5. Remove the system data for the other nodes that display a 550 or a 578 error.

All nodes previously in this system must have a node status of Candidate and have no errors listed against them.
6. Resolve any hardware errors until the error condition for all nodes in the system is **None**.
7. Ensure that all nodes in the system of nodes to be recovered display a status of candidate.

Results

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

Running system recovery by using the service assistant

You can use the service assistant or use the **satask t3recovery** command to start recovery when all node canisters that were members of the system are online and have candidate status. You can use the

service assistant to start recovery when all nodes that were members of the system are online and are in candidate status.

Before you begin

Note: Ensure that the web browser is not blocking pop-up windows. If it does, progress windows cannot open.

Before you begin this procedure, read the recover system procedure introductory information; see Chapter 6, “Recover system procedure,” on page 57.

For IBM Storage Virtualize for Public Cloud, you must manually detach volumes and remove tags IBM-SV-cluster-id and IBM-SV-cluster-name for all volumes that are managed by the system.

About this task



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

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

Important: If you have a mixed system configuration that contains different models of enclosures with FlashSystem 9500 enclosures, all recovery operations must be completed on node canisters within the FlashSystem 9500 enclosures to avoid possible licensing errors that prevent a successful recovery of the system.

If the system has USB encryption, run the recovery from any node or node canister in the system that has a USB flash drive that is inserted which contains the encryption key.

If the system contains an encrypted cloud account that uses USB encryption, a USB flash drive with the system master key must be present in the configuration node before the cloud account can move to the online state. This requirement is necessary when the system is powered down, and then restarted.

If the system has key server encryption, note the following items before you proceed with the T3 recovery.

- Run the recovery on a node that is attached to the key server. The keys are fetched remotely from the key server.
- Run the recovery procedure on a node that is not hardware that is replaced or node that is rescued. All of the information that is required for a node to successfully fetch the key from the key server resides on the node's file system. If the contents of the node's original file system are damaged or no longer exist (rescue node, hardware replacement, file system that is corrupted, and so on), then the recovery fails from this node.

If the system uses both USB and key server encryption, providing either a USB flash drive or a connection to the key server (only one is needed, but both will work also) will unlock the system.

If you use USB flash drives to manage encryption keys, the T3 recovery causes the connection to a cloud service provider to go offline if the USB flash drive is not inserted into the system. To fix this issue, insert the USB flash drive with the current keys into the system.

If you use key servers to manage encryption keys, the T3 recovery causes the connection to a cloud service provider to go offline if the key server is offline. To fix this issue, ensure that the key server is online and available during T3 recovery.

If you use both key servers and USB flash drives to manage encryption keys, the T3 recovery causes the connection to a cloud service provider to go offline if none of the key providers are available. To fix this issue, ensure that either the key server is online or a USB flash drive is inserted into the system (only one is needed, but both will work also) during T3 recovery.

Note: Each individual stage of the recovery procedure can take significant time to complete, depending on the specific configuration.

Procedure

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

If you do not know the IP address or if it was not configured, configure the service address in the following way:

- Use the technician port to connect to the service assistant and configure a service address on the node. For more information, see [“Procedure: Accessing the service assistant from the technician port” on page 40](#)

2. Point your browser to the service IP address of one of the node canisters.
3. Log on to the service assistant. For more information, see [“Accessing the service assistant” on page 14](#).
4. 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](#).

5. Follow the online instructions to select a node to perform this action and then complete the recovery system procedure.

- a) Click **Prepare for Recovery**.

The system searches for the most recent backup file and scans quorum disk. If this step is successful, **Preparation Status: Prepare complete** is displayed on the bottom of the page.

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



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

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



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

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

- d) If the quorum time and backup date are correct, click **Recover** to recreate the system.

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

Results

For any nodes that display error code 550 or 578, ensure that all nodes in the system are visible and all the recommended actions are completed before you place them into candidate status. To place a node into candidate status, remove system information for that node canister. Do not run the recovery procedure on different node canisters in the same system. If any nodes display error code 550 or 578, remove system information to place them into candidate status. Do not run the recovery procedure on different nodes in the same system; this restriction includes remote systems.

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

- T3 successful

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

- T3 recovery completed with errors

T3 recovery that is completed with errors: One or more of the volumes are offline because fast write data was in the cache. To bring the volumes online, see [“Recovering from offline volumes by using the CLI” on page 63](#) for details.

- T3 failed

Call the support center. Do not attempt any further action.

Verify that the environment is operational by completing the checks that are provided in [“What to check after running the system recovery”](#) on page 64.

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

If the recovery completes with offline volumes, go to [“Recovering from offline volumes by using the CLI”](#) on page 63.

Recovering from offline volumes by using the CLI

If a Tier 3 recovery procedure completes with offline volumes, then it is likely that the data that is in the write-cache of the node canisters is lost during the failure that caused all of the node canisters to lose the block storage system cluster state. You can use the command-line interface (CLI) to acknowledge that there was data that is lost from the write-cache and bring the volume back online to attempt to deal with the data loss.

About this task

If you run the recovery procedure but there are offline volumes, you can complete the following steps to bring the volumes back online. Some volumes might be offline because of write-cache data loss or metadata loss during the event that led all node canisters to lose cluster state. Any data that is lost from the write-cache cannot be recovered. These volumes might need extra recovery steps after the volume is brought back online.

Note: If you encounter errors in the event log after you run the recovery procedure that is related to offline arrays, use the fix procedures to resolve the offline array errors before you fix the offline volume errors.

Important: For systems that are using data reduction pools, contact [IBM support](#) for assistance in recovering offline volumes.

Example

Complete the following steps to recover an offline volume after the recovery procedure is completed:

1. Delete all IBM FlashCopy function mappings and replication relationships that use the offline volumes.
2. If there are corrupted volumes in a data reduction pool, the user must run the **recovervdiskbysystem** command to recover all volumes.

Note: Use this command only under the supervision of IBM Support personnel.

3. If there are corrupted volumes in a pool, and the volumes are thin-provisioned or compressed, run the following command:

```
repairsevdiskcopy vdisk_name | vdisk_id
```

This command brings the volume back online so that you can attempt to deal with the data loss.

Note: If running the **repairsevdiskcopy** command does not start the repair operation, then use the **recovervdisk** command.

4. If the volume is not a thin-provisioned or compressed volume, and it is outside of a data reduction pool, then run the **recovervdiskbysystem** command. This brings all corrupted volumes back online so that you can attempt to deal with the data loss.
5. Refer to [“What to check after running the system recovery”](#) on page 64 for what to do with volumes that are corrupted by the loss of data from the write-cache.
6. Re-create all FlashCopy mappings and replication relationships that use the volumes.

What to check after running the system recovery

Several tasks must be completed before you use the system.

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

- Parity arrays are likely syncing to restore parity; they do not have redundancy when this operation proceeds.
- Because there is no redundancy in this process, bad blocks might be created where data is not accessible.
- Parity arrays might be marked as corrupted. This identification indicates that the extent of lost data is wider than in-flight I/O; to bring the array online, the data loss must be acknowledged.
- RAID 6 arrays that were degraded before the system recovery might require a full restore from backup. For this reason, it is important to have at least a capacity match spare available.

Be aware of these differences about the recovered configuration:

- FlashCopy mappings are restored as "idle_or_copied" with 0% progress. Both volumes are restored to their original I/O groups. If FlashCopy mappings to volumes exist in a Safeguarded backup location, those FlashCopy mappings are not restored during system recovery. These FlashCopy mappings are typically created automatically by an external scheduler. FlashCopy mappings to new Safeguarded backups are created at the next scheduled time that is defined by the Safeguarded policy.
- System recovery makes current Safeguarded backups invalid. Therefore, as part of the recovery process, all Safeguarded backups are deleted. After recovery is completed, the external scheduler will create new Safeguard backups based on the policy that was created when Safeguarded Copy function was configured.
- Snapshots are not recovered (the VDisks used for snapshots will be deleted asynchronously). The clone and thin-clone volumes are recovered to be volumes of the correct type in a volume group, the volumes are corrupted. When recover VDisk commands are run against these volumes, the corrupted status is fixed but the data is not recovered because the snapshots are no longer available. It is for the user to delete and re-create these volume groups when they have the appropriate snapshots for the repopulation.
- As part of the recovery process, all snapshots are deleted and the current snapshot becomes invalid. After the process is complete, snapshots are created based on the schedule of the policy that was defined when the snapshot was configured.

Note: Snapshots might not be created if volumes are still offline or corrupted in the source volume group.

- As part of recovery process, CHAP secret for iSCSI host can be entered manually and host should be able to make logins.
- As part of the recovery process, partnerships with remote systems are recovered. However, as the system ID of the recovered system has changed, the partnership configuration on remote systems must be updated with the new system ID. To update the new ID on each remote system, the partnership must be in a stopped state. The **chpartnership -newclusterid** CLI command must then be used on each remote system to update the partnership. Once the system ID has been updated the partnerships can be started.
- The system recovers the replication configuration for any volume groups that use policy-based replication. However, a full resynchronization of volume groups occurs when replication is restarted. For more information, see 4040.

- The system ID is different. Any scripts or associated programs that refer to the system-management ID of the 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.
- Volumes with cloud snapshots that were enabled before the recovery need to have the cloud snapshots manually re-enabled.
- If hardware was replaced before the recovery, the SSL certificate might not be restored. If it is not restored and previously used a CA-signed certificate, then the system creates an internally signed certificate signed by the internal root CA. Even if the previous certificate was signed by a trusted third-party CA, the system creates an internally signed certificate. If a self-signed certificate was previously used before the system recovery, then the system creates a self-signed certificate with 30 days of validity.
- The system time zone might not be restored.
- Any volumes being formatted as a system failure occurs are set to the "formatting_corrupt" state by a system recovery and are taken offline. The **recovervdisk** CLI command must be used to recover the volume, synchronize it with a synchronized copy, and bring it back online.
- After the system recovery process completes, the disks are initially set to entire real-capacity. When I/O resumes, the capacity is determined, and is adjusted to reflect the correct value.

Similar behavior occurs when you use the `-autoexpand` option on volumes. The real capacity of a disk might increase slightly, caused by the same kind of behavior that affects compressed volumes. Again, the capacity shrinks down as I/O to the disk is resumed.

- Distributed RAID 1 rebuild in place synchronizes data between data strip mirrors, where possible. This synchronization can be observed through the **lsarraymemberprogress** command.
- If the system recovery occurs during a nondisruptive system migration, recovery of system data is dependent on the point in the migration process when the system recovery action occurred. For more information, see Verifying migration volumes after a system recovery.

Before you use the volumes, complete the following tasks.

- Start the host systems.
- Manual actions might be necessary on the hosts to trigger them to rescan for devices. You can complete this task by disconnecting and reconnecting the Fibre Channel cables to each host bus adapter (HBA) port.
- Verify that all mapped volumes can be accessed by the hosts.
- Run file system consistency checks.

Note: Any data that was in the system write cache at the time of the failure is lost.

- Run file system consistency checks.
- Run the application consistency checks.

For VMware Virtual Volumes (vVols), complete the following tasks.

- After you confirm that the T3 completed successfully, restart the embedded VASA provider service using the command: **satask restartservice -service nginx**.
- Rescan the storage providers from using the vSphere Client.
 - Select **vCSA > Configure > Storage Providers**, select the storage provider and click on the rescan action.

For VMware Virtual Volumes (vVols), also be aware of the following information.

- FlashCopy mappings are not restored for vVols. The implications are as follows.
 - The mappings that describe the VM's snapshot relationships are lost. However, the Virtual Volumes that are associated with these snapshots still exist, and the snapshots might still appear on the vSphere Client. This outcome might have implications on your VMware back up solution.

- Do not attempt to revert to snapshots.
- Use the vSphere Client to delete any snapshots for VMs on a vVol data store to free up disk space that is being used unnecessarily.
- The targets of any outstanding 'clone' FlashCopy mappings might not function as expected (even if the vSphere Client previously reported recent clone operations as complete). For any VMs, which are targets of recent clone operations, complete the following tasks.
 - Complete data integrity checks as is recommended for conventional volumes.
 - If clones do not function as expected or show signs of corrupted data, take a fresh clone of the source VM to ensure that data integrity is maintained.

Verifying migration volumes after a system recovery

If system recovery occurred while volumes were being migrated between two systems with non-disruptive volume migration, volumes are recovered based on where in the migration process that the system recovery occurred.

After a system recovery, use the following information to determine recovery actions for volumes that are associated with non-disruptive volume migration.

Table 6. Volume migration after a system recovery		
When the system recovery occurred during Non-disruptive volume migration	System recovery actions	User actions
System recovery on the source system before the switch	• The migration relationship is deleted on the source system. • The partnership is removed on the source system.	1. Delete the auxiliary volume from the target system in the migration relationship with the management GUI or the rmvolume command. 2. Recover the source volume, potentially from tape. 3. If you want to recover all the data in migration relationships that are <i>consistent_synchronized</i> at the time of the system recovery, contact support. This recovery can be a lengthy process.
System recovery on source system after the switch	• The migration relationship is deleted on the source system. • The partnership is removed on the source system. • The UUID of the source volume is changed. • Any host mappings to the source volume in the migration relationship are removed.	1. Delete the source volume with the management GUI or the rmvolume command. 2. Delete the migration relationship on the target system with the management GUI or the rmrelationship command.

Table 6. Volume migration after a system recovery (continued)

When the system recovery occurred during Non-disruptive volume migration	System recovery actions	User actions
System recovery on the target system before the switch	• The migration relationship is deleted on the target system. • The partnership is removed on the target system. • The UUID target volume is changed. • Any host mappings to the target volume are removed.	1. Delete the target volume with the management GUI or rmvolume command. 2. Remove the migration relationship from the source system with the management GUI or the rmrelationship command.
System recovery on the target system after the switch	• The migration relationship is deleted on the target system. • The partnership is removed on the target system.	1. Delete the source volume with the management GUI or the rmvolume command. 2. Recover the target volume, potentially from tape. 3. If you want to recover all the data in migration relationships that are <i>consistent_synchronized</i> at the time of the system recovery, contact support. This recovery can be a lengthy process.

Chapter 7. Backing up and restoring the system configuration

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

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

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

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

Note:

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

Configuration backup and restore process

Configuration backup is the process of extracting configuration settings from a system and writing it to disk. The configuration restore process uses backup configuration data files for the system to restore a specific system configuration. Restoring the system configuration is an important part of a complete backup and disaster recovery solution.

Only the data that describes the system configuration is backed up. You must back up your application data by using the appropriate backup methods.

To enable routine maintenance, the configuration settings for each system are stored on each node. If power fails on a system or if a node in a system is replaced, the system configuration settings are automatically restored when the repaired node is added to the system. To restore the system configuration in a disaster (if all nodes in a system are lost simultaneously), plan to back up the system configuration settings to tertiary storage. You can use the configuration backup functions to back up the system configuration.

The virtualization map is stored on the quorum disks of external MDisks, and is accessible to every node. If the system has only one control enclosure, the virtualization map can be stored on internal drive quorum disks.

For complete disaster recovery, regularly back up the business data that is stored on volumes at the application server level or the host level.

Backing up the system configuration using the CLI

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

Before you begin

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

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

About this task

The backup feature of the **svconfig** CLI command is designed to back up information about your system configuration, such as volumes, storage pools, and nodes. Any data on the volumes is *not* backed up. Application data must be backed up using appropriate methods.

Configuration data must be regularly backed up, and after any significant changes to the system configuration.

Note: The system automatically creates a backup of the configuration data each day at 1 AM. This backup is written to `/dumps/svconfig/svc.config.backup.json` on the configuration node.

Use these instructions to generate a manual backup at any time. If a severe failure occurs, both the configuration of the system and application data might be lost. The backup of the configuration data can be used to restore the system configuration to the exact state it was in before the failure, except partnerships and replication policies. In some cases, it might be possible to automatically recover the application data. This restore can be attempted with the Recover System Procedure, also known as a Tier 3 (T3) procedure. To restore the system configuration without attempting to recover the application data, use the Restoring the System Configuration procedure, also known as a Tier 4 (T4) recovery. Both of these procedures require a recent backup of the configuration data.

Complete the following steps to backup your configuration data:

Procedure

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

```
svconfig backup
```

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

```
CMMVC6155I SVCCONFIG processing completed successfully
```

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/svconfig` directory of the configuration node canister.

[Table 7 on page 70](#) describes the three files that are created by the backup process:

Table 7. Files created by the backup process	
File name	Description
<code>svc.config.backup.json</code>	Contains your configuration data.
<code>svc.config.backup.sh</code>	Contains the names of the commands that were issued to create the backup of the system.
<code>svc.config.backup.log</code>	Contains details about the backup, including any reported errors or warnings.

Note: The `svc.config.backup.sh` file name is not created until the prepare for recover and restore process is complete.

2. 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 SVCCONFIG processing completed successfully
```

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

3. Keep backup copies of the files outside the system to protect them against a system hardware failure. Copy the backup files off the system to a secure location; use scp command line. For example:

```
pscp -unsafe superuser@cluster_ip:/dumps/svcconfig/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.

Restoring the system configuration

Use this procedure to restore the system configuration only if the recover system procedure fails or if the data that is stored on the volumes is not required. This procedure is also known as Tier 4 (T4) recovery.

Before you begin

Important: Before running a T4 procedure, contact [IBM support](#) for assistance.

This configuration restore procedure is designed to restore information about your configuration, such as volumes, storage pools, and nodes. The data that you wrote to the volumes is not restored. To restore the data on the volumes, you must restore the data backed-up on a separate storage. Therefore, you must have a backup of this data before you follow the configuration recovery process.

Note: All configuration related to replication and FlashSystem grid is not restored. The previous replication settings are included in the configuration backup for reference.

If the system uses encryption and USB flash drives to manage encryption keys, then at least 3 USB flash drives need to be installed in the node canister USB ports for the configuration restore to work. The 3 USB flash drives must be inserted into the node canister or enclosure from which the configuration restore commands are run. Any USB flash drives in other nodes or enclosures (that might become part of the system) are ignored. On systems with fewer than 3 USB ports, encryption must be enabled manually later on in the recovery. On these systems, follow the instructions that are displayed on screen to manually enable encryption during step 17 when the configuration restore is prepared. If you are not recovering an encrypted transparent cloud tiering configuration, the USB flash drives do not need to contain any keys. They are for generation of new keys as part of the restore process. If you are recovering an encrypted transparent cloud tiering configuration, the USB flash drives must contain the previous set of keys to allow the current encrypted data to be unlocked and re-encrypted with the new keys.

During recovery, the system creates a new internal root CA and system certificates. If the system uses Thales key servers to manage encryption keys, the new root certificate must be exported and installed on the key servers before the configuration restore operation prepares successfully. If the system uses IBM key servers, then the new system certificate must be exported and installed on the key servers. If the previous system was using a certificate that is signed by a third-party CA, then it might also be necessary to get the new system's certificate signed.

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

- The Security Administrator role is associated with your username 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.
- No zoning changes were made on the Fibre Channel fabric that would prevent communication between the system and any storage controllers that are present in the configuration.
- If the system uses encryption, the recovery procedure generates new encryption keys. Make sure that key servers are online and USB flash drives are installed in the system. Any existing encryption key files on USB flash drives correspond to the previous system and are no longer required unless using transparent cloud tiering.

- If the system uses encryption with transparent cloud tiering, at least one USB flash drive that contains the encryption key file for the previous system must be installed in the system.

One of the nodes previously in I/O group zero must perform the restoring the system configuration. For example, **property name="IO_group_id" value="0"**. The remaining nodes canister must 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, canister 1 might perform the restoration.

Before you begin, hardware recovery must be complete. The following hardware must be operational: hosts, system enclosures, internal flash drives, and expansion enclosures (if applicable), the Ethernet network, the SAN fabric, and any external storage systems (if applicable).

After you finish, you must manually lock the superuser account again if the superuser account was locked before backing up and restoring the system configuration. To lock the superuser account, enter the CLI command:

```
chuser -lock superuser
```

After you finish, if the system is using multifactor authentication then the new system certificate must be exported and installed as a signer certificate in IBM Security Verify, by using the system's id_alias as the friendly name. To export the new system certificate, enter the CLI command:

```
chsystemcert -export
```

To find the id_alias of the system, enter the CLI command:

```
lssystem | grep id_alias
```

About this task

You must regularly back up your configuration data and your application data to avoid data loss. If a system is lost after a severe failure occurs, both configuration for the system and application data is lost. You must restore the system to the exact state it was in before the failure, and then recover the application data.

During the restore process, the nodes and the storage enclosures are restored to the system, and then the MDisk and the arrays are re-created and configured. If multiple storage enclosures are involved, the arrays and MDisk are restored on the proper enclosures based on the enclosure IDs.

Important:

- There are two phases during the restore process: prepare and execute. You must not change the fabric or system between these two phases.
- For systems that contain nodes that are attached to external controllers virtualized by iSCSI, all nodes must be added into the system before you restore your data. Additionally, the system mkp settings and iSCSI storage ports must be manually reapplied before you restore your data.
- For VMware vSphere Virtual Volumes (sometimes referred to as vVols) environments, after a T4 restoration, Virtual Volumes must be enabled on the system again and reconfigured by using the VMware vCenter.
- 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.
- If the system uses encryption, the restore procedure generates new encryption keys. Make sure that key servers are online and USB flash drives are installed in the system. Any existing encryption key files on USB flash drives correspond to the previous system and are no longer required, unless using transparent cloud tiering.
- If the system uses encryption with transparent cloud tiering, the USB flash drive and key servers containing the encryption key for the previous system must be installed in the system before recovery.

If the original encryption key is not available, and the system data is encrypted in the cloud provider, then the data in the cloud is not accessible. If the system contains an encrypted cloud account that is configured with both USB and key server encryption, the encryption keys from both need to be available at the time of a T4 recovery.

- If the system contains an encrypted cloud account that uses USB encryption, a USB flash drive with the system master key must be present in the configuration node before the cloud account can move to the online state. This requirement is necessary when the system is powered down, and then restarted.
- If a system management IP is configured on a port with MTU 9000, then during T4 recovery, follow this procedure:
 1. Create a system using the other IP address and port than the actual system management IP address and port.
 2. Change the MTU of the ethernet port to 9000 on the actual port of the system management IP address.
 3. Configure the system management IP address on the actual port of the system management IP address.
 4. Delete the first system management IP address using **rmip** CLI command.
- After a T4 recovery, cloud accounts are in an offline state. It is necessary to re-enter the authentication information to bring the accounts back online.
- If you use USB flash drives to manage encryption keys, the T4 recovery causes the connection to a cloud service provider to go offline if the USB flash drive is not inserted into the system. To fix this issue, insert the USB flash drive with the current keys into the system.
- If you use key servers to manage encryption keys, the T4 recovery causes the connection to a cloud service provider to go offline if the key server is offline. To fix this issue, ensure that the key server is online and available during T4 recovery.
- If you use both key servers and USB flash drives to manage encryption keys, the T4 recovery causes the connection to a cloud service provider to go offline if the key server is offline. To fix this issue, ensure that both the key server is online and a USB flash drive is inserted into the system during T4 recovery.
- After a T4 recovery, volumes with cloud snapshots that were enabled before the recovery need to have the cloud snapshots manually reenabled.
- If Fibre Channel port masks are required on the system, these port masks need to be manually reconfigured after a T4 recovery.
- The T4 recovery creates a new root CA, so the root certificate must be exported from the system and added to trust stores.
- FlashCopy mappings and host mappings that have inconsistent ownership groups are not restored.
- If you have an SNMP server that is configured with a certificate in the trust store, then these settings are not recovered on a T4 and the trust store must be manually set-up again after the recovery process.
- SNMP passphrases are not recovered during a T4 and these will have to be manually set back after the recovery by using the GUI or **chsnmpserver** command.

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, follow this procedure.
 - 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) From the **Home** page, put the node canister into service state if it is not already in that state.
 - d) Select **Manage System**.
 - e) Click **Remove System Data**.
 - f) Confirm that you want to remove the system data when prompted.
 - g) Exit service state from the **Home** page. The 550 or 578 errors are removed, and the node appears as a candidate node.
 - h) 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 by 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. Complete 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**.
 3. Create a system by using the technician port.
 4. 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.

Once the setup wizard finishes, make no other configuration changes.



Warning: If you use the management GUI for the initial setup to restore the system configuration, check if a default call home email user was created. If it was created, delete the default call home email-user for the T4 system recovery to proceed successfully.

5. If you set up email notification in the setup wizard, you must now remove that email user and server so that the original configuration can be restored.

Issue the following CLI command to remove the new email user:

```
rmemailuser 0
```

Issue the following CLI command to remove the new email server:

```
rmemailserver 0
```

6. From the management GUI, click **Access > Users** and configure an SSH key for the superuser.
7. By default, the newly initialized system is created in the storage layer. The layer of the system is not restored automatically from the configuration backup JSON file. If the system you are restoring was previously configured in the replication layer, you must change the layer manually now. Refer to the System layers topic that is located under Product overview in the IBM Documentation for your product for more information.
8. If the clustered system was previously configured as replication layer, then use the **chsystem** command to change the layer setting.
9. Identify the configuration backup file that you want to restore.

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

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

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

- a) From the management GUI, click **Settings > Support > Support Package**.
- b) Expand **Manual Upload Instructions** and select **Download Support Package**.
- c) On the **Download New Support Package or Log File** page, select **Download Existing Package**.
- d) For each node (canister) in the system, complete the following steps:
 - i) Select the node to operate on from the selection box at the top of the table.
 - ii) Find all the files with names that match the pattern `svc.config.*.json*`.
 - iii) Select the files and click **Download** to download them to your computer.

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

10. Copy onto the system the JSON backup file from which you want to restore.

```
pscp full_path_to_identified_svc.config.file
superuser@cluster_ip:/dumps/svcconfig/svc.config.backup.json
```

11. If the system contains any iSCSI storage controllers, these controllers must be detected manually now. The nodes that are connected to these controllers, the iSCSI port IP addresses, and the iSCSI storage ports must be added to the system before you restore your data.

Note: If the system contains only Fibre Channel storage controllers, proceed to the next step.

Note: For a stretched topology, after you run the **addnode** command, change the sites of all of the nodes added in the system. For example,

```
chnodecanister -site site_id node_id/node_name
```

- a) Run the following command to change the replication layer.

```
chsystem -layer replication
```

- b) To restore iSCSI initiator port configuration, use the **chportethernet** command. All the ethernet ports capable of iSCSI backend i.e. storage flag set to yes needs to be configured. Find out the MTU of ethernet port for which "storage" is set to yes in the `node_ethernet_port` section from the backup configuration file and restore the port MTU.

- i) To restore port MTU, determine `iogrp_id` from the `node_id`, `port_id`, and `mtu` of the port for which "storage" is set to yes from the configuration backup file, and run the following command:

```
chportethernet -iogrp iogrp_name_or_id -mtu mtu port_id
```

where *mtu* is the MTU of the port, *iogrp_name_or_id* is the name or ID of the I/O group, *port_id* is the ID of the port.

Complete step i for all (earlier configured) Ethernet ports that belong to storage in the `node_ethernet_port` section from the backup configuration file.

- c) To restore iSCSI initiator port IP addresses, use the **mkip** command. All the IP addresses belonging to the storage portset (i.e. portset3) needs to be configured. Find out the IP addresses whose "portset_name" property matches with portset3 in the `node_ethernet_ip` section from the backup configuration file and restore the IP addresses.

- i) To restore IP address, determine `port_ID`, `node_name`, `portset_name`, IP address, prefix and vlan of the IP address that belongs to `portset3` from the configuration backup file, and run the following command:

```
mkip -node node_id_or_name -port port_id -portset portset_id_or_name -ip ip_address -prefix prefix -vlan vlan
```

Complete step i for all (earlier configured) IP ports that belong to `portset3` in the `node_ethernet_ip` section from the backup configuration file.

- d) Next, detect, and add the iSCSI storage port candidates by using the **`detectiscsistorageportcandidate`** and **`addiscsistorageport`** commands. Make sure that you detect the iSCSI storage ports and add these ports in the same order as you see them in the configuration backup file. If you do not follow the correct order, it might result in a T4 failure. Step c.i must be followed by steps c.ii and c.iii. You must repeat these steps for all the iSCSI sessions that are listed in the backup configuration file exactly in the same order.

- i) To detect iSCSI storage ports, determine `src_port_id`, `IO_group_id` (optional, not required if the value is 255), `target_ipv4/target_ipv6` (the target IP that is not blank is required), `iscsi_user_name` (not required if blank), `iscsi_chap_secret` (not required if blank), and `site` (not required if blank) from the configuration backup file, run the following command:

```
svctask detectiscsistorageportcandidate -srcportid src_port_id -iogrp IO_group_id -targetip/targetip6 target_ipv4/target_ipv6 -username iscsi_user_name -chapsecret iscsi_chap_secret -site site_id_or_name
```

Where `src_port_id` is the source Ethernet port ID of the configured port, `IO_group_id` is the I/O group ID or name being detected, `target_ipv4/target_ipv6` is the IPv4/IPv6 target iSCSI controller IP address, `iscsi_user_name` is the target controller username being detected, `iscsi_chap_secret` is the target controller chap secret being detected, and `site_id_or_name` is the specified id or name of the site being detected.

- ii) Match the discovered `target_iscsiname` with the `target_iscsiname` for this particular session in the backup configuration file by running the **`lsiscsistorageportcandidate`** command, and use the matching index to add iSCSI storage ports in step c.iii.

Run the **`svcinfo lsiscsistorageportcandidate`** command and determine the id field of the row whose `target_iscsiname` matches with the `target_iscsiname` from the configuration backup file. This is your **`candidate_id`** to be used in step c.iii.

- iii) To add the iSCSI storage port, determine `IO_group_id` (optional, not required if the value is 255), `site` (not required if blank), `iscsi_user_name` (not required if blank in backup file), and `iscsi_chap_secret` (not required if blank) from the configuration backup file, provide the `target_iscsiname_index` matched in step c.ii, and then run the following command:

```
addiscsistorageport -iogrp iogrp_id -username iscsi_user_name -chapsecret iscsi_chap_secret -site site_id_or_name candidate_id
```

Where `iogrp_id` is the I/O group ID or name that is added, `iscsi_user_name` is the target controller user name that is being added, `iscsi_chap_secret` is the target controller chap secret being added, and `site_id_or_name` specified the ID or name of the site that is being added.

- iv) If the configuration is a stretched system, the controller name and site needs to be restored. To restore the controller name and site, determine `ccontroller_name` and controller `site_id/name` from the backup json file by matching the `inter_WWPN` field with the newly added iSCSI controller, and then run the following command:

```
chcontroller -name controller_name -site site_id/name controller_id/name
```

Where `controller_name` is the name of the controller from the backup json file, `site_id/name` is the ID or name of the site of iSCSI controller from the backup json file, and `controller_id/name` is the ID or current name of the controller.

12. If the system is using Lightweight Directory Access Protocol (LDAP) as the remote authentication service with an administrator password configured, the password must be restored manually before you restore your data. The following example shows the command to configure the LDAP administrator user name and password:

```
svctask chldap -username ldap_username -password 'administrator_password'
```

13. . If the system is using multifactor authentication with IBM Security Verify, the OpenID Connect and API credentials must be restored manually before you restore your data. The following example shows the command to configure the parameters Open ID Client ID, Open ID Client Secret, API Client ID and API Client Secret:

```
svctask chauthmultifactorverify -openidclientid 'clientid' -openidclientsecret 'clientsecret' -clclientid 'clientid' -clclientsecret 'clientsecret'
```

14. If the system is using multifactor authentication with Duo Security, the OpenID Connect and API credentials must be restored manually before you restore your data. The following example shows the command to configure the parameters Open ID Client ID, Open ID Client Secret, Integration key and Secret key:

```
chauthmultifactorduo -hostname 'host_name' -integrationkey 'unix_application_integration_key' -secretkey 'unix_application_secret_key' -openidclientid 'web_SDK_client_id' -openidclientsecret 'web_SDK_client_secret'
```

15. If the system is using single sign-on, the OpenID Connect credentials must be restored manually before you restore your data. The following example shows the command to configure the parameters Client ID and Client Secret:

```
svctask chauthsinglesignon -clientid 'clientid' -clientsecret 'clientsecret'
```

16. If the system has key server encryption, the new certificate must be exported by using the **chsystemcert -export** command, and then installed on all key servers in the correct device group before you run the T4 recovery. The device group that is used is the one in which the previous system was defined. It might also be necessary to get the new system's certificate signed.
17. Issue the following CLI command to compare the current configuration with the backup configuration data file:

```
svcconfig restore -prepare
```

This CLI command creates a log file in the /dumps/svcconfig directory of the configuration node. The name of the log file is svc.config.restore.prepare.log.

Note: It can take up to a minute for each 256-MDisk batch to be discovered. If you receive error message CMMVC6200W for an MDisk after you enter this command, all the managed disks (MDisks) might not be discovered yet. Allow a suitable time to elapse and try the **svcconfig restore -prepare** command again.

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

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

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

20. Check the log file for errors.

- If you find errors, correct the condition that caused the errors and reissue the command. You must correct all errors before you can proceed to step “21” on page 77.
- If you need assistance, contact the support center.

21. Issue the following CLI command to restore the configuration:

```
svcconfig restore -execute
```

This CLI command creates a log file in the /dumps/svcconfig directory of the configuration node. The name of the log file is svc.config.restore.execute.log.

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

```
pscp superuser@cluster_ip:/dumps/svcconfig/svc.config.restore.execute.log  
full_path_for_where_to_copy_log_files
```

23. Open the log file from the server where the copy is now stored.
24. Check the log file to ensure that no errors or warnings occurred.

Note: You might receive a warning that states that a licensed feature is not enabled. This message means that after the recovery process, the current license settings do not match the previous license settings. The recovery process continues normally and you can enter the correct license settings in the management GUI later.

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

```
IBM_MTM:your_cluster_name:superuser>
```

25. After the configuration is restored, verify that the quorum disks are restored to the MDisks that you want by using the **lsquorum** command. To restore the quorum disks to the correct MDisks, issue the appropriate **chquorum** CLI commands.

Note: If IP Quorum was enabled on the system, it is not recovered automatically as the system certificate is regenerated. It is necessary to manually re-enable IP Quorum by downloading a java application from the **Settings>System>IP Quorum** tab in the GUI, and then installing the application on the host server.

26. Run **svcconfig** command to delete backup configuration files.
27. If the system is using two person integrity, after restore, enable two person integrity manually. It cannot be enabled automatically. To enable manually, see Two person integrity .

What to do next

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

```
svcconfig clear -all
```

Deleting backup configuration files by using the CLI

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

About this task

Complete the following steps to delete backup configuration files:

Procedure

Issue the following CLI command to erase all of the files that are stored in the /dumps/svcconfig directory:

```
svcconfig clear -all
```

Chapter 8. Enabling concurrent maintenance

To enable concurrent maintenance, a system is configured as pairs of nodes. Each pair is called an I/O group. If one system node is shut down or disconnected for maintenance, the other node can keep the I/O group operational.

With concurrent maintenance, all field-replaceable units (FRUs) can be removed, replaced, and tested on one system node while the network and host systems are powered on and doing productive work through the second node.



Attention: Do not remove the power from any nodes unless the procedures instruct you to do so.

Verify that concurrent maintenance is enabled before you shut down a node that is part of a system or when you delete the node from a system. To do so, complete the following checks.

1. Confirm that no volumes have dependencies on the node by completing these steps:
 - a. In the management GUI, select **Monitoring > System**.
 - b. On the **System -- Overview** page, use the directional arrow near the enclosure that contains the node canister to open the **Enclosure Details** page.
 - c. Under **Rear View** of the system, right click the canister and select **Dependent Volumes** from the **Actions** menu you display all volumes that become unavailable to hosts if the canister is powered off. You can also use the node parameter with the **lsdependentvdisks** CLI command to view dependent volumes.

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

2. Ensure that the host multipathing device drivers can fail over to the partner node.

Some host multipathing device drivers take a while to update after changes are made on the fabric. Do not shut down a node or delete the node from a cluster if the partner node in the I/O group to which the node belongs has not been online for more than 30 minutes.

If possible, check the status of the host multipathing device drivers before you shut down a node to ensure that the device drivers can fail over to the partner node.

When you shut down the node, follow the procedure that is described in the Procedure: Powering off a node canister page of the relevant Hardware Guide.



Attention: Do not power off any expansion enclosures when you power off a node.

When you delete a node from the clustered system, retain the node information that is described in Deleting a node from a clustered system by using the management GUI. This information helps you avoid data corruption when you add the node back to the system. The topic describes how to ensure that the multipathing device driver does not rediscover any paths that are manually removed. Other considerations about dependent volumes are also provided.

For more information about working with dependent volumes, see the following topics:

- **Dependent volumes**
- Chapter 10, “Deleting a node from a system by using the CLI,” on page 83 and the **lsnodecandidate (SAN Volume Controller)** command description
- **I/O group**

Chapter 9. Removing a control enclosure and its expansion enclosures

You can use the management GUI and the command-line interface (CLI) to remove the control enclosure and all its attached expansion enclosures.

About this task

Important: Before you remove a control enclosure and its expansion enclosures, note that all FlashCopy mappings on the control enclosure are lost if their dependencies are not resolved before the controller is removed.

1. Make sure that your volume data is backed up.
2. Determine the volumes, hosts, and arrays that are affected by the removal:
 - a. All the volumes that belong to the I/O group and the hosts that require access to them.
 - b. Which arrays have member drives in the enclosures.
 - c. The list of volumes that depend on these arrays.

You can use the management GUI and command-line interface (CLI) filtering capabilities to obtain this information.

3. Decide which volumes are still necessary, and plan for their migration. Consider both available storage pool capacity and the system configuration limits.
4. If FlashCopy copies with their associated bitmaps are on the control enclosure that you plan to remove, you must first resolve the dependencies. Otherwise, the bitmaps are lost, which causes target volumes to go offline. For help, refer to the **rmnodecanister** command before you remove the second node canister.

The procedure provides the necessary steps for removing volumes.

Procedure

1. Delete any volumes that are no longer needed and are dependent on the enclosure that you plan to remove.
 2. Move any remaining volumes that are accessed through the I/O group to a different I/O group.
Any hosts that use these volumes must be reconfigured so they retain access.
 3. Migrate any remaining volume data from the arrays in the I/O group:
 - Migrate the entire volume copy to another storage pool. This method is preferable when most of the arrays in the current storage pool are from the I/O group.
 - Migrate the volume data extents to other arrays or managed disks (MDisks) within the current storage pool. Deleting an array forces migration of all volume extents that are allocated on it to other arrays in the pool. You can also migrate extents for each volume individually. This method is preferable when most of the storage capacity in the pool is provided by other I/O groups or external controllers.
- Important:** Wait for the migrations to complete. You can use the management GUI to monitor the migration progress.
4. Delete any remaining arrays that are formed from drives in the I/O group.
If any arrays still show dependent volumes, repeat the previous step and migrate the data for these volumes.
 5. Mark all the drives (including any configured as spare drives) in the enclosures to be removed as unused.
 6. Ensure no I/O is running to any volumes that are from the I/O group of which the enclosure is part.

7. Use the **chiogrp** command to put the I/O group into maintenance mode.
8. Unmanage and remove enclosures by using the management GUI. Remove the expansion enclosures first and the control enclosure last. Select **Monitoring > System**. On the **System - Overview** page, select the directional arrow next to the enclosure that you are removing to open the **Enclosure Details** page. Select **Enclosure Actions > Remove**.
9. Remove the node canisters by using the **rmnodecanister** command or the management GUI.
10. The control enclosure and attached expansions are now removed from the system. Remove power from the enclosures to shut them down.
11. Use the **chiogrp** command to take the I/O group out of maintenance mode.

Chapter 10. Deleting a node from a system by using the CLI

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

Before you begin

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

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



Attention:

- If you are removing a single node and the remaining node in the I/O group is online, the data can be exposed to a single point of failure if the remaining node fails.
- If both nodes in the I/O group are online and the volumes are already degraded before you delete the node, redundancy to the volumes is already degraded. If the **-force** option is used, removing a node might result in loss of access to data, and data loss might occur.
- Removing the last node destroys the system. Before you delete the last node in the system, ensure that you want to destroy the system.
- When you delete a node, you remove all redundancy from the I/O group. As a result, new or existing failures can cause I/O errors on the hosts. The following failures can occur:
 - Host configuration errors
 - Zoning errors
 - Multipathing software configuration errors
- If you are deleting the last node in an I/O group and volumes are assigned to the I/O group, you cannot delete the node from the system if the node is online. You must back up or migrate all data that you want to save before you delete the node. If the node is offline, you can delete the node.
- To take the specified node offline immediately without flushing the cache or ensuring that data loss does not occur, run the **rmnode** command with the **-force** parameter. The **-force** parameter forces continuation of the command even though any node-dependent volumes will be taken offline. Use the **-force** parameter with caution; access to data on node-dependent volumes will be lost.
- To delete a node that is in the service state and that has an associated spare node, you must specify the **-deactivatespare** parameter with the **rmnode** command.

About this task

Complete these steps to delete a node:

Procedure

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

```
lsvdisk -filtervalue IO_group_name=name
```

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

b) Issue the following CLI command to list the hosts that this volume is mapped to:

```
lsvdiskhostmap vdiskname/identification
```

Where *vdiskname/identification* is the name or identification of the volume.

Note: If volumes are assigned to this I/O group that contain data that you want to continue to access, back up the data or migrate the volumes to a different (online) I/O group.

2. Issue this CLI command to delete a node from the clustered system:



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

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

```
rmnode node_name_or_identification
```

Where *node_name_or_identification* is the name or identification of the node.

Note: Before a node is removed, the command checks for any node-dependent volumes that would go offline. If the node that you selected to delete contains a flash drive that has dependent volumes, volumes that use the flash drive go offline and become unavailable if the node is deleted. To maintain access to volume data, mirror these volumes before you remove the node. To continue removing the node without mirroring the volumes, specify the **force** parameter.

Chapter 11. Moving expansion enclosures

To move an expansion closure from one control enclosure to another, complete the following steps.

Before you begin

Review the following procedure. If the expansion enclosure that you want to move is not at the end of a SAS chain, you might need a longer pair of SAS cables to complete the procedure. In that case, ensure that you have two SAS cables of suitable length before you start this procedure.

Procedure

1. Delete any volumes that are no longer needed and that depend on the enclosure that you plan to remove.
2. Delete any remaining arrays that are formed from drives in the expansion enclosure.
Any data in those arrays is automatically migrated to other managed disks in the pool if there is enough capacity.
3. Wait for data migration to complete.
4. Mark all the drives (including any configured as spare or candidate drives) in the enclosures to be removed as unused.
5. Unmanage and remove the expansion enclosure by using the management GUI. Select **Monitoring** > **System**. On the **System - Overview** page, select the directional arrow next to the enclosure that you are removing to open the **Enclosure Details** page. Select **Enclosure Actions** > **Remove**.
Do not proceed until the enclosure removal process completes successfully.
6. On the I/O group that contains the expansion enclosure that you want to remove, enter the following command to put the I/O group into maintenance mode:

```
chiogrp -maintenance yes iogroup_name_or_id
```
7. If the expansion enclosure that you want to move is at the end of a SAS chain, complete the following steps to remove the enclosure from the SAS chain.
 - a) Disconnect the SAS cable from port 1 of canister 1 and canister 2.
The enclosure is now disconnected from the system.
 - b) Disconnect the other ends of the SAS cables from the previous enclosure in the SAS chain.
The previous enclosure is now the end of the SAS chain. Proceed to step “9” on page 86.
8. If the expansion enclosure is not at the end of a SAS chain, complete the following steps to remove the enclosure from the SAS chain.
 - a) Disconnect the SAS cable from port 2 of canister 1 of the expansion enclosure that you want to move.
 - b) Disconnect the other end of the same SAS cable from port 1 of canister 1 of the next expansion enclosure in the SAS chain.
 - c) Disconnect the SAS cable from port 1 of canister 1 of the expansion enclosure that you want to move.
 - d) Reroute the cable that was disconnected in the previous step and connect it to port 1 of canister 1 of the next expansion enclosure in the SAS chain.
Do not continue until you complete this connection step.
 - e) Disconnect the SAS cable from port 2 of canister 2 of the expansion enclosure that you want to move.
 - f) Disconnect the other end of the same SAS cable from port 1 of canister 2 of the next expansion enclosure in the SAS chain.

- g) Disconnect the SAS cable from port 1 of canister 2 of the expansion enclosure that you want to move.
 - h) Reroute the cable that was disconnected in the previous step and connect it to port 1 of canister 2 of the next expansion enclosure in the SAS chain.
9. Take the I/O group out of maintenance mode by entering the following command.

```
chiogrp -maintenance no iogroup_name_or_id
```

- 10. Check the event log for any errors and fix those errors as needed.
- 11. Disconnect the power from the expansion enclosure that you want to move.
- 12. Remove the expansion enclosure from the rack along with its two power cables and two SAS cables.
- 13. Add the expansion enclosure to the other system, near to its new control enclosure.

Chapter 12. Shutting down a system by using the CLI

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

Procedure

To power off your system, complete the following steps.

1. Determine which hosts have access to volumes on this system by running the **lshostvdiskmap** command.
2. Stop input/output (I/O) to the system from each host that is listed in step “1” on page 87.

Note: Failure to stop host I/O can result in failed I/O operations being reported to your host operating systems.

3. Shut down the system by using this command:

```
stopssystem
```

4. Wait for the power light-emitting diodes (LEDs) on all nodes or control enclosures to flash at 1 Hz, indicating that the shutdown operation has completed.




Figure 2. Location of the power LED on a node canister

- 1 Power
- 2 Status
- 3 Fault

5. Disconnect the power cords from both power supplies in each node or control enclosure.
6. Disconnect the power cords from both power supplies in each expansion enclosure.

