



NVIDIA DGX OS 5.0

Release Notes for Version 5.0.0

Table of Contents

| | |
|---|-----------|
| NVIDIA DGX OS 5.0, Version 5.0.0, Release Notes..... | 3 |
| About Release 5.0..... | 3 |
| Delivery and Update Mechanisms..... | 4 |
| Initial 5.0 Release..... | 4 |
| Update Advisement..... | 4 |
| Version History..... | 6 |
| DGX OS 5.0 Software Content..... | 7 |
| Package Versions in Version 5.0.0..... | 7 |
| DGX Server Firmware Version Reference..... | 7 |
| Downgrading Firmware for Mellanox ConnectX-4 Cards..... | 8 |
| Checking the Device Type..... | 8 |
| Downgrading the Firmware..... | 9 |
| Known Issues..... | 10 |
| Fixed Issues..... | 10 |
| Known Software Issues..... | 11 |
| Known DGX A100 System Issues..... | 12 |
| Known DGX-2 System Issues..... | 17 |
| Known DGX-1 System Issues..... | 18 |
| Known Issues Related to Ubuntu / Linux Kernel..... | 20 |
| Known Limitations..... | 21 |
| Appendix A. Third Party License Notice..... | 25 |

NVIDIA DGX OS 5.0, Version 5.0.0, Release Notes

This document describes the NVIDIA® DGX™ OS 5.0 software package. While DGX OS 5.0 is a unified ISO image that can be installed on DGX server as well as DGX Station, this document focuses on DGX servers. For information specific to DGX Station, refer to the [DGX OS release notes for DGX Station](#).

About Release 5.0

The following are the primary features of DGX OS 5.0:

- „ First release to support all NVIDIA DGX servers as well as DGX Station
- „ Ubuntu 20.04 LTS
- „ NVIDIA GPU driver Release 450.
 - Supports the CUDA Toolkit up to 11.0 natively, or newer versions via the compatibility module.
- „ Added rootfs encryption option, configurable during the re-imaging process.
- „ Added option to password protect the GRUB menu, configurable during the first boot process.
- „ Updated NVSM
 - Added support for custom drive partitioning
 - Added monitoring of firmware health
- „ Includes MOFED 5.1.

Delivery and Update Mechanisms

Initial 5.0 Release

DGX OS 5.0, version 5.0.0, is provided as an ISO release and is also available as an “over-the-network” update, which requires an internet connection and ability to access the NVIDIA public repositories.

You can download the ISO image file from the [NVIDIA DGX OS 5.0 Release announcement](#) on the NVIDIA Enterprise Support portal (Requires an [NVIDIA Enterprise Support](#) account).

Refer to the DGX OS 5 User Guide (<https://docs.nvidia.com/dgx/dgx-os-5-user-guide/index.html>) for the following instructions.

- How to re-image the system with the ISO image
- How to upgrade from the NVIDIA repositories
- How to install the software on air-gapped systems

Update Advisement

- **IMPORTANT:** This release is a major DGX OS release and incorporates the following updates.

- Ubuntu 20.04 LTS
- Mellanox OFED 5.1

Customers are advised to consider these updates and any effect they may have on their application. For example, some MOFED-dependent applications may be affected.

A best practice is to upgrade on select systems and verify that your applications work as expected before deploying on more systems.

- **KVM Support on DGX-2**

This release does not support the Linux Kernel-based Virtual Mode (KVM) on DGX-2 systems. Customers that require this feature should stay with the latest DGX Server OS 4.x release.

- **Update DGX OS on DGX A100 prior to updating VBIOS**

DGX A100 systems running DGX OS earlier than version 4.99.8 should be updated to the latest version before updating the VBIOS to version 92.00.18.00.0 or later (via the DGX A100 firmware update container version 20.05.12.x). Failure to do so will result in the GPUs not getting recognized.

- **NGC Containers**

In conjunction with DGX OS 5.0, customers should update their NGC containers to container release 20.10¹ if they are using multi-node training. For all other use cases, refer to the NCG Container compatibility matrix

(<https://docs.nvidia.com/deeplearning/frameworks/support-matrix/#framework-matrix-2020>)

• **Ubuntu Security Updates**

Customers are responsible for keeping the DGX server up to date with the latest Ubuntu security updates using the 'apt full upgrade' procedure. See the [Ubuntu Wiki Upgrades](#) web page for more information. Also, the Ubuntu Security Notice site (<https://usn.ubuntu.com/>) lists known Common Vulnerabilities and Exposures (CVEs), including those that can be resolved by updating the DGX OS software.

¹ See the *NVIDIA Deep Learning Frameworks* documentation website (<http://docs.nvidia.com/deeplearning/frameworks/index.htm>) for information on the latest container releases as well as <https://docs.nvidia.com/deeplearning/frameworks/user-guide/index.html> for instructions on how to access them.

Version History

This section lists the changes made in each released version of DGX OS 5.0. See [DGX OS Software Content](#) for the software component list and versions.

Version 5.0.0

- Changes since Version 4.99.11/4.6.0:
 - Ubuntu 20.04 LTS
 - Updated NVSM to 20.07.40
 - Updated DCGM to 2.0.13
 - Updated Docker Engine to 19.03.13
 - Updated NVIDIA Container Toolkit components (see [Package Versions in Version 5.0.0](#))
 - Updated Mellanox OFED to 5.1-2.4.6.0

- Changes in the network repository since the release of the DGX OS 5 version 5.0.0 ISO image.

| Date | Change | | | | | | | | | | |
|-------------------|--|------|------------------|-------|--|------------|--|------------|------------|------------|------------|
| February 23, 2021 | <ul style="list-style-type: none"> - Updated NVSM to 20.09.17 - Improved the mechanism used in upgrading from DGX OS 4.x to 5.x. | | | | | | | | | | |
| December 11, 2020 | <p>MOFED MLNX5.1-2.5.8.0</p> <p>When the update is made, the Mellanox FW updater updates the ConnectX card firmware as follows:</p> <table border="0" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Card</th> <th style="text-align: left;">Firmware Version</th> </tr> </thead> <tbody> <tr> <td colspan="2" style="border-top: 1px dashed black; padding-top: 5px;">-----</td> </tr> <tr> <td>ConnectX-4</td> <td>12.28.2006 (You may need to force a "downgrade". See Downgrading Firmware for Mellanox ConnectX-4 Cards)</td> </tr> <tr> <td>ConnectX-5</td> <td>16.28.4000</td> </tr> <tr> <td>ConnectX-6</td> <td>20.28.4000</td> </tr> </tbody> </table> <p>docker-ce 19.03.14 (Addresses CVE-2020-15257)</p> | Card | Firmware Version | ----- | | ConnectX-4 | 12.28.2006 (You may need to force a "downgrade". See Downgrading Firmware for Mellanox ConnectX-4 Cards) | ConnectX-5 | 16.28.4000 | ConnectX-6 | 20.28.4000 |
| Card | Firmware Version | | | | | | | | | | |
| ----- | | | | | | | | | | | |
| ConnectX-4 | 12.28.2006 (You may need to force a "downgrade". See Downgrading Firmware for Mellanox ConnectX-4 Cards) | | | | | | | | | | |
| ConnectX-5 | 16.28.4000 | | | | | | | | | | |
| ConnectX-6 | 20.28.4000 | | | | | | | | | | |

DGX OS 5.0 Software Content

The following tables provides version information for software included in the DGX OS ISO image as well as software installed on the system after performing an update from the repositories.

Package Versions in ISO Version 5.0.0

The following table shows the version information for software included in the DGX OS 5.0, version 5.0.0. See the section [Version 5.0.0](#) for the list of repository updates made since the ISO release.

| Component | Version |
|-------------------------------------|---|
| GPU Driver | 450.80.02 |
| NVIDIA Container Toolkit | libnvidia-container1 1.3.0-1 libnvidia-container-tools 1.3.0-1 nvidia-container-runtime 3.4.0-1 nvidia-container-toolkit 1.3.0-1 nvidia-docker2 2.5.0-1 |
| Ubuntu | 20.04 LTS |
| Ubuntu kernel | 5.4.0-52-generic ² |
| Docker Engine | 19.03.13 |
| NVIDIA System Health Monitor (NVSM) | NVSM 20.07.40 |
| Data Center GPU Management (DCGM) | 2.0.13 |
| NVIDIA System Tools | 20.09-1 |
| Mellanox OFED | MLNX 5.1-2.4.6.0 |

DGX Server Firmware Version Reference

See the appropriate release notes for the firmware versions available at the time of this DGX OS release.

- [DGX A100 System Firmware Update Container Version 20.05.12.3.](#)
- [DGX-2 System Firmware Update Container Version 20.01.25](#)
- [DGX-1 System Firmware Update Container Version 19.04.1](#)

² The resulting kernel version on your system may be a later version depending on when the repository update is performed.

Downgrading Firmware for Mellanox ConnectX-4 Cards

DGX OS 5.0.0 provides the `mlnx-fw-updater` package version 5.1-2.4.6.0 which automatically installs firmware version 12.28.2040 on ConnectX-4 devices.

Since 12.28.2006 is the recommended firmware version, on December 15 the updater package has been updated to install version 12.28.2006. However, if the firmware has already been updated to 12.28.2040, the updater will not install the downlevel firmware version since a newer version is already installed.

In this case, you will need to force the downgrade as explained in this section.

Checking the Device Type

Use the `mlxfwmanager` tool to verify whether ConnectX-4 devices are installed on the DGX system.

```

:~$ sudo mlxfwmanager

Querying Mellanox devices firmware ...
Device #1:
-----
Device Type:      ConnectX4
Part Number:     MCX455A-ECA_Ax
Description:     ConnectX-4 VPI adapter card; EDR IB (100Gb/s) and
100GbE; single-port QSFP28; PCIe3.0 x16; ROHS R6
PSID:           MT_2180110032
PCI Device Name: /dev/mst/mt4115_pciconf1
Base GUID:      248a070300945e60
Versions:
Current         Available
FW              12.28.2040     N/A
PXE             3.6.0102       N/A
UEFI            14.21.0017     N/A

```


Downgrading the Firmware

If the output indicates that ConnectX-4 devices are installed, then downgrade the firmware as follows.

1. Determine the correct firmware package name.
 - a). Switch to the `/opt/Mellanox/mlnx-fw-updater/firmware` directory (where the updater installs the firmware files) and list the contents.

```
:/opt/mellanox/mlnx-fw-updater/firmware$ ls
```

- b). Identify the correct package from the output.

```
mlxfwmanager_sriov_dis_x86_64_4115  mlxfwmanager_sriov_dis_x86_64_4119
mlxfwmanager_sriov_dis_x86_64_4123  mlxfwmanager_sriov_dis_x86_64_4127
mlxfwmanager_sriov_dis_x86_64_41686 mlxfwmanager_sriov_dis_x86_64_4117
mlxfwmanager_sriov_dis_x86_64_4121  mlxfwmanager_sriov_dis_x86_64_4125
mlxfwmanager_sriov_dis_x86_64_41682
```

2. Execute the firmware package using the `-f` flag.

```
:/opt/mellanox/mlnx-fw-updater/firmware$ sudo
./mlxfwmanager_sriov_dis_x86_64_4115 -f
```

The software queries the current firmware and then updates (downgrades) the firmware.

```
Querying Mellanox devices firmware ...
...
-----
Found 2 device(s) requiring firmware update...

Device #1: Updating FW ...
Initializing image partition - OK
Writing Boot image component - OK
Done

Device #2: Updating FW ...
Initializing image partition - OK
Writing Boot image component - OK
Done
```

3. Reboot the system to allow the updates to take effect.

```
$ sudo reboot
```

Known Issues

This chapter captures the issues related to the DGX OS software or DGX hardware at the time of the software release.

Fixed Issues

The following issues are fixed in 5.0.0.

- [All DGX systems]: When starting the DCGM service, a version mismatch error message similar to the following will appear:

```
[78075.772392] nvidia-nvswitch: Version mismatch, kernel version 450.80.02 user version 450.51.06
```
- [All DGX systems]: When issuing `nvsm show health`, the `nvsmhealth_log.txt` log file reports that the `/proc/driver/` folders are empty.
- [DGX A100]: The Mellanox software that is included in the DGX OS installed on DGX A100 system does not automatically update the Mellanox firmware as needed when the Mellanox driver is installed.
- [DGX A100]: `nvsm stress-test` Does not Stress the System if MIG is Enabled
Reported in 4.99.10
- [DGX A100]: With eight U.2 NVMe drives installed, the `nvsm-plugin-pcie` service reports `ERROR: Device not found in mapping table` for the additional four drives (for example, in response to `systemctl status nvsm*`).
Reported in 4.99.11
- [DGX A100]: When starting the Fabric Manager service, the following error is reported: `detected NVSwitch non-fatal error 10003 on NVSwitch pci`
Reported in 4.99.9
- [DGX-1]: When starting the DCGM service, a version mismatch error message similar to the following will appear:

```
[78075.772392] nvidia-nvswitch: Version mismatch, kernel version 450.80.02 user version 450.51.06
```


Reported in 4.6

Known Software Issues

The following are currently open issues with the software and that apply to all DGX systems (DGX-1, DGX-2, DGX A100).

- [Running NGC Containers Older than 20.10 May Produce "Incompatible MOFED Driver" Message](#)
- [NVSM Reports "Unknown" for Number of logical CPU cores on non-English system](#)

Running NGC Containers Older than 20.10 May Produce "Incompatible MOFED Driver" Message

Issue

DGX OS 5.0 incorporates Mellanox OFED 5.1 for high performance multi-node connectivity. Support for this version of OFED was added in NGC containers 20.10, so when running on earlier versions (or containers derived from earlier versions), a message similar to the following may appear.

```
ERROR: Detected MOFED driver 5.1-2.4.6, but this container has version
4.6-1.0.1.
    Unable to automatically upgrade this container.
    Multi-node communication may be unreliable or may result in
crashes with this version.
    This incompatibility will be resolved in an upcoming release.
```

Explanation

For applications that rely on OFED (typically those used in multi-node jobs), this is an indication that an update to NGC containers 20.10 or greater is required. For all other applications, this error can be ignored.

NVSM Reports "Unknown" for Number of logical CPU cores on non-English system

Issue

On systems set up for a non-English locale, the `nvsm show health` command lists the number of logical CPU cores as Unknown.

```
Number of logical CPU cores [None]..... Unknown
```

Explanation

This issue will be resolved in a later version of the DGX OS software.

Known DGX A100 System Issues

The following are known issues specific to the DGX-A100 server.

- [cuMemFree CUDA API Performance Regression](#)
- [`nvsm show alerts` Reports NVSwitch PCIe Link Width Degraded](#)
- [AMD Crypto Co-processor is not Supported](#)
- [Erroneous Insufficient Power Error May Occur for PCIe Slots](#)
- [USB Errors are Logged When Shutting Down the System](#)
- [BMC is not Detectable After Restoring BMC to Default](#)
- [NVSM Enumerates NVSwitches as 8-13 Instead of 0-5](#)
- [A System with Encrypted rootfs May Fail to Boot if one of the M.2 drives is Corrupted](#)

`nvsm show health` Reports Firmware as Not Authenticated

Issue

When issuing `nvsm show health`, the output shows CEC firmware components as Not Authenticated, even when they have passed authentication.

Example:

```
CEC:
  CEC Version: 3.5
  EC_FW_TAG0: Not Authenticated
  EC_FW_TAG1: Not Authenticated
  BMC FW authentication state: Not Authenticated
```

Explanation

The message can be ignored and does not affect the overall `nvsm health` output status.

cuMemFree CUDA API Performance Regression

Issue

Reported in 4.99.10

In cases when NVLINK peers are enabled, there is a performance regression of cuMemFree CUDA API.

Explanation

The cuMemFree API is usually used during application teardown and is discouraged from being used in performance-critical paths, so the regression should not impact application end-to-end performance.

`nvsm show alerts` Reports NVSwitch PCIe Link Width Degraded

Issue

NVSM raises alerts of Severity=Warning against PCIe links between NVSwitch and the Draco switch. The alert states "PCIe link width degraded" - the PCIe link width is expected to be x4 while the actual link width is x2.

There are six pairs of the PCIe links, so NVSM raises six such alerts in this condition.

Explanation

The Broadcom firmware for the synthetic switch advertises the Draco switch has PCIe link width capability of x4. This synthesized information is not reflecting the hardware capability which is of width x2. NVSM raises alerts based on this incorrect information.

This issue will be resolved with updated firmware to be provided in the DGX A100 Firmware Update Container after version 20.05.12.3. See the [DGX A100 Firmware Update Container release notes](#) for the latest firmware status.

AMD Crypto Co-processor is not Supported

Issue

Reported in 4.99.9

The DGX A100 currently does not support the AMD Cryptograph Co-processor. When booting the system, you may see the following error message in the syslog:

```
ccp initialization failed
```

Explanation

Even if the message does not appear, CCP is still not supported. The SBIOS makes zero CCP queues available to the driver, so CCP cannot be activated.

Erroneous Insufficient Power Error May Occur for PCIe Slots

Issue

Reported in 4.99.9

The DGX A100 server reports "Insufficient power" on PCIe slots when network cables are connected.

Explanation

This may occur with optical cables and indicates that the calculated power of the card + 2 optical cables is higher than what the PCIe slot can provide.

The message can be ignored.

USB Errors are Logged When Shutting Down the System

Issue

Reported 4.99.9

When rebooting the system, "USB 3-1-port1" error messages appear on the console. This occurs even with no physical USB flash drive plugged in, or without the BMC ISO image mounted.

Explanation

This issue will be addressed in a future DGX OS release. The error messages can be ignored as the system will still boot.

BMC is not Detectable After Restoring BMC to Default

Issue

Reported in 4.99.8

After using the BMC Web UI dashboard to restore the factory defaults (**Maintenance > Restore Factory Defaults**), the BMC can no longer be detected and the system is rendered unusable.

Explanation

Do not attempt to restore the factory defaults using the BMC Web UI dashboard.

NVSM Enumerates NVSwitches as 8-13 Instead of 0-5

Issue

Reported in 4.99.9

NVSM commands that list the NVSwitches (such as `nvsm show nvswitches`) will return the switches with 8-13 enumeration.

Example:

```
nvsm show /systems/localhost/nvswitches
/systems/localhost/nvswitches
Targets:
  NVSwitch10
  NVSwitch11
  NVSwitch12
  NVSwitch13
  NVSwitch8
  NVSwitch9
```

Explanation

Currently, NVSM recognizes NVSwitches as graphics devices, and enumerates them as a continuation of the GPU 0-7 enumeration.

A System with Encrypted rootfs May Fail to Boot if one of the M.2 drives is Corrupted

Issue

Reported in 4.99.9

On systems with encrypted rootfs, if one of the M.2 drives is corrupted, the system stops at the BusyBox shell when booting.

Explanation and Workaround

The inactive RAID array (due to the corrupted M.2 drive) is not getting converted to a degraded RAID array.

To work around, perform the following within the BusyBox.

1. Issue the following.

```
$ mdadm --run /dev/md?*
```

2. Wait a few seconds for the RAID and crypt to be discovered.

3. Exit.

```
$ exit
```


Known DGX-2 System Issues

The following are known issues specific to the DGX-2 server.

- [Applications that call the cuCTXCreate API may experience a performance drop](#)
- [NVSM Does not Detect Downgraded GPU PCIe Link](#)

Applications that Call the cuCTXCreate API May Experience a Performance Drop

Issue

When some applications call cuCtxCreate, cuGLCtxCreate, or cuCtxDestroy, there may be a drop in performance.

Explanation

This occurs with Ubuntu 20.04, but not previous versions. The issue affects applications that perform graphics/compute interoperations or have a plugin mechanism for CUDA, where every plugin creates its own context, or video streaming applications where computations are needed. Examples include ffmpeg, Blender, simpleDrvRuntime, and cuSolverSp_LinearSolver.

This is not expected to impact deep learning training.

NVSM Does not Detect Downgraded GPU PCIe Link

Issue

If the GPU PCIe link is downgraded to Gen1, NVSM still reports the GPU health status as OK.

Explanation and Resolution

The NVSM software currently does not check for this condition. The check will be added in a future software release.

Known DGX-1 System Issues

The following are known issues specific to the DGX-1 server.

- [Applications that call the cuCTXCreate API may experience a performance drop](#)
- [Forced Reboot Hangs the OS](#)
- [Software Power Cap Not Reported Correctly by nvidia-smi](#)

Applications that Call the cuCTXCreate API May Experience a Performance Drop

Issue

When some applications call `cuCtxCreate`, `cuGLCtxCreate`, or `cuCtxDestroy`, there may be a drop in performance.

Explanation

This occurs with Ubuntu 20.04, but not previous versions. The issue affects applications that perform graphics/compute interoperations or have a plugin mechanism for CUDA, where every plugin creates its own context, or video streaming applications where computations are needed. Examples include `ffmpeg`, `Blender`, `simpleDrvRuntime`, and `cuSolverSp_LinearSolver`.

This is not expected to impact deep learning training.

Forced Reboot Hangs the OS

Issue

When issuing `reboot -f` (forced reboot), I/O error messages appear on the console and then the system hangs.

The system reboots normally when issuing `reboot`.

Resolution

This issue will be resolved in a future version of the DGX OS.

Software Power Cap Not Reported Correctly by nvidia-smi

Issue

On DGX-1 systems with Pascal GPUs, nvidia-smi does not report Software Power Cap as "Active" when clocks are throttled by power draw.

Explanation and Resolution

This issue is with nvidia-smi reporting and not with the actual functionality.

Known Issues Related to Ubuntu / Linux Kernel

The following are known issues related to the Ubuntu OS or the Linux kernel that affect the DGX server.

- [System May Slow Down When Using `mpirun`](#)

System May Slow Down When Using `mpirun`

Issue

Customers running Message Passing Interface (MPI) workloads may experience the OS becoming very slow to respond. When this occurs, a log message similar to the following would appear in the kernel log:

```
kernel BUG at /build/linux-fQ94TU/linux-4.4.0/fs/ext4/inode.c:1899!
```

Explanation and Workaround

Due to the current design of the Linux kernel, the condition may be triggered when `get_user_pages` is used on a file that is on persistent storage. For example, this can happen when `cudaHostRegister` is used on a file path that is stored in an ext4 filesystem. DGX systems implement `/tmp` on a persistent ext4 filesystem.



NOTE: If you performed this workaround on a previous DGX OS software version, you do not need to do it again after updating to the latest DGX OS version.

In order to avoid using persistent storage, MPI can be configured to use shared memory at `/dev/shm` (this is a temporary filesystem).

If you are using Open MPI, then you can solve the issue by configuring the Modular Component Architecture (MCA) parameters so that `mpirun` uses the temporary file system in memory.

For details on how to accomplish this, see the Knowledge Base Article [DGX System Slows Down When Using `mpirun`](#) (requires login to the [NVIDIA Enterprise Support portal](#)).

Known Limitations

This section lists known limitations and other issues that will not be fixed.

- [\[DGX A100\]: Hot-plugging of Storage Drives is not Supported](#)
- [\[DGX A100\]: Syslog Contains Numerous "SM LID is 0, maybe no SM is running" Error Messages](#)
- [\[DGX-2\] Serial Over LAN Does not Work After Cold Resetting the BMC](#)
- [\[DGX-2\] Some BMC Dashboard Quick Links Appear Erroneously](#)
- [\[DGX-2\] Applications Cannot be Run Immediately Upon Powering on the DGX-2](#)
- [\[DGX-1\] Script Cannot Recreate RAID Array After Re-inserting a Known Good SSD](#)

[DGX A100]: Hot-plugging of Storage Drives is not Supported

Issue

Hot-plugging or hot-swapping one of the storage drives might result in system instability or incorrect device reporting.

Workaround and Resolution

Turn off the system before removing and replacing any of the storage drives.

[DGX A100]: Syslog Contains Numerous "SM LID is 0, maybe no SM is running" Error Messages

Issue

The system log (`/var/log/syslog`) contains multiple "SM LID is 0, maybe no SM is running" error message entries.

Explanation and Workaround

This issue is the result of the `srp_daemon` within the Mellanox driver. The daemon is used to discover and connect to InfiniBand SCSI RDMA Protocol (SRP) targets.

If you are not using RDMA, then disable the `srp_daemon` as follows.

```
$ sudo systemctl disable srp_daemon.service
$ sudo systemctl disable srptools.service
```

[DGX-2] Serial Over LAN Does not Work After Cold Resetting the BMC

Issue

After performing a cold reset on the BMC (`ipmitool mc reset cold`) while serial over LAN (SOL) is active, you cannot restart a SOL session.

Workaround

To re-active SOL, either

- Reboot the system, or
- Kill and then restart the process as follows.
 - a) Identify the Process ID of the SOL TTY process by running the following.


```
ps -ef | grep "/sbin/agetty -o -p -- \u --keep-baud 115200,38400,9600 ttyS0 vt220"
```
 - b) Kill the process.


```
kill <PID>
```

 where `<PID>` is the Process ID returned by the previous command.
 - c) Either wait for the cron job to respawn the process or manually restart the process by running


```
/sbin/agetty -o -p -- \u --keep-baud 115200,38400,9600 ttyS0 vt220
```

[DGX-2] Some BMC Dashboard Quick Links Appear Erroneously

Issue

On the BMC dashboard, the following Quick Links appear by mistake and should not be used.

- Maintenance->Firmware Update
- Settings->NvMeManagement->NvMe P3700Vpd Info

[DGX-2] Applications Cannot be Run Immediately Upon Powering on the DGX-2

Issue

When attempting to run an application that uses the GPUs immediately upon powering on the DGX-2 system, you may encounter the following error.

`CUDA_ERROR_SYSTEM_NOT_READY`

Explanation and Workaround

The DGX-2 uses a fabric manager service to manage communication between all the GPUs in the system. When the DGX-2 system is powered on, the fabric manager initializes all the GPUs. This can take approximately 45 seconds. Until the GPUs are initialized, applications that attempt to use them will fail.

If you encounter the error, wait and launch the application again.

[DGX-1] Script Cannot Recreate RAID Array After Re-inserting a Known Good SSD

Issue

When a good SSD is removed from the DGX-1 RAID 0 array and then re-inserted, the script to recreate the array fails.

Explanation and Workaround

After re-inserting the SSD back into the system, the RAID controller sets the array to offline and marks the re-inserted SSD as `Unconfigured_Bad (UBad)`. The script will fail when attempting to rebuild an array when one or more of the SSDs are marked `UBad`.

To recreate the array in this case,

1. Set the drive back to a good state.

```
# sudo /opt/MegaRAID/storcli/storcli64  
/c0/e<enclosure_id>/s<drive_slot> set good
```

2. Run the script to recreate the array.

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

Appendix A. Third Party License Notice

This NVIDIA product contains third party software that is being made available to you under their respective open source software licenses. Some of those licenses also require specific legal information to be included in the product. This section provides such information.

msecli

The msecli utility (<https://www.micron.com/products/solid-state-storage/storage-executive-software>) is provided under the following terms:

Micron Technology, Inc. Software License Agreement

PLEASE READ THIS LICENSE AGREEMENT ("AGREEMENT") FROM MICRON TECHNOLOGY, INC. ("MTI") CAREFULLY: BY INSTALLING, COPYING OR OTHERWISE USING THIS SOFTWARE AND ANY RELATED PRINTED MATERIALS ("SOFTWARE"), YOU ARE ACCEPTING AND AGREEING TO THE TERMS OF THIS AGREEMENT. IF YOU DO NOT AGREE WITH THE TERMS OF THIS AGREEMENT, DO NOT INSTALL THE SOFTWARE.

LICENSE: MTI hereby grants to you the following rights: You may use and make one

(1) backup copy the Software subject to the terms of this Agreement. You must maintain all copyright notices on all copies of the Software. You agree not to modify, adapt, decompile, reverse engineer, disassemble, or otherwise translate the Software. MTI may make changes to the Software at any time without notice to you.

In addition MTI is under no obligation whatsoever to update, maintain, or provide new versions or other support for the Software.

OWNERSHIP OF MATERIALS: You acknowledge and agree that the Software is proprietary property of MTI (and/or its licensors) and is protected by United States copyright law and international treaty provisions. Except as expressly provided herein, MTI does not grant any express or implied right to you under any patents, copyrights, trademarks, or trade secret information. You further acknowledge and agree that all right, title, and interest in and to the Software, including associated proprietary

rights, are and shall remain with MTI (and/or its licensors). This Agreement does not convey to you an interest in or to the Software, but only a limited right to use and copy the Software in accordance with the terms of this Agreement. The Software is licensed to you and not sold.

DISCLAIMER OF WARRANTY: THE SOFTWARE IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND. MTI EXPRESSLY DISCLAIMS ALL WARRANTIES EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, NONINFRINGEMENT OF THIRD PARTY RIGHTS, AND ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. MTI DOES NOT WARRANT THAT THE SOFTWARE WILL MEET YOUR REQUIREMENTS, OR THAT THE OPERATION OF THE SOFTWARE WILL BE UNINTERRUPTED OR ERROR-FREE. FURTHERMORE, MTI DOES NOT MAKE ANY REPRESENTATIONS REGARDING THE USE OR THE RESULTS OF THE USE OF THE SOFTWARE IN TERMS OF ITS CORRECTNESS, ACCURACY, RELIABILITY, OR OTHERWISE. THE ENTIRE RISK ARISING OUT OF USE OR PERFORMANCE OF THE SOFTWARE REMAINS WITH YOU. IN NO EVENT SHALL MTI, ITS AFFILIATED COMPANIES OR THEIR SUPPLIERS BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, INCIDENTAL, OR SPECIAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION, OR LOSS OF INFORMATION) ARISING OUT OF YOUR USE OF OR INABILITY TO USE THE SOFTWARE, EVEN IF MTI HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Because some jurisdictions prohibit the exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

TERMINATION OF THIS LICENSE: MTI may terminate this license at any time if you are in breach of any of the terms of this Agreement. Upon termination, you will immediately destroy all copies the Software.

GENERAL: This Agreement constitutes the entire agreement between MTI and you regarding the subject matter hereof and supersedes all previous oral or written communications between the parties. This Agreement shall be governed by the laws of the State of Idaho without regard to its conflict of laws rules.

CONTACT: If you have any questions about the terms of this Agreement, please contact MTI's legal department at (208) 368-4500.

By proceeding with the installation of the Software, you agree to the terms of this Agreement. You must agree to the terms in order to install and use the Software.

Mellanox (OFED)

MLNX OFED (<http://www.mellanox.com/>) is provided under the following terms:

Copyright (c) 2006 Mellanox Technologies.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED.

IN NO EVENT SHALL THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Notice

THE INFORMATION IN THIS DOCUMENT AND ALL OTHER INFORMATION CONTAINED IN NVIDIA DOCUMENTATION REFERENCED IN THIS DOCUMENT IS PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE INFORMATION FOR THE PRODUCT, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the product described in this document shall be limited in accordance with the NVIDIA terms and conditions of sale for the product.

THE NVIDIA PRODUCT DESCRIBED IN THIS DOCUMENT IS NOT FAULT TOLERANT AND IS NOT DESIGNED, MANUFACTURED OR INTENDED FOR USE IN CONNECTION WITH THE DESIGN, CONSTRUCTION, MAINTENANCE, AND/OR OPERATION OF ANY SYSTEM WHERE THE USE OR A FAILURE OF SUCH SYSTEM COULD RESULT IN A SITUATION THAT THREATENS THE SAFETY OF HUMAN LIFE OR SEVERE PHYSICAL HARM OR PROPERTY DAMAGE (INCLUDING, FOR EXAMPLE, USE IN CONNECTION WITH ANY NUCLEAR, AVIONICS, LIFE SUPPORT OR OTHER LIFE CRITICAL APPLICATION). NVIDIA EXPRESSLY DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY OF FITNESS FOR SUCH HIGH RISK USES. NVIDIA SHALL NOT BE LIABLE TO CUSTOMER OR ANY THIRD PARTY, IN WHOLE OR IN PART, FOR ANY CLAIMS OR DAMAGES ARISING FROM SUCH HIGH RISK USES.

NVIDIA makes no representation or warranty that the product described in this document will be suitable for any specified use without further testing or modification. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to ensure the product is suitable and fit for the application planned by customer and to do 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 does not accept any 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.

Other than the right for customer to use the information in this document with the product, no other license, either expressed or implied, is hereby granted by NVIDIA under this document. Reproduction of information in this document is permissible only if reproduction is approved by NVIDIA in writing, is reproduced without alteration, and is accompanied by all associated conditions, limitations, and notices.

Trademarks

NVIDIA, the NVIDIA logo, DGX, DGX-1, and DGX-2 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.

