IBM System Storage SAN Volume Controller

Procedures for Replacing 2145-4F2 Nodes with 2145-CF8 Nodes

Version 1.0
December, 2009

Scope and Objectives

This document provides procedures for either non-disruptively or disruptively, replacing 2145-4F2 nodes in an existing SAN Volume Controller with 2145-CF8 nodes while upgrading to SAN Volume Controller version 5.1.0. It provides important information to avoid problems while following the procedures.

Table of Contents

Overview

SAN Volume Controller 2145-4F2 nodes are not supported in SAN Volume Controller version 5.1.0 or later. To upgrade your SAN Volume Controller cluster to version 5.1.0 or later, you must replace all 2145-4F2 nodes with a node type that is supported by this software version. The latest SAN Volume Controller node is the model 2145-CF8 node, which can only be installed in clusters running SAN Volume Controller version 5.1.0 or later. To replace 2145-4F2 nodes with 2145-CF8 nodes, you cannot use the standard procedures; you must use the procedures described in this document.

The procedures described in this document implement node replacements while simultaneously upgrading the SAN Volume Controller software to version 5.1.0 or later. During a software upgrade each node in turn goes offline and installs the new software version. When the node restarts it uses the new software version, but continues to communicate with the rest of the cluster as if it were using the old software version, therefore the cluster continues to operate even though nodes are at different software versions. When all the nodes have been updated the cluster switches to operate at the new software version. Normally the software upgrade proceeds automatically from node to node; these procedures however make use of a new facility which allows the user to control when each node is upgraded. When the 2145-4F2 node is replaced by a 2145-CF8 node the 2145-CF8 node has the 5.1.0 or later code installed, but it operates in upgrade mode so it can still operate while the remainder of the cluster is running software version 4.3.1.

This document provides two alternate sets of instructions:

1. A non-disruptive procedure for replacing all 2145-4F2 nodes with 2145-CF8 nodes. This procedure simultaneously upgrades the cluster to SAN Volume Controller Version 5.1.0 or later without disrupting any host I/O.
2. A disruptive procedure for replacing all 2145-4F2 nodes, with 2145-CF8 nodes. This procedure simultaneously upgrades the cluster to SAN Volume Controller version 5.1.0 or later during a maintenance window with all host applications stopped. This procedure reduces the time required to complete the procedure, and removes the risk of accidental loss of access due to incorrect maintenance

Neither procedure requires changes to your SAN environment. The 2145-CF8 node that is being installed uses the same worldwide node name (WWNN) as the node that it is replacing, therefore no SAN zoning or disk controller LUN masking changes are required.

If necessary, node types other than 2145-4F2 can also be upgraded to 2145-CF8 nodes during either procedure.

ATTENTION: To replace 2145-4F2 nodes with 2145-8F2, 2145-8F4, 2145-8G4 or 2145-8A4 nodes, do not use the procedures defined in this document. You must use the standard procedures defined in the *SAN Volume Controller Software Installation and Configuration Guide*, SC23-6628, which is available at www.ibm.com/storage/support/2145.

The SAN Volume Controller 2145-CF8 node fibre-channel ports will be connected to the identical physical ports on the fabric(s) that the replaced node was connected to. The 2145-CF8 node's fibre-channel ports cannot operate at 1 Gb/s. Therefore these procedures can only be used to replace SAN

Volume Controller nodes that are currently connected to fibre-channel switch ports that support 2 Gb/s or greater. If the SAN Volume Controller nodes you are replacing are currently connected to 1 Gb/s switch ports, you must upgrade the switch port connections **before** starting these procedures.

The software version of the cluster must be 4.3.1.10 or later to run these procedures. If your cluster is not at SAN Volume Controller version 4.3.1.10 or later you must use the normal upgrade procedure to install version 4.3.1.10 or later before continuing with these procedures.

Both procedures are non-trivial and should only be performed by an experienced SAN Volume Controller administrator. Contact your IBM representative if you would like support with these procedures.

How to use this document:

1. Review the non-disruptive and disruptive procedures, and decide which is most appropriate for your environment.
2. Ensure that you have met all prerequisites, as defined in the *Upgrade prerequisites* section on page 7.
3. Follow the *Document the current configuration* section on page 9 to complete *Appendix A: Configuration table* on page 37.
4. Follow the steps defined in *Initial steps to perform before beginning either procedure* section on page 11.
5. Follow the steps as defined in the chosen procedure.

Revision History

The following changes have been made to this document:

Version	Date	Modifications
1.0	11 Dec 2009	Initial version

Upgrade prerequisites

The non-disruptive and disruptive upgrade procedures both require that the following prerequisite conditions have been met:

- The fibre-channel switch ports must support 2, 4 or 8 Gb/s connections; 1 Gb/s is not supported by 2145-CF8 nodes. If any of the nodes being replaced are currently connected to a 1 Gb/s SAN fabric, you must move the nodes to new, compatible, switch ports before starting this procedure.

- Do not start this procedure unless you have sufficient replacement hardware to replace every 2145-4F2 node in the cluster.

- The cluster software must be at SAN Volume Controller version 4.3.1.10 or later before starting these procedures. If the cluster software is at a version earlier than 4.3.1.10 you must download the latest version of the 4.3.1 code and install it on the cluster using the normal upgrade procedures specified in the *SAN Volume Controller Software Installation and Configuration Guide, SC23-6628*, which is available at www.ibm.com/storage/support/2145.

- All nodes that are configured in the cluster are accessible and online, and all I/O groups contain two nodes.

- All errors in the cluster error log have been addressed and marked as fixed.

- The SAN Volume Controller cluster's service mode IP address is correctly configured.

- The cluster 'admin' password is known. If it is not known, it can be changed in SAN Volume Controller version 4.3.1 by running the command 'svctask chcluster –admpwd <*new password*>'.

- The SAN Volume Controller configuration has been backed up via the CLI or GUI and the configuration file has been saved.

- Before beginning the node upgrade procedure, verify that there are no known issues with the current cluster environment by downloading, installing and running the latest *"SAN Volume Controller Software Upgrade Test Utility"* from: http://www.ibm.com/support/docview.wss?rs=591&uid=ssg1S4000585

- Verify that all VDisks, MDisks and storage systems are online and that none are in a state of "degraded" or "offline". If any are degraded or offline, resolve this issue before continuing or a loss of access to data can occur when running the procedure
    - Issue the following commands from the CLI. If any of the commands return an offline or degraded status, resolve the problems with the MDisks or VDisks before continuing.

svcinfo lsvdisk -filtervalue "status=degraded"
svcinfo lsvdisk -filtervalue "status=offline"

svcinfo lsmdisk -filtervalue "status=degraded"
svcinfo lsmdisk -filtervalue "status=offline"

- For each storage system (controller) that is attached to the SAN Volume Controller, issue the following command to verify that each controller's degraded status is "no". If any storage system's degraded status value is "yes", resolve the issue before continuing.

  svcinfo lscontroller *<object_id or object_name>*

  Where *<object_id or object_name>* is the name, or ID, of the storage system (controller) that you want to view.

Document the current configuration

Fill in Appendix A: Configuration table on page 37 with information about all cluster nodes. The following steps 1-3 detail how to collect this information using the SAN Volume Controller CLI. The information is also available from the San Volume Controller Console (GUI) interface.

1. Issue the following command from the command-line interface (CLI):

   svcinfo lsnode

   For all nodes, record the following columns from the svcinfo lsnode command output in *Appendix A: Configuration table* on page 37.

Column from svcinfo lsnode	Column from configuration table
"id" and "name"	Node ID / Name
"IO_group_id" and "IO_group_name"	IO Group ID / Name
"config_node"	Config Node (Yes/No)
"hardware"	Original Hardware Type

2. Issue the following command from the CLI for every node in the cluster (as identified in step 1) and locate the field called "front_panel_id". Record the value in the column labeled "Front Panel ID" for the corresponding node. A node's front panel ID can be found on a sticker located on the front of the node; you can use this to match each physical node to the *node_name or node_id* that you plan to replace.

   svcinfo lsnodevpd <*node_name or node_id*>

   Where <*node_name or node_id*> is the name or ID of the 2145 node identified in step 1.

3. Perform the following steps to record the WWNN for each node:

   a. Issue the following command from the CLI:

      svcinfo lsnode <*node_name or node_id*>

      Where <*node_name or node_id*> is the name or ID of the node that you are listing the WWNN for.

   b. Record the last five characters of the WWNN in the column labeled "WWNN (last 5 characters)" of the corresponding node. These five digits are displayed on the node's front panel later in the procedure.

   The column labeled "WWNN of replacement node (last 5 characters)" will be completed later.

4. Determine the nodes you are replacing with 2145-CF8s, and complete the “Replacing node Yes/No?” column.

Initial steps to perform before beginning either procedure

1. If you are unable to resolve any of the prerequisite conditions that are not met, do not continue without consulting with IBM support.

2. Review all of the following steps, with special attention to the items in RED throughout this document, before actually performing the procedure.

3. Download the latest SAN Volume Controller upgrade package (5.1.0.2 or later) from http://www.ibm.com/support/docview.wss?rs=591&uid=ssg1S1001707. Note the software version on the "Appendix A: Configuration table" on page 37.

4. Rename and upload the SAN Volume Controller upgrade package onto the configuration node for use later on in the upgrade procedure. You will need to repeat this step if any nodes that contain this package are upgraded or rebooted.

   Important Note:

   If you have to repeat the upload of the package, you MUST use exactly the same software version. DO NOT download another SAN Volume Controller upgrade package after starting the upgrade procedures.

   a. Make a copy of the IBM2145_INSTALL_X.X.X.X package you downloaded and rename the file **IBM2145_REMOTE_UPGRADE**

   b. Use the standard procedures to upload this package onto the SAN Volume Controller cluster, as if preparing for a software upgrade.

   c. Record the front panel ID of the configuration node. The IBM2145_REMOTE_UPGRADE package was uploaded to this node.

5. The following steps are not required, but are best practices for both procedures:
   a. Stop Metro and/or Global Mirroring during the procedure.
   b. Stop TPC performance data gathering and probes of this SAN Volume Controller cluster.
   c. Stop the Call Home service.
   d. Stop any automated scripts/jobs that may try and access the cluster during the procedure.
   e. Avoid running any flash copies during the procedure.
   f. Avoid running any VDisk migrations or formats during the procedure.
   g. Check the console "Manage/View Progress" panel to ensure that no operations are running.

6. Prior to performing this procedure, all 2145-CF8 nodes and uninterruptible power supply (2145 UPS-1U) units should ideally be powered on and allowed to fully charge for at least 24 hours to ensure that there are no hardware errors when you attempt to use the 2145-CF8 node in the cluster. The 2145-CF8 nodes must be attached to a new 2145

UPS-1U unit. Wherever possible, install the nodes into their final rack locations before beginning the procedure. Connect the 2145-CF8 nodes to the 2145 UPS-1U according to the *SAN Volume Controller Model 2145-CF8 Hardware Installation Guide*, which is available at www.ibm.com/storage/support/2145. Do NOT attach the fibre-channel or Ethernet cables to the nodes at this time.

Note:

If you plan to reuse the replaced nodes to create a test cluster or to add to another cluster, you must ensure that the WWNNs are set to a unique number on your SAN. Document the original WWNN of the replacement 2145-CF8 nodes so that the removed nodes can be set to use these WWNNs. Then each node will still have a unique WWNN. Failure to do this can lead to duplicate WWNNs and WWPNs, causing unpredictable SAN problems.

Non-disruptive upgrade procedure

This non-disruptive procedure replaces 2145-4F2 nodes with 2145-CF8 nodes while upgrading the SAN Volume Controller cluster to version 5.1.0 without any interruption to host I/O operations. There is a small amount of risk during the hardware replacement procedure, due to the single point of failure introduced while a node is being replaced. To perform this procedure during a maintenance period with all host operations stopped, see the ***Disruptive* upgrade procedure – upgrading clusters with 2145-4F2 nodes to SAN Volume Controller version 5.1.0** on page 25**.**

Important Notes:

Before beginning the non-disruptive upgrade procedure, it is important to consider the following points:

a. If the configuration node changes, because the original one is removed from the cluster, you will have to reconnect to the CLI or Console.

b. If the node containing the IBM2145_REMOTE_UPGRADE package is upgraded or rebooted the package will be deleted. You will have to repeat step 4 on page 11 to upload the package again.

c. The cluster will be performing a software upgrade for the duration of this procedure, and only the CLI commands and Console tasks permitted during a normal software upgrade are permitted. You should not attempt to perform any other configuration task while performing this procedure.

d. Do not power off a SAN Volume Controller node that you are replacing by using the front panel unless you are specifically instructed to do so. Otherwise, always power off nodes by using the command line interface as described in the procedure.

To perform the non-disruptive procedure:

1. Decide the order in which you will replace the nodes. Using the configuration chart to help you, the order should be decided so that:
   - One node in each I/O group is replaced first and then the second node in each I/O group.
   - If the configuration node is being replaced, it is the last node to be replaced.

   Then continue with step 2 for the first node to be replaced. You will repeat steps 2 to 20 for each node being replaced. For each node use the front panel ID and Node Name recorded in the configuration chart to locate the physical node that will be replaced. The front panel ID is on a label on the front of the node. The node's name can be displayed on the front panel by pressing the down button until the "Node:" display is shown.

2. Issue the following CLI command to delete the node that is being replaced from the cluster and I/O group:

svctask rmnode <*node_name or node_id*>

Where <*node_name or node_id*> is the name or ID of the node.

Note: If all the other nodes in the cluster have been replaced, the cluster should upgrade shortly after the node is deleted.

3. Issue the following CLI command to ensure that the node you deleted, using the svctask rmnode command, is no longer a member of the cluster:

   svcinfo lsnode

   Verify that the removed node is not listed in the command output.

4. Label and disconnect the four fibre-channel cables and the Ethernet cable from the deleted node, using the following image as a guide to identifying the fibre-channel cables. Do not connect the cables to the replacement 2145-CF8 node yet.

   Important:
   You must reconnect the cables on the new 2145-CF8 node exactly as they were connected on the removed node. At the back of the node, the fibre-channel ports are numbered 1-4 from left to right. However, note that there are no numerical markings on the ports themselves. The cables must be reconnected to the same port number that they were removed from, or the fibre-channel port IDs will change, which can impact host access to VDisks or cause problems adding the new node to the cluster. Do not move any fibre-channel cables between ports in the switch or director either.

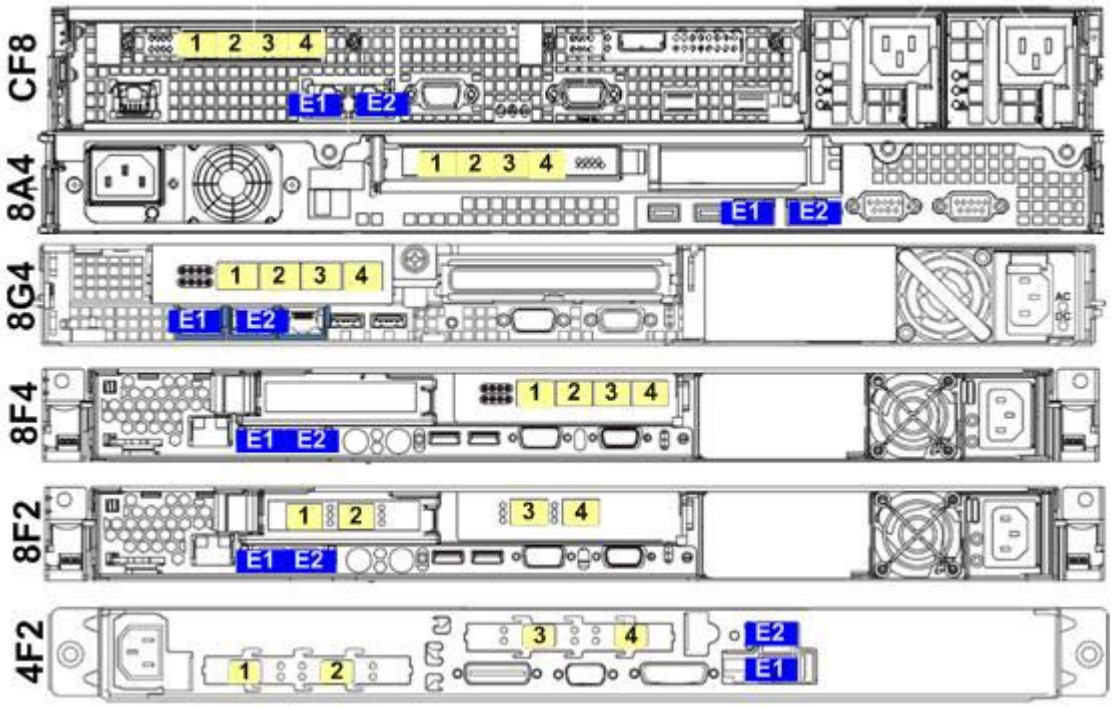

5. Unless you are going to discard the deleted node immediately, after removing the cables change the node's WWNN so that it can never conflict with the WWNN of the replacement 2145-CF8 node. Change the node's WWNN using the instructions in "Appendix B: Changing a Node's WWNN" on page 39. Check that the initial WWNN for the node matches the value recorded in "Appendix A: Configuration table", set the WWNN to "FFFFF".

6. Power off the node using the power button on the front panel

7. If desired, remove the node and the connected uninterruptible power supply from the rack.
   Important: A 2145-4F2 may be connected to a 2145 UPS. This 2U high unit can be connected to two 2145-4F2s. If you are replacing a 2145-4F2, check the type of uninterruptible power supply it is connected to; if it is connected to a 2145 UPS check to see if another node is still powered by it. Ensure you do not power the 2145 UPS off, or disconnect the other node's cables, if it is connected to two active nodes.
   Important: Be careful not to disturb any other nodes while removing the node. If a node accidentally powers off your host applications may not be able to access data and you may have to perform a cluster recovery.

8. If the nodes are not yet installed in their final location, install the replacement 2145-CF8 node and its 2145 UPS-1U uninterruptible power supply in the rack using the *SAN Volume Controller Hardware Installation Guide*, which is shipped with the node and available at www.ibm.com/storage/support/2145. Connect the 2145-CF8 node to the 2145 UPS-1U using the power and serial cable assembly. Do NOT connect any fibre-channel cables at this time.
   Important: Be careful not to disturb any other nodes while performing the node replacement. If a node accidentally powers off your host applications may not be able to access data you may have to perform a cluster recovery.

9. Connect the Ethernet cable you removed from the replaced node to Ethernet port 1 on the 2145-CF8 node.

10. Ensure that the replacement 2145-CF8 node is powered on with the Ethernet cable connected and the fibre-channel cables **disconnected**. Once the node has booted, you might receive Node Error 558 "The 2145 cannot see the fibre-channel fabric or the fibre-channel card port speed might be set to a different speed than the fibre-channel fabric"; this message is expected because the node was booted with no fibre cables connected.

    a. **Note:** Error 550, "Cannot form a cluster due to a lack of cluster resources", occurs if the 2145-CF8 node is recognized as part of a SAN Volume Controller cluster. This error can occur if the 2145-CF8 node had been added to another cluster and its state was not cleared before removing it. Delete the configuration information using the "Delete Cluster?" option on the front panel:

        i. Press the down button until you see "Node Name"
        ii. Press the right button until you see "Create Cluster?"
        iii. Press the select button.
        iv. The front panel will now display "Delete Cluster?"

v. Press and hold the up button.
vi. Press and release the select button.
vii. Release the up button.

11. Use the instructions in "Appendix B: Changing a Node's WWNN" on page 39 to record the WWNN of the new 2145-CF8 node and to change the WWNN to match the WWNN of the replaced node. Record the WWNN of the new 2145-CF8 in the column labeled "WWNN of replacement node (last 5 characters)", and find the WWNN of the replaced node in "Appendix A: Configuration table".

12. Ensure that the configuration node of the cluster still contains the IBM2145_REMOTE_UPGRADE package by running the command
    svcinfo lssoftwaredumps

    If the configuration node has changed, as it would have done if the original configuration node has been replaced or upgraded, the package will have to be uploaded again.

13. Connect the fibre-channel cables to the same port numbers on the new 2145-CF8 node that they were originally connected to on the replaced node.

    Important: **DO NOT** connect the new 2145-CF8 nodes to different ports at the switch or director; this would cause port IDs to change, which could impact host access to VDisks or cause problems when adding new 2145-CF8 nodes to the cluster. Moving the node cables to faster ports on the switch or director is a separate process that must be planned independently of this procedure.

14. If all the other nodes in the cluster have been replaced, every node will now have the version 5.1 code installed and the cluster should have upgraded to the new software version. If there are nodes in the cluster that have not been replaced go on to step 15 below. Otherwise run the command:
    svcinfo lscluster <cluster name>
    and check the "code_level" value. If it is no longer 4.3.1.X, skip to step 17 because any node added will automatically install the software version running on the cluster. If the cluster software version is still 4.3.1.X, contact IBM support for further assistance.

15. Upgrade the 2145-CF8 node to the SAN Volume Controller version 5.1.0.2 (or later) package that you downloaded as part of the "Initial steps to perform before beginning either procedure" on page 11 and renamed IBM2145_REMOTE_UPGRADE by following these instructions:

    Important:
    You must have exactly the same IBM2145_REMOTE_UPGRADE package on every node in the cluster. Do not change software versions midway through this process.

    a. On the front panel of the node, press the down button until the "Version:" panel is displayed. This shows the V.R.M.F (e.g. 5.1.0.0) of the software that is installed on the node.

b. If the version of V.R.M.F of the node matches the version of V5.1.0 software to be installed on the cluster, go to Step 16.

c. Press the right navigation button until the "Build:" panel displays. Then press and hold the up button, press and release the select button, and then release the up button.

d. This causes the SAN Volume Controller node to log onto the fabric and to discover the upgrade package IBM2145_REMOTE_UPGRADE package that you previously uploaded to the configuration node.

e. The node front panel will display "Shutting Down". It will then restart and upgrade the software on the node, and the front panel will cycle through a number of Booting codes (defined in the *IBM System Storage SAN Volume Controller: Troubleshooting Guide*, which is available at www.ibm.com/storage/support/2145).

f. Part way through the upgrade, the Front panel will display "Booting 384: *<front panel ID>*" for five seconds. Write down the front panel ID shown. The front panel ID displayed is the ID of the node that is providing the upgrade package to this node.

g. If the front panel ID displayed on the Booting 384 panel does not match the front panel ID of the configuration node you uploaded the IBM2145_REMOTE_UPGRADE file to, there is more than one node on the fibre-channel fabric that contains a copy of the IBM2145_REMOTE_UPGRADE file. Locate the node with the reported front panel ID and remove the redundant software package using the command "svctask cleardumps –prefix /home/admin/upgrade/IBM2145_REMOTE_UPGRADE *<node id or name>*". You may need to connect a CLI to a different cluster to run the command. Once the redundant software package has been removed, restart step 15 to retry the node upgrade.

h. When the upgrade is completed, the front panel will display "Cluster: <blank>". If a "Node Error:" is displayed, use the standard maintenance procedures to fix the error.

i. Repeat step a. to verify that the software level running on the node matches the version of software downloaded during the "Initial steps to perform before beginning either procedure" stage.

j. If the software level has not changed or has changed to an unexpected V.R.M.F., the upgrade has either failed to discover a suitable upgrade package, or it has discovered an unexpected upgrade package. Verify or correct the following requirements and then retry the procedure:

   i. The WWNN of the node is set correctly.
   ii. The special upgrade package that you uploaded as part of "Initial steps to perform before beginning either procedures" still exists on that node (Run the 'svcinfo lssoftwaredumps' command to check).
   iii. The fabric zoning allows the node to log in to the node that contains the

special upgrade package that you uploaded during the "Initial steps to perform before beginning either procedure" steps.

iv. The only node in the entire SAN fabric that contains an upgrade file named IBM2145_REMOTE_UPGRADE is the one used during the "Initial steps to perform before beginning either procedure" steps.

k. If you are unable to complete the software upgrade successfully, contact IBM support for assistance.

16. The 2145-CF8 is not currently a member of the cluster. Its upgrade mode must be set so that when it is added to a cluster it does not try to install the software version running on the cluster; instead it joins the cluster without changing its current software version, as if the cluster were upgrading to that level. To set the upgrade mode :

    a. On the front panel of the new 2145-CF8 node, press the down button until the "Node:" panel is displayed. Then use the right or left navigation button to display the "Status:" panel.

    b. Press and hold the up button, press and release the select button and then release the up button to select the upgrade mode. The panel will display "Set Upgrade?"

    c. Press Select to confirm and enter the upgrade mode. The front panel will display "Node Error: 950 Upgrade". This message is expected and simply indicates that the node is in upgrade mode. Any other hardware errors that occur, such as "Node Error: 558", will be displayed instead of "Node Error: 950 Upgrade". Any other Node Errors should be fixed to ensure that the node can be added to the cluster.

17. Issue the following CLI command to verify that the node is available as a candidate to add to the cluster:

    svcinfo lsnodecandidate

18. Add the 2145-CF8 node to the cluster, to the same I/O group that the replaced node was in. For command usage details, see the **svctask addnode** command in the *IBM System Storage SAN Volume Controller: Command-Line Interface User's Guide*, which is available at www.ibm.com/storage/support/2145. For example:

    svctask addnode -wwnodename 50050768010<*wwnn_arg*> -iogrp <*iogroup_name*> -name <*node name*>

    The parameter values that need to be used are those recorded for the replaced node in Appendix A: Configuration table*; w*here <*wwnn_arg*> is the content of the "WWNN (last 5 characters)" column (with no space between the preceding numbers and the WWNN digits), <*iogroup_name*> is the name from the "IO Group ID / Name" column and <node name> is the name from the "Node ID / Name" column.
    Note: If the replaced node had a default name, with the format "nodeN" (where N is a number), the –name <node name> parameter should not be specified, the SAN Volume Controller will allocate a different default name for the node. You can modify the name

when you have completed this procedure.

Note: If you skipped steps 15 and 16, when you run the addnode command the software on the node is automatically updated to match the cluster version. In this case, the addnode command takes longer than usual because the node has to reboot and upgrade.

Run the command

svcinfo lsnode

and verify the node is in the online state. Allow 15 minutes for the node to transition from the adding state to the online state.

If the upgrade mode was not set when a 2145-CF8 node was added to the version 4.3.1 cluster, the node will stay in the adding state forever in the svcinfo lsnode output. Delete the node from the cluster using the svctask rmnode command, and resume from step 16.

19. Verify that all VDisks for the I/O group are back online and are no longer degraded. If the node replacement process is being done without I/O, you must still wait approximately 30 minutes to verify that the replacement 2145-CF8 node is back online and available before you upgrade the other node in the I/O group. The commands required to perform these checks are available in the prerequisites section.

Important Notes:

- You do not have to reconfigure the host multipathing device drivers because the replacement 2145-CF8 node uses the same WWNN and WWPNs as the replaced node. The multipathing device drivers should detect the recovery of paths that are available to the replacement node.

- The host multipathing device drivers typically take up to 30 minutes to recover the paths. Therefore, do not upgrade the remaining node in the I/O group for at least 30 minutes after successfully upgrading the first node in the I/O group. During the interim, you can upgrade other nodes in other I/O groups.

20. Review the documentation that is provided with your multipathing device driver for information on how to query paths to ensure that all paths have been recovered before proceeding to the next step. For SDD, the command is “datapath query device”. Check several of your critical hosts and at least one host that is running each operating system to ensure that the paths have recovered before replacing the second node in the I/O group. Otherwise, a loss of access to host data could result.

21. Repeat steps 2 through 20 for each node in the cluster that you are replacing with 2145-CF8 nodes. Follow the order you determined in step 1

22. If all the nodes have been replaced skip to step 34 on page 23. Otherwise you must manually upgrade the nodes that have not been replaced. Decide the order in which you will upgrade the nodes. Using the configuration chart to help you, the order should be decided so that:

- One node in each I/O group is upgraded first and then the second node in each I/O group.
- If the configuration node is being upgraded, it is the last node to be upgraded.

Then continue with step 23 for the first node to be upgraded. You will repeat steps 23 to 32 for each node being upgraded.

23. Issue the following CLI command to delete the node from the cluster and I/O group:

    svctask rmnode <*node_name or node_id*>

    Where <*node_name or node_id*> is the name or ID of the node.

    Note: If all the other nodes in the cluster have been replaced or upgraded, the cluster should upgrade shortly after the node is deleted. At this point the upgrade has completed, and the last remaining node only needs to be added to the cluster.

24. Issue the following CLI command to ensure that the node you deleted, using the svctask rmnode command, is no longer a member of the cluster:

    svcinfo lsnode

    Verify that the deleted node is not listed in the command output.

25. If all the other nodes in the cluster have been replaced or upgraded, every node will now have the version 5.1 code installed and the cluster should have upgraded to the new software version. If there are nodes in the cluster that have not been replaced or upgraded go on to step 26 below. Otherwise run the command:
    svcinfo lscluster <cluster name>
    and check the "code_level" value. If it is no longer 4.3.1.X, skip to step 29 because any node added will automatically install the software version running on the cluster. If the cluster software version is still 4.3.1.X, contact IBM support for further assistance.

26. Check that the configuration node of the cluster still contains the IBM2145_REMOTE_UPGRADE package by running the command
    svcinfo lssoftwaredumps

    If the configuration node has changed, as it would have done if the original configuration node has been replaced or upgraded, the package will have to be uploaded again.

27. Upgrade the node with the SAN Volume Controller version 5.1.0.2 (or later) package that you downloaded during the "Initial steps to perform before beginning either procedure" on page 11 steps.

Important:
You must have exactly the same IBM2145_REMOTE_UPGRADE package on every node in the cluster. Do not change software versions midway through this process.

a. Press the down navigation button until the "Version" panel displays.

b. Press the right navigation button until the "Build:" panel displays. Then press and hold the up button, press and release the select button, and then release the up button.

c. This causes the SAN Volume Controller node to log onto the fabric and to discover the upgrade package IBM2145_REMOTE_UPGRADE package that you previously uploaded to the configuration node.

d. The node front panel will display "Shutting Down". It will then restart and upgrade the software on the node, and the front panel will cycle through a number of Booting codes (defined in the *IBM System Storage SAN Volume Controller: Troubleshooting Guide*, which is available at www.ibm.com/storage/support/2145).

e. Part way through the upgrade, the Front panel will display "Booting 384: <*front panel ID*>" for five seconds. Write down the front panel ID shown. The front panel ID displayed is the ID of the node that is providing the upgrade package to this node.

f. When the upgrade is completed, the front panel will display "Cluster: <blank>". If a "Node Error:" is displayed, press the down button to display the "Cluster:" panel.

g. If the front panel ID displayed on the Booting 384 panel does not match the front panel ID of the configuration node you uploaded the IBM2145_REMOTE_UPGRADE file to, there is more than one node on the fibre-channel fabric that contains a copy of the IBM2145_REMOTE_UPGRADE file. Locate the node with the reported front panel ID and remove the redundant software package using the command "svctask cleardumps –prefix /home/admin/upgrade/IBM2145_REMOTE_UPGRADE <*node id or name*>". You may need to connect a CLI to a different cluster to run the command. Once the redundant software package has been removed, restart step 27 to retry the node upgrade.

h. On the front panel of the node, press the down button until the "Version:" panel is displayed. This shows the V.R.M.F (e.g. 5.1.0.2) of the software that is installed on the node.

i. If the software level has not changed or has changed to an unexpected V.R.M.F., the upgrade has either failed to discover a suitable upgrade package, or it has discovered an unexpected upgrade package. Verify or correct the following requirements and then retry the procedure:

    i. The WWNN of the node is set correctly.
    ii. The special upgrade package that you uploaded during the "Initial steps to perform before beginning either procedure" steps still exists on that node (Run the 'svcinfo lssoftwaredumps' command to check).
    iii. The fabric zoning to allows the node to log in to the node that contains the special upgrade package that you uploaded during the prerequisite steps.
    iv. The only node in the entire SAN fabric that contains an upgrade file named

IBM2145_REMOTE_UPGRADE is the one used during the "Initial steps to perform before beginning either procedure" steps.

j. If you are unable to complete the software upgrade successfully, contact IBM support for assistance.

28. The node is not currently a member of the cluster. Its upgrade mode must be set so that when it is added to a cluster it does not try to install the software version running on the cluster; instead it joins the cluster without changing its current software version, as if the cluster were upgrading to that level. To set the upgrade mode:

    a. On the front panel of the node, press the down button until the "Node:" panel is displayed. Then use the right or left navigation button to display the "Status:" panel.

    b. Press and hold the up button, press and release the select button and then release the up button to select the upgrade mode. The panel will display "Set Upgrade?"

    c. Press Select to confirm and enter the upgrade mode. The front panel will display "Node Error: 950 Upgrade". This message is expected and simply indicates that the node is in upgrade mode. Any other hardware errors that occur, such as "Node Error: 558", will be displayed instead of "Node Error: 950 Upgrade". Any other Node Errors should be fixed to ensure that the node can be added to the cluster.

29. Issue the following CLI command to verify that the node is available as a candidate to add to the cluster:

    svcinfo lsnodecandidate

30. Add the node to the cluster, to the same I/O group that it was originally in. For command usage details, see the **svctask addnode** command in the *IBM System Storage SAN Volume Controller: Command-Line Interface User's Guide*, which is available at www.ibm.com/storage/support/2145. For example:

    svctask addnode -wwnodename 50050768010<*wwnn_arg*> -iogrp <*iogroup_name*> -name <*node name*>

    The parameter values that need to be used are those recorded for the replaced node in Appendix A: Configuration table*; w*here <*wwnn_arg*> is the content of the "WWNN (last 5 characters)" column (with no space between the preceding numbers and the WWNN digits), <*iogroup_name*> is the name from the "IO Group ID / Name" column and <node name> is the name from the "Node ID / Name" column.
    Note: If the replaced node had a default name, with the format "nodeN" (where N is a number), the –name <node name> parameter should not be specified, the SAN Volume Controller will allocate a different default name for the node. You can modify the name when you have completed this procedure.

    Note: If you skipped steps 27 and 28 because this is the last node to upgrade, when you run the addnode command the software on the node is automatically updated to match the

cluster version. In this case, the addnode command takes longer than usual because the node has to reboot and upgrade.

Run the command

svcinfo lsnode

and verify the node is in the online state. Allow 15 minutes for the node to transition from the adding state to the online state.

Also verify that the node is running the correct software version. On the front panel press the down button a number of times until the "Version:" panel is displayed and check that the version displayed matches the version of the software you downloaded.

If the upgrade mode was not set when the node was added to the 4.3.1 cluster, the node will downgrade and rejoin the cluster at version 4.3.1. Delete the node from the cluster and resume from step 25 on page 20.

31. Verify that all VDisks for the I/O group are back online and are no longer degraded. If the node replacement process is being done without I/O, you must still wait approximately 30 minutes to verify that the node is back online and available before you upgrade the other node in the I/O group. The commands required to perform these checks are available in the prerequisites section.

Important Notes:

- The host multipathing device drivers typically take up to 30 minutes to recover the paths. Therefore, do not upgrade the remaining node in the I/O group for at least 30 minutes after successfully upgrading the first node in the I/O group. During the interim, you can upgrade other nodes in other I/O groups.

32. Review the documentation that is provided with your multipathing device driver for information on how to query paths to ensure that all paths have been recovered before proceeding to the next step. For SDD, the command is "datapath query device". Check several of your critical hosts and at least one host that is running each operating system to ensure that the pathing has recovered before replacing the second node in the I/O group. Otherwise, a loss of access to host data could result.

33. Repeat steps 23 on page 20 through 32 on page 23 for each node in the cluster that you are not replacing. Follow the order you determined in step 22 on page 19.

34. Verify that the SAN Volume Controller cluster is operating correctly.
    a. Run the command:
       svcinfo lscluster <cluster name>
       and check the code_level value matches the version you are upgrading to. If it does not, contact IBM Support.

    b. Check the error log for any unfixed errors and resolve them using the standard procedures.

c. Issue the following commands from the CLI. If any of the commands return an offline or degraded status, resolve the problems with the MDisks or VDisks before continuing.

   - svcinfo lsvdisk -filtervalue "status=degraded"
   - svcinfo lsvdisk -filtervalue "status=offline"

   - svcinfo lsmdisk -filtervalue "status=degraded"
   - svcinfo lsmdisk -filtervalue "status=offline"

d. For each storage system (controller) that is attached to the SAN Volume Controller, issue the following command to verify that each controller's degraded status is "no". If any storage system's degraded status value is "yes", resolve the issue before continuing.

   - svcinfo lscontroller <object_id or object_name>

   Where <object_id or object_name> is the name, or ID, of the storage system (controller) that you want to view.

35. To restore your environment to its fully operational state, restart any operations that you stopped before starting this procedure. For example:

    - Restart Metro Mirror and Global Mirror
    - Restart TPC performance data gathering and cluster probes
    - Restart the Call Home service
    - Restart any automated scripts and jobs

Notes on reusing replaced nodes:

If you skipped step 7 and you are reusing the replaced nodes in another cluster, these nodes retain configuration information from the original cluster because they were powered off before being deleted. When you power these nodes on again, the original cluster name will appear on the front panel. You might also receive error code "550 - Cannot form a cluster due to a lack of cluster resources". Both of these situations are expected – use the front panel on each node to delete the original cluster name.

ATTENTION: You must ensure unique WWNNs to avoid the problems caused by duplicate WWNNs. Power on the nodes *before attaching them to a SAN* and change the WWNN value to FFFFF or to a unique value. If the replaced nodes are still attached to the SAN when you power them on and you have not changed the WWNN, power off the nodes immediately. Do not add any node with a WWNN ending with "FFFFF" to a SAN Volume Controller cluster.

Disruptive upgrade procedure

This *disruptive* procedure replaces 2145-4F2 nodes with 2145-CF8 nodes while upgrading the SAN Volume Controller cluster to version 5.1.0 *with complete cessation of all host I/O operations*. For the non-disruptive procedure, see "*Non-disruptive* upgrade procedure" on page 13.

During this procedure you will separate the cluster into two halves. Each half contains one node from each I/O group. The nodes in one half of the cluster are then deleted from the cluster; this leaves the minimum number of nodes that can continue to operate as a cluster online, although there is no redundancy. The nodes in the half of the cluster that have been deleted from the cluster are replaced with 2145-CF8 nodes (if necessary). Then version 5.1 software is installed on the nodes and they are set into the special upgrade mode. These nodes are then added back to the cluster, they will continue to operate the version 5.1.0 software in upgrade mode, while the cluster continues to run at version 4.3.1.

When the nodes in the half cluster that are in the special upgrade mode have been added back to the cluster the nodes in the other half of the cluster are deleted from the cluster. When the final node in this half is deleted the remaining nodes, which are all in upgrade mode, can complete the upgrade and the half cluster will start to operate at version 5.1.

Any of the nodes that have now been deleted from the cluster that need to be replaced are now replaced. There is no need to download the version 5.1 software to these nodes. When this set of nodes are re-added to the cluster, which is now running at version 5.1, they will automatically run at that version.

Whenever nodes are replaced during this procedure, all the nodes in the cluster are powered off. This protects the cluster from accidental removal of power from a node.

Important Notes:

Before beginning the node replacement procedures, it is important to consider the following notes:

a. Verify that all the prerequisites listed in "Upgrade prerequisites" on page 7 have been met, and the initial steps listed in "Initial steps to perform before beginning either procedure" on page 11 have been performed before starting this procedure.

b. Only power off a node via the front panel when specifically instructed to do so; otherwise, always power off nodes using the command line interface as described in the procedure.

c. Only work on the hardware when all the nodes are powered off. This protects the integrity of the SAN Volume Controller system. Note, however, that there may be other equipment in the rack that must not be disturbed.

To perform the disruptive procedure:

1. Stop ALL host I/O operations to the cluster.

2. Using the information recorded in "Appendix A: Configuration table" on page 37, select one node from each I/O group to upgrade first, label these as Node Set A. Do not select the configuration node in this group.

   If any of the 2U high 2145 UPSs are installed, check whether any uninterruptible power supply connects to two nodes. If they do either select both the nodes in this group or select neither of them. Keeping them in the same group makes the hardware replacement easier. To see if an uninterruptible power supply connects to two nodes run the command:

   svcinfo lsnode

   and see if the same "UPS_serial_number" is shown for two nodes.

   Exactly half the nodes should have been selected. Check carefully that you have not selected both nodes from an I/O group; if you have you will destroy some of the cluster configuration when you run these procedures. On the Configuration table, label the other nodes as Node Set B. These procedures will refer to these selections as Node Set A and Node Set B.

   Identify each of the nodes selected in Node Set A. For each node use the front panel ID and Node Name recorded in the configuration chart to locate the physical node. The front panel ID is on a label on the front of the node. The node's name can be displayed on the front panel by pressing the down button until the "Node:" display is shown.

3. For each of the nodes that you selected in Node Set A in step 1, perform the following commands:

   - Issue the following CLI command to delete this node from the cluster and I/O group:

     svctask rmnode <*node_name or node_id*>

     Where <*node_name or node_id*> is the name or ID of the node that you want to delete.

   - Issue the following CLI command to ensure that the node you deleted, using the svctask rmnode command, is no longer a member of the cluster:

     svcinfo lsnode

     Verify that the deleted node is not listed in the command output.

4. If any of the nodes in Node Set A are going to be replaced, and you are not going to immediately discard the node perform the instructions in "Appendix B: Changing a Node's WWNN" on page 39 to change each nodes WWNN to "FFFFF"

Important: Before continuing, clearly record or mark all nodes that you have deleted from the cluster.

5. Power off all nodes in Node Set A using the front panel power button. Be careful not to power off an incorrect node.

6. Verify that the SAN Volume Controller write cache has been successfully flushed to the backend storage systems by running the following command:

   ```
   svcinfo lsvdisk -nohdr | while read id name rest; do svcinfo lsvdisk $id | while read key value ; do if [ "$key" == "fast_write_state" ] ; then if [ "$value" != "empty" -a "$copy" == "n/a" ] ; then echo "Vdisk $name ($id) - cache state $value"; fi; fi; if [ "$key" == "copy_id" ] ; then copy=$value; fi; done;done
   ```

   If the command returns no results, all cache data has been successfully flushed to the backend storage arrays. The cache has been successfully flushed if the fast_write_state of a VDisk is "empty". If the command returns one or more VDisks with a fast_write_state of "not_empty", continue running the command until the cache data for those VDisks has been successfully flushed. Investigate any other fast_write_state values thoroughly before continuing. If the cache has not been completely flushed in 30 minutes, contact IBM support.

7. Power off the remainder of the cluster (Node Set B). Powering off all nodes puts SAN Volume Controller into a safe state while hardware is replaced. For example, if an incorrect cable is removed, SAN Volume Controller is not compromised. Use the following command to power off the remainder of the cluster:

   ```
   svctask stopcluster
   ```

8. When all of the cluster nodes are powered off, replace all nodes in Node Set A that are being replaced with 2145-CF8 nodes. For each node being replaced:

   - Record and mark the fibre-channel cables with the SAN Volume Controller node port number (1-4) and remove the cables from the back of the node. You must reconnect the cables to the new 2145-CF8 node exactly as they were connected to the removed node. At the back of the node, the fibre-channel ports are numbered 1-4 from left to right. However, note that there are no numerical markings on the ports themselves. The cables must be reconnected to the same port number that they were removed from, or the fibre-channel port IDs will change, which can impact host access to VDisks or cause problems adding the new node to the cluster. Do not move any fibre-channel cables between ports in the switch or director either.

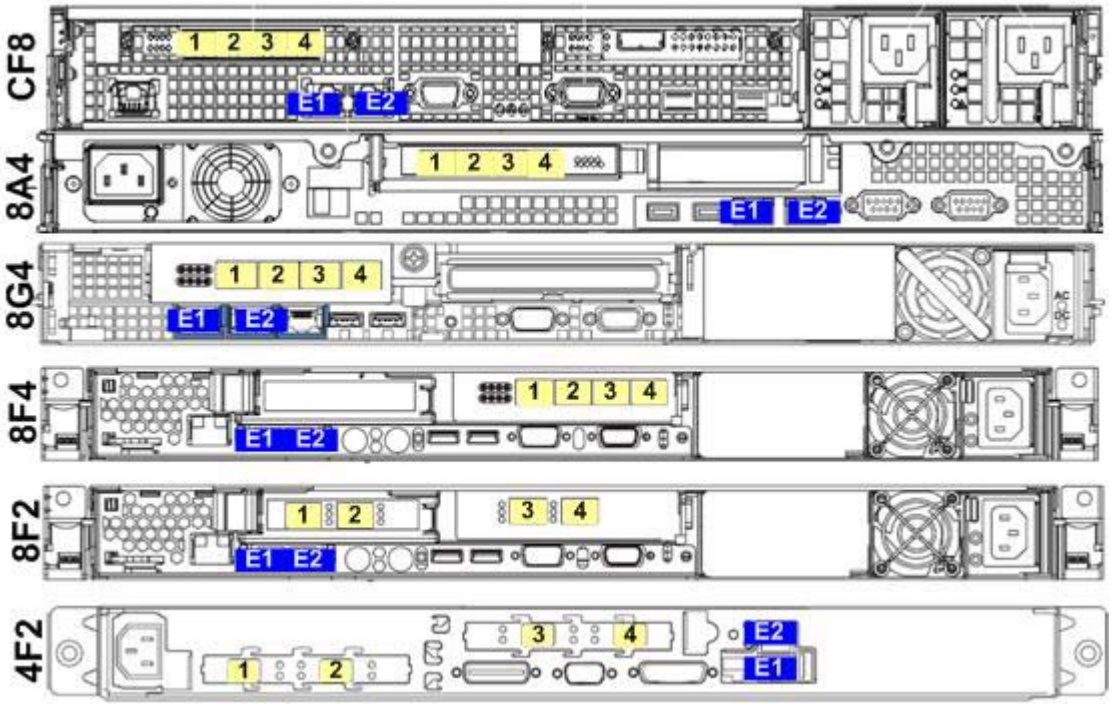

- Remove the nodes being replaced and the uninterruptible power supplies from the rack.
  Important: A 2145-4F2 node may be connected to a 2145 UPS. This 2U high unit can be connected to two 2145-4F2 nodess. If you are replacing a 2145-4F2 node, check the type of uninterruptible power supply it is connected to; if it is connected to a 2145 UPS check to see if another node is in Node Set B and still requires power from it and do NOT remove the 2145 UPS if it does.

- Install the replacement nodes and uninterruptible power supplies in the rack and connect the power, serial, Ethernet and Fibre Channel cables to the nodes according to the *SAN Volume Controller Hardware Installation Guide*, which is shipped with the nodes and available at www.ibm.com/storage/support/2145.

Important:

- You must connect the fibre-channel cables to the same port that each cable was connected to on the replaced node.
- You must connect the Ethernet cable to the E1 port on the 2145-CF8 node.

9. Power on all nodes using the front panel power buttons. Nodes from Node Set B, which were not deleted from the cluster, and have therefore not been upgraded, should power on and reform the cluster. If the nodes in Node Set B do not reform the cluster, use the standard service procedure given in the *SAN Volume Controller Troubleshooting Guide,* which is available at www.ibm.com/storage/support/2145, to resolve the problem.

10. For all new 2145-CF8 nodes, use the instructions in "Appendix B: Changing a Node's WWNN" on page 39 to record the WWNN of the new 2145-CF8 node and to change the WWNN to match the WWNN of the replaced node. In "Appendix A: Configuration table"

record the WWNN of the new 2145-CF8 in the column labeled "WWNN of replacement node (last 5 characters)", and find the WWNN of the replaced node in the column labeled "WWNN (last 5 characters)". As you might be changing a number of nodes ensure you select the correct row from the table.

11. Reconnect your command line session, and then follow the instructions in "Initial steps to perform before beginning either procedure" on page 11 to re-upload the **IBM2145_REMOTE_UPGRADE** file. The copy uploaded previously would have been deleted when the configuration node was restarted.

12. Upgrade all nodes in Node Set A with the SAN Volume Controller version 5.1.0 package that you downloaded and renamed IBM2145_REMOTE_UPGRADE. Repeat the following sub-steps for all nodes in Node Set A.

    Important:
    You must use the same IBM2145_REMOTE_UPGRADE package on every node. Do not change software versions midway through this step.

    a. From the front panel of the node, press the down button until the "Version:" panel is displayed. The SAN Volume Controller V.R.M.F (e.g. 5.1.0.2) that is installed on the node is displayed.

    b. If the V.R.M.F of the node matches the version of SAN Volume Controller software that you are installing on the cluster, repeat this step (12) for the next node in Node Set A.

    c. Press the right button until the "Build:" panel is displayed. Then press and hold the up button, press and release the select button, then release the up button.

    d. This will cause the node to log onto the fabric and to try to discover the upgrade package that you uploaded in step 11.

    e. The node front panel will display "Shutting Down". It will then restart and upgrade the software, and the front panel will cycle through a number of booting codes that are defined in the *SAN Volume Controller Troubleshooting Guide*, which is available at www.ibm.com/storage/support/2145.

    f. Part way through the installation, the Front panel will display "Booting 384 Node: <front panel id>" for five seconds. The front panel ID indicates the node that is providing the upgrade package to this node.

    g. If the front panel ID does not match the expected front panel ID, then there is more than one node on the fibre-channel fabric that contains a copy of the IBM2145_REMOTE_UPGRADE file. Let the upgrade complete, then locate the node with the reported front panel ID and remove the extraneous software package by running the command: "svctask cleardumps –prefix /home/admin/upgrade/IBM2145_REMOTE_UPGRADE <node id or name>"

You may need to connect a CLI to a different cluster to run the command. Return to step 12(a) on page 29 and repeat the code upgrade.

h. When the upgrade is completed, the front panel will display "Cluster: <blank>". If a "Node Error:" is displayed, use the normal procedures detailed in the *SAN Volume Controller Troubleshooting Guide*, which is available at www.ibm.com/storage/support/2145, to resolve the error.

i. From the front panel of the node, press the down button until the "Version:" panel is displayed. The SAN Volume Controller V.R.M.F (e.g. 5.1.0.2) that is installed on the node is displayed. Verify that the node is running the expected software level.

j. If the software level has not changed or has changed to an unexpected version, then the upgrade has either failed to discover a suitable upgrade package or it has discovered an unexpected upgrade package. Verify the following requirements and then retry the procedure:

   i. The WWNN of the node is set correctly.
   ii. The special upgrade package that was uploaded in step 11 on page 29 still exists on that node ('svcinfo lssoftwaredumps' can be used to check this).
   iii. The fabric contains zoning that allows this node to log in to the node that contains the special upgrade package.
   iv. The only node in the entire SAN fabric that contains an upgrade file named IBM2145_REMOTE_UPGRADE is the one used in step 11 on page 29.

k. If you are unable to upgrade the software successfully, contact IBM support.

13. Set the upgrade mode on all nodes that you selected to be in Node Set A in step 1 on page 26. Repeat the following sub-steps for all nodes in Node Set A:

    a. From the front panel of the node, press the down button until the "Node:" panel is displayed, then use the right or left navigation button to display the "Status:" panel.

    b. Press and hold the up button, press and release the select button and then release the up button to select the upgrade mode. The panel should display "Set Upgrade?"

    c. Press Select to confirm and enter the upgrade mode. The front panel should display "Node Error: 950 Upgrade". This message is expected and simply indicates that the node is in the upgrade mode. Any other hardware errors that exist, such as "Node Error: 558", will be displayed instead of "Node Error: 950 Upgrade". Any other Node Errors should be fixed to ensure that the node can be added to the cluster.

14. Issue the following CLI command to verify that all the nodes that you selected in step 1 are candidates to add to the cluster:

    svcinfo lsnodecandidate

    All the nodes in Node Set A should be listed.

Important:

- If the WWNN does not match the original node's WWNN exactly as recorded in Appendix A: Configuration table, you must repeat step 10 on page 28.

15. Add each of the nodes in Node Set A to the cluster, to the same I/O group that it was originally in. For command usage details, see the **svctask addnode** command in the *IBM System Storage SAN Volume Controller: Command-Line Interface User's Guide*, which is available at www.ibm.com/storage/support/2145. For example:

    svctask addnode -wwnodename 50050768010<*wwnn_arg*> -iogrp <*iogroup_name*> -name <*node name*>

    The parameter values that need to be used are those recorded for the node in Appendix A: Configuration table on page 37*; w*here <*wwnn_arg*> is the content of the "WWNN (last 5 characters)" column (with no space between the preceding numbers and the WWNN digits), <*iogroup_name*> is the name from the "IO Group ID / Name" column and <node name> is the name from the "Node ID / Name" column.
    Note: If the replaced node had a default name, with the format "nodeN" (where N is a number), the –name <node name> parameter should not be specified, the SAN Volume Controller will allocate a different default name for the node. You can modify the name when you have completed this procedure.

    Important:

    Ensure you add the node to the same I/O group as the node it replaced. If you do not you might have data corruption.

    If the upgrade mode was not set when a 2145-CF8 node was added to the 4.3.1 cluster, the node will stay in the adding state forever in the svcinfo lsnode output. Delete the node from the cluster using the svctask rmnode command, and resume from step 12 on page 29 for this node only.

    If the upgrade mode was not set when a node that is not a 2145-CF8 is added to the 4.3.1 cluster, the node will downgrade and rejoin the cluster at version 4.3.1. Delete the node from the cluster, and then resume from step 12 on page 29 for just this node.

16. Verify that all VDisks are back online and are no longer degraded.

17. The cluster now has half its nodes upgraded. The initial steps performed on Node Set A are now performed on Node Set B. For each of the nodes that you selected in Node Set B in step 1, perform the following commands, delete the configuration node last:

    - Issue the following CLI command to delete this node from the cluster and I/O group:

      svctask rmnode <*node_name or node_id*>

      Where <*node_name or node_id*> is the name or ID of the node that you want to delete.

- Issue the following CLI command to ensure that the node you deleted, using the svctask rmnode command, is no longer a member of the cluster:

  svcinfo lsnode

  Verify that the deleted node is not listed in the command output.

18. When the last node in Node Set B was deleted, the remainder of the cluster should have completed the software upgrade to the version of 5.1.0 code downloaded. Reconnect your CLI session to the cluster and run the command svcinfo lscluster <cluster name> and check the "code_level" value matches the version you are upgrading to. If it does not contact IBM Support.

19. If any of the nodes in Node Set B are going to be replaced, and you are not going to immediately discard the node perform the instructions in "Appendix B: Changing a Node's WWNN" on page 39 to change each nodes WWNN to "FFFFF".

Important: Before continuing, clearly record or mark all nodes that you have deleted from the cluster.

20. Power off all nodes in Node Set B using the front panel power button. Be careful not to power off an incorrect node.

21. Verify that the SAN Volume Controller write cache has been successfully flushed to the backend storage systems by running the following command:

    svcinfo lsvdisk -nohdr | while read id name rest; do svcinfo lsvdisk $id | while read key value ; do if ["$key" == "fast_write_state"] ; then if ["$value" != "empty" -a "$copy" == "n/a"] ; then echo "Vdisk $name ($id) - cache state $value"; fi; fi; if ["$key" == "copy_id"] ; then copy=$value; fi; done;done

    If the command returns no results, all cache data has been successfully flushed to the backend storage arrays. The cache has been successfully flushed if the fast_write_state of a VDisk is "empty". If the command returns one or more VDisks with a fast_write_state of "not_empty", continue running the command until the cache data for those VDisks has been successfully flushed. Investigate any other fast_write_state values thoroughly before continuing. If the cache has not been completely flushed in 30 minutes, contact IBM support.

22. Power off the remainder of the cluster (Node Set A) by issuing the following command. Powering off all nodes puts SAN Volume Controller into a safe state while hardware is replaced. For example, if an incorrect cable is removed, SAN Volume Controller is not compromised.

    svctask stopcluster

23. When all of the cluster nodes are powered off, replace those nodes in Node Set B that are being replaced with 2145-CF8 nodes:

- Record and mark the fibre-channel cables with the SAN Volume Controller node port number (1-4) and remove the cables from the back of the node. You must reconnect the cables to the new 2145-CF8 node exactly as they were connected to the removed node. At the back of the node, the fibre-channel ports are numbered 1-4 from left to right. However, note that there are no numerical markings on the ports themselves. The cables must be reconnected to the same port number that they were removed from, or the fibre-channel port IDs will change, which can impact host access to VDisks or cause problems adding the new node to the cluster. Do not move any fibre-channel cables between ports in the switch or director either.

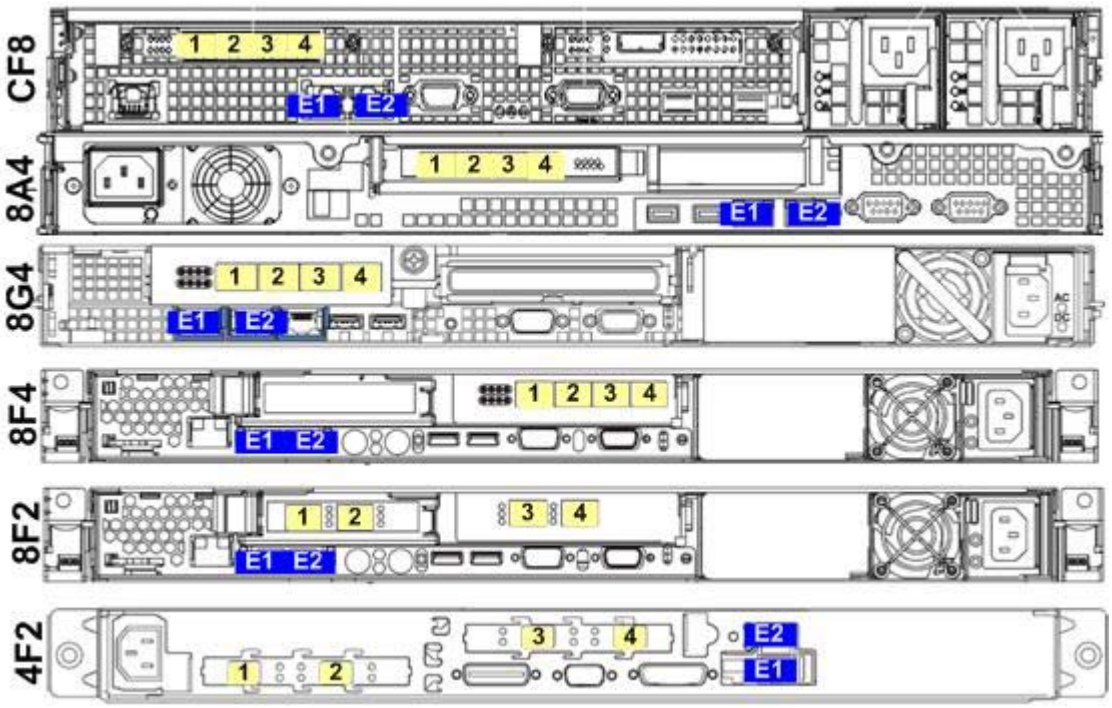

- Remove the nodes being replaced and the uninterruptible power supplies from the rack. Any 2145 UPSs can be safely removed at this time because they can only be connected to 2145-4F2 nodes.

- Install the replacement nodes and uninterruptible power supplies in the rack and connect the power, serial, Ethernet and Fibre Channel cables to the nodes according to the *SAN Volume Controller Hardware Installation Guide*, which is shipped with the nodes and available at www.ibm.com/storage/support/2145.

Important:

- You must connect the fibre-channel cables to the same port that each cable was connected to on the replaced node.
- You must connect the Ethernet cable to the E1 port on the 2145-CF8 node.

24. Power on all nodes using the front panel power buttons. Nodes from Node Set A, which were not deleted from the cluster, should power on and reform the cluster. If these nodes do

not reform the cluster, use the standard service procedures from the *SAN Volume Controller: Troubleshooting Guide*, which is available at www.ibm.com/storage/support/2145, to resolve the problem.

25. For all new 2145-CF8 nodes in Node Set B, use the instructions in "Appendix B: Changing a Node's WWNN" on page 39 to record the WWNN of the new 2145-CF8 node and to change the WWNN to match the WWNN of the replaced node. In "Appendix A: Configuration table" record the WWNN of the new 2145-CF8 in the column labeled "WWNN of replacement node (last 5 characters)", and find the WWNN of the replaced node in the column labeled "WWNN (last 5 characters)". Because you might be changing a number of nodes ensure you select the correct row from the table.

26. After all hardware changes are complete, add each node in Node Set B to the cluster. Because the cluster has already upgraded, there is no need to set the upgrade mode. The nodes will install the correct software level and restart before joining the cluster. Add each of the nodes in Node Set B to the cluster, to the same I/O group that it was originally in. For command usage details, see the **svctask addnode** command in the *IBM System Storage SAN Volume Controller: Command-Line Interface User's Guide*, which is available at www.ibm.com/storage/support/2145. For example:

    svctask addnode -wwnodename 50050768010<*wwnn_arg*> -iogrp <*iogroup_name*> -name <*node name*>

    The parameter values that need to be used are those recorded for the node in Appendix A: Configuration table on page 37*; w*here <*wwnn_arg*> is the content of the "WWNN (last 5 characters)" column (with no space between the preceding numbers and the WWNN digits), <*iogroup_name*> is the name from the "IO Group ID / Name" column and <node name> is the name from the "Node ID / Name" column.
    Note: If the replaced node had a default name, with the format "nodeN" (where N is a number), the –name <node name> parameter should not be specified, the SAN Volume Controller will allocate a different default name for the node. You can modify the name when you have completed this procedure.

    Important:

    Ensure you add the node to the same I/O group as the node it replaced. If you do not you might have data corruption.

27. Verify that the SAN Volume Controller cluster is operating correctly.
    a. Run the command:
       svcinfo lscluster <cluster name>
       and check the code_level value matches the version you are upgrading to. If it does not contact IBM Support.

    b. Check the error log for any unfixed errors and resolve them using the standard procedures.

    c. Issue the following commands from the CLI. If any of the commands return an offline or degraded status, resolve the problems with the MDisks or VDisks before continuing.

svcinfo lsvdisk -filtervalue "status=degraded"
svcinfo lsvdisk -filtervalue "status=offline"

svcinfo lsmdisk -filtervalue "status=degraded"
svcinfo lsmdisk -filtervalue "status=offline"

d. For each storage system (controller) that is attached to the SAN Volume Controller, issue the following command to verify that each controller's degraded status is "no". If any storage system's degraded status value is "yes", resolve the issue before continuing.

svcinfo lscontroller <object_id or object_name>

Where <object_id or object_name> is the name, or ID, of the storage system (controller) that you want to view.

28. To restore your environment to its fully operational state, restart any operations that you stopped before starting the procedure. For example:

- Restart Metro Mirror and Global Mirror
- Restart TPC performance data gathering and cluster probes
- Restart the Call Home service
- Restart any automated scripts and jobs
- Restart host I/O.

Appendix A: Configuration table

Node ID / Name	IO Group ID / Name	Front Panel ID	WWNN (last 5 characters)	Original Hardware Type	Config Node (Yes/No)	WWNN of replacement node (last 5 characters)	Node being replaced (Yes/No)	Upgrade / replacement order	Node Set (A / B) (disruptive procedure)

Appendix B: Changing a Node's WWNN

You will have to change a node's WWNN a number of times during both the Non-Disruptive and the Disruptive procedure described in this document. A replacement node must be given the same WWNN as the node it replaces. If you do not do this the host applications may not be able to access their data through the SAN Volume Controller.

If you are not going to immediately discard the replaced nodes, you must reset its WWNN so that you do not have duplicate WWNNs on your SAN. Setting the WWNN to FFFFF indicates it is not valid. You might also set the replaced node's WWNN to the value the 2145-CF8 node that replaced it originally had.

A WWNN is 16 characters long. The first 11 characters cannot be changed for a SAN Volume Controller node. These instructions guide you through setting the last five characters.

When you are told to use these steps to change the WWNN of a node, the new WWNN is indicated. You must ensure the node's WWNN is set to exactly this value. In some cases you will also be directed to record the node's original WWNN, or to check that the node has the expected original WWNN. This WWNN is shown at the end of step a. below.

a. On the front panel of the node, press the down button until the "Node:" panel is displayed, and use the right or left navigation button to display the "Node WWNN:" panel. The last five characters of the current WWNN are shown on the second line.

b. Press and hold the down button; press and release the select button; and release the down button. Line one should display "Edit WWNN:" and line two should display the last five characters of the WWNN. The first character of the WWNN is highlighted.

   Note: When changing the WWNN you might receive error 540, "An Ethernet port has failed on the 2145" and/or error 558, "The 2145 cannot see the fibre-channel fabric or the fibre-channel card port speed might be set to a different speed than the fibre-channel fabric". Both messages are expected when the node is booted with no fiber cables connected and no LAN connection. However, if this error occurs while you are editing the WWNN, you have to reenter edit mode by returning to step a.

c. Press the up or down button to increment or decrement the character that is highlighted, change the character until it matches the required value.

   Note: The characters wrap from F to 0 or from 0 to F.

   Press the right navigation button to move to the next field or the left navigation button to return to the previous field and repeat for each character. When the WWNN exactly matches the required value, go to the next step.

d. Press the select button to retain the characters that you have updated and return to the "Node WWNN:" panel.

e. Ensure the WWNN has changed by repeating step a.