

**IBM System Storage SAN Volume Controller**



# **Service Guide - Errata**

**Version 4.2.0**  
**August 15, 2007**

GC26-7901-01-Errata



---

# Contents

<b>Introduction</b> .....	<b>5</b>
Who should use this guide .....	5
Last Update .....	5
Change History .....	5
<b>Chapter 1. SAN Volume Controller Overview</b> .....	<b>7</b>
Cluster identification .....	7
<b>Chapter 2. Redundant ac power switch</b> .....	<b>9</b>
Preparing your redundant ac power environment .....	9
Example mains wiring diagram .....	9
<b>Chapter 7. Diagnosing problems with the SAN Volume Controller, the uninterruptible power supply, and the master console</b>	<b>11</b>
Defining cluster error codes .....	11
Understanding the node rescue codes .....	14
<b>Chapter 8. Maintenance analysis procedures</b> .....	<b>17</b>
MAP 5000: Start .....	17
Using the SAN Volume Controller Console to power off a node .....	18
Using the SAN Volume Controller CLI to power off a node .....	20
<b>Chapter 9. Removing and replacing parts</b> .....	<b>23</b>
Removing the service controller .....	23
Replacing the SAN Volume Controller 2145-8G4 SATA cable assembly with backplate .....	23
Replacing the frame assembly .....	23
Replacing the SAN Volume Controller 2145-8G4 fibre channel adapter assembly .....	25



## Introduction

This guide provides errata information that pertains to release 4.2.0 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 15, 2007

---

### Change History

The following revisions have been made to this document:

Revision Date	Sections Modified
August 15, 2007	New publication

*Table 1: Change History*



# Chapter 1. SAN Volume Controller Overview

*The following correction should be noted.*

---

## Cluster identification

*Page 63. Replace this section with the following:*

A SAN Volume Controller cluster is identified by its IP address.

This address is used to access the cluster when using the SAN Volume Controller graphical user interface or the command-line interface. When a node has been assigned to a cluster, you can display the cluster IP address on the front panel by selecting **Cluster** from the menu.

### Configuration Node

Only one node within a SAN Volume Controller cluster can be set to use the cluster IP address. This node then acts as the focal point for all configuration, and other, requests made from the SAN Volume Controller Console application or the CLI. This node is known as the **configuration node**.

If the configuration node is stopped, or fails, the remaining nodes in the cluster will collaborate to decide which node will take on the role of configuration node. The new configuration node will bind the cluster IP address to its Ethernet port. It broadcasts this new mapping so any connections to the cluster configuration interface can continue.

The new configuration node broadcasts the new IP address mapping using the Address Resolution Protocol (ARP). Some switches need configuring to forward the ARP packet on to other devices on the subnetwork. You should ensure all Ethernet devices are configured to pass on unsolicited ARP packets.

If the ARP packet is not forwarded a device will lose its connection to the SAN Volume Controller Cluster. If the device is on the same subnetwork as the cluster this should resolve itself quickly as devices regenerate their address resolution caches very regularly. If the device is not on the same subnetwork it may take some hours for the gateway's address resolution cache to be refreshed. In this case you can restore the connection by establishing a command line connection to the cluster from a terminal which is on the same subnetwork, and then starting a secure copy to the device that has lost its connection.





## Chapter 2. Redundant ac power switch

*The following correction should be noted.*

---

### Preparing your redundant ac power environment

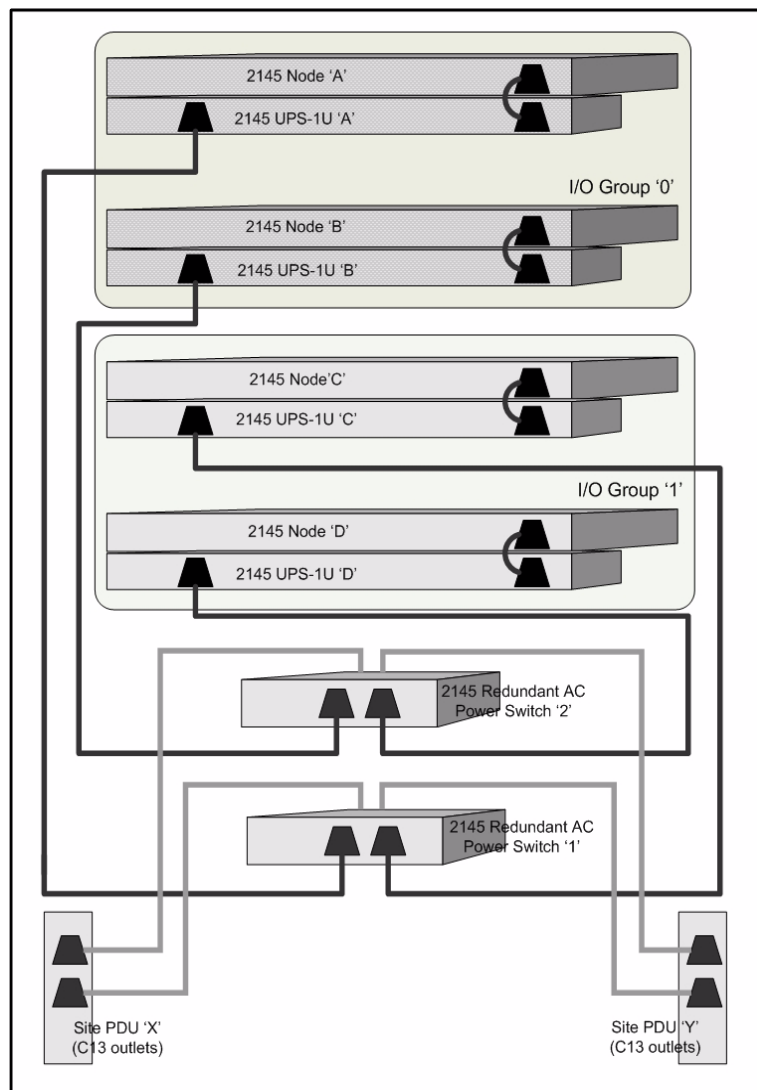
*Page 65. Correction to the redundant ac power switch power requirements.*

The redundant ac power switch requires two independent power sources, provided through two, rack mounted, power distribution units (PDUs). The PDUs must have IEC320-C13 outlets.

---

### Example mains wiring diagram

*Page 66. Add this new section.*



*Four Node SAN Volume Controller, with Redundant AC Power feature, mains wiring*

The diagram shows the mains power connections in an example installation. This example shows a four node cluster, there are two I/O groups; I/O group '0' contains nodes 'A' and 'B', I/O group '1' contains nodes 'C' and 'D'. In this example only two redundant ac power switches are used; each powers one node in each I/O group. [For maximum redundancy four redundant ac power switches would be used, each powering a single node.]

Note that the diagram is included to show an example of the cable connections; it does not show a suggested physical location for the components.

## Chapter 7. Diagnosing problems with the SAN Volume Controller, the uninterruptible power supply, and the master console

*The following corrections should be noted.*

---

### Defining cluster error codes

*The following new cluster error codes have been defined.*

*Page 195.*

#### **1627 The 2145 cluster has insufficient redundancy in its disk controller connectivity.**

##### **Explanation**

The SAN Volume Controller has detected that it does not have sufficient redundancy in its connections to the disk controllers. This means that another failure in the SAN could result in loss of access to the application data. SAN Volume Controller should be operated in an environment where there is redundancy in all its connections to disk controllers. This allows it to continue operating when there is a failure in one of the SAN components.

To provide good redundancy a SAN Volume Controller cluster should be configured so that:

- each node is able to access each disk controller through two or more different initiator ports on the node.
- each node is able to access each disk controller through two, or more, different controller target ports. **Note:** there are some disk controllers that only provide a single target port.
- each node is able to access each disk controller target port through at least one initiator port on the node.

If there are no higher priority errors being reported this error usually indicates a problem with the SAN design, a problem with the SAN zoning or a problem with the disk controller.

If there are unfixed, higher priority, errors that relate to the SAN or to disk controllers these should be fixed before attempting to resolve this error as they may be indicating the reason for the lack of redundancy. Error codes that must be fixed first are:

1210	Local FC port excluded
1230	Login has been excluded

**Note:** This error can be reported if the required action, to rescan the fibre-channel network for new MDisks, has not been made after a deliberate reconfiguration of a disk controller or the SAN zoning.

The 1627 error code is reported for a number of different error ids. The different error ids indicate the area where there is a lack of redundancy; the data reported within an error log entry indicates where the condition was found.

The meaning of the different error ids is shown below. For each error id the most likely reason for the condition is given, if the problem is not found within the suggested areas investigate the configuration and state of all the SAN components (switches, controllers, disks, cables and SAN Volume Controller) to determine where there is a single point of failure.

**Error id Meaning**

- 010040 A disk controller is only accessible from a single node port.  
A node has detected that it only has a connection to the disk controller through exactly one initiator port and more than one initiator port is operational.  
The error data indicates the device WWNN and the WPN of the connected port.  
A zoning issue or a fibre channel connection hardware fault may be responsible for this condition.
- 010041 A disk controller is only accessible through a single port on the controller.  
A node has detected that it is only connected to exactly one target port on a disk controller, and more than one target port connection is expected.  
The error data indicates the WWPN of the disk controller port that is connected.  
A zoning issue, a disk controller hardware issue or a fibre channel connection hardware fault may be responsible for this condition.
- 010042 Only a single port on a disk controller is accessible from every node in the cluster.  
Only a single port on a disk controller is accessible to every node when there are multiple ports on the controller that could be connected.  
The error data indicates the WWPN of the disk controller port that is connected.  
A zoning issue, a disk controller hardware issue or a fibre channel connection hardware fault may be responsible for this condition.

### **Error id Meaning**

- 010043 A disk controller is accessible through only half, or less, of the previously configured controller ports.  
Although there may still be multiple ports accessible on the disk controller it is possible that a hardware component of the controller has failed, or one of the SAN fabrics has failed, so the system has been reduced to a single point of failure.  
The error data indicates a port on the disk controller that is still connected as well as a list of controller ports that are expected but that are not connected.  
A disk controller hardware issue, or a switch hardware or zoning issue, or a cable fault may be responsible for this condition.
- 010044 A disk controller is not accessible from a node  
A node has detected that it has no access to a disk controller. The controller is still available to the partner node in the SAN Volume Controller I/O group so its data is still accessible to the host applications.  
The error data indicates the WWNN of the missing disk controller.  
A zoning or cabling error may be responsible for this condition.

### **Action**

1. Check the error id and data for a more detailed description of the error.
2. Check if there has been an intentional change to the SAN zoning or a disk controller configuration that reduces the SAN Volume Controller's access to the indicated disk controller. If this is the case continue with step 8.
3. Use the SAN Volume Controller Console application or the CLI command `svcinfolsfabric` to check all disk controller WWPNs are being reported as expected.
4. Check all disk controller WWPNs are zoned appropriately for use by the SAN Volume Controller cluster.
5. Check for any unfixed errors on the disk controllers.
6. Check all fibre channel cables are connected to the correct ports at each end.
7. Check for failures in the fibre channel cables and connectors.
8. If you believe the issues are resolved, use the SAN Volume Controller Console application or the CLI command `svctask detectmdisk` to rescan the fibre channel network for changes to the MDisks.  
**Note:** you should not attempt to detect MDisks until you believe a problem has been fixed, or you may mask an issue.
9. Mark the error that you have just repaired as fixed. The SAN Volume Controller will revalidate the redundancy and will report another error if there is not sufficient redundancy.
10. Go to MAP 5700: Repair verification.

Possible Cause-FRUs or other:

- None

### **Related tasks**

“Marking errors as fixed” on page 12 of the Service Guide.

“Viewing the fibre-channel fabric connections” on page 13 of the Service Guide.

“Rescanning the fibre-channel network for new MDisks” on page 23 of the Service Guide.

*Page 206.*

## **2600 The 2145 was unable to send an e-mail**

### **Explanation**

The SVC cluster has attempted to send an e-mail in response to an event, but there was no acknowledgement that it was successfully received by the SMTP mail server. It may have failed because the cluster was not able to connect to the configured SMTP server or they may have been a rejection by the server or a timeout. It could be that the SMTP server is not running or is not configured correctly, or it could be that the SVC is not configured correctly.

This error is not logged by the test e-mail function because it responds immediately with a result code.

### **Action**

1. Check the SMTP e-mail server is active.
2. Check the SAN Volume Controller e-mail configuration and ensure the SMTP server TCP/IP address and port are correct.
3. Validate the fix by sending a test e-mail.
4. Mark the error that you have just repaired as fixed.
5. Go to MAP 5700: Repair verification.

Possible Cause-FRUs or other:

- None

### **Related tasks**

Marking errors as fixed

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

*Page 206.*

## **2601 Error detected while sending an e-mail**

### **Explanation**

The SVC cluster has attempted to send an e-mail in response to an event, but an error has been detected. The 2145 is unable to determine if the e-mail has been sent and will attempt to resend it. The problem may be an issue with the SMTP server or the 2145 configuration; it may also be due to a fail-over of the configuration node.

This error is not logged by the test e-mail function because it responds immediately with a result code.

### **Action**

1. If there are higher priority unfixed errors in the log fix these errors first.
2. Check the SMTP e-mail server is active.
3. Check the SAN Volume Controller e-mail configuration and ensure the SMTP server TCP/IP address and port are correct.
4. Validate the fix by sending a test e-mail.
5. Mark the error that you have just repaired as fixed.
6. Go to MAP 5700: Repair verification.

Possible Cause-FRUs or other:

- None

### **Related tasks**

Marking errors as fixed

You can use the SAN Volume Controller Console to mark errors as fixed for the cluster error log. This action is only necessary if you fix an error without using the online maintenance procedures. The online procedures automatically mark an error as fixed after a successful repair.

---

## **Understanding the node rescue codes**

*Page 237. Code 350 “The 2145 cannot find a donor node”, replace the **Action** section with the following:*

### **Action**

If the progress bar has been stopped for more than two minutes, perform the following steps:

1. Check the correct fibre-channel cables are securely connected to the 2145.
2. Ensure that at least one other, suitable, donor node is operational and is connected to the same fibre-channel network. A node is a suitable donor node if the level of SAN Volume Controller software it is running supports the model type of the node being rescued.
3. Check the fibre-channel zoning allows connection between the node and a suitable donor node.
4. Perform the problem determination procedures for the network.





# Chapter 8. Maintenance analysis procedures

The following correction should be noted.


## MAP 5000: Start

Page 258. Replace Step 9

9. (from step 8)

**Is any SAN Volume Controller front panel showing an error?** There is an error if any of the following conditions are true for the node:

- none of the LEDs on the front panel are on and the front panel display is blank.
- the node is a SAN Volume Controller 2145-4F2 and the check LED on the front panel is on.
- the node is a SAN Volume Controller 2145-8F2, a SAN Volume Controller 2145-8F4 or a SAN Volume Controller 2145-8G4 and the error LED on the front panel is on.



**NO** Go to step 10

**YES** The service controller for the SAN Volume Controller has failed.

- a. Check that the service controller that is indicating an error is correctly installed. If it is, replace the service controller.
- b. Go to “MAP 5700: Repair verification”

Page 261. Add a new step. Go to this step if the answer to step 21 is NO.

21a.(from step 21)

**Is there a node that is not a member of a cluster?** You can tell if a node is not a member of a cluster by checking the front panel cluster menu. If **Cluster:** is displayed but no cluster name is shown, the node is not a member of a cluster. (The cluster name will be on the second line of the front panel display if the current language font allows a two line display, otherwise it can be displayed by pressing the select button.)

**NO** Go to step 22

**YES** The node is not a member of a cluster. The node may have been deleted from the cluster during a maintenance procedure and has not been added back into the cluster. Check if the node should be part of the

cluster (each I/O group in the cluster should contain two nodes). If necessary add the node back into the cluster, ensuring the node is restored to the same I/O group it was deleted from.

---

## Using the SAN Volume Controller Console to power off a node

Page 286. The following update should be noted.

This topic describes how to use the SAN Volume Controller Console to power off a node. Perform the following steps to use the SAN Volume Controller Console to power off a node:

1. Sign on to the master console as an administrator and then launch the SAN Volume Controller Console for the cluster that you are servicing.
2. Click **Work with Nodes** -> **Nodes** in the My Work pane and click on the submenu that opens. The Viewing Nodes panel is displayed. Find the node that you are about to shut down and write down the name of the I/O group it belongs to. Confirm that the other node in the I/O group is Online.

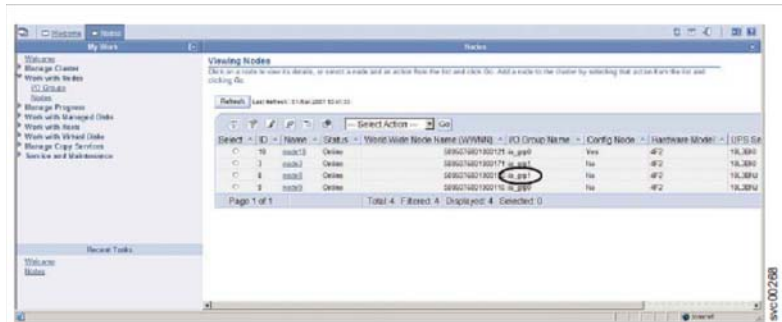


Figure 45. Viewing Nodes

If the node that you want to power off is shown as Offline, then the node is not participating in the cluster and is not processing I/O requests. In these circumstances, you must use the power button on the node to switch the node off.

If the node that you want to power off is shown as Online but the other node in the I/O group is not Online, powering off the node impacts all the hosts that are submitting I/O to the VDisks that are managed by the I/O group. Try to get the other node in the I/O group Online before continuing.

3. Click **Work with Virtual Disks** -> **Virtual Disks** in the My Work pane and then click in the submenu that opens. The filtering virtual disks panel is displayed.



Figure 46. Filtering Virtual Disks

4. Open the **I/O Group** drop down menu and select the I/O group that you wrote down in step 2 on page 286 for the node. Then click **OK**. The list of virtual disks in the I/O group is displayed. Check that the status of each virtual disk in the I/O group is Online; you might need to view more than one page.

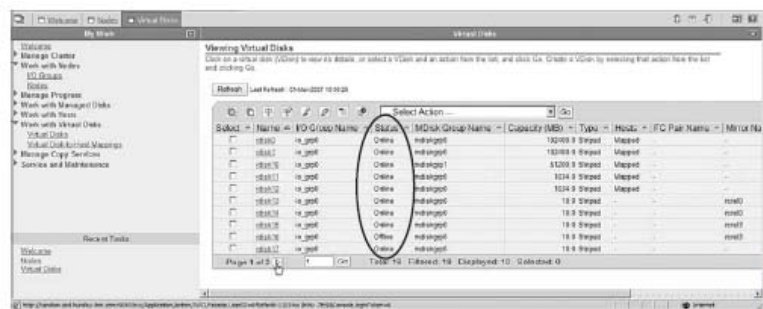


Figure 47. Virtual disks status

If any VDisks are shown as degraded, only one node in the I/O is processing I/O requests for that VDisk. If that node is powered off, it impacts all the hosts that are submitting I/O to the degraded VDisks.

If any virtual disks are degraded and you believe this might be because the partner node in the I/O group has been powered off recently, wait until a refresh of the screen shows all the virtual disks online. All the virtual disks should be online within thirty minutes of the partner node being powered off.

Try to get all VDisks that are being used by hosts Online before continuing.

5. If you have decided it is okay to continue and power off the node, click **Work with Nodes** -> **Nodes** in the My Work pane. The list of nodes in the cluster is displayed. Select the node that you are powering off and select **Shut Down a Node** from the drop down list.

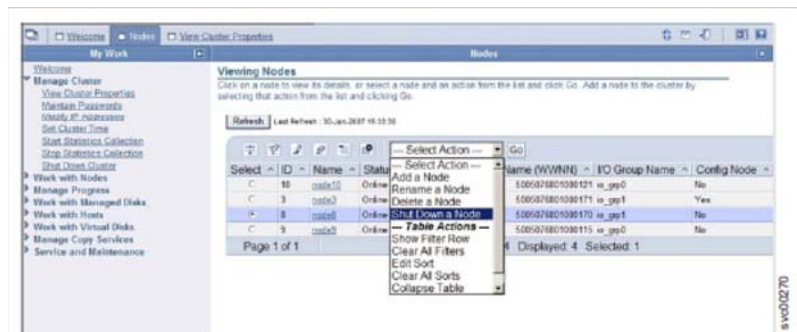


Figure 48. Shut Down a Node

During the shut down, the node saves its data structures to its local disk and destages all the write data held in cache to the SAN disks; this processing can take several minutes.

At the end of this process, the node powers off.

## Using the SAN Volume Controller CLI to power off a node

Page 288. The following update should be noted.

This topic describes how to use the SAN Volume Controller CLI to power off a node.

1. Issue the **svcinfolnode** CLI command to display a list of nodes in the cluster and their properties. Find the node that you are about to shut down and write down the name of the I/O group it belongs to. Confirm that the other node in the I/O group is online.

```
svcinfolnode -delim :
id:name:UPS_serial_number:WWNN:status:IO_group_id:
IO_group_name:config_node: UPS_unique_id
1:group1node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202378101C0D188
2:group1node2:10L3ANF:5005076801000009:online:0:io_grp0:no:202378101C0D1796
3:group2node1:10L3ASH:5005076801000001:online:1:io_grp1:no:202378101C0D18D8
4:group2node2:10L3ANF:50050768010000F4:online:1:io_grp1:no:202378101C0D1796
```

If the node you want to power off is shown as Offline, the node is not participating in the cluster and is not processing I/O requests. In these circumstances, you must use the power button on the node to switch the node off.

If the node you want to power off is shown as Online, but the other node in the I/O group is not Online, powering off the node impacts all the hosts that are submitting I/O to the VDisks that are managed by the I/O group. Try to get the other node in the I/O group Online before continuing.

2. Issue the **svcinfolsvdisk** CLI command to list the Vdisks managed by the I/O Group that you wrote down for the node in step 1.

```
svcinfolsvdisk -filtervalue IO_group_name=io_grp0 -delim :
0:mainvdisk1:0:io_grp0:online:0:mainmdiskgroup: 512.0GB:striped:::
1:bkpvdisk1:1:io_grp0:online:1:bkpmdiskgroup: 512.0GB:striped:::
```

The list of virtual disks in the I/O group is displayed. Ensure that the status of each virtual disk in the I/O group is online.

If any VDisks are shown as degraded, only one node in the I/O is processing I/O requests for that VDisk. If that node is powered off, it affects all the hosts that are submitting I/O to the degraded VDisks.

If any virtual disks are degraded and you believe this might be because the partner node in the I/O group has recently been powered off, wait until a refresh of the screen shows all the virtual disks online. All the virtual disks should be online within 30 minutes of the partner node being powered on.

Try to get all VDisks that are being used by hosts Online before continuing.

3. If you have decided that it is okay to continue and that you can power off the node, issue the **svctask stopcluster -node <name>** CLI command to power off the node. Ensure that you use the **-node** parameter, as you do not want to power off the whole cluster:

```
svctask stopcluster -node group1node2
Are you sure that you want to continue with the shut down? yes
```

4. During the shut down, the node saves its data structures to its local disk and destages all the write data held in the cache to the SAN disks; this process can take several minutes.
5. At the end of this process, the node powers off.



## Chapter 9. Removing and replacing parts

*The following correction should be noted.*

---

### Removing the service controller

*Page 342. Add the following Note to this section.*

**Note:** When replacing a service controller it is normal to change the World Wide Node Name (WWNN) of the new service controller to match that of the one being replaced. If this is the case, you will now have two service controllers with the same WWNN. You should clearly label the service controller that was removed, indicating that its WWNN is now a duplicate and that it must not be connected to a SAN before its WWNN is reset.

---

### Replacing the SAN Volume Controller 2145-8G4 SATA cable assembly with backplate

*Page 372. Replace Step 4 with the following:*

4. Reconnect the power cable to the power supply backplane card.

Reconnect the two blue signal cables to the system board. The lefthand cable goes to the front system board connector and the right hand cable goes to the rear system board connector. See the figure below.



---

### Replacing the frame assembly

*Page 393. Note the following changes.*

The SAN Volume Controller 2145-8F2, SAN Volume Controller 2145-8F4, and the SAN Volume Controller 2145-8G4 frame assembly must be replaced when the system board fails or when replacing other system board components fails to isolate the error.

**Note:** If the frame assembly is being replaced following a service procedures which replaced the service controller, you need to take special precautions to ensure the node can be re-added to the cluster. If any required updates to the World Wide Node Name (WWNN) stored in the service controller have already been made then you should continue with the node replace

instructions given below. If the node is being replaced because the new service controller did not fix the original problem, and the WWNN of the service controller could not be updated, then you should use the original service controller while performing the frame replacement. This means the WWNN will be correct.

If you have to use a new service controller when replacing the frame assembly you can use the following procedure to add the node back into the cluster. This would replace step 10. below.

- a. Start the node. Use your SAN monitoring tools to determine the WWNN of the node.
- b. Rezone the switch to allow one port of the replaced node to connect to a port of an existing node.
- c. Perform a node rescue.
- d. Set the replaced nodes WWNN to be the same as the node it replaced.
- e. Rezone the switch to its original settings.
- f. Add the node back into the cluster.
- g. Clearly label the original service controller to show its WWNN has been re-used, and that the part should not be used without first changing the WWNN.

Perform the following steps to replace the frame assembly:

1. Make a note of the seven-character system serial number that is on the serial number label on the front of the node. For the SAN Volume Controller 2145-8F2, SAN Volume Controller 2145-8F4, and the SAN Volume Controller 2145-8G4 models only: If you cannot read the serial number or if you want to verify that it is correct, you can also find it in the node vital product data (VPD). Perform the following steps to find the serial number in the VPD:
  - a. Start the SAN Volume Controller application.
  - b. Display the VPD for the failed node.
  - c. Note the system serial number.
2. Remove the SAN Volume Controller from the rack.
3. Remove the top cover.
4. Remove the fibre-channel adapter cards from the frame assembly that you are replacing.
5. Remove the service controller from the frame assembly you are replacing.
6. Install the fibre-channel adapter cards into the new frame assembly.
7. Install the service controller into the new frame assembly.
8. Replace the top cover.
9. Install the SAN Volume Controller in the rack.
10. Start the node using the node rescue procedure.

**Note:** It is essential that you perform the next steps to restore the original machine serial number. Failure to do this might invalidate the customer's warranty or service agreement.



11. If you are performing this repair as part of a directed maintenance procedure, you will be prompted to type the machine serial number that you noted above. Otherwise, perform the following steps:

- a. Delete the failed node from the cluster.
- b. Add the repaired node to the cluster.
- c. Start the command-line interface.
- d. Issue the following command:

```
svcservicetask writesernum -sernum nodeserialnumber
nodename
```

where *nodeserialnumber* is the number that you noted previously and *nodename* is the name of the repaired node that you added to the cluster. The `svcservicetask writesernum -sernum nodeserialnumber nodename` command writes the machine serial number to the SAN Volume Controller system board.

**Note:** The node will restart as soon as the serial number has been written to it.

- e. Write the serial number, noted previously, on the blank serial number label on the front of the node.

---

## Replacing the SAN Volume Controller 2145-8G4 fibre channel adapter assembly

*Page 401. Replace step 1*

1. As you start inserting the adapter, align the edge connector on the low-profile adapter [1] with the connector on the riser-card assembly [3]. Make sure that the adapter snaps into the riser-card securely. Press the riser-card edge connector firmly into the system board connector [4].

