IBM Storage Virtualize for SAN Volume Controller and FlashSystem Family

Troubleshooting Guide - Version 863

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 6, 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

Tables

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

© Copyright IBM Corp. 2024

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.

Troubleshooting

Troubleshooting procedures help you diagnose problems.

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.

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 **svcconfig** 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 known as a **cron** backup and is written to `/dumps/svc.config.cron.xml_`*`serial`* 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.

© Copyright IBM Corp. 2024

Complete the following steps to backup your configuration data:

Procedure

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

   ```
   svcconfig backup
   ```

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

   ```
   CMMVC6155I SVCCONFIG processing completed successfully
   ```

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

   Table 1 on page 6 describes the three files that are created by the backup process:

   Table 1. Files created by the backup process

File name	Description
`svc.config.backup.xml`	Contains your configuration data.
`svc.config.backup.sh`	Contains the names of the commands that were issued to create the backup of the system.
`svc.config.backup.log`	Contains details about the backup, including any reported errors or warnings.

2. Check that the **svcconfig 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:/tmp/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, local Metro Mirror information, local Global Mirror information, 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: The policy-based replication configuration will not be 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 14 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 i.

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 MDisks and the arrays are re-created and configured. If multiple storage enclosures are involved, the arrays and MDisks 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 mkip 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. See Implementing vVols Replication.
- 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.
- 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 passphareses 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. Any other control enclosures and expansion enclosures in other I/O groups are added to the system later.

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 XML 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. For configurations with more than one I/O group, add the rest of the control enclosures into the clustered system by using the **addcontrolenclosure** CLI command.

   The remaining enclosures are added in the appropriate order based on the previous **IO_group_id** of its node canisters. The following example shows the command to add a control enclosure to I/O group 2.

   ```
   svctask addcontrolenclosure -sernum SVT5M48 -iogrp 2
   ```

10. Identify the configuration backup file that you want to restore.

    The file can be either a local copy of the configuration backup XML 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.*.xml*`.

       iii) Select the files and click **Download** to download them to your computer.

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

11. Copy onto the system the XML backup file from which you want to restore.

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

12. 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 or HyperSwap® 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) To add these nodes, determine the panel name, node name, and I/O groups of any such nodes from the configuration backup file. To add the nodes to the system, run the following command:

```
svctask addcontrolenclosure -iogrp iogrp_name_or_id -sernum enclosure_serial_number
-site site_id
```

    Where *enclosure_serial_number* is the serial number of the control enclosure, *iogrp_name_or_id* is the name or ID of the I/O group to which you want to add this node, and *site_id* is the numeric site value (1 or 2) of the control enclosure.

    b) Run the following command to change the replication layer.

```
chsystem -layer replication
```

    c) 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.

    d) 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.

    e) 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), i*scsi_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 IPv4/IPv6 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 HyperSwap or stretched system, the controller name and site needs to be restored. To restore the controller name and site, determine c*ontroller_name* and controller *site_id/name* from the backup xml 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 xml file, *site_id/name* is the ID or name of the site of iSCSI controller from the backup xml file, and *controller_id/name* is the ID or current name of the controller.

13. 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'
```

14. If required, follow the steps for exporting the system certificate as described in Configuring multifactor authentication with IBM Security Verify. 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' -cliclientid 'clientid' -cliclientsecret 'clientsecret'
```

15. 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'
```

16. 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'
```

17. 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.
18. 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 `/tmp` directory of the configuration node. The name of the log file is `svc.config.restore.prepare.log`.

    Note: It can take up to a minute for each 256-MDisk batch to be discovered. If you receive error message `CMMVC6200W` for an MDisk after you enter this command, all the managed disks (MDisks) might not be discovered yet. Allow a suitable time to elapse and try the **svcconfig restore -prepare** command again.
19. Issue the following command to copy the log file to another server that is accessible to the system:

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

20. Open the log file from the server where the copy is now stored.
21. 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 “22” on page 13.
    - If you need assistance, contact the support center.
22. Issue the following CLI command to restore the configuration:

```
svcconfig restore -execute
```

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

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

24. Open the log file from the server where the copy is now stored.
25. 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>
```

26. 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.
27. See **svcconfig** to delete backup configuration files.
28. If the system is using two person integrity, after restore, enable two person integrity manually. It cannot be enabled automatically. To enable manually, see Configuring two person integrity (TPI).

What to do next

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

```
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

1. Issue the following command to log on to the system:
2. Issue the following CLI command to erase all of the files that are stored in the /tmp directory:

```
svcconfig clear -all
```

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 XML configuration backup file. For IBM Storage Virtualize for Public Cloud, the saved configuration data is in the system disk and the latest XML 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 XML 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 XML 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 16.
- All hardware errors are fixed. See "Fix hardware errors" on page 17.
- All node canisters have candidate status. Otherwise, see step "1" on page 15.
- 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 17.
- 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 24.

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 17.
    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 **svcconfig 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 **svcconfig 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.

   Note: Run the procedure on one system in a fabric at a time. Do not run the procedure on different nodes 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 17.
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 “Recover system procedure” on page 14.

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 Storwize® V7000 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.

Run the recovery from any nodes in the system; the nodes must not participate in any other 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.
2. Point your browser to the service IP address of one of the node canisters.
3. Log on to the service assistant.
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 complete the recovery 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 20 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 21.

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 20.

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 Metro Mirror or Global Mirror relationships that use the offline volumes.

   Note: Do not delete HyperSwap relationships for HyperSwap 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 21 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 Metro Mirror or Global Mirror 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.
- 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 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.
- 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 Resolving synchronization errors in policy-based replication.
- 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 the partnership must be in a stopped state. The **chpartnership -newclusterid** CLI command must be used on each remote system to update the partnership data. All Metro Mirror or Global Mirror relationships and consistency groups using the partnership must be removed before running this command. Once the system ID has been updated the partnerships can be started.
- 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.
- Except for active-active relationships and consistency groups for HyperSwap volumes, remote-copy relationships and consistency groups are not restored and must be re-created manually.

- Volumes with cloud snapshots that were enabled before the recovery need to have the cloud snapshots manually reenabled.
- 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 Global Mirror secondary volumes on the recovered system might have inconsistent data if replication I/O from the primary volume is cached on the secondary system at the point of the disaster. A full synchronization is required when re-creating and restarting these relationships.
- Any volumes that are while being formatted as a system failure occurs are set to the "formating_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” on page 24.

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 2. 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 *consistent_synchronized* 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 **rmrcrelationship** command.

Table 2. 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 **rmrcrelationship** 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 *consistent_synchronized* at the time of the system recovery, contact support. This recovery can be a lengthy process.

Related information

Migrating data between systems non-disruptively

Unable to add large storage provider on VMware vCenter

When you configure VMware Virtual Volumes and try to add a large VASA certificate through VMware vCenter, the files are not added.

If registration of EVP as the VASA Provider on VMware vCenter fails, do the following:

1. Check the output of the command **lseventlog** to see if there is an entry with the message *VASA Provider registration failed*:

```
IBM_FlashSystem:Cluster_9.71.24.59:superuser>lseventlog9000003         220927151013
cluster              Cluster_9.71.24.59         message no    989050              VASA
Provider registration failed
```

2. If such an entry exists, run the command **lseventlog <message-id>** to check the details:

```
IBM_FlashSystem:Cluster_9.71.24.59:superuser>lseventlog 9000003sequence_number 9000003
.....
event_id 989050
event_id_text VASA Provider registration failed
......
sense1 05 0100 00 00 00 00 00 00 00 00 00 00 00 00 00
......
sense8 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
IBM_FlashSystem:Cluster_9.71.24.59:superuser>
```

3. If the initial two values in the **sense1** data correspond to **05 01**, then it indicates that VMware vCenter has sent multiple certificates in the register message to EVP, which could not be added in a single truststore due to the large size.

In this scenario, the Storage Admin must complete the following steps to proceed further:

1. The admin must have a machine where they can securely copy the certificate file from IBM Storage Virtualize, which has those multiple certificates received from VMware vCenter in the register message. The certificates are available in the IBM Storage Virtualize config node in file `/dumps/vmware-vasa-certs`. Therefore, the following command can be run from an external machine to securely copy the certificate file to the local /tmp directory `[15:31:23] 78e01rm-1:/tmp# scp superuser@9.71.24.59:/dumps/vmware-vasa-certs /tmp`
2. This certificate file must be split manually into two or more smaller files with different names `vmware-vasa-certs-1, vmware-vasa-certs-2`, etc. It must be ensured that one entire certificate (including the line having the text **BEGIN CERTIFICATE** as well as the line having the text **END CERTIFICATE**) is in a single file and the final size of each file is less than 12 KB.
3. These multiple certificate files must then be securely copied back to the IBM Storage Virtualize config node `[15:31:23] 78e01rm-1:/tmp# scp /tmp/vmware-vasa-certs-*superuser@9.71.24.59:/dumps/`
4. The storage admin must now create multiple trust stores using the split certificate files copied to IBM Storage Virtualize

```
IBM_FlashSystem:Cluster_9.71.24.59:superuser> svctask mktruststore-file /dumps/vmware-vasa-certs-1 -vasa on
IBM_FlashSystem:Cluster_9.71.24.59:superuser> svctask mktruststore -file /dumps/vmware-vasa-certs-2 -vasa on
```

5. The admin must also create the file `/dumps/bypass-truststore` to trigger the EVP to bypass the truststore creation as it is already created manually. This file must be created on an external machine and then copied to the IBM Storage Virtualize config node using **scp**.

```
[15:31:23] 78e01rm-1:/tmp# touch /tmp/bypass-truststore
[15:31:23] 78e01rm-1:/tmp# scp /tmp/bypass-truststore superuser@9.71.24.59:/dumps
```

6. Now register the EVP again as VASA Provider on VMware vCenter. It should register successfully.