



NVIDIA DGX OS SERVER RELEASE 4.4

DA-08260-441_v01 | April 2020

Release Notes and Update Guide




TABLE OF CONTENTS

NVIDIA DGX OS Server 4.4 Release Notes	3
About Release 4.4	3
Delivery and Update Mechanisms	4
Initial 4.4 Release	4
Patch Updates.....	4
Update Advisement.....	5
Version History	6
DGX OS Server Software Content.....	7
Package Versions After Getting the 4.4.1 Patch Update	7
Package Versions in the 4.4.0 ISO Image	8
DGX Server Firmware Version Reference	9
DGX-2 / DGX-2H Firmware.....	9
DGX-1 with Tesla P100, Tesla V100 Firmware.....	9
Updating the Software	10
Preparing for Updating the Software	10
Connecting to the DGX server Console	10
Verifying the DGX Server Connection to the Repositories	11
Performing the Updates	12
Update Path Instructions.....	12
Updating from Release 4.1 and Later	13
Updating from Release 4.0 (4.0.1 or later).....	15
Updating from 3.1.x.....	17
Known Issues	20
Resolved Issues	20
Known Software Issues	21
Known DGX-2 System Issues	22
NVIDIA KVM Issues	24
Known DGX-1 System Issues	26
Known Issues Related to Ubuntu / Linux Kernel.....	28
Appendix A. Third Party License Notice	31

NVIDIA DGX OS SERVER 4.4 RELEASE NOTES

This document describes the NVIDIA® DGX™ OS Server Release 4.4 software package.

ABOUT RELEASE 4.4

The following are the primary new features of DGX OS Server Release 4.4 since Release 4.3:

- ▶ Updated NVIDIA GPU driver to [418.126.02](#)
- ▶ Updated NVSM software to version [20.01.15](#)
- ▶ Added the ability to disable and re-enable CPU mitigations
- ▶ Added support for PXE boot
- ▶ Added support for Mellanox CX6 cards on DGX-1 (Tesla V100)

DELIVERY AND UPDATE MECHANISMS

Initial 4.4 Release

DGX OS Server Release 4.4, version 4.4.0, is provided as an ISO image available from NVIDIA Enterprise Support in the event the server needs to be re-imaged. Version 4.4.0 is also provided as an “over-the-network” update, which requires an internet connection and ability to access the NVIDIA public repositories.

Refer to the *DGX-2 User Guide* (<https://docs.nvidia.com/dgx/dgx2-user-guide/index.html>) and *DGX-1 User Guide* (<https://docs.nvidia.com/dgx/dgx1-user-guide/index.html>) for the following instructions.

- ▶ How to re-image the system with the ISO image
- ▶ How to install the software on air-gapped systems

Patch Updates

To deliver security updates and critical bug fixes quickly, NVIDIA has changed how DGX OS Server software updates are distributed.

Previously, all DGX OS software updates were full releases and were available to all users without the need to opt in. Starting with DGX OS 4.1.1, NVIDIA introduced patch updates to DGX OS software in addition to the existing full DGX OS releases.

Patch updates are package upgrades to individual components in DGX OS software that are delivered through the DGX OS update repository. You must opt in to receive patch updates. If you don't opt in to patch updates, you will get these updates in the next full release, which aggregates the changes in all patch updates since the previous full release. Therefore, you can choose between always having the latest updates or reducing the frequency of software updates to your system by updating only when NVIDIA issues a full DGX OS release.

To obtain patch updates to Release 4.4, be sure you have first opted in to the patch updates by installing the update package. See [Update Path Instructions](#) for details. If you already opted in to patch updates for Release 4.1, 4.2, or 4.3, then you will automatically receive patch updates for Release 4.4 when performing an update.

For more information about full DGX OS Server releases and patch updates, and how they are identified, see [DGX OS Server Release Number Scheme](#).

UPDATE ADVISEMENT

▶ NVIDIA GPU Cloud Containers

In conjunction with DGX OS Server v4.4, customers should update their NVIDIA GPU Cloud containers to the latest container release¹.

▶ 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/dgx/index.htm>) for information on the latest container releases as well as <https://docs.nvidia.com/deeplearning/dgx/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 Release 4.4. See [DGX OS Server Software Content](#) for the software component list and versions.

Path Update 4.4.1

- ▶ Updated DGX KVM package.
Fixed a bug that caused GPU VMs to fail to launch.

Version 4.4.0

- ▶ Initial Release 4.4 version.
- ▶ Changes since Version 4.3.0:
 - Updated NVIDIA GPU driver to [418.126.02](#)
 - Updated NVSM software to version [20.01.15](#)
 - Updated Docker-CE to [19.03.5](#)

Includes native support for running GPU-accelerated containers which eliminates the need to use the NVIDIA Container Runtime for Docker (`nvidia-docker2` package).



NOTE: While the `nvidia-docker2` package is included in this release, support for it may be discontinued in future DGX OS releases.

- Updated the NVIDIA Container Toolkit to 3.1.4
- Updated the Mellanox OFED to MLNX 4.7-3.2.9.0.
- Added the ability to disable and re-enable CPU mitigations
- Added support for PXE boot
- Added support for Mellanox CX6 cards on DGX-1 (Tesla V100)

DGX OS SERVER SOFTWARE CONTENT

The following tables provides version information for software included in the DGX OS Server ISO image as well as software installed on the system after getting subsequent updates.

Package Versions After Getting the 4.4.1 Patch Update

The following table shows the version information for software included in the DGX OS Server patch update 4.4.1.

Component	Version
GPU Driver	418.126.02
NVIDIA Container Toolkit	3.1.4-1
Ubuntu	18.04.3 LTS
Ubuntu kernel	4.15.0-76-generic ²
Docker Engine	19.03.5-3
NVIDIA System Health Monitor (NVSM)	NVSM 20.01.15
Data Center GPU Management (DCGM)	1.7.2
Mellanox OFED	MLNX 4.7-3.2.9.0
KVM Package Components (DGX-2 only)	Version
dgx-kvm-sw	19.07.0
dgx-kvm-host-utils	20.03.0
dgx-kvm-host-conf	19.04.0
qemu	2.11+dfsg-1ubuntu7.15-Nvidia.19.02.1
dgx-kvm-image-4-3-0	dgx-kvm-image-4-3-0-191126-f1bbec.0_amd64.deb
dgx-kvm-image-4-2-0	dgx-kvm-image-4-2-0-190925-5ea264.0_amd64.deb
dgx-kvm-image-4-1-1	dgx-kvm-image-4.1.1-190814-c14357.0_amd64.deb
dgx-kvm-image-4-1-0	dgx-kvm-image-4-1-0-190424-2c5230.6_amd64.deb
dgx-kvm-image-4-0-6	dgx-kvm-image-4-0-6-190424-6df14a.6_amd64.deb
dgx-kvm-image-4-0-5	dgx-kvm-image-4-0-5-190222-a27277.6_amd64.deb
dgx-kvm-image-4-0-4	dgx-kvm-image-4-0-4-181116-cc1114.6_amd64.deb
dgx-kvm-image-4-0-3	dgx-kvm-image-4-0-3-181023-f812a7.6_amd64.deb
dgx-kvm-image-4-0-2	dgx-kvm-image-4-0-2-180918-d47d38.6_amd64.deb

² If updating over-the-network, your kernel version may be a later version depending on when the update is performed.

Package Versions in the 4.4.0 ISO Image

The following table shows the version information for software included in the DGX OS Server ISO image for version 4.4.0.

Component	Version
GPU Driver	418.126.02
NVIDIA Container Toolkit	3.1.4-1
Ubuntu	18.04.3 LTS
Ubuntu kernel	4.15.0-76-generic ³
Docker Engine	19.03.5-3
NVIDIA System Health Monitor (NVSM)	NVSM 20.01.15
Data Center GPU Management (DCGM)	1.7.2
Mellanox OFED	MLNX 4.7-3.2.9.0
KVM Package Components (DGX-2 only)	Version
dgx-kvm-sw	19.07.0
dgx-kvm-host-utils	19.11.1
dgx-kvm-host-conf	19.04.0
qemu	2.11+dfsg-1ubuntu7.15-Nvidia.19.02.1
dgx-kvm-image-4-3-0	dgx-kvm-image-4-3-0-191126-f1bbec.0_amd64.deb
dgx-kvm-image-4-2-0	dgx-kvm-image-4-2-0-190925-5ea264.0_amd64.deb
dgx-kvm-image-4-1-1	dgx-kvm-image-4.1.1-190814-c14357.0_amd64.deb
dgx-kvm-image-4-1-0	dgx-kvm-image-4-1-0-190424-2c5230.6_amd64.deb
dgx-kvm-image-4-0-6	dgx-kvm-image-4-0-6-190424-6df14a.6_amd64.deb
dgx-kvm-image-4-0-5	dgx-kvm-image-4-0-5-190222-a27277.6_amd64.deb
dgx-kvm-image-4-0-4	dgx-kvm-image-4-0-4-181116-cc1114.6_amd64.deb
dgx-kvm-image-4-0-3	dgx-kvm-image-4-0-3-181023-f812a7.6_amd64.deb
dgx-kvm-image-4-0-2	dgx-kvm-image-4-0-2-180918-d47d38.6_amd64.deb

³ If updating over-the-network, your kernel version may be a later version depending on when the update is performed.

DGX SERVER FIRMWARE VERSION REFERENCE

The following tables show the firmware and BIOS versions for the DGX hardware at the time of this release. Information is provided for reference purposes.

DGX-2 / DGX-2H Firmware

Component	Version
BMC	V01.05.12
SBIOS	V0.24
VBIOS DGX-2 DGX-2H	88.00.6B.00.01 88.00.6B.00.08
InfiniBand FW	16.26.4012 (updated as part of the DGX OS update)
PSU FW	2.7
OS Drive Samsung	CXV8601Q
Data Drive Micron	101008R0
GPU Sled FPGA	3.1

DGX-1 with Tesla P100, Tesla V100 Firmware

Component	Version
BMC	3.36.30
SBIOS	S2W_3A10
VBIOS DGX-1 Tesla P100 DGX-1 Tesla V100 (16 GB) DGX-1 Tesla V100 (32 GB)	86.00.41.00.05 88.00.18.00.01 88.00.80.00.04
PSU	00.03.07
SSD	GXM1103Q

UPDATING THE SOFTWARE


These instructions explain how to update the DGX OS server software through an internet connection to the NVIDIA public repository. The process updates a DGX system image to the latest versions of the entire DGX software stack, including the drivers.


Perform the updates using commands on the DGX server console.

PREPARING FOR UPDATING THE SOFTWARE

Connecting to the DGX server Console

Connect to the DGX server console using either a direct connection or a remote connection through the BMC.

 **NOTE:** SSH can be used to perform the update. However, if the Ethernet port is configured for DHCP, there is the potential that the IP address can change after the DGX server is rebooted during the update, resulting in loss of connection. If this happens, connect using either a direct connection or through the BMC to continue the update process.

 **WARNING:** Connect directly to the DGX server console if the DGX is connected to a 172.17.xx.xx subnet.

DGX OS Server software installs Docker CE which uses the 172.17.xx.xx subnet by default for Docker containers. If the DGX server is on the same subnet, you will not be able to establish a network connection to the DGX server.

Refer to the appropriate DGX-1 or DGX-2 User Guide for instructions on how to change the default Docker network settings after performing the update.

Direct Connection

1. Connect a display to the VGA connector and a keyboard to any one of the USB ports.
2. Power on the DGX server.

Remote Connection through the BMC

Refer to the appropriate user guide (DGX-1 or DGX-2) for instructions on establishing a remote connection to the BMC.

Verifying the DGX Server Connection to the Repositories

Before attempting to perform the update, verify that the DGX server network connection can access the public repositories and that the connection is not blocked by a firewall or proxy.

On DGX-1 Systems if Upgrading from Version 2.x.

Enter the following on the DGX-1 system.

```
$ wget -O f1-changelogs http://changelogs.ubuntu.com/meta-release-lts
$ wget -O f2-archive
http://archive.ubuntu.com/ubuntu/dists/xenial/Release
$ wget -O f3-usarchive
http://us.archive.ubuntu.com/ubuntu/dists/xenial/Release
$ wget -O f4-security
http://security.ubuntu.com/ubuntu/dists/xenial/Release
$ wget -O f5-download
https://download.docker.com/linux/ubuntu/dists/xenial/Release
$ wget -O f6-international
http://international.download.nvidia.com/dgx/repos/dists/xenial/Release
```

All the **wget** commands should be successful and there should be six files in the directory with non-zero content.

On DGX-2 and DGX-1 Systems

Enter the following on the DGX system

```
$ wget -O f1-changelogs http://changelogs.ubuntu.com/meta-release-lts
$ wget -O f2-archive
http://archive.ubuntu.com/ubuntu/dists/bionic/Release
$ wget -O f3-usarchive
http://us.archive.ubuntu.com/ubuntu/dists/bionic/Release
$ wget -O f4-security
http://security.ubuntu.com/ubuntu/dists/bionic/Release
$ wget -O f5-download
http://download.docker.com/linux/ubuntu/dists/bionic/Release
$ wget -O f6-international
http://international.download.nvidia.com/dgx/repos/bionic/dists/bionic/Release
```

All the **wget** commands should be successful and there should be six files in the directory with non-zero content

PERFORMING THE UPDATES

Update Path Instructions

Follow the instructions corresponding to your current DGX OS server software.

Updating from Release 4.1 and later

Follow the instructions at [Updating from Release 4.1 and later](#).

Updating from Release 4.0 (Version 4.0.1 or later only)

Follow the instructions at [Updating from 4.0.1 \(or Later\)](#).

Updating from Release 3.1

Follow the instructions at [Updating from Release 3.1](#).

Updating from Release 2.x

- a) Update from Release 2.x to the latest Release 3.1 as described in the [DGX OS 3.1.8 Release Notes](#).
- b) [Update from Release 3.1](#).

Updating from Release 4.1 and Later

See the section [Connecting to the DGX Console](#) for guidance on connecting to the console to perform the update.



CAUTION: These instructions update all software for which updates are available from your configured software sources, including applications that you installed yourself. If you want to prevent an application from being updated, you can instruct the Ubuntu package manager to keep the current version. For more information, see [Introduction to Holding Packages](#) on the Ubuntu Community Help Wiki.

Update Instructions

1. If you have not already done so, verify that your DGX system can access the public repositories as explained in [Verifying the DGX Server Connection to the Repositories](#).

2. Update the list of available packages and their versions.

```
$ sudo apt update
```

3. Review the packages that will be updated.

```
$ sudo apt full-upgrade -s
```

To prevent an application from being updated, instruct the Ubuntu package manager to keep the current version. See [Introduction to Holding Packages](#).

4. Upgrade to version 4.4.0.

This step will update your system to 4.4.1 if you have already opted in to patch updates.

```
$ sudo apt full-upgrade
```

- Answer any questions that appear.
 - Most questions require a Yes or No response. When asked to select the grub configuration to use, select the current one on the system.
 - Other questions will depend on what other packages were installed before the update and how those packages interact with the update.
 - If a message appears indicating that `nvidia-docker.service` failed to start, you can disregard it and continue with the next step. The service will start normally at that time.
5. Reboot the system.

Obtaining Patch Updates

If you have not already opted in to patch updates and wish to update beyond version 4.4.0, perform the following additional instructions.

6. Update the list of available packages and their versions.

```
$ sudo apt update
```

7. Install the update package.

```
$ sudo apt install -y dgx-bionic-updates-repo
```

8. Update to the latest 4.4 patch update.

```
$ sudo apt full-upgrade
```

Recovering from an Interrupted or Failed Update

If the script is interrupted during the update, such as from a loss of power or loss of network connection, then restore power or restore the network connection, whichever caused the interruption.

- ▶ If the system encounters a kernel panic after you restore power and reboot the DGX-2, you will not be able to perform the over-the-network update. You will need to re-image the DGX-2 with the latest image (see the [DGX-2 User Guide](#) for instructions) and then perform the network update.

If you are successfully returned to the Linux command line, continue following the instructions from step 2 in the [Updating from Release 4.1 or 4.2](#) update instructions

Updating from Release 4.0 (4.0.1 or later)

For Release 4.0, only updates from versions 4.0.1 and later are supported with these instructions. To update from version 4.0.0, you must re-image the system.

See the section [Connecting to the DGX Console](#) for guidance on connecting to the console to perform the update.

CAUTION: These instructions update all software for which updates are available from your configured software sources, including applications that you installed yourself. If you want to prevent an application from being updated, you can instruct the Ubuntu package manager to keep the current version. For more information, see [Introduction to Holding Packages](#) on the Ubuntu Community Help Wiki.

Update Instructions

1. If you have not already done so, verify that your DGX system can access the public repositories as explained in [Verifying the DGX Server Connection to the Repositories](#).

2. Update the list of available packages and their versions.

```
$ sudo apt update
```

3. Install the 4.1.0 components from the repository.

```
$ sudo apt install -y dgx-bionic-r418+cuda10.1-repo
```

4. Update the new list of packages and their versions.

```
$ sudo apt update
```

5. Review the packages that will be updated.

```
$ sudo apt full-upgrade -s
```

To prevent an application from being updated, instruct the Ubuntu package manager to keep the current version. See [Introduction to Holding Packages](#).

6. Upgrade to version 4.4.0.

This step will update your system to 4.4.1 if you have already opted in to patch updates.

```
$ sudo apt full-upgrade
```

- Answer any questions that appear.

- Most questions require a Yes or No response. When asked to select the grub configuration to use, select the current one on the system.
 - Other questions will depend on what other packages were installed before the update and how those packages interact with the update.
 - If a message appears indicating that `nvidia-docker.service` failed to start, you can disregard it and continue with the next step. The service will start normally at that time.
7. Reboot the system.

Obtaining Patch Updates

If you have not already opted in to patch updates and wish to update beyond version 4.4.0, perform the following additional instructions.

8. Update the list of available packages and their versions.

```
$ sudo apt update
```

9. Install the update package.

```
$ sudo apt install -y dgx-bionic-updates-repo
```

10. Update to the latest 4.4 patch update.

```
$ sudo apt full-upgrade
```

Recovering from an Interrupted or Failed Update

If the script is interrupted during the update, such as from a loss of power or loss of network connection, then restore power or restore the network connection, whichever caused the interruption.

- ▶ If the system encounters a kernel panic after you restore power and reboot the DGX-2, you will not be able to perform the over-the-network update. You will need to re-image the DGX-2 with the latest image (see the [DGX-2 User Guide](#) for instructions) and then perform the network update.

If you are successfully returned to the Linux command line, continue following the instructions from step 2 in the [Updating from Version 4.0.1 \(or Later\)](#) update instructions.

Updating from 3.1.x

See the section [Connecting to the DGX Console](#) for guidance on connecting to the console to perform the update.



CAUTION: These instructions update all software for which updates are available from your configured software sources, including applications that you installed yourself. If you want to prevent an application from being updated, you can instruct the Ubuntu package manager to keep the current version. For more information, see [Introduction to Holding Packages](#) on the Ubuntu Community Help Wiki.

Update Instructions

The following instructions update your system to DGX OS 4.4.0.

1. If you have not already done so, verify that your DGX-1 system can access the public repositories as explained in [Verifying the DGX Server Connection to the Repositories](#).

2. Update the list of available packages and their versions.

```
$ sudo apt update
```

3. Install any updates.

```
$ sudo apt -y full-upgrade
```

4. Install dgx-release-upgrade.

```
$ sudo apt install -y dgx-release-upgrade
```

5. Begin the update process.

```
$ sudo dgx-release-upgrade
```

If you are using a proxy server, then add the **-E** option to keep your proxy environment variables.

Example:

```
$ sudo -E dgx-release-upgrade
```

6. After starting the update process, respond to the presented options as follows:
 - Press **y** if you are logged in to the DGX server remotely through secure shell (SSH) and are asked if you want to continue running under SSH.

```
Continue running under SSH?
```

```
This session appears to be running under ssh. It is not recommended
to perform a upgrade over ssh currently because in case of failure
it is harder to recover.
```

```
If you continue, an additional ssh daemon will be started at port
'1022'.
```

```
Do you want to continue?
```

```
Continue [yN]
```

An additional sshd daemon is started.

Press **Enter** in response to the following message.

```
Starting additional sshd
```

```
To make recovery in case of failure easier, an additional sshd will
be started on port '1022'. If anything goes wrong with the running
ssh you can still connect to the additional one.
```

```
If you run a firewall, you may need to temporarily open this port.
As this is potentially dangerous it's not done automatically. You
can open the port with e.g.:
```

```
'iptables -I INPUT -p tcp --dport 1022 -j ACCEPT'
```

```
To continue please press [ENTER]
```

- Press **Enter** in response to the message warning you that third-party sources are disabled.

```
Third party sources disabled
```

```
Some third party entries in your sources.list were disabled. You
can re-enable them after the upgrade with the 'software-properties'
tool or your package manager.
```

```
To continue please press [ENTER]
```

- Press **N** if prompted about `dgx.list` configuration choices.

```
Configuration file '/etc/apt/sources.list.d/dgx.list'
```

```
==> Modified (by you or by a script) since installation.
```

```
==> Package distributor has shipped an updated version.
```

```
What would you like to do about it ? Your options are:
```

```
Y or I : install the package maintainer's version
```

```
N or O : keep your currently-installed version
```

```
D      : show the differences between the versions
```

```
Z      : start a shell to examine the situation
```

```
The default action is to keep your current version.
```

```
*** dgx.list (Y/I/N/O/D/Z) [default=N] ?
```

- When prompted to resolve other configuration files, evaluate the changes before accepting the package maintainer's version, keeping the local version, or manually

resolving the difference. You are also asked to confirm that you want to remove obsolete packages.

7. At the prompt to confirm starting the upgrade, press **Y** to begin.

```
Do you want to start the upgrade?
...

Installing the upgrade can take several hours. Once the download has
finished, the process cannot be canceled.

Continue [yN]  Details [d]
```

8. Press **Y** to proceed with the final reboot.

```
System upgrade is complete.

Restart required

To finish the upgrade, a restart is required.
If you select 'y' the system will be restarted.

Continue [yN]
```

After this reboot, the update process will take several minutes to perform some final installation steps.

Perform the steps in the next section to update to DGX OS 4.4.1.

Obtaining Patch Updates

To update beyond version 4.4.0 via patch updates when they are available, perform the following additional instructions.

1. Update the list of available packages and their versions.

```
$ sudo apt update
```

2. Install the update package.

```
$ sudo apt install -y dgx-bionic-updates-repo
```

3. Update to the latest patch update.

```
$ sudo apt full-upgrade
```

KNOWN ISSUES

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

RESOLVED ISSUES

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

Issues Resolved in DGX OS 4.4.1

- ▶ *[DGX-2]: GPU VMs fail to launch.*

Issues Resolved in DGX OS 4.4.0

- ▶ *[DGX-2]: NVSM Does not Clear md1 Storage Alert After Rebuilding RAID 0*

KNOWN SOFTWARE ISSUES

The following are known issues with the software.

- ▶ [NVSM Reports "Unknown" for Number of logical CPU cores on non-English system](#)
- ▶ [Apparmor Profile May not Work with Some Containers](#)
- ▶ [InfiniBand Bandwidth Drops for KVM Guest VMs](#)

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
```

Resolution

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

Apparmor Profile May not Work with Some Containers

Issue

Apparmor is enabled in this version of the DGX OS Server, with Docker generating a default profile. The default profile may or may not work with your containers.

Workaround

If there is a conflict with your containers, then either

- ▶ Disable Apparmor, or
- ▶ Provide a custom Apparmor profile and include it in the `docker run` command.

InfiniBand Bandwidth Drops for KVM Guest VMs

Issue

The InfiniBand bandwidth when running on multi-GPU guest VMs is lower than when running on bare metal.

Explanation

Currently, performance when using GPUDirect within a guest VM will be lower than when used on a bare-metal system.

KNOWN DGX-2 SYSTEM ISSUES

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

- ▶ [Serial Over LAN Does not Work After Cold Resetting the BMC](#)
- ▶ [NVSM Does not Detect Downgraded GPU PCIe Link](#)
- ▶ [Applications Cannot be Run Immediately Upon Powering on the DGX-2](#)
- ▶ [Hot-plugging of Storage NVMe Drives is not Supported](#)
- ▶ [Some BMC Dashboard Quick Links Appear Erroneously](#)

The following are known issues related to NVIDIA KVM.

- ▶ [Full Health Test on Faulty GPU Fails](#)
- ▶ [Logfile Setup Error When Creating a VM](#)
- ▶ [nvidia-vm vmshow Command Does Not Work for Running VMs](#)

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.
 - c) 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"
```
 - d) Kill the process.


```
kill <PID>
```

 where `<PID>` is the Process ID returned by the previous command.
 - e) 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
```

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.

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.

Hot-plugging of Storage NVMe Drives is not Supported

Issue

Hot-plugging or hot-swapping one of the storage non-volatile memory express (NVMe) drive might result in system instability or incorrect device reporting.

Workaround and Resolution

Turn off the system before removing and replacing any of the storage NVMe drives. This will be resolved in a future software update.

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

Resolution

These quick links will be removed from the menu in a future BMC firmware release.

NVIDIA KVM Issues

Full Health Test on Faulty GPU Fails

Issue

Running `nvidia-vm health-check` fails when run on a faulty GPU within the VM.

Example:

```
$ nvidia-vm health-check --gpu-count 1 --gpu-index 0 --fulltest run
```

The output includes the error

```
"TypeError: object of type 'c_dcgmDiagErrorDetail_t' has no len()".
```

This is due to an issue with DCGM, and does not occur when run on bare metal.

Workaround

Check the GPU health database.

```
$ nvidia-vm health-db show
```

The output should identify the GPUs that are faulty or marked with "non-ok" state. Exclude those GPUs while launching KVM guest VMs.

Logfile Setup Error When Creating a VM

Issue

The following error may appear while creating a VM:

```
..Error setting up logfile: No write access to directory  
/home/$USER/.cache/virt-manager
```

Workaround

To avoid the error, remove the `/home/$USER/.cache/virt-manager` directory after installing KVM packages or before running the first `nvidia-vm` command.

nvidia-vm vmshow Command Does Not Work for Running VMs

Issue

When running `nvidia-vm vmshow`, the information for running guest VMs is reported as “Unknown”.

Resolution

This issue will be resolved in a future release or patch update.

KNOWN DGX-1 SYSTEM ISSUES

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

- ▶ [Forced Reboot Hangs the OS](#)
- ▶ [Script Cannot Recreate RAID Array After Re-inserting a Known Good SSD](#)
- ▶ [Software Power Cap Not Reported Correctly by nvidia-smi](#)
- ▶ [GPUs Cannot be Reset While the System is Running](#)

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 server.

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
```

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.

GPUs Cannot be Reset While the System is Running

Issue

You will not be able to reset the GPUs while the system is running.

Workaround

If an issue occurs which causes the GPUs to hang or if they need to be reset, you must reboot the system.

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.

- ▶ [Serial Over LAN Does not Work After Cold Resetting the BMC](#)
- ▶ [System May Slow Down When Using mpirun](#)
- ▶ [PKCS Errors Appear When the System Boots](#)

Serial Over LAN Does not Work After Cold Resetting the BMC

Issue

Under Ubuntu 18.04, 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.
 - f) 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"
```
 - g) Kill the process.


```
kill <PID>
```

 where <PID> is the Process ID returned by the previous command.
 - h) 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
```

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

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.

Workaround



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](#)).

PKCS Errors Appear When the System Boots

Issue

When the DGX system boots, “PKCS#7 signature not signed with a trusts key” messages appear on the console and system logs.

Explanation

DGX OS Server installs Ubuntu 18.04, which checks all kernel modules for signatures even though Secure Boot is not enabled. Since the NVIDIA drivers are not part of the Ubuntu kernel, the drivers will be flagged with the message when the system boots.

This does not affect the system nor indicate a problem with system software.

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.