

IBM System Storage SAN Volume Controller



Troubleshooting Guide - Errata

Version 4.3.1
July 14, 2009

GC27-2227-00-Errata

Contents

Introduction	5
Who should use this guide	5
Last Update	5
Change History	5
Chapter 3. Using the SAN Volume Controller Console and CLI	7
Recovering from offline VDisks	7
Recovering from offline VDisks after an I/O group has failed	7
Recovering a node and returning it to the cluster	9
Chapter 6. Diagnosing Problems	13
Understanding cluster error codes	13
1700 Unrecovered Metro Mirror or Global Mirror relationship	13
1895 Unrecovered FlashCopy mappings	14
Troubleshooting the SAN Volume Controller Console	14

Introduction

This guide provides errata information that pertains to release 4.3.1 of the *IBM System Storage SAN Volume Controller Troubleshooting Guide*.

This guide contains the corrections and additions on a per chapter basis. The chapter numbers in this guide correspond directly with the chapter numbers in the *Service Guide* supplied with your SAN Volume Controller.

Who should use this guide

This errata should be used by anyone using the using the *IBM System Storage SAN Volume Controller Troubleshooting Guide*. You should review the errata contained within this guide and note the details with respect to the copy of the *Service Guide* supplied with your SAN Volume Controller.

Last Update

This document was last updated: July 14, 2009

Change History

The following revisions have been made to this document:

Revision Date	Sections Modified
December 19, 2008	New publication
July 14, 2009	Troubleshooting the SAN Volume Controller Console correction added. Minor clarification made to Recovering from Offline VDIs. Minor clarification made to Actions for error code 1700.

Table 1: Change History

Chapter 3. Using the SAN Volume Controller Console and CLI

The following corrections should be noted.

Recovering from offline VDIsks

This section has been improved. Replace the topics titled “Recovering from offline VDIsks”, “Recovering from offline VDIsks using the CLI”, “Recovering a node and returning it to the cluster using the CLI”, “Recovering VDIsks” and “Recovering offline VDIsks using the CLI” with the following sections.

Page 84 to page 88.

Recovering from offline VDIsks after an I/O group has failed

You will need to recover offline virtual disks (VDIsks) after an I/O group has failed.

There are circumstances when the VDisk data that is being held in the write cache is lost. These scenarios cause the VDisk to go offline and their fast write state will be **corrupt**.

The circumstances that cause write cache data to be lost are those that cause a failure of both nodes in an I/O group. When the first node failed it may, or may not, have had chance to write its state and cache data to its local disk. Even if it does write its state and cache data to disk, it must be assumed it is out of date by the time the second node fails.

After the first node fails the remaining node in the I/O group will start the failover process and will attempt to destage the write cache data to the storage controllers. If during this destage it too fails, and is unable to complete the write of its state and cache data to local disk (say because of a power supply failure in the node), then the data in the write cache has been lost.

The service procedures to repair and add the nodes back into the cluster should be followed. If there were more than one I/O group in the cluster, the cluster will have continued to operate and the nodes can be deleted from the cluster and added back in. If this were the only I/O group in the cluster you will need to contact your IBM service representative.

Once the nodes are operating within the cluster again any VDIsks that have not lost cache data will be online. If a VDisk has lost cache data it will remain offline, and their fast write state will be **corrupt**. The VDIsks can be recovered and brought back online using the Console or CLI commands described below. The data on the recovered VDIsks must not be trusted as it is not possible to tell which data was not written from the cache. Therefore the VDisk data should be recovered. The SAN Volume Controller VDisk mirroring or Global Mirror or Metro Mirror options can be used to maintain and recover the VDisk data. Local procedures can also be used.

When the VDIsks are recovered it is not possible to recover the FlashCopy mappings and Metro Mirror and Global Mirror relationships that were active in the I/O group. Errors will be logged to indicate if any FlashCopy mappings or Metro Mirror and Global Mirror relationships are not recovered. The Directed Maintenance Procedures or the service actions shown in the SAN Volume Controller Troubleshooting Guide should be followed to recreate the FlashCopy mappings and Metro Mirror and Global Mirror relationships and add them to their consistency groups.

Note: Updated service procedures for the unrecovered FlashCopy mappings and Metro Mirror and Global Mirror relationships errors, 1895 and 1700, are also contained in this errata.

Recovering offline VDIs using the Console

You can recover offline virtual disks (VDIs) using the SAN Volume Controller Console.

You can recover one or more corrupted VDIs and VDI copies. This task assumes that you have already launched the SAN Volume Controller Console. Perform the following steps to recover VDIs and VDI copies:

1. Click Work with Virtual Disks -> Virtual Disks in the portfolio. The Viewing Virtual Disks panel is displayed.
2. Select the offline VDIs with a fast write state of corrupt. To acknowledge data loss for all the selected virtual disks with a fast_write_state of corrupt, select Recover VDIs from the task list and click Go. The Recovering VDIs panel is displayed.

Note: If any VDI is space-efficient or has space-efficient copies, the action starts the space-efficient repair process.

Note: If any VDI is mirrored, the action starts the resynchronization process.

3. Verify that the VDIs and VDI copies have completed recovery by monitoring the VDI Recovery Results panel.

Note: Remember to ensure that the data on the recovered VDIs whose fast write state was corrupt is consistent, and to check for error log entries indicating the FlashCopy mappings or Metro Mirror or Global Mirror relationships need recovering.

Recovering offline VDIs using the CLI

You can recover offline virtual disks (VDIs) using the command-line interface (CLI).

Perform the following steps to recover offline VDIs:

1. Issue the following CLI command to list all VDIs that are offline and belong to an I/O group, enter:

```
svcinfolsvdisk -filtervalue IO_group_name=
IOGRPNAME/ID:status=offline
```

where *IOGRPNAME/ID* is the name or ID of the I/O group that failed.

2. To acknowledge data loss for all virtual disks in an I/O group with a fast_write_state of corrupt and bring them back online:

```
recovervdiskbyiogrp io_group_id | io_group_name
```

where *io_group_id | io_group_name* is the name or ID of the I/O group.

Note: If any VDI is space-efficient or has space-efficient copies, the recovervdiskbyiogrp command starts the space-efficient repair process.

Note: If any VDI is mirrored, the recovervdiskbyiogrp command starts the resynchronization process.

3. Verify that the VDIs and VDI copies have completed recovery by checking the results from the **svcinfolsvdisk**, **svcinfol**

lsrepairvdiskcopyprogress and **svcinfo lsrepairvdiskcopyprogress** commands.

Note: Remember to ensure that the data on the recovered VDisks whose fast write state was corrupt is consistent, and to check for error log entries indicating the FlashCopy mappings or Metro Mirror or Global Mirror relationships need recovering.

Recovering a node and returning it to the cluster

If a node fails you must recover it and return it to the cluster

You should fix any hardware problem before attempting to return the node to the cluster. If the node is repaired by replacing the front panel module or the node is replaced with another node, then you should ensure the WWNN for the node does not change. Follow the hardware replacement instructions or the instructions to change a node's WWNN in *Chapter 5. Using the Front Panel of the SAN Volume Controller* of the SAN Volume Controller Troubleshooting Guide.

If the node that failed was able to write its state to its internal disk when it failed, and if it is able to read that state information when it restarts it will automatically return to the cluster when it is restarted.

If the node that failed was unable to write its state to its internal disk when it failed, or if it is unable to read that state information when it restarts, the node will fail to rejoin the cluster and will display node error 578.

After the node is repaired or replaced you can use either the SAN Volume Controller Console or the Command Line Interface (CLI) to return the node to the cluster.

Returning a node to the cluster using the Console

Perform the following steps to return a node to the cluster using the SAN Volume Controller Console:

1. Click **Work with Nodes** -> **Nodes** in the portfolio. The **Viewing Nodes** panel is displayed.
2. Verify that the node is offline.
3. Select the offline node.
4. Select **Delete a Node** from the task list and click **Go**. The **Deleting Node** from **Cluster** panel is displayed.
5. Click **Yes**.
6. Add the node back into the cluster.
 - a. From the **Viewing Nodes** panel, select **Add a Node** from the task list and click **Go**. The **Adding a Node to a Cluster** panel is displayed.
 - b. Select the node from the list of candidate nodes and select the I/O group from the list. Optionally enter a node name for this node.
 - c. Click **OK**.
7. Verify that the node is online by refreshing the **Viewing Nodes** panel.

Note: If the panel does not refresh, close the panel and reopen it.

Returning a node to the cluster using the CLI

Perform the following steps to return a node to the cluster using the CLI:

1. Issue the following command to verify that the node is offline:

```
svcinfo lsnode
```

2. Issue the following command to remove the old instance of the offline node from the cluster:

```
svctask rmnode nodename/id
```

Where *nodename/id* is the name or ID of the node.

3. Issue the following command to verify that the node can be seen on the fabric:

```
svcinfo lsnodecandidate
```

4. Issue one of the following commands to add the node back into the cluster:

```
svctask addnode -panelname NNNNNN -iogrp  
IOGRPNAME/ID [-name NODENAME]
```

Where NNNNNN is the front panel name, IOGRPNAME/ID is the I/O group name or ID and NODENAME is the name of the node.

Or:

```
svctask addnode -wwnodename WWNN -iogrp IOGRPNAME/ID  
[-name NODENAME]
```

Where WWNN is the worldwide node name, IOGRPNAME/ID is the I/O group name or ID and NODENAME is the name of the node.

5. Issue the following command to verify that the node is online:

```
svcinfo lsnode
```

After returning the node to the cluster using either of the methods described above, if the node is repaired by replacing the front panel module or the node is repaired by replacing it with another node and the WWNN was not reset to that of the original node, then the following additional steps are required to detect the node using the new WWNN:

- a. You may need to modify the zoning set on your SAN fabric switches to make the node accessible to other nodes in the cluster, to your host application servers and to your disk controllers.
- b. On your host application servers you must discover the new paths and check that each device identifier presents the correct number of paths. For example, if you are using the system device driver (SDD), the device identifier is referred to as the virtual path (vpath) number. See the IBM System Storage Multipath Subsystem Device Driver User's Guide or the documentation that is provided with your multipathing driver for details about how to dynamically reconfigure and add device identifiers for the given host operating system.
- c. You might also need to modify the configuration of your disk controllers. If your controller uses a mapping technique to present its RAID arrays or partitions to the cluster, you must modify the port groups that belong to the cluster because the WWNN or WWPNS of the node have changed.

Attention: If more than one I/O group is affected, ensure that you are adding the node to the same I/O group from which it was removed. Failure to do this can result in data corruption. Use the information that was recorded when the node was originally added to the cluster. This can avoid a possible data corruption exposure if the node must be removed

from and re-added to the cluster. If you do not have access to this information, call the IBM Support Center to add the node back into the cluster without corrupting the data. If you are adding the node into the cluster for the first time, you must record the following information:

- Node serial number
- WWNN
- All WWPNs
- I/O group that contains the node

Chapter 6. Diagnosing Problems

The following corrections should be noted.

Understanding cluster error codes

The service actions for the following error codes have been enhanced.

Page 308

1700 Unrecovered Metro Mirror or Global Mirror relationship

Explanation

This error might be reported during the recovery action for the failure of one or more complete I/O groups. The error is reported because some Metro Mirror or Global Mirror relationships, whose control data is stored by the I/O group, were active at the time of the failure and the current state of the relationship could not be recovered.

Action

To fix this error it is necessary to delete all of the relationships that could not be recovered and then recreate the relationships.

1. Check that all higher priority errors are fixed. In particular check that all nodes in the I/O group are online and that all VDIs managed by the I/O group are online. See “Recovering from offline VDIs after an I/O group has failed” in Chapter 3.
2. Note the I/O group index against which the error is logged.
3. List all of the Metro Mirror and Global Mirror relationships that have either a master or an auxiliary VDisk in this I/O group. Use the VDisk view to determine which VDIs in the I/O group you noted have a relationship defined.
4. Note the details of the Metro Mirror and Global Mirror relationships that are listed so that they can be recreated. Note the direction of the relationship. Also note the consistency group, if any, the Metro Mirror or Global Mirror relationships belong to.
5. Delete all of the Metro Mirror and Global Mirror relationships that are listed. Note: The error will automatically be marked as “fixed” once the last relationship on the I/O group is deleted. New relationships should not be created until the error is fixed.
6. Using the details noted in step 4, recreate all of the Metro Mirror and Global Mirror relationships that you just deleted. Do not add the relationships back into the consistency groups at this stage.
Note: You are able to delete a Metro Mirror or Global Mirror relationship from either the master or auxiliary cluster; however, you must recreate the relationship on the master cluster. Therefore, it might be necessary to go to another cluster to complete this service action.
7. Start the Metro Mirror and Global Mirror relationships just created and allow them to resynchronize. To resynchronize the most important Metro Mirror and Global Mirror relationships first, you may wish to do this in batches.

8. If required, when the resynchronization has completed, add the Metro Mirror or Global Mirror relationship back into its consistency group.

Possible Cause-FRUs or other:

- None

page 313

1895 Unrecovered FlashCopy mappings

Explanation

This error might be reported after the recovery action for a cluster failure or a complete I/O group failure. The error is reported because some FlashCopies, whose control data is stored by the I/O group, were active at the time of the failure and the current state of the mapping could not be recovered.

Action

To fix this error it is necessary to delete all of the FlashCopy mappings on the I/O group that failed.

1. Check that all higher priority errors are fixed. In particular check that all nodes in the I/O group are online and that all VDIs are online. See “Recovering from offline VDIs after an I/O group has failed” in Chapter 3.
2. Note the I/O group index against which the error is logged.
3. List all of the FlashCopy mappings that are using this I/O group for their bitmaps. You should get the detailed view of every possible FlashCopy ID. Note the IDs of the mappings whose IO_group_id matches the ID of the I/O group against which this error is logged.
4. Note all the details of the FlashCopy mappings that are listed so that they can be recreated.
5. Delete all of the FlashCopy mappings that are listed. Note: The error will automatically be marked as “fixed” once the last mapping on the I/O group is deleted. New mappings cannot be created until the error is fixed.
6. Using the details noted in step 3, recreate all of the FlashCopy mappings that you just deleted. The copies can be started when required.

Possible Cause-FRUs or other:

- None

Troubleshooting the SAN Volume Controller Console

The Updating SSL certificate instructions have been corrected.

p333, step 4 should be

4. Issue the command `cimconfig -s sslCertificateFilePath=C:\“Program Files”\IBM\svconsole\cimon\certificate\ssl.cert -p`