IBM System Storage SAN Volume Controller

Service Guide - Errata

Version 4.3.0
August 8, 2008

GC26-7901-03-Errata

Contents

Introduction

This guide provides errata information that pertains to release 4.3.0 of the *IBM System Storage SAN Volume Controller Service 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 Service 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: August 8, 2008

Change History

The following revisions have been made to this document:

Revision Date	Sections Modified
June 10, 2008	New publication
June 20, 2008	Added updated error 1230 and Replacing SSH keys
August 8, 2008	Added updates to errors 1160, 1320 and 2030 Updated OVER SPEC Light Path service action

Table 1: Change History

Chapter 7. Diagnosing problems with the SAN Volume Controller

The following corrections should be noted.

Understanding the boot codes

The following boot code has been corrected.

Page 216

205 A problem with output overload was reported by the uninterruptible power supply

Explanation

A problem with output overload has been reported by the uninterruptible power supply 2145 UPS. The Mode Indicator on the 2145 UPS front panel is illuminated red.

Action

Ensure that only one or two 2145s are receiving power from the 2145 UPS. Also ensure that no other devices are connected to the 2145 UPS.

Disconnect the 2145-4F2(s) from the 2145 UPS. If the Mode Indicator is now green, on the disconnected 2145 UPS, reconnect the 2145-4F2s one at a time to determine which one causes the overload. Then, on the 2145-4F2 which caused the problem, in the sequence shown, exchange the FRUs for new FRUs. See "Possible Cause-FRUs or other" after the last action in this section.

If the Mode Indicator is still red with all outputs disconnected, replace the 2145 UPS electronics assembly. If the Mode Indicator is still red, replace the 2145 UPS assembly.

Possible Cause-FRUs or other:

2145-4F2 power cable assembly (45%)

2145-4F2 power supply assembly (45%)

2145 UPS electronics assembly (9%)

2145 UPS assembly (1%)

Defining cluster error codes

The following cluster error codes have been corrected.

Page 297

1160 The output load on the 2145 UPS exceeds the specification

Explanation

The 2145 UPS is reporting that too much power is being drawn from it. The power overload warning LED, which is above the load level indicators, on the 2145 UPS will be on.

Action

1. Determine the 2145 UPS that is reporting the error from the error event data. Perform the following steps on just this uninterruptible power supply.
2. Check that the 2145 UPS is still reporting the error. If the power overload warning LED is no longer on go step 6.

3. Ensure that only 2145s are receiving power from the uninterruptible power supply. Ensure that there are no switches or disk controllers that are connected to the 2145 UPS.
4. Remove each connected 2145 input power in turn, until the output overload is removed.
5. Exchange the FRUs for new FRUs in the sequence shown, on the overcurrent 2145. See “Possible Cause-FRU or other” after the last action in this section.
6. Check node status. If all nodes show a status of “online”, mark the error that you have just repaired “fixed”. If any nodes do not show a status of “online”, go to start MAP. If you return to this step, contact your support center to resolve the problem with the 2145 UPS.
7. Go to repair verification MAP.

Page 313

1230 Login excluded

Explanation

A port to port fabric connection (login) between the SVC node and either a controller or another SVC cluster has had excessive errors, it has therefore been excluded, and will not be used for I/O.

Action

Determine the remote system, it may be either a controller or an SVC cluster. Check the error log for other 1230 errors. Ensure all higher priority errors are fixed.

This error event is usually the result of a fabric problem. If possible use the fabric switch or other fabric diagnostic tools to pinpoint the link and/or port that is giving errors. If there are error events for this node to a number of different controllers and/or clusters, then it is probably the node to switch link that is giving errors. Unless there are other contrary indications, replace the cable between the switch and remote system first.

1. Determine the most likely FRU from the fabric analysis, is this FRU has recently been replaced while resolving a 1230 error, choose the next most likely FRU that has not been replaced so far. Exchange the FRU for new FRU.
2. Mark the error as fixed. If the FRU replacement has not fixed the issue, the error will be logged again, although, depending on the severity of the problem, it may not be re-logged immediately.
3. Start a cluster discovery operation to recover the login by rescanning the fibre-channel network.
4. Check the status of the disk controller or remote cluster. If the status is not “good”, go to the Start MAP.
5. Go to repair verification MAP

Possible Cause-FRUs or other:

- Fibre-channel cable, switch to remote port (30%)
- Switch or remote device SFP or adapter (30%)
- Fibre-channel cable, local port to switch (30%)
- SAN Volume Controller SFP (9%)

- SAN Volume Controller fibre channel adapter (1%)

Note: the first two FRUs are not SAN Volume Controller FRUs.

Page 313

1320 A disk I/O medium error has occurred

Explanation

A disk I/O medium error has occurred.

Action

1. Check whether the VDisk the error is reported against is mirrored. If it is, check if there is a "1870 Mirrored VDisk offline because a hardware read error has occurred" error relating to this VDisk in the error log. Also check if one of the mirror copies is synchronising. If all these tests are true then you must delete the VDisk Copy that is not synchronized from the VDisk. Check that the VDisk is online before continuing with the following actions. Wait until the medium error is corrected until trying to re-create the VDisk mirror.
2. If the medium error was detected by a read from a host, ask the customer to rewrite the incorrect data to the block LBA that is reported in the host systems SCSI sense data. If an individual block cannot be recovered it will be necessary to restore the VDisk from backup. (If this error has occurred during a migration, the host system does not notice the error until the target device is accessed.)
3. If the medium error was detected during a mirrored VDisk synchonization, the block might not be being used for host data. The medium error must still be corrected before the mirror can be established. It may be possible to fix the block that is in error using the disk controller or host tools. Otherwise it will be necessary to use the host tools to copy the VDisk content that is being used to a new VDisk. Depending on the circumstances this new VDisk can be kept and mirrored, or the original VDisk can be repaired and the data copied back again.
4. Check the managed disk status. If all managed disks show a status of "online", mark the error that you have just repaired as "fixed". If any managed disks do not show a status of "online", go to start MAP. If you return to this step, contact your support center to resolve the problem with the disk controller.
5. Go to repair verification MAP.

Page 322

1610 Too many copied media errors on managed disk

Explanation

SAN Volume Controller maintains a virtual medium error table for each MDisk. This is a list of LBAs on the managed disk that contain invalid data and cannot be read. The virtual medium error table is of fixed length. This error event indicates that the system attempted to add an entry to a table but could not because it is already full.

There are two circumstances when an entry will be added to the virtual medium error table:

1. FlashCopy, data migration and mirrored VDisk synchronization copy data from one managed disk extent to another. If the source extent contains either a virtual medium error or the RAID controller reports a real medium error, then the system creates a matching virtual medium error on the target extent.
2. The mirrored VDisk validate and repair process has the option to create virtual medium errors on sectors that do not match on all VDisk copies. Normally no differences, or very few differences are expected, however if the copies have mistakenly been marked as synchronized then a large number of virtual medium errors can be created.

Action

You should ensure all higher priority errors are fixed before attemping to resolve this error.

Check if the excessive number of virtual medium errors occured because of a mirrored disk validate and repair which created medium errors for differences or if they were created because of a copy operation and follow the correct option shown below.

1. If the virtual medium errors occured because of a mirrored disk validate and repair which created medium errors for differences, then also check if the VDisk copies should have been fully synchronized.

   If the copies have been synchonized there should only be a few virtual medium errors created by the validate and repair process. In this case it might be possible to rewite just the data that was not consistent on the copies using the local data recovery processes.

   If the copies had not been synchronized it is likely that there are now a large number of medium errors on all the VDisk copies. Even if the virtual medium errors are only believed to be for blocks that have never been written, it is important to clear the virtual medium errors so as not to inhibit other operations. To recover the data for all these virtual medium errors it is likely that the VDisk will have to be recovered from a backup using a process that rewrites all the VDisks sectors.
2. If the virtual medium errors have been created by a copy operation, it is best practice to correct any medium errors on the source VDisk, and not to propagate the medium errors to copies of the disk. Fixing higher priority errors in the error log will have corrected the medium error on the source disk. Once the medium errors have been fixed the copy operation must be run again to clear the virtual medium errors from the target VDisk. It may be necessary to repeat a sequence of copies if copies have been made of already copied medium errors.

As an alternative, which does not address the root cause, deleting the VDisks on the target managed disk that have the virtual medium errors will reduce the number of virtual medium error entries in the MDisk table; migrating the VDisk to a different managed disk will also delete entries in the MDisk table (but will create more on the MDisk the VDisk is migrated to).

Possible Cause-FRUs or other:

- None

Related tasks

"Marking error as fixed"

"MAP 5700: Repair verification"

Page 340

2030 Software error

Explanation

The 2145 software has restarted due to a problem in the cluster, on a disk system or on the fibre-channel fabric.

Action

1. Collect the software dump file(s) generated at the time the error was logged on the cluster.
2. Contact your product support center to investigate and resolve the problem.
3. Ensure that the software is at the latest supported level on the cluster and on the disk systems.
4. Use the available SAN monitoring tools to check for any problems on the fabric
5. Mark the error that you have just repaired as fixed.
6. Go to repair verification MAP.

Possible Cause-FRU or other:

- Your support center might indicate a FRU based on their problem analysis (2%)

Other:

- 2145 software (48%)
- Enclosure/controller software (25%)
- Fibre-channel switch or switch configuration (25%)

Replacing the client SSH private key known to the SAN Volume Controller software

The following correction have been made. This correction is included in some releases of the Guide.

Page 352.

You can replace the client SSH private key that is known to the SAN Volume Controller software.

Attention: If you have successfully contacted other SAN Volume Controller clusters, you will break that connectivity if you replace the client SSH private key that is known to the SAN Volume Controller software.

Perform the following steps to replace the client SSH private key:

1. Sign off the SAN Volume Controller Console.
2. Stop the CIM Agent service. Go to **Start > Programs > IBM System Storage SAN Volume Controller > Stop CIMOM Service**.
3. Perform the following steps to copy the client SSH private key into the appropriate SAN Volume Controller Console directory:
   a. Open a command prompt window.

b. Issue the following command:
   ```
   copy filename C:\Program Files\IBM\svcconsole\cimom\icat.ppk
   ```
   Where *filename* is the path and file name of the client SSH private key.

4. Restart the CIM Agent service. Go to **Start > Programs > IBM System Storage SAN Volume Controller > Start CIMOM Service**.
5. Log on to the SAN Volume Controller Console.
6. Click **Clusters** in the portfolio.
7. Check the status of the cluster.

Chapter 8. Using maintenance analysis procedures

The following corrections should be noted.

MAP 5800: Light Path for 2145-8G4

The following service procedure has been extended.

Page 420

OVER SPEC diagnostic panel LED

Action

Replace FRUs in the following order until the problem is fixed -

1. Power supply (48%)
2. Replace the power backplane (48%)
3. System board (2%)