

# NVIDIA DGX A100 System

Service Manual

## Table of Contents

Chapter 1. Introduction	1
1.1. Customer-replaceable Components	1
1.2. Recommended Tools	2
1.3. Customer Support	2
1.4. Running the Pre-flight Test	2
Chapter 2. Front Fan Module Replacement	4
2.1. Front Fan Module Replacement Overview	4
2.2. Identifying the Failed Fan Module	4
2.3. Replacing and Returning the Front Fan Module	6
Chapter 3. Power Supply Replacement	8
3.1. Power Supply Replacement Overview	8
3.2. Identifying the Failed Power Supply	8
3.3. Replacing the Power Supply	10
Chapter 4. Motherboard Tray - Accessing in Place	12
4.1. Accessing the Motherboard Tray	12
4.2. Replacing the Motherboard Tray	14
Chapter 5. Motherboard Tray - Removal and Installation	17
5.1. Removing the Motherboard Tray	17
5.2. Reinstalling the Motherboard Tray	20
Chapter 6. U.2 NVMe Cache Drive Upgrade from 4 to 8	25
6.1. U.2 NVMe Cache Drive Upgrade Overview	25
6.2. Determining the NVMe Drive to Order	25
6.3. Installing the Optional NVMe Drives	26
Chapter 7. U.2 NVMe Cache Drive Replacement	29
7.1. U.2 NVMe Cache Drive Replacement Overview	29
7.2. Identifying the Failed U.2 NVMe	29
7.3. Replacing the U.2 NVMe Drive	31
Chapter 8. U.2 NVMe Cache Drive Upgrade to 7.68 TB Drives	33
8.1. U.2 NVMe Cache Drive 7.68 TB Upgrade Overview	33
8.2. Installing the 7.68 TB NVMe Drives	33
Chapter 9. U.2 NVMe Cache Drive Post-Installation Tasks	36
9.1. Recreating the Cache RAID 0 Volume	36
9.2. Returning the NVMe Drive	37

Chapter 10. M.2 NVMe Boot Drive Replacement	38
10.1. M.2 NVMe Boot Drive Replacement Overview	
10.2. Identifying the Failed M.2 NVMe	
10.3. Replacing the M.2 NVMe Drive	
10.4. Rebuilding the Boot Drive RAID 1 Volume	42
10.5. Returning the NVMe Drive	43
Chapter 11. M.2 Boot Drive Riser Assembly Replacement	
11.1. M.2 Boot Drive Riser Assembly Replacement Overview	
11.2. Determining a Failed M.2 NVMe Riser Assembly	
11.3. Replacing the M.2 NVMe Riser Assembly	
11.4. Returning the Riser Assembly	
Chapter 12. DIMM Replacement	
12.1. DIMM Replacement Overview	48
12.2. Identifying the Failed DIMM	48
12.3. Replacing the DIMM	49
Chapter 13. DIMM Upgrade	53
13.1. DIMM Upgrade Overview	53
13.2. Upgrading the DIMM	53
Chapter 14. Network Card Replacement	58
14.1. Network Card Replacement Overview	58
14.2. Identifying the Failed Network Card	58
14.3. Replacing the Vertical Network Card	
14.4. Replacing the Horizontal Network Card	60
Chapter 15. Adding the Optional Dual-port Horizontal Network Card	65
15.1. Dual-port Network Card Upgrade Overview	65
15.2. Adding the Horizontal Network Card	65
Chapter 16. Updating the Mellanox Network Card Firmware	70
Chapter 17. Front Console Board Replacement	71
17.1. Front Console Board Replacement Overview	71
17.2. Replacing the Front Console Board	
Chapter 18. Motherboard Tray Battery Replacement	74
18.1. Motherboard Tray Battery Replacement Overview	74
18.2. Replacing the Motherboard Tray Battery	74
Chapter 19. Trusted Platform Module Replacement	
19.1. Trusted Platform Module Replacement Overview	
' 19.2. Replacing the Trusted Platform Module	
-	

Chapter 20. Removing and Attaching the Bezel	. 81
Chapter 21. Installing the Rack Mount Kit	. 83
21.1. Installing the Rails	83
21.2. Installing the Cage Nuts	86

### List of Figures

Figure 1.	NVMe Drives: PCIe to Slot Mapping	30

### List of Tables

Table 1.	Network Car	rd Slot IDs	
----------	-------------	-------------	--

# Chapter 1. Introduction

This document contains instructions for replacing NVIDIA DGX<sup>™</sup> A100 system components. Be sure to familiarize yourself with the NVIDIA Terms & Conditions documents before attempting to perform any modification or repair to the DGX A100 system. These Terms & Conditions for the DGX A100 system can be found through the <u>NVIDIA DGX Systems Support</u> page.

Contact NVIDIA Enterprise Support to obtain an RMA number for any system or component that needs to be returned for repair or replacement. When replacing a component, use **only** the replacement supplied to you by NVIDIA.

## 1.1. Customer-replaceable Components

List of customer-replaceable components in the NVIDIA DGX A100.

You can obtain the following components for replacement in your data center.

- ▶ Fan Modules
- Power Supplies
- ▶ Cache (U.2) NVMe Drives
- Boot (M.2) NVMe Drive
- Boot (M.2) Riser Assembly
- ▶ DIMMs
- Network Card (vertical, single or dual port)
- Network Card (horizontal, dual port)
- ▶ Front Console Board
- ▶ Battery
- Trusted Platform Module (TPM)
- Bezel
- Rack Mount Kit

Contact NVIDIA Enterprise Support for replacement instructions and guidance for specific components if those instructions are not included in this document.

# 1.2. Recommended Tools

List of recommended tools needed to service the NVIDIA DGX A100.

- Laptop
- USB key with tools and drivers
- ▶ USB key imaged with the DGX Server OS ISO
- Screwdrivers (Phillips #1 and #2, small flat head)
- ► KVM Crash Cart
- Anti-static wrist strap
- Masking tape or label maker
- ▶ Tie wraps or VELCRO<sup>®</sup> Brand for cable management
- Box cutter
- Black Permanent Marker
- Packing materials

# 1.3. Customer Support

Contact NVIDIA Enterprise Support for assistance in reporting, troubleshooting, or diagnosing problems with your DGX A100 system. Also contact NVIDIA Enterprise Support for assistance in installing or moving the DGX A100 system.

For details on how to obtain support, visit the NVIDIA Enterprise Support web site (<u>https://www.nvidia.com/en-us/support/enterprise/</u>).

# 1.4. Running the Pre-flight Test

Instructions for running the DGX stress test.

NVIDIA recommends running the pre-flight stress test before putting a system into a production environment or after servicing. You can specify running the test on the GPUs, CPU, memory, and storage, and also specify the duration of the tests.

To run the tests, use NVSM.

#### Syntax:

\$ sudo nvsm stress-test [--usage] [--force] [--no-prompt] [<test>...] [DURATION]

For help on running the test, issue the following.

 $\sim$  \$ sudo nvsm stress-test --usage

#### Recommended Command

The following command runs the test on all supported components (GPU, CPU, memory, and storage), and takes approximately 20 minutes.

 $^{-\$}$  sudo nvsm stress-test --force

# Chapter 2. Front Fan Module Replacement

### 2.1. Front Fan Module Replacement Overview

This is a high-level overview of the steps needed to replace the front fan modules.

- 1. Identify the failed front fan module through the BMC or the fan module LED and submit a service ticket to NVIDIA Enterprise Support.
- 2. Get a replacement from NVIDIA Enterprise Support.
- 3. Remove the failed fan module using the fan numbering diagram as a reference.
- 4. Insert the new fan module.
- 5. Confirm the new fan module is working correctly through the BMC or NVSM. sudo nvsm show health
- 6. Return the bad fan module using the packaging from the new fan module.

# 2.2. Identifying the Failed Fan Module

There are several ways to determine the faulty fan module to replace.



### Viewing the Fan Module LED

Look for the lit fault LED on the upper right corner of the faulty fan module.



### Using the BMC Dashboard and NVSM

- 1. Identify the faulty fan module using the BMC dashboard.
  - a). Log on to the BMC.
  - b). Click **Sensor** from the left navigation menu, then review the **Normal Sensors** section.

♣ SPD_FAN_SYSD_F	3654 Rpm
♣ SPD_FAN_SYS0_R	3248 Rpm
♣ SPD_FAN_SYS1_F	3654 Rpm
♣ SPD_FAN_SYS1_R	3248 Rpm
♣ SPD_FAN_SYS2_F	3654 Rpm
♣ SPD_FAN_SYS2_R	3248 Rpm
♣ SPD_FAN_SYS3_F	3654 Rpm
♣ SPD_FAN_SYS3_R	3248 Rpm
♣ SPD_FAN_SYS4_F	4408 Rpm
♣ SPD_FAN_SYS4_R	3886 Rpm
♣ SPD_FAN_SYS5_F	4408 Rpm
♣ SPD_FAN_SYSS_R	3886 Rpm
♣ SPD_FAN_SYS6_F	4408 Rpm
♣ SPD_FAN_SYS6_R	3944 Rpm
♣ SPD_FAN_SYS7_F	4405 Rpm
♣ SPD_FAN_SYS7_R	3886 Rpm

There are two fans in the fan module, identified by SPD\_FAN\_SYSn\_F and SPD\_FAN\_SYSn\_R, where *n* is the module ID. If either fan fails, then the entire module must be replaced.

2. Use NVSM to confirm the fan issue.

#### \$ sudo nvsm show fans

In the output, look for the 'unhealthy' status for the same fan.

### 2.3. Replacing and Returning the Front Fan Module

- 1. Remove the new fan module from its packaging and be ready to install it.
- 2. Remove the failed fan module by pressing on the release button on the top of the module and pulling on the handle.



3. Quickly insert the new fan module, observing that the handle release mechanism is facing up.

**CAUTION:** Replace the fan module within 30 seconds to prevent overheating of the system components.

- 4. Confirm that the fan module is healthy working properly by
  - Verifying that the fan module fault LED is not lit.
  - Viewing the state of the fan module on he BMC dashboard.
  - Using NVSM (sudo nvsm show fans)
- 5. Use packaging to pack up the bad fan and follow the shipping instructions to return the bad fan to NVIDIA Enterprise Support.

# Chapter 3. Power Supply Replacement

This chapter describes how to replace one of the DGX A100 system power supplies (PSUs).

# 3.1. Power Supply Replacement Overview

This is a high-level overview of the steps needed to replace a power supply.

- 1. Identify failed power supply through the BMC and submit a service ticket.
- 2. Get replacement power supply from NVIDIA Enterprise Support.
- 3. Identify the power supply using the diagram as a reference and the indicator LEDs.
- 4. Remove the power cord from the power supply that will be replaced.
- 5. Replace the failed power supply with the new power supply.
- 6. Insert the power cord and make sure both LEDs light up green (IN/OUT).
- 7. Use the BMC to confirm that the power supply is working correctly.
- 8. Ship back the failed unit to NVIDIA Enterprise Support using the packaging provided.

## 3.2. Identifying the Failed Power Supply

### Identifying the Failed Power Supply from the Back

If physical access to the system is available, you can identify a failed PSU by the inspecting the LEDs on the power supply when the system is powered on.



Both LEDs should be solid green. If either of the LEDs are not green or if they are blinking, contact NVIDIA Enterprise Support to troubleshoot the issue.

### Identifying the Failed Power Supply from the Console

There are several ways to identify the failed PSU from the DGX A100 console.

Use the NVSM CLI as follows.

#### \$ sudo nvsm show psus

The output shows information for each PSU. Look for any that do not report Status\_Health=OK.

• View the PSU status from the BMC.

Click **Sensor** from the left side menu and inspect the PSU information from the **Normal Sensors** section.

PSU0 STATUS		Presence Detected
PSU1 STATUS		Presence Detected
PSU2 STATUS		Presence Detected
PSU3 STATUS		Presence Detected
PSU4 STATUS		Presence Detected
PSU5 STATUS		Presence Detected
B PWR_PSU0	182.000 Watts	
B PWR_PSU1	182.000 Watts	
B PWR_PSU2	351.000 Watts	
B PWR_PSU3	351.000 Watts	
B PWR_PSU4	260.000 Watts	
PWR_PSU5	351.000 Watts	
I TEMP_PSU0	26.000 °C	
I TEMP_PSU1	27.000 °C	
L TEMP_PSU2	30.000 °C	
L TEMP_PSU3	29.000 °C	
L TEMP_PSU4	29.000 °C	
L TEMP_PSU5	27.000 °C	

Use ipmitool.

#### \$ sudo ipmitool sdr |grep -i psu

Look for power supplies with no temperature reading or an output reading close to or equal to zero.

Both NVSM and the BMC identify each power supply as PSUx, where x is from 0 to 5. The following diagram shows the physical location of each PSU.



#### Determining the Manufacturer

Important: All PSUs in the system must be from the same manufacturer.

Issue the following to determine the PSU manufacturer.

\$ sudo nvsm show /chassis/localhost/power/PSUX

Where *X* corresponds to the PSU identifier. The following examples uses PSU0, and shows that the manufacturer is "Delta".

```
$ sudo nvsm show /chassis/localhost/power/PSU0
```

```
/chassis/localhost/power/PSU0
Properties:
FirmwareVersion = 01.05.01.05.01.05
LastPowerOutputWatts = 312
Manufacturer = Delta
MemberId = PSU0
Model = ECD16010092
Name = PSU0
Oem PSU Error = <NOT SET>
PowerSupplyType = AC
SerialNumber = DTHTCP200807M
Status Health = OK
Status State = Present
Targets:
Verbs:
cd
show
```

Obtain the replacement PSU (of the same manufacturer) from NVIDIA Enterprise Support.

## 3.3. Replacing the Power Supply

- 1. Be sure you have obtained the replacement PSU and that you have saved the packaging to use when sending back the failed PSU.
- 2. Determine whether you need to shut down the system.
  - If the three remaining PSUs are working and energized, then you do not need to shut down power to the DGX A100 system.
  - If fewer than three PSUs are working and energized, then shut down power to the DGX A100 system.
- 3. Unlock the power cord and then unplug it from the PSU to be replaced.

You may need to dislodge the power cord from the retaining clip.

- 4. Remove the PSU.
  - a). Push on the green tab to release the lock.



b). Pull on the black handle to remove the PSU from the chassis.



- 5. Install the new power supply.
  - a). Insert the new power supply into the chassis and push it all the way in, making sure that the green locking mechanism engages.
  - b). Plug in the power cord and lock it in place.
  - c). If needed, power on the system.
- 6. Confirm the installation by
  - ▶ Viewing the PSU status from the BMC dashboard->**Sensors** page.
  - Running nvsm show health to confirm all power supplies are healthy.

Pack the old power supply and ship it back to NVIDIA Enterprise Support.

# Chapter 4. Motherboard Tray -Accessing in Place

You will need to access the motherboard tray in order to service the following components. This process provides access to the motherboard components while the motherboard remains attached to the server.

- M.2 NVMe drives
- M.2 module riser card
- Motherboard tray battery
- Single-port or dual port CX-6 PCI network adapter card

## 4.1. Accessing the Motherboard Tray

1. Loosen the two motherboard thumbscrews and then pull the handles out to eject the motherboard tray.





2. Pull the motherboard tray out of the system until it locks, then loosen the two thumbscrews holding the lid in place.



3. Lift the rear section of the motherboard lid.



## 4.2. Replacing the Motherboard Tray

1. Close the lid to the motherboard tray.



2. Tighten the two thumbscrews and then push the motherboard tray into the system.



3. Close the handles to secure the motherboard tray in place.



4. Tighten the motherboard tray thumbscrews to complete the motherboard insertion.



# Chapter 5. Motherboard Tray -Removal and Installation

You will need to completely remove the motherboard tray from the server in order to service the following components.

DIMMs (either adding or replacing)

# 5.1. Removing the Motherboard Tray

1. Loosen the two motherboard thumbscrews and then pull the handles out to eject the motherboard tray.





2. Pull the motherboard tray out of the system until it locks, then press the two buttons on the top of the lid to release the tray and finish pulling the tray out of the system.



Place the tray on a solid, flat work surface.

3. Loosen two rear thumbscrews on the motherboard lid.



4. Loosen the two front thumbscrews on the motherboard tray lid.



5. Lift the lid off of the tray and set aside.



6. Remove all three air baffles to allow access to the DIMMs.



# 5.2. Reinstalling the Motherboard Tray

1. Reinstall the three air baffles.



Replace and secure the lid.
 a). Install the lid.



b). Tighten the rear thumbscrews



c). Tighten the front thumbscrews.



3. Slide the motherboard tray into the slot, open the tray handles, and then continue pushing the motherboard tray in.



4. Close the handles to secure the motherboard tray in place.



5. Tighten the motherboard tray thumbscrews to complete the motherboard insertion.



# Chapter 6. U.2 NVMe Cache Drive Upgrade from 4 to 8

If you require more drive cache capacity, you can increase the cache by adding four more NVMe drives to the DGX A100 system.



**Important:** Upgrade using only U.2 NVMe drives of the same density (capacity) as the existing four drives.

### 6.1. U.2 NVMe Cache Drive Upgrade Overview

This is a high-level overview of the steps needed to upgrade the DGX A100 system's cache size.

- 1. Identify the density (capacity) of the currently installed NVMe drives.
- 2. Place an order for additional four NVME drives from NVIDIA Sales.
- 3. Power off the system.
- 4. Remove the air baffles in the blank drive slots.
- 5. Install the NVMe drives in the DGX A100 system.
- 6. Power on the system.
- 7. Re-initialize the /raid filesystem to recognize all eight drives.
- 8. Confirm the system is healthy by running nvsm show health.

# 6.2. Determining the NVMe Drive to Order

Identify the drives in the RAID volume by issuing the following.
 \$ sudo nvme list

2. Determine the drive size (3.84 TB or 7.68 TB) of the currently installed drives, and order drives of same density (capacity) from the NVIDIA Sales team.

**Note:** If 3.84 TB drives are installed but you want to use or add 7.68 TB drives, refer to <u>U.2</u> <u>NVMe Cache Drive Upgrade to 7.68 TB Drives</u> for instructions.

# 6.3. Installing the Optional NVMe Drives

- 1. Be sure you have obtained the additional drives.
- 2. Back up any critical data to a network shared volume or some other means of backup.
- 3. Power off the system using the power button or through the console.
- 4. Remove the blank filler modules from slots 1, 3, 5, and 7.



a). Push the lever release button (on the right side of the lever) to unlock the lever.



b). Pull the lever to remove the blank filler module.



5. Move the drives in slots 4 and 6 and install in slots 1 and 3.

a). Push the lever release button (on the right side of the lever) to unlock the lever.



b). Pull the lever to remove the module.



- c). Unlock the release lever and then slide the drive into the slot **until the front face is flush with the other drives**.
- d). Close the lever and lock it in place.



- 6. Install the new NVMe drives in slots 4, 5, 6, and 7.
  - a). Unlock the release lever and then slide the drive into the slot **until the front face is flush with the other drives**.
  - b). Close the lever and lock it in place.



7. Power on the system.

Perform the tasks describes in the chapter <u>U.2 NVMe Cache Drive Post-Installation Tasks</u>.

# Chapter 7. U.2 NVMe Cache Drive Replacement



Replace only with U.2 NVMe drives of the same manufacturer, model, and density (capacity) as

## 7.1. U.2 NVMe Cache Drive Replacement Overview

This is a high-level overview of the procedure to replace a cache Non-Volatile Memory Express (NVMe) drive.

1. Identify the failed U.2 NVMe drive.

the existing or replaced drives.

- 2. Order a replacement from NVIDIA Enterprise Support.
- 3. Use nvsm to prepare the drive for removal look for the white LED.
- 4. Replaced the failed NVMe drive.
- 5. Rebuild the RAID volume and remount the /raid partition.
- 6. Confirm the system is healthy by running nvsm show health.
- 7. Ship the failed unit back to NVIDIA Enterprise Support using the provided packaging.

# 7.2. Identifying the Failed U.2 NVMe

### Identifying the Failed NVMe from the Front

If physical access to the system is available, you can identify a failed drive by the illuminated amber LED .





Ready to remove (WHITE) Drive activity (GREEN) Fault (AMBER)

### Identifying the Failed NVMe from the Console

To identify the failed NVMe drive from the DGX A100 console, enter the following and then look for drive alerts in the output to identify the failed drive.

#### \$ sudo nvsm show health

The command returns the PCIe bus ID. Refer to the following figure to find the slot ID that corresponds to the PCIe bus ID for the faulty drive.

#### Figure 1. NVMe Drives: PCIe to Slot Mapping



Alternatively, you can use the BMC dashboard to access the Sensor screen, the IPMI event log, and the System log to identify issues with the U.2 drives.

**Note:** The PCIe bus IDs for slots 6 and 7 depend on the firmware version.

### Identifying the NVMe Manufacturer and Model

Enter the following, replacing x with the number corresponding to the Linux device name for the failed drive.

```
$ sudo nvsm show /systems/localhost/storage/drives/nvmeXn1
```

#### Example output:

```
/systems/localhost/storage/drives/nvme5n1
Properties:
    Capacity = 3840755982336
```
```
BlockSizeBytes = 7501476528
SerialNumber = 174719FCF9F1
PartNumber = N/A
Model = Micron 9200 MTFDHAL3T8TCT
Revision = 100\overline{0}07H0
Manufacturer = Micron Technology Inc
Status_State = Enabled
Status_Health = OK
Name = Non-Volatile Memory Express
MediaType = SSD
IndicatorLED = N/A
EncryptionStatus = N/A
HotSpareType = N/A
Protocol = NVMe
NegotiatedSpeedsGbs = 0
Id = 5
```

Determine the manufacturer and model from the 'Model' entry in the output, and then request a replacement NVMe from NVIDIA Enterprise Support, specifying this information.

## 7.3. Replacing the U.2 NVMe Drive

- 1. Be sure you have requested and obtained the replacement drive from NVIDIA Enterprise Support.
- 2. Back up any critical data to a network shared volume or some other means of backup.
- 3. Power off the system using the power button.
- 4. Remove the NVMe drive.
  - a). Push the lever release button (on the right side of the lever) to unlock the lever.



b). Pull the lever to remove the module.



- 5. Replace the new NVMe drive in the same slot.
  - a). Unlock the release lever and then slide the drive into the slot **until the front face is flush with the other drives**.
  - b). Close the lever and lock it in place.



6. Power on the system.

Perform the tasks describes in the chapter <u>U.2 NVMe Cache Drive Post-Installation Tasks</u>.

## Chapter 8. U.2 NVMe Cache Drive Upgrade to 7.68 TB Drives

If you require more drive cache capacity, you can increase the cache by replacing the existing 3.84 TB drives with 7.68 TB drives.

#### 8.1. U.2 NVMe Cache Drive 7.68 TB Upgrade Overview

This is a high-level overview of the steps needed to upgrade the DGX A100 system's cache size.

- 1. Place an order for the 7.68 TB U.2 NVMe drives from NVIDIA Sales.
- 2. Power off the system.
- 3. Remove the existing components.
  - ▶ Remove all 3.84 TB cache drives.
  - If you are also upgrading from four to eight cache drives, then remove the air baffles in the blank drive slots.
- 4. Install the 7.68 TB NVMe drives in the DGX A100 system.
- 5. Power on the system.
- 6. Re-initialize the /raid filesystem to recognize all eight drives.
- 7. Confirm the system is healthy by running nvsm show health.

## 8.2. Installing the 7.68 TB NVMe Drives

- 1. Be sure you have obtained the additional drives.
- 2. Back up any critical data to a network shared volume or some other means of backup.
- 3. Power off the system using the power button or through the console.
- 4. Remove all 3.84 TB NVMe drives.

**Important:** You must remove all 3.84 TB drives so that there is no mix of 3.84 TB and 7.68 TB drives in the same system.

a). Push the lever release button (on the right side of the lever) to unlock the lever.



b). Pull the lever to remove the module.



5. (Optional) If you are also upgrading from four to eight drives, remove the blank filler modules from slots 1, 3, 5, and 7.



a). Push the lever release button (on the right side of the lever) to unlock the lever.



b). Pull the lever to remove the blank filler module.



- 6. Install the 7.68 TB NVMe drives.
  - a). Unlock the release lever and then slide the drive into the slot **until the front face is flush with the other drives**.
  - b). Close the lever and lock it in place.



7. Power on the system.

Perform the tasks describes in the chapter <u>U.2 NVMe Cache Drive Post-Installation Tasks</u>.

## Chapter 9. U.2 NVMe Cache Drive Post-Installation Tasks

This chapter describes the tasks that are typically needed after replacing a U.2 NVME drive or upgrading from 4 to 8 drives.

## 9.1. Recreating the Cache RAID 0 Volume

- 1. Power on the system and log in.
- 2. Confirm that all expected drives are visible.

```
$ sudo nvme list
```

Two boot drives and either four or eight cache drives should be visible, depending on how many are installed in the system.

3. If the previously installed cache drives (which are capable of self-encryption) are locked with an access key, then unlock them.

```
$ sudo nv-disk-encrypt disable
```

4. Re-create the cache volume and the /raid filesystem.

```
$ sudo nvsm start /systems/localhost/storage/volumes/rebuild
```

Output (with appropriate responses to prompts)

- 5. If the volume is to be locked with an access key, then re-enable drive encryption.
  - ▶ To enter drive passwords yourself, issue the following.

#### \$ sudo nv-disk-encrypt init -g

The software prompts you to enter a password for the vault, and then a password for each eligible SED.

Passwords must consist of only upper-case letters, lower-case letters, digits, and/or the following special-characters: ~ :  $0 \% ^ + = _$ ,

\$ sudo nv-disk-encrypt lock

 To allow the encryption software to randomly generated the passwords, issue the following.

\$ sudo nv-disk-encrypt init -k <your-vault-password> -g -r

The vault password must consist of only upper-case letters, lower-case letters, digits, and/or the following special-characters: ~ :  $0 \% ^ + = _$ ,

\$ sudo nv-disk-encrypt lock

\$ sudo nvsm enable /systems/localhost/storage/volumes/mdl/encryption

- 6. Issue the following to confirm the volume is healthy and that the system is healthy.
  - \$ sudo nvsm show storage
  - \$ sudo nvsm show health

Make sure that the drive firmware is up to date. Refer to the <u>DGX A100 firmware release notes</u> for information on the latest firmware for the U.2 NVMe drive.

### 9.2. Returning the NVMe Drive

Use the packaging from the new drive and follow the instructions that came with the package to ship the old drive back to NVIDIA Enterprise Support.

**Note:** If your organization has purchased a media retention policy, you may be able to keep failed drives for destruction. Check with NVIDIA Enterprise Support on the status of the policy for specifics.

## Chapter 10. M.2 NVMe Boot Drive Replacement

#### 10.1. M.2 NVMe Boot Drive Replacement Overview

This is a high-level overview of the procedure to replace a boot drive.

- 1. With the help of NVIDIA Enterprise Support, determine which M.2 drive needs to be replaced.
- 2. Get replacement from NVIDIA Enterprise Support.
- 3. Power down the system.
- 4. Label all cables and unplug them from the motherboard tray.
- 5. Slide motherboard out until it locks in place.
- 6. Open rear compartment and pull out the M.2 riser card with both M.2 disks attached.
- 7. Replace the failed M.2 device on the riser card.
- 8. Install the M.2 riser card with both M.2 disks.
- 9. Close the rear motherboard compartment and then slide the motherboard back into the system.
- 10. Plug in all cables using the labels as a reference.
- 11. Power on the system.
- 12. Confirm the M.2 RAID 1 mirror is synchronizing.
- 13. Ship back the failed unit to NVIDIA Enterprise Support using the packaging provided.

### 10.2. Identifying the Failed M.2 NVMe

The DGX A100 system automatically sets the failed M.2 drive offline when it detects the failure.

- 1. Identify which of the M.2 drives has failed (nvme0n1 or nvme1n1).
- \$ sudo nvsm show health
- 2. You can confirm this by issuing the following.

#### \$ sudo mdadm -D /dev/md0

Normally, the output would show both drives (nvme0 and nvme1) in an active sync state. The following example output shows only nvme1 in active sync, indicating that nvme0n1 is the failed drive.

Number Major Minor RaidDevice State 0 259 2 0 active sync /dev/nvmeln1p2 - 0 0 1 removed

3. Make a note of the device name for the failed drive (nvme0 or nvme1) and the device name for the good drive (nvme0 or nvme1).

You will need this information when rebuilding the RAID 1 array after replacing the drive.

4. Obtain the replacement from NVIDIA Enterprise Support.

### 10.3. Replacing the M.2 NVMe Drive

Before attempting to replace one of the M.2 NVMe drives, be sure to have performed the following:

- Determined the location ID of the faulty M.2 NVMe drive.
- Obtained the replacement M.2 NVMe drive and have saved the packaging for use when returning the faulty drive.



**CAUTION**: **Static Sensitive Devices**: - Be sure to observe best practices for electrostatic discharge (ESD) protection. This includes making sure personnel and equipment are connected to a common ground, such as by wearing a wrist strap connected to the chassis ground, and placing components on static-free work surfaces.

- 1. Back up any critical data to a network shared volume or some other means of backup.
- If not already done, mark the drive as failed, then remove the failed drive from the array by issuing the following (replacing x with the failed drive identifier 0 or 1).
   \$ sudo mdadm --manage /dev/md0 --fail /dev/nvmeXn1
  - \$ sudo mdadm --manage /dev/md0 --remove /dev/nvmeXn1
- 3. Power down the system.
- 4. Label all network, monitor, and USB cables connected to the motherboard tray for easy identification when reconnecting.
- 5. Unplug all power cords, and all network, monitor, and USB cables.
- 6. Remove the motherboard tray.

Refer to the instructions in the section <u>Accessing the Motherboard Tray</u>.

7. Remove the M.2 riser card from the motherboard tray by lifting the riser assembly.



8. Identify the failed M.2 module and remove it from the riser card.a). Determine the location of the drive to remove.



b). Using a Phillips #1 screwdriver, loosen the black screw that secures the drive in place.



**Note:** The screw is not a captive screw and can drop. Be careful when loosening the screw to avoid dropping and losing the screw.



c). Carefully lift the M.2 drive away from the riser.



d). Pull the drive to disconnect from the connector on the riser board, then insert the new drive into the connector on the riser board.



e). Place the drive against the card and secure by tightening the screw using a Phillips #1 screwdriver.



9. Install the assembled module on the motherboard by inserting the riser card in its slot.



10. Close the motherboard tray lid and then install the motherboard tray.

Refer to the instructions in the section <u>Replacing the Motherboard Tray</u>.

11. Connect all the cables to the motherboard tray, connect all power cords, then power on the system.

Rebuild the RAID 1 array according to the instruction in the section <u>Rebuilding the Boot Drive</u> <u>RAID 1 Volume</u>.

#### 10.4. Rebuilding the Boot Drive RAID 1 Volume

After replacing a faulty M.2 OS drive, you must rebuild the RAID 1 array.

- 1. If you have not already done so, boot the DGX A100 system and log in.
- 2. Rebuild the boot drive mirror.

In the following steps, replace X with the number that corresponds to the replaced drive. If you did not note this information when identifying the failed drive, then follow the instructions in the first step of <u>Identifying the Failed M.2 Drive</u>.

a). Start the rebuild process.

```
\$ sudo nvsm start /systems/localhost/storage/volumes/md0/rebuild/
```

b). Enter the device name of the spare (replaced) drive when prompted - either nvme0n1 or nvme1n1, depending on which drive was replaced.

After entering **y** at the prompt to start the RAID 1 rebuild, the "Initiating rebuild ..." message appears.

```
/systems/localhost/storage/volumes/md0/rebuild started at 2018-10-12
15:27:26.525187
Initiating RAID-1 rebuild on volume md0...
0.0% [\ ]
```

If this message remains at "Initiating RAID-1 rebuild" for more than 30 seconds, then there is a problem with the rebuild process. In this case, make sure the name of the replacement drive is correct and try again.

The RAID 1 rebuild process should take about 1 hour to complete.

Make sure that the drive firmware is up to date, Refer to the <u>DGX A100 firmware release notes</u> for information on the latest firmware for the M.2 NVMe drive.

#### 10.5. Returning the NVMe Drive

Use the packaging from the new drive and follow the instructions that came with the package to ship the old drive back to NVIDIA Enterprise Support.

# **Note:** If your organization has purchased a media retention policy, you may be able to keep failed drives for destruction. Check with NVIDIA Enterprise Support on the status of the policy for specifics.

## Chapter 11. M.2 Boot Drive Riser Assembly Replacement

This chapter applies when both M.2 OS drives need to be replaced. In that case, a replacement riser assembly (which includes both M.2 NVMe drives) should be ordered.

### 11.1. M.2 Boot Drive Riser Assembly Replacement Overview

This is a high-level overview of the procedure to replace the boot drive riser assembly.

- 1. Confirm that both M.2 drives cannot be reached and need to be replaced.
- 2. Get replacement M.2 riser assembly from NVIDIA Enterprise Support.
- 3. Power down the system.
- 4. Label all motherboard tray cables and unplug them.
- 5. Slide out the motherboard tray and open the motherboard tray lid.
- 6. Pull out the M.2 riser assembly with both M.2 disks attached.
- 7. Install the new M.2 riser assembly with both M.2 disks attached..
- 8. Close the lid on the motherboard tray.
- 9. Slide the motherboard tray into the system.
- 10. Plug in all cables using the labels as a reference.
- 11. Power on the system.
- 12. Re-install the OS and confirm the system is healthy.
- 13. Ship back the failed unit to NVIDIA Enterprise Support using the packaging provided.

#### 11.2. Determining a Failed M.2 NVMe Riser Assembly

The following are the conditions for which NVIDIA Enterprise Support may instruct the M.2 riser assembly be replaced:

- ▶ The DGX A100 cannot be booted.
- The boot drives cannot be seen from the SBIOS.
- The system indicates that the boot drives are not available when booting from the ISO image.
- Loss of communication with the M.2 boot drives.
- ▶ The M.2 riser assembly was damaged.

#### 11.3. Replacing the M.2 NVMe Riser Assembly

Before attempting to replace the M.2 NVMe riser assembly, be sure you have obtained the replacement assembly and have saved the packaging for use when returning the faulty riser assembly.



**CAUTION**: **Static Sensitive Devices**: - Be sure to observe best practices for electrostatic discharge (ESD) protection. This includes making sure personnel and equipment are connected to a common ground, such as by wearing a wrist strap connected to the chassis ground, and placing components on static-free work surfaces.

- Power down the system.
   You will likely need to use the BMC console.
- 2. Label all cables connected to the motherboard tray for easy identification when reconnecting.
- 3. Remove the motherboard tray.

Refer to the instructions in the section <u>Accessing the Motherboard Tray</u>.

4. Remove the M.2 riser card from the motherboard tray by lifting the riser assembly.



5. Install the assembled module on the motherboard by inserting the riser card in its slot.



- Close the motherboard tray lid and then install the motherboard tray.
   Refer to the instructions in the section <u>Replacing the Motherboard Tray</u>.
- 7. Connect all the cables to the motherboard tray.
- Re-install the DGX OS server software.
   See the <u>DGX A100 User Guide</u> for detailed instructions.

## 11.4. Returning the Riser Assembly

Use the packaging from the new riser assembly and follow the instructions that came with the package to ship the old riser assembly back to NVIDIA Enterprise Support.

Ę

**Note:** If your organization has purchased a media retention policy, you may be able to keep failed drives for destruction. Check with NVIDIA Enterprise Support on the status of the policy for specifics.

## Chapter 12. DIMM Replacement

#### 12.1. DIMM Replacement Overview

This is a high-level overview of the procedure to replace a dual inline memory module (DIMM) on the DGX A100 system.

- 1. Use the nvsm health commands to identify the failed DIMM
- 2. Get a replacement DIMM from NVIDIA Enterprise Support.
- 3. Shut down the system.
- 4. Label all motherboard tray cables and unplug them.
- 5. Remove the motherboard tray and place on a solid flat surface.
- 6. Remove the motherboard tray lid.
- 7. Use the reference diagram on the lid of the motherboard tray to identify the failed DIMM.
- 8. Replace the bad DIMM with the new one.
- 9. Close the lid on the motherboard tray.
- 10. Insert the motherboard tray into the system.
- 11. Plug in all cables using the labels as a reference.
- 12. Power on the system.
- 13. Verify that all DIMMs are now healthy with nvsm.

### 12.2. Identifying the Failed DIMM

1. From the console, run the following nvsm command to identify memory alerts.

\$ sudo nvsm show /systems/localhost/memory/alerts

Alerts will appear under the Target section. For example. Targets:

alert0

2. Get specific information about the memory alert. The following example obtains information for alert0. \$ sudo nvsm show /systems/localhost/memory/alerts/alert0 Inspect the component\_id = line to determine the DIMM ID. The following example shows a DIMM ID of A1.

```
Properties:
system_name = ....
component_id = CPU1_DIMM_A1
...
```

The output provides other information about the alert that can be provided to NVIDIA Enterprise Support.

3. Determine the DIMM manufacturer.

\$ sudo nvsm show memory

4. Request the replacement DIMM from NVIDIA Enterprise Support, specifying the manufacturer.

## 12.3. Replacing the DIMM

Before attempting to replace any of the dual inline memory modules (DIMMs), be sure to have performed the following:

- Determined the location ID of the faulty DIMM needing replacement as explained in <u>Identifying the Failed DIMM</u>. The location ID is an alpha-numeric designator, such as A0, A1, B0, B1, etc.
- Obtained the replacement DIMM and have saved the packaging for use when returning the faulty DIMM.



**CAUTION**: **Static Sensitive Devices**: - Be sure to observe best practices for electrostatic discharge (ESD) protection. This includes making sure personnel and equipment are connected to a common ground, such as by wearing a wrist strap connected to the chassis ground, and placing components on static-free work surfaces.

- 1. Power down the system.
- 2. Label all cables connected to the motherboard tray for easy identification when reconnecting.
- 3. Remove the motherboard tray.

Refer to the instructions in the section <u>Removing the Motherboard Tray</u>.

4. Using the diagram label on the lid as a guide, locate the faulty DIMM to be replaced.



#### Top of diagram = front of server

- 5. Remove the DIMM.
  - a). Press down on the side latches at both ends of the DIMM socket to push them away from the DIMM. This should unseat the DIMM from the socket.



b). Pull the DIMM straight up to remove it from the socket.



- 6. Carefully insert the replacement DIMM.
  - a). Make sure the socket latches are open.
  - b). Position the DIMM over the socket, making sure that the notch on the DIMM lines up with the key in the slot, then press the DIMM down into the socket until the side latches click in place.
  - c). Make sure that the latches are up and locked in place.



7. Install the three air baffles, replace the motherboard tray lid and then install the motherboard tray.

Refer to the instructions in the section <u>Reinstalling the Motherboard Tray</u>.

- 8. Connect all the cables to the motherboard tray.
- 9. Install all the power cords.
- 10. Power on the system and log in.
- 11.Confirm that the system is healthy.
  - \$ sudo nvsm show health
  - \$ sudo nvsm show /systems/localhost/memory/alerts
    There should be no new alerts listed.
- 12. Ship the bad DIMM back to NVIDIA Enterprise Support.

## Chapter 13. DIMM Upgrade

You can upgrade the system memory by installing 16 additional DIMMs.

## 13.1. DIMM Upgrade Overview

This is a high-level overview of the procedure to add 16 additional dual-inline memory modules (DIMMs) on the DGX A100 system.

- 1. Obtain the memory upgrade (16 DIMMs) from NVIDIA Sales.
- 2. Shut down the system.
- 3. Label all motherboard tray cables and unplug them.
- 4. Remove the motherboard tray and place on a solid flat surface.
- 5. Remove the motherboard tray lid.
- 6. Use the reference diagram on the lid of the motherboard tray to identify the empty DIMM locations (spots with air baffles installed).
- 7. Replace the air baffles with the new DIMMs.
- 8. Close the lid on the motherboard tray.
- 9. Insert the motherboard tray into the system.
- 10. Plug in all cables using the labels as a reference.
- 11. Power on the system.
- 12. Verify that all DIMMs as well as the system are healthy using nvsm.

## 13.2. Upgrading the DIMM



**CAUTION**: **Static Sensitive Devices**: - Be sure to observe best practices for electrostatic discharge (ESD) protection. This includes making sure personnel and equipment are connected to a common ground, such as by wearing a wrist strap connected to the chassis ground, and placing components on static-free work surfaces.

- 1. Power down the system.
- 2. Label all cables connected to the motherboard tray for easy identification when reconnecting.

3. Remove the motherboard tray and air baffles.

Refer to the instructions in the section <u>Removing the Motherboard Tray</u>.

4. Using the diagram label on the lid as a guide, locate the DIMMs to be installed during the upgrade.





5. Remove the DIMM air baffles.

Press down on the side latches at both ends of the air baffle to eject the module from the slot, then pull the air baffle out of the slot.



6. Remove 8 DIMMs from CPU-1 slots I1, J1, K1, L1, M, N1, O1, and P1

Press down on the side latches at both ends of the DIMM to eject the module from the slot, then pull the DIMM out of the slot.



- 7. Carefully insert the 8 DIMMS that you just removed into CPU-0 slots A0, B0, C0, D0, E0, F0, G0, and H0.
  - a). Make sure the socket latches are open.
  - b). Position the DIMM over the socket, making sure that the notch on the DIMM lines up with the key in the slot, then press the DIMM down into the socket until the side latches click in place.



- c). Make sure that the latches are up and locked in place.
- 8. Install the new DIMMs from the upgrade kit to CPU-1 slots I0, I1, J0, J1, K0, K1, L0, L1, M0, M1, N0, N1, O0, O1, P0, and P1.
  - a). Make sure the socket latches are open.
  - b). Position the DIMM over the socket, making sure that the notch on the DIMM lines up with the key in the slot, then press the DIMM down into the socket until the side latches click in place.



- c). Make sure that the latches are up and locked in place.
- 9. Install the three motherboard air baffles, replace the motherboard tray lid and then install the motherboard tray.

Refer to the instructions in the section <u>Reinstalling the Motherboard Tray</u>.

- 10. Connect all the cables to the motherboard tray.
- 11. Install all the power cords.
- 12. Power on the system and log in.
- 13. Confirm that the total memory is now 2 TB.

\$ lsmem
Total online memory: 21

14. Confirm that the system is healthy.

\$ sudo nvsm show health

## Chapter 14. Network Card Replacement

#### 14.1. Network Card Replacement Overview

This is a high-level overview of the procedure to replace one or more network cards on the DGX A100 system.

- 1. Use the nvsm show commands to identify the failed network card.
- 2. Get a replacement card from NVIDIA Enterprise Support.
- 3. Shut down the system.
- 4. Label all motherboard tray cables and unplug them.
- 5. Remove the motherboard tray and open the lid.
- 6. Locate the failed network card and remove it.
- 7. Insert the new card into the slot and secure with the screw.
- 8. Close the lid on the motherboard tray, then insert the tray into the system.
- 9. Plug in all cables using the labels as a reference.
- 10. Power on the system.
- 11. Verify that the network card is healthy using nvsm.

#### 14.2. Identifying the Failed Network Card

Before attempting to replace any of the network cards, be sure to have performed the following:

1. Issue the following.

#### \$ sudo nvsm show health

2. Match the PCIe bus ID for the failed card with the slot ID using the following diagram.



Note the slot ID for ordering and replacing.

3. Order the appropriate card type as indicated in the following table, then follow the corresponding replacement instructions.

Table 1.Network Card Slot IDs

Slot ID	Card Type	Replacement Instructions
0-3, 6-9	Single port	Replacing the Vertical Network Card
4	Dual-port vertical	Replacing the Vertical Network Card
5	Dual-port horizontal	<u>Replacing the Horizontal Network</u> <u>Card</u>

## 14.3. Replacing the Vertical Network Card

Before attempting to replace any of the network cards, be sure to have performed the following:

Determined the location ID of the faulty network card needing replacement.

See Identifying the Failed Network Card.

Obtained the replacement network card and have saved the packaging for use when returning the faulty component.

![](_page_64_Picture_11.jpeg)

**CAUTION**: **Static Sensitive Devices**: - Be sure to observe best practices for electrostatic discharge (ESD) protection. This includes making sure personnel and equipment are connected to a common ground, such as by wearing a wrist strap connected to the chassis ground, and placing components on static-free work surfaces.

- 1. Power down the system.
- 2. Label all network, monitor, and USB cables connected to the motherboard tray for easy identification when reconnecting.
- 3. Unplug all power cords, and all network, monitor, and USB cables.

4. Remove the motherboard tray.

Refer to the instructions in the section <u>Accessing the Motherboard Tray</u>.

5. Lift the network card off the motherboard and replace with the new network card.

![](_page_65_Picture_4.jpeg)

6. Install the motherboard tray lid and then install the motherboard tray.

Refer to the instructions in the section <u>Replacing the Motherboard Tray</u>.

- 7. Connect all the cables to the motherboard tray, connect all power cords, then power on the system and log in.
- 8. Confirm that the system is healthy.

#### \$ sudo nvsm show health

There should be no new alerts listed.

Verify that the firmware is up to date according to the instructions in <u>Updating the Mellanox</u> <u>Network Card Firmware</u>.

#### 14.4. Replacing the Horizontal Network Card

Before attempting to replace any of the network cards, be sure to have performed the following:

Determined the location ID of the faulty network card needing replacement.

See Identifying the Failed Network Card.

 Obtained the replacement network card have saved the packaging for use when returning the faulty component.

![](_page_66_Picture_2.jpeg)

**CAUTION**: **Static Sensitive Devices**: - Be sure to observe best practices for electrostatic discharge (ESD) protection. This includes making sure personnel and equipment are connected to a common ground, such as by wearing a wrist strap connected to the chassis ground, and placing components on static-free work surfaces.

- 1. Power down the system.
- 2. Label all network, monitor, and USB cables connected to the motherboard tray for easy identification when reconnecting.
- 3. Unplug all power cords, and all network, monitor, and USB cables.
- 4. Remove the motherboard tray.

Refer to the instructions in the section <u>Accessing the Motherboard Tray</u>.

- 5. Unlock the horizontal network card.
  - a). Loosen the black thumbscrew that secures the PCIe card locking mechanism in place.

![](_page_66_Picture_11.jpeg)

b). Open the locking mechanism by turning 90 degrees or more.

![](_page_67_Picture_1.jpeg)

- 6. Replace the card.
  - a). Pull the network card out of the riser card slot.

![](_page_67_Picture_4.jpeg)

b). Replace the old network card with the new one.

![](_page_67_Picture_6.jpeg)

c). Install the network card into the riser card slot.

![](_page_68_Picture_1.jpeg)

- 7. Lock the network card in place.
  - a). Close the locking mechanism by turning it back into its slot.

![](_page_68_Picture_4.jpeg)

b). Tighten the black thumb screw to secure the card in place.

![](_page_69_Figure_1.jpeg)

8. Install the motherboard tray lid and then install the motherboard tray.

Refer to the instructions in the section <u>Replacing the Motherboard Tray</u>.

- 9. Connect all cables back into the network card ports.
- 10. Power on the system and log in.
- 11.Confirm that the system is healthy.

#### **\$ sudo nvsm show health**

There should be no new alerts listed.

Verify that the firmware is up to date according to the instructions in <u>Updating the Mellanox</u> <u>Network Card Firmware</u>.

## Chapter 15. Adding the Optional Dualport Horizontal Network Card

The DGX A100 comes with a vertical dual-port network card. You can expand the ports by adding a horizontal dual-port network card for slot 5.

#### 15.1. Dual-port Network Card Upgrade Overview

This is a high-level overview of the procedure to install the optional horizontal dual-port network card on the DGX A100 system.

- 1. Obtain the new card from NVIDIA Sales.
- 2. Shut down the system.
- 3. Label all motherboard tray cables and unplug them.
- 4. Slide out the motherboard tray until it locks and then open the lid.
- 5. Remove the PCI slot filler for slot 5.
- 6. Insert the new card into the slot and secure with the screw.
- 7. Close the lid on the motherboard tray, then insert the tray into the system.
- 8. Plug in all cables using the labels as a reference.
- 9. Power on the system.
- 10. Verify that the network card is healthy using nvsm show health.

## 15.2. Adding the Horizontal Network Card

Be sure you have obtained the horizontal dual-port network card.

![](_page_70_Picture_16.jpeg)

**CAUTION**: **Static Sensitive Devices**: - Be sure to observe best practices for electrostatic discharge (ESD) protection. This includes making sure personnel and equipment are

connected to a common ground, such as by wearing a wrist strap connected to the chassis ground, and placing components on static-free work surfaces.

- 1. Power down the system.
- 2. Label all network, monitor, and USB cables connected to the motherboard tray for easy identification when reconnecting.
- 3. Unplug all power cords, and all network, monitor, and USB cables.
- 4. Remove the motherboard tray.

Refer to the instructions in the section <u>Accessing the Motherboard Tray</u>.

- 5. Unlock the PCI EMI shield.
  - a). Loosen the black thumbscrew that secures the locking mechanism in place.

![](_page_71_Picture_9.jpeg)

b). Open the locking mechanism by turning 90 degrees or more.

![](_page_71_Picture_11.jpeg)

Remove the EMI shield.
 Slide the EMI shield to the left to release it from the PCI slot on the riser card.


7. Install the new network card, inserting it into the slot on the riser card.



8. Lock the network card in place.



a). Close the locking mechanism by turning it back into its slot.

b). Tighten the black thumb screw to secure the card in place.



9. Install the motherboard tray lid and then install the motherboard tray.

Refer to the instructions in the section <u>Replacing the Motherboard Tray</u>.

- 10. Connect all cables back into the network card ports.
- 11. Power on the system and log in.
- 12. Confirm that the system is healthy.

#### \$ sudo nvsm show health

There should be no new alerts listed.

Verify that the firmware is up to date according to the instructions in <u>Updating the Mellanox</u> <u>Network Card Firmware</u>.

### Chapter 16. Updating the Mellanox Network Card Firmware

After replacing or installing the Mellanox ConnectX cards, make sure the firmware on the cards is up to date.

- 1. Confirm the OS is updated to the latest version; this will ensure the latest tested and supported firmware is downloaded.
  - \$ sudo apt update
  - \$ sudo apt full-upgrade
- 2. Run the firmware update script as follows:

\$ sudo /opt/mellanox/mlnx-fw-updater/mlnx\_fw\_updater.pl

If the firmware version in the new card is more recent than the one available in the operating system, the firmware should be downgraded to a supported version by issuing the following:

- \$ sudo /opt/mellanox/mlnx-fw-updater/mlnx\_fw\_updater.pl --force-fw-update
- Reboot the system for the firmware update to take effect.
  \$ sudo reboot
- 4. Once the system comes up, make sure the firmware versions are all the same by issuing:
  \$ cat /sys/class/infiniband/mlx5\_\*/fw\_ver

# Chapter 17. Front Console Board Replacement

#### 17.1. Front Console Board Replacement Overview

This is a high-level overview of the procedure to replace the front console board on the DGX A100 system.

- 1. Unpack the new front console board.
- 2. Shut down the system.
- 3. Using a Phillips #2 screwdriver, release the captive screws.
- 4. Pull the front console board out of the system.
- 5. Insert the new front console board.
- 6. Tighten the screws.
- 7. Power on the system and confirm the ports work.
- 8. Ship the failed unit back to NVIDIA Enterprise Support using the provided packaging.

## 17.2. Replacing the Front Console Board

A front console board malfunction can be determined in a few ways.

- No display or connectivity occurring after plugging in a keyboard and monitor to the front of the system, or
- The USB ports do not work, or
- The front temperature sensor does not provide a temperature reading.

Raise a ticket with NVIDIA Enterprise Services to request a replacement.

When the new board arrives, unpack it and keep the packaging to use for sending back the old board.



**CAUTION**: **Static Sensitive Devices**: - Be sure to observe best practices for electrostatic discharge (ESD) protection. This includes making sure personnel and equipment are connected to a common ground, such as by wearing a wrist strap connected to the chassis ground, and placing components on static-free work surfaces.

- 1. Power down the system.
- 2. Remove the bezel.
- 3. Replace the front console board.
  - a). Using a Phillips #2 screwdriver, loosen the two captive screws that secure the front console board.



b). Replace the front console board.



c). Tighten the screws.



- 4. Confirm functionality.
  - a). Power on the system.
  - b). Issue the following to confirm the temperature sensor is working properly. \$ sudo nvsm show health
- 5. Return the old module to NVIDIA Enterprise Services.

# Chapter 18. Motherboard Tray Battery Replacement

#### 18.1. Motherboard Tray Battery Replacement Overview

This is a high-level overview of the procedure to replace the DGX A100 system motherboard tray battery.

- 1. Get a replacement battery type CR2032.
- 2. Shut down the system.
- 3. Label all motherboard cables and unplug them.
- 4. Slide out the motherboard tray.
- 5. Open the motherboard tray IO compartment.
- 6. Pull out the M.2 riser card with both M.2 disks attached.
- 7. Pull out the dual port Mellanox card (slot 4).
- 8. Replace the battery on the motherboard.
- 9. Install the dual port Mellanox card (slot 4).
- 10. Install the M.2 riser card with both M.2 disks attached.
- 11. Close the lid IO compartment on the motherboard tray.
- 12. Slide the motherboard tray into the system.
- 13. Plug in all cables using the labels as a reference.
- 14. Power on the system.
- 15. Confirm the system is healthy by running  $\mathtt{nvsm}$  show <code>health</code>

#### 18.2. Replacing the Motherboard Tray Battery

A battery failure can be determined in a few ways.

- "Invalid configuration" will appear on your screen.
- Setup appears on your screen before booting.
- "Press F1 to continue" appears on the console.
- A Clock Error or Clock Message appears on your screen.
- The system clock loses time and date.

Call NVIDIA Enterprise Support to confirm that the battery is the right component to replace. The CR2032 battery is not provided by NVIDIA, but can be purchased from a convenience store.



**CAUTION**: **Static Sensitive Devices**: - Be sure to observe best practices for electrostatic discharge (ESD) protection. This includes making sure personnel and equipment are connected to a common ground, such as by wearing a wrist strap connected to the chassis ground, and placing components on static-free work surfaces.

- 1. Power down the system.
- 2. Label all cables connected to the motherboard tray for easy identification when reconnecting.
- 3. Unplug the cables.
- 4. Remove the motherboard tray.

Refer to the instructions in the section <u>Accessing the Motherboard Tray</u>.

5. Remove the M.2 riser card, the IO card, and the air baffle.

These components need to be removed to gain access to the battery.

a). Confirm the locations of the components to remove using the following diagram.



b). Lift off the IO card, the M.2.riser, and the air baffle and set them on a solid surface.



- 6. Replace the battery.
  - a). Locate the battery, using the following image as a guide.



b). Use a small flat-head screwdriver or similar thin tool to gently lift the battery from the battery holder.



c). Replace the battery with a new CR2032, installing it in the battery holder.



7. Re-insert the IO card, the M.2 riser card, and the air baffle into their respective slots.



8. Replace the motherboard tray.

Refer to the instructions in the section <u>Replacing the Motherboard Tray</u>.

- 9. Connect all the cables and power cords to the motherboard tray.
- 10. Apply power to the system and then log in.
- 11. Restore the date on the system.
  - a). Set the date.
    - \$ sudo date [MMDDhhmm[[CC]YY][.ss]]
  - b). Sync the date and time to the hardware real time clock.
    \$ sudo hwclock -w
  - c). Reset the BMC
    - \$ sudo ipmitool mc reset cold
- 12. Confirm that the time and date on the system are updated.
  - \$ sudo nvsm show health

# Chapter 19. Trusted Platform Module Replacement

#### 19.1. Trusted Platform Module Replacement Overview

This is a high-level overview of the procedure to replace the trusted platform module (TPM) on the DGX A100 system.

- 1. If enabled, disable drive encryption.
- 2. Shut down the system.
- 3. Label all motherboard tray cables and unplug them.
- 4. Slide out the motherboard tray and open the motherboard tray I/O compartment.
- 5. Replace the TPM on the motherboard.
- 6. Close the I/O compartment lid on the motherboard tray.
- 7. Slide the motherboard tray into the system.
- 8. Plug in all cables using the labels as a reference.
- 9. Power on the system.
- 10. Verify that the network card is healthy using nvsm show health.
- 11. If the data drives need to be protected, then enable encryption.

# 19.2. Replacing the Trusted Platform Module



**CAUTION: Static Sensitive Devices**: - Be sure to observe best practices for electrostatic discharge (ESD) protection. This includes making sure personnel and equipment are connected to a common ground, such as by wearing a wrist strap connected to the chassis ground, and placing components on static-free work surfaces.

- 1. Obtain a new Trusted Platform Module (TPM) from NVIDIA.
- 2. If data drives are encrypted, then disable encryption.

\$ sudo nv-disk-encrypt disable

**Note:** The TPM2 OS package must be installed and TPM enabled in the SBIOS. Refer to the chapter *Managing the DGX A100 Self-Encrypting Drives* in the <u>NVIDIA DGX A100 User Guide</u> for more information.

3. Power down the system.

E

- 4. Label all cables connected to the motherboard tray for easy identification when reconnecting.
- 5. Unplug all power cords.
- 6. Unplug all network, monitor, and USB cables.
- Remove the motherboard tray.
  Refer to the instructions in the section Removing the Motherboard Tray.
- 8. Replace the TPM by pulling it up vertically and then inserting the replacement in its slot.
- Close the motherboard tray lid and reinsert the mother board tray and secure. Refer to the instructions in the section <u>Reinstalling the Motherboard Tray</u>.
- 10. Connect all the cables to the motherboard tray.
- 11. Install all the power cords.

E

- 12. Power on the system and log in.
- 13. If data drives were encrypted prior to replacing the TPM, encrypt the drives.

**Note:** The TPM2 OS package must be installed and TPM enabled in the SBIOS. Refer to the chapter *Managing the DGX A100 Self-Encrypting Drives* in the <u>NVIDIA DGX A100 User Guide</u> for more information.

The following is an example command for enabling drive encryption:

\$ sudo nv-disk-encrypt init -g -r -k <your vault password>

- 14. Confirm that the system is healthy.
  - \$ sudo nvsm show health

# Chapter 20. Removing and Attaching the Bezel



1. Grab the bezel on both sides by the side handles, then pull directly away from the system to disengage from the magnetic latch.



2. To replace the bezel, align the bezel alignment pins with the chassis, then let the magnetic latch complete the attachment of the bezel.



#### Chapter 21. Installing the Rack Mount Kit

#### 21.1. Installing the Rails

Follow these instructions to install the DGX A100 server rack mount kit. The rack mount kit acts as a shelf in the rack, it does not allow the system to be moved once installed. All components are serviceable from the front or rear, so this movement is not necessary.



ltem	Description
1	Bottom lip to support the DGX A100.
2	Prongs that enter the holes of the rack. Holes can be square or round.
3	Metal clips that hold the rail in place.
4	Thread for securing the rack mount kit to the rack. Two types of screws are provided.
	Flat head: Must be used along with the provided washers at the front of the rack so that the DGX A100 can be installed flush to the rack.
	Pan or round head: Must be used at the rear of the rack.



- Use a Phillips screwdriver to assist in mounting the rails to the rack.
- Follow any designations on the slide rail to determine front/back and left-side/right-side positioning against the rack.
- 1. Align the bottom lip of the left or right rail to the bottom of the first rack unit for the server.
- 2. Attach the front of the rail to the rack.

Push the metal tab on the rail and then insert the two spring-loaded prongs into the holes on the front rack post.



On square-holed racks, make sure the prongs are completely inserted into the hole by confirming that the spring is fully extended.

The metal clip helps secure the rail in place while installing the rail to the rear of the rack.

3. Extend the rail to the rear post of the rack and secure the rail to the post.

Make sure that rail is level and the attachment on the rear post is at the same rack unit as the front.

a). Insert the spring-loaded prongs into the holes on the rear rack post.



- b). On square-holed racks, make sure the prongs are completely inserted into the hole by confirming that the spring is fully extended.
- 4. Repeat these steps for the other rail.
- 5. Secure the rails to the rack using the provided screws.

Install the four screws in the bottom holes of the rack kit.



Do not use the top holes, as those will be used for the captive screw in the chassis.

- 6. Confirm that the rails have been installed properly.
  - Both sides are installed on the same rack unit and are horizontal to each other.
  - The bottom lip is at the same height on all four posts.
  - ▶ The metal clips are properly attached.
  - ▶ Four screws are installed flat head on the front and pan head on the back.

# 21.2. Installing the Cage Nuts

The DGX A100 server is secured to the rack using four captive screws - one at each corner of the front of the unit.

- If your rack has round holes with 10-32 threads, then the screws will attached directly to the rack mounting holes.
- If your rack has square holes, then you need to install the included cage nuts to provide attachment points for mounting the DGX A100 server to the front of the rack.

The number of cage nuts to install depends on the type of rack that is used. On most racks, cage nuts are needed only for the top screws. The bottom screws will attach to the threads in the rails. The image on the left side of the following diagram shows this type. Green arrows indicate the general area for the cage nuts.

The image on the right side shows a rack with square posts such that the rail does not meet the front-most part of the rack, In this rack, cage nuts need to be installed at the bottom holes as well, as indicated by the green arrows. All four cage nuts are also needed if you have installed a rack shelf instead of rack rails.



Use the provided template to determine the exact location for installing the cage nuts.



1. Place the template so that the bottom of the template rests on the rail lip (or at the same level as the rail lip).

The template is double sided so it can be used as a reference on the left and right posts of the rack.

- 2. Install the cage nuts at the positions indicated on the template.
  - ▶ Rail kits attached to Type A racks require two (2) cage nuts installed; top positions only.
  - Rail kits attached to Type B racks require four (4) cage nut installed; both top and bottom positions.

#### Copyright

© 2022 NVIDIA Corporation & Affiliates. All rights reserved.