IBM System Storage SAN Volume Controller

Troubleshooting Guide - Errata

Version 6.2.0
June 3, 2011

GC27-2284-01 Errata

Contents

Introduction

This guide provides errata information that pertains to release 6.2.0 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 *SAN Volume Controller Model Troubleshooting Guide* which is available as an online file. The information provided in these errata also applies to the Information Center content.

Who should use this guide

This errata should be used by anyone 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 *IBM System Storage SAN Volume Controller Troubleshooting Guide* downloaded to your system.

Last Update

This document was last updated: June 3, 2011.

Change History

The following revisions have been made to this document:

Revision Date	Sections Modified
June 3, 2011	New publication

Table 1: Change History

Chapter 9. MAP 6001: Replace offline SSD in a RAID 0 array

The text below replaces all of MAP 6001
Page 272ff

MAP 6001: This procedure replaces a solid-state drive (SSD) that has failed while it is still a member of a storage pool.

You must check the offline reason before continuing with this procedure, only replace the drive if it has failed.

If you are not familiar with these maintenance analysis procedures (MAPs), first read Chapter 9, “Using the maintenance analysis procedures”.

This map applies to models with internal solid-state drives (SSDs). Be sure that you know which model you are using before you start this procedure. To determine which model you are working on, look for the label that identifies the model type on the front of the node.

Attention:

1. Back up your SAN Volume Controller configuration before you begin these steps.
2. If the drive use property is member and the drive must be replaced, contact IBM support before taking any actions.

Only perform the following steps if a drive in a RAID 0 (striped) array has failed.

1. Record the properties of all volume copies, MDisks (arrays) and storage pools that are dependent on the failed drive.

a. Identify the drive ID and the error sequence number with status equals offline and use equals the failed using the **lsdrive** CLI command. Review the offline reason using the **lseventlog <seq_no>** command.
b. Obtain detailed information about the offline drive or drives using the **lsdrive <drive_id>** CLI command. Record the mdisk_id, mdisk_name, node_id, node_name and slot_id for each offline drive.
c. Obtain the storage pools (MDisk groups) of the failed drives using **the lsmdisk <mdisk_id>** command for each MDisk identified by step 1.b. Continue with the following steps replacing all the failed dives in one of the storage pools. Note the node, slot and id of the selected drives.
d. Determine all the MDisks in the storage pool using the **lsmdisk –filtervalue mdisk_grp_id=<grp id>** command. Note which MDisks are internal (ctrl_type equals 4) and which contain SSDs (ctrl_type equals 6).
e. Find the volumes with extents in the storage pool using the **lsmdiskmember <mdisk_id>** CLI command for each MDisk found in step1. d. It is likely that the same volumes will be returned for each MDisk.

f. Record all the properties on each volume listed by step 1.e. using the **lsvdisk <vdisk_id>** CLI command. For each volume check if it has online vdisk copies which indicate it is mirrored. This information is used in step 9.
g. Obtain a list of all of drives in each internal MDisk in the storage pool using the **lsdrive –filtervalue mdisk_id=<mdisk_id>** command. This information is used in step 8.
h. Record all the properties of all the MDisks in the storage pool using **the lsmdisk <mdisk_id>** command. This information is used in step 8.
i. Record all the properties of the storage pool using the lsmdiskgrp <mdiskgrp_id> command. This information is used in step 7.

Note: If a listed volume has a mirrored, online and in sync copy, you can recover the copied volume data from the copy. All the data on the unmirrored volumes will be lost and will need to be restored from backup.

2. Delete the storage pool using **rmmdiskgrp -force <mdiskgrp id>**.
All MDisks and volume copies in the storage pool are also deleted. If any of the volume copies were the last in sync copy of a volume, all the copies that are not in sync are deleted as well even if they are not in the storage pool.

3. Using the drive ID that you recorded in step 1.c, set the use property of the drive to unused using the **chdrive** command.
chdrive -use unused <id of offline drive>
The drive is removed from the drive listing.

4. Follow the physical instructions to replace or remove a drive. See the "Replacing a SAN Volume Controller 2145-CG8 solid-state drive (SSD)" documentation or the "Removing a SAN Volume Controller 2145-CG8 solid-state drive (SSD)" documentation to find out how to perform the procedures.

5. A new drive object is created with the use attribute set to unused. This action might take several minutes.
Obtain the ID of the new drive using the **lsdrive** CLI command.

6. Change the use property for the new drive to candidate.
chdrive -use candidate <drive id of new drive>

7. Create a new storage pool with the same properties as the deleted storage pool. Use the properties that you recorded in step 1.i.
mkmdiskgrp -name <mdiskgrp name as before> -ext <extent size as before>

8. Recreate all MDisks that were previously in the storage pool using the information from steps 1.g and 1.h.
For internal RAID 0 MDisks use the command:
mkarray -level raid0 -drive <list of drive IDs> -name <mdisk_name> <mdiskgrp id or name>

where **-name <mdisk_name>** is optional, but you can use the parameter to make the new array have the same MDisk name as the old array.
Use the **addmdisk** command for external MDisks or the **mkarray** command for non RAID 0 MDisks.

9. For all the volumes which had online, in sync, mirrored volume copies before the MDisk group was deleted, add a new volume copy in the new storage pool to restore redundancy using the command **addvdiskcopy -mdiskgrp <mdiskgrp id> -vtype striped -easytier <on or off as before> <vdisk_id>**

10. For any volumes which did not have an online, in sync, mirrored copy before, you must recreate the volume and restore the data from a back-up or other means.

11. Mark the drive error as fixed using the error sequence number from step 1. **cherrstate -sequencenumber <error_sequence_number**>.

