IBM® Installation Instruction

Title: Adding PCI adapters to SAN Volume Controller Model DH8 or SV1

Written by:	Gareth Nicholls	28 Jul 2017	SG Development, Hursley
Approved by:	Chen Gu	4 Aug 2017	Global Service Planning, China
Reviewed by:	Alex Ainscow	4 Aug 2017	SG Development, Hursley
Checked by:	Alexander Watson	2 Aug 2017	SG Development, Hursley
Revised by:			

(P) 01YM121A

1

Before installation (Sections 1 through 8)

1.0 Machines affected

These feature codes apply to any SAN Volume Controller model 2145-DH8 or 214x-SV1 (2145-SV1 or 2147-SV1) engine for which the upgrade was ordered.

The SAN Volume Controller 2147-SV1 is physically identical to the 2145-SV1, but includes three-year Enterprise Class Support and warranty.

Feature	Description	Maximum per DH8	Maximum per SV1	FRU part number
AH10	4 port 8Gbps Fibre Channel Card with four shortwave SFPs	3 (4 after 7.6.0.3)	0	31P1702
AH11	2 port 16Gbps Fibre Channel Card with two short-wave SFPs	4	0	00RY004
AH12	4 port 10Gbps Ethernet Card with four SFP+	1	1	00AR319
AH13	4 port 12Gbps SAS Card. Required for SAS expansion enclosure attachment	1	1	64P8485
AH14	4 port 16Gbps Fibre Channel adapter with four short-wave SFPs	4	4	00WY983
AH1A	Compression accelerator card. Requires feature AH1B in a DH8.	2	2	00AR065

2

2.0 Related bills of material (BOMs) or engineering changes (ECs)

2.1 Prerequisite

The IBM Spectrum Virtualize software on the DH8 must be at the version shown in the following table or newer.

214x-SV1 is not supported on versions of the software prior to 7.7.1.0.

Feature	Description	Minimum Software level
AH10	4 port 8Gbps Fibre Channel Card with four shortwave SFPs	7.3.0.0 (7.6.0.3 for 4)
AH11	2 port 16Gbps Fibre Channel Card with two short-wave SFPs	7.4.0.0
AH12	4 port 10Gbps Ethernet Card with four SFP+	7.3.0.0
AH13	4 port 12Gbps SAS Card. Required for Model 24F attachment	7.3.0.0
AH14	4 port 16Gbps Fibre Channel adapter with four short-wave SFPs	7.6.0.0
AH1A	Compression accelerator card. Requires feature AH1B in a DH8.	7.3.0.0

Feature code AH1B Additional microprocessor and 32GB memory, must be fitted in the 2145-DH8 node for adapters cards to function in PCI express expansion slots 4 to 6 provided by PCI express riser card assembly 2 (the one nearest to the power supplies). Feature code AH1A Compression accelerator card will only function in PCI express slot 4 and 6 for DH8. All SV1 have the required microprocessors and memory fitted as standard.

Users must either use the localfcportmask/remotefcportmask or be careful with zoning if using more than 4 fabric ports per SVC node. Fabric ports are Fibre Channel (FC) or Fibre Channel over ethernet (FCoE) ports which are *physically 10* Gbps Ethernet ports.

2.2 Concurrent

The following table identifies which Feature codes or models can be added at the same time as installing a PCI express adapter feature in a San Volume controller 2145-DH8 or 2145/2147-SV1 up to the maximum number shown.

Feature or Model	Description	Max add to DH8	Max add to SV1
2145-12F	3.5-inch SAS drive expansion enclosure	0 to 20	0 to 20
2145-24F	2.5-inch SAS drive expansion enclosure	0 to 20	0 to 20
214x-92F	92 SAS drive expansion enclosure	0 to 8	0 to 8
AH10	4 port 8Gb FC adapter card with four SW SFPs	4	0
AH11	2 port 16Gbps Fibre Channel Card with two short-wave SFP Transceivers	4	0
AH12	4 port 10Gb Ethernet adapter card with four SFPs	1	1

3

AH13	4 port 12Gb SAS adapter card. Connects up to 2 SAS drive expansion enclosures	1	1
AH14	4 port 16Gbps Fibre Channel Card with four short-wave SFP Transceivers	4	4
AH1A	Compression accelerator card. Requires feature AH1B for 2145-DH8	2	2
AH1T	Two 8Gbps Fibre Channel long-wave SFP Transceivers		
ACHU	Two 16Gbps Fibre Channel long-wave SFP Transceivers		

2.3 Companion

The other Feature codes that can be used with the PCI adapter card features are:

Feature or Model	Description	Max per DH8	Max per SV1
214x-92F	92 SAS drive expansion enclosure	0 to 8	0 to 8
214x-12F 984x-12F	3.5-inch SAS drive expansion enclosure	0 to 20	0 to 20
2145-24F	2.5-inch SAS drive expansion enclosure (DH8 only)	0 to 20	0 to 20
AH10	4 port 8Gb FC Card with four short-wave SFPs	4	0
AH11	2 port 16Gbps Fibre Channel Card with two short-wave SFP Transceivers	4	0
AH12	4 port 10Gb Ethernet Card with four SFPs	1	1
AH13	4 port 12Gb SAS adapter. Required for SAS drive expansion attachment	1	1
AH14	4 port 16Gbps Fibre Channel Card with four short-wave SFP Transceivers	4	4
AH1A	Compression accelerator card. Requires feature AH1B for 2145-DH8	2	2
AH1B	Additional microprocessor and 32GB memory	1	In base
AH1T	Two 8Gbps Fibre Channel long-wave SFP Transceivers		
ACHU	Two 16Gbps Fibre Channel long-wave SFP Transceivers		

Each pair of DH8 or SV1 with an AH13 4 port 12Gbps SAS adapter fitted can attach up to twenty 2U SAS drive expansion enclosures (24F or 12F) or up to 8 5U 92 SAS drive expansion enclosures (92F).
The Fibre Channel long-wave SFP transceivers can be used for metro mirror and global mirror links between systems on different sites.

3.0 Bills of Material to be installed

Any number and combination of the following up to the maximum quantity per node shown:

Feature	Description	Maximum per DH8	Maximum per SV1	FRU part number
AH10	4 port 8Gbps Fibre Channel adapter with four Shortwave SFPs	4	0	31P1702
AH11	2 port 16Gbps Fibre Channel adapter with two short-wave SFPs	4	0	00RY004
AH12	4 port 10Gbps Ethernet adapter with four SFPs Used for Fibre Channel over Ethernet (FCoE), iSCSI and IP replication.	1	1	00AR319
AH13	4 port 12Gbps SAS adapter. Required for Model 12F and 24F attachment.	1	1	64P8485
AH14	4 port 16Gbps Fibre Channel adapter with four short-wave SFPs	4	4	00WY983
AH1A	Compression accelerator. Requires feature AH1B for 2145-DH8	2	2	00AR065

4.0 Preparation

__ 1.   Read and understand the purpose of this installation instruction.

__ 2.   Make sure the customer has provided the following plans with reference to the Planning section of the SAN Volume controller Knowledge Center:

A plan of which PCI express expansion slot each of the PCI express adapter cards should be installed in
A plan of which 10 Gbps Ethernet switch ports should be connected to each port of any new 10Gbps Ethernet adapter card
A plan of what IP addresses will be assigned to each port of any new 10Gbps Ethernet Adapter
A plan of which of his Fibre Channel switch ports should be connected to each port of any new Fibre Channel Adapters
A plan of which WWPN each of the FC or FCoE ports will be able to communicate with
A plan for the Storage Area Network Zoning and the local port mask and remote port mask settings

__ 3.   Check that the customer's plans are supported by the software level of the system. Make sure that you have the latest level of this document because future versions of the software might support new hardware configuration options such as more Host Interface Cards, more expansion enclosures. Here are the checks to make. Check the following as appropriate:

Is each of the adapter cards supported in the slot planned? See the following table. Note for DH8 an existing Fibre Channel adapter in PCIe slot 2 can be moved to PCIe slot 1 so that a new 10 Gbps Ethernet adapter can be installed in PCIe slot 2. This will only be necessary if PCIe Slot 5 cannot be used because the additional microprocessor will not be fitted. The World Wide Port Names (WWPN) of the Fibre Channel adapter will change if it is moved from PCIe slot 2 to slot 1. The Storage Area Network zoning might need to be changed and a rescan of the multi pathing on each host might be needed.
If there will be more than 4 fabric ports per node (FC and FCoE) then has the customer planned to avoid more than 16 node fabric logins being received by each node from other nodes?   See: http://www.ibm.com/support/knowledgecenter/STPVGU_7.5.0/com.ibm.storage.svc.console.750.doc/svc_planning_morethanfourfabricports.html
For DH8 if a compression accelerator will be fitted is the additional microprocessor already fitted or will it be fitted at the same time?
If a SAS adapter will be added and if PCIe slot 3 already contains an adapter card then it must be discarded or moved to slot 5 so that the SAS adapter can be fitted in slot 3.

SVC 2145-DH8 adapter slot placement table

PCIe slot	Cards supported	PCIe slot	Cards supported
1	8 Gbps Fibre Channel or 16 Gbps Fibre Channel	4	Compression Accelerator
2	8 Gbps Fibre Channel or 16 Gbps Fibre Channel or 10 Gbps Ethernet	5	8 Gbps Fibre Channel or 16 Gbps Fibre Channel or 10 Gbps Ethernet
3	12 Gbps SAS or 8 Gbps Fibre Channel or 16 Gbps Fibre Channel	6	Compression Accelerator

To fit the SAS adapter in the only slot that it will work in (slot 3) there might already be an adapter in slot 3 which must be moved to slot 5 or discarded.

These are the WWPNs of the ports of an adapter in slot 3 on DH8.

Port 1	500507680c**31**xxxx
Port 2	500507680c**32**xxxx
Port 3	500507680l**33**xxxx
Port 4	500507680l**34**xxxx

These are the WWPNs of the ports of an adapter in slot 5 on DH8.

Port 1	500507680c**51**xxxx
Port 2	500507680c**52**xxxx
Port 3	500507680c**53**xxxx
Port 4	500507680c**54**xxxx

If you move a network adapter to slot 5 and the port cabling with it then the customer might need to change the fabric zoning because of the WWPN changes.

See "2.1 Prerequisite" on page 2, for the minimum software level needed to support each type of adapter.
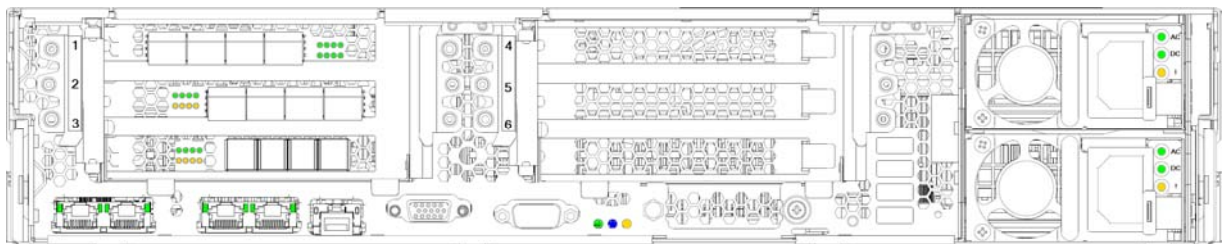
Figure 1. DH8 – Rear view showing the PCI express expansion slots 1 to 6

PN 01YM121A 7 of 38	EC M15426 20170728				

©

Copyright IBM Corp.2017

SVC 214x-SV1 node adapter slot placement table

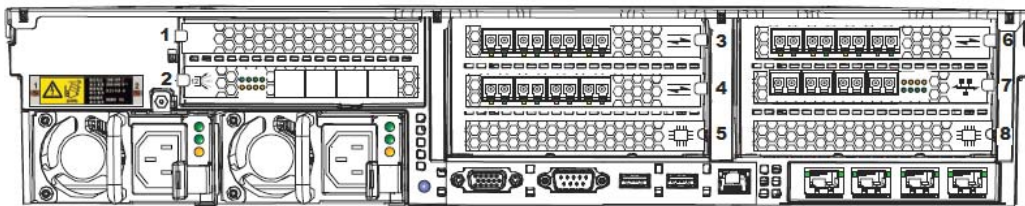| PCIe slot	Cards supported	PCIe slot	Cards supported	PCIe slot	Cards supported
1	Unused	3	16 Gbps Fibre Channel or 10 Gbps Ethernet	6	16 Gbps Fibre Channel or 10 Gbps Ethernet
2	12 Gbps SAS	4	16 Gbps Fibre Channel or 10 Gbps Ethernet	7	16 Gbps Fibre Channel or 10 Gbps Ethernet
		5	Compression Accelerator or 12 Gbps SAS	8	Compression Accelerator

Figure 2. SV1 – Rear view showing the PCI express expansion slots

__ 4. For DH8 additional care is needed if a 4-port adapter is to be replaced by a 2-port adapter in the same PCIe expansion slot, or if a 2-port adapter is to be re-placed by a 4-port adapter. For more information search **ibm.com/support** for the following documents:

Upgrading From a 8Gbps SVC Node to a 16Gbps SVC 2145-DH8 node

Upgrading a 2-port 16Gbps HBA to a 4-port 16Gbps HBA on a SVC 2145-DH8 node with a legacy WWNN

__ 5. **Important**: The customer has ordered upgrades for one or more DH8 or SV1. Install the upgrades only to those machines that are listed by serial number in the order.

__ 6. It is probable that more than one DH8, SV1 in the system is being up-graded. If so, plan the order in which you upgrade the nodes so that the two nodes in an I/O group are not upgraded one immediately after the other. This allows more time for the host systems to recover their connections to the system and reduces the chance of loss of access to volumes. Allow at least 30 minutes between finishing up-grading the first node in an input output group (I/O group) until staring to upgrade the second node in the same I/O group.

__ 7. The customer might also have ordered SAS drive expansion enclosures to attach to a node pair and some SAS drive features.

__ 8. The customer provides access to the management GUI and command line.

PN 01YM121A 8 of 38	EC M15426 20170728				

Note: Cooperate with any customer who might want to supervise the use of the management GUI.

__ 9.   Wear an anti-static wrist strap during any procedure that involves removing or replacing parts inside the node. In addition, use a static-protected mat or work surface, when possible.

5.0 Program Updates

None

6.0 Purpose and Description

6.1 Purpose

The purpose of these installation instructions is to install and test one or more PCI express adapter features as shown in this table.

Feature	Description	Maximum add per DH8	Maximum add per SV1
AH10	4 port 8Gbps Fibre Channel adapter with four short-wave SFPs	4	0
AH11	2 port 16Gbps Fibre Channel adapter with two short-wave SFPs	4	0
AH12	4 port 10Gbps Ethernet adapter with four SFPs	1	1
AH13	4 port 12Gbps SAS adapter. Required for Model 24F attachment	1	1
AH14	4 port 16Gbps Fibre Channel adapter with four short-wave SFPs	4	4
AH1A	Compression accelerator. Requires feature AH1B for 2145-DH8	2	2

6.2 Description

Use these installation instructions to upgrade a DH8 or SV1 by installing one or more PCI express adapter card features.

7.0 Installation time

Concurrent Installation

SSR Hours	Machine (Subsystem) Hours	System Hours	No. of SSRs
1.2 for 1 card	0.0	0.0	1
1.4 for 2 cards	0.0	0.0	1

1.6 for 3 cards	0.0	0.0	1
1.8 for 4 cards	0.0	0.0	1
2.0 for 5 cards	0.0	0.0	1
2.2 for 6 cards	0.0	0.0	1

Subsystem hours means the estimated total outage time that the customer does not have access to the data.

8.0 Special tools and/or materials required

None.

Installation (Sections 9 through 12)

9.0 Safety

Observe all normal safety precautions. Read any safety notices provided with the adapter card feature.

10.0 Details of installation

Use these instructions to upgrade a DH8 or SV1 node by installing PCI express adapter cards.

The upgrade procedure involves several steps:

1. Prepare for the installation.
2. Power off the node. This must be performed without disrupting the operation of the clustered system.
3. Optionally remove the node from the rack.
4. Install the PCI express adapter cards.
5. Replace the node in the rack.
6. Cable up new ports if required.
7. Test the new configuration.
8. Return the node to the customer.

10.1 Preparing for the installation

__ 10.   If the DH8 or SV1 node has not been installed in a rack and powered on yet, go to section 10.6 Installing the PCI adapter cards.

__ 11.   Do not continue with this upgrade on a node that has an existing problem. If the node fault LED is on solid in the front panel, follow the procedures in the *Troubleshooting Guide* to fix the problem before continuing with these instructions.

Figure 3 shows the controls and indicators on the front panel of the DH8.

Figure 4 shows the controls and indicators on the front panel of the SV1.

PN 01YM121A	EC M15426				
11 of 38	20170728				

©
Copyright IBM Corp.2017
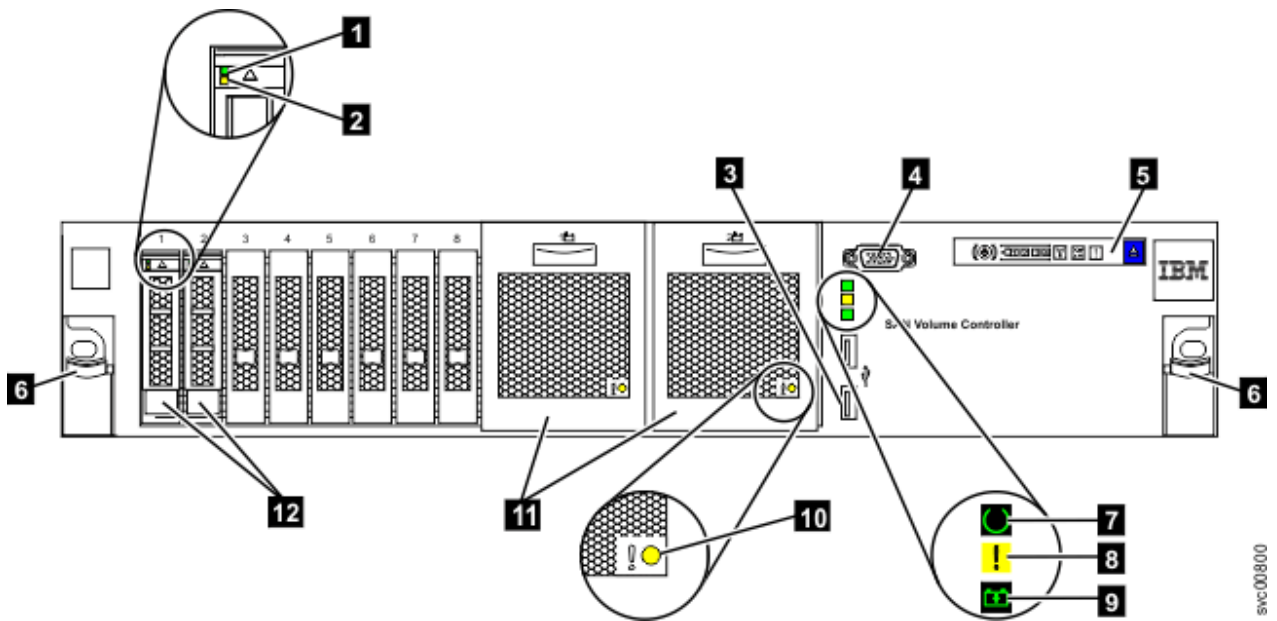
Figure 3. DH8 front panel

[1] Hard disk drive activity LED   [2] Hard disk drive status LED
[3] USB port   [4] Video connector
[5] Operator-information panel   [6] Rack release latch
[7] Node status LED   [8] Node fault LED
[9] Battery status LED   [10] Battery fault LED
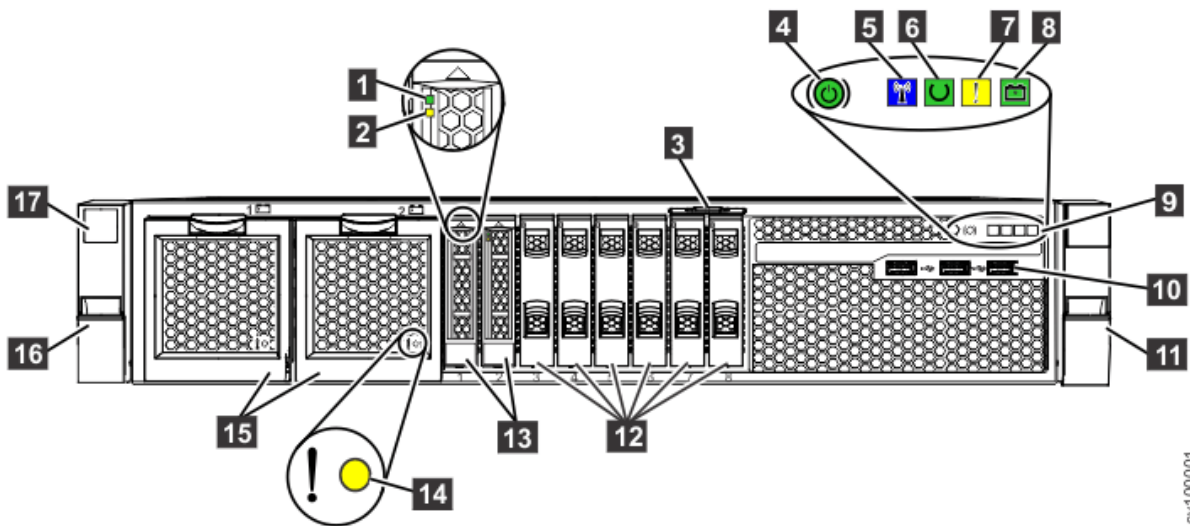[11] Batteries   [12] Hard disk drives (boot drives)

Figure 4. SV1 front panel

[1] Boot drive activity LED
[2] Boot drive status LED

©

12

[3] Pullout tab showing 11S serial number

[4] Power-control button and power-on LED

[5] Identify LED

[6] Node Status LED

[7] Node fault LED

[8] Battery Status LED

[9] Operator-information panel

[10] Front USB ports 1-3

[11] Right side latch (releases chassis to slide out on rails)

[12] Drive slot fillers (no empty slots permitted)

[13] Boot drives

[14] Battery fault LED

[15] Batteries

[16] Left side latch (releases chassis to slide out on rails)

__ 12. If the number of FC ports will change or the PCIe expansion slots that the FC adapters are in will change then it is best to check if the local FC port mask or remote FC port mask for the nodes is anything other than all on. Then set those port masks to all on, before you power off the first node, if they are not already so. This will prevent a loss of node to node communications should FC port numbering change. (Use the CLI commands lssystem, and chsystem with the -localfcportmask or –remotefcportmask command line switch.)

__ 13. If the customer has already powered the node off, determine the state of the node before it was powered off, and decide if any problems must be fixed before continuing with the upgrade.

10.2 Before powering off the node

__ 14. To determine if the node is a member of a system, look at the node status LED on the front panel (**7** in Figure 3 for DH8 and **6** in Figure 4 for SV1). If the **Node status LED** is not on solid then the node is not a member of a system.

If the node is not a member of a system, turn off the node by pressing the power button on the operator panel and immediately releasing it; then continue with section "10.5 Removing the node from the rack".

If the node is a member of an active clustered system, you must ensure that turning off the node that you are upgrading will not disrupt access from the host applications to the volumes.

Note: An I/O group continues to operate and to service all I/O requests when a single node is offline. The system operates in a degraded state, however, with reduced performance and reduced resilience to failures. Do not continue this procedure if any nodes in the system are offline, if any node errors are displayed, or node fault LEDs are on or if unfixed errors are in the system event log. Continuing the installation in such instances might result in data loss or loss of service.

PN 01YM121A	EC M15426				
13 of 38	20170728				

Important: Do not turn off a node if any of the volumes are not online. Turning off a node when its volumes are not online might lead to loss of access to those volumes. One reason that volumes might not be online is because the partner node in the I/O group was recently turned off and this node has not yet had time to write all of its cached data. Therefore, following this procedure to verify that the volumes are online is important.

Work with the customer to check the volume status using the management GUI.

__ 15. **Important:** Look at the front panel of the node, and confirm that the node fault LED is not lit. If the node fault LED is on then follow the procedures in the *Troubleshooting Guide* to fix the error. This is important as the error might mean that the system is not able to maintain access to the backend data when the node is powered off.

The customer should open the management GUI for the system in a browser and sign on with their user name and password.

__ 16. Click **Monitoring > System** in the management GUI. A picture that shows the nodes in the system and which I/O group they are in is displayed. Confirm that the status of every node is Online. If any node is not online, follow the procedures in the *Troubleshooting Guide* to bring each node online.

__ 17. Click **Monitoring > System.** Select the node that you intend to shutdown and write down the node name and the name of the I/O group to which the node belongs. Then select **Actions** and click **Show Dependent Volumes**.

__ 18. There should not be any volumes listed as dependent on this node but if there are then investigate the dependencies to resolve any problems before continuing with this installation.

__ 19. If possible, have the customer confirm that the multi-pathing device driver on each of the hosts accessing volumes managed by this I/O group is able to failover to the partner node in the I/O group when this node is turned off. This check must be performed on each of the connected hosts; the check performed varies depending on the multi-pathing device driver being used on the host.

If it is not feasible to check the status of each host, check if the partner node in the I/O group of the node being upgraded was last started less than 30 minutes ago. If so, wait until at least 30 minutes after the partner node was restarted to continue. Waiting gives the host system multi-pathing device drivers time to re-establish connections through the partner node. If the drivers have not had time to re-establish connections, hosts lose access to the backend data.

10.3 Powering off a node which has software level 8.1.0.0 or after

__ 20.   Select the node in the management GUI and choose the action to service the node. Or use this command on the system cli:

swapnode -service <node_id>

__ 21.   It takes up to a minute for the node to go to service status. If there is a suitable spare node then this will take over from the service status node and change to online_spare status.

__ 22.   **Power off the Node** using the service assistant GUI or by pressing the power button and releasing it immediately.

__ 23.   Wait until the node power-on LED is either flashing or off. See **1** in Figure 5 for DH8 and **1** in Figure 6 for SV1.

10.4 Powering off a node which has software level older than 8.1.0.0

__ 24.   Use the service assistant GUI to put the node into service status or use this command on the service cli:

startservice <panelname>

__ 25.   It takes up to a minute for the node to go to service status. The node status LED will change from on to blinking.

__ 26.   **Power off the Node** using the service assistant GUI or by pressing the power button and releasing it immediately.

__ 27.   Wait until the node power-on LED is either flashing or off. See **1** in Figure 5 for DH8 and **1** in Figure 6 for SV1.
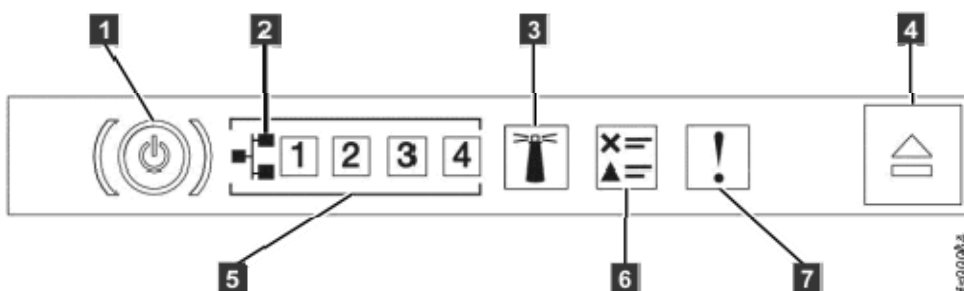
Figure 5. Power LED location on DH8 operator information panel

PN 01YM121A	EC M15426				
15 of 38	20170728				

©

Copyright IBM Corp.2017
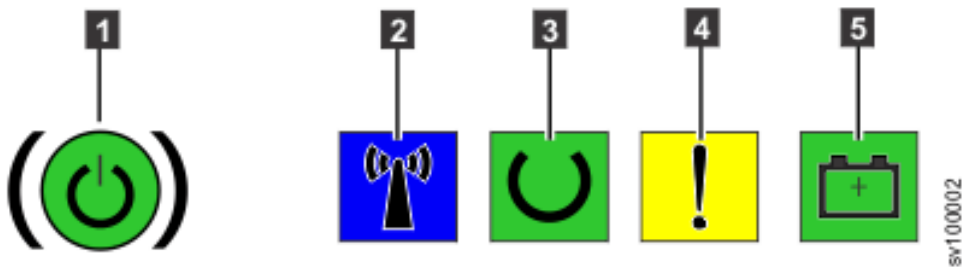
Figure 6. Power LED location on SV1 operator information panel

10.5 Removing the node from the rack

You can perform this procedure without removing the node from the rack, as long as the cable management arm, or similar, has been used and there is sufficient access to work on the node when it is extended from the rack.

To perform this action on the node without removing it from the rack, carefully extend the node out of the rack on its rails, disconnect the power cable, and go to section 10.4 installing the PCI adapter cards.

__ 28.   If you are going to remove the node from the rack then carefully label the port number position of each Fibre Channel cable, Ethernet cable and SAS cable; then disconnect the Fibre Channel cables, Ethernet cables and SAS cables from the rear panel. See Figure 7. DH8 rear view of connectors Figure 7. DH8 rear view of connectors or Figure 8. SV1 rear view of connectors.

Figure 7. DH8 rear view of connectors

16
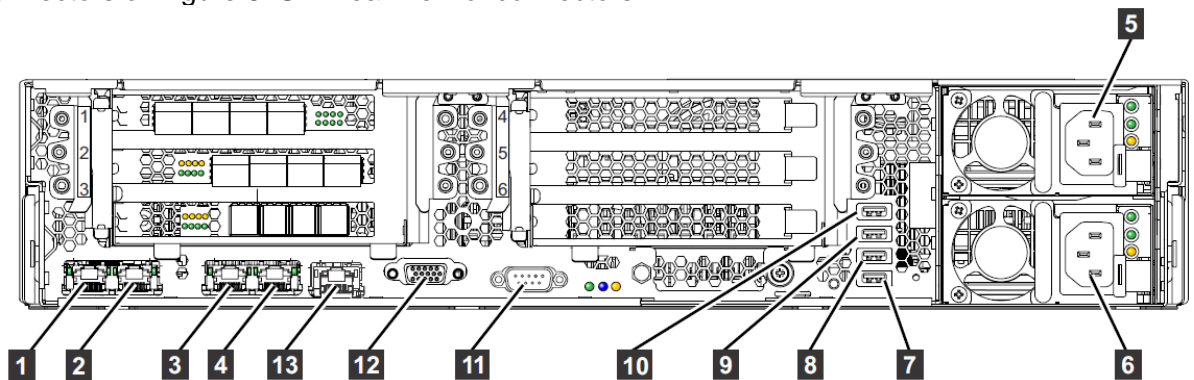
1 1 Gbps Ethernet port 1
2 1 Gbps Ethernet port 2
3 1 Gbps Ethernet port 3
4 Technician port (Ethernet)
5 Power supply 2
6 Power supply 1
7 USB 6
8 USB 5
9 USB 4
10 USB 3
11 Serial
12 Video
13 Unused Ethernet port
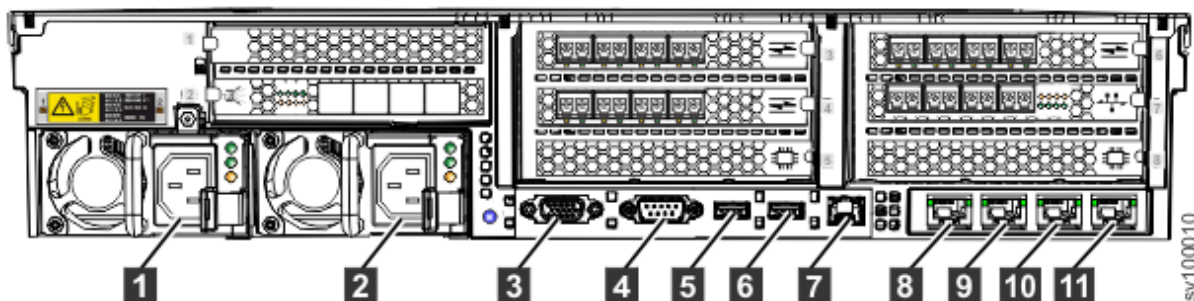
Figure 8. SV1 rear view of connectors

1 Power supply 1
2 Power supply 2
3 Video port
4 Serial port
5 Rear USB port 1
6 Rear USB port 2
7 Unused Ethernet port
8 10 Gbps Ethernet port 1
9 10 Gbps Ethernet port 2
10 10 Gbps Ethernet port 3
11 Technician port (Ethernet)

   Note: There could be 8 Gbps Fibre Channel, 16 Gbps Fibre Channel, 10 Gbps Ethernet and 12 Gbps SAS adapter cards in the PCI express expansion slots with cables connected to the ports on them. You must ensure that the Fibre Channel cables and the Ethernet cables and any SAS cables that you unplugged from the rear of the node are plugged back into the same port at the end of the procedure. Otherwise, unexpected problems might occur.

__ 29.   Disconnect the power cords from the rear panel. See Figure 7. DH8 rear view of connectors or Figure 8. SV1 rear view of connectors.

Note: Be careful not to dislodge any cables or power cords from any other DH8 or SV1 nodes in the clustered system during this process or when sliding the node out of the rack.

__ 30.   Depress the blue tabs on either side of the DH8 or SV1 and carefully pull the node forward until the rails click into a locked position.

__ 31.   Optionally, remove the DH8 or SV1 from its rails.

Position yourself, or a mechanical lift, to support the weight of the node when it is no longer held by the slides; then push forward the locking levers (Figure 9. Removing the DH8 from the rack or Figure 10. Removing the SV1 from the rack) on both sides of the node.
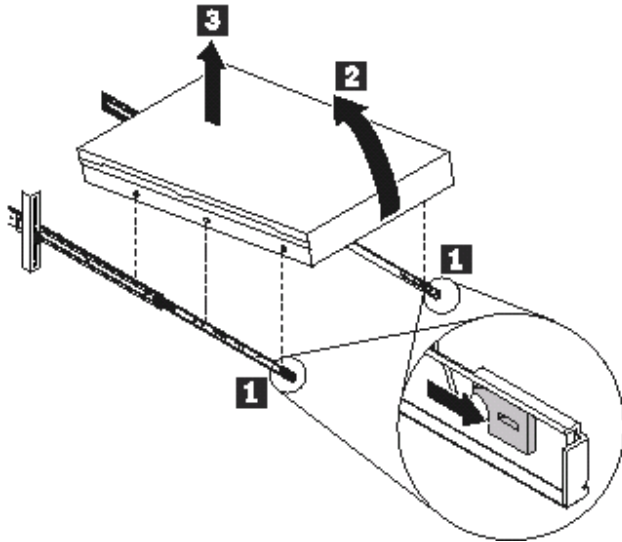
Figure 9. Removing the DH8 from the rack

PN 01YM121A	EC M15426				
18 of 38	20170728				
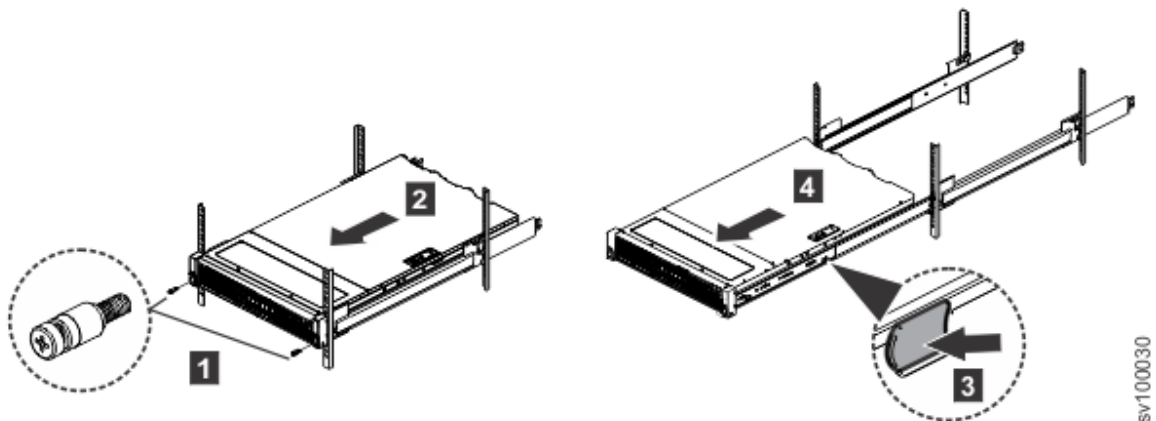
Figure 10. Removing the SV1 from the rack

Lift up the front of the node, and then pull the node out of the rack.

10.6 Installing the PCI adapter cards

__ 32.   If you did not remove the DH8 or SV1 node from the rack then you must now unplug both power cables from the power supply units at the rear of the DH8 or SV1 and unplug any signal cables that are connected to any existing PCI Express adapter cards in the PCI express riser card assemblies that you need to install new adapter cards in.

Note: There could be 8 Gbps Fibre Channel, 16 Gbps Fibre Channel, 10 Gbps Ethernet and 12 Gbps SAS adapter cards in the PCI express expansion slots with cables connected to the ports on them. You must ensure that the Fibre Channel cables and the Ethernet cables and any SAS cables that you unplugged from the rear of the node are plugged back into the same port at the end of the procedure. Otherwise, unexpected problems might occur.

__ 33.   While wearing your ESD wrist strap, press down firmly on the blue tab on the top (near the front of the DH8) or release the black latch (near the front of the SV1) of the cover, lift the latch and slide the cover toward the rear of the DH8 or SV1 until the cover has disengaged from the chassis. See Figure 11. Removing the cover from the DH8 or Figure 12. Removing the cover from the SV1.

PN 01YM121A	EC M15426				
19 of 38	20170728				
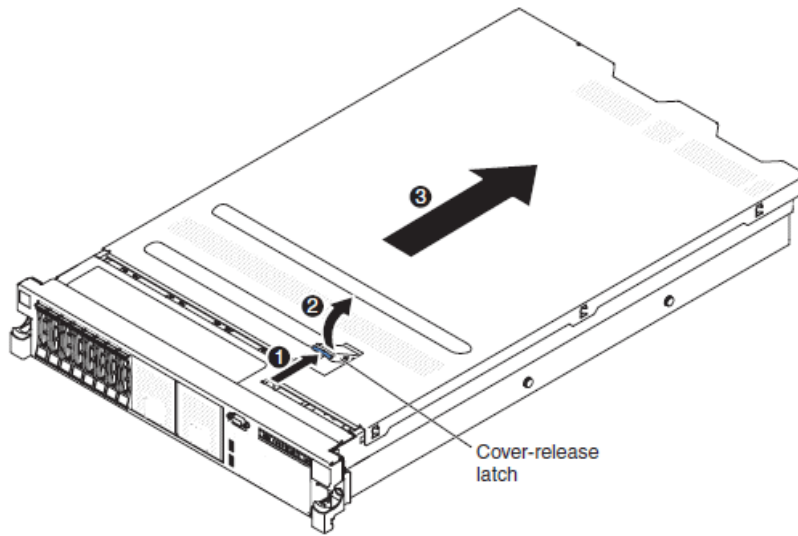
Figure 11. Removing the cover from the DH8

Figure 12. Removing the cover from the SV1

__ 34.   Lift the cover off the DH8 or SV1 and put it in a safe place. Note that for DH8 there is a servicing information label on the inside of this top cover that you may find useful for quick reference.

__ 35.   Remove one or more of the riser card assemblies as needed.
1. Grasp the riser-card assembly at the front tab and rear edge and lift it to remove it from the slot.
2. Place the riser-card assembly on a flat, static-protective surface.
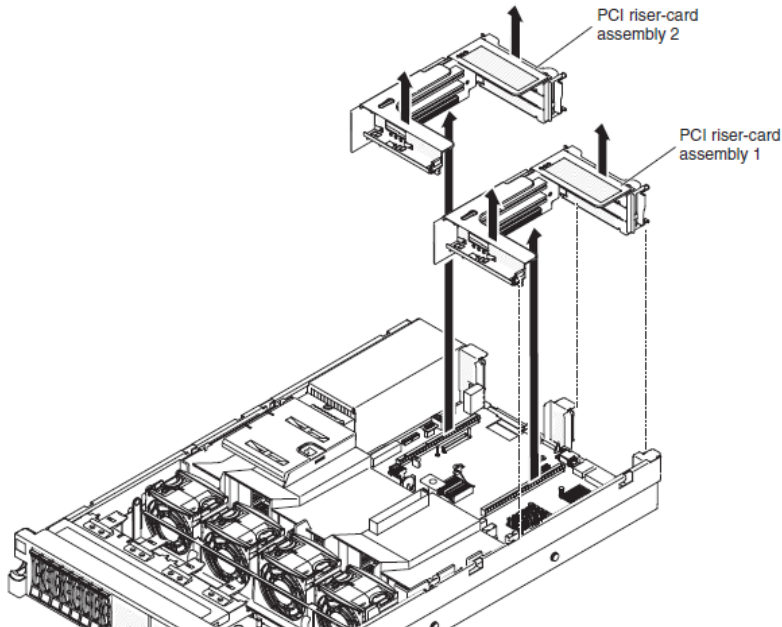
Figure 13. Removing the PCI riser-card assemblies from the DH8
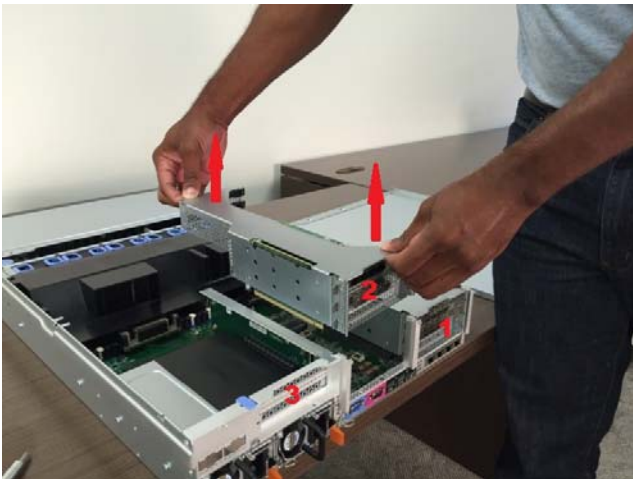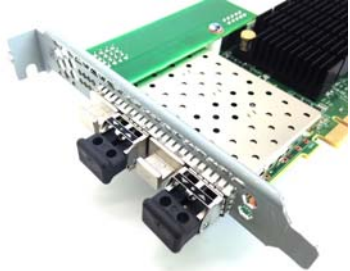
Figure 14. Removing the riser-card assemblies from the SV1

__ 36.   For DH8 if you need to move an existing Fibre Channel adapter from PCIe slot 2 to PCIe slot 1 so that a new 10 Gbps Ethernet adapter can be installed in PCIe slot 2, then move it now. This will only be necessary if PCIe Slot 5 cannot be used because the additional microprocessor is not fitted. Note that the World Wide Port Names (WWPN) of the Fibre Channel adapter will change if it is moved from PCIe slot 2 to slot 1. See *Table 1. Port numbers and corresponding WWPNs*, at the end of this document.

PN 01YM121A 21 of 38	EC M15426 20170728				

©

Copyright IBM Corp.2017

__ 37.　　　For DH8 if you need to move an existing fabric adapter from PCIe slot 3 so that a 12 Gbps SAS adapter can be fitted in PCIe slot 3, then move it now. Note that the World Wide Port Names (WWPN) of the fabric adapter will change if it is moved from PCIe slot 3 to another slot (e.g. PCIe slot 5)

__ 38.　　　Locate the new parts to be installed.

31P1702	8 Gbps Fibre Channel adapter With 4 short-wave SFP	 Figure 15. 8 Gbps Fibre Channel adapter card
00RY004	2 port 16Gbps Fibre Channel adapter with two short-wave SFPs	 Figure 16. 2 port 16 Gbps Fibre Channel adapter
00WY983	4 port 16Gbps Fibre Channel adapter with four short-wave SFPs	 Figure 17. 4 port 16 Gbps Fibre Channel adapter
64P8485 or 01AC573	SAS adapter	 Figure 18. 12 Gbps SAS adapter card

22

00AR319	10 Gbps Ethernet adapter With 4 SFP	 Figure 19. 10 Gbps Ethernet adapter card
00AR065 or 01EJ490	Compression Accelerator	 Figure 20. Compression accelerator card

__ 39.   Put each new PCI adapter card in the PCI express slot as required by the customer. For DH8, PCI expansion slots 1, 2, and 3 are in PCI riser-card assembly 1, and slots 4, 5, and 6 are in PCI riser-card assembly 2. For SV1, PCI expansion slots 1 and 2 are in PCI riser-card-assembly 3, 4 and 5 are in PCI riser-card assembly 2 and PCI expansion slots 6, 7 and 8 are in PCI riser-card assembly 3.

a) It is best to fit the adapter without SFP modules fitted.

b) Align the adapter with the PCI connector on the riser card and the guide on the external end of the riser-card assembly.

c) Press the adapter firmly into the PCI connector on the riser card.

d) Fit the SFP modules if need.

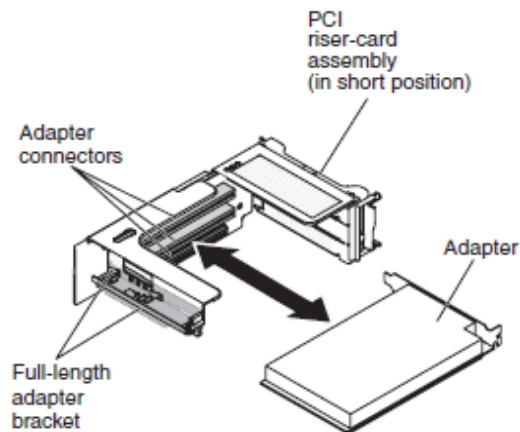| PN 01YM121A 23 of 38	EC M15426 20170728				

23

Figure 21. Installing a card in a PCI riser-card assembly

__ 40.   Align the PCI riser-card assembly with the selected PCI connector on the system board:

__ 41.   Carefully fit the two alignment slots on the side of the assembly onto the two alignment brackets on the side of the chassis; align the rear of the assembly with the guides on the rear of the DH8 or SV1.

__ 42.   Press down on the assembly. Make sure that the riser-card assembly is fully seated in the riser-card connector on the system board.
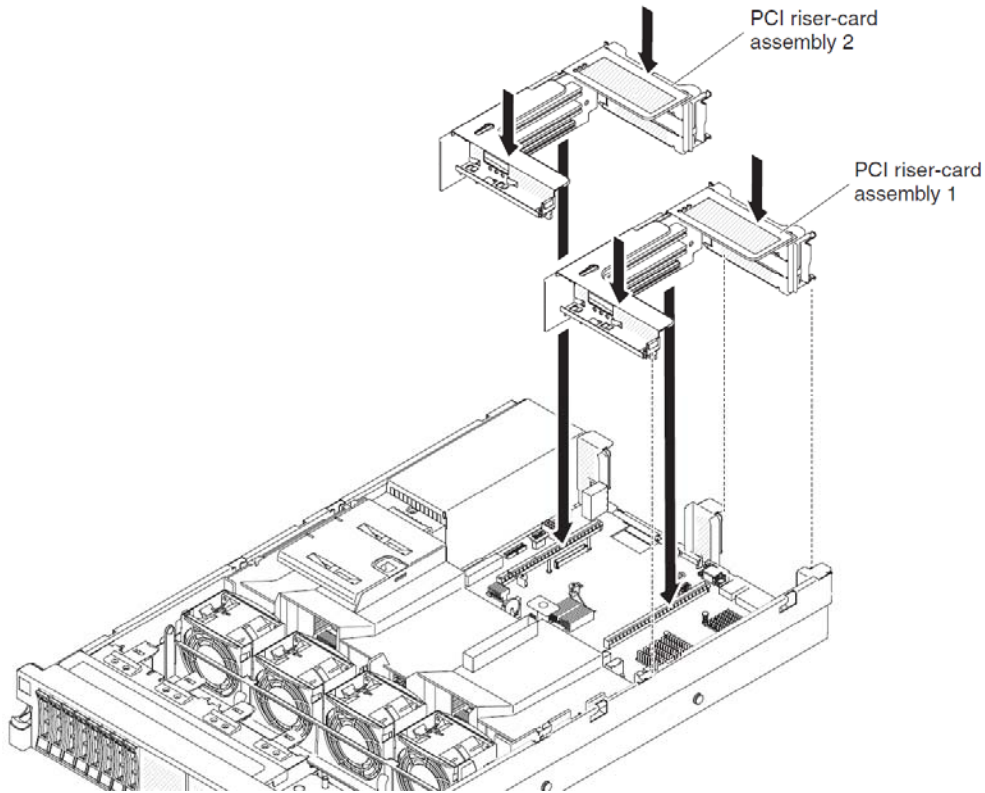
Figure 22. Replacing the DH8 PCI riser-card assemblies

Figure 23. Replacing the SV1 PCI riser-card assemblies

__ 43.   Position the cover on top of the DH8 or SV1 and slide it forward. Make
sure that the cover correctly engages all the inset tabs on the DH8 or SV1. See Fig-

PN 01YM121A	EC M15426				
25 of 38	20170728				

25
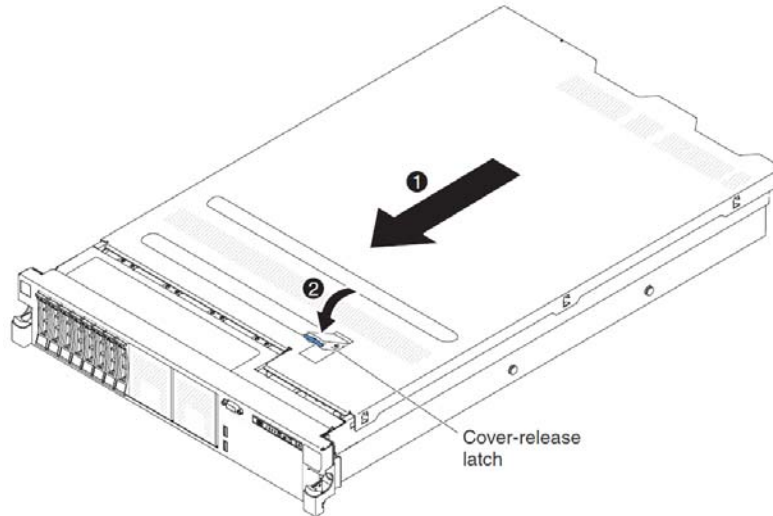
Figure 24. Replacing the DH8 cover

Figure 25. Replacing the SV1 cover

10.7 Replacing the node in the rack

__ 44.   Install the DH8 or SV1 onto the rails if you removed it from the rails to perform this procedure.

__ 45.   Lift the blue tabs on the sides of both rails and slide the node fully into the rack.

__ 46.   Reconnect the cables to the DH8 or SV1 node, ensuring that each data cable is reconnected to the same port from which it was removed. See Figure 7. DH8 rear view of connectors or Figure 8. SV1 rear view of connectors.

PN 01YM121A	EC M15426				
26 of 38	20170728				

10.7 Power on the node

__ 47.   Turn on the DH8 or SV1 node by pressing and releasing the power button once. See **1** in Figure 5. Power LED location on DH8 operator information panel or Figure 6. Power LED location on SV1 operator information panel.

__ 48.   Allow a few minutes for the node to come online. If the node status LED comes on solid then jump to 10.8 Testing the new adapters.

__ 49.   If node status LED is not on solid and the node error LED is on then access the service assistant GUI via the technician port of that node to see the node error code.

__ **50.**   If the node error is 550 which means **A clustered system cannot be formed because of a lack of clustered system resources** then refer to Recovering from node error 550 when the software level is 8.1.0.0 or after or Recovering from node error 550 when the software level is older than 8.1.0.0 in the Appendix.

__ **51.**   Look up any other node error in the knowledge center to see what you need to do.

10.8 Testing the new adapters

__ 52.   Use the management GUI to complete the next 3 steps.

__ 53.   Confirm that the new PCI adapter cards have been detected by selecting **Monitoring > System** in the management GUI, clicking on the node that has just been added, selecting the **Hardware** tab, then scrolling down to the **Adapters** section. Each adapter should be listed and the value **valid** should be shown as **yes**.

__ 54.   For SV1 only: Note that if the default port mapping values for the SV1 have ever been changed, and a card is inserted into a slot that is not currently mapped you will get a node error 842. Follow the instructions in the Knowledge Center to map the new ports to the desired WWPNs using the satask chvpd -fcportmap command on the CLI.

__ 55.   There might be a `1199` error code in the event log which shows that the hardware configuration has changed and is valid. If so, then right mouse click on the 1199 error code and select `Run Fix Procedure`. This will allow you to verify that the configuration change is as expected so that the hardware can be initialized.

__ *56.*   If this is the second node to be upgraded in an IO group, and there is a new SAS drive expansion enclosure to be connected to it, then connect the expansion enclosure now. See *Installing an optional SAS expansion enclosure* in the Information Center or the *Hardware Installation Guide.*

10.9 Return the DH8 or SV1 to the customer

__ 57.   Return the DH8 or SV1 to the customer. Explain to the customer that the adapters (and optionally drives) are installed, but they need to configure any drives into arrays and storage pools.

For DH8, the World Wide Port Names (WWPN) of the Fibre Channel adapter will have changed if it was moved from one PCIe slot to another. See *Table 1. Port numbers and corresponding WWPNs*, at the end of this document. The Storage Area Network zoning might need to be changed and a rescan of the multi pathing on each host might be needed.

The local FC port mask and partner FC port masks might need to be set after all nodes in the system have been serviced, if FC port numbering has changed. Refer them to "**Planning for more than four fabric ports**" in the Knowledge Center.

PN 01YM121A	EC M15426				
28 of 38	20170728				

After installation (Sections 11 through 12)

11.0 Parts disposition

__ 58.   All removed parts remain the property of the customer. Give all removed parts to the customer for disposition, including the riser card and the blanking plate that were removed prior to the installation of the adapters.

Important: Ask the customer to keep the old riser card safe as it must be reinstalled if the SAS adapter assembly is removed without another adapter assembly replacing it. Otherwise the node may not initialize correctly after booting up.

12.0 Machine records

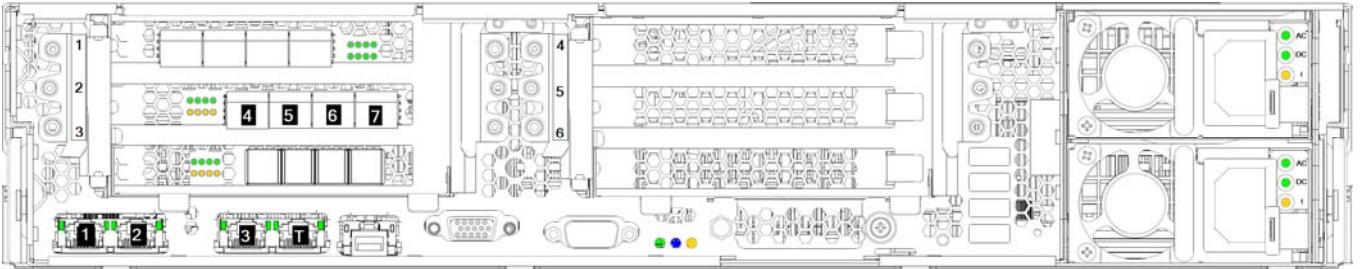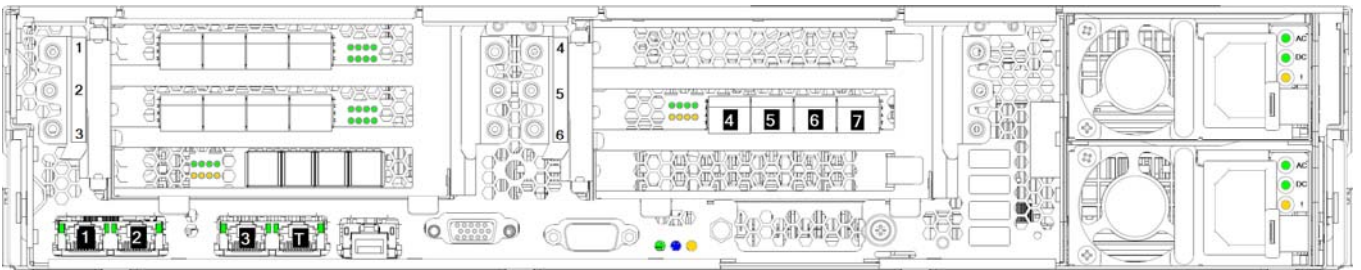__ 59.   Report any installation or quality problems according to existing procedures.

Appendix

DH8 Ethernet port numbers

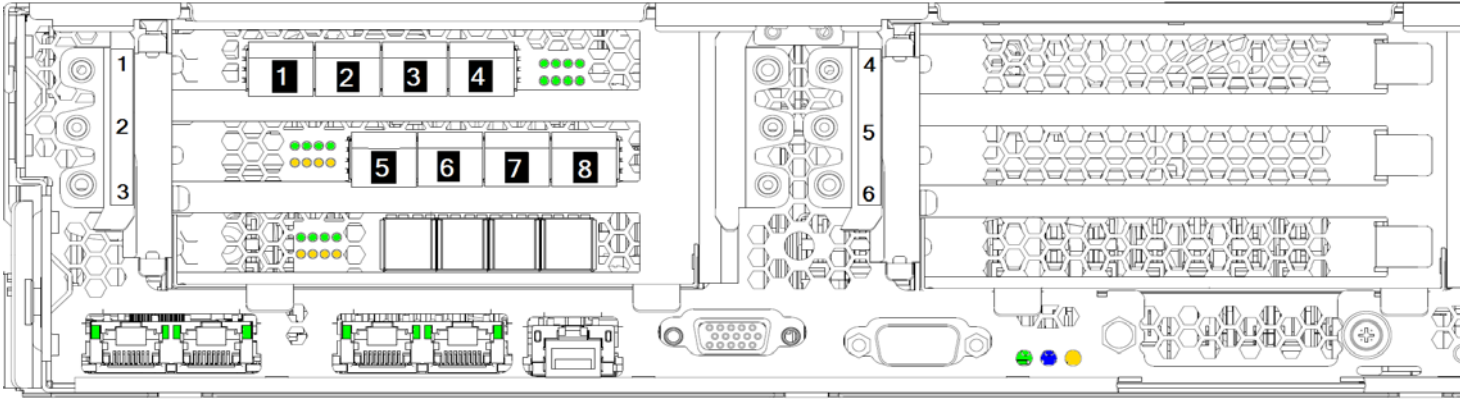When the 10 Gbps Ethernet adapter card is in PCI express expansion slot 2:

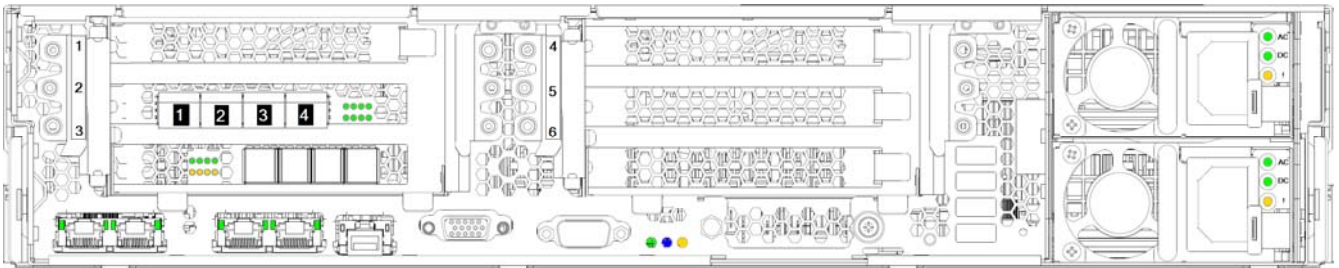When the 10 Gbps Ethernet adapter card is in PCI express expansion slot 5:
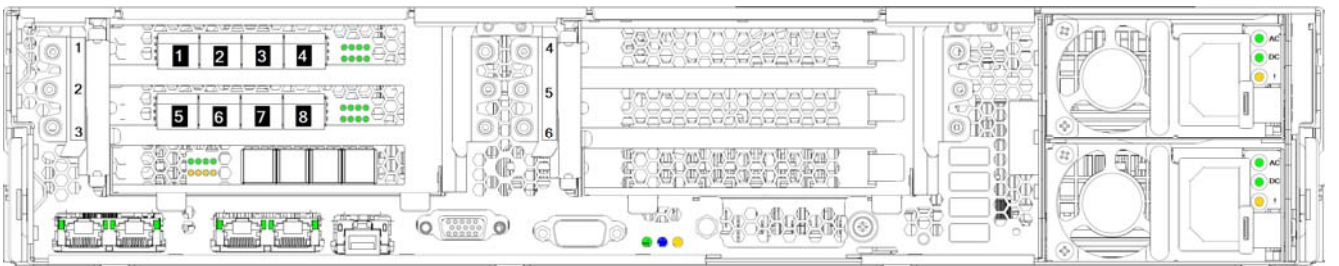
DH8 Fibre Channel port numbers

When a 8 Gbps Fibre Channel adapter card is in PCI express expansion slot 1 and a 10 Gbps Ethernet adapter in slot 2 which can be used for Fibre Channel over Ethernet (FCOE):

When the 8 Gbps Fibre Channel adapter card is in PCI express expansion slot 2:
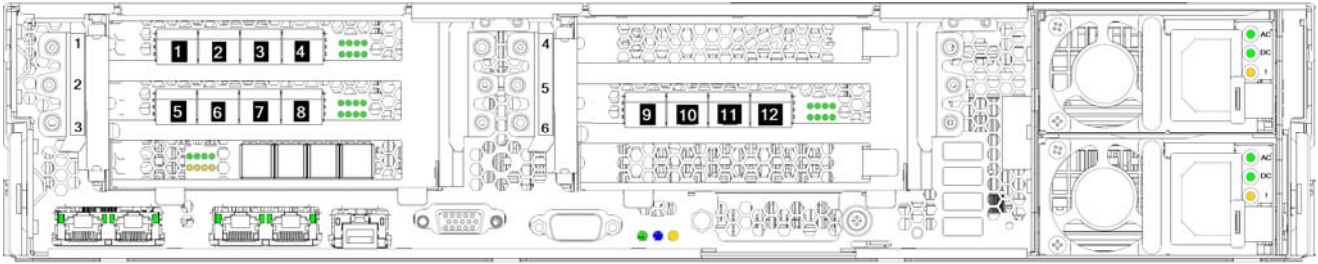
When the 8 Gbps Fibre Channel adapter cards are in PCI express expansion slots 1 and 2:
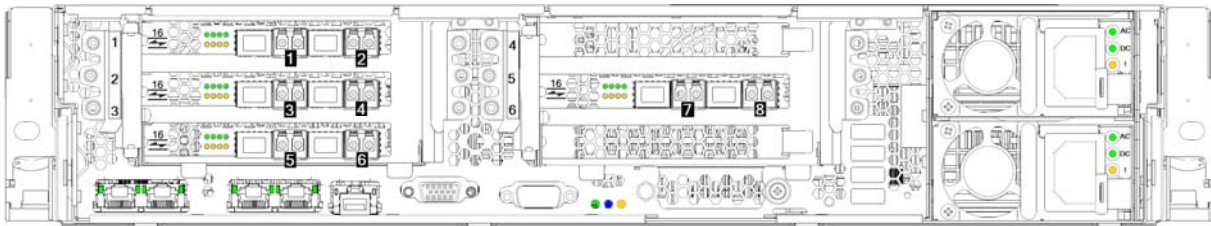
When the 8 Gbps Fibre Channel adapter cards are in PCI express expansion slots 1,2 and 5:
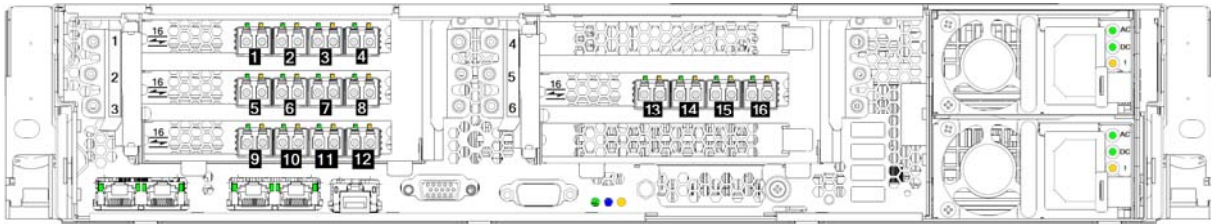
The 2 port,16 Gbps Fibre Channel adapter cards are in PCI express expansion slots 1,2,3 and 5:

The 4 port,16 Gbps Fibre Channel adapter cards are in PCI express expansion slots 1,2,3 and 5:

DH8 World Wide Port Names (WWPN)

When the World Wide Node Name (WWNN) has not been changed since it was set at IBM manufacturing:

WWPNs are in the 500507680c**SPXXXX** format, where

- **S** is the hexadecimal slot number,

- **P** is the hexadecimal port number, and

- **XXXX** is the sequence number of the SAN Volume Controller 2145-DH8 node assigned during manufacture.

As slots 1, 2, 3 and 5 are the only slots in a DH8 that can have an I/O card installed in them, and each I/O card has 4 ports, 16 WWPNs are possible. Table 1 lists these WWPNs.

Table 1. Port numbers and corresponding WWPNs

PCI slot number	Adapter port number	WWPN
1	1	500507680c**11**xxxx

Table 1. Port numbers and corresponding WWPNs

PCI slot number	Adapter port number	WWPN
1	2	500507680c**12**xxxx
1	3	500507680c**13**xxxx
1	4	500507680c**14**xxxx
2	1	500507680c**21**xxxx
2	2	500507680c**22**xxxx
2	3	500507680c**23**xxxx
2	4	500507680c**24**xxxx
3	1	500507680c**31**xxxx
3	2	500507680c**32**xxxx
3	3	500507680**133**xxxx
3	4	500507680**134**xxxx
5	1	500507680c**51**xxxx
5	2	500507680c**52**xxxx
5	3	500507680c**53**xxxx
5	4	500507680c**54**xxxx

When upgrading to a DH8 node from an older node model and when the WWNN of the DH8 is set to the WWNN of the replaced node then the legacy WWPN will be used so that there is no need for cabling, zoning or host configuration changes.

The I/O adapter in the new node must be the same as those in the older node. Because a 4-port Fibre Channel I/O adapter is installed in slot 1 of all nodes in previous versions of SAN Volume Controller, the same 4-port adapter must also be installed in the first slot of the SAN Volume Controller 2145-DH8 node. If a second I/O card is installed in the legacy node, the same adapter must be installed in slot 2 of the SAN Volume Controller 2145-DH8 node.

Table 2. Port numbers and corresponding WWPNs with a legacy WWNN (801)

PCI slot number	Adapter port number	WWPN
1	1	500507680**140**xxxx
1	2	500507680**13**0xxxx
1	3	500507680**11**0xxxx
1	4	500507680**12**0xxxx
2	1	500507680**15**0xxxx
2	2	500507680**16**0xxxx
2	3	500507680**17**0xxxx
2	4	500507680**18**0xxxx

Table 2. Port numbers and corresponding WWPNs with a legacy WWNN (801)

| PCI slot number	Adapter port number	WWPN
3	1	500507680190xxxx
3	2	5005076801a0xxxx
3	3	5005076801b0xxxx
3	4	5005076801c0xxxx
5	1	500507680158xxxx
5	2	500507680168xxxx
5	3	500507680178xxxx
5	4	500507680188xxxx

Important: Avoid using a 2 port adapter in PCIe expansion slot 1 because that will cause the third and fourth Fibre Channel ports to have an unexpected WWPN as shown in table 3.

Table 3. Port numbers and corresponding WWPNs with a legacy WWNN (801) when the adapter in slot 1 has only 2 ports

PCI slot number	Adapter port number	WWPN
1	1	500507680140xxxx
1	2	500507680130xxxx
2	1	500507680180xxxx
2	2	500507680170xxxx
2	3	500507680170xxxx
2	4	500507680180xxxx

SV1 Ethernet port numbers

When the 10 Gbps Ethernet adapter card is in PCI express expansion slot 7 (example only – see 4.0 Preparation for details of supported slots in SV1):
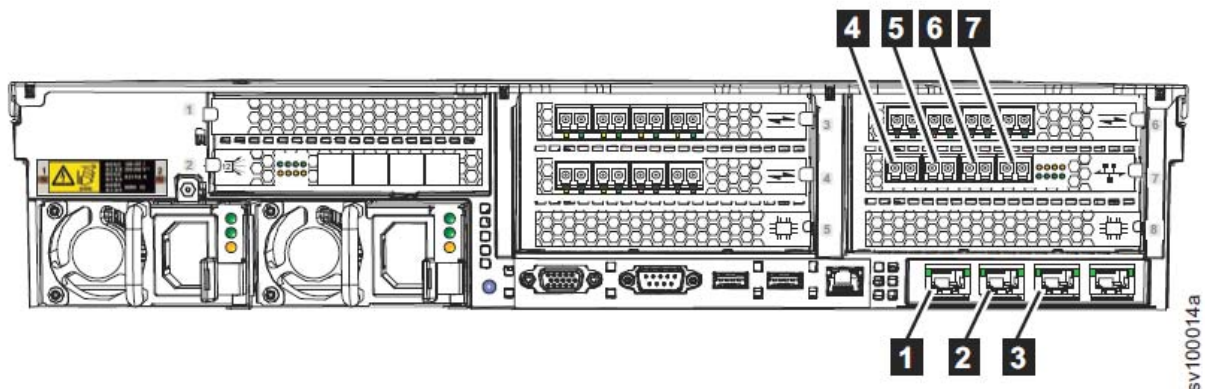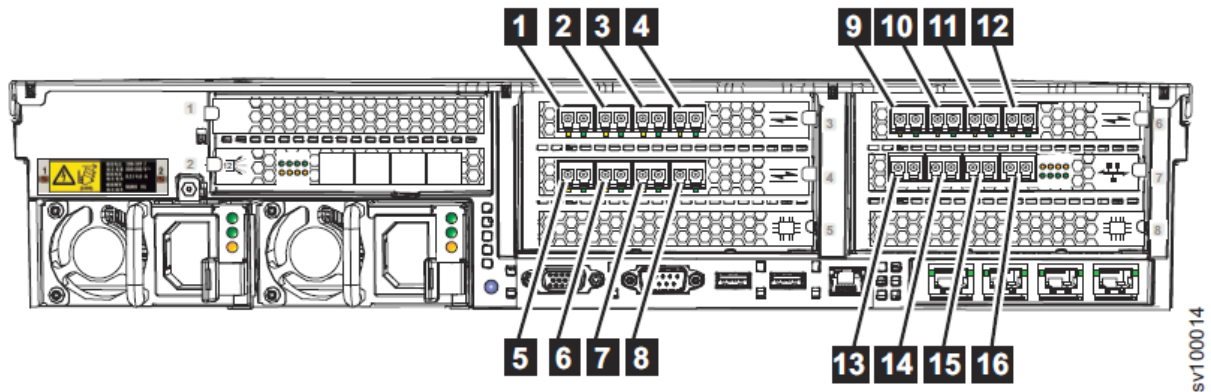
SV1 Fibre Channel port numbers

When the 4 port,16 Gbps Fibre Channel adapter cards are in PCI express expansion slots 3,4 and 6. A 10Gbps Ethernet adapter is in PCI express expansion slot 7 (example only, see 4.0 Preparation for details of supported slots in SV1):

SV1 WWPNs

When the World Wide Node Name (WWNN) has not been changed since it was set at IBM manufacturing:

The WWNN has the following format, where xxxx is a unique number that was assigned when the node was manufactured:

500507680f00xxxx.

The worldwide port names (WWPNs) are derived from the WWNN of the node in which the ports are installed. The WWPNs are also derived from Fibre Channel IO port ID of the port. Each Fibre Channel IO port ID can be assigned to any Fibre Channel or Fibre Channel over Ethernet port using the satask chvpd -fcportmap command. The default assignment is shown in the table below. Note that the WWPN for a particular FC IO port ID is fixed, only the physical location of that FC IO port ID can be changed.

			Default Mapping	
FC IO port ID	Physical WWPN	NPIV WWPN	PCI slot	Adapter port ID
1	500507680f01xxxx	500507680f81xxxx	3	1
2	500507680f02xxxx	500507680f82xxxx	3	2
3	500507680f03xxxx	500507680f83xxxx	3	3
4	500507680f04xxxx	500507680f84xxxx	3	4
5	500507680f05xxxx	500507680f85xxxx	4	1
6	500507680f06xxxx	500507680f86xxxx	4	2
7	500507680f07xxxx	500507680f87xxxx	4	3

8	500507680f**08**xxxx	500507680f**88**xxxx	4	4
9	500507680f**09**xxxx	500507680f**89**xxxx	6	1
10	500507680f**0a**xxxx	500507680f**8a**xxxx	6	2
11	500507680f**0b**xxxx	500507680f**8b**xxxx	6	3
12	500507680f**0c**xxxx	500507680f**8c**xxxx	6	4
13	500507680f**0d**xxxx	500507680f**8d**xxxx	7	1
14	500507680f**0e**xxxx	500507680f**8e**xxxx	7	2
15	500507680f**0f**xxxx	500507680f**8f**xxxx	7	3
16	500507680f**10**xxxx	500507680f**90**xxxx	7	4

Note that if the default port mapping values for the SV1 have ever been changed, and a card is inserted into a slot that is not currently mapped you will get a node error 842. Follow the instructions in the Knowledge Center to map the new ports to the desired WWPNs using the satask chvpd -fcportmap command on the CLI.

PN 01YM121A 36 of 38	EC M15426 20170728				

Recovering from node error 550 when the software level is 8.1.0.0 or after

If an updated node has node error 550 shown in the Service Assistant GUI via the technician port, and the fiber channel cabling is correct then this might be because the adaptor that was being used to communicate with the other nodes has just been replaced with a different type of adapter. e.g. when an AH11 was replaced by an AH14.

- Use the service assistant GUI to put this node back to candidate status

    o by selecting **Remove System Data** in the **Manage System** page.

    o or use the following command in the service CLI of the node

        ▪ satask leavecluster -force <panelname>

- check that the node status is candidate.

- Make the candidate node take over from the offline node:

    o Select the offline node in the management GUI and chose the action to replace this node.

    o Or use this command on the system CLI

        1. swapnode -replace <offline node_id/name>

    o The candidate status node should change to online status after five minutes or more. If it does not then the old procedure can be used with a new

    -deactivatespare switch if a spare node, which has online_spare status, has taken over from the offline node.

        1. rmnode [-deactivatespare] <offline node id>

        2. addnode -panelname <panelname of candidate> -iogrp X [-site n]

Recovering from node error 550 when the software level is older than 8.1.0.0

If an updated node has node error 550 shown in the Service Assistant GUI via the technician port, and the fiber channel cabling is correct then this might be because the adaptor that was being used to communicate with the other nodes has just been replaced with a different type of adapter. e.g. when an AH11 was replaced by an AH14.

- use the service assistant GUI to put this node back to candidate status

  - by selecting **Remove System Data** in the **Manage System** page.

- or use the following command in the service CLI of the node:

  - satask leavecluster -force \<panelname>

- check that the node has candidate status.

- Make the candidate node take over from the offline node:

  1. rmnode \<offline node id>

  2. addnode -panelname \<panelname of candidate> -iogrp X [-site n]