



DGX-2 System

Service Manual

Table of Contents

Chapter 1. Introduction.....	1
1.1. NVIDIA Enterprise Support Portal.....	1
1.2. NVIDIA Enterprise Support Email.....	2
1.3. NVIDIA Enterprise Support - Local Time Zone Phone Numbers.....	2
Chapter 2. Front Fan Module Replacement.....	3
2.1. Front Fan Module Replacement Overview.....	3
2.2. Identifying the Failed Fan Module.....	3
2.3. Replacing and Returning the Front Fan Module.....	4
Chapter 3. U.2 NVMe Cache Drive Replacement.....	7
3.1. U.2 NVMe Cache Drive Replacement Overview.....	7
3.2. Identifying the Failed U.2 NVMe.....	7
3.3. Replacing the U.2 NVMe Drive.....	9
Chapter 4. U.2 NVMe Cache Drive Upgrade from 8 to 16.....	11
4.1. U.2 NVMe Cache Drive Upgrade Overview.....	11
4.2. Identifying the NVMe Drive Manufacturer.....	11
4.3. Installing the Optional NVMe Drives.....	12
Chapter 5. U.2 NVMe Cache Drive Post-Installation Tasks.....	13
5.1. Recreating the Cache RAID 0 Volume.....	13
5.2. Confirming the Volume is Ready.....	13
5.3. Enabling the Temperature Sensor.....	14
5.4. Returning the NVMe Drive/Riser Assembly.....	14
Chapter 6. M.2 NVMe Boot Drive Replacement.....	15
6.1. M.2 NVMe Boot Drive Replacement Overview.....	15
6.2. Identifying the Failed M.2 NVMe.....	15
6.3. Replacing the M.2 NVMe Drive.....	16
6.4. Rebuilding the Boot Drive RAID 1 Volume.....	19
6.5. Returning the NVMe Drive/Riser Assembly.....	20
Chapter 7. M.2 Boot Drive Riser Assembly Replacement.....	21
7.1. M.2 Boot Drive Riser Assembly Replacement Overview.....	21
7.2. Determining a Failed M.2 NVMe Riser Assembly.....	21
7.3. Replacing the M.2 NVMe Riser Assembly.....	22
7.4. Returning the NVMe Drive/Riser Assembly.....	24
Chapter 8. Power Supply Replacement.....	25
8.1. Power Supply Replacement Overview.....	25

8.2. Identifying the Failed Power Supply.....	25
8.3. Replacing the Power Supply.....	26
Chapter 9. Power Supply Carrier Replacement.....	28
9.1. Power Supply Carrier Replacement Overview.....	28
9.2. Identifying the Failed Power Supply Carrier.....	28
9.3. Replacing the Power Supply Carrier.....	29
9.4. Verifying the PSU Carrier is Working.....	33
Chapter 10. DIMM Replacement.....	34
10.1. DIMM Replacement Overview.....	34
10.2. Identifying the Failed DIMM.....	34
10.3. Replacing the DIMM.....	35
Chapter 11. ConnectX-5 Card Replacement.....	38
11.1. ConnectX-5 Card Replacement Overview.....	38
11.2. Replacing the ConnectX-5 Card.....	38
11.3. Verifying the ConnectX-5 Cards.....	44
Chapter 12. Dual-port ConnectX-5 PCI Card/PCI Riser Replacement.....	48
12.1. Dual-port ConnectX-5 Card Replacement Overview.....	48
12.2. Replacing the Dual-Port ConnectX-5 PCI Card.....	49
Chapter 13. Optional Dual-port ConnectX-5 PCI Card Installation.....	54
13.1. Optional Dual-port ConnectX-5 Card Installation Overview.....	54
13.2. Installing the Optional Dual-Port ConnectX-5 PCI Card.....	54
Chapter 14. Removing and Attaching the Bezel.....	59
14.1. Removing the Bezel.....	59
14.2. Attaching the Bezel.....	60
Chapter 15. Removing and Attaching the EMI Shield.....	62
15.1. Removing the EMI Shield.....	62
15.2. Attaching the EMI Shield.....	63
Chapter 16. Front Console Board Replacement.....	66
16.1. Front Console Board Replacement Overview.....	66
16.2. Replacing the Front Console Board.....	66
Chapter 17. Motherboard Tray Battery Replacement.....	70
17.1. Motherboard Tray Battery Replacement Overview.....	70
17.2. Replacing the Motherboard Tray Battery.....	70
Chapter 18. I/O Tray Removal and Installation.....	74
18.1. I/O Tray Replacement Overview.....	74
18.2. Replacing the I/O Tray.....	74

Chapter 19. Motherboard Tray Removal and Installation.....	81
19.1. Removing the Motherboard Tray.....	81
19.2. Installing the Motherboard Tray.....	83
Chapter 20. Identifying the Component Manufacturer.....	86

List of Figures

Figure 1. NVMe Drives: PCIe to Slot Mapping	9
---------------------------------------------------	---

Chapter 1. Introduction

This document contains instructions for replacing NVIDIA® DGX-2™ System components. Be sure to familiarize yourself with the NVIDIA Terms & Conditions documents before attempting to perform any modification or repair to the DGX-2 System. These Terms & Conditions for the DGX-2 System can be found through the [NVIDIA DGX Systems Support](#) 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.

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

- ▶ Front Fan Modules
- ▶ Cache (U.2) NVMe Drives
- ▶ Power Supplies
- ▶ Power Supply Carrier
- ▶ Power Supply Carrier Fan
- ▶ Boot (M.2) NVMe Drives
- ▶ Boot Drives Riser Assembly
- ▶ DIMMs
- ▶ Motherboard Tray Battery
- ▶ PCIe Riser Assembly
- ▶ ConnectX-5 Network Adapter Card
- ▶ I/O Expander Tray
- ▶ Front Console Board
- ▶ EMI Shield
- ▶ Bezel

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

1.1. NVIDIA Enterprise Support Portal

The best way to file an incident is to log on to the [NVIDIA Enterprise Support](#) portal.

1.2. NVIDIA Enterprise Support Email

You can also send an email to enterprisesupport@nvidia.com.

1.3. NVIDIA Enterprise Support - Local Time Zone Phone Numbers

Visit [NVIDIA Enterprise Customer Support \(https://www.nvidia.com/en-us/support/enterprise/\)](https://www.nvidia.com/en-us/support/enterprise/)

Chapter 2. Front Fan Module Replacement

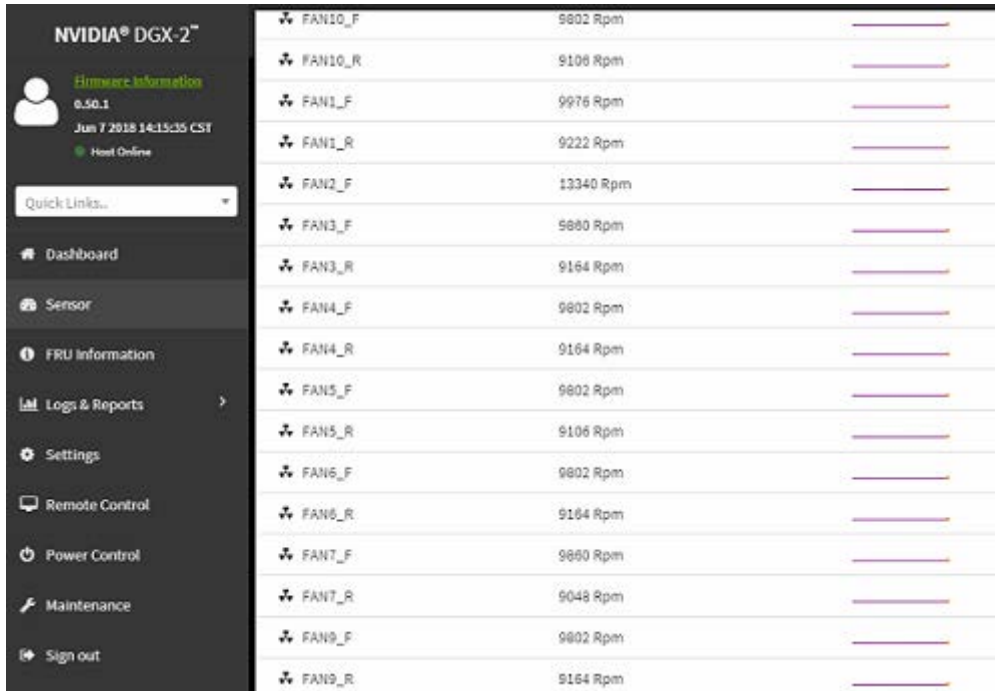
2.1. Front Fan Module Replacement Overview

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 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.
`nvsm show health`
6. Return the bad fan module using the packaging from the new fan module.

2.2. Identifying the Failed Fan Module

1. Log on to the BMC.
2. Click **Sensor** from the left navigation menu, then review the **Normal Sensors** section.



There are two fans in the fan module, identified by `FAN n _F` or `FAN n _R`, where n is the module ID. If either fan fails, then the entire module must be replaced.

3. Use NVSM to confirm the fan issue.

```
$ sudo nvsm show health
```


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. Locate the failed fan module on the physical system using the following diagram.



3. Remove the failed front fan module by pressing on the tab and pulling on the handle.



4. Quickly insert the new fan module, observing that the handle release mechanism is facing up and the rear connector is facing down.



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

5. Confirm that the fan module is working properly by verifying on the BMC and by using NVSM (`nvsmd show health`) to confirm the replaced fan is healthy.
6. Use packaging to pack up the bad fan and follow the shipping instructions to return the bad fan to NVIDIA Enterprise Support

Chapter 3. U.2 NVMe Cache Drive Replacement

3.1. U.2 NVMe Cache Drive Replacement Overview

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



CAUTION: Hot-swapping of the NVMe drives is not supported. Be sure to turn the system off before replacing a failed drive.

1. Identify the failed Non-Volatile Memory Express (NVMe) drive.
2. Get replacement from NVIDIA Enterprise Support.
3. Power down the system and then remove the failed NVMe drive.
4. Insert the new NVMe drive.
5. Power on the DGX-2 System.
6. Rebuild the RAID volume and remount the `/raid` partition.

3.2. 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 blinking red LED as illustrated in the following example.



Identifying the Failed NVMe from the Console

To identify the failed NVMe drive from the DGX-2 console, enter the following and then look for a missing entry from the output.

```
$ sudo mdadm -D /dev/md1
```

Number	Major	Minor	RaidDevice	State	
0	259	8	0	active sync	/dev/nvme9n1
1	259	13	1	active sync	/dev/nvme5n1
2	259	7	2	active sync	/dev/nvme6n1
3	259	10	3	active sync	/dev/nvme3n1
4	259	12	4	active sync	/dev/nvme2n1
5	259	11	5	active sync	/dev/nvme7n1
6	259	9	6	active sync	/dev/nvme8n1
7	259	6	7	active sync	/dev/nvme4n1


The list should include device names from `nvme2n1` through `nvme9n1` for systems with 8 NVMe drives, and from `nvme0n1` through `nvme15n1` for systems with 16 NVMe drives.

To map the device name to the physical slot ID, enter the following, where X corresponds to the missing device name.

```
$ ls -l /dev/disk/by-path |grep nvmeX |cut -d'|' -f3
```

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



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 = 100007H0
  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.

3.3. Replacing the U.2 NVMe Drive

1. Be sure you have obtained the replacement drive.
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 by squeezing the levers on the handle and pulling the drive out.

5. Replace the new NVMe drive in the same slot by fully inserting it and making sure it clicks into place.
6. Power on the system.

Perform the tasks describes in the chapter [U.2 NVMe Cache Drive Post-Installation Tasks](#).

Chapter 4. U.2 NVMe Cache Drive Upgrade from 8 to 16

If you require more than 30 terabytes (useable 28 terabytes) for cache, you can increase the cache to 60 terabytes (useable 56 terabytes) by adding eight more NVMe drives to the DGX-2 System.

4.1. U.2 NVMe Cache Drive Upgrade Overview

This is a high-level overview of the steps needed to upgrade the DGX-2 System's cache size.

1. Identify the manufacturer and model of the of currently installed NVMe drives.
2. Place an order for additional eight NVME drives.
3. Power off the system.
4. Install the NVMe drives in the DGX-2 System.
5. Power on the system.
6. Re-initialize the `/raid` filesystem to recognize all 16 drives.

4.2. Identifying the NVMe Drive Manufacturer

1. Identify the drives in the RAID volume.

```
$ sudo nvsm show /systems/localhost/storage/volumes/md1
Properties
...
  Drives = [ nvme2n1, nvme3n1, nvme4n1, nvme5n1, nvme6n1, nvme7n1, nvme8n1,
             nvme9n1 ]
...
```

2. Select one of the drives listed in the `Drives=` entry from the previous command, and then enter the following command, where X corresponds to the drive that you selected.

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

Example showing an excerpt of the output for `nvme5n1`.

```


/systems/localhost/storage/drives/nvme5n1
Properties:
  Capacity = 3840755982336
  BlockSizeBytes = 7501476528
  SerialNumber = 174719FCF9F1
  PartNumber = N/A
  Model = Micron_9200_MTFDHAL3T8TCT
  Revision = 100007H0
  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

```

3. Determine the manufacturer (Samsung or Micron) and model from the `Model=` entry in the output, and then order the additional drives from NVIDIA Enterprise Support, specifying the manufacturer and model.

4.3. 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.
4. Remove the blank filler disks from slots 2, 3, 6, 7, 10, 11, 14, and 15.



Squeeze the levers on the handle and pull the blank filler disks out.

5. Install the additional eight NVMe drives in slots 2, 3, 6, 7, 10, 11, 14, and 15.
6. Power on the system.

Perform the tasks describes in the chapter [U.2 NVMe Cache Drive Post-Installation Tasks](#).

Chapter 5. 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 8 to 16 drives.

5.1. Recreating the Cache RAID 0 Volume

1. Stop cachefilesd.

```
$ sudo systemctl stop cachefilesd
```

2. Umount /raid and stop raid-0.

```
$ sudo umount -f /raid
```

```
$ sudo mdadm --stop /dev/md1
```

3. Run the script to rebuild the RAID volume.

```
$ sudo /usr/bin/configure_raid_array.py -c -f
```

Press **Y** at any questions.

4. When completed, confirm that the /raid volume is mounted.

```
$ df -hl /raid
```

The /dev/md1 filesystem should be mounted on /raid with size 28 TB or 56 TB, depending on whether 8 or 16 drives are installed.

5.2. Confirming the Volume is Ready

1. Confirm the storage devices and volumes in the system are healthy using the following command.

```
$ sudo nvsm show systems/localhost/storage/volumes/md1
```

2. Verify status_Health=OK and that the numbers of drives listed in Drives = is as expected.

3. Confirm that the drives are now available.

```
$ sudo mdadm -D /dev/md1
```

If the drive manufacturer is Micron, perform the steps in [Enabling the Temperature Sensor](#).

5.3. Enabling the Temperature Sensor

The steps in this section need to be followed only for Micron NVMe drives.

1. Verify the need to enable temperature reading for the installed NVMe drives by running `ipmitool`.

```
$ sudo ipmitool sdr | grep -i temp | grep -i -e nvme*temp -e temp_u2
```

2. If any of the NVMe drives do not show a temperature reading, then enable the SMBUS on all the drives.

a). Establish root role before running the script.

```
$ sudo su
```

b). Run the following script.

```
:user# for drives in `nvme list|grep Micron | cut -d' ' -f1 |sed 's/..$//'`;
do /opt/MicronTechnology/MicronMSECLI/msecli -M -k 1 -n $drives ;
done
```

c). Exit out of root role.

```
:user# exit
```

3. Confirm that temperature reading for the replaced drive is enabled by running `ipmitool`.

```
$ sudo ipmitool sdr | grep -i temp | grep -i -e nvme*temp -e temp_u2
```

5.4. Returning the NVMe Drive/Riser Assembly

Use the packaging from the new drive/riser assembly and follow the instructions that came with the package to ship the old drive/riser assembly 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 6. M.2 NVMe Boot Drive Replacement

6.1. 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. Remove the motherboard tray and place on a solid flat surface.
6. Remove the motherboard tray lid.
7. Pull out the M.2 riser card with both M.2 disks attached.
8. Replace the failed M.2 device on the riser card.
9. Install the M.2 riser card with both M.2 disks.
10. Close the lid on the motherboard tray.
11. Insert the motherboard tray into the system.
12. Plug in all cables using the labels as a reference.
13. Power on the system.
14. Confirm the RAID 1 array is being rebuilt.

6.2. Identifying the Failed M.2 NVMe

The DGX-2 System automatically sets the failed M.2 drive offline when it detects the failure.

1. From the console, run the following command to identify the failed drive.

```
$ 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 nvme0 is the failed drive.

Number	Major	Minor	RaidDevice	State	
0	259	2	0	active sync	/dev/nvme1n1p2
-	0	0	1	removed	

2. 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.

3. Run the following command to determine the location of the failed boot drive, replacing X with the number corresponding to the device name of the failed drive.

```
$ ls -l /dev/disk/by-path |grep nvmeX |cut -d':' -f3
```

The output will be either '01' or '05'. Be sure to note this number as you will need it when performing the replacement.

4. Identify the manufacturer and model for the M.2 drive by running the following command on the healthy drive, where X corresponds to the healthy drive, and inspecting the Manufacturer = and Model = line.

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

5. Provide the vendor name for the drive when ordering the replacement and then obtain the replacement from NVIDIA Enterprise Support.

6.3. 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.
2. Power down the system.
3. Label all cables connected to the motherboard tray for easy identification when reconnecting.
4. Remove the motherboard tray.

Refer to the instructions in the section [Removing the Motherboard Tray](#).

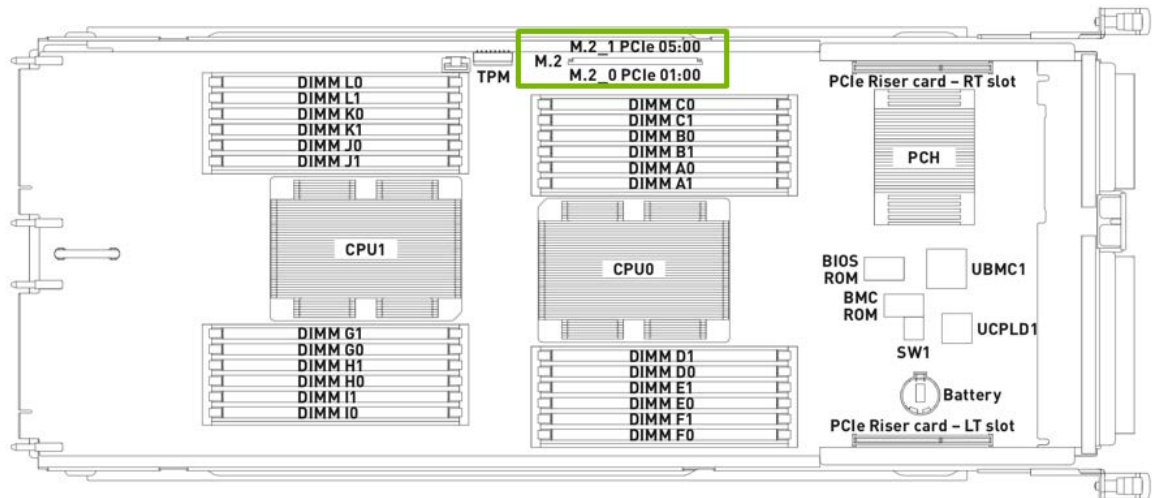
5. Remove the M.2 modules and the riser card from the motherboard tray by pushing on the clip to release the riser.




6. Identify the failed M.2 module and remove it from the riser card by loosening the screw with a Philips #2 screwdriver.



Use the label on the motherboard tray lid to help identify the M.2_0 module and the M.2_1 module.



7. Insert the new M.2 module and secure it with the screw to the riser card.
8. Install the assembled module on the motherboard by inserting the the riser card in its slot.



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

Refer to the instructions in the section [Installing the Motherboard Tray](#).

10. Connect all the cables to the motherboard tray.

Rebuild the RAID 1 array according to the instruction in the section [Rebuilding the Boot Drive RAID 1 Volume](#).

6.4. Rebuilding the Boot Drive RAID 1 Volume

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

1. Turn the DGX-2 System on.

The rebuilding process should begin automatically upon system boot.

2. Log in and then confirm that the RAID 1 array is being rebuilt.

```
$ sudo mdadm -D /dev/md0
```

- ▶ If the RAID 1 array is still in the process of being rebuilt, the output will include the following line.

```
Rebuilt Status : XX% complete
```

- ▶ If the RAID 1 array rebuilding process is completed, the output will show both drives in 'active sync' state and you can skip the remaining steps.

3. If the rebuilding process did not start automatically, then rebuild the array manually.

In the following steps, replace X with the number that corresponds to the replaced drive, and Y with the number that corresponds to the drive that was not replaced (the surviving drive). If you did not note this information when identifying the failed drive, then follow the instructions in the first step of [Identifying the Failed M.2 Drive](#).

- a). Start an NVSM CLI interactive session and switch to the storage target.

```
$ sudo nvsm
nvsm-> cd /systems/localhost/storage
```

- b). Start the rebuilding process and be ready to enter the device name of the replaced drive.

```
nvsm(/systems/localhost/storage)-> start volumes/md0/rebuild
PROMPT: In order to rebuild this volume, a spare drive
        is required. Please specify the spare drive to
        use to rebuild md0.
Name of spare drive for md0 rebuild (CTRL-C to cancel): nvmeXn1
WARNING: Once the volume rebuild process is started, the
        process cannot be stopped.
Start RAID-1 rebuild on md0? [y/n] y
```

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% [\ ]
```

After about 30 seconds, the "Rebuilding RAID-1 ..." message should appear.

```
/systems/localhost/storage/volumes/md0/rebuild started at 2018-10-12
15:27:26.525187
Rebuilding RAID-1 rebuild on volume md0...
31.0% [=====/] ]
```

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.

6.5. Returning the NVMe Drive/Riser Assembly

Use the packaging from the new drive/riser assembly and follow the instructions that came with the package to ship the old drive/riser assembly 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 7. 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.

7.1. 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 from NVIDIA Enterprise Support.
3. Power 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. Pull out the M.2 riser assembly with both M.2 disks attached.
8. Install the new M.2 riser assembly with both M.2 disks attached..
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. Re-install the OS.

7.2. 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-2 cannot be booted.
- ▶ The boot drives cannot be seen from the SBIOS.
- ▶ Loss of communication with the M.2 boot drives.
- ▶ The M.2 riser assembly was damaged.


7.3. 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.

1. 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 [Removing the Motherboard Tray](#).
4. Remove the M.2 modules and the riser card from the motherboard tray by pushing on the clip to release the riser.



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



6. Install the motherboard tray lid and then install the motherboard tray.
Refer to the instructions in the section [Installing the Motherboard Tray](#).
7. Connect all the cables to the motherboard tray.

8. Re-install the DGX OS server software.

See the [DGX-2 User Guide](#) for detailed instructions.

7.4. Returning the NVMe Drive/Riser Assembly

Use the packaging from the new drive/riser assembly and follow the instructions that came with the package to ship the old drive/riser assembly 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 8. Power Supply Replacement

This chapter describes how to replace one of the DGX-2 System power supplies (PSUs).

8.1. 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. Remove the failed power supply.
6. Insert new power supply.
7. Insert the power cord and make sure both LEDs light up green (IN/OUT).
8. Use the BMC to confirm that the power supply is working correctly.

8.2. 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 inspecting the LEDs on the power supply.



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 a couple of ways to identify the failed PSU from the DGX-2 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`.

- You can also log into the BMC, then click **Sensor** from the left side menu and inspect the PSU information from the **Normal Sensors** section.

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.



Identifying the Power Supply Manufacturer


Enter the following NVSM CLI command to see the manufacturer of the PSUs in the system..

```
$ sudo nvsm show psus |grep Manufacturer
```


Request a replacement PSU from NVIDIA Enterprise Support, specifying this information.

8.3. 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 five remaining PSUs are working and energized, then you do not need to shut down power to the DGX-2 System..
 - If fewer than five PSUs are working and energized, then you do need to shut down power to the DGX-2 System.
3. Unplug the power cable 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 blue tab to release the lock.



b). Pull on the 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 blue locking mechanism engages.
 - b). Plug in the power cord and attach the retaining clip.
 - 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 the health of the system.

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

Chapter 9. Power Supply Carrier Replacement

This chapter describes how to replace a failed DGX-2 System power supply carrier.

The power supply carrier can fail due to a power distribution board failure, or a bad carrier fan.

9.1. Power Supply Carrier Replacement Overview

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


1. Identify failed power supply carrier using the BMC and submit a service ticket.
2. Get replacement power supply carrier from NVIDIA Enterprise Support.
3. Identify the power supply carrier using the diagram as a reference.
4. Power off the system.
5. Remove the power cords from the three power supplies in the carrier to be replaced.
6. Remove the failed power supply carrier and place on a solid, stable work surface.
7. Move the power supplies to the new carrier.
8. Insert new power supply carrier and secure in place with the thumbscrew.
9. Insert the power cords and make sure the LEDs light up green (IN/OUT).
10. Use the BMC to confirm that the power supply carrier, power supplies, and fans are working correctly.
11. Power on the system.

9.2. Identifying the Failed Power Supply Carrier

1. Log on to the BMC.
2. Click **Sensor** from the left navigation menu and review the PDB entries.

NVIDIA® DGX-2™		
<div> <div> <div> <div>Enterprise Information</div> <div>6.50.1</div> <div>Jun 7 2018 14:13:35 CEST</div> <div>Host Online</div> </div> </div> <div>Quick Links</div> <div> Dashboard Sensor Firmware Information Logs & Reports Settings Remote Control Power Control Maintenance Sign out </div> </div>		
Sensor Name	Reading	
FAN0_F	3132 Rpm	
FAN0_PDB	10880 Rpm	
FAN0_R	3016 Rpm	
FAN1_F	3654 Rpm	
FAN1_PDB	10472 Rpm	
FAN1_R	3422 Rpm	
FAN2_F	3654 Rpm	
FAN2_PDB	11016 Rpm	
FAN2_R	3422 Rpm	
FAN3_F	3654 Rpm	
FAN3_PDB	10676 Rpm	

The following diagram shows the PSU carrier location corresponding to each PSU carrier fan (FAN_x_PDB).



If necessary, work with NVES to identify the failed power supply carrier – this could be due to a power distribution board failure.

3. Request a new power supply carrier from NVES.
4. When the replacement arrives, unpack the item and save the packaging.

9.3. Replacing the Power Supply Carrier

1. Identify a solid work surface where the components can be rested for the procedure.
2. Power off system before replacing power supply carrier.

3. Unplug the three power cords connected to the power supply carrier cage.
4. Remove the power supply carrier.
 - a). Release the power supply carrier by turning the indicated thumbscrew.

The following diagram show the the right power supply carrier as an example.



- b). Use the chrome handle to pull out the power supply carrier.



Important: The module will be heavy as it holds three power supplies



5. Move the power supply units to the new carrier.
- a). Pull the power supplies out of the old carrier.



- b). Insert the power supplies into the new carrier.



6. Replace the power supply carrier.
 - a). Insert the power supply carrier into the chassis.



- b). Tighten the thumbscrew to secure the power supply carrier.



9.4. Verifying the PSU Carrier is Working

This section describes the steps needed to verify that the PSU carrier has been replaced correctly.

1. Plug in the three power cords that were previously unplugged.
2. Confirm power supply LEDs light up (IN and OUT).
3. Log on to the BMC.
4. Go to sensor information and confirm the new power supply carrier is operational.

Power distribution board, PDB fans and power supplies should be active and working.

Sensor Name	Reading
FAN0_F	3132 Rpm
FAN0_PDB	10880 Rpm
FAN0_R	3016 Rpm
FAN1_F	3654 Rpm
FAN1_PDB	10472 Rpm
FAN1_R	3422 Rpm
FAN2_F	3654 Rpm
FAN2_PDB	11016 Rpm
FAN2_R	3422 Rpm
FAN3_F	3654 Rpm
FAN3_PDB	10676 Rpm

5. Power on the system.

Ship back the power supply carrier in the packaging that the new one arrived in.

Chapter 10. DIMM Replacement

10.1. DIMM Replacement Overview

This is a high-level overview of the procedure to replace a dual inline memory module (DIMMs) on the DGX-2 System.

1. Use the `nvsm show` 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`.

10.2. 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 dmidecode -t memory | grep Manufacturer | tail -1
```

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

10.3. Replacing the DIMM

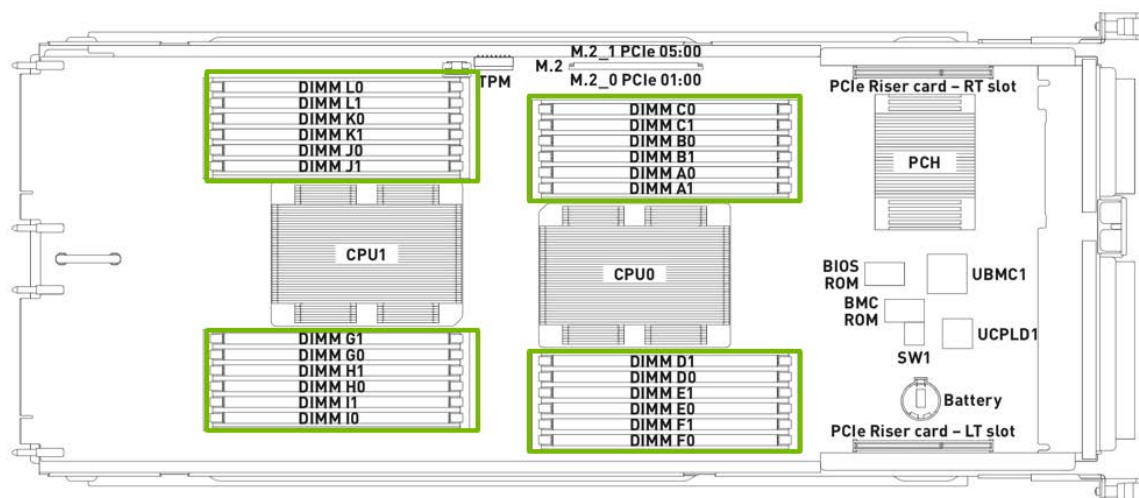
Before attempting to replace any of the dual inline memory modules (DIMMs), be sure to have performed the following:

- ▶ Determined the location ID of the faulty DIMM needing replacement as explained in [Identifying the Failed DIMM](#). 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 [Removing the Motherboard Tray](#).
4. Using the following diagram as a guide, locate the faulty DIMM to be replaced.



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 motherboard tray lid and then install the motherboard tray.
Refer to the instructions in the section [Installing the Motherboard Tray](#).
8. Connect all the cables to the motherboard tray.
9. Power on the system and log in.
10. Confirm that the system is healthy.

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

There should be no new alerts listed.

Chapter 11. ConnectX-5 Card Replacement

11.1. ConnectX-5 Card Replacement Overview

This is a high-level overview of the procedure to replace one or more Mellanox ConnectX-5 cards on the DGX-2 System.

1. Use the `nvsm show` commands to identify the failed ConnectX-5 card.
2. Get a replacement ConnectX-5 card from NVIDIA Enterprise Support.
3. Shut down the system.
4. Label all I/O tray cables and unplug them.
5. Remove the I/O tray and open the lid.
6. Locate the failed ConnectX-5 card, then remove the screw that attaches the card and remove the card.
7. Insert the new card into the slot and secure with the screw.
8. Close the lid on the I/O 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 ConnectX-5 card is healthy using `nvsm`.

11.2. Replacing the ConnectX-5 Card

Before attempting to replace any of the ConnectX-5 cards, be sure to have performed the following:

- Determined the location ID of the faulty ConnectX-5 card needing replacement.
Run `nvsm show health` to identify the bad card. Note the PCIe bus ID and slot number.
- Obtained the replacement ConnectX-5 card have saved the packaging for use when returning the faulty component.



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 the network ports (0-7) connected to the I/O tray for easy identification when reconnecting.




3. Remove the cables.
4. Remove the I/O tray.
 - a). Loosen the two green I/O tray screws with a Philips #2 screwdriver and pull the levers outward to release the tray.




- b). Pull the I/O tray out of the system and place it on a solid, flat work surface.




CAUTION: Exercise care when removing the tray as it is long and heavy, and do not handle the module from the rear connectors.




5. Remove the I/O tray lid.
 - a). Loosen the black screws and then push the lid towards you to release the lid.




b). Lift the lid upward.




6. Do the following for each card that needs to be replaced.
 - a). To assist in locating the card to remove, refer to the service label that maps the PCIe bus ID to the slot number.



b). Remove the screw that secures the card, then pull the card out of the slot.




- c). Insert the replacement card and secure with the screw removed from the previous step
- 7. Install the I/O tray.
 - a). Replace the I/O tray lid by placing it over the module using the guiding pins and grooves.




- b). Slide the lid back so that the black screws engage with the tray, then tighten the black screws to secure the lid.



c). Push the I/O tray back into the system.



d). Close the levers toward the center, making sure the connectors engage with the midplane, then tighten the thumbscrews by hand or with a Phillips #2 screwdriver.



8. Connect all cables back into the ConnectX-5 card ports.
9. Power on the system and log in.
10. Confirm that the system is healthy.

```
$ sudo nvsm show health
```

There should be no new alerts listed.

11.3. Verifying the ConnectX-5 Cards

This section describes the steps needed to verify that the ConnectX-5 cards has been replaced correctly.

1. With the DGX-2 turned on, verify that the card was installed correctly and is recognized by the system.

```
$ lspci | grep -i mellanox
```

The output should show all installed Mellanox cards, including the dual port (and optional dual port) cards.

Example

```
35:00.0 Infiniband controller: Mellanox Technologies MT27800 Family [ConnectX-5]
3a:00.0 Infiniband controller: Mellanox Technologies MT27800 Family [ConnectX-5]
58:00.0 Infiniband controller: Mellanox Technologies MT27800 Family [ConnectX-5]
5d:00.0 Infiniband controller: Mellanox Technologies MT27800 Family [ConnectX-5]
86:00.0 Ethernet controller: Mellanox Technologies MT27800 Family [ConnectX-5]
86:00.1 Ethernet controller: Mellanox Technologies MT27800 Family [ConnectX-5]
b8:00.0 Infiniband controller: Mellanox Technologies MT27800 Family [ConnectX-5]
bd:00.0 Infiniband controller: Mellanox Technologies MT27800 Family [ConnectX-5]
e1:00.0 Infiniband controller: Mellanox Technologies MT27800 Family [ConnectX-5]
```

```
e6:00.0 Infiniband controller: Mellanox Technologies MT27800 Family [ConnectX-5]
```

The dual port cards are identified by bus ID 86. Look for all other cards. If eight cards (those other than bus ID 86) are not reported, then the card was not installed properly and should be reseated.

If a card other than the officially supported Mellanox family of adapters appears, contact NVIDIA Enterprise Support.

2. Verify the firmware version.

```
$ cat /sys/class/infiniband/mlx5*/fw_ver
```

Example output:

```
12.23.1020
12.23.1020
12.23.1020
12.23.1020
12.23.1020
12.23.1020
12.23.1020
12.23.1020
```

The latest InfiniBand firmware version supported for each DGX OS Server release is as follows:

- Release 4.x: Firmware version 12.23.1020

3. If you need to update the firmware, follow these steps:

- a). Initiate the firmware update.

```
$ sudo /opt/mellanox/mlnx-fw-updater/mlnx_fw_updater.pl
```

The script will check the firmware version of each card and update where needed.

If the firmware is updated for any card, you will need to reboot the system for the changes to take effect.

- b). Reboot the system if instructed.

- c). After rebooting the system, verify that all the Mellanox InfiniBand cards are using the current firmware.

```
$ cat /sys/class/infiniband/mlx5*/fw_ver
12.23.1020
12.23.1020
12.23.1020
12.23.1020
12.23.1020
12.23.1020
12.23.1020
12.23.1020
12.23.1020
```

4. Verify the physical port state for the InfiniBand cards.

```
$ ibstat
```

In the output text, verify that the Physical State for each card with a cable connection is **LinkUp** and that the port for the card is configured with a GUID. The following example output shows one card in a non-connected state, and the remaining cards in a connected state. Relevant text is highlighted in bold.

```
CA 'mlx5_0'
CA type: MT4119
Number of ports: 1
Firmware version: 12.23.1020
Hardware version: 0
Node GUID: 0x248a0703000de288
System image GUID: 0x248a0703000de288
```

```

Port 1:
  State: Down
  Physical state: Polling
  Rate: 10
  Base lid: 65535
  LMC: 0
  SM lid: 0
  Capability mask: 0x2651e848
  Port GUID: 0x248a0703000de288
  Link layer: InfiniBand
CA 'mlx5_1'
  CA type: MT4119
  Number of ports: 1
  Firmware version: 12.23.1020
  Hardware version: 0
  Node GUID: 0x248a0703000de26c
  System image GUID: 0x248a0703000de26c
  Port 1:
    State: Initializing
    Physical state: LinkUp
    Rate: 100
    Base lid: 65535
    LMC: 0
    SM lid: 0
    Capability mask: 0x2651e848
    Port GUID: 0x248a0703000de26c
    Link layer: InfiniBand
CA 'mlx5_2'
  CA type: MT4119
  Number of ports: 1
  Firmware version: 12.23.1020
  Hardware version: 0
  Node GUID: 0x248a0703001effde
  System image GUID: 0x248a0703001effde
  Port 1:
    State: Initializing
    Physical state: LinkUp
    Rate: 100
    Base lid: 65535
    LMC: 0
    SM lid: 0
    Capability mask: 0x2651e848
    Port GUID: 0x248a0703001effde
    Link layer: InfiniBand
CA 'mlx5_3'
  CA type: MT4119
  Number of ports: 1
  Firmware version: 12.23.1020
  Hardware version: 0
  Node GUID: 0x7cfe900300118f22
  System image GUID: 0x7cfe900300118f22
  Port 1:
    State: Initializing
    Physical state: LinkUp
    Rate: 100
    Base lid: 65535
    LMC: 0
    SM lid: 0
    Capability mask: 0x2651e848
    Port GUID: 0x7cfe900300118f22
    Link layer: InfiniBand
CA 'mlx5_4'
  CA type: MT4119
  Number of ports: 1
  Firmware version: 12.23.1020
  Hardware version: 0
  Node GUID: 0x7cfe900300118f26

```

```

System image GUID: 0x7cfe900300118f26
Port 1:
    State: Initializing
    Physical state: LinkUp
    Rate: 100
    Base lid: 65535
    LMC: 0
    SM lid: 0
    Capability mask: 0x2651e848
    Port GUID: 0x7cfe900300118f26
    Link layer: InfiniBand
CA 'mlx5_5'
CA type: MT4119
Number of ports: 1
Firmware version: 12.23.1020
Hardware version: 0
Node GUID: 0x7cfe900300118f25
System image GUID: 0x7cfe900300118f25
Port 1:
    State: Initializing
    Physical state: LinkUp
    Rate: 100
    Base lid: 65535
    LMC: 0
    SM lid: 0
    Capability mask: 0x2651e848
    Port GUID: 0x7cfe900300118f25
    Link layer: InfiniBand
CA 'mlx5_6'
CA type: MT4119
Number of ports: 1
Firmware version: 12.23.1020
Hardware version: 0
Node GUID: 0x7cfe900300118f24
System image GUID: 0x7cfe900300118f24
Port 1:
    State: Initializing
    Physical state: LinkUp
    Rate: 100
    Base lid: 65535
    LMC: 0
    SM lid: 0
    Capability mask: 0x2651e848
    Port GUID: 0x7cfe900300118f24
    Link layer: InfiniBand
CA 'mlx5_7'
CA type: MT4119
Number of ports: 1
Firmware version: 12.23.1020
Hardware version: 0
Node GUID: 0x7cfe900300118f23
System image GUID: 0x7cfe900300118f23
Port 1:
    State: Initializing
    Physical state: LinkUp
    Rate: 100
    Base lid: 65535
    LMC: 0
    SM lid: 0
    Capability mask: 0x2651e848
    Port GUID: 0x7cfe900300118f23
    Link layer: InfiniBand

```

See the *Switching Between InfiniBand and Ethernet* section of the [NVIDIA DGX-2 User Guide](#) for instructions on switching the port to InfiniBand or Ethernet, if required.

Chapter 12. Dual-port ConnectX-5 PCI Card/PCI Riser Replacement

The system comes with a dual-port Mellanox ConnectX-5 card that is configured to work in Ethernet mode. The card is installed in a PCI riser assembly. The following steps outline how to replace either the card alone, or the entire PCI riser assembly.

12.1. Dual-port ConnectX-5 Card Replacement Overview

This is a high-level overview of the procedure to replace the dual-port Mellanox ConnectX-5 PCI card or PCI riser assembly on the DGX-2 System.

1. Use the `nvsm show health` commands to verify an issue with the dual-port ConnectX-5 card.
2. Obtain the replacement parts - either the dual-port ConnectX-5 card or the PCI riser assembly - 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 work surface.
6. Remove the right-side PCI card riser.
7. Replace the PCI card if you are only replacing the card itself.
8. Replace the right-side PCI card riser.
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 the ConnectX-5 card is working.

12.2. Replacing the Dual-Port ConnectX-5 PCI Card



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. Identify the failed card by running nvsm.


```
$ sudo nvsm show health
```
2. If the failed component is the Mellanox dual-port card located at PCIe bus 86:00, obtain a replacement part from NVIDIA Enterprise Services.
3. If replacing the card alone, unpack it upon receipt and confirm that it comes with a low-profile bracket.

Replacement Instructions


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 [Removing the Motherboard Tray](#).

5. Remove the right PCI card riser.
 - a). Release the right PCI card riser by turning the right black screw.




b). Remove the right PCI riser card from the motherboard tray.




6. Replace the dual-port PCI card (if applicable).


a). Loosen and remove the screw that secures the PCI card to the riser.



b). Pull the old card out of the riser and install the new card into the riser.




c). Replace and tighten the screw that secures the PCI card to the riser.



7. Install the right PCI riser.

a). Replace the right PCI riser card on the motherboard tray.



b). Tighten the black screw on the right PCI card riser.



8. Replace the motherboard tray.

Refer to the instructions in the section [Installing the Motherboard Tray](#).

9. Connect all the cables to the motherboard tray.

10. Apply power to the system.

11. Confirm that the PCI card is visible from the system.

```
$ sudo lspci |grep 86\:00
86:00.0 Ethernet controller: Mellanox Technologies MT27800 Family [ConnectX-5]
86:00.1 Ethernet controller: Mellanox Technologies MT27800 Family [ConnectX-5]
```

12. Confirm that the system is healthy.

```
$ sudo nvsm show health
```

13. Verify basic connectivity to the network.

Verify mount points are available (if mounted over the ConnectX-5 card).

Consult the DGX-2 User Guide for instructions on reconfiguring network interfaces, if necessary.

Chapter 13. Optional Dual-port ConnectX-5 PCI Card Installation

You can install an optional dual-port Mellanox ConnectX-5 card in the blank PCI riser assembly. If you purchased the optional PCI card before installing the system, then be sure to install the card when installing the system. The following steps outline how to install the card, but you can also use the instructions for guidance in replacing the card or replacing the entire PCI riser assembly.

13.1. Optional Dual-port ConnectX-5 Card Installation Overview

This is a high-level overview of the procedure to install the optional dual-port Mellanox ConnectX-5 PCI card on the DGX-2 System.

1. Remove the motherboard tray and place on a solid, flat work surface.
2. Remove the left-side PCI card riser.
3. Install the optional dual-port ConnectX-5 network card.
4. Replace the left-side PCI card riser.
5. Insert the motherboard tray into the system.

13.2. Installing the Optional Dual-Port ConnectX-5 PCI Card



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.

Installation Instructions


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 [Removing the Motherboard Tray](#).

5. Remove the left PCI card riser.
 - a). Release the left PCI card riser by turning the left black screw.




- b). Remove the left PCI riser card from the motherboard tray.




6. Install the dual-port PCI card.
 - a). Loosen and remove the screw that secures the metal bracket slot cover.




- b). Install the new card into the riser, then replace and tighten the screw that secures the PCI card to the riser..



7. Install the left PCI riser.
 - a). Replace the left PCI riser card on the motherboard tray.



- b). Tighten the black screw on the left PCI card riser.



8. Replace the motherboard tray.

Refer to the instructions in the section [Installing the Motherboard Tray](#).

Chapter 14. Removing and Attaching the Bezel



14.1. Removing the Bezel

1. Pull out at the top of the bezel to release from the magnetic attachment.




2. Pull the bezel up to release the bezel from the pins that act as pivot points.



14.2. Attaching the Bezel

1. Attach the bottom of the bezel to the pins at the bottom of the front ears on either side of the system.



2. Pivot the bezel up to connect to the magnetic attachments.



Chapter 15. Removing and Attaching the EMI Shield

15.1. Removing the EMI Shield

1. Release the thumbscrews on the EMI shield.



2. Press down on the tabs to release the EMI shield and then pull on the black handles to release the shield.



3. Remove the shield to expose the NVIDIA NVLink bridge cards.



15.2. Attaching the EMI Shield

1. Press down on the locking tabs to allow the shield to be installed.



2. Install the shield so that the edges of the shield are flush with the chassis.



3. Tighten the thumbscrews.



Chapter 16. Front Console Board Replacement

16.1. Front Console Board Replacement Overview

This is a high-level overview of the procedure to replace the front console board on the DGX-2 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

16.2. Replacing the Front Console Board

A front console board malfunction can be determined by either

- ▶ No display or connectivity occurring after plugging in a keyboard and monitor to the front of the system, or
- ▶ The BMC system event log indicating a front temperature sensor failure.


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 front console board.
 - a). Loosen the two captive screws that secure the front console board.



- b). Pull the console board out of the system.



3. Install the new front console board.
 - a). Insert the new console board.



- b). Secure the board by tightening both captive screws.
- 4. Confirm functionality.
 - a). Power on the system.
 - b). Confirm from the BMC that the outside temperature sensor reading is available.
 - c). Confirm that the VGA output and USB ports work, using a KVM or crash cart.
- 5. Return the old module to NVIDIA Enterprise Services.

Chapter 17. Motherboard Tray Battery Replacement

17.1. Motherboard Tray Battery Replacement Overview

This is a high-level overview of the procedure to replace the DGX-2 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. Remove the motherboard tray, place on a solid, flat surface.
5. Remove the left PCI riser card.
6. Use the reference diagram on the lid of the motherboard tray to identify the battery placement.
7. Replace the dead battery with the new one.
8. Replace the left PCI riser card.
9. Insert the motherboard tray into the system.
10. Plug in all cables using the labels as a reference.
11. Apply power to the system and then boot the system.
12. Verify the date and time on the system BIOS.

17.2. Replacing the Motherboard Tray Battery




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 [Removing the Motherboard Tray](#).

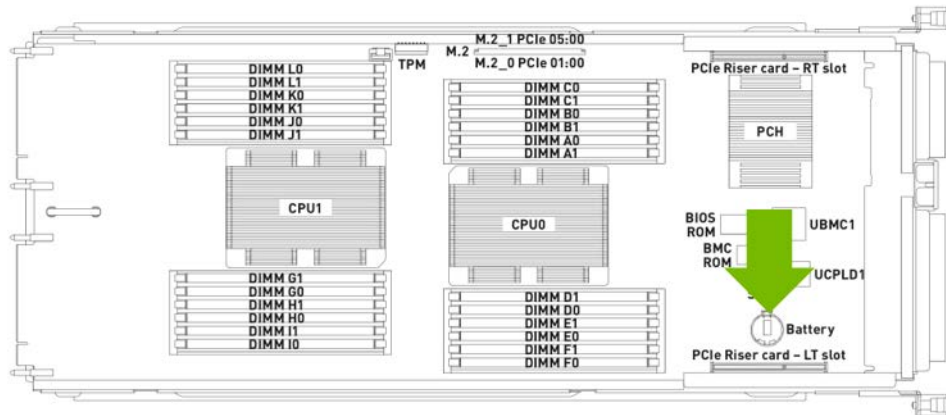
5. Remove the left PCI card riser by turning the left black screw.




6. Remove left PCI riser card from the motherboard tray.




7. Replace the battery.
 - a). Identify the battery receptacle.



- b). Remove the battery.
 - c). Make sure the replacement is a CR2032 3V lithium battery, then install the new battery with the + side up.
8. Replace left PCI riser card on the motherboard tray.



9. Tighten the black screw on the left PCI card riser.



10. Replace the motherboard tray.

Refer to the instructions in the section [Installing the Motherboard Tray](#).

11. Connect all the cables to the motherboard tray.

12. Apply power to the system and then log in.

13. 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
```

Chapter 18. I/O Tray Removal and Installation

18.1. I/O Tray Replacement Overview

This is a high-level overview of the procedure to replace the I/O tray on the DGX-2 System.

1. Get a replacement I/O tray from NVIDIA Enterprise Support.
2. Shut down the system.
3. Label all I/O tray cables and unplug them.
4. Remove the I/O tray, place on a solid, flat surface, and then open the lid.
5. Remove all ConnectX-5 cards by removing the screw that attaches the card and then removing the card.
6. Install all the ConnectX-5 cards into the new I/O tray.
7. Close the lid on the I/O 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 all ConnectX-5 cards are healthy and accessible using `nvsm health`.

18.2. Replacing the I/O Tray



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 the network ports (0-7) connected to the I/O tray for easy identification when reconnecting.




3. Remove the cables.
4. Remove the I/O tray.
 - a). Loosen the two green I/O tray screws with a Philips #2 screwdriver and pull the levers outward to release the tray.




- b). Pull the I/O tray out of the system and place it on a solid, flat work surface.




CAUTION: Exercise care when removing the tray as it is long and heavy, and do not handle the module from the rear connectors.



5. Remove the I/O tray lid.
 - a). Loosen the black screws and then push the lid towards you to release the lid.



b). Lift the lid upward.



6. Move all the ConnectX-5 cards from the old tray to the new tray.

Perform the following for each card.


a). Remove the screw that secures the card, then pull the card out of the slot.



b). Insert the card into the new I/O tray and secure with the screw removed from the previous step.

7. Install the I/O tray.


a). Replace the I/O tray lid by placing it over the module using the guiding pins and grooves.




- b). Slide the lid back so that the black screws engage with the tray, then tighten the black screws to secure the lid.



- c). Push the I/O tray back into the system.



- d). Close the levers toward the center, making sure the connectors engage with the midplane, then tighten the thumbscrews by hand or with a Phillips #2 screwdriver.



8. Confirm the I/O tray replacement.
- a). Connect all cables back into the ConnectX-5 card ports.

- b). Power on the system and log in.
- c). Confirm that the system is healthy.

```
$ sudo nvsm show health
```

There should be no new alerts listed.

- 9. Return the old I/O tray using the packaging from the new tray.


Chapter 19. Motherboard Tray Removal and Installation

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

- ▶ M.2 NVMe drives
- ▶ M.2 module riser card
- ▶ Motherboard tray battery
- ▶ Dual port CX5 PCI network adapter card
- ▶ DIMMs

19.1. Removing the Motherboard Tray


1. Loosen the two motherboard screws with a Philips #2 screwdriver and pull out and down on the levers to release the tray.




2. Pull motherboard tray out of the system and place on a work surface.




CAUTION: The motherboard tray is heavy. At least two people are required to move the motherboard tray.



3. Press on both clips at the sides of the tray and then push the lid back.




4. Lift the lid off of the motherboard tray.




19.2. Installing the Motherboard Tray

1. Align the guiding pins on the lid to the grooves on the motherboard tray chassis while lowering the tray lid to the chassis.




2. Push the lid towards the PCI cards to lock the lid in place.




Make sure you hear the click from the clips to indicate that the lid is locked in place.

3. Push the motherboard tray into its slot on the DGX-2 System



4. Once the tray is pushed in all the way, push up on the levers to complete the engagement with the chassis and finalize the insertion, then secure by tightening the thumbscrews.



Chapter 20. Identifying the Component Manufacturer

Some NVIDIA DGX-2 components are sourced from more than one manufacturer. When replacing a faulty component, be sure the replacement component is from the same manufacturer as the faulty component being replaced.

You can use the NVSM CLI to determine the manufacturer for key components as explained in the following sections of this document.

- ▶ [How to identify the U.2 NVMe manufacturer](#)
- ▶ [How to identify the M.2 NVMe manufacturer](#)
- ▶ [How to identify the PSU manufacturer](#)
- ▶ [How to identify the DIMM manufacturer](#)

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, DGX, DGX-1, DGX-2, and DGX Station are trademarks and/or registered trademarks of NVIDIA Corporation in the United States and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2020 NVIDIA Corporation. All rights reserved.