Restoring the system configuration

Use this procedure in the following situations: only if the recover procedure has failed or if the data that is stored on the volumes is not required.

Before you begin

This configuration restore procedure is designed to restore information about your configuration, such as volumes, local Metro Mirror information, local Global Mirror information, storage pools, and nodes. All the data that you have written to the volumes is not restored. To restore the data on the volumes, you must restore application data from any application that uses the volumes on the clustered system as storage separately. Therefore, you must have a backup of this data before you follow the configuration recovery process.

About this task

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

Important: There are two phases during the restore process: prepare and execute. You must not change the fabric or system between these two phases.

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

To restore your configuration data, follow these steps:

Procedure

1. Verify that all nodes are available as candidate nodes before you run this recovery procedure. You must remove errors 550 or 578 to put the node in candidate state. For all nodes that display these errors, perform the following steps:
   a. Point your browser to the service IP address of one of the nodes, for example, https://node_service_ip_address/service/.
   b. Log on to the service assistant.
   c. From the **Home** page, put the node 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 using the service assistant IP address and by physically verifying the LEDs for the hardware components.
2. Verify that all nodes are available as candidate nodes with blank system fields. Perform the following steps on one node in each control enclosure:
   a. Connect to the service assistant on either of the nodes in the control enclosure.
   b. Select **Configure Enclosure**.

c. Select the **Reset the system ID** option. Do not make any other changes on the panel.

d. Click **Modify** to make the changes.

3. Use the initialization tool that is available on the USB flash drive to a create a new Storwize® V7000 system. Select the Initialize a new Storwize V7000 (block system only) option from the **Welcome** panel of the initialization tool.
4. In a supported browser, enter the IP address that you used to initialize the system and the default superuser password (`passw0rd`).
5. At this point the setup wizard is shown. Be aware of the following items:

   a. Accept the license agreements.

   b. Set the values for the system name, date and time settings, and the system licensing. The original settings are restored during the configuration restore process.

   c. Verify the hardware. Only the control enclosure on which the clustered system was created and directly attached expansion enclosures are displayed. Any other control enclosures and expansion enclosures in other I/O groups will be added to the system.

   d. On the **Configure Storage** panel, deselect **Yes, automatically configure internal storage now**. Any internal storage configuration is recovered after the system is restored.
6. Optional: 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, so if the system you are restoring was previously configured in the replication layer, you must change the layer manually now. Refer to Metro Mirror and Global Mirror partnerships for more information.
8. For configurations with more than one I/O group add the rest of the control enclosures into the clustered system.

   a. From the management GUI, select **Monitoring > System Details**.

   b. Select the system name in the tree.

   c. Go to **Actions > Add Enclosures > Control and Expansions** .

   d. Continue to follow the on-screen instructions to add the control enclosures. Decline the offer to configure storage for the new enclosures when asked if you want to do so.
9. Identify the configuration backup file that you want to restore from.

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

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

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

   For each node in the system:

   a. From the management GUI, click **Settings** > **Support**.

   b. Click **Show full log listing**.

   c. Select the node to operate on from the selection box at the top of the table.

   d. Find the file name that begins with `svc.config.cron.xml`.

   e. Double-click the file to download the file to your computer.

   The XML files contain a date and time that can be used to identify the most recent backup. After you identify the backup XML file that is to be used when you restore the system, rename the file to svc.config.backup.xml.
10. Issue the following CLI command to remove all of the existing backup and restore configuration

files that are located on your configuration node in the `/tmp` directory: **svcconfig clear -all**

11. Copy the XML backup file that you wish to restore from back onto the system.

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

12. 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 have been discovered yet. Allow a suitable time to elapse and try the **svcconfig restore -prepare** command again.

13. 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
    ```

14. Open the log file from the server where the copy is now stored.
15. Check the log file for errors.
    - If there are errors, correct the condition that caused the errors and reissue the command. You must correct all errors before you can proceed to step 16.
    - If an error indicates that the system layer will not be restored, then return to 7, configure the layer setting correctly, and then continue the restore process from 10.
    - If you need assistance, contact the IBM® Support Center.
16. 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`.

17. 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
    ```

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

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

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

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

20. 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.

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
```

Parent topic: Backing up and restoring the system configuration

Related concepts:
Object naming
Metro Mirror and Global Mirror partnerships

Related tasks:
Using the initialization tool

Related reference:
Recover system procedure
help
restore

Feedback