FRU Replacement Guide

HGST Active Archive System SA-7000

September 2015

1ET0036

Revision 1.1

Long Live Data ™ | www.hgst.com

Copyright

HGST, a Western Digital company
3403 Yerba Buena Road
San Jose, CA 95135
Produced in the United States

One MB is equal to one million bytes, one GB is equal to one billion bytes, one TB equals 1,000GB (one trillion bytes) and one PB equals 1,000TB when referring to storage capacity. Usable capacity will vary from the raw capacity due to object storage methodologies and other factors.

References in this publication to HGST products, programs or services do not imply that HGST intends to make these available in all countries in which HGST operates.

Product information is provided for information purposes only and does not constitute a warranty.

Information is true as of the date of publication and is subject to change. Actual results may vary. This publication is for general guidance only. Photographs may show design models.

Contents

List of Figures

List of Tables

1 About this Guide

Topics:

This guide provides procedures for customer replaceable units and field replaceable units of the HGST Active Archive System.

1.1 Conventions

Element	Sample Notation
OS shell or Q-Shell commands (user input)	`rm -rf /tmp`
OS shell or Q-Shell system output	`Installation successful!`
Commands longer than one line are split with "\"	`q.dss.manage.setPermissions('/manage', \ [....])`
User-supplied values	*`ManagementNodeVirtualIPAddress`* or `<ManagementNodeVirtualIPAddress>`
File and directory names	The file `aFile.txt` is stored in `/home/user`.
Any graphical user interface label	Click **OK**.
Keyboard keys and sequences	To cancel the operation, press `Ctrl+c`.
Menu navigation in a GUI	Navigate to **Dashboard** > **Administration** > **Hardware** > **Servers**.

1.2 Storage Notations

Convention	Prefix	Size (bytes)
KB	kilobyte	1,000
KiB	kibibyte	1,024
MB	megabyte	1,000,000
MiB	mebibyte	1,048,567
GB	gigabyte	1,00,000,000
GiB	gibibyte	1,073,741,824
TB	terabyte	1,000,000,000,000
TiB	tibibyte	1,099,511,627,776

- Sizes of disks are expressed with *SI prefixes* (kilo, mega, tera, peta, exa)
- Space, size of partitions and file systems are expressed with the *binary prefixes* (kibi, mebi, tebi, pebi, exbi)
- A comma (",") is used for digit grouping, for example 1,000 is 1 thousand.
- A period (".") is used as decimal mark, for example 12.5 %.

1.3 Admonitions

Type	Usage
Note:	Indicates extra information that has no specific hazardous or damaging consequences.
Tip:	Indicates a faster or more efficient way to do something.
Caution:	Indicates an action that, if taken or avoided, may result in hazardous or damaging consequences.
Warning:	Indicates an action that, if taken or avoided, may result in data loss or unavailability.

1.4 Related Documents

For more information about the Active Archive System, please consult the following documents:

- The *HGST Active Archive System Administration Guide* explains how to use the Active Archive System interfaces for executing system management, monitoring, and analytics tasks.
- The *HGST Active Archive System API Guide* provides a reference for the Active Archive System S3 API.
- The *HGST Active Archive System FRU Replacement Guide* provides procedures for replacing hardware components of the Active Archive System.
- The *HGST Active Archive System Installation Guide* provides instructions for the installation of the Active Archive System in the data center, and its initial bringup.
- The *HGST Active Archive System Release Notes* provide important information about changes, new features, and known limitations.
- The *HGST Active Archive System Site Requirements Document* contains data center requirements for the Active Archive System.
- The *HGST Active Archive System Troubleshooting Guide* provides help for issues you might encounter.
- The *HGST Active Archive System Upgrade Guide* provides instructions for software and firmware updates, and system expansion.

For the latest or online version of any of these documents, visit http://www.hgst.com/support.

1.5 Weight

Rack:

The following table displays the weight of the Active Archive System:

Hardware	Dimensions (Width x Height x Depth)
Active Archive System	2,250 lbs. 1,020 kg.

Table 1: Active Archive System Weight

Note: The weight mentioned previous is the total unpacked weight after delivery.

Controller (SM 1028U-TR4T+):

The following table displays the weight of the Controller:

Hardware	Dimensions (Width x Height x Depth)
Controller	Net weight is 26lbs. Gross weight is 41 lbs **Note:** The gross weight of the controller is based on the combined weight of the server, accessories kit, rail kit, and packaging

Table 2: Active Archive System Weight

Storage (SM 1018R-WC0R):

The following table displays the weight of the Storage server:

Hardware	Dimensions (Width x Height x Depth)
Storage server	Net weight is 25lbs. Gross weight is 40lbs **Note:** The gross weight of the storage server is based on the combined weight of the server, accessories kit, rail kit, and packaging

Table 3: Active Archive System Weight

2 For More Information

Topics:

This chapter provides points of contact for the Active Archive System.

2.1 Points of Contact

For further assistance with the Active Archive System, contact Elastic Storage Platforms support. Please be prepared to provide the following information: serial number (S/N), product name, model number, and a brief description of the issue.

Telephone:

Region	Telephone Numbers	Support Hours and Additional Information
United States/International	1-408-717-7766	24 hours a day, 7 days a week
North America	1-844-717-7766	24 hours a day, 7 days a week Toll-free

Email:
support@hgst.com

Website:
www.hgst.com/support

3 Controller Node Replaceable Units

Topics:

This section provides replacement procedures for the following parts in a Controller Node:

- Chassis
- HDD
- SSD
- PSU

3.1 Warnings

Caution: Opening or removing the system cover when the system is powered on may expose you to a risk of electric shock.

When replacing items from the inside of the chassis, ensure that you take precautions to prevent electrostatic discharge (ESD).

Warning: Replace only one disk at a time on the Management Node.

3.2 Chassis Replacement Procedure

The Controller Node chassis is a *SuperMicro DP 1U Server, 1028*. Replacing the chassis replaces its NICs, CPU, memory, motherboard, and fans, **but not its disks**.

Prerequisites

- Obtain a replacement Controller Node chassis from HGST.
- Obtain the virtual IP address of the Management Node.
- Obtain the IP addresses of the other (non failed) Controller Nodes.
- Obtain the admin password for the CMC.
- Obtain the root password.
- Fill in as much of the work table as possible before starting this procedure.

Required Tools

- Ladder
- Long Phillips-head screwdriver

Time Estimate: 3 hours.

Figure 1: Overview of Controller Node Chassis Replacement

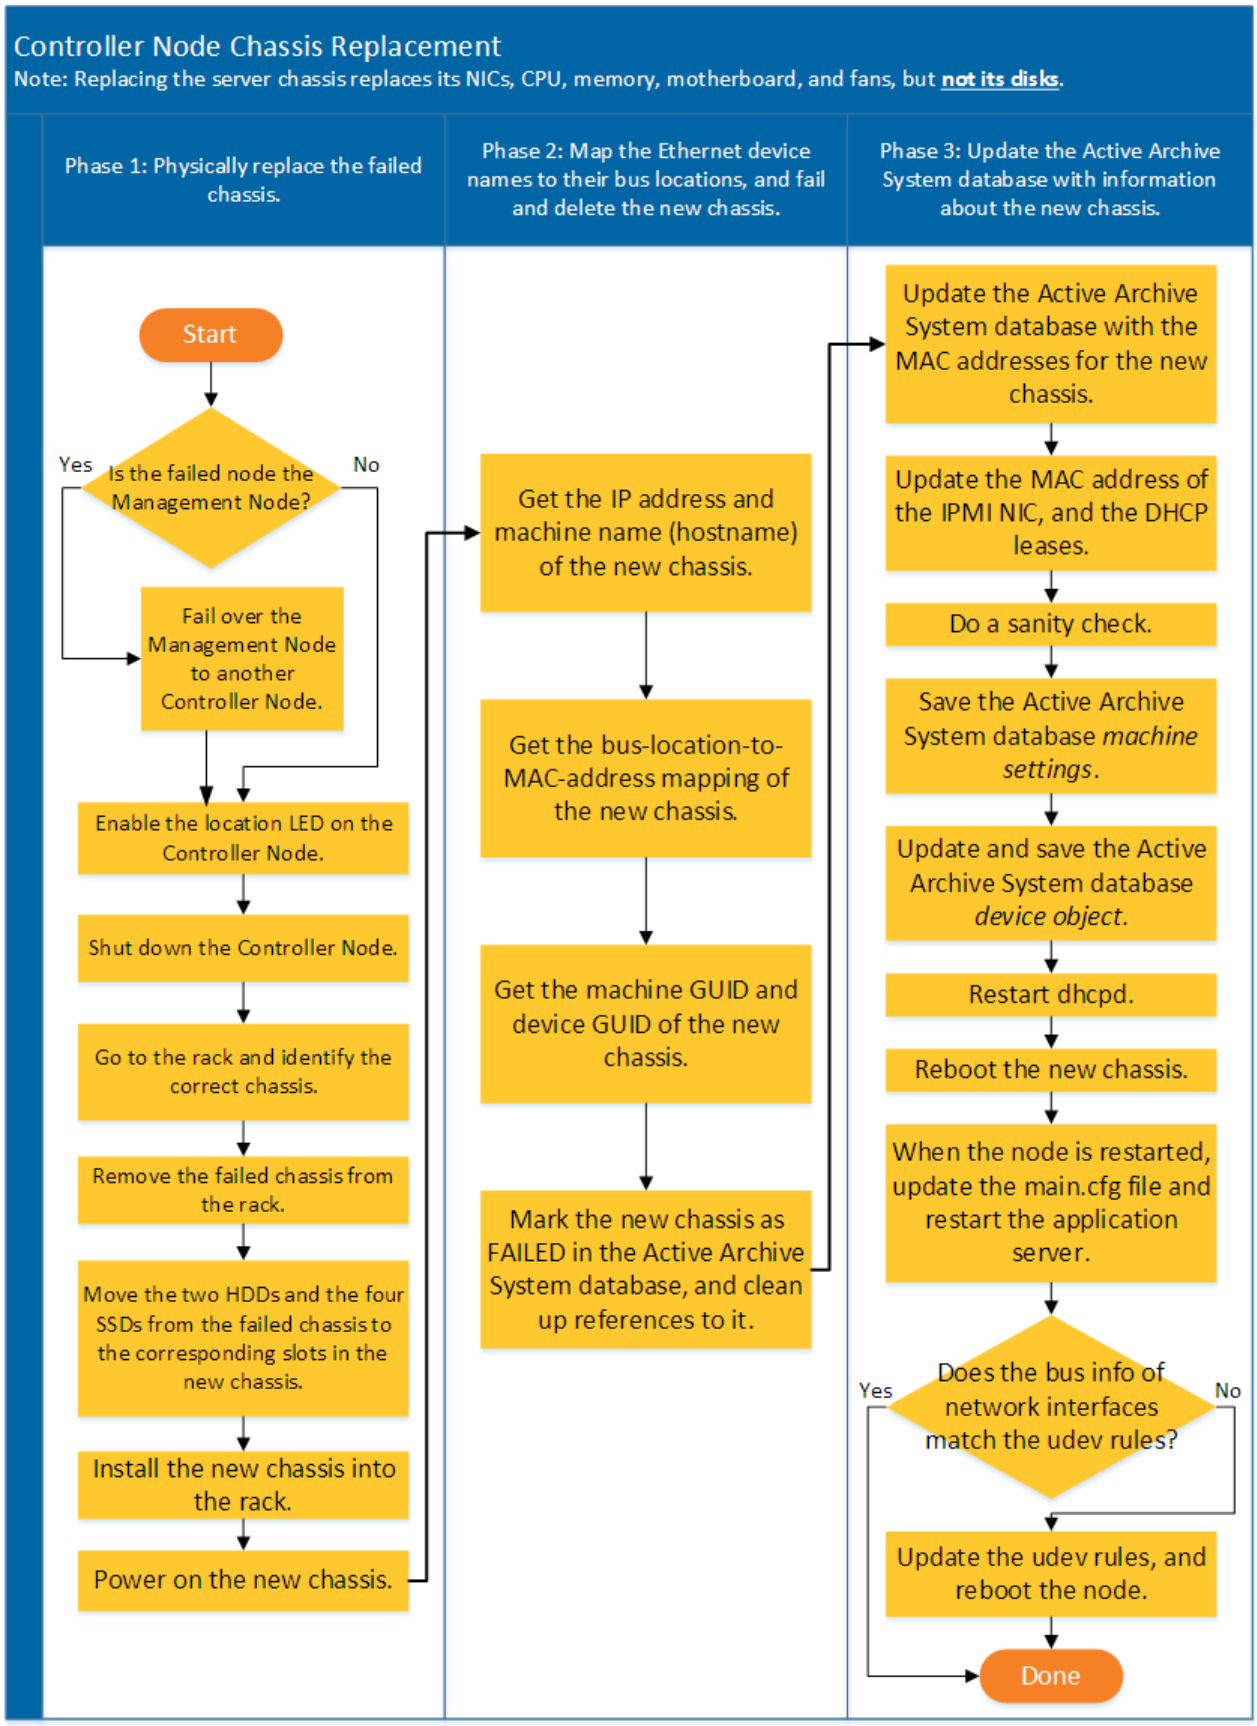

A work table is provided at the end of this section for your convenience, to store all of the information needed for a chassis replacement.

To replace a Controller Node chassis, proceed as follows:

1. 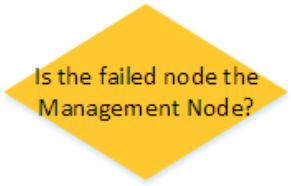

   If the failed chassis is the Management Node chassis, fail over the Management Node to another Controller Node.

   a) Open an SSH session to any Controller Node.

      You must obtain the IP addresses of the Controller Node ahead of time.

   b) Use the following command to determine the virtual IP address of the Management Node.

      ```
      grep dmachine.amplistor.com /etc/hosts | grep -v 127.0.0.1 | awk '{print $1}'
      ```

      The output of this command is the virtual IP address of the Management Node. For example,

      ```
      172.16.63.154
      ```

   c) Open an SSH session to the Management Node using the virtual IP address obtained in the previous substep.

   d) Exit the OSMI menu.
      The Linux prompt appears.

   e) Copy or write down the hostname in the Linux prompt.

   f) Log into the CMC.

   g) Navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Controller Nodes**, and select the failed Controller Node.

   h) Compare the hostname of the failed node, as displayed in the CMC, to the hostname you saved from substep e.

   i) If the failed node is the Management Node, fail over the Management Node to another Controller Node.

      For instructions on how to fail over the Management Node, see *Managing Hardware* in the *HGST Active Archive System Administration Guide*.

2. 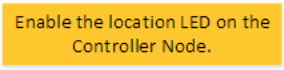

   Enable the location LED on the Controller Node.

   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Controller Nodes**.

   b) Select the correct node.

   c) In the **Commands** pane, click **Location LED On**.
      A blue LED on its front and back panels is now blinking.

3. 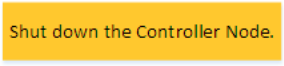

   Shut down the Controller Node from the CMC.

   Note: Save the node's hostname in your worktable under **Original Hostname of Node**.

   ---

   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Controller Nodes**.

   b) Select the desired Controller Node.

c) In the **Commands** pane, click **Shutdown**.

Figure 2: The Shutdown Button in the Commands Pane

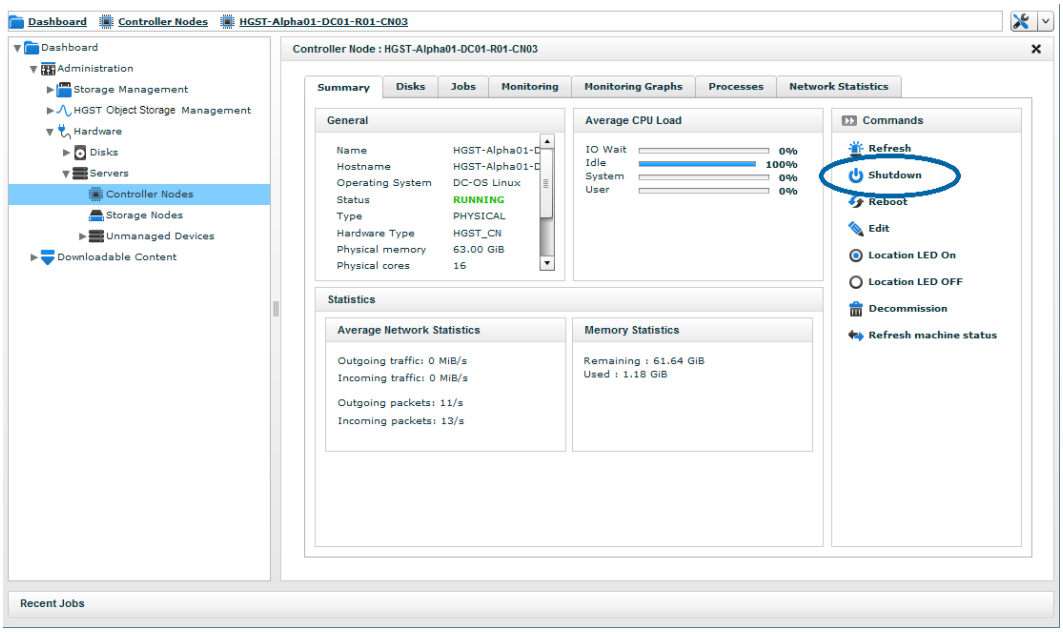

d) Wait for the **Status** field to change to **DONE**.

4. Go to the rack and identify the correct chassis.

Go to the rack and identify the correct chassis by the blinking blue LED on its front and back panels.

5. Remove the failed chassis from the rack.

Remove the failed chassis from the rack.

a) At the front of the chassis, loosen the rack mounting screws.

b) At the back of the chassis, disconnect the five network cables connected to the ports labeled as follows in the image below.

Important: Pull very gently on the pull tabs of the SFP+ cables, otherwise they might break.

Note: Check that the cables are labeled correctly, so that you can put them back in the same order.

---

a. N1 (remove the SFP+ optical transceiver attached to the cable)

b. N2

c. N3

d. N4 (remove the SFP+ optical transceiver attached to the cable)

e. **M2**

Figure 3: Controller Node, Back

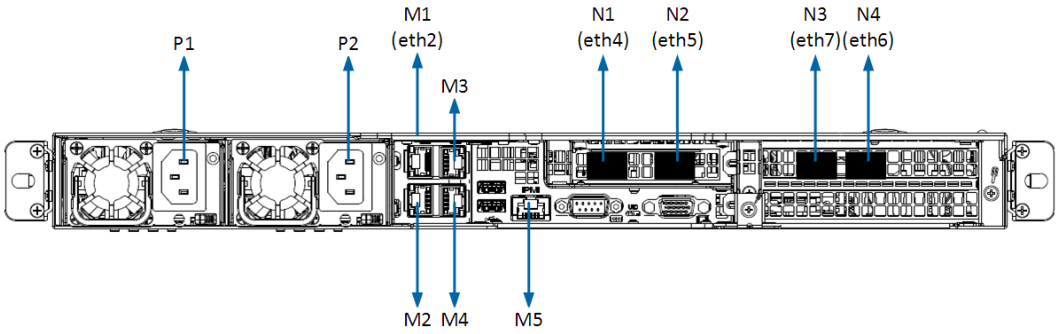

c) At the back of the chassis, disconnect the two power cords.

In the image above, the power cords are connected to the PSUs labeled `P1` and `P2`.

d) At the back of the chassis, depress both rail slides inwards, and push the chassis towards the front of the rack, until the rails pass the safety catch (about 3"/7.6cm).

e) At the front of the chassis, slowly slide the chassis out until you reach the *pull-safety* at the midway point (you will hear a soft clicking sound, and feel the chassis "catch" on the rails).

f) Disengage the *pull-safety* on both sides of the chassis and slide it out until the split line of the two top covers.

Push the *pull-safety* on one side up, and the *pull-safety* on the other side down.

g) Continue to slowly slide the chassis out until you reach the *pull-safety* at the end point, and disengage it as you did the earlier one.

h) Safely unmount the chassis from the rack and place it on a table.

Caution: A Controller Node chassis weighs about 50lbs. Ensure that you have sufficient manpower to handle it safely.

Warning: Once you pull the chassis past the *pull-safety*, do not leave it hanging in the rack. Otherwise, the rack rails may be damaged permanently.

---

6. Move the two HDDs and the four SSDs from the failed chassis to the corresponding slots in the new chassis.

Move the two HDDs and the four SSDs from the failed chassis to the **exact corresponding slots** in the new chassis.

Tip: Write down the disk serial number and slot location so that you can double-check that each disk is seated in the correct slot post installation into the new chassis.

---

a) Remove each disk from its slot in the front bay of the failed chassis.

b) Install the disk into the corresponding slot in the new chassis.

Figure 4: Controller Node, Front

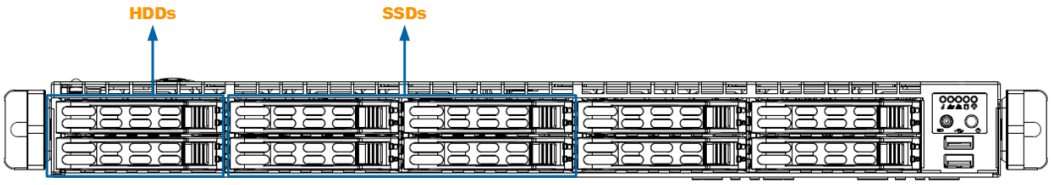

7. Install the new chassis into the rack.

Install the new chassis into the rack.

a) Mount the new chassis onto the rack slides and slide it into the rack.

Caution: Mounting the chassis is a two person task.

b) Tighten the rack mounting screws to secure the chassis to the rack.

c) Reconnect the five network cables to the chassis ports.

The network cables are labeled.

- Connect the cable labeled `CNx.N1.SW1.Nxx` (with SFP+ optical transceiver attached) to the port labeled **N1** in the image below.
- Connect the cable labeled `CNx.N2.SW1.Nxx` to the port labeled **N2** in the image below.
- Connect the cable labeled `CNx.N3.SW2.Nxx` to the port labeled **N3** in the image below.
- Connect the cable labeled `CNx.N4.SW1.Nxx` (with SFP+ optical transceiver attached) to the port labeled **N4** in the image below.
- Connect the cable labeled `CNx.M2.SW1.Nxx` to the port labeled **M2** in the image below.

Figure 5: Controller Node, Back

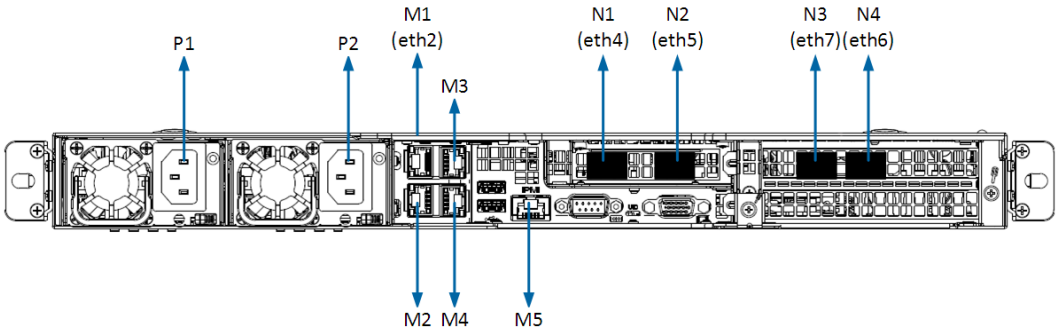

d) Reconnect the power cords.

8. Get the MAC address of the new chassis IPMI NIC from the BIOS.

a) Connect a VGA monitor and USB keyboard to the new chassis.

b) Power on the new chassis.

The power button is located on the chassis front control panel.

c) At power up, press `Del` to enter into the system BIOS.

d) In the system BIOS, navigate to **IPMI** > **BMC Network Configuration** > **.**

e) Record **Station MAC Address** in your work table under **IPMI MAC Address of Node**.

f) Exit the BIOS without saving any changes by pressing the `ESC`.

The boot process continues.

g) Disconnect the VGA monitor and USB keyboard from the new chassis.

9. Get the IP address and machine name (hostname) of the new chassis.

Get the IP address and machine name (hostname) of the new chassis.

a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Unmanaged Devices** > **Uninitialized**.
The new chassis appears in the list of uninitialized devices. This indicates that it has started successfully.
b) Write the value of **Name** into your work table, under **Temporary Hostname of Node**.
c) Write the value of **Name** without the `PM-` prefix into your work table, under **MAC Address of Node**.
d) Write the **IP address** into your work table, under **Temporary IP Address of Node**.

Figure 6: Uninitialized Nodes

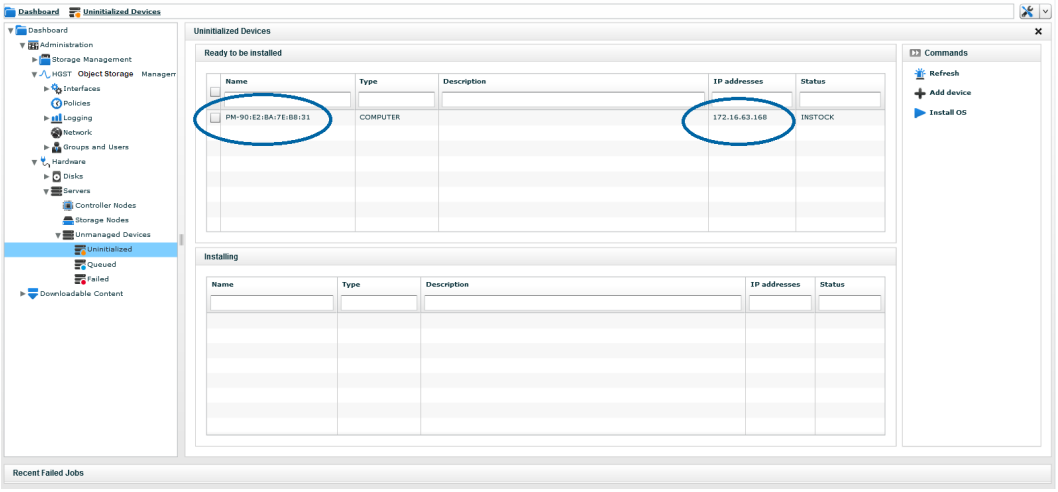

10. Get the bus-location-to-MAC-address mapping of the new chassis.

Get the bus-location-to-MAC-address mapping of the new chassis.

a) From the Management Node, open an SSH session to the **new IP address of the Controller Node** obtained in the previous step.
Log in with username `root` and password `rooter`.
b) At the Linux prompt, run the following command:

```
for add in `ls /sys/devices/pci*/*/*/net/*/address`; do echo $add; cat
 $add; done
```

The output of this command is similar to the example below.

```
/sys/devices/pci0000:00/0000:00:01.0/0000:01:00.0/net/eth0/address
00:25:90:fd:e8:7c
/sys/devices/pci0000:00/0000:00:01.0/0000:01:00.1/net/eth2/address
00:25:90:fd:e8:7d
/sys/devices/pci0000:00/0000:00:01.0/0000:01:00.2/net/eth3/address
00:25:90:fd:e8:7e
/sys/devices/pci0000:00/0000:00:01.0/0000:01:00.3/net/eth5/address
00:25:90:fd:e8:7f
/sys/devices/pci0000:80/0000:80:01.0/0000:81:00.0/net/eth1/address
90:e2:ba:7c:5a:fc
/sys/devices/pci0000:80/0000:80:01.0/0000:81:00.1/net/eth4/address
90:e2:ba:7c:5a:fd
/sys/devices/pci0000:80/0000:80:02.0/0000:82:00.0/net/eth6/address
```

```
90:e2:ba:7c:5d:a4
/sys/devices/pci0000:80/0000:80:02.0/0000:82:00.1/net/eth7/address
90:e2:ba:7c:5d:a5
root@nfsROOT:~#
```

The output of this command shows the serial bus path (for example, `0000:81:00.1`) and the new MAC address (for example, `90:e2:ba:7c:5a:fc`).

Tip: As an alternative to the command above, you can use the command below to print only the serial bus paths and MAC addresses in uppercase.

```
for add in `ls /sys/devices/pci*/*/*/net/*/address`; do echo -en
 "`echo $add|sed 's/\// /g' | awk '{print $5}'`\t"; cat $add|tr 'a-f'
 'A-F'; done
```

c) Fill in the serial bus path in ascending order in the **Serial Bus Path** column of the work table.

d) Fill in the MAC address corresponding to the serial bus path in ascending order in the **MAC Address on the New Chassis** column of the work table.

For the sample output from the step above, the work table would look like this:

Serial Bus Path	MAC Address on the New Chassis	Ethernet Port Name
0000:01:00.0	00:25:90:fd:e8:7c	eth0
0000:01:00.1	00:25:90:fd:e8:7d	eth1
0000:01:00.2	00:25:90:fd:e8:7e	eth2
0000:01:00.3	00:25:90:fd:e8:7f	eth3
0000:81:00.0	90:e2:ba:7c:5a:fc	eth4
0000:81:00.1	90:e2:ba:7c:5a:fd	eth5
0000:82:00.0	90:e2:ba:7c:5d:a4	eth6
0000:82:00.1	90:e2:ba:7c:5d:a5	eth7

Table 4: Work Table with Sample MAC Addresses and Serial Bus Paths

e) Close the SSH session to the Controller Node.
You are now back in the SSH session to the Management Node.

11. Get the machine GUID and device GUID of the new chassis.

Get the machine GUID and device GUID of the new chassis.

a) On the Management Node, start the Q-Shell:

```
/opt/qbase3/qshell
```

b) Create a cloudAPI connection.

```
cloudapi = i.config.cloudApiConnection.find('main')
```

c) Retrieve the machine GUID for the **new chassis**, using the value of **Temporary Hostname of Node** in **uppercase**, from the work table, for *hostname_of_new_node* in the command below:

```
machine_guid = cloudapi.machine.find(name='hostname_of_new_node')['result'][0]
```

For example,

```
machine_guid = cloudapi.machine.find(name='PM-90:E2:BA:7E:B8:31')['result'][0]
```

d) Retrieve the device GUID using the machine GUID you obtained from the previous step.

```
dg = cloudapi.machine.list(machineguid=machine_guid)['result'][0]['deviceguid']
```

e) Sanity check: print the value of `dg`.
For example,

```
dg
'd951f6d9-7104-470d-8c97-ecf52d57c7b5'
```

12. Mark the new chassis as FAILED in the Active Archive System database, and clean up references to it.

Mark the new chassis as **FAILED** in the Active Archive System database, and clean up references to it.

The Active Archive System created a new **INSTOCK** node in its database for the new chassis. If you do not mark the new chassis as **FAILED** in the database, you are in effect adding a new node rather than replacing an existing node's chassis. Therefore, you must remove the **INSTOCK** node by following the steps below.

a) Mark the new chassis as **FAILED** in the Active Archive System database:
Execute this command on the Management Node:

```
cloudapi.device.updateModelProperties(dg, \
status=str(q.enumerators.devicestatustype.FAILED))
```

The new chassis now appears under the **FAILED** list in the CMC, and is removed from the **Unmanaged Devices** list.

Figure 7: The New Chassis Appears Under the FAILED List in the CMC

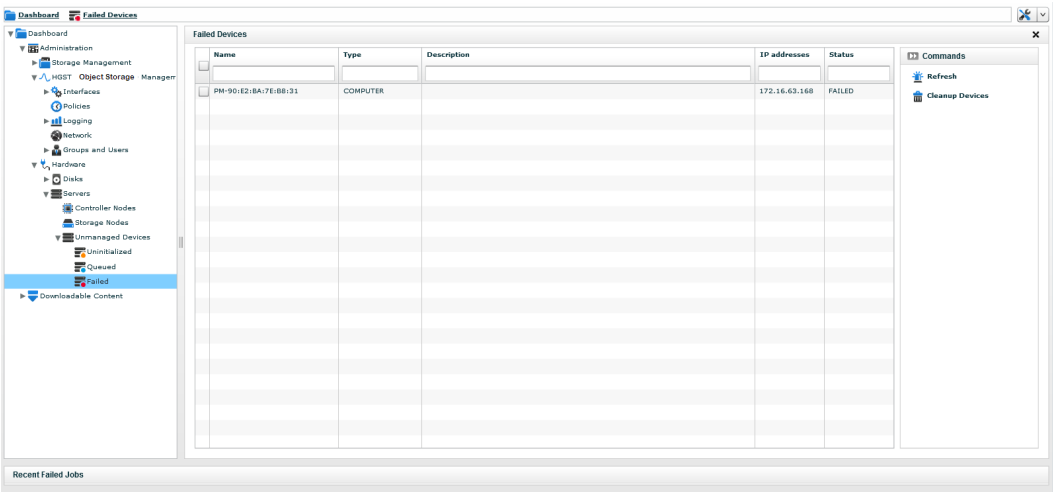

b) From the Management Node, clean up references to the new chassis in the Active Archive System database.

In the command below, replace *MAC_ADDRESS* with the value you wrote in the work table for **MAC Address of Node**.

Important: Use capital letters for the MAC address.

```
q.amplistor.cleanupMachine('MAC_ADDRESS')
```

For example,

```
In [14]: q.amplistor.cleanupMachine('90:E2:BA:7E:B8:31')
Out[14]: True
```

This command takes about 10 seconds to complete.

c) Do a sanity check.

Refresh the screen by clicking **Refresh** in the **Commands** pane. Check that the new chassis is no longer in the **FAILED** list.

13. Update the Active Archive System database with the MAC addresses for the new chassis.

Update the Active Archive System database with the MAC addresses for the new chassis.

a) From the Management Node, create a cloudAPI connection.

```
cloudapi = i.config.cloudApiConnection.find('main')
```

b) From the Management Node, get the machine GUID using your work table value for **Original Hostname of Node**.

Note: Use upper case for *HOSTNAME_OF_OLD_NODE*.

```
machine_guid = cloudapi.machine.find(name='HOSTNAME_OF_OLD_NODE')\
['result'][0]
```

For example,

```
machine_guid = cloudapi.machine.find(name='HGST-Alpha02-DC01-R02-CN01')\
['result'][0]
```

c) From the Management Node, get the machine object.

```
machine = cloudapi.machine.getObject(machine_guid)
```

d) Display all the Ethernet port names (`ethN`) that are registered:
For example,

```
In [4]: print machine.nics[0].name
eth0
In [5]: print machine.nics[1].name
eth2
In [6]: print machine.nics[2].name
eth3
In [7]: print machine.nics[3].name
eth5
In [8]: print machine.nics[4].name
eth1
In [9]: print machine.nics[5].name
eth4
```

```
In [10]: print machine.nics[6].name
eth6
In [11]: print machine.nics[7].name
eth7
In [12]: print machine.nics[8].name
BMC
```

e) Write the *index* of the above `machine.nics[index].name` value into the work table in column **NIC Array ID**, in the row corresponding to `eth`*N*.

For the sample output from the step above, the work table would look like this:

Serial Bus Path	MAC Address on the New Chassis	Ethernet Port Name	NIC Array ID
`0000:01:00.0`	`00:25:90:fd:e8:7c`	`eth0`	`0`
`0000:01:00.1`	`00:25:90:fd:e8:7d`	`eth1`	`4`
`0000:01:00.2`	`00:25:90:fd:e8:7e`	`eth2`	`1`
`0000:01:00.3`	`00:25:90:fd:e8:7f`	`eth3`	`2`
`0000:81:00.0`	`90:e2:ba:7c:5a:fc`	`eth4`	`5`
`0000:81:00.1`	`90:e2:ba:7c:5a:fd`	`eth5`	`3`
`0000:82:00.0`	`90:e2:ba:7c:5d:a4`	`eth6`	`6`
`0000:82:00.1`	`90:e2:ba:7c:5d:a5`	`eth7`	`7`
IPMI	See **IPMI MAC Address of Node** in the work table.	`BMC`	`8`

Table 5: Work Table with Sample Ethernet Port Names and NIC Array IDs

f) Update the database entry for `machine.nics[N].hwaddr` with the corresponding MAC address for `eth`*N* from your work table.

Important: Use capital letters for the MAC address.

```
machine.nics[0].hwaddr = 'NEW_MAC_ADDRESS_FOR_ETHN'
```

For example,

```
machine.nics[0].hwaddr = '00:25:90:FD:E8:7C'
machine.nics[1].hwaddr = '00:25:90:FD:E8:7E'
machine.nics[2].hwaddr = '00:25:90:FD:E8:7F'
machine.nics[3].hwaddr = '90:E2:BA:7C:5A:FD'
machine.nics[4].hwaddr = '00:25:90:FD:E8:7D'
machine.nics[5].hwaddr = '90:E2:BA:7C:5A:FC'
machine.nics[6].hwaddr = '90:E2:BA:7C:5D:A4'
machine.nics[7].hwaddr = '90:E2:BA:7C:5D:A5'
```

g) Update the database entry for `machine.nics[8].hwaddr` with the corresponding IPMI MAC address from your work table, under **MAC Address of Node**.

Important: Use capital letters for the MAC address.

For example,

```
machine.nics[8].hwaddr = '0C:C4:7A:36:8B:12'
```

14. Update the MAC address of the IPMI NIC, and the DHCP leases.

Update the MAC address of the IPMI NIC, and the DHCP leases.

a) Log into the CMC.
b) Navigate to the CMC's view of the node whose chassis you have replaced.
c) Select to that node's **Summary** tab.
d) Write the IPMI IP address, as shown in the **General** section, in your work table, under **IPMI IP Address of Node**.

Figure 8: The Old IPMI IP Address of the Node

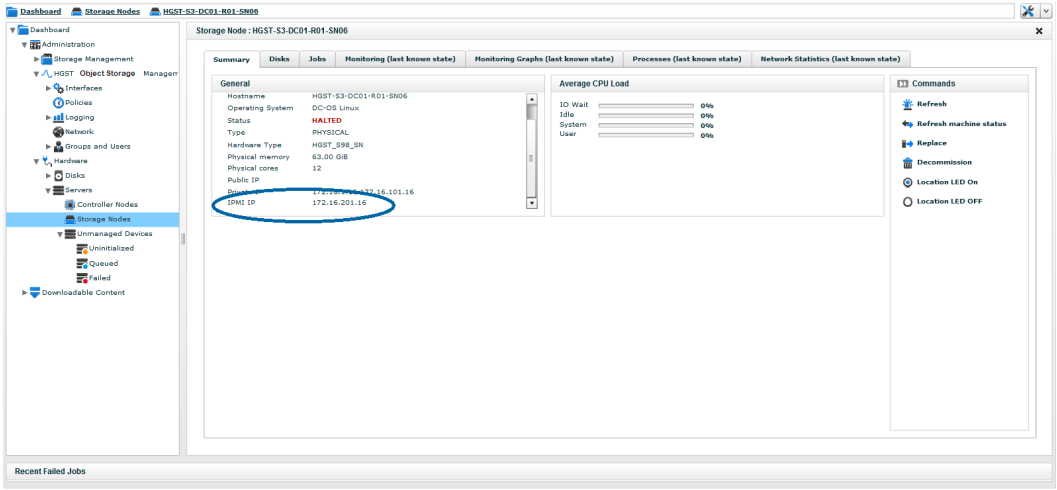

e) Leave the current SSH session as is. Open a new SSH session on the Management Node.
f) Open `/opt/qbase3/cfg/dhcpd/dhcpd.leases` with your text editor.
g) Search for the IPMI IP address (obtained in substep d) in the file.
The section containing the IPMI IP address looks like the following example:

```
host 457f495a-80b7-4125-862b-5f87d9121cfa {
  dynamic;
  hardware ethernet 0c:c4:7a:36:8b:12:;
  fixed-address 172.16.201.16;
  group "pmachines";
}
```

h) Change the `hardware ethernet` value to the new IPMI MAC address **in lowercase** from your work table, under **IPMI MAC Address of Node**.

> **Note:** Use lowercase only.

For example,

```
hardware ethernet 0c:c4:7a:36:8b:12:;
```

i) Save and close the file.
j) Exit the new SSH session.

15. Do a sanity check.

Do a sanity check to verify that you have updated the new MAC addresses correctly.

Compare the output of the command below to your work table.

```
In [9]: for nic in machine.nics: nic.name; nic.hwaddr
   ...:
Out[9]: 'eth0'
Out[9]: '00:25:90:FD:E8:7C'
Out[9]: 'eth2'
Out[9]: '00:25:90:FD:E8:7E'
Out[9]: 'eth3'
Out[9]: '00:25:90:FD:E8:7F'
Out[9]: 'eth5'
Out[9]: '90:E2:BA:7C:5A:FD'
Out[9]: 'eth1'
Out[9]: '00:25:90:FD:E8:7D'
Out[9]: 'eth4'
Out[9]: '90:E2:BA:7C:5A:FC'
Out[9]: 'eth6'
Out[9]: '90:E2:BA:7C:5D:A4'
Out[9]: 'eth7'
Out[9]: '90:E2:BA:7C:5D:A5'
Out[9]: 'BMC'
Out[9]: '0C:C4:7A:36:8B:12'
```

16. Save the Active Archive System database *machine settings*.

Save the Active Archive System database *machine settings*.

```
q.drp.machine.save(machine)
```

17. Update and save the Active Archive System database *device object*.

Update and save the Active Archive System database *device object*.

a) Get the device object.

```
device = cloudapi.device.getObject(machine.deviceguid)
```

b) Update the MAC address of the chassis with the value you saved in the work table under **MAC Address of Node**.

Note: Use capital letters for the MAC address.

```
device.nicports[0].hwaddr = 'NEW_MAC_ADDRESS'
```

For example,

```
In [12]: device.nicports[0].hwaddr='90:E2:BA:7E:B8:31'
```

c) Save the device object.

```
q.drp.device.save(device)
```

18. Restart dhcpd.

Restart dhcpd.

```
In [14]: q.manage.dhcpd.restart()
```

```
Stopping dhcpd...
dhcpd is halted
Starting dhcpd...
dhcpd is running
```

19. Reboot the new chassis.

Reboot the new chassis.

a) Open an SSH session to the node whose chassis you have just replaced.

Use the IP address you saved in the work table under **Original IP Address of Node**.

The Linux prompt appears.

b) At the Linux prompt, run the following command:

```
reboot
```

20. When the node is restarted, update the main.cfg file and restart the application server.

When the node is restarted, update the `main.cfg` file and restart the application server.

a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Controller Nodes**.
A list of Controller Nodes appears in the CMC.

b) Click the Controller Node whose chassis you have just replaced.

Identify the correct Controller Node by its hostname: it now matches the **Original Hostname of Node** value you recorded in the worktable. This value is typically of the format *SystemID*-DC*nn*-R*nn*-CN*nn*.

c) Identify the IP addresses listed in the **Private IP** field.

d) Open an SSH session to the Controller Node, using any one of the IP addresses you obtained from substep c, and exit the OSMI menu.
The Linux prompt appears.

e) At the Linux prompt on the Controller Node, open the file `/opt/qbase3/cfg/qconfig/main.cfg` with your text editor.
The file has a section that looks like this:

```
[main]
lastlogcleanup = 1428960577
domain = somewhere.com
nodetype = CPUNODE
nodename = 90E2BA7EB831
logserver_loglevel = 6
logserver_port = 9998
logserver_ip = 127.0.0.1
qshell_firstrun = False
machineguid = fc635662-5247-45b1-ab66-d0abe8e60712
```

f) Replace the value after `nodename =` with the new MAC address from your work table, under **MAC Address of Node**.

Note: The MAC address must be **in uppercase and without colons**. For example, `00:25:90:3B:C1:72` must be typed as `0025903BC172`.

---

g) Save and close the configuration file.

h) Start the Q-Shell.

```
/opt/qbase3/qshell
```

i) Restart the application server:

```
In [1]: q.manage.applicationserver.restart()
 Restarting applicationserver Applicationserver...
 Stopping applicationserver Applicationserver...
 Applicationserver is still running, waiting for 5 more seconds
 Applicationserver is still running, waiting for 4 more seconds
 Starting applicationserver Applicationserver...
```

j) Exit the Q-Shell.

```
quit()
```

21. Update the udev rules, and reboot the node.

Verify that the bus information of the network interfaces matches the `udev` rules.

a) Run the following command:

> **Tip:** Check the hardware paths in the command below, as they might be different on the new chassis.

```
root@HGST-S3-DC01-R01-CN03:~#  for add in `ls /sys/devices/pci*/*/*/net/*/
address`; do echo $add; cat $add; done
```

For example, the output of the above command looks like this:

```
/sys/devices/pci0000:00/0000:00:01.0/0000:01:00.0/net/eth0/address
00:25:90:fd:e8:7c
/sys/devices/pci0000:00/0000:00:01.0/0000:01:00.1/net/eth2/address
00:25:90:fd:e8:7d
/sys/devices/pci0000:00/0000:00:01.0/0000:01:00.2/net/eth3/address
00:25:90:fd:e8:7e
/sys/devices/pci0000:00/0000:00:01.0/0000:01:00.3/net/eth5/address
00:25:90:fd:e8:7f
/sys/devices/pci0000:80/0000:80:01.0/0000:81:00.0/net/eth1/address
90:e2:ba:7c:5a:fc
/sys/devices/pci0000:80/0000:80:01.0/0000:81:00.1/net/eth4/address
90:e2:ba:7c:5a:fd
/sys/devices/pci0000:80/0000:80:02.0/0000:82:00.0/net/eth6/address
90:e2:ba:7c:5d:a4
/sys/devices/pci0000:80/0000:80:02.0/0000:82:00.1/net/eth7/address
90:e2:ba:7c:5d:a5
```

> **Tip:** As an alternative to the command above, you can use the command below to print only the serial bus paths and MAC addresses in uppercase.
>
> ```
> for add in `ls /sys/devices/pci*/*/*/net/*/address`; do echo -en
>  "`echo $add|sed 's/\// /g' | awk '{print $5}'`\t"; cat $add|tr 'a-f'
>  'A-F'; done
> ```

b) Compare the output of the command above to the contents of the file `/etc/udev/rules.d/70-persistent-net.rules`.

For example, the contents of this file look like this:

```
root@HGST-S3-DC01-R01-CN03:~# cat /etc/udev/rules.d/70-persistent-net.rules
SUBSYSTEM=="net", ACTION=="add", KERNELS=="0000:01:00.0", KERNEL=="eth*",
 NAME="eth0"
```

```
SUBSYSTEM=="net", ACTION=="add", KERNELS=="0000:01:00.1", KERNEL=="eth*",
 NAME="eth2"
SUBSYSTEM=="net", ACTION=="add", KERNELS=="0000:01:00.2", KERNEL=="eth*",
 NAME="eth3"
SUBSYSTEM=="net", ACTION=="add", KERNELS=="0000:01:00.3", KERNEL=="eth*",
 NAME="eth5"
SUBSYSTEM=="net", ACTION=="add", KERNELS=="0000:81:00.0", KERNEL=="eth*",
 NAME="eth1"
SUBSYSTEM=="net", ACTION=="add", KERNELS=="0000:81:00.1", KERNEL=="eth*",
 NAME="eth4"
SUBSYSTEM=="net", ACTION=="add", KERNELS=="0000:82:00.0", KERNEL=="eth*",
 NAME="eth6"
SUBSYSTEM=="net", ACTION=="add", KERNELS=="0000:82:00.1", KERNEL=="eth*",
 NAME="eth7"
```

If they do not match, update `/etc/udev/rules.d/70-persistent-net.rules` to match the output of the command above, then reboot the node again.

```
reboot
```

The chassis replacement procedure is done.

Warning: Be very careful when recording and updating MAC addresses. A mistake may render the new chassis unusable.

Item	Value
Virtual IP Address of the Management Node: Get this value as instructed in the *Administrator Interfaces* chapter in the *HGST Active Archive System Troubleshooting Guide*	
IP Addresses of the Other (Non Failed) Controller Nodes:	
Original Hostname of Node: Get this value from the CMC **before** you shut down the failed node. This value is typically of the format *SystemID*-DC*nn*-R*nn*-CN*nn*.	
Temporary IP Address of Node: The CMC displays this value **after** the new chassis is installed.	
Temporary Hostname of Node: The CMC displays this value **after** the new chassis is installed. This value is of the format PM-*MAC_ADDRESS*.	
MAC Address of Node:	
IPMI IP Address of Node:	
IPMI MAC Address of Node:	

Serial Bus Path	MAC Address on the New Chassis	Ethernet Port Name	NIC Array ID
0000:01:00.0		eth0	
0000:01:00.1		eth1	
0000:01:00.2		eth2	
0000:01:00.3		eth3	
0000:81:00.0		eth4	
0000:81:00.1		eth5	
0000:82:00.0		eth6	
0000:82:00.1		eth7	
IPMI			

Table 6: Work Table for Controller Node Chassis Replacement

3.3 Hard Disk Drive Replacement Procedure

The HDD on a Controller Node is an *HGST 1TB SATA 6 Gb/sec drive*. It is a front-bay drive. It is hot swappable after being decommissioned in the CMC.

Prerequisites

- Decommission the faulty drive in the CMC. For more information, see *Managing Hardware* in the *HGST Active Archive System Administration Guide*.
- Obtain a replacement HDD from HGST.

Required Tools

- Ladder
- Long Phillips-head screwdriver

Time Estimate: 40 minutes.

To replace an HDD, proceed as follows:

Warning: Replace only one disk at a time on the Controller Node.

---

1. Obtain details about the decommissioned disk from the CMC.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Disks** > **Decommissioned**.

   Figure 9: Decommissioned Disks in the CMC

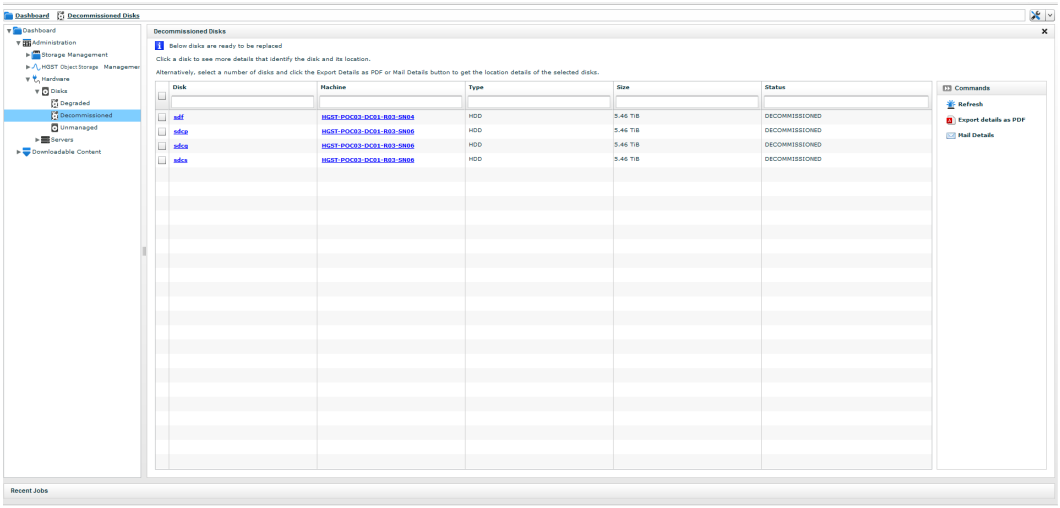

   b) Click the desired decommissioned disk.
   The *decommissioned disk details* are displayed.
   c) (Optional) Right-click anywhere in the decommissioned disk details, select **Print** > **Print to PDF**.
   The *decommissioned disk details* contains the following information, which you will need to refer to later:

   - The device name
   - The model type and serial number
   - A *drive map* showing the exact location of the disk (name of data center, rack, and node).
   - The node type
   - The MAC addresses of the node

- The current status of the node

Figure 10: Decommissioned Disk Details from the CMC

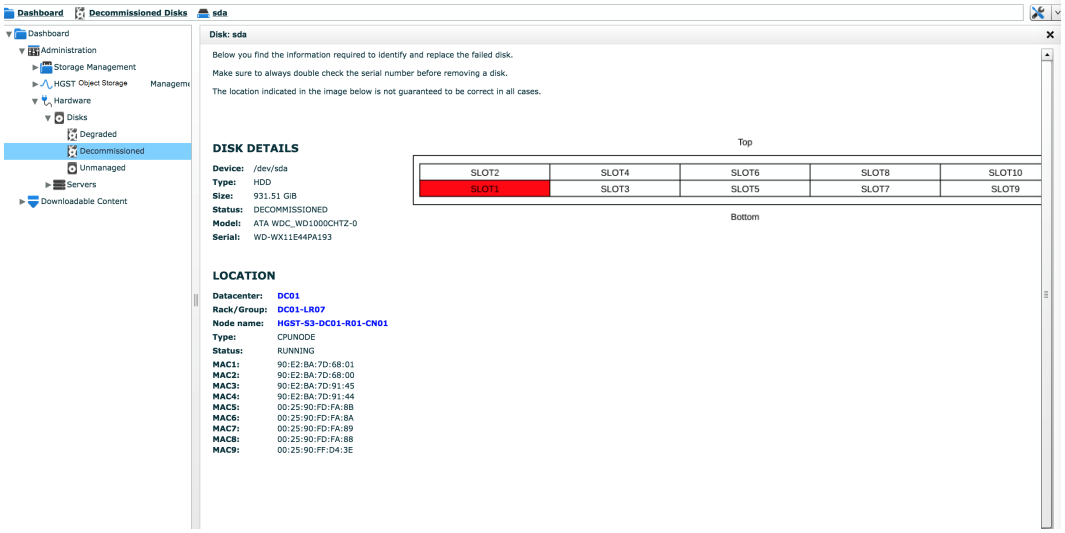

2. Enable the location LED on the Controller Node.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Controller Nodes**.
   b) Select the correct node.
   c) In the **Commands** pane, click **Location LED On**.
   A blue LED on its front and back panels is now blinking.
3. Go to the rack and identify the correct chassis by the blinking blue LED on its front and back panels.
4. Replace the decommissioned HDD.
   a) Identify the decommissioned HDD by using the drive map.
   The HDD is a front-bay drive.

Figure 11: Controller Node, Front

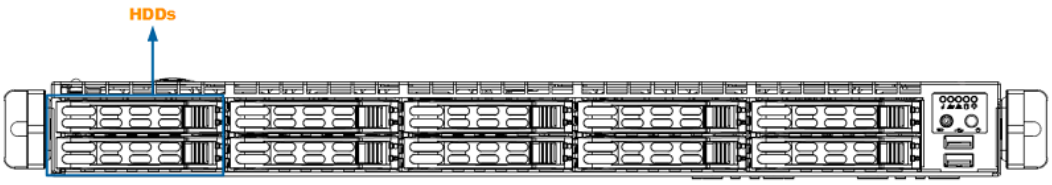

b) Press the release button on the drive carrier of the decommissioned HDD to extend the *drive carrier handle*.

Figure 12: Removing a Drive Carrier

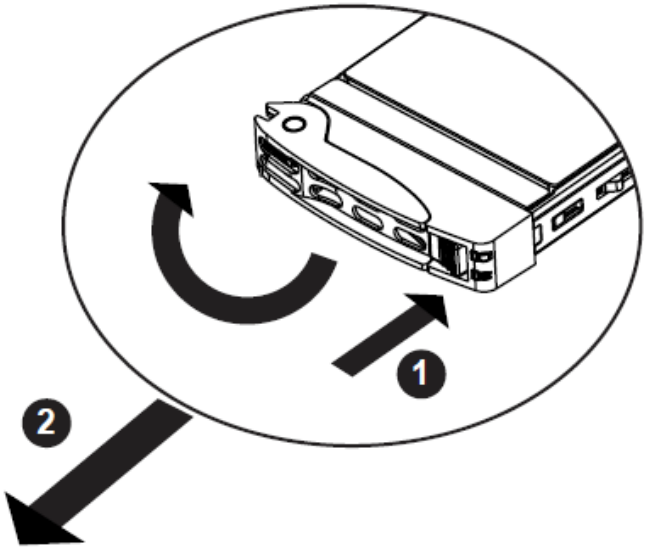

c) Pull drive carrier out of the front bay using the drive carrier handle.
d) Compare the serial number on the HDD to the serial number specified in the decommissioned disk details to confirm the that you have the correct HDD.
e) Unscrew the drive carrier from the decommissioned HDD.
f) Screw the drive carrier onto the replacement HDD.
g) Install the replacement HDD into the same slot that the decommissioned HDD was using.
A blue LED will blink for a moment.

5. Disable the location LED on the Controller Node.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Controller Nodes**.
   b) Select the desired Controller Node.
   c) In the **Commands** pane, click **Location LED Off**.
6. Confirm that the Active Archive System correctly determines the purpose for the new disk.
   a) Wait 15 minutes.
   b) In the CMC, navigate to **Dashboard** > **Administration** > **HGST Object Storage Management** > **Logging** > **Events**.
   c) In the **Events** list, check to see that a new empty disk has been detected.
   d) In the **Jobs** list, check to see that an `Initializing new disk` job has been triggered.
   It may take about 2 minutes for the job to appear.
   e) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers**.
   f) Select the desired node.
   g) Select the **Disks** tab.
   h) Wait for the physical drive that has been replaced, as well as the logical disks, to change status from a red icon to a green icon.

   Note: The physical drive that has been replaced, as well as the logical disks, may take up to 40 minutes to change status.

   The `Initializing new disk` job has completed successfully when the number of degraded disks decreases by 1.
7. If the disk still shows up in the **Degraded** or **Unmanaged** list, you must manually specify the purpose of the new disk:

a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Disks** > **Unmanaged**.
b) Select the new disk, and in the **Commands** pane, click **Repurpose**.
c) In the **Use As** field, select **Replacement Disk**.

> **Note:** You can only select **Replacement Disk** when there is a decommissioned disk. If there are no decommissioned disks, you can only select **Additional Disk** as the purpose for the disk.

d) In the **Replacement For** field, select the decommissioned disk that you want to replace.
e) Click **Next** to start the repurposing.
An `Initializing new disks on` *`node_name`* job starts.

If You Replaced the Wrong Disk

See *Troubleshooting*.

3.4 Solid State Disk Replacement Procedure

The a solid state disk (SSD) on a Controller Node is an *Intel DCS3500 Series 240 GB, SATA 6 Gb/s*. It is hot swappable after being decommissioned in the CMC.

Prerequisites

- **Caution:** If the faulty SSD has postgresql database on it, you must perform a failover first!

  To check whether the faulty SSD is the one hosting the postgresql database, do the following:

  1. In the CMC, navigate to **Dashboard** > **Administration** > **Storage Management** > **MetaStores** > **env_metastore**.
  2. See if the device name is part of env_metastore (in other words, if the device name appears under the **SSD Disk** column).

  To perform a failover, see *Managing Hardware* in the *HGST Active Archive System Administration Guide*.

- Decommission the faulty SSD in the CMC. For more information, see *Managing Hardware* in the *HGST Active Archive System Administration Guide* .
- Obtain a replacement SSD from HGST.

Required Tools

- Ladder
- Long Phillips-head screwdriver

Time Estimate: 40 minutes.

To replace an SSD, proceed as follows:

Warning: Replace only one disk at a time on the Controller Node.

1. Obtain details about the decommissioned disk from the CMC.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Disks** > **Decommissioned**.

   Figure 13: Decommissioned Disks in the CMC

   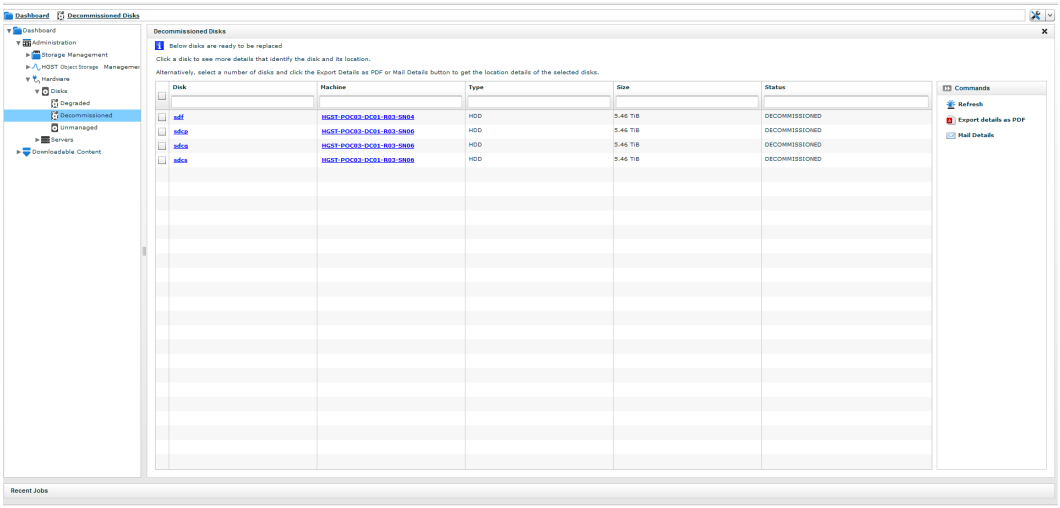

   b) Click the desired decommissioned disk.
   The *decommissioned disk details* are displayed.
   c) (Optional) Right-click anywhere in the decommissioned disk details, select **Print** > **Print to PDF**.

The *decommissioned disk details* contains the following information, which you will need to refer to later:

- The device name
- The model type and serial number
- A *drive map* showing the exact location of the disk (name of data center, rack, and node).
- The node type
- The MAC addresses of the node
- The current status of the node

Warning: The CMC identifies the incorrect slot for failed SSDs on Controller Nodes.

The image in the *decommissioned disk details* for SSDs is mislabeled: when it highlights slot 9, the decommissioned SSD is actually located in slot 5; when it highlights slot 10, the decommissioned SSD is actually located in slot 6.

Figure 14: Decommissioned Disk Details from the CMC

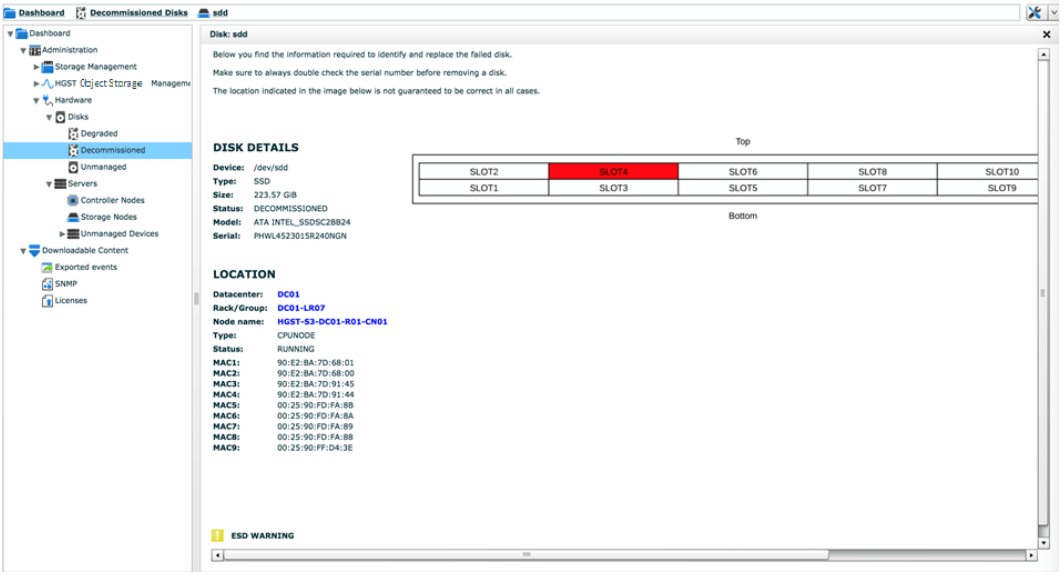

2. Enable the location LED on the Controller Node.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Controller Nodes**.
   b) Select the correct node.
   c) In the **Commands** pane, click **Location LED On**.
   A blue LED on its front and back panels is now blinking.
3. Go to the rack and identify the correct chassis by the blinking blue LED on its front and back panels.
4. Replace the decommissioned disk.
   a) Identify the decommissioned SSD by using the drive map.
   The SSD is a front-bay drive.

Figure 15: Controller Node, Front

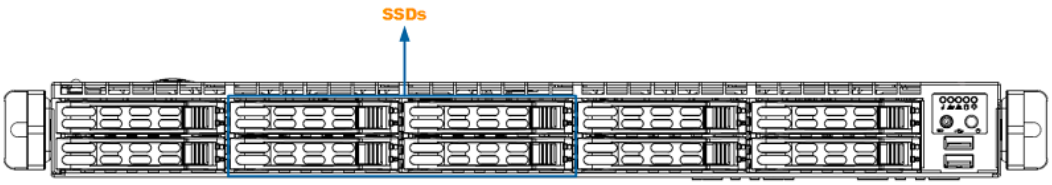

b) Press the release button on the drive carrier of the decommissioned SSD to extend the *drive carrier handle*.

Figure 16: Removing a Drive Carrier

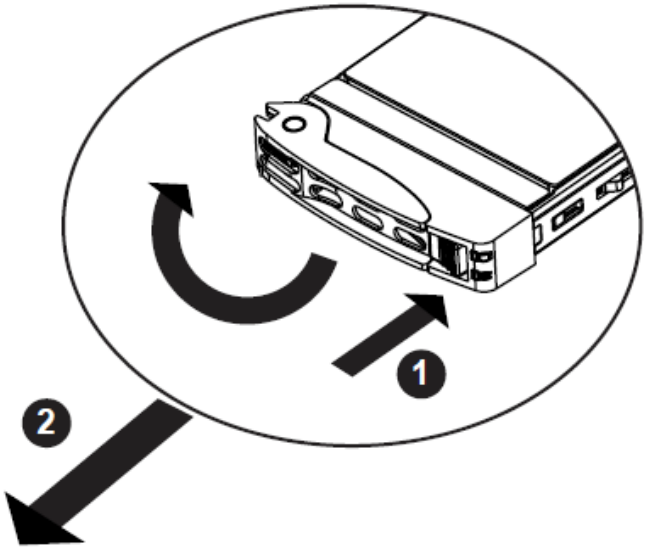

c) Pull drive carrier out of the front bay using the drive carrier handle.
d) Compare the serial number on the SSD to the serial number specified in the decommissioned disk details to confirm the that you have the correct SSD.
e) Unscrew the drive carrier from the decommissioned SSD.
f) Screw the drive carrier onto the replacement SSD.
g) Install the replacement SSD into the same slot that the decommissioned SSD was using.
A blue LED will blink for a moment.

5. Disable the location LED on the Controller Node.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Controller Nodes**.
   b) Select the desired Controller Node.
   c) In the **Commands** pane, click **Location LED Off**.
6. Confirm that the Active Archive System correctly determines the purpose for the new disk.
   a) Wait 15 minutes.
   b) In the CMC, navigate to **Dashboard** > **Administration** > **HGST Object Storage Management** > **Logging** > **Events**.
   c) In the **Events** list, check to see that a new empty disk has been detected.
   d) In the **Jobs** list, check to see that an `Initializing new disk` job has been triggered.
   It may take about 2 minutes for the job to appear.
   e) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers**.
   f) Select the desired node.
   g) Select the **Disks** tab.
   h) Wait for the physical drive that has been replaced, as well as the logical disks, to change status from a red icon to a green icon.

   Note: The physical drive that has been replaced, as well as the logical disks, may take up to 40 minutes to change status.

   The `Initializing new disk` job has completed successfully when the number of degraded disks decreases by 1.
7. If the disk still shows up in the **Degraded** or **Unmanaged** list, you must manually specify the purpose of the new disk:

a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Disks** > **Unmanaged**.
b) Select the new disk, and in the **Commands** pane, click **Repurpose**.
c) In the **Use As** field, select **Replacement Disk**.

> **Note:** You can only select **Replacement Disk** when there is a decommissioned disk. If there are no decommissioned disks, you can only select **Additional Disk** as the purpose for the disk.

d) In the **Replacement For** field, select the decommissioned disk that you want to replace.
e) Click **Next** to start the repurposing.
An `Initializing new disks on` *`node_name`* job starts.

If You Replaced the Wrong Disk

See *Troubleshooting*.

3.5 Power Supply Unit Replacement Procedure

The power supply units (PSU) on a Controller Node are redundant hot-swappable *SuperMicro 1U 750w 74mm, Platinum*.

Prerequisites

- Obtain a replacement PSU from HGST.
- Ensure that the other PSU connected to this node is working, before pulling out the defective one.

Required Tools

- Ladder

Time Estimate: 5 minutes.

To replace a PSU proceed as follows:

1. Enable the location LED on the Controller Node.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Controller Nodes**.
   b) Select the correct node.
   c) In the **Commands** pane, click **Location LED On**.
   A blue LED on its front and back panels is now blinking.
2. Go to the rack and identify the correct chassis by the blinking blue LED on its front and back panels.
3. On the back of the rack, identify the failed PSU on the node identified in the previous step.

   In the image below, the PSUs are labeled `P1` and `P2`. The faulty PSU has an amber LED illuminated.

   Figure 17: Controller Node, Back, with PSU Status LEDs Highlighted

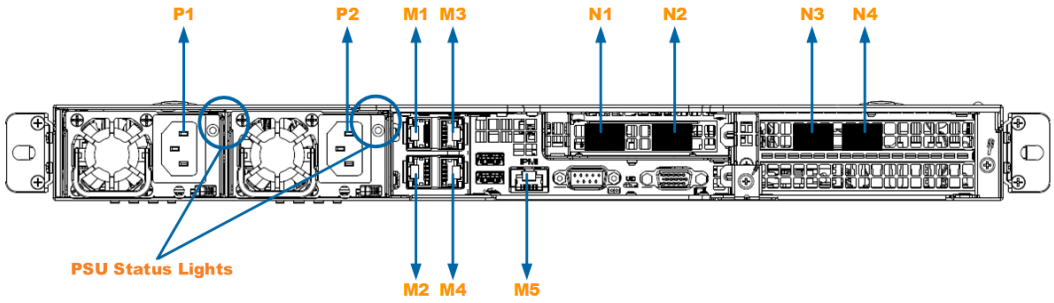

4. Remove the failed PSU.
   a) Disconnect the power cord from the failed PSU **only**.
   b) Push the red release tab towards the power connector of the failed PSU.
   c) Pull the PSU out of the node using the grab handle.
5. Install the replacement PSU.
   a) Push the replacement PSU into the Controller Node, and listen for a click.
   b) Connect the power cord to the replacement PSU.
6. Disable the location LED on the Controller Node.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Controller Nodes**.
   b) Select the desired Controller Node.
   c) In the **Commands** pane, click **Location LED Off**.

3.6 SFP+ Cable Replacement Procedure

The SFP+ cable connected to a Controller Node is a *10G SFP+ SFP+ 3m DAC Cable, 30AWG*.

Prerequisites

None.

Required Tools

None.

To replace an SFP+ cable, proceed as follows:

1. Identify the faulty SFP+ cable.
   a) Go to the back of the rack.
   b) Look for the Controller Node with a faulty SFP+ cable.
      Each Controller Node has an SFP+ cable with an SFP+ optical transceiver attached in ports **N1** and **N4** as shown in the figure below. The faulty cable has an amber LED illuminated on its metal connector.

      Figure 18: Controller Node SFP+ Ports

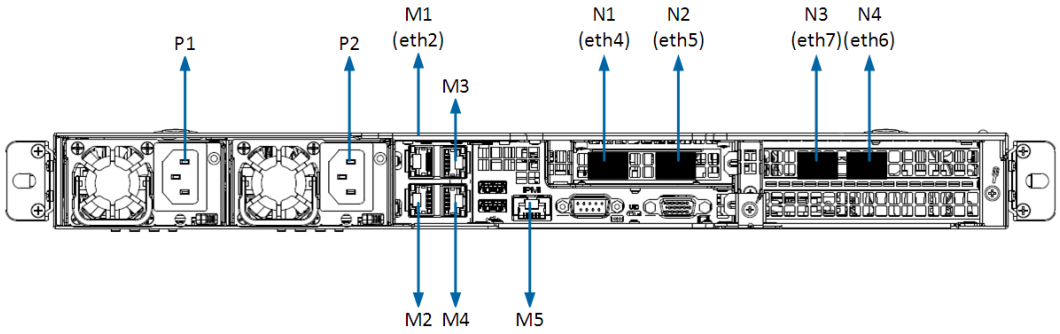

2. Remove the faulty SFP+ cable from the Controller Node.
   a) Pull the SFP+ optical transceiver out of the Controller Node.
   b) Unlatch the faulty SFP+ cable from the SFP+ optical transceiver by pulling very gently on its pull tab. Once the latch is disengaged, the cable is loose.
   c) While pulling gently on its pull tab, grasp its metal connector or cord to pull the cable out of its port.

   Warning: Do not pull the cable out by its pull tab, because the pull tab might break.

3. Install the new SFP+ cable.
   a) Connect the new SFP+ cable to the SFP+ optical transceiver you removed from the Controller Node port.
      The cable is reseated properly (the latch is engaged) when you hear a click.
   b) Push the SFP+ optical transceiver into the same Controller Node port.
4. Verify that the amber LED on the SFP+ optical transceiver is off.

4 Storage Node Replaceable Units

Topics:

This section provides replacement procedures for the following parts in a Storage Node:

- Chassis
- HDD
- PSU

4.1 Warnings

Caution:

- Opening or removing the system cover when the system is powered on may expose you to a risk of electric shock.
- When replacing items from the inside of the chassis, ensure that you take precautions to prevent electrostatic discharge (ESD).
- A Storage Node weighs about 43lbs. Ensure sufficient manpower to handle it safely.

4.2 Chassis Replacement Procedure

The Storage Node chassis is a *SuperMicro UP 1U Server, 1018*. Replacing the chassis replaces its NICs, CPU, memory, motherboard, and fans, **but not its disks**.

Prerequisites

- Obtain a replacement Storage Node chassis from HGST.
- Obtain the virtual IP address of the Management Node.
- Obtain the admin password for the CMC.
- Obtain the root password.
- Fill in as much of the work table as possible before starting this procedure.

Required Tools

- Ladder
- Long Phillips-head screwdriver

Time Estimate: 3 hours.

Figure 19: Overview of Storage Node Chassis Replacement

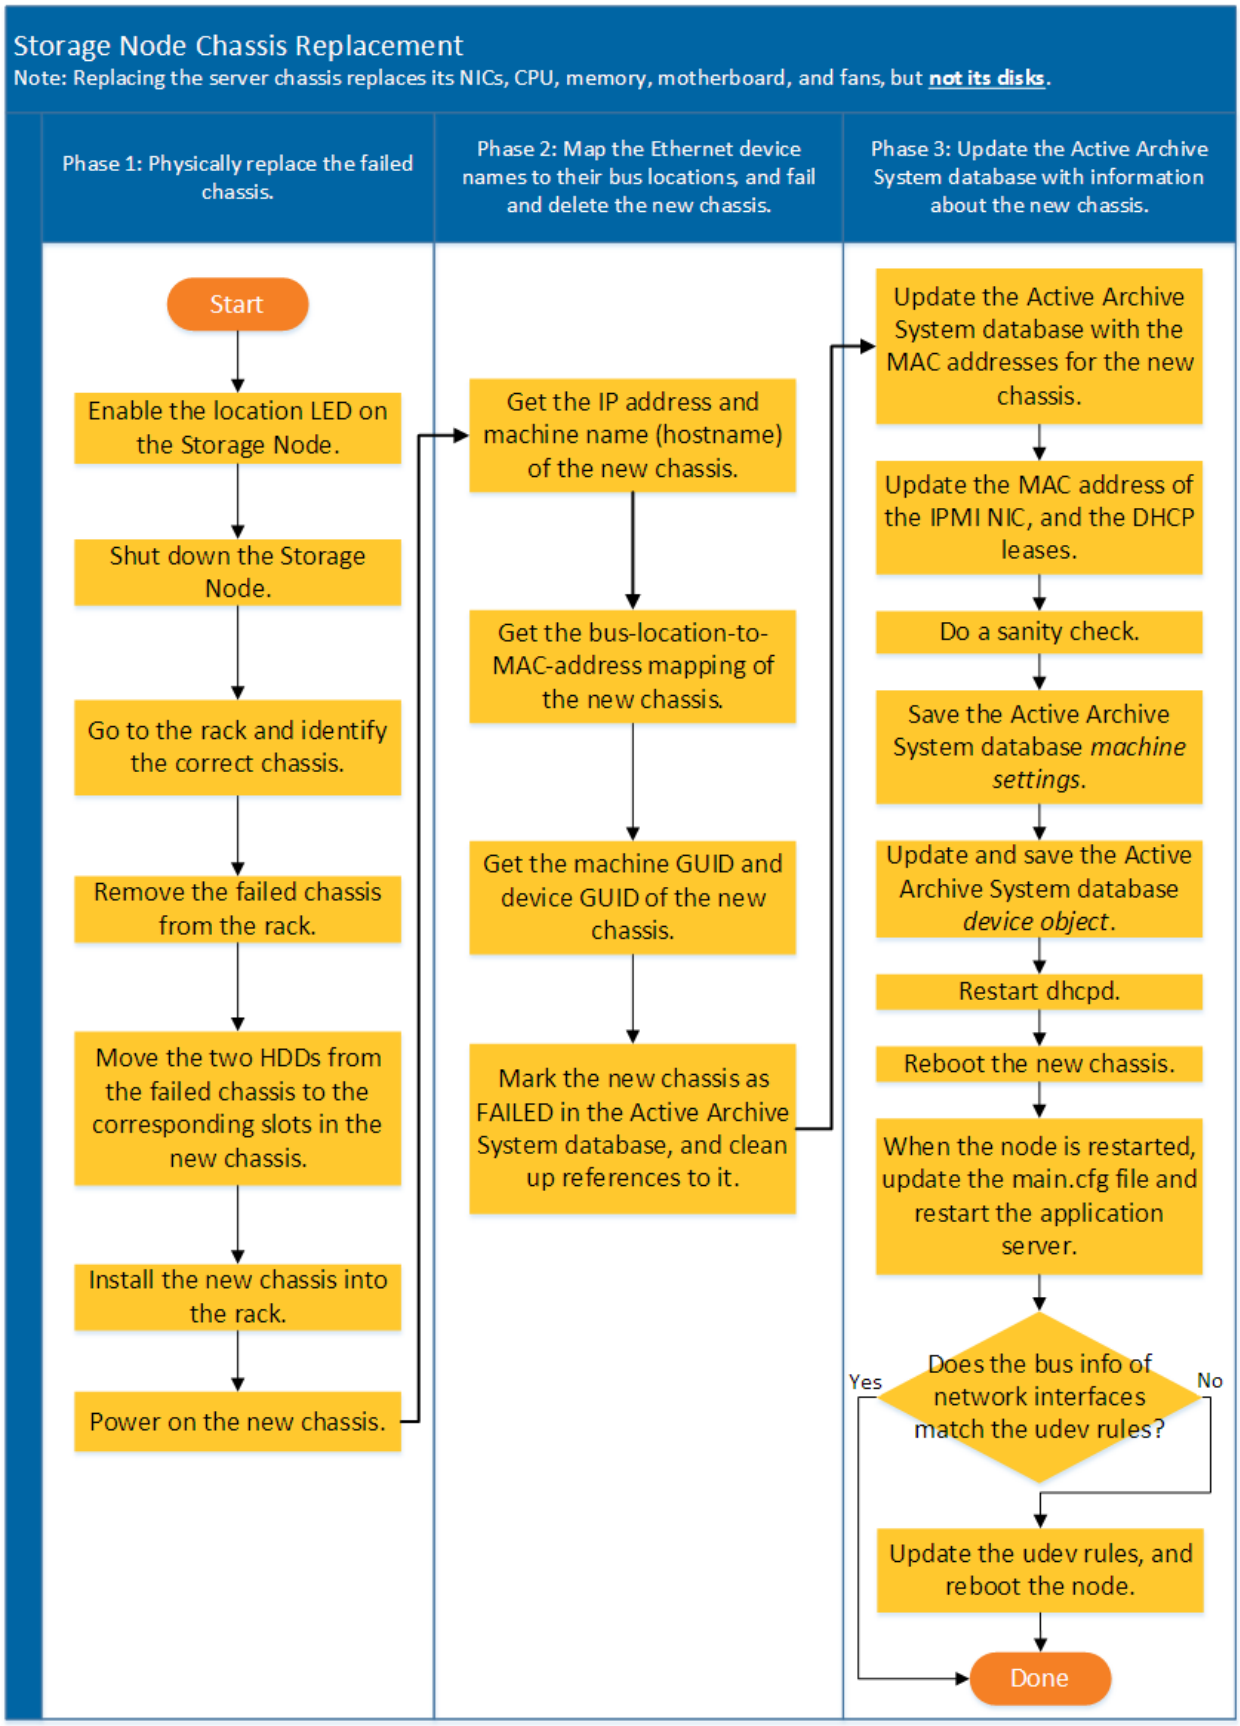

A work table is provided at the end of this section for your convenience, to store all of the information needed for a chassis replacement.

To replace a Storage Node chassis, proceed as follows:

1. Enable the location LED on the Storage Node.

Enable the location LED on the Storage Node.

Tip: The sample outputs shown for this procedure are from Storage Node 6.

a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Storage Nodes**.
b) Select the correct node.
c) In the **Commands** pane, click **Location LED On**.

2. Shut down the Storage Node.

Shut down the Storage Node from the CMC.

Note: Save the node's hostname in your worktable under **Original Hostname of Node**.

a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Storage Nodes**.
b) Select the desired Storage Node.

Figure 20: A Storage Node Pane in the CMC

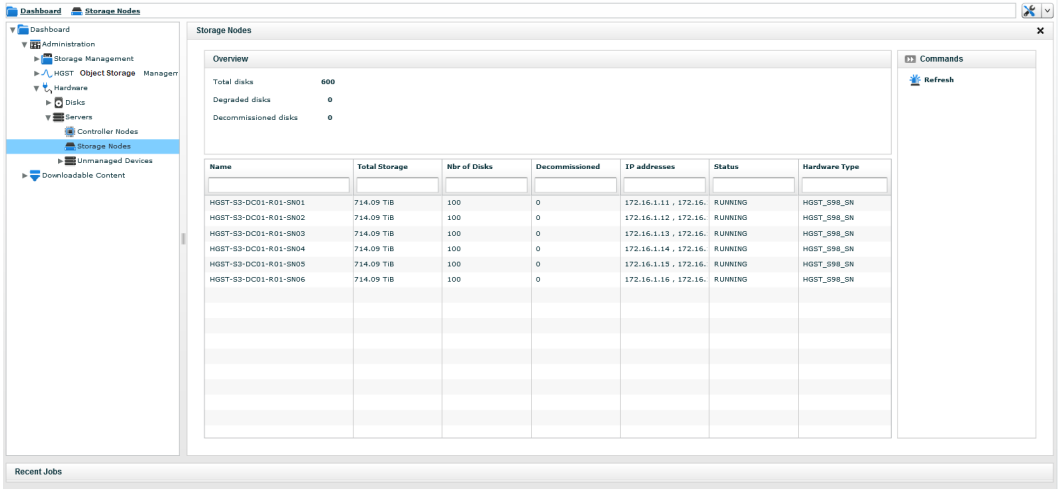

c) In the **Commands** pane, click **Shutdown**.

Figure 21: The Shutdown Button in the Commands Pane

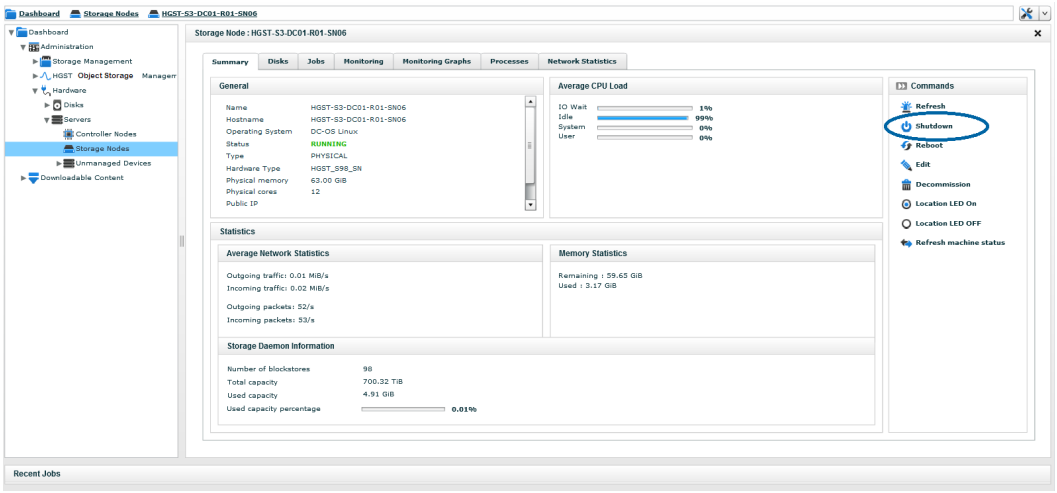

d) Wait for the **Status** field to change to **DONE**.

Warning: Even if all LEDs are off, you must still wait until the CMC shows **DONE** in the **Status** field.

All I/O to the Storage Enclosure Basic attached to this Storage Node is now quiesced.

3. Go to the rack and identify the correct chassis.

Go to the rack and identify the correct chassis by the blinking blue LED on its front and back panels.

4. Remove the failed chassis from the rack.

Remove the failed chassis from the rack.

a) At the front of the chassis, loosen the rack mounting screws.

b) At the back of the chassis, disconnect the three network cables connected to the ports labeled as follows in the image below.

Note: Check that the cables are labeled correctly, so that you can put them back in the same order.

---

a. **N1** (Public network #1)
b. **N2** (Public network #1)
c. **M2** (IPMI)
d. **S1** (Storage network #1)

e. **S2** (Storage network #2)

Figure 22: Storage Node, Back

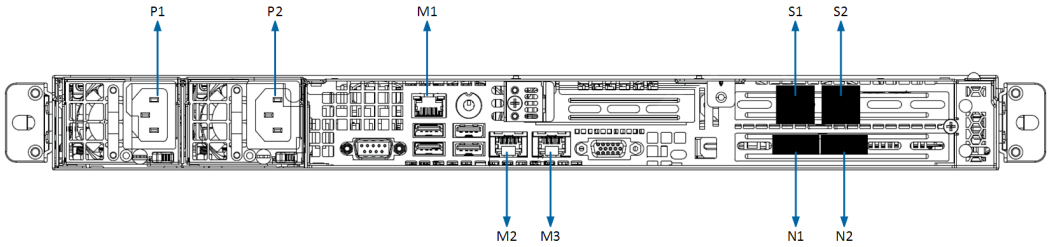

c) At the back of the chassis, disconnect the miniSAS cables.

Unlatch the miniSAS cables by pulling very gently on their pull tabs. While pulling gently on its pull tab, grasp its metal connector or cord to pull the cable out of its port.

Warning: Do not pull the cable out by its pull tab, because the pull tab might break.

Observe the amber LED on the paired Storage Enclosure Basic indicating loss of connection.

d) At the back of the chassis, disconnect the two power cords.

In the image above, the power cords are connected to the PSUs labeled `P1` and `P2`.

e) At the front of the chassis, slowly slide the chassis out until you reach the *pull-safety* at the midway point (you will hear a soft clicking sound, and feel the chassis "catch" on the rails).

f) Disengage the *pull-safety* on both sides of the chassis and slide it out until the split line of the two top covers.

Push the *pull-safety* on one side up, and the *pull-safety* on the other side down.

g) Continue to slowly slide the chassis out until you reach the *pull-safety* at the end point, and disengage it as you did the earlier one.

h) Safely unmount the chassis from the rack and place it on a table.

Caution: A Storage Node chassis weighs about 43lbs. Ensure that you have sufficient manpower to handle it safely.

Warning: Once you pull the chassis past the *pull-safety*, do not leave it hanging in the rack. Otherwise, the rack rails may be damaged permanently.

5. Move the two HDDs from the failed chassis to the corresponding slots in the new chassis.

Move the two HDDs from the failed chassis to the **exact corresponding slots** in the new chassis.

Tip: Write down the disk serial number and slot location so that you can double-check that each disk is seated in the correct slot post installation into the new chassis.

a) Remove each disk from its slot in the front bay of the failed chassis.

b) Install the disk into the corresponding slot in the new chassis.

Figure 23: Storage Node, Front

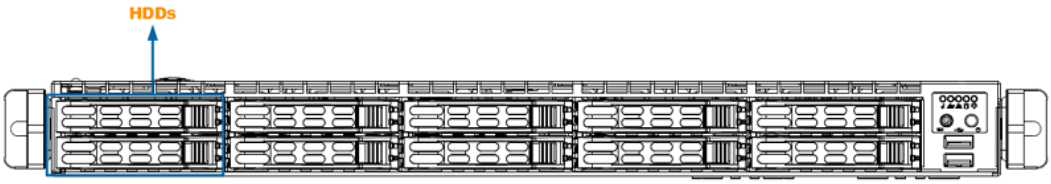

6. Install the new chassis into the rack.

Install the new chassis into the rack.

a) Mount the new chassis onto the rack slides and slide it into the rack.

> **Caution:** Mounting the chassis is a two person task.

b) Tighten the rack mounting screws to secure the chassis to the rack.

c) Reconnect the three network cables to the chassis ports.

The network cables are labeled.

a. Connect the cable labeled `SNx.N1.SW1.Nxx` to the port labeled **N1** in the image below.

b. Connect the cable labeled `SNx.N2.SW2.Nxx` to the port labeled **N2** in the image below.

c. Connect the cable labeled `SNx.M2.SW1.Nxx` to the port labeled **M2** in the image below.

Figure 24: Storage Node, Back

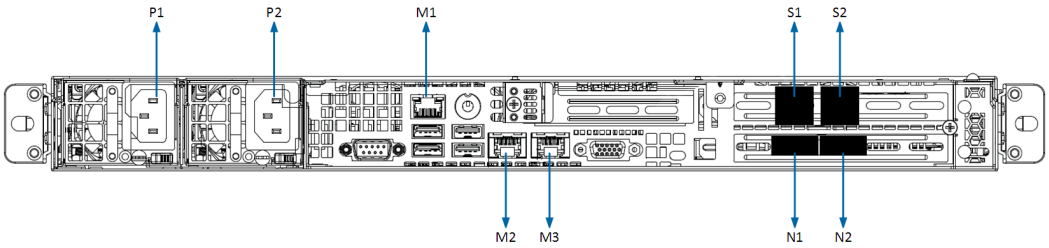

d) Reconnect the miniSAS cables.

The miniSAS cables are labeled.

- Connect the cable labeled `SNx.S1.DAx.SA` to the port labeled **S1** in the image below.
- Connect the cable labeled `SNx.S2.DAx.SB` to the port labeled **S2** in the image above.

e) Reconnect the power cords.

7. Get the MAC address of the new chassis IPMI NIC from the BIOS.

a) Connect a VGA monitor and USB keyboard to the new chassis.

b) Power on the new chassis.

The power button is located on the chassis front control panel.

c) At power up, press `Del` to enter into the system BIOS.

d) In the system BIOS, navigate to **IPMI** > **BMC Network Configuration** > **.**

e) Record **Station MAC Address** in your work table under **IPMI MAC Address of Node**.

f) Exit the BIOS without saving any changes by pressing the `ESC`.

The boot process continues.

g) Disconnect the VGA monitor and USB keyboard from the new chassis.

8. Get the IP address and machine name (hostname) of the new chassis.

Get the IP address and machine name (hostname) of the new chassis.

a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Unmanaged Devices** > **Uninitialized**.
The new chassis appears in the list of uninitialized devices. This indicates that it has started successfully.
b) Write the value of **Name** into your work table, under **Temporary Hostname of Node**.
c) Write the value of **Name** without the `PM-` prefix into your work table, under **MAC Address of Node**.
d) Write the **IP address** into your work table, under **Temporary IP Address of Node**.

Figure 25: Uninitialized Nodes

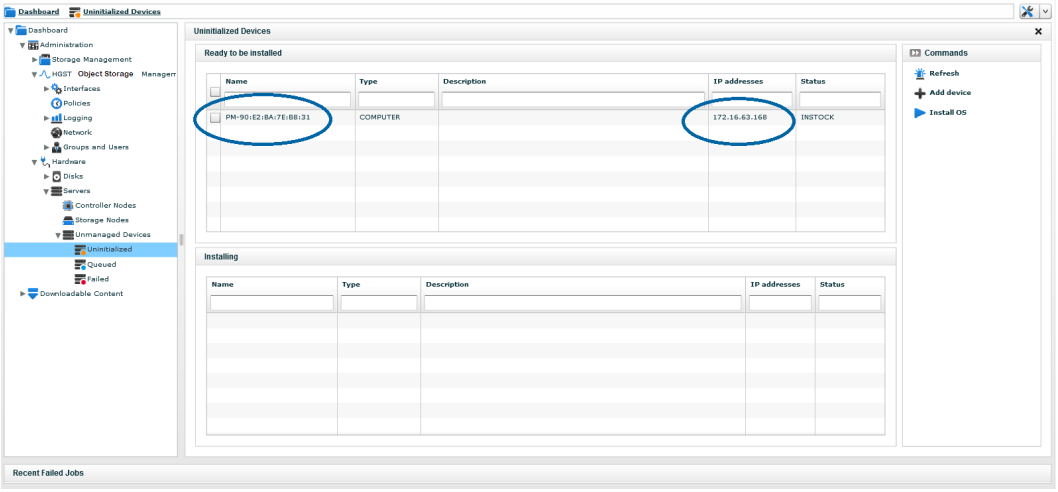

9. Get the bus-location-to-MAC-address mapping of the new chassis.

Get the bus-location-to-MAC-address mapping of the new chassis.

a) From the Management Node, open an SSH session to the **new IP address of the Storage Node** obtained in the previous step.
Log in with username `root` and password `rooter`.
b) At the Linux prompt, run the following command:

```
for add in `ls /sys/devices/pci*/*/*/net/*/address`; do echo $add; cat
 $add; done
```

The output of this command is similar to the example below.

```
/sys/devices/pci0000:00/0000:00:02.0/0000:02:00.0/net/eth1/address
90:e2:ba:7e:b8:30
/sys/devices/pci0000:00/0000:00:02.0/0000:02:00.1/net/eth3/address
90:e2:ba:7e:b8:31
/sys/devices/pci0000:00/0000:00:1c.4/0000:07:00.0/net/eth0/address
0c:c4:7a:33:38:10
/sys/devices/pci0000:00/0000:00:1c.4/0000:07:00.1/net/eth2/address
0c:c4:7a:33:38:11
root@nfsROOT:~#
```

The output of this command shows the serial bus path (for example, `0000:02:00.1`) and the new MAC address (for example, `90:e2:ba:7e:b8:31`).

Tip: As an alternative to the command above, you can use the command below to print only the serial bus paths and MAC addresses in uppercase.

```
for add in `ls /sys/devices/pci*/*/*/net/*/address`; do echo -en
 "`echo $add|sed 's/\// /g' | awk '{print $5}'`\t"; cat $add|tr 'a-f'
 'A-F'; done
```

c) Fill in the serial bus path in ascending order in the **Serial Bus Path** column of the work table.

d) Fill in the MAC address corresponding to the serial bus path in ascending order in the **MAC Address on the New Chassis** column of the work table.

For the sample output from the step above, the work table would look like this:

Serial Bus Path	MAC Address on the New Chassis	Ethernet Port Name
0000:02:00.0	90:e2:ba:7e:b8:30	eth0
0000:02:00.1	90:e2:ba:7e:b8:31	eth1
0000:07:00.0	0c:c4:7a:33:38:10	eth2
0000:07:00.1	0c:c4:7a:33:38:11	eth3

Table 7: Work Table with Sample MAC Addresses and Serial Bus Paths

e) Close the SSH session to the Storage Node.
You are now back in the SSH session to the Management Node.

10. Get the machine GUID and device GUID of the new chassis.

Get the machine GUID and device GUID of the new chassis.

a) On the Management Node, start the Q-Shell:

```
/opt/qbase3/qshell
```

b) Create a cloudAPI connection.

```
cloudapi = i.config.cloudApiConnection.find('main')
```

c) Retrieve the machine GUID for the **new chassis**, using the value of **Temporary Hostname of Node** in **uppercase**, from the work table, for *hostname_of_new_node* in the command below:

```
machine_guid = cloudapi.machine.find(name='hostname_of_new_node')['result'][0]
```

For example,

```
machine_guid = cloudapi.machine.find(name='PM-90:E2:BA:7E:B8:31')['result'][0]
```

d) Retrieve the device GUID using the machine GUID you obtained from the previous step.

```
dg = cloudapi.machine.list(machineguid=machine_guid)['result'][0]['deviceguid']
```

e) Sanity check: print the value of `dg`.
For example,

```
dg
'd951f6d9-7104-470d-8c97-ecf52d57c7b5'
```

11. Mark the new chassis as FAILED in the Active Archive System database, and clean up references to it.

Mark the new chassis as **FAILED** in the Active Archive System database, and clean up references to it.

The Active Archive System created a new **INSTOCK** node in its database for the new chassis. If you do not mark the new chassis as **FAILED** in the database, you are in effect adding a new node rather than replacing an existing node's chassis. Therefore, you must remove the **INSTOCK** node by following the steps below.

a) Mark the new chassis as **FAILED** in the Active Archive System database:
Execute this command on the Management Node:

```
cloudapi.device.updateModelProperties(dg, \
status=str(q.enumerators.devicestatustype.FAILED))
```

The new chassis now appears under the **FAILED** list in the CMC, and is removed from the **Unmanaged Devices** list.

Figure 26: The New Chassis Appears Under the FAILED List in the CMC

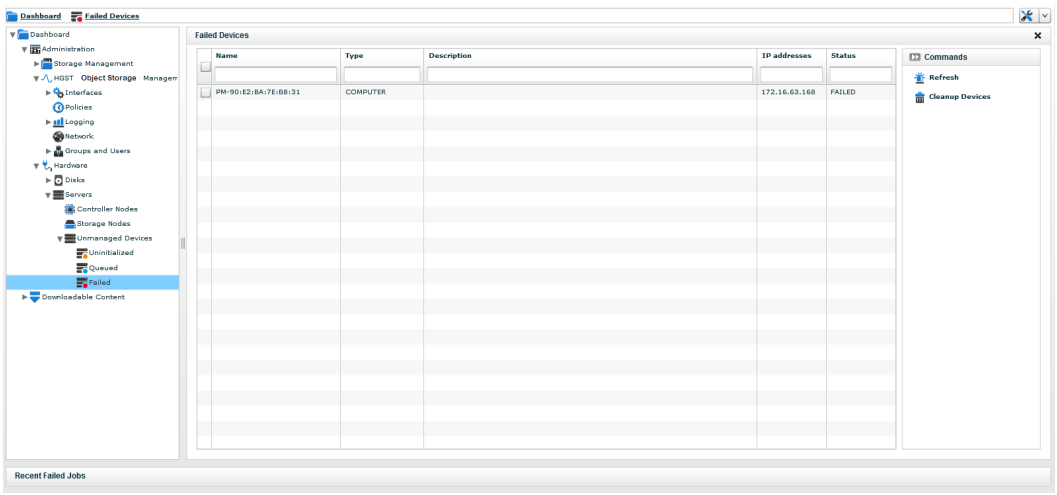

b) From the Management Node, clean up references to the new chassis in the Active Archive System database.

In the command below, replace *MAC_ADDRESS* with the value you wrote in the work table for **MAC Address of Node**.

> **Important:** Use capital letters for the MAC address.

```
q.amplistor.cleanupMachine('MAC_ADDRESS')
```

For example,

```
In [14]: q.amplistor.cleanupMachine('90:E2:BA:7E:B8:31')
Out[14]: True
```

This command takes about 10 seconds to complete.

c) Do a sanity check.

Refresh the screen by clicking **Refresh** in the **Commands** pane. Check that the new chassis is no longer in the **FAILED** list.

12. Update the Active Archive System database with the MAC addresses for the new chassis.

Update the Active Archive System database with the MAC addresses for the new chassis.

a) From the Management Node, create a cloudAPI connection.

```
cloudapi = i.config.cloudApiConnection.find('main')
```

b) From the Management Node, get the machine GUID using your work table value for **Original Hostname of Node**.

> **Note:** Use upper case for *HOSTNAME_OF_OLD_NODE*.

```
machine_guid = cloudapi.machine.find(name='HOSTNAME_OF_OLD_NODE')\
['result'][0]
```

For example,

```
machine_guid = cloudapi.machine.find(name='HGST-S3-DC01-R01-SN06')['result'][0]
```

c) From the Management Node, get the machine object.

```
machine = cloudapi.machine.getObject(machine_guid)
```

d) Display all the Ethernet port names (`ethN`) that are registered:
For example,

```
In [3]: machine = cloudapi.machine.getObject(machine_guid)
In [4]: print machine.nics[0].name
eth1
In [5]: print machine.nics[1].name
eth3
In [6]: print machine.nics[2].name
eth2
In [7]: print machine.nics[3].name
eth0
In [8]: print machine.nics[4].name
BMC
```

e) Write the *index* of the above `machine.nics[index].name` value into the work table in column **NIC Array ID**, in the row corresponding to `ethN`.

For the sample output from the step above, the work table would look like this:

Serial Bus Path	MAC Address on the New Chassis	Ethernet Port Name	NIC Array ID
0000:02:00.0	90:e2:ba:7e:b8:30	eth0	3
0000:02:00.1	90:e2:ba:7e:b8:31	eth1	0
0000:07:00.0	0c:c4:7a:33:38:10	eth2	2
0000:07:00.1	0c:c4:7a:33:38:11	eth3	1
IPMI	See **IPMI MAC Address of Node** in the work table.	BMC	4

Table 8: Work Table with Sample Ethernet Port Names and NIC Array IDs

f) Update the database entry for `machine.nics[N].hwaddr` with the corresponding MAC address for `ethN` from your work table.

> **Important:** Use capital letters for the MAC address.
>
> ```
> machine.nics[0].hwaddr = 'NEW_MAC_ADDRESS_FOR_ETHN'
> ```

For example,

```
machine.nics[0].hwaddr = '90:E2:BA:7E:B8:31'
machine.nics[1].hwaddr = '0C:C4:7A:33:38:11'
machine.nics[2].hwaddr = '0C:C4:7A:33:38:10'
machine.nics[3].hwaddr = '90:E2:BA:7E:B8:30'
```

g) Update the database entry for `machine.nics[4].hwaddr` with the corresponding IPMI MAC address from your work table, under **MAC Address of Node**.

> **Important:** Use capital letters for the MAC address.

For example,

```
machine.nics[4].hwaddr = '0C:C4:7A:36:8B:12'
```

13. Update the MAC address of the IPMI NIC, and the DHCP leases.

Update the MAC address of the IPMI NIC, and the DHCP leases.

a) Log into the CMC.

b) Navigate to the CMC's view of the node whose chassis you have replaced.

c) Select to that node's **Summary** tab.

d) Write the IPMI IP address, as shown in the **General** section, in your work table, under **IPMI IP Address of Node**.

Figure 27: The Old IPMI IP Address of the Node

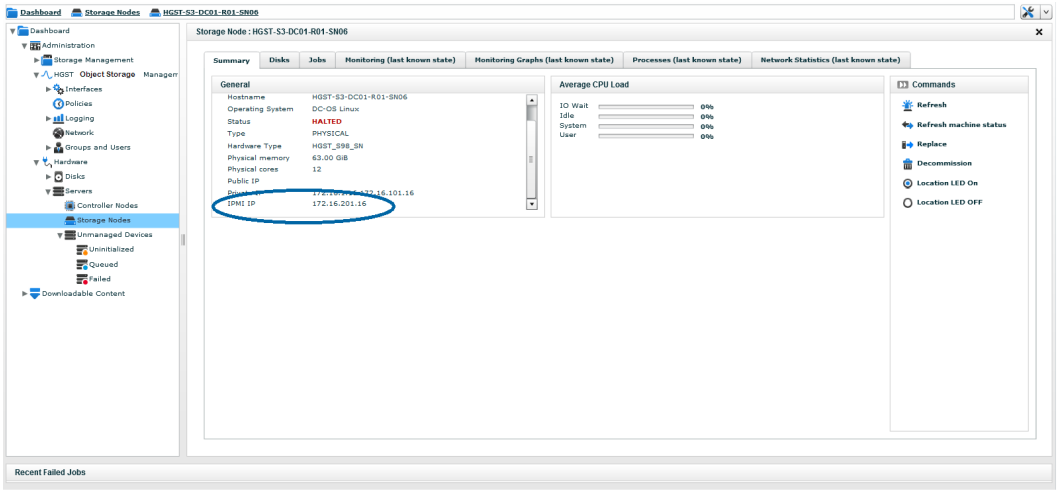

e) Leave the current SSH session as is. Open a new SSH session on the Management Node.

f) Open `/opt/qbase3/cfg/dhcpd/dhcpd.leases` with your text editor.

g) Search for the IPMI IP address (obtained in substep d) in the file.

The section containing the IPMI IP address looks like the following example:

```
host 457f495a-80b7-4125-862b-5f87d9121cfa {
  dynamic;
  hardware ethernet 0c:c4:7a:36:8b:12:;
  fixed-address 172.16.201.16;
  group "pmachines";
}
```

h) Change the `hardware ethernet` value to the new IPMI MAC address **in lowercase** from your work table, under **IPMI MAC Address of Node**.

Note: Use lowercase only.

For example,

```
hardware ethernet 0c:c4:7a:36:8b:12:;
```

i) Save and close the file.

j) Exit the new SSH session.

14. Do a sanity check.

In your **previous** SSH session, do a sanity check to verify that you have updated the new MAC addresses correctly. Compare the output of the command below to your work table.

```
In [9]: for nic in machine.nics: nic.name; nic.hwaddr
   ...:
Out[9]: 'eth1'
Out[9]: '90:E2:BA:7E:B8:31'
Out[9]: 'eth3'
Out[9]: '0C:C4:7A:33:38:11'
Out[9]: 'eth2'
Out[9]: '0C:C4:7A:33:38:10'
Out[9]: 'eth0'
Out[9]: '90:E2:BA:7E:B8:30'
Out[9]: 'BMC'
Out[9]: '0C:C4:7A:36:8B:12'
```

15. Save the Active Archive System database *machine settings*.

Save the Active Archive System database *machine settings*.

```
q.drp.machine.save(machine)
```

16. Update and save the Active Archive System database *device object*.

Update and save the Active Archive System database *device object*.

a) Get the device object.

```
device = cloudapi.device.getObject(machine.deviceguid)
```

b) Update the MAC address of the chassis with the value you saved in the work table under **MAC Address of Node**.

Note: Use capital letters for the MAC address.

```
device.nicports[0].hwaddr = 'NEW_MAC_ADDRESS'
```

For example,

```
In [12]: device.nicports[0].hwaddr='90:E2:BA:7E:B8:31'
```

c) Save the device object.

```
q.drp.device.save(device)
```

17. Restart dhcpd.

Restart dhcpd.

```
In [14]: q.manage.dhcpd.restart()
 Stopping dhcpd...
 dhcpd is halted
 Starting dhcpd...
 dhcpd is running
```

18. Reboot the new chassis.

Reboot the new chassis.

a) From the Management Node, open an SSH session to the new chassis using the IP address from the work table, under **IP Address of Node**.
The Linux prompt appears.

b) At the Linux prompt, run the following command:

```
reboot
```

Once rebooted, if you log into the CMC, you can see that the chassis now has the correct hostname as shown in the figure below.

Figure 28: The Node with a Rebooted Chassis as Seen on the CMC

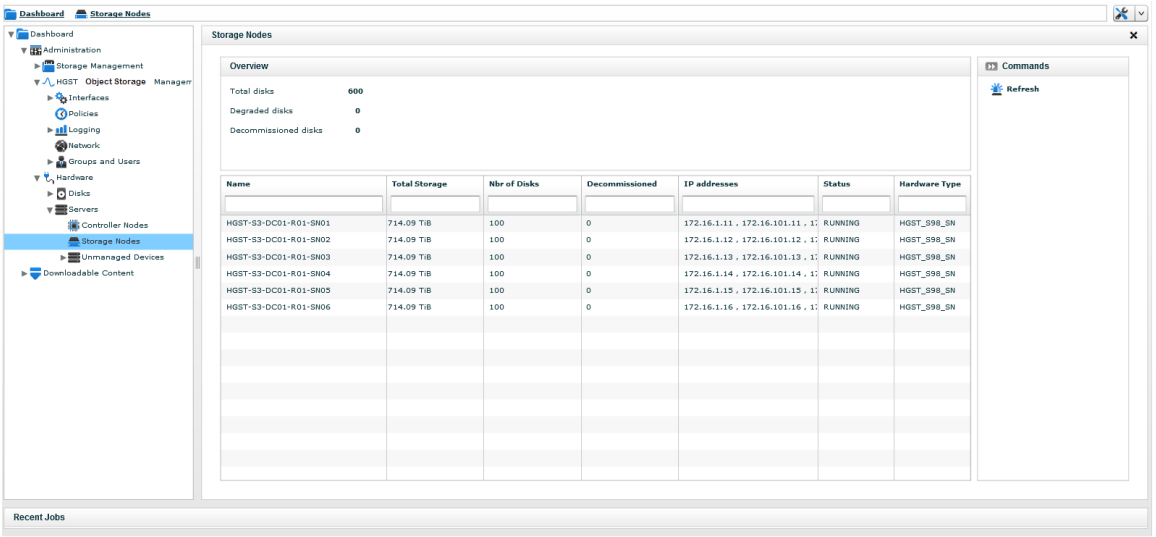

19. When the node is restarted, update the main.cfg file and restart the application server.

When the node is restarted, update the `main.cfg` file and restart the application server.

a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Storage Nodes**.
A list of Storage Nodes appears in the CMC.

b) Click the Storage Node whose chassis you have just replaced.

Identify the correct Storage Node by its hostname: it now matches the **Original Hostname of Node** value you recorded in the worktable. This value is typically of the format *SystemID*-DC*nn*-R*nn*-SN*nn*.

c) Identify the IP addresses listed in the **Private IP** field.

d) Open an SSH session to the Management Node, and exit the OSMI menu.
The Linux prompt appears.

e) Open an SSH session to the Storage Node, using any one of the IP addresses you obtained from substep c.
The Linux prompt appears.

f) At the Linux prompt on the Storage Node, open the file `/opt/qbase3/cfg/qconfig/main.cfg` with your text editor.
The file has a section that looks like this:

```
[main]
lastlogcleanup = 1428960577
domain = somewhere.com
nodetype = STORAGENODE
nodename = 90E2BA7EB831
logserver_loglevel = 6
logserver_port = 9998
logserver_ip = 127.0.0.1
qshell_firstrun = False
machineguid = fc635662-5247-45b1-ab66-d0abe8e60712
```

g) Replace the value after `nodename =` with the new MAC address from your work table, under **MAC Address of Node**.

> **Note:** The MAC address must be **in uppercase and without colons**. For example, `00:25:90:3B:C1:72` must be typed as `0025903BC172`.

h) Save and close the configuration file.

i) Start the Q-Shell.

```
/opt/qbase3/qshell
```

j) Restart the application server:

```
In [1]: q.manage.applicationserver.restart()
 Restarting applicationserver Applicationserver...
 Stopping applicationserver Applicationserver...
 Applicationserver is still running, waiting for 5 more seconds
 Applicationserver is still running, waiting for 4 more seconds
 Starting applicationserver Applicationserver...
```

k) Exit the Q-Shell.

```
quit()
```

20. Update the udev rules, and reboot the node.

Verify that the bus information of the network interfaces matches the `udev` rules.

a) Run the following command:

Tip: Check the hardware paths in the command below, as they might be different on the new chassis.

```
root@HGST-S3-DC01-R01-SN06: for add in `ls /sys/devices/pci*/*/*/net/*/
address`; do echo $add; cat $add; done
```

For example, the output of the above command looks like this:

```
/sys/devices/pci0000:00/0000:00:02.0/0000:02:00.0/net/eth0/address
90:e2:ba:7e:b8:30
/sys/devices/pci0000:00/0000:00:02.0/0000:02:00.1/net/eth1/address
90:e2:ba:7e:b8:31
/sys/devices/pci0000:00/0000:00:1c.4/0000:07:00.0/net/eth2/address
0c:c4:7a:33:38:10
/sys/devices/pci0000:00/0000:00:1c.4/0000:07:00.1/net/eth3/address
0c:c4:7a:33:38:11
```

Tip: As an alternative to the command above, you can use the command below to print only the serial bus paths and MAC addresses in uppercase.

```
for add in `ls /sys/devices/pci*/*/*/net/*/address`; do echo -en
 "`echo $add|sed 's/\// /g' | awk '{print $5}'`\t"; cat $add|tr 'a-f'
 'A-F'; done
```

b) Compare the output of the command above to the contents of the file `/etc/udev/rules.d/70-persistent-net.rules`.

For example, the contents of this file look like this:

```
root@HGST-S3-DC01-R01-SN06:~# cat /etc/udev/rules.d/70-persistent-net.rules
SUBSYSTEM=="net", ACTION=="add", KERNELS=="0000:02:00.0", KERNEL=="eth*",
 NAME="eth0"
SUBSYSTEM=="net", ACTION=="add", KERNELS=="0000:02:00.1", KERNEL=="eth*",
 NAME="eth1"
SUBSYSTEM=="net", ACTION=="add", KERNELS=="0000:07:00.0", KERNEL=="eth*",
 NAME="eth2"
SUBSYSTEM=="net", ACTION=="add", KERNELS=="0000:07:00.1", KERNEL=="eth*",
 NAME="eth3"
```

If they do not match, update `/etc/udev/rules.d/70-persistent-net.rules` to match the output of the command above, then reboot the node again.

```
reboot
```

Sanity check: you can observe that the Storage Enclosure Basic LEDs are now solid green. In addition, the CMC shows the node status as **RUNNING**.

Figure 29: Storage Node Status in the CMC

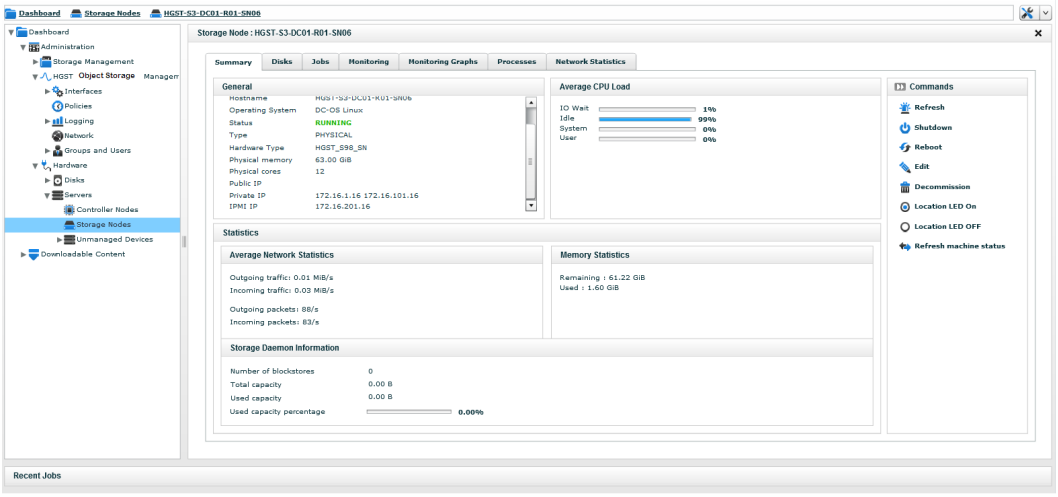

The chassis replacement procedure is done.

Warning: Be very careful when recording and updating MAC addresses. A mistake may render the new chassis unusable.

Item	Value
Virtual IP Address of the Management Node: Get this value as instructed in the *Administrator Interfaces* chapter in the *HGST Active Archive System Troubleshooting Guide*	
Original Hostname of Node: Get this value from the CMC **before** you shut down the failed node. This value is typically of the format `SystemID-DCnn-Rnn-SNnn`.	
Temporary IP Address of Node: The CMC displays this value **after** the new chassis is installed.	
Temporary Hostname of Node: The CMC displays this value **after** the new chassis is installed. This value is of the format `PM-MAC_ADDRESS`.	
MAC Address of Node:	
IPMI IP Address of Node:	
IPMI MAC Address of Node:	

Serial Bus Path	MAC Address on the New Chassis	Ethernet Port Name	NIC Array ID
`0000:01:00.0`		`eth0`	

Serial Bus Path	MAC Address on the New Chassis	Ethernet Port Name	NIC Array ID
0000:01:00.1		eth1	
0000:01:00.2		eth2	
0000:01:00.3		eth3	
IPMI			

Table 9: Work Table for Storage Node Chassis Replacement

4.3 Hard Disk Drive Replacement Procedure

The HDD on a Storage Node is an *HGST 500 GB SATA 6 Gb/sec drive*. It is hot swappable after being decommissioned in the CMC.

Prerequisites

- Decommission the faulty drive in the CMC. For more information, see *Managing Hardware* in the *HGST Active Archive System Administration Guide*.
- Obtain a replacement HDD from HGST.

Required Tools

- Ladder
- Long Phillips-head screwdriver

Time Estimate: 40 minutes.

To replace an HDD, proceed as follows:

1. Obtain details about the decommissioned disk from the CMC.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Disks** > **Decommissioned**.

   Figure 30: Decommissioned Disks in the CMC

   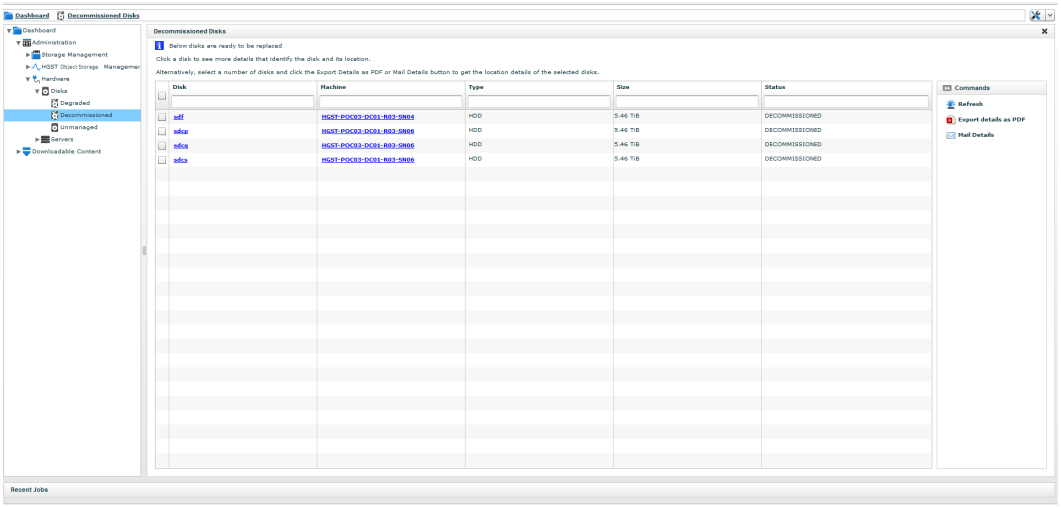

   b) Click the desired decommissioned disk.
   The *decommissioned disk details* are displayed.
   c) (Optional) Right-click anywhere in the decommissioned disk details, select **Print** > **Print to PDF**.
   The *decommissioned disk details* contains the following information, which you will need to refer to later:

   - The device name.
   - The model type and serial number.
   - A *drive map* showing the exact location of the disk (name of data center, rack, and node).
   - The node type.
   - The MAC addresses of the node.

- The current status of the node.

Figure 31: Drive Map Showing a Decommissioned Drive on a Storage Node

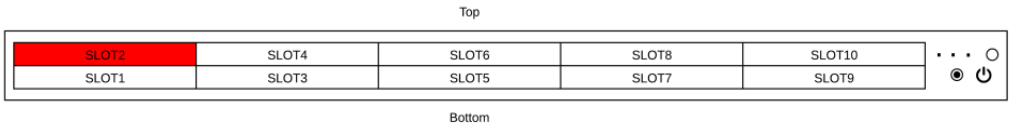

2. Enable the location LED on the Storage Node.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Storage Nodes**.
   b) Select the correct node.
   c) In the **Commands** pane, click **Location LED On**.
3. Go to the rack and identify the correct chassis by the blinking blue LED on its front and back panels.
4. Replace the decommissioned HDD.
   a) Identify the decommissioned HDD by using the drive map.
   The HDD is a front-bay drive.

Figure 32: Storage Node, Front

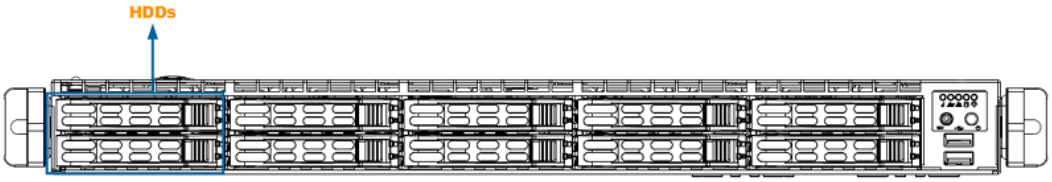

   b) Press the release button on the drive carrier of the decommissioned HDD to extend the *drive carrier handle*.

Figure 33: Removing a Drive Carrier

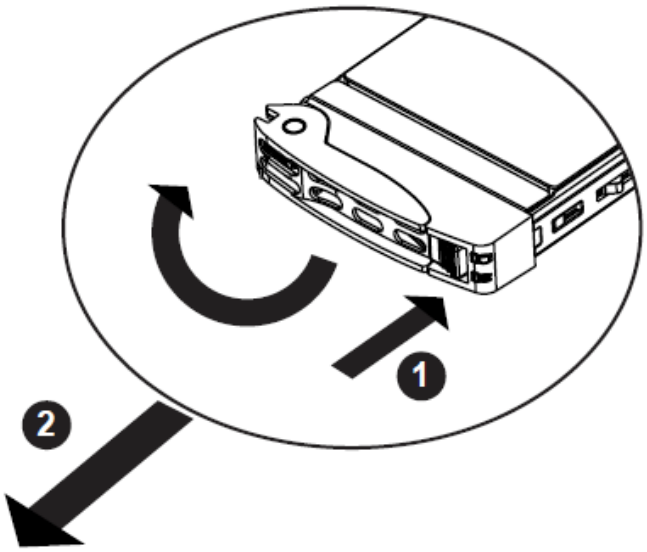

   c) Pull drive carrier out of the front bay using the drive carrier handle.
   d) Compare the serial number on the HDD to the serial number specified in the decommissioned disk details to confirm the that you have the correct HDD.
   e) Unscrew the drive carrier from the decommissioned HDD.
   f) Screw the drive carrier onto the replacement HDD.
   g) Install the replacement HDD into the same slot that the decommissioned HDD was using.

A blue LED will blink for a moment.

5. Disable the location LED on the Storage Node.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Storage Nodes**.
   b) Select the desired Storage Node.
   c) In the **Commands** pane, click **Location LED Off**.
6. Confirm that the Active Archive System correctly determines the purpose for the new disk.
   a) Wait 15 minutes.
   b) In the CMC, navigate to **Dashboard** > **Administration** > **HGST Object Storage Management** > **Logging** > **Events**.
   c) In the **Events** list, check to see that a new empty disk has been detected.
   d) In the **Jobs** list, check to see that an `Initializing new disk` job has been triggered.
      It may take about 2 minutes for the job to appear.
   e) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers**.
   f) Select the desired node.
   g) Select the **Disks** tab.
   h) Wait for the physical drive that has been replaced, as well as the logical disks, to change status from a red icon to a green icon.

      Note: The physical drive that has been replaced, as well as the logical disks, may take up to 40 minutes to change status.

   The `Initializing new disk` job has completed successfully when the number of degraded disks decreases by 1.
7. If the disk still shows up in the **Degraded** or **Unmanaged** list, you must manually specify the purpose of the new disk:
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Disks** > **Unmanaged**.
   b) Select the new disk, and in the **Commands** pane, click **Repurpose**.
   c) In the **Use As** field, select **Replacement Disk**.

      Note: You can only select **Replacement Disk** when there is a decommissioned disk. If there are no decommissioned disks, you can only select **Additional Disk** as the purpose for the disk.

   d) In the **Replacement For** field, select the decommissioned disk that you want to replace.
   e) Click **Next** to start the repurposing.
      An `Initializing new disks on` *`node_name`* job starts.

If You Replaced the Wrong Disk

See *Troubleshooting*.

4.4 Power Supply Unit Replacement Procedure

The power supply units (PSU) on a Controller Node are redundant hot-swappable *SuperMicro 1U 750w 74mm, Platinum*.

Prerequisites

- Obtain a replacement PSU from HGST.
- Ensure that the other PSU connected to this node is working, before pulling out the defective one.

Required Tools

- Ladder

Time Estimate: 5 minutes.

To replace a PSU proceed as follows:

1. Enable the location LED on the Storage Node.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Storage Nodes**.
   b) Select the correct node.
   c) In the **Commands** pane, click **Location LED On**.
2. Go to the rack and identify the correct chassis by the blinking blue LED on its front and back panels.
3. On the back of the rack, identify the failed PSU on the node identified in the previous step.

   In the image below, the PSUs are labeled `P1` and `P2`.

   Figure 34: Storage Node, Back, with PSU Status LEDs Highlighted

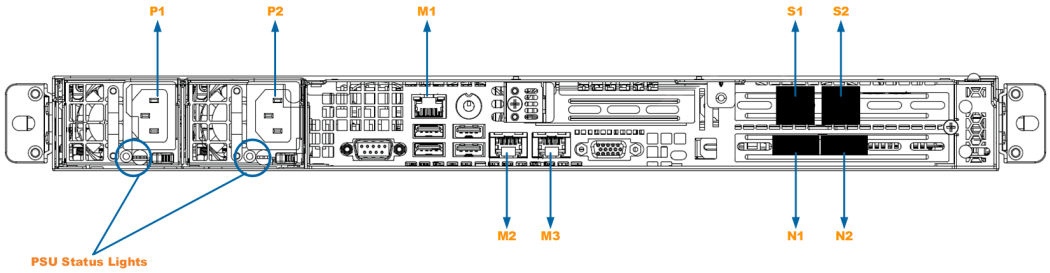

4. Remove the failed PSU.
   a) Disconnect the power cord from the failed PSU **only**.
   b) Push the red release tab towards the power connector of the failed PSU.
   c) Pull the PSU out of the node using the grab handle.
5. Install the replacement PSU.
   a) Push the replacement PSU into the Storage Node, and listen for a click.
   b) Connect the power cord to the replacement PSU.
6. Disable the location LED on the Storage Node.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Storage Nodes**.
   b) Select the desired Storage Node.
   c) In the **Commands** pane, click **Location LED Off**.

4.5 MiniSAS 3M or 6M Cable Replacement Procedure

The miniSAS cable connecting a Storage Node to its paired Storage Enclosure Basic is a *3M or 6M 12G miniSAS, miniSAS HD, 28 AWG*.

Prerequisites

None.

Required Tools

None.

To replace a miniSAS cable, proceed as follows:

1. Identify the faulty miniSAS cable.
   a) Go to the back of the rack.
   b) Look for the Storage Node with a faulty miniSAS cable.

   Each Storage Node has a miniSAS cable in ports **S1** and **S2** as shown in the figure below. The faulty cable has an amber LED illuminated on its metal connector.

   Figure 35: Storage Node MiniSAS Ports

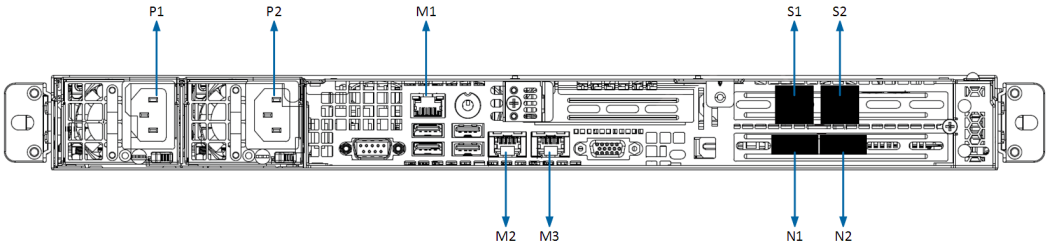

2. Remove the faulty miniSAS cable from the Storage Node.
   a) Unlatch the faulty miniSAS cable at its Storage Node end by pulling very gently on its pull tab.
   b) While pulling gently on its pull tab, grasp its metal connector or cord to pull the cable out of its port.

   Warning: Do not pull the cable out by its pull tab, because the pull tab might break.

3. Remove the faulty miniSAS cable from the **I/O canister** of the Storage Enclosure Basic.
   a) Disengage the latch on the faulty miniSAS cable at its Storage Enclosure Basic end by pulling very gently on its pull tab.

   Figure 36: Removing the MiniSAS Cables

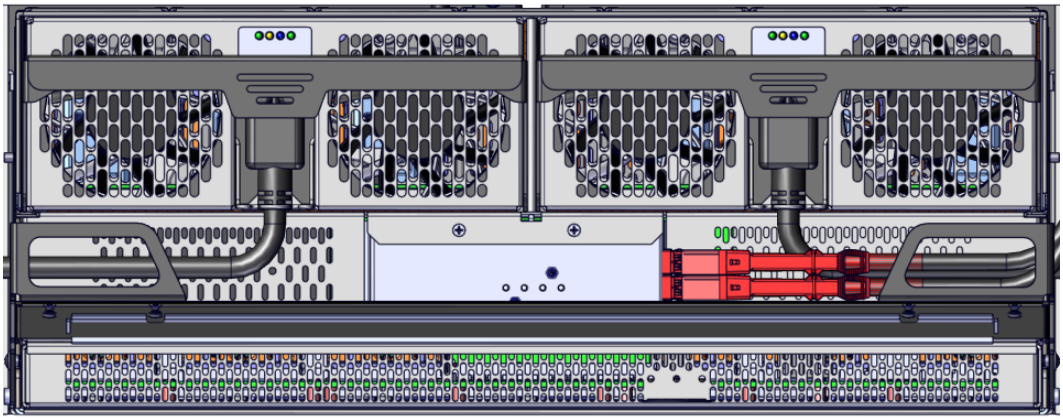

Note:

- The miniSAS cables are marked in red.
- Take note of which miniSAS cable came from which port to ensure that they are plugged in correctly when reassembling.

b) Once the latch is disengaged and the cable is loose, grasp its metal connector or cord (not its pull tab) to pull it out of its port.

4. Install the new miniSAS cable.

Tip: If you are replacing the 6M cable, install it over the top of the rack for ease of replacement.

a) Connect the new miniSAS cable into the same Storage Node port.

The cable is reseated properly (the latch is engaged) when you hear a click.

5. Verify that the amber LED on the new miniSAS connector is off.

5 Storage Interconnect Replaceable Units

Topics:

This section provides replacement procedures for the following parts in a Storage Interconnect:

- Chassis
- Fan
- PSU

5.1 Warnings

Caution: All data on the Active Archive System is unavailable during repair.

5.2 Switch Replacement Procedure

Prerequisites

- Obtain a replacement Storage Interconnect from HGST.
- Obtain the switch configuration file from HGST.

Required Tools

- Ladder
- Long Phillips-head screwdriver

Time Estimate: Unknown.

Important: If Storage Interconnect 1, (`SWA`, lower), is down, then IPMI is down.

To replace a Storage Interconnect, proceed as follows:

1. Edit the switch configuration file you obtained from HGST as follows:
   a) Modify the IP address for the syslog target to match your environment.
      The syslog target is the Management Node.
   b) Modify the IP address for VLAN to match your environment.
   c) Modify the time and time zone to match your environment.
2. Load your modified configuration file into the new Storage Interconnect as follows:
   a) Log into the new Storage Interconnect.
      Replace *`switchport_IP`* with the default IP address of the new Storage Interconnect:
      For switch 1 (lower), the default IP address is `192.168.123.123`.
      For switch 2 (upper), the default IP address is `192.168.123.123`.
      The default username is `admin`. There is no password.

```
telnet switchport_IP
```

b) Type `enable` to enter the Privileged EXEC command mode.

```
(Routing)>enable
```

c) Type `configure` to enter the Privileged CONF command mode.

```
(Routing)#configure
(Routing)  (Config)#
```

d) Copy the entire contents of your modified switch configuration file, and paste it into the telnet session.

e) Save the new configuration.

```
(Routing)#write memory
Are you sure you want to save? (y/n) y
Config file 'startup-config' created successfully.
Configuration Saved!
(Routing)#
```

This operation may take a few minutes. Management interfaces are not available during this time.

f) Verify that the settings on the switch match your configuration file.

```
(Routing)#show startup-config
```

3. Remove the two blanking plates from the top front of the rack.
4. Identify the faulty Storage Interconnect. It is a TOR switch.

Figure 37: Location of Storage Interconnect Switches: [1]

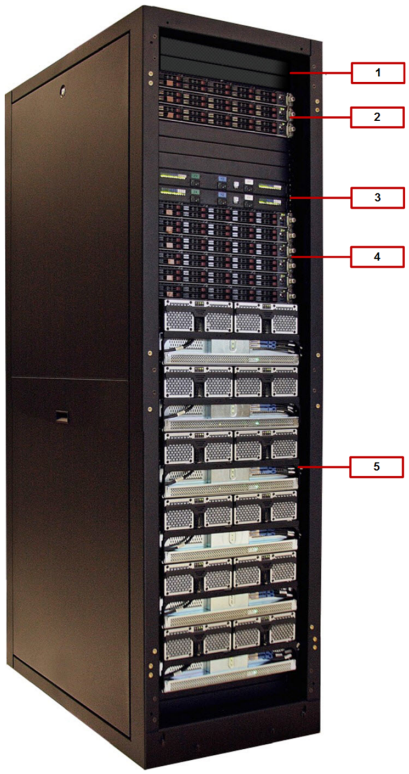

5. Power off the faulty Storage Interconnect:

   Disconnect the power cords on the back of the faulty Storage Interconnect.
6. Remove the faulty Storage Interconnect from the rack.
   a) Disconnect the labeled network cables from the Storage Interconnect.
   b) Unscrew the mounting screws from the Storage Interconnect.

c) From the front of the rack, push the Storage Interconnect out of the rack. Move around to the back of the rack and slide it out of the rack.

7. Install the new Storage Interconnect into the rack.

a) Slide the new Storage Interconnect into the rack.

b) Tighten new Storage Interconnect screws to secure it to the rack.

c) Reconnect the network cables to the new Storage Interconnect.

The network cables are labeled.

a. Connect the cable labeled `SN1.Nx.SWx.`**`N01`** to the port labeled **01** in the image below.

b. Connect the cable labeled `SN2.Nx.SWx.`**`N03`** to the port labeled **03** in the image below.

c. Connect the cable labeled `SN3.Nx.SWx.`**`N05`** to the port labeled **05** in the image below.

d. Connect the cable labeled `SN4.Nx.SWx.`**`N07`** to the port labeled **07** in the image below.

e. Connect the cable labeled `SN5.Nx.SWx.`**`N09`** to the port labeled **09** in the image below.

f. Connect the cable labeled `SN6.Nx.SWx.`**`N11`** to the port labeled **11** in the image below.

g. Connect the cable labeled `CN1.Nx.SWx.`**`N33`** to the port labeled **33** in the image below.

h. Connect the cable labeled `CN2.Nx.SWx.`**`N35`** to the port labeled **35** in the image below.

i. Connect the cable labeled `CN3.Nx.SWx.`**`N37`** to the port labeled **37** in the image below.

j. (**For SW1 only**) And so on (refer to the Signal Cabling Scheme below).

Figure 38: Switch Port Reservations

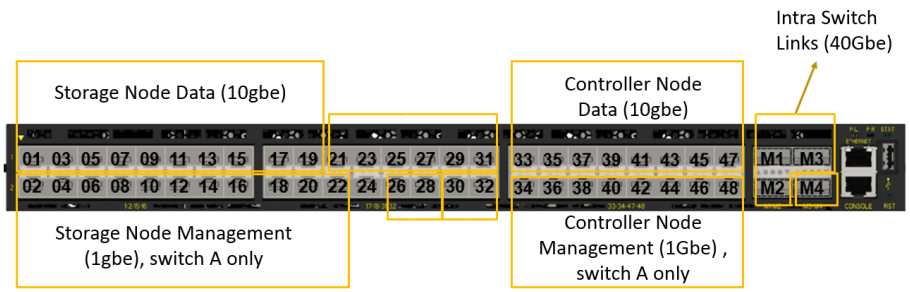

Figure 39: Signal Cabling Scheme

	Local Device	Local Port	Wrap Label (1"x1")	Peer Port	Peer Device	Length (m)	Cable P/N
Management Cables	CN3	MGT	CN3.M2:SW1.N38	Port 38	SW1	1	G400-10079-01
	CN2	MGT	CN2.M2:SW1.N36	Port 36	SW1	1	
	CN1	MGT	CN1.M2:SW1.N34	Port 34	SW1	1	
	SN6	MGT	SN6.M2:SW1.N12	Port 12	SW1	1.5	
	SN5	MGT	SN5.M2:SW1.N10	Port 10	SW1	1.5	
	SN4	MGT	SN4.M2:SW1.N08	Port 08	SW1	1.5	
	SN3	MGT	SN3.M2:SW1.N06	Port 06	SW1	1.5	
	SN2	MGT	SN2.M2:SW1.N04	Port 04	SW1	1.5	
	SN1	MGT	SN1.M2:SW1.N02	Port 02	SW1	1.5	
10G DAC Cables	CN3	NIC1	CN3.N2:SW1.N37	Port 37	SW1	1	R0893-C0005-01
		NIC2	CN3.N3:SW2.N37	Port 37	SW2	1	
	CN2	NIC1	CN2.N2:SW1.N35	Port 35	SW1	1	
		NIC2	CN2.N3:SW2.N35	Port 35	SW2	1	
	CN1	NIC1	CN1.N2:SW1.N33	Port 33	SW1	1	
		NIC2	CN1.N3:SW2.N33	Port 33	SW2	1	
	SN6	NIC1	SN6.N1:SW1.N11	Port 11	SW1	1.5	
		NIC2	SN6.N2:SW2.N11	Port 11	SW2	1.5	
	SN5	NIC1	SN5.N1:SW1.N09	Port 09	SW1	1.5	
		NIC2	SN5.N2:SW2.N09	Port 09	SW2	1.5	
	SN4	NIC1	SN4.N1:SW1.N07	Port 07	SW1	1.5	
		NIC2	SN4.N2:SW2.N07	Port 07	SW2	1.5	
	SN3	NIC1	SN3.N1:SW1.N05	Port 05	SW1	1.5	
		NIC2	SN3.N2:SW2.N05	Port 05	SW2	1.5	
	SN2	NIC1	SN2.N1:SW1.N03	Port 03	SW1	1.5	
		NIC2	SN2.N2:SW2.N03	Port 03	SW2	1.5	
	SN1	NIC1	SN1.N1:SW1.N01	Port 01	SW1	1.5	
		NIC2	SN1.N2:SW2.N01	Port 01	SW2	1.5	

8. Reconnect the power cords to the new Storage Interconnect.
The Storage Interconnect begins to power up as soon as the power cables are connected.

9. Reattach the two blanking plates onto the top front of the rack.

5.3 Fan Replacement Procedure

The switch has redundant hot-swappable fans.

Prerequisites

- Obtain a replacement fan from HGST.

Required Tools

- Ladder
- Long Phillips-head screwdriver

Time Estimate: Unknown.

This section describes how to replace a fan of the Storage Interconnect.

To replace the fan, do the following:

1. Remove the two blanking plates from the top front of the rack.
2. Identify the Storage Interconnect that has the faulty fan.

Figure 40: Location of Storage Interconnect Switches: [1]

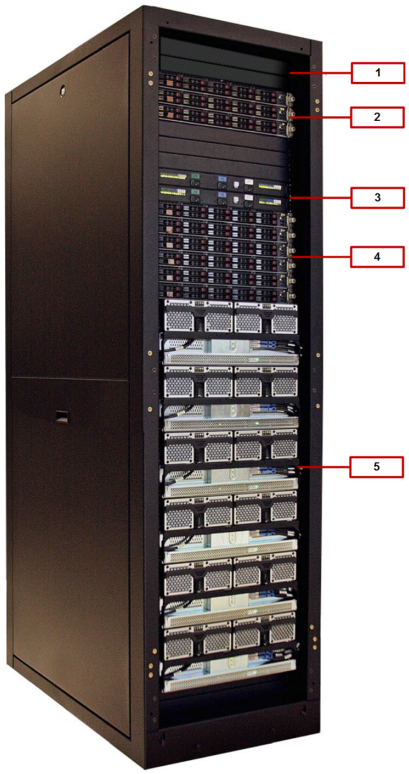

3. Identify the faulty fan.
4. Squeeze the two tabs inward on the faulty fan and pull the fan out.
5. Push the new fan into the same slot.
6. Reattach the two blanking plates onto the top front of the rack.

5.4 Power Supply Unit Replacement Procedure

The Storage Interconnect has redundant hot-swappable power supply units (PSUs).

Prerequisites

- Obtain a replacement PSU from HGST.

Required Tools

- Ladder
- Long Phillips-head screwdriver

Time Estimate: Unknown.

To replace a PSU, proceed as follows:

1. Remove the two blanking plates from the top front of the rack.
2. Identify the Storage Interconnect that has the faulty PSU.

Figure 41: Location of Storage Interconnect Switches: [1]

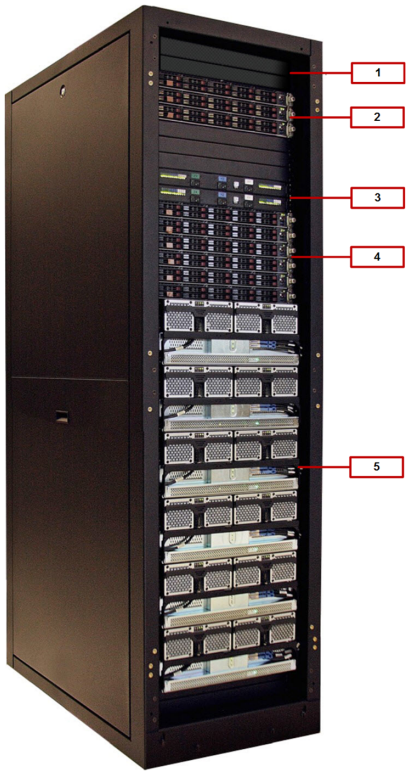

3. On the front side, identify the faulty PSU (amber colored LED).
4. Unplug the power cable of the faulty PSU.
5. Push the blue release tab towards the power connector and use the grab handle to remove the faulty PSU.
6. Push the new PSU into the Storage Interconnect.
7. Attach the power cable to the new PSU.
8. Reattach the two blanking plates onto the top front of the rack.

5.5 SFP+ 1G Module Replacement Procedure

The SFP+ 1G module connected to a Storage Interconnectswitch is an *Intel E10GSFPSR FTLX8571D3BCV-IT*..

Prerequisites

None.

Required Tools

None.

To replace an SFP+ 1G module on a Storage Interconnect, proceed as follows:

1. Remove the two blanking plates from the top front of the rack.
2. Identify the Storage Interconnect that has the faulty SFP+ 1G module.

   Figure 42: Location of Storage Interconnect Switches: [1]

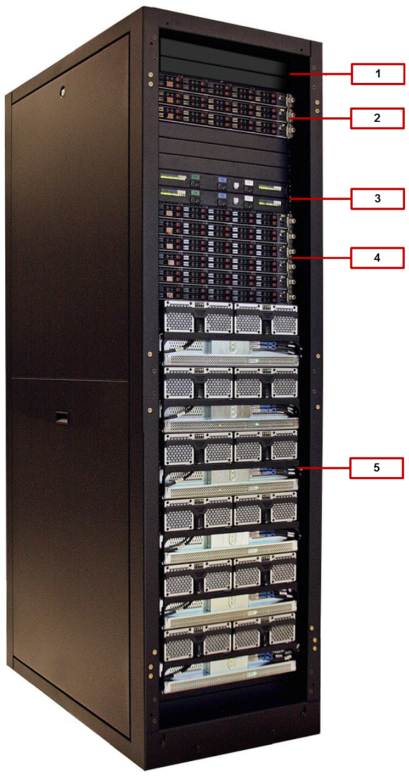

   a) Go to the back of the rack.
   b) Look for the Storage Interconnect whose SFP+ 1G module has an amber LED illuminated on its metal connector.
3. Remove the faulty SFP+ 1G module from the Storage Interconnect.
   a) Make a note of which Storage Interconnect port the SFP+ 1G module is connected to.
   b) Pull the SFP+ 1G module out of the Storage Interconnect.
   c) Disconnect the LC-to-LC optical cable from the SFP+ 1G module.
4. Install the new SFP+ 1G module.
   a) Connect the LC-to-LC optical cable you removed from the faulty SFP+ 1G module to the new SFP+ 1G module. The cable is reseated properly (the latch is engaged) when you hear a click.
   b) Push the SFP+ 1G module into the same Storage Interconnect port.
5. Verify that the network/activity LEDs on the new SFP+ 1G module are illuminated/active.
6. Reattach the two blanking plates onto the top front of the rack.

6 Power Distribution Unit Replaceable Units

Topics:

This section provides replacement procedures for a power distribution unit (PDU):

- PDU

6.1 Warnings

Caution: During the replacement of a PDU, the Active Archive System is running on a single power source.

6.2 Power Distribution Unit Replacement Procedure

Prerequisites

- Obtain a replacement power distribution unit from HGST.

Required Tools

- Ladder
- Long Phillips-head screwdriver

To replace a Power Distribution Unit (PDU) in a rack, proceed as follows:

Note: PDU01 is located on the right side when facing the back of the rack (the side of switch ports).

1. Shut down the entire rack.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers**.
   b) Select each of the six Storage Nodes, and in the right pane, click **Shutdown**.
   c) Shutdown Controller Node 3, then Controller Node 2, then Controller Node 1 (in other words, shut down the Management Node last).
2. Go to the data center and identify the defective PDU.
3. Remove the defective PDU from the rack.
   a) Disconnect the external power cables from the rack.
   b) Unscrew the green ground wire from the rack of the desired PDU.
   c) Disconnect the internal power cables from the desired PDU.
   d) Unscrew the rear support bracket of the desired PDU.
   e) Unscrew the mounting screws at the front of the desired PDU.
   f) Pull the desired PDU out of the rack. Be careful to feed the external power cable along with the desired PDU out of the rack.

   Caution: Two people should be used to unmount the PDU.

   Warning: Once you pull the PDU past the *pull-safety*, do not leave the PDU hanging in the rack. Otherwise, the rack rails may be damaged permanently.
4. Install the replacement PDU into the rack.

a) Feed the external power cable of the replacement PDU into the rack as you slide the replacement PDU into the slides of the rack.

Caution: Two people should be used to mount the PDU.

---

b) Screw the front mounting screws of the PDU to secure it.
c) Screw the green rack ground wire to the rack.
d) Screw the rear support bracket onto the desired PDU.
e) Connect the internal power cables to the desired PDU.

5. Power on the entire rack.

a) Connect the external power cords of the rack to two different power distribution networks.
The rack begins to power up as soon as the power cords are connected. The intelligent programmable PDUs control the bring-up sequence.
b) Confirm that all nodes power on in the right order. There is a short gap between each segment:

a. Network switches
b. Controller Nodes
c. Storage Enclosure Basic
d. Storage Nodes

c) Log into the CMC.
d) Wait until the CMC displays the status of the Management Node as **RUNNING**; in other words, its startup is complete.
e) Verify that the CMC dashboard indicates that the system status is good:

Disk Safety is **5**.
Controller Nodes indicate the correct number are **UP**.
Storage Nodes indicate the correct number are **UP**.
MetaStores indicate the correct number are **OK**.
Disks displays the correct number for your system, and none are degraded or decommissioned.
No status indicator is red.

f) Verify that the CMC displays the status of at least 5 Storage Nodes as **RUNNING**:
Navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Storage Nodes**. Check the **Status** field.

7 Storage Enclosure Basic Field Replaceable Units

Topics:

This section provides replacement procedures for the following parts in a Storage Enclosure Basic:

- HDD
- Rear Fan
- PSU
- I/O Canister

7.1 Visual Indicator and Field Replaceable Units Locations

The Storage Enclosure Basic displays the following visual indicators:

- I/O Canister
  - 1 Green enclosure OK LED
  - 1 Blue enclosure Identify LED
  - 1 Amber enclosure Fault LED
- SAS Riser (one set of LEDs per host port, 2 sets total)
  - 1 Green Link OK LED
  - 1 Amber Identify/Fault LED
- Fan (one set of LEDs per fan, 3 sets total)
  - 1 Amber Fan Identify/Fault LED
- PSU (one set of LEDs per PSU, 2 sets total)
  - 1 Green AC OK LED
  - 1 Green DC OK LED
  - 1 Blue PSU Identify LED
  - 1 Amber PSU Fault LED
- Sled (one set of LEDs per sled, 7 sets total)
  - 1 Amber Sled Identify/Fault LED
  - 1 Amber Drive Identify/Fault LED per drive (14 per sled)

The following diagram displays the visual indicators for the I/O canister, sled, and the rear fans in the Storage Enclosure Basic:

Figure 43: System Enclosure Information

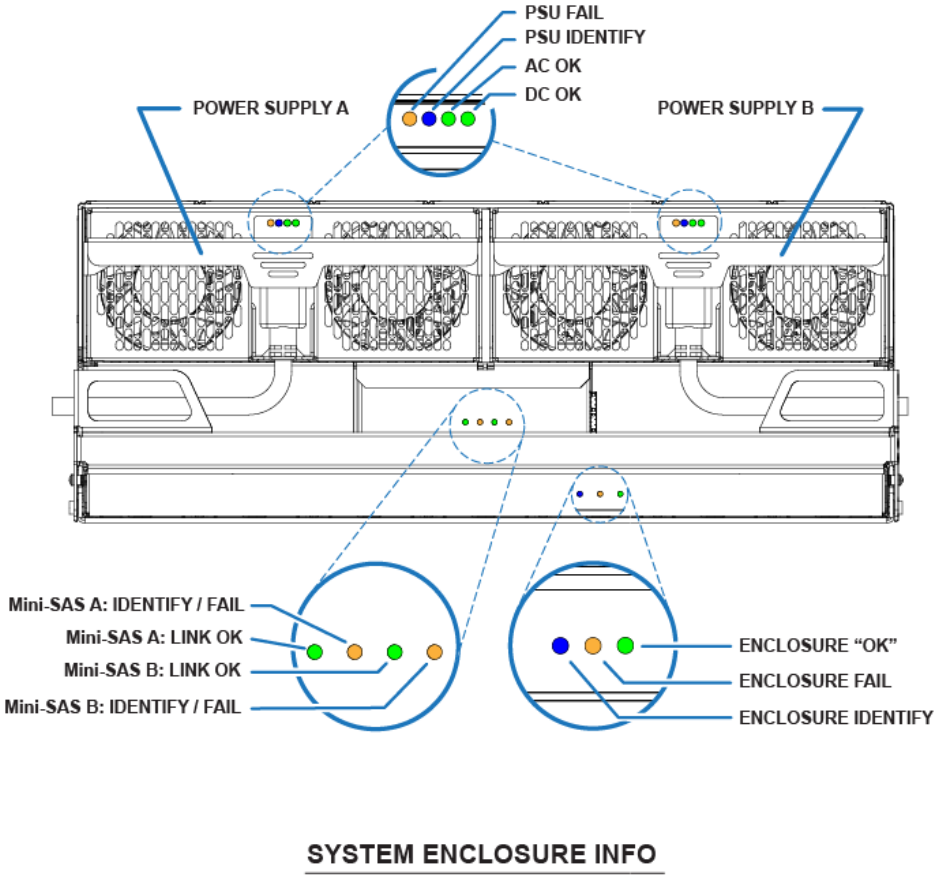

The following diagrams display the physical locations of the various FRUs and visual indicators in the Storage Enclosure Basic:

Storage Enclosure Basic

Figure 44: Rear Fan Order

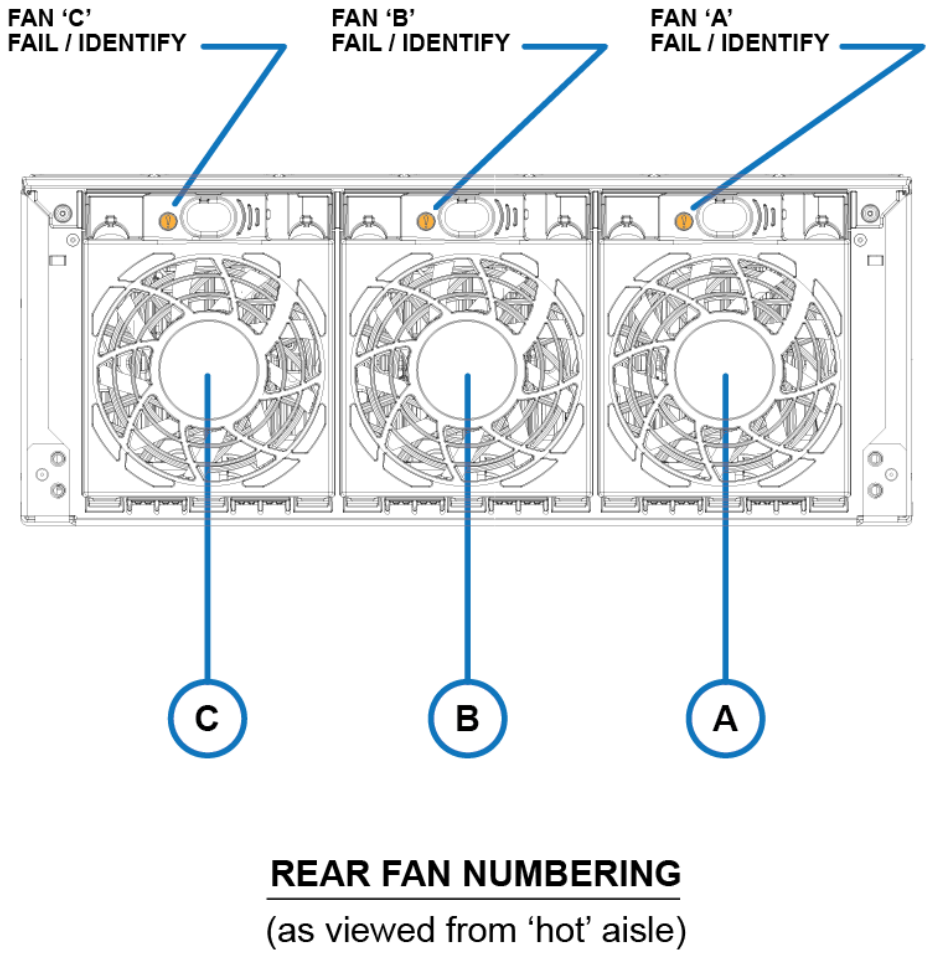

Figure 45: Sled HDD Order

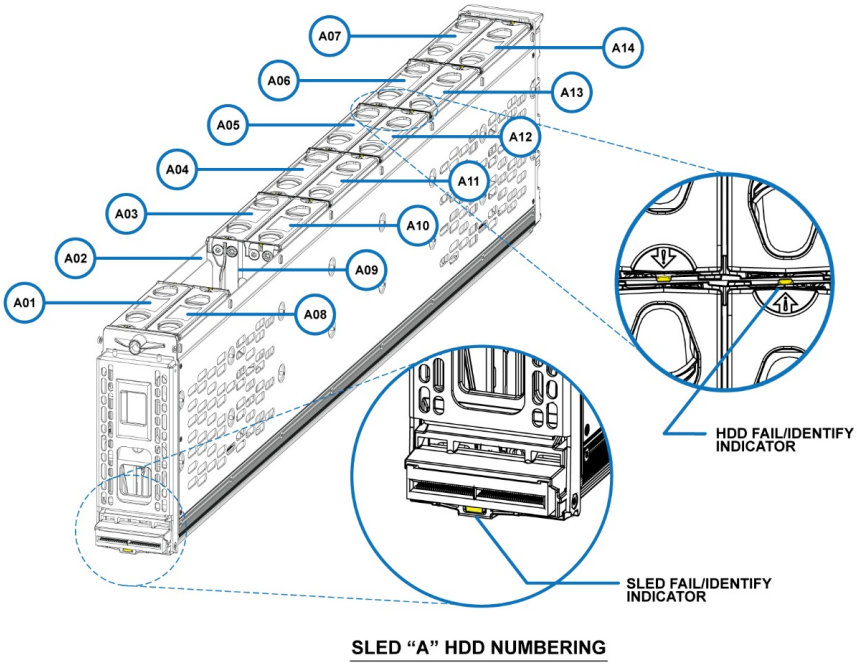

7.2 Hard Disk Drive Replacement Procedure

Based on the configuration of a Storage Enclosure Basic, an enclosure contains seven sleds.

Prerequisites

- If you are replacing multiple drives on different enclosures, temporarily configure the MetaStores as read only. Do not attempt to replace multiple drives on different enclosures unless you first make the MetaStores read only. For instructions, see *Managing Storage* in the *HGST Active Archive System Administration Guide*.
- Decommission the faulty drive in the CMC. For more information, see *Managing Hardware* in the *HGST Active Archive System Administration Guide*.
- Obtain a replacement HDD from HGST Support.

Required Tools

- None

Time Estimate: 1 hour

To replace a hard disk drive, do the following:

Note:

- Ensure that you store all removed parts in a safe location while replacing the FRU.
- The hard disk drive FRU is defined as the hard drive with the drive carrier. This means that you need to replace the hard drive along with the carrier.

Warning: If you are replacing more than one hard disk drive, wait in between each replacement until the CMC displays a **Disk Safety** of 5.

---

1. Obtain details about the decommissioned disk from the CMC.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Disks** > **Decommissioned**.

   Figure 46: Decommissioned Disks in the CMC

   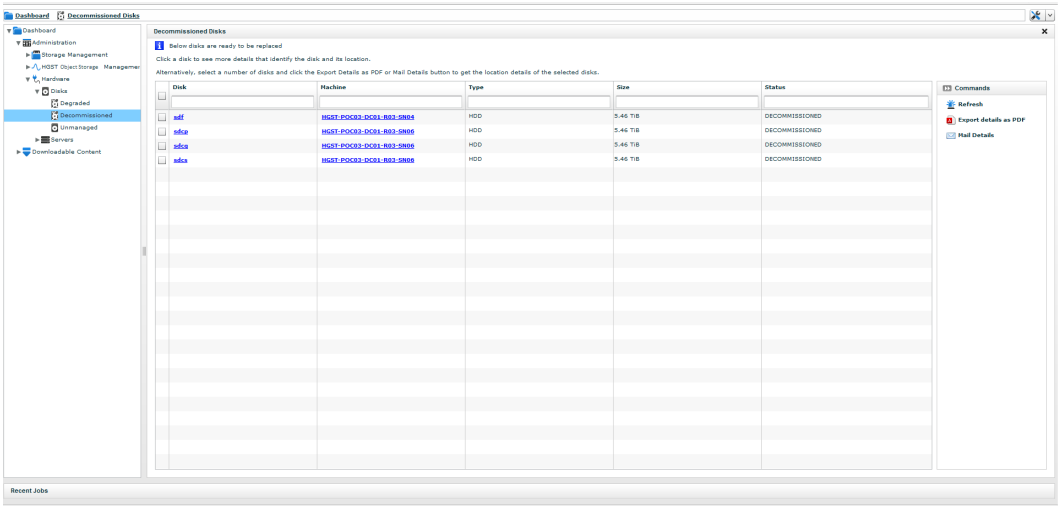

   b) Click the desired decommissioned disk.
   A PDF is displayed.
   c) Click **Save** or **Print**.
   The PDF contains *decommissioned disk details*, which you will need refer to:
   - The **device name**

- The model type and serial number
- A *drive map* showing the exact location of the disk (name of data center, rack, and node)
- The type of node (`CPUNODE`/`STORAGENODE`)
- The hostname of the node (**Node name**)
- The MAC addresses of the node
- The current status of the node

Figure 47: Decommissioned Disk Details from the CMC

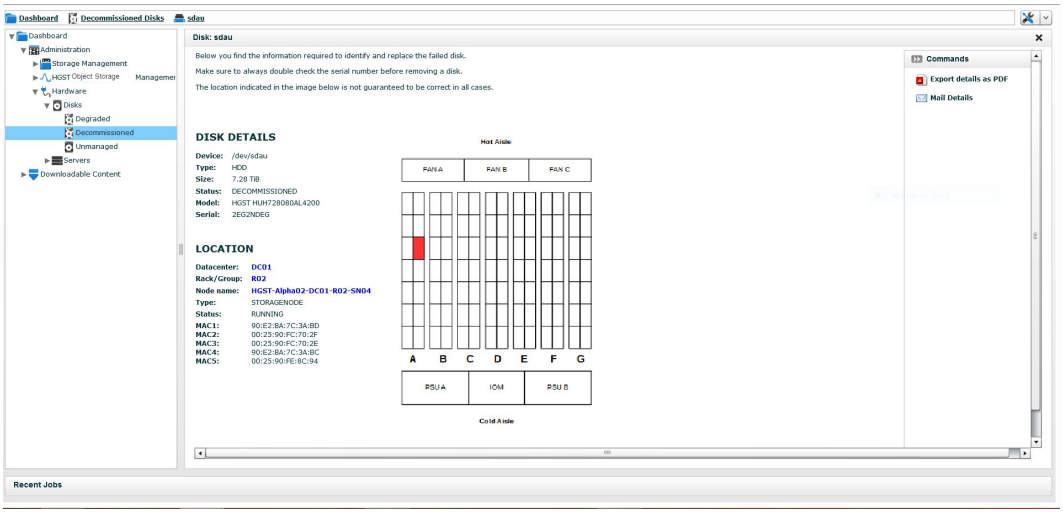

2. Get the IP address of the Storage Node.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Storage Nodes**.
   b) Select the correct Storage Node based on the decommissioned disk details.
   c) In the **Commands** pane, click **Location LED On**.
3. Enable the location LED on the Storage Node.
   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Storage Nodes**.
   b) Select the correct node.
   c) In the **Commands** pane, click **Location LED On**.
4. Enable the blue blinking location LED on the decommissioned drive.
   a) Open an SSH session to any Controller Node.
   The OSMI menu appears.
   b) Exit the OSMI menu.
   The Linux prompt appears.
   c) Open an SSH session to the Storage Node that is paired with the Storage Enclosure Basic containing the decommissioned drive.
   d) Start the Q-Shell.

```
/opt/qbase3/qshell

Welcome to qshell

?          -> Introduction to features.
help()     -> python help system.
object?    -> Details about 'object'.
object??   -> Extended details about 'object'.

Type q. and press [TAB] to list qshell library
Type i. and press [TAB] to list interactive commands
```

e) Identify the *drive slot number* based on the **device name** in the decommissioned disk details you obtained in Step 1.

> **Note:** Replace *hostname_of_storage_node* with the value of **Node name** in the decommissioned disk details, and replace *device_name* with the value of **Device** in the decommissioned disk details (but remove the /dev/ prefix).

```
In [1]:api = i.config.cloudApiConnection.find('main')

In [2]:mguid = api.machine.find(name='hostname_of_storage_node')['result'][0]

In [3]:print(api.disk.list(machineguid=mguid, name='device_name')['result'][0]
['bus_location'])
EXP_SLOT_69
```

In the example above, the **drive slot number** is 69.

f) Exit the Q-Shell.

```
quit()
```

g) Identify the *sg number* of the Storage Enclosure Basic.

```
lsscsi -g | grep PEAK
[0:0:42:0]   enclosu HGST     PIKES PEAK       0109  -         /dev/sg42
```

In the example above, the sg number is /dev/sg42.

h) Issue the command to the Storage Enclosure Basic to enable ("set") the LED on the drive identified by the **drive slot number** obtained in substep e.

> **Important:** Subtract 1 from the **drive slot number**, because the index starts from 0.
>
> For example, if the **drive slot number** is 69, you must use 68 in the command below.

```
sg_ses /dev/sg42 -p2 -I0,68 -S ident
```

i) Log out of the Storage Node.

j) Log out of the Controller Node.

5. Shut down the Storage Node from the CMC.

> **Caution:** Shut down **only** the Storage Node that is paired with the Storage Enclosure Basic containing the FRU.

a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Storage Nodes**.

b) Select the desired Storage Node.

Figure 48: A Storage Node Pane in the CMC

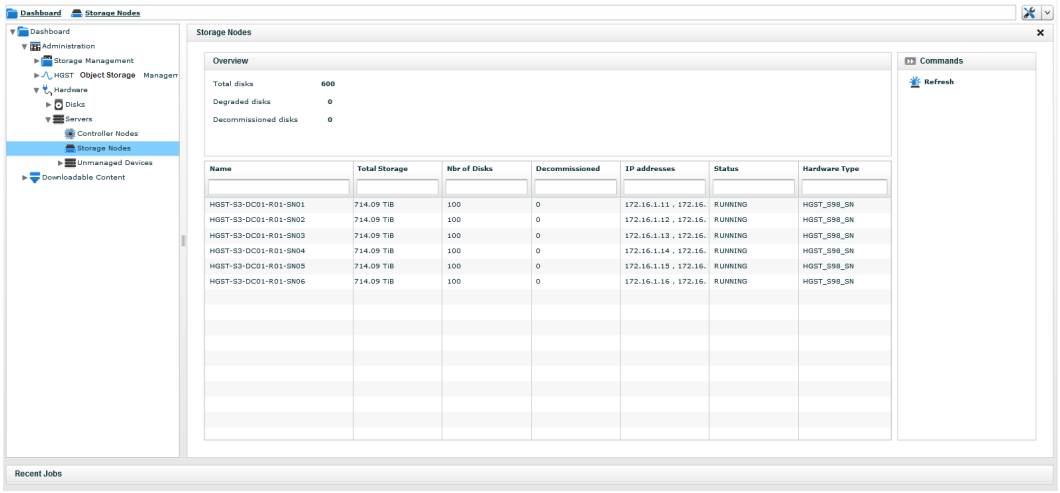

c) In the **Commands** pane, click **Shutdown**.

Figure 49: The Shutdown Button in the Commands Pane

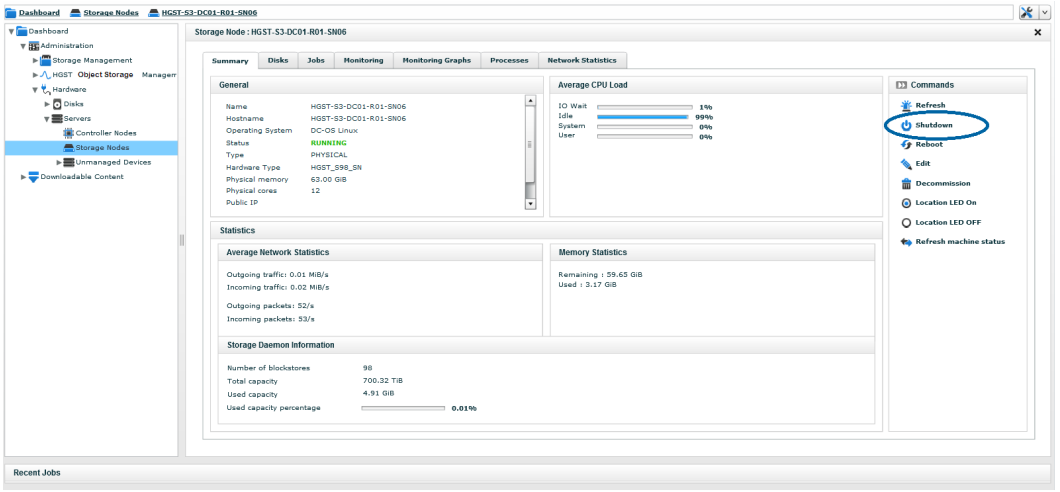

d) Wait for the **Status** field to change to **DONE**.

Warning: Even if all LEDs are off, you must still wait until the CMC shows **DONE** in the **Status** field.

All I/O to the Storage Enclosure Basic attached to this Storage Node is now quiesced.

6. Go to the rack and identify the correct chassis by the blinking blue LED on its front and back panels.
7. Locate the enclosure that contains the failed hard disk drive.

Note: The enclosure containing the failed drive will have a flashing blue identification LED.

---

8. Unplug the power cables by lifting the **power cord retention bale** and carefully removing the power cord from the **power supply**.

Note: Repeat for both the A and B power supplies.

Figure 50: Removing the Power Cords

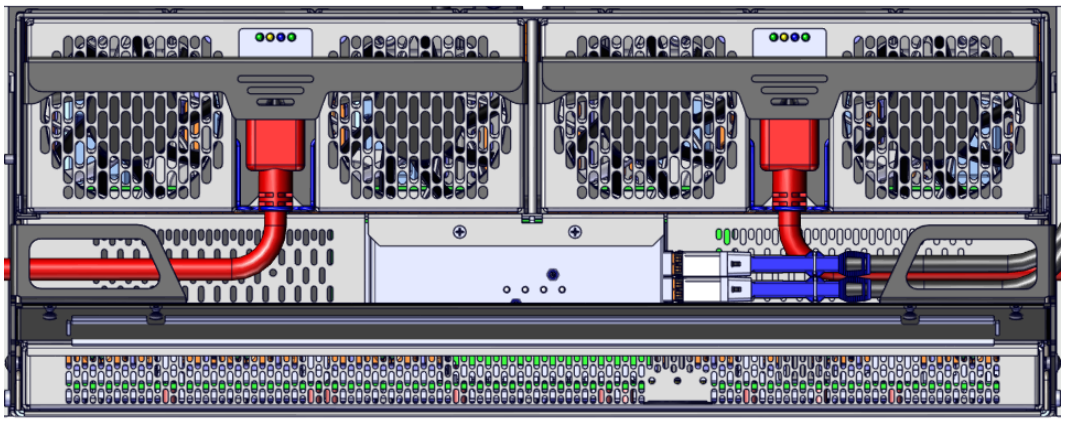

Note:

- Power cords marked in red.
- Cord retention bale marked in blue.

9. Unplug the miniSAS cables from the **I/O canister**.

Figure 51: Removing the MiniSAS Cables

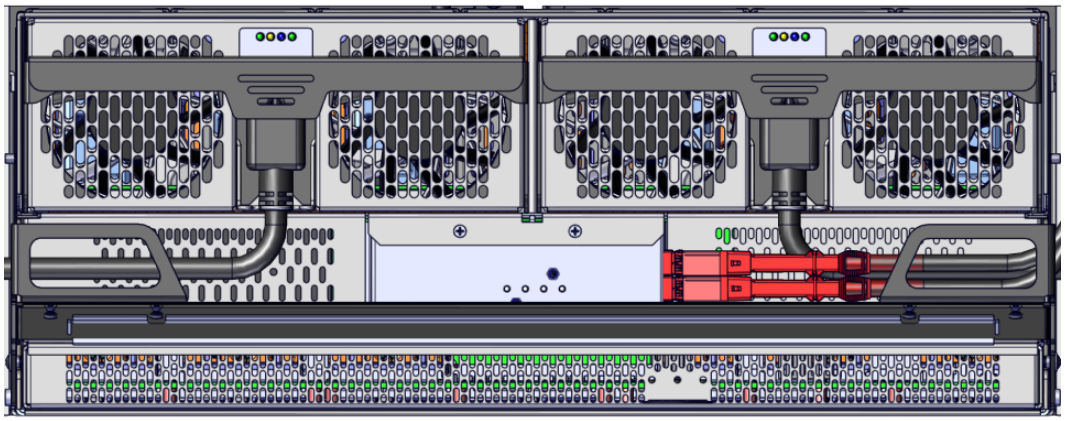

Note:

- The miniSAS cables are marked in red.
- Take note of which miniSAS cable came from which port to ensure that they are plugged in correctly when reassembling.

10. Unlock the I/O canister from the enclosure by pulling the **latch handle** out and away from the I/O canister.

Figure 52: Unlocking the I/O Canister

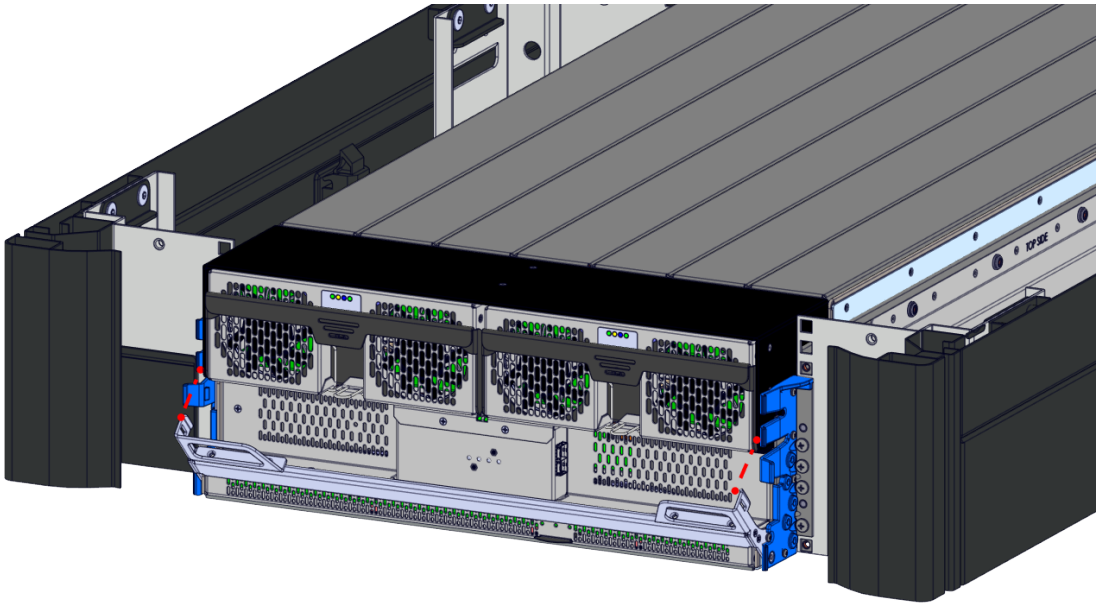

11. Slide the **I/O canister** away from the **chassis**.

Figure 53: Removing the I/O Canister

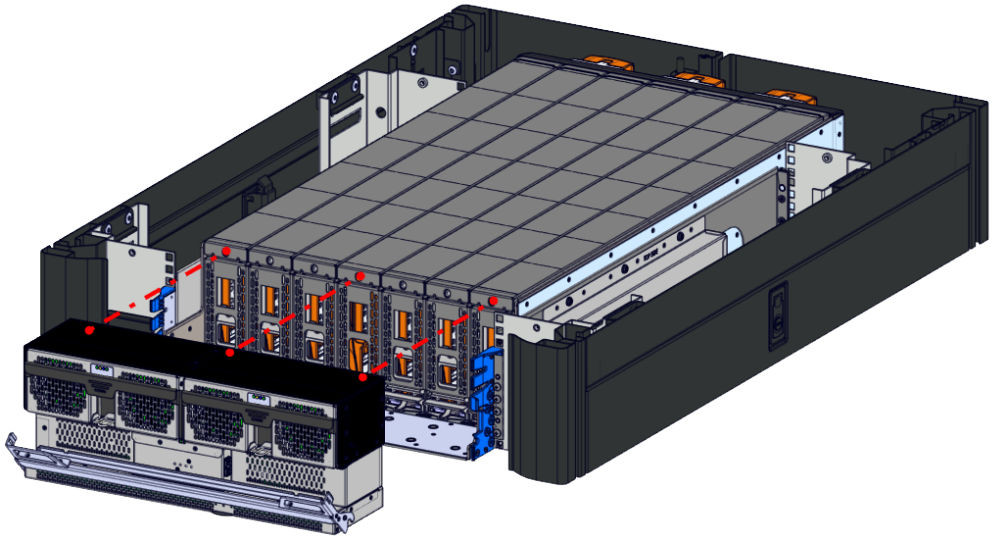

12. Locate the failed sled by identifying the Sled Fail/Identify indicator is blinking **amber**.

Figure 54: Sled HDD Order

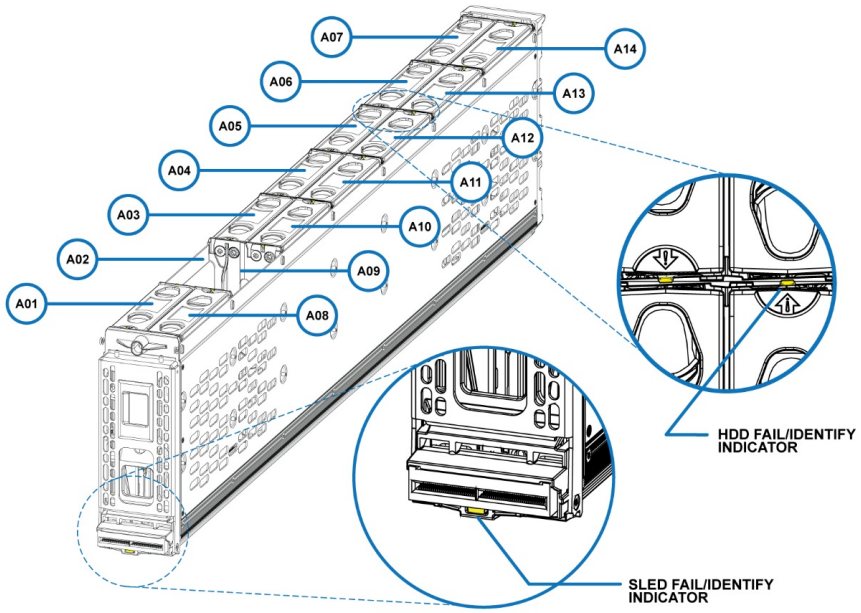

13. On the front of the first sled, depress the **latch mechanism button** and pull it until it is at a 45 degree angle.

Figure 55: Sled Release Button

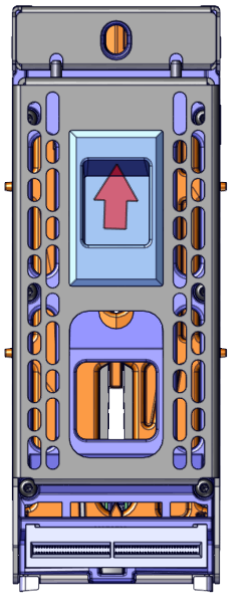

Figure 56: Sled Release at 45 Degrees

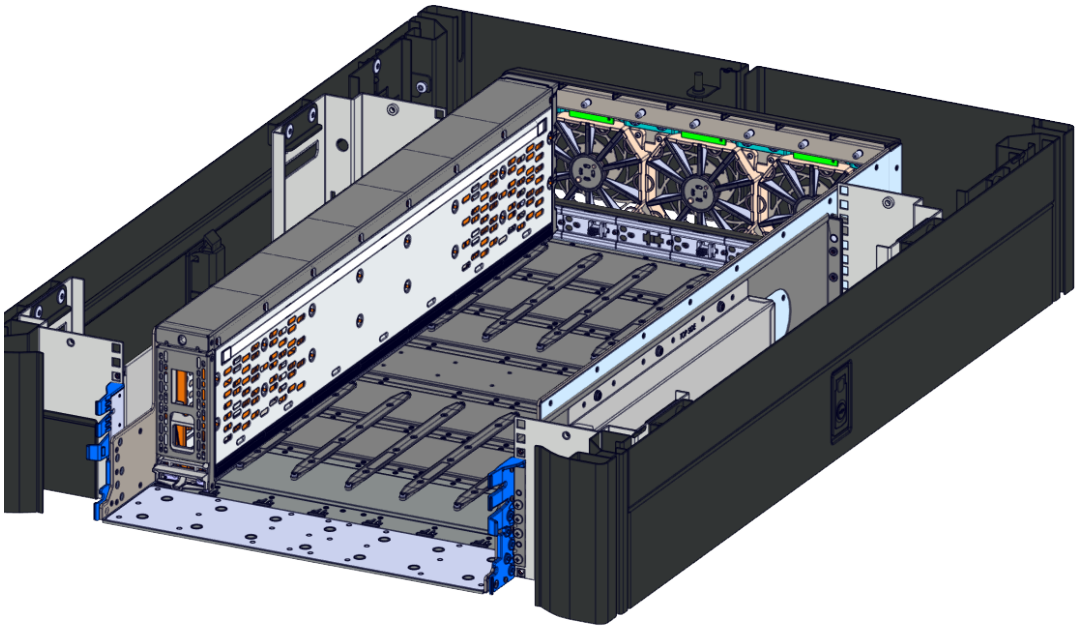

14. Remove the sled in need of a new hard disk drive, out of the chassis.

Figure 57: Removing the Sled

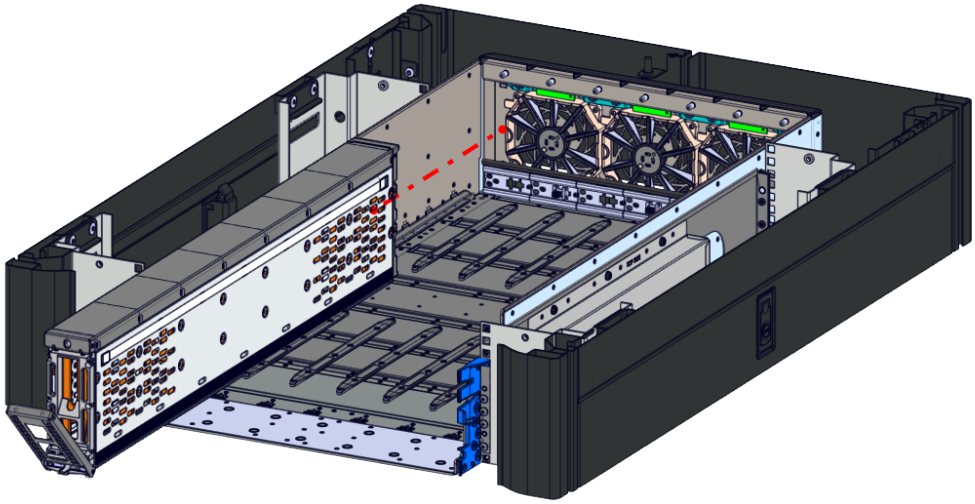

Note:

- Ensure that you remove and replace the sleds and sled blanks in the same order.
- Repeat the two previous steps until all of the sleds, in need of replacement, have been removed from the chassis.

15. On the sled, slide the drive cover forward and up until the cover has been removed.

Figure 58: Sled Cover

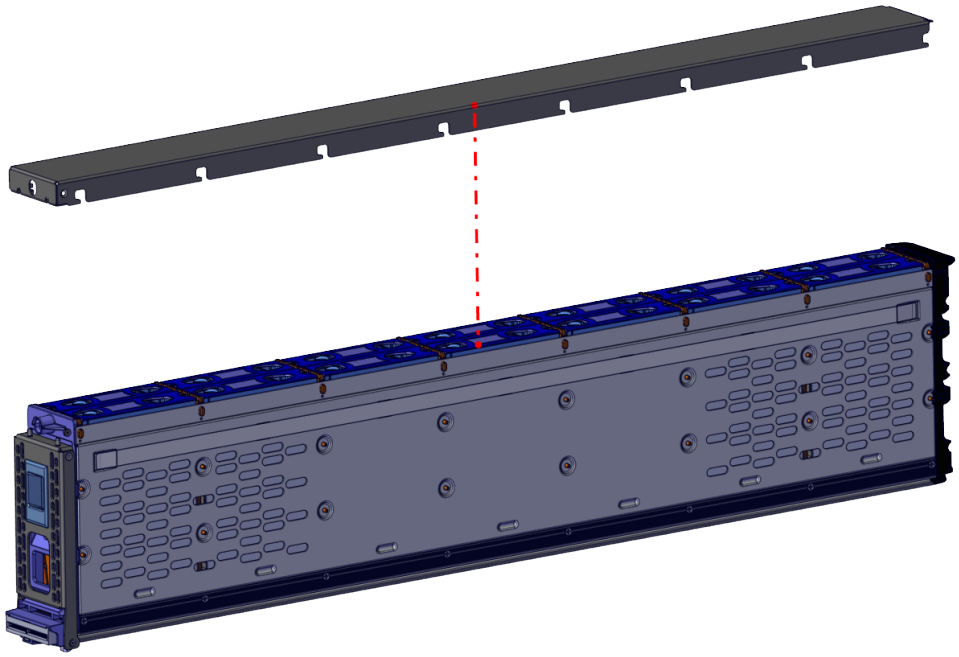

16. Identify the drive to be replaced by referring to the drive map you obtained from the CMC and the amber LED.

Tip: The correct drive is the one whose blue arrow is pointing at the illuminated amber LED.

17. To remove the hard disk drive, on the hard drive carrier, depress the two buttons and remove the drive.

Figure 59: Hard Disk Drive Carrier Buttons

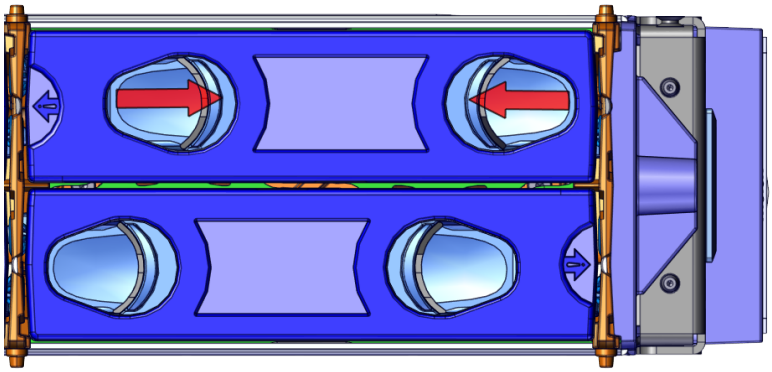

Figure 60: Removing the Hard Disk Drive with Carrier

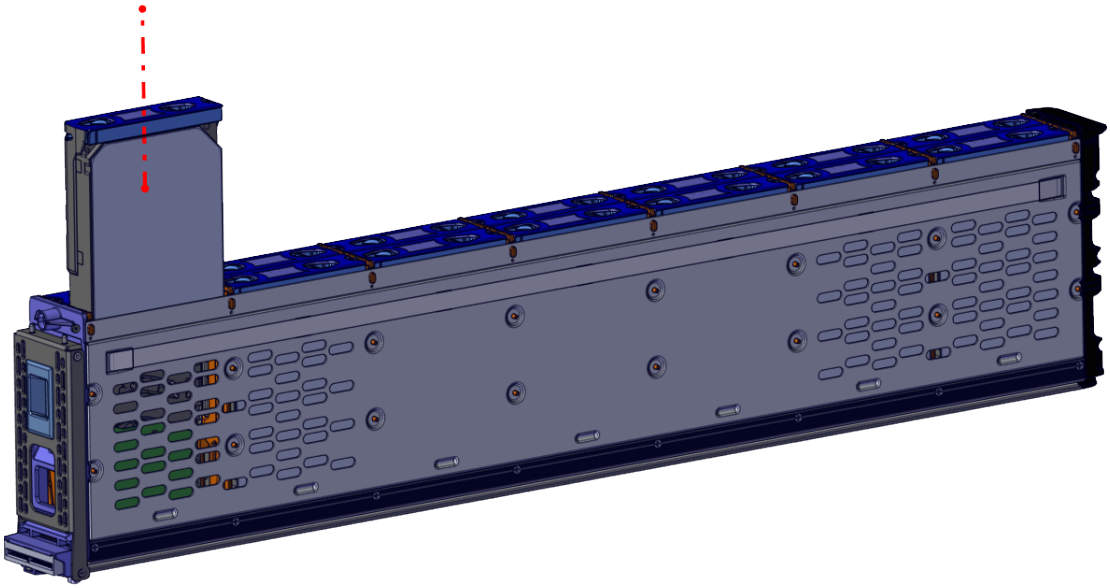

18. Install the replacement hard disk drive with carrier in the reverse order that you removed it.
19. Install the remaining enclosure components in the reverse order that you removed them.
20. Re-connect the enclosure to the power cords.
21. On the PSUs, identify that the AC and DC LEDs display green indicators.
22. Power on the Storage Node.
    The power button is located on the chassis front control panel.
23. Disable the identification LED on the Storage Enclosure Basic.
    a) Open an SSH session to any Controller Node.
    The OSMI menu appears.
    b) Exit the OSMI menu.
    The Linux prompt appears.

c) Open an SSH session to the Storage Node that is paired with this Storage Enclosure Basic.
d) Disable ("clear") the drive LED using the **drive slot number** you obtained in Step 4e.

> **Important:** Subtract 1 from the **drive slot number**, because the index starts from 0.
>
> For example, if the **drive slot number** is `69`, you must use `68` in the command below.

```
sg_ses /dev/sg42 -p2 -I0,68 -C ident
```

24. Disable the location LED on the Storage Node.
a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Storage Nodes**.
b) Select the desired Storage Node.
c) In the **Commands** pane, click **Location LED Off**.

25. Confirm that the Active Archive System correctly determines the purpose for the new disk.
a) Wait 15 minutes.
b) In the CMC, navigate to **Dashboard** > **Administration** > **HGST Object Storage Management** > **Logging** > **Events**.
c) In the **Events** list, check to see that a new empty disk has been detected.
d) In the **Jobs** list, check to see that an `Initializing new disk` job has been triggered.
It may take about 2 minutes for the job to appear.
e) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers**.
f) Select the desired node.
g) Select the **Disks** tab.
h) Wait for the physical drive that has been replaced, as well as the logical disks, to change status from a red icon to a green icon.

> **Note:** The physical drive that has been replaced, as well as the logical disks, may take up to 40 minutes to change status.

The `Initializing new disk` job has completed successfully when the number of degraded disks decreases by 1.

26. If the disk still shows up in the **Degraded** or **Unmanaged** list, you must manually specify the purpose of the new disk:
a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Disks** > **Unmanaged**.
b) Select the new disk, and in the **Commands** pane, click **Repurpose**.
c) In the **Use As** field, select **Replacement Disk**.

> **Note:** You can only select **Replacement Disk** when there is a decommissioned disk. If there are no decommissioned disks, you can only select **Additional Disk** as the purpose for the disk.

d) In the **Replacement For** field, select the decommissioned disk that you want to replace.
e) Click **Next** to start the repurposing.
An `Initializing new disks on` *`node_name`* job starts.

If You Replaced the Wrong Disk

See *Troubleshooting*.

7.3 Power Cord Replacement Procedure

Required Tools

- None

To replace a power cord, do the following:

Time Estimate: 7 minutes

1. Unplug the failed power cord by lifting the **power cord retention bale** and carefully removing the power cord from the **power supply**.

   Figure 61: Replacing the Failed Power Cord

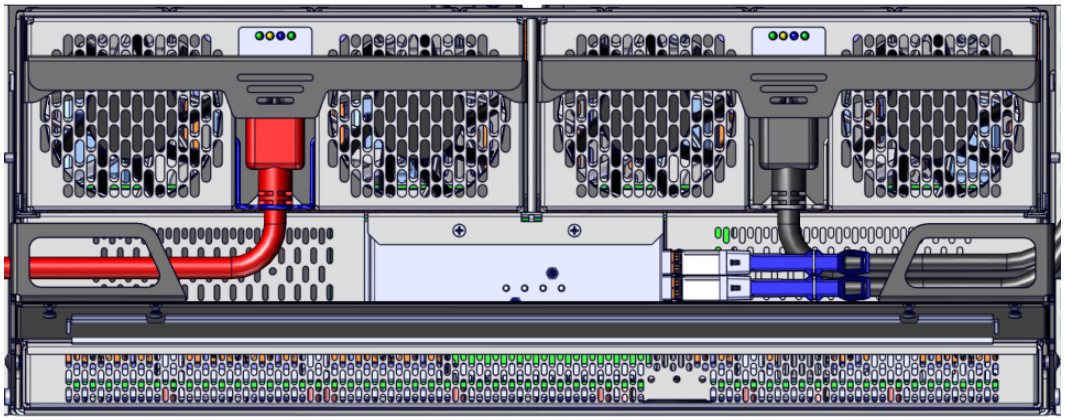

   Note: The power cord is marked in red and the power cord retention bale is marked in blue.

2. Do the following to remove the failed power cord from the rack:
   a) Disconnect the failed power cord from the server.
   b) From the Enclosure end, pull the power cord through the rail kit cable guides.
   c) Pull the power cord up through the side of the rack rail.
   d) Pull the power cord through the top of the rack.
3. **Note:** Ensure the new power cord is installed in the same location as the failed power cord.

   Do the following to install the new power cord into the rack:
   a) Run the new power cord through the top of the rack.
   b) Pull the power cord down through side of the rack rail.
   c) From the Enclosure end, pull the power cord through the rail kit cable guides.
   d) Connect the power cable to the server.
4. Pull the power cord through I/O module cable guides.
5. Plug the power cord into the power supply unit.
6. To secure the power cord, press the power cord retention bale into the I/O module.

7.4 MiniSAS Cable Replacement Procedure

Required Tools

- None

To replace a miniSAS cable, do the following:

Time Estimate: 7 minutes

1. From the I/O module, unplug the failed miniSAS cable.

   Figure 62: Replacing the Failed MiniSAS

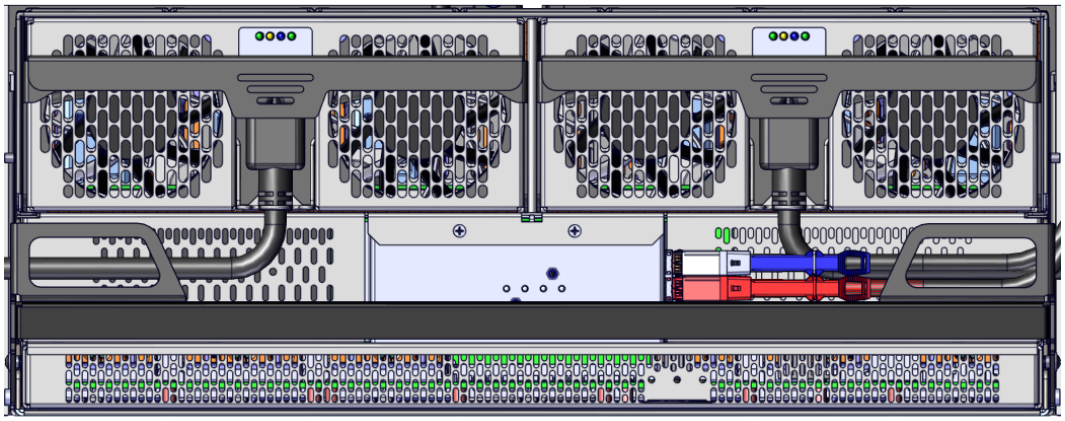

   Note: The miniSAS cable is marked in red.

2. Do the following to remove the failed miniSAS cable from the rack:
   a) Disconnect the failed miniSAS cable from the storage node.
   b) From the enclosure end, pull the miniSAS cable through the rail kit cable guides.
   c) Pull the miniSAS cable up through the side of the rack rail.
3. **Note:** Ensure the new miniSAS cable is installed in the same location as the failed miniSAS cable.

   Do the following to install the new miniSAS cable into the rack:
   a) From the storage node, run the new miniSAS cable through the cable guides.
   b) Pull the miniSAS cable down through side of the rack rail.
   c) From the enclosure end, pull the miniSAS cable through the rail kit cable guides.
   d) Connect the miniSAS cable to the server.
4. Pull the miniSAS cable through I/O module cable guides.
5. Plug the miniSAS cable into the I/O module.

7.5 Rear Fan Replacement Procedure

Based on the configuration of a Storage Enclosure Basic, an enclosure can contain three, five, or seven sleds. The system is designed for fail-in-place operational mode.

Required Tools

- None

Time Estimate: 2 minutes

To replace a rear fan, do the following:

Note:

- Ensure that you store all removed parts in a safe location while replacing the FRU.
- The rear fans are hot-swappable. The enclosure does not need to be powered down in order to replace them.

---

1. From the rear of the chassis, remove the failed rear fan by depressing the **release button** on the top right of the fan.

   Figure 63: Fan Release Button

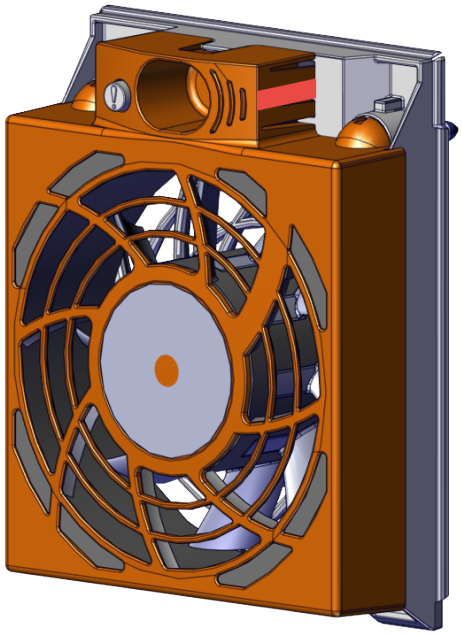

Note: Fan release button highlighted in red.

---

2. Rotate the top of the fan away from the chassis until the fan pins clear the connectors on the chassis.

Note: Repeat the previous step until all of the fans in need of replacement have been removed.

Figure 64: Rear Fan

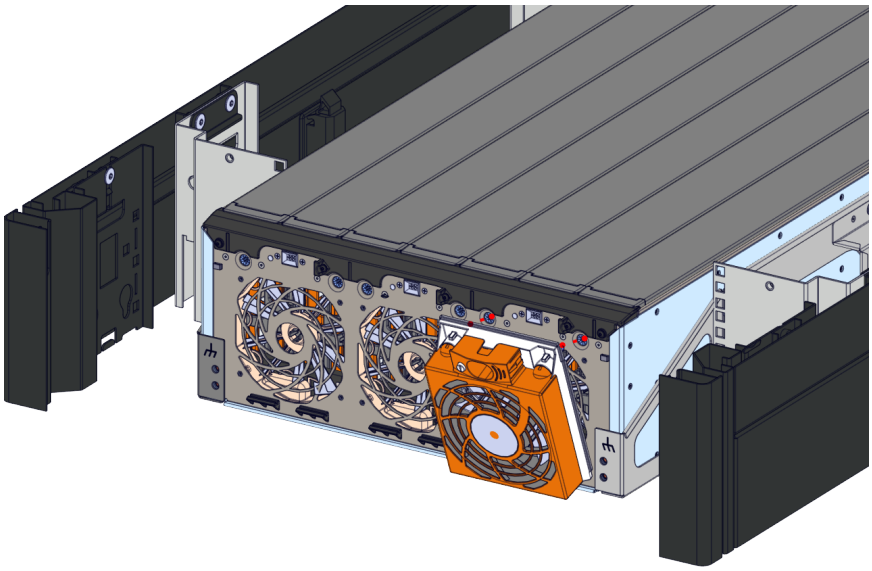

3. Remove the fan from the fan rubber bumpers on the chassis.
4. Install the replacement fan in the reverse order that you removed it.

7.6 Power Supply Unit Replacement Procedure

Required Tools

- None

Time Estimate: 3 minutes

To replace a power supply unit, do the following:

Note: Ensure that you store all removed parts in a safe location while replacing the FRU.

1. Shut down the Storage Node from the CMC.

   Caution: Shut down **only** the Storage Node that is paired with the Storage Enclosure Basic containing the FRU.

   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Storage Nodes**.
   b) Select the desired Storage Node.

   Figure 65: A Storage Node Pane in the CMC

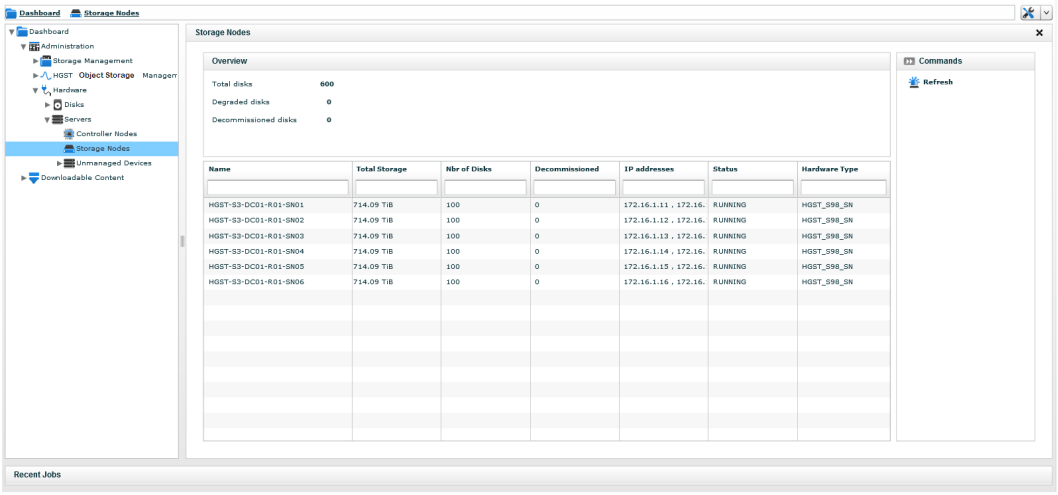

   c) In the **Commands** pane, click **Shutdown**.

   Figure 66: The Shutdown Button in the Commands Pane

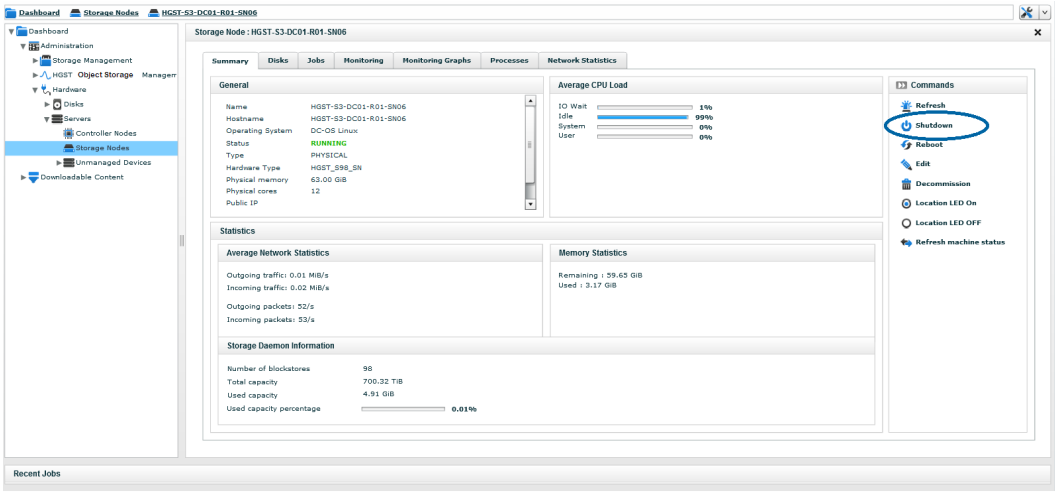

d) Wait for the **Status** field to change to **DONE**.

Warning: Even if all LEDs are off, you must still wait until the CMC shows **DONE** in the **Status** field.

All I/O to the Storage Enclosure Basic attached to this Storage Node is now quiesced.

2. Lift the cord retention bale and unplug the power cord from the failed power supply unit.

Figure 67: Removing the Power Cord

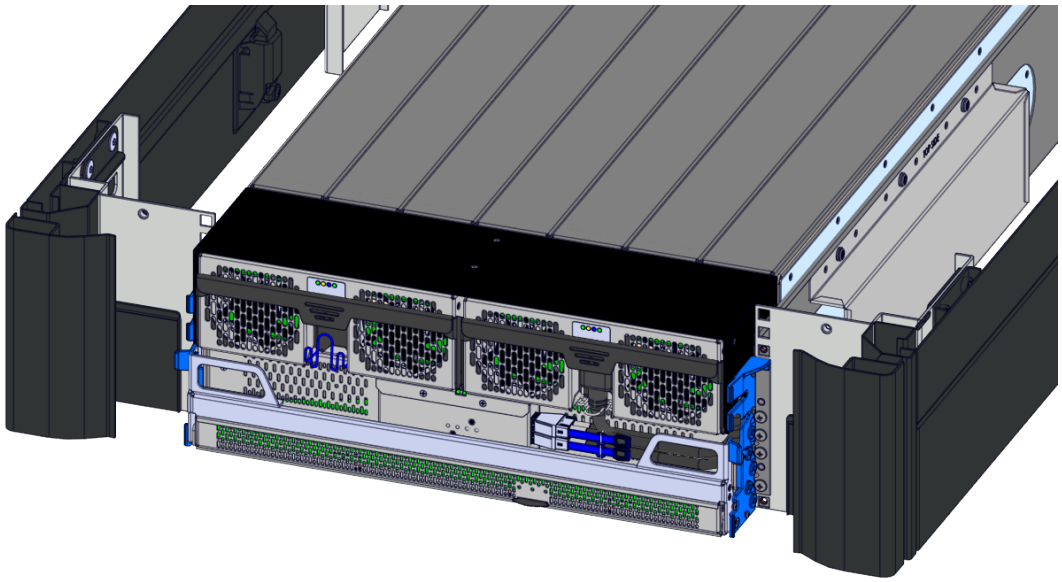

Note:

- Cord retention bale marked in blue.
- If you are removing power supply A, you do not need to remove the miniSAS cables. To remove power supply B, it is recommended that you remove the miniSAS cables for ease of replacement. To remove the miniSAS cables, pull the **blue tab** and remove the cable from the port. Repeat for both miniSAS cables as necessary.

3. Unlock the failed power supply unit by pulling the **latch handle** out and away from the I/O canister.

Note:

- The power supply unit latch handle should be at 45° when removed.

- Repeat this step for the remaining power supply unit if necessary.

Figure 68: Removing the Power Supply Unit

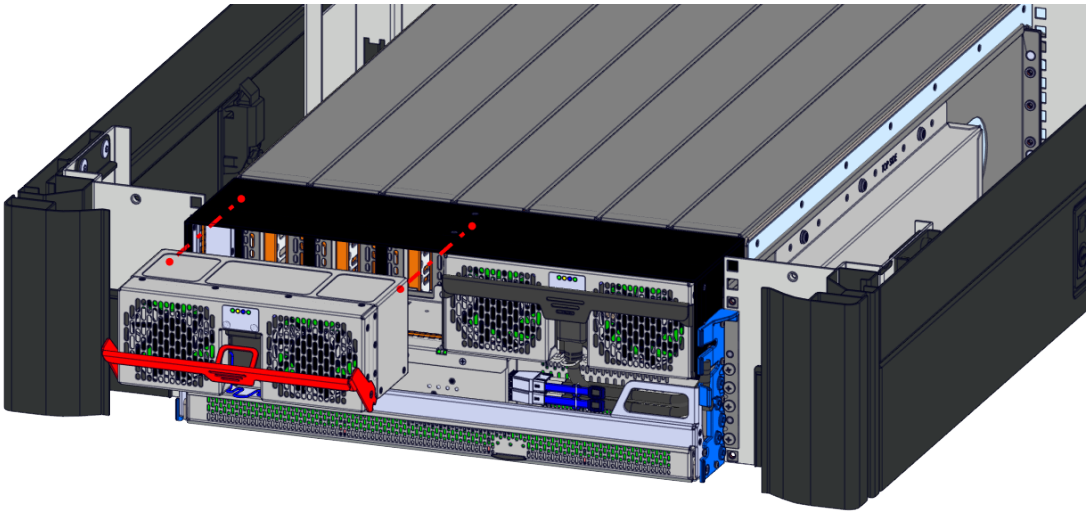

4. Remove the power supply unit until free of the I/O canister.
5. Install the replacement power supply unit.
6. Reconnect the miniSAS cables.
7. Plug the power cord back into the replaced power supply.
8. Power on the Storage Node.
   The power button is located on the chassis front control panel.

7.7 I/O Canister Replacement Procedure

Required Tools

- None

Time Estimate: 5 to 7 minutes

To replace the I/O Canister, do the following:

Note:

- Ensure that you store all removed parts in a safe location while replacing the FRU.
- Ensure you are wearing an ESD wrist strap to complete the replacement of the I/O canister.

---

1. Shut down the Storage Node from the CMC.

   Caution: Shut down **only** the Storage Node that is paired with the Storage Enclosure Basic containing the FRU.

   ---

   a) In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Servers** > **Storage Nodes**.
   b) Select the desired Storage Node.

   Figure 69: A Storage Node Pane in the CMC

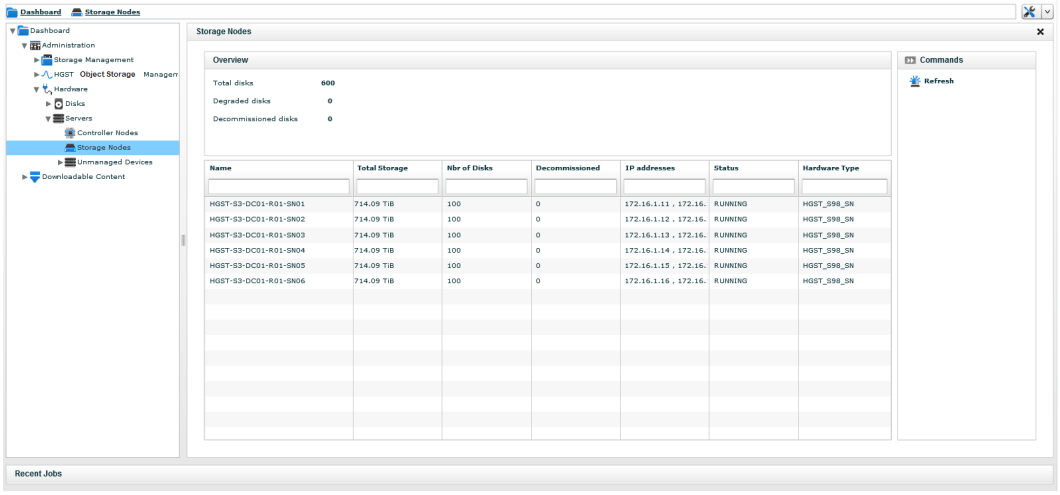

c) In the **Commands** pane, click **Shutdown**.

Figure 70: The Shutdown Button in the Commands Pane

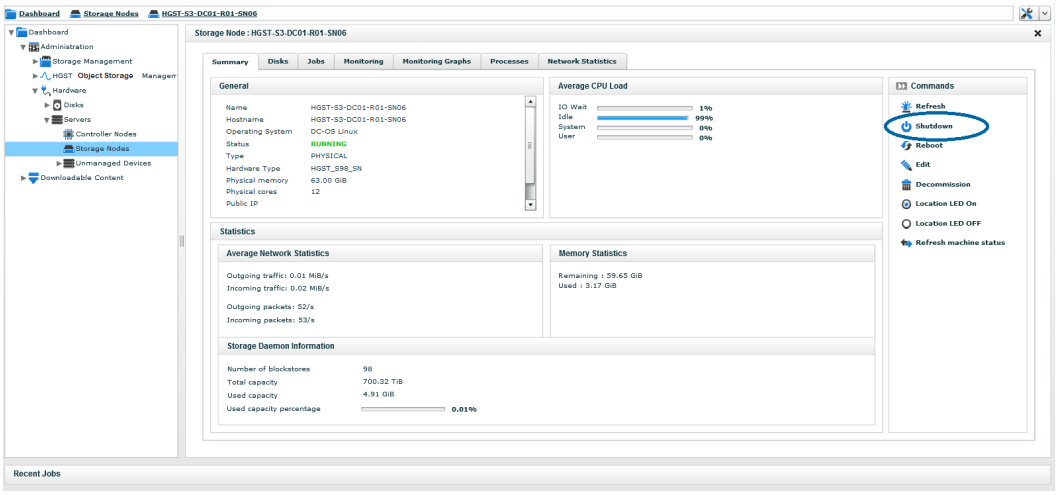

d) Wait for the **Status** field to change to **DONE**.

> **Warning:** Even if all LEDs are off, you must still wait until the CMC shows **DONE** in the **Status** field.

All I/O to the Storage Enclosure Basic attached to this Storage Node is now quiesced.

2. Identify the Storage Enclosure Basic that contains the failed I/O canister.

> **Note:** To identify the failed I/O Canister, verify that the amber light is blinking.

3. Remove the miniSAS cables by pulling the blue tab and remove the cable from the port.

4. Unplug the power cables by lifting the **power cord retention bale** and carefully removing the power cord from the **power supply**.

Figure 71: Removing the MiniSAS Cables

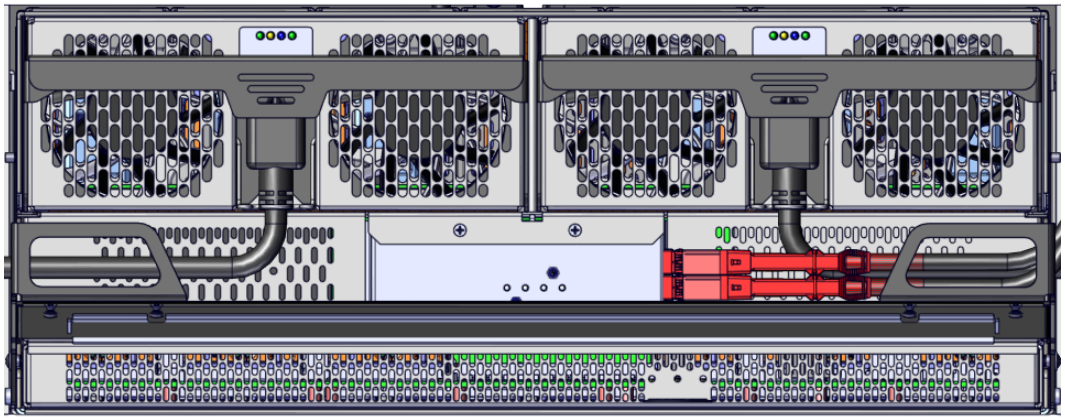

Note:

- The miniSAS cables are marked in red.
- Take note of which miniSAS cable came from which port to ensure that they are plugged in correctly when reassembling.

- Repeat for both the A and B power supplies.

Figure 72: Removing the Power Cords

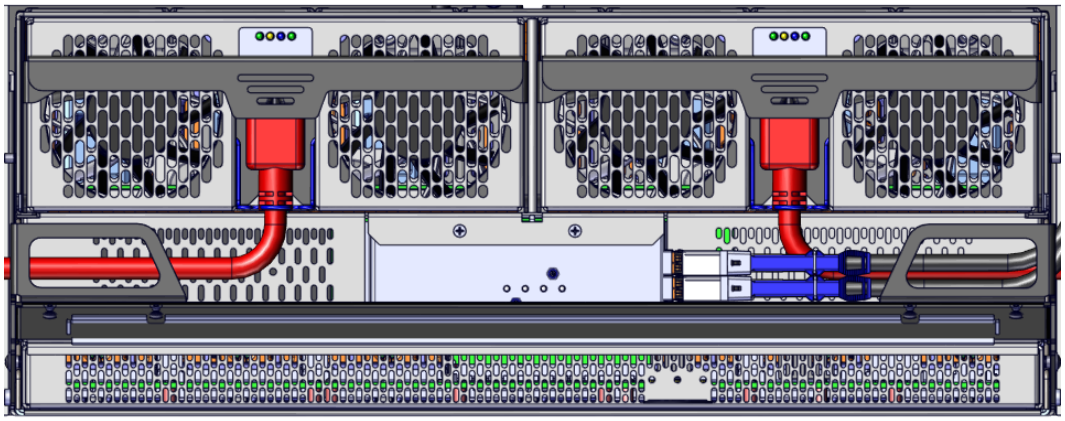

Note:

- Power cords marked in red.
- Cord retention bale marked in blue.

5. Wait approximately 30 seconds after the I/O canister is unplugged to continue with the replacement procedure.

6. With your palms facing up, place the **pointer** and **middle** finger into the **latch handle sides**.

Figure 73: Latch Handle Identification

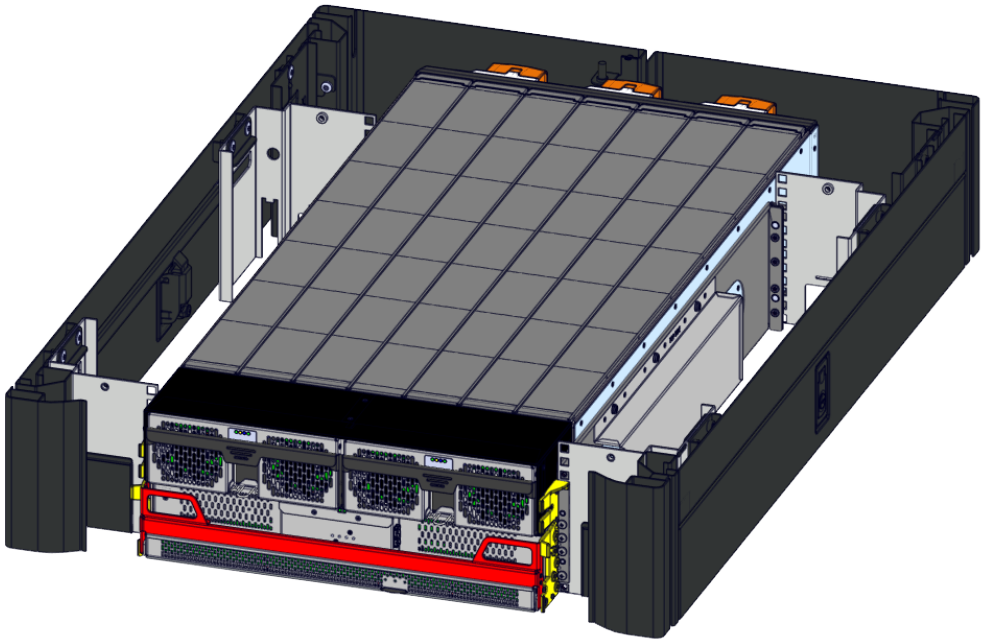

Note:

- Latch handle marked in red.
- Rack ears marked in yellow.

7. With your **thumbs** on the **rack ears**, pull the **latch handle sides** and push on the **rack ear release**.

8. Pull the **latch handle** until clear of the **rack ear latch**.

Figure 74: Latch Handle Clear of Rack Ear

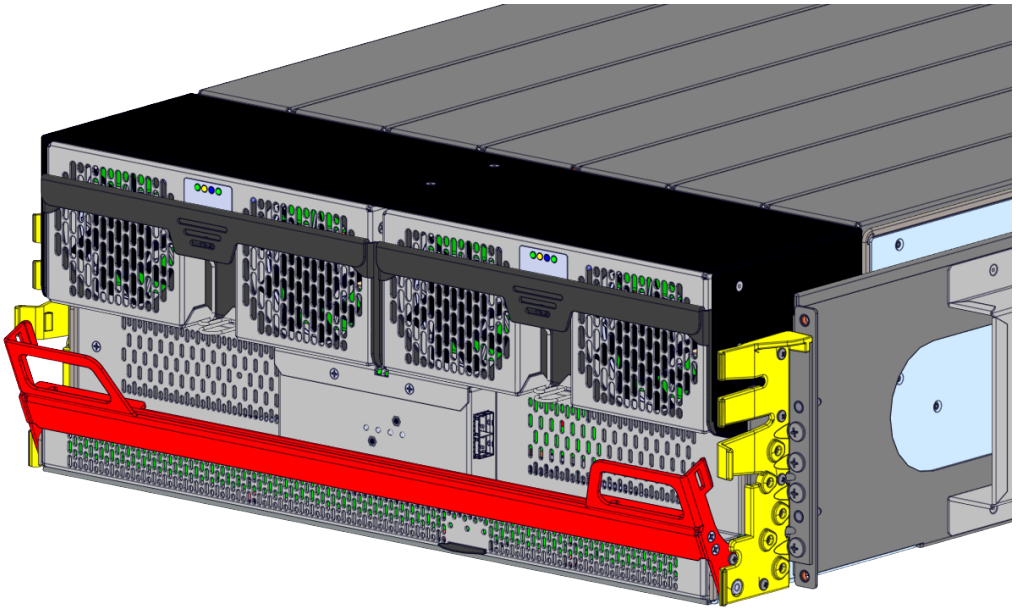

Note:

- Latch handle marked in red.
- Rack ears marked in yellow.

---

9. Completely remove the miniSAS cables and power cords from the I/O canister.
10. With your palms facing up, reposition your hands so that your thumbs are on the outside and your fingers cradle the bottom and rear of either side of the I/O canister .
11. Slowly pull the **I/O canister** away from the **chassis**.

Warning: The I/O canister is very back-heavy. Ensure that you are fully supporting the component during the removal.

Figure 75: Removing the I/O Canister

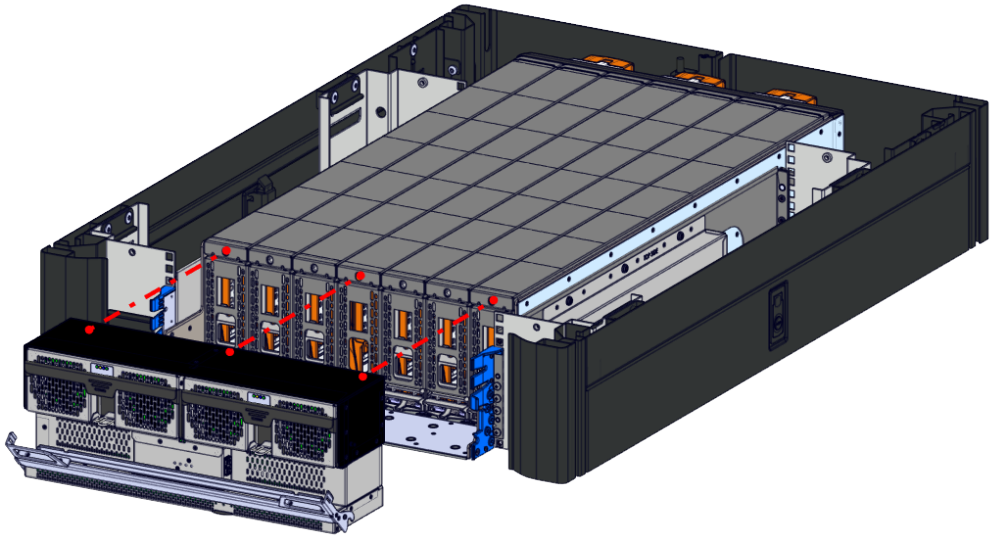

12. Install the replacement I/O canister in the reverse order that you removed it.

    Note: Ensure that the I/O canister is center properly and press firmly to ensure you are able to latch the replacement I/O canister.

13. Reconnect the miniSAS cables.
14. Reconnect the power cords.
15. Verify that the two green LEDs on both power supply A and B are illuminated.
16. Power on the Storage Node.

    The power button is located on the chassis front control panel.

A Troubleshooting

Topics:

This chapter provides troubleshooting tips.

A.1 Field Replaceable Units

Problem	Recommended Action
The PostgreSQL partition has failed, or a NIC has failed on the Management Node.	Fail over the CMC. **Warning:** When you are upgrading your setup, do not execute a failover. First complete the upgrade before you start the failover. To execute a failover, follow the instructions in *Managing Hardware* in the *HGST Active Archive System Administration Guide*.
The wrong disk was replaced.	If you accidentally replace the wrong disk, it shows up in the CMC as an unmanaged disk. An unmanaged disk is a newly installed disk that the Active Archive System cannot determine a purpose for (in other words, whether it is a replacement disk or really a new disk). **Warning:** Adding disks to the Active Archive System or changing the configuration of any hardware in the Active Archive System is not supported. Please contact HGST Support for more information. Correct this problem as follows: 1. Physically remove the new disk, and replace it with the disk that was accidentally removed. 2. In the CMC, navigate to **Dashboard** > **Administration** > **Hardware** > **Disks** > **Unmanaged**. 3. Select the new disk, and in the **Commands** pane, click **Delete**. When you first remove the disk through the CMC, the disk will most likely be added again by the monitoring agent before you can actually remove the disk from the node. If this happens, repeat the steps above to delete the disk again.
You shut down a node in order to replace it or something in it, but when you powered on the new/fixed node, it did not boot or was not detected by the CMC.	Connect a monitor to the node's VGA port, and a keyboard to its USB port. Restart the node. Observe any error messages that it outputs.

Active Archive System Glossary

A

AC	**Alternating Current**
ACMA	**Australian Communications and Media Authority**

Top of A | Top of Glossary

B

BIOS	**Basic Input/Output System**
BIS	**Business Information System**
BIST	**Built-In Self-Test**
BMC	**Baseboard Management Controller**
BOM	**Bill of Materials**
BSMI	**Bureau of Standards, Metrology and Inspection**

Top of B | Top of Glossary

C

CDB	**Computer Data Bus**
CLI	**Command Line Interface**
CS	**Climate Saver**

Top of C | Top of Glossary

D

DC	**Direct Current**

Top of D | Top of Glossary

E

EC	**Engineering Change**
EEPROM	**Electrically Erasable Programmable Read-Only Memory**
EMC	Electromagnetic Compatibility
EMI	**Electromagnetic Interference**
ESD	**Electrostatic Discharge**
EVPD	**Enable Vital Product Data**

Top of E | Top of Glossary

F

FCC	**Federal Communications Commission**
FRU	**Field Replaceable Unit**
FW	**Firmware**

Top of F | Top of Glossary

G

GBE	**Gigabit Ethernet**
GPIO	**General-Purpose Input/Output**
GUI	**Graphical User Interface**

Top of G | Top of Glossary

H

HD	**Hard Drive**
HDD	**Hard Disk Drive**

Top of H | Top of Glossary

I

ICT	**In-circuit Test**
IEC	**International Electrotechnical Commission**
I/O	**Input/Output**
IOC	**Input/Output Controller**
IOM	**I/O Module**
IPMI	**Intelligent Platform Management Interface**

Top of I | Top of Glossary

J

JBOD	**Just a Bunch of Disks**

Top of J | Top of Glossary

K

KVALITET	Spell out acronym here

Top of K | Top of Glossary

L

LED	**Light-Emitting Diode**
LPC	**Low Pin Count**
LPH	**Low Profile Hybrid**
LUN	**Logical Unit Number**

Top of L | Top of Glossary

M

MAC	**Media Access Control**
miniSAS	**Mini Statistical Analysis System**

Top of M | Top of Glossary

O

OS	**Operating System**
OUI	**Organizationally Unique Identifier**

Top of O | Top of Glossary

P

PCB	**Printed Circuit Boards**
PCI	**Peripheral Component Interconnect**
PDB	**Power Distribution Board**
PDU	**Power Distribution Unit**
PMBus	**Power Management Bus**
POST	**Power On Self Test**
PSU	**Power Supply Unit**
PHY	**Physical Layer**
PWM	**Pulse-Width Modulation**
PWR	**Power**

Top of P | Top of Glossary

R

ROC	**Recovery Oriented Computing**
RTC	**Real Time Clock**

Top of R | Top of Glossary

S

SAS **Serial Attached SCSI**

SATA **Serial Advanced Technology Attachment**

SBB **Storage Bridge Bay**

SCSI **Small Computer System Interface**

SDK **Software Development Kit**

SEP **SCSI Enclosure Processor**

SES **SCSI Enclosure Services**

SMART **Self-Monitoring, Analysis and Reporting Technology**

SMP **Server Message Block**

SMB **Server Message Block**

SPI **Serial Peripheral Interface**

SSP **Serial SCSI Protocol**

Top of S | Top of Glossary

T

TCA **Telecommunications Computing Architecture**

Top of T | Top of Glossary

U

UART **Universal Asynchronous Receiver/Transmitter**

Top of U | Top of Glossary

V

VBOD **Virtualized Bunch of Disks**

VPD **Vital Product Data**

Top of V | Top of Glossary

W

WOL **Wake On LAN**

WOS **Wake On SAS**

Top of W | Top of Glossary

X

XDP **XML Data Package**

Top of X | Top of Glossary

Index

T

W