



NVIDIA DGX OS SERVER VERSION 4.0.6

DA-08260-406_v02 | August 2019

Release Notes and Update Guide




TABLE OF CONTENTS

NVIDIA DGX OS Server Release Notes for Version 4.0.6	3
Update Advisement.....	4
About Release 4.0	4
Version History	5
Version 4.0.6	5
Version 4.0.5	5
Version 4.0.4	6
Version 4.0.3	6
Version 4.0.2	6
Version 4.0.1	6
DGX OS Server Software Content	7
Base Package Versions	7
KVM Package Versions (DGX-2 only)	7
DGX Server Firmware Version Reference	8
DGX-2 / DGX-2H Firmware.....	8
DGX-1 with Tesla P100, Tesla V100 Firmware	8
Known Issues	9
Resolved Issues.....	9
Known Software Issues.....	9
Known DGX-2 System Issues	11
Known DGX-1 System Issues	20
Known Issues Related to Ubuntu / Linux Kernel	23
Updating to Version 4.0.6.....	25
Update Path Instructions	25
Connecting to the DGX server Console	26
Verifying the DGX Server Connection to the Repositories	27
Updating from 4.0.1 (or later) to 4.0.6.....	28
Update Instructions	28
Recovering from an Interrupted or Failed Update	29
Appendix A. Third Party License Notice	30

NVIDIA DGX OS SERVER RELEASE NOTES FOR VERSION 4.0.6

This document describes version 4.0.6 of the NVIDIA® DGX™ OS Server Release 4.0 software and update package.

DGX OS Server v4.0.6 is provided as an ISO image which is available from NVIDIA Enterprise Support in the event the server needs to be re-imaged.

The software is also provided as an “over-the-network” update, which requires an internet connection and ability to access the NVIDIA public repository.

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

UPDATE ADVISEMENT

► NVIDIA GPU Cloud Containers

In conjunction with DGX OS Server v4.0.6, 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.

ABOUT RELEASE 4.0

The following are the primary features of the DGX OS Server Release 4.0:

- First release to support the NVIDIA DGX-2 System
- Ubuntu 18.04 LTS
- NVIDIA GPU Driver Release 410
 - Supports CUDA 10.0
- Includes the NVIDIA Container Runtime for Docker (nvidia-docker2)

This release replaces the Docker Engine Utility for NVIDIA GPUs (nvidia-docker) with the NVIDIA Container Runtime for Docker (nvidia-docker2).
- New NVSM health monitoring software framework

Replaces `nvsysinfo` and `nvhealth`

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

Version 4.0.6

- ▶ Updated NVSM software (nvhealth updated to 19.01-8)
- ▶ Docker updated to version 18.09.4-ce
- ▶ Updated KVM software
 - CPU Core-affinity feature updated to support 2 and 4 GPU VMs
 - Enable external network connectivity (public IP) for VM using macvtap
 - Enable Host<->VM network connectivity based on isolated (private) network

See [DGX OS Server Software Content](#) for software component list and versions.

Version 4.0.5

- ▶ Added support for NVIDIA DGX-2H systems, including both bare-metal and KVM modes.
- ▶ Updated to Ubuntu 18.04.2 LTS.
- ▶ Updated the NVIDIA GPU driver to version 410.104.

Includes security updates. For details, see [Security Bulletin: NVIDIA GPU Display Driver - February 2019](#).

See also the *Restricting Access to GPU Performance Counters* section of the Linux Readme for driver version 410.104.

- ▶ Enabled Linux crash dump mechanism.
See the [DGX-2 User Guide](#) for instructions on enabling/disabling DGX crash dump capabilities.
- ▶ Docker CE updated to 18.09.2.
Includes security updates. For details, see [DGX Software Security Update: runc Vulnerability – CVE-2019-5736](#) (requires login to NVIDIA Enterprise Support).
- ▶ NVSM components updated to version 19.02 (See [DGX OS Server Software Content](#) for details).
Includes the ability to control MaxQ/MaxP power settings to accommodate power-limited server racks. See the [DGX-2 User Guide](#) for instructions.
- ▶ Updated KVM² Software.

² NVIDIA KVM is available only on DGX-2 systems. It is not available on DGX-1 systems.

- KVM software (dgx-kvm-sw) updated to version 19.02.0
- KVM image (dgx-kvm-image) updated to version 4.0.5
- Enabled CPU core-affinity (vThread) for 8 and 16 GPU VMs

Version 4.0.4

- ▶ Added support for NVIDIA DGX-1 systems.
Sets default Ubuntu IO scheduler from **CFQ** to **deadline**.
- ▶ Updated NVIDIA GPU driver to version 410.79.

Version 4.0.3

- ▶ Updated NVIDIA GPU driver to version 410.72.
- ▶ Updated other software components.
 - DCGM updated to version 1.5.3
 - Docker updated to version 18.06.1-ce
 - NVSM components updated to version 18.10
- ▶ Updated KVM Software.
 - KVM software (dgx-kvm-sw) updated to version 18.10.2.
 - KVM image (dgx-kvm-image) updated to version 4-0-3.
 - Added FS-Cache support for guest VMs.
 - Added multi-queue support for logical drives.
 - Added multi-queue support for virtual networking.
 - Added NUMA tuning.
 - Added CPU tuning (emulatorpin).

Version 4.0.2

- ▶ Updated NVIDIA GPU driver to version 410.48.

Version 4.0.1

- ▶ Initial release for the NVIDIA DGX-2 System
Does not support DGX-1 systems.
- ▶ NVIDIA GPU driver version 410.47
- ▶ See [About Release 4.0](#) for additional features list

DGX OS SERVER SOFTWARE CONTENT

The following tables provides version information for software included in the DGX OS Server ISO image.

Base Package Versions

Component	Version
DGX OS Server	4.0.6
GPU Driver	410.104
NVIDIA Container Runtime for Docker	2.0.3
Ubuntu	18.04.2 LTS
Ubuntu kernel	4.15.0-47 LTS
Docker CE	18.09.4
NVIDIA System Health Monitor (NVSM)	nvsm-cli 19.02.1 nvsm-dshm 19.02.2 nvsm-apis 19.02.5 nvhealth 19.01-8 nvsysinfo 19.01-3
Data Center GPU Management (DCGM)	1.5.9
Mellanox OFED	MLNX 4.4-2.0.1.0

KVM Package Versions (DGX-2 only)

Component	Version
dgx-kvm-sw	19.05.0
dgx-kvm-host-utils	19.05.0
dgx-kvm-host-conf	19.04.0
qemu	2.11+dfsg-1ubuntu7.9~Nvidia.19.02.1
dgx-kvm-image-4.0.6	dgx-kvm-image-4.0.6-190424-6df14a.0.qcow2

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	V01.04.03
SBIOS	V0.22
VBIOS	
DGX-2	88.00.6B.00.01
DGX-2H	88.00.6B.00.08
InfiniBand FW	16.23.1020
PSU FW	2.7
OS Drive	
Samsung (first source)	CXV8601Q
Samsung (second source)	EDA7202Q
Data Drive	
Micron (first source)	101008R0
Samsung (second source)	EDA5202Q

DGX-1 with Tesla P100, Tesla V100 Firmware

Component	Version
BMC	3.30.30
SBIOS	S2W_3A08
VBIOS	
DGX-1 Tesla P100	86.00.41.00.05
DGX-1 Tesla V100 (16 GB)	88.00.18.00.01
DGX-1 Tesla V100 (32 GB)	88.00.80.00.04
PSU	00.03.07
SSD	GXM1103Q

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.

- ▶ *(DGX-2)(KVM) cachefilesd Service Fails After Switching from KVM Mode Back to Bare-metal*

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.

- ▶ [nvidia-vm create Command Does Not Retrieve Guest VM IP Addresses Running 4.0.4 and Older VM Images](#)
- ▶ [NVSM Issues Volume Alerts When Converting Between 8 and 16 SSDs](#)
- ▶ [NVSM Does Not Show Complete BMC Version](#)
- ▶ [RAID 1 Rebuild Progress Bar May Freeze](#)
- ▶ [NVSM Does Not Show Alerts for Degraded RAID 1 Array](#)
- ▶ [NVSM Does Not Show Alerts for Modified EFI Directory on Boot Drive](#)
- ▶ [Logfile Setup Error When Creating a VM](#)
- ▶ [KVM VMs May Become Inaccessible After Host Reboot](#)
- ▶ [Update to Version 4.0.x May Fail Due to Missing DGX GPG Key](#)
- ▶ [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](#)
- ▶ [Storage NVMe Removal May Result in Removal of Different NVMe Drive](#)
- ▶ [BMC SNMP Community String Limitations](#)
- ▶ [Some BMC Dashboard Quick Links Appear Erroneously](#)
- ▶ [Long Boot Time](#)

nvidia-vm create Command Does Not Retrieve Guest VM IP Addresses Running 4.0.4 and Older VM Images

Issue

After issuing the `nvidia-vm create` command, the IP address portion of the output hangs for a while before timing out. This occurs when creating VMs based on KVM guest image versions 4.0.5 and earlier.

Workaround

KVM guest images version 4.0.4 and earlier do not incorporate the `qemu-guest-agent` software and consequently will not respond to the `virsh domifaddr` request issued by the `nvidia-vm` tool. The guest VM is still created, though.

If you create a guest VM based on one of these images, press [Ctrl+C] to escape from the long pause attempting to get the IP address, then issue `virsh console <vm>` to get the IP address.

NVSM Issues Volume Alerts When Converting Between 8 and 16 SSDs

Issue

When the drive configuration changes, such as after adding or removing 8 SSDs to or from an 8 SSD or 16 SSD system, respectively, NVSM generates critical volume alerts.

Explanation

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

NVSM Does Not Show Complete BMC Version

Issue

The BMC version that appears from an `nvsm` command is not the full version.

Example:

```
$ sudo nvsm show health
<snip>
BMC Firmware Revision [1.04].....
<snip>
```

The actual BMC version in this case is 1.04.03.

Explanation

The BMC version returned by `nvsm` follows the format defined in the [IPMI specification](#), which includes a major and minor component version. Sub-versions beyond that are not included.

To see the full BMC version, use the BMC dashboard.

RAID 1 Rebuild Progress Bar May Freeze

Issue

When using NVSM CLI to rebuild the RAID 1 array, the progress bar may freeze - for example, at 0%.

Workaround

The rebuild is occurring even though the progress bar does not indicate it. You can confirm this with the following command:

```
$ 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.

NVSM Does Not Show Alerts for Degraded RAID 1 Array

Issue

If one of the RAID 1 OS drives becomes corrupted or fails, the RAID goes in degraded mode but NVSM does not show an alert. The `nvsm show` command reports the RAID array as healthy.

Workaround

NVSM raises the correct alerts after a system reboot.

NVSM Does Not Show Alerts for Modified EFI Directory on Boot Drive

Issue

If the EFI directory of one of the RAID 1 OS drives is inadvertently modified, the system will boot off the good drive but NVSM does not show an alert. The `nvsm show` command reports the drive as healthy.

Explanation and Recovery

The EFI directory is used to hold the UEFI boot file. The ESP monitor will not be aware of changes to the directory name and will not generate an alert.

This will be resolved in the next DGX OS release.

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.

KVM VMs May Become Inaccessible After Host Reboot

Issue

Rebooting the KVM host while VMs are running can corrupt the `/raid` volume used by the VM. While the KVM host attempts to shut down the VMs during a reboot, it may fail to shut down all of them, resulting in an inaccessible VM.

Workaround

Manually shut down all active VMs before rebooting/shutting down the KVM host.

Update to Version 4.0.x May Fail Due to Missing DGX GPG Key

Issue

If an over-the-network update was previously performed on the DGX-2 server, for example, to update from version 4.0.1 to 4.0.3, then updating to a later version may fail due to a missing DGX GPG key with the following messages:

```
...
Err:1 http://international.download.nvidia.com/dgx/repos/bionic bionic
InRelease
  The following signatures couldn't be verified because the public key
is not available: NO_PUBKEY ECDB3674629C85F2
...
```

```
W: GPG error: http://international.download.nvidia.com/dgx/repos/bionic
bionic InRelease: The following signatures couldn't be verified because
the public key is not available: NO_PUBKEY ECDB3674629C85F2
E: The repository
'http://international.download.nvidia.com/dgx/repos/bionic bionic
InRelease' is not signed.
```

Explanation and Recovery

A recent change in the Ubuntu keyring update causes the DGX GPG key to be overwritten during the DGX OS update. Perform the following before attempting to update again.

1. Download the dgx-repo-keys package.

```
wget
https://international.download.nvidia.com/dgx/repos/bionic/pool/multive
rse/d/dgx-repo-keys/dgx-repo-keys_2.0_amd64.deb
```

2. Install the dgx-repo-keys package.

```
sudo dpkg -i ./dgx-repo-keys_2.0_amd64.deb
```

3. Confirm that the DGX GPG key has been installed:

```
apt-key list

/etc/apt/trusted.gpg.d/dgx-key.gpg
-----
pub   rsa4096 2016-06-07 [SC]
      7835 264B 12C3 DA96 5D0E  1078 ECDB 3674 629C 85F2
uid           [ unknown] dgx-cosmos-support <dgx-cosmos-
support@nvidia.com>
sub   rsa4096 2016-06-07 [E]
```

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.

Storage NVMe Removal May Result in Removal of Different NVMe Drive

Issue

When attempting to remove access to an NVMe drive using the following command,

```
echo 1 > /sys/class/nvme/nvmeX/device/remove
```

where X specifies which NVMe drive to remove, other NVMe drives may get removed.

Workaround and Resolution

This is the result of an issue in the NVMe driver and will be resolved in a future software update. To work around, shut down the system before removing the NVMe drive.

BMC SNMP Community String Limitations

Issue

The DGX-2 BMC has the following SNMP Community String limitations:

- No support for SNMPv3
- No SNMP configuration controls in the BMC dashboard
- No support for setting RO and RW permissions from the command line `ipmitool`.

Resolution

This will be resolved in a future BMC firmware release.

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.

Long Boot Time

Issue

The DGX-2 System boot process can take over five minutes from the start of the system boot until the login prompt.

Explanation and Workaround

Part of the boot process involves setting up of the PXE boot feature. If you do not need PXE boot functionality, you can shorten the boot time by disabling PXE boot for all installed Mellanox cards as follows.



NOTE: This disables the ability to perform PXE boot from the ConnectX 5 network adapters.

1. Start the Mellanox Software Tools driver.

```
$ sudo mst start
```

2. Determine the Mellanox cards that have PXE boot enabled.

```
$ sudo mlxconfig query|grep -i -e "device\|EXP_ROM_UEFI_x86_ENABLE"
```

Look for any Mellanox devices that show EXP_ROM_UEFI_x86_ENABLE as True as highlighted in the following example.

```
Device #5:
Device type:    ConnectX5
Device:         /dev/mst/mt4119_pciconf4
                EXP_ROM_UEFI_x86_ENABLE                True(1)
```

In this example, PXE boot needs to be disabled for /dev/mst/mt4119_pciconf4.

3. Disable PXE boot for each required card.

Example:

```
$ sudo mlxconfig -d /dev/mst/mt4119_pciconf4 set
EXP_ROM_UEFI_x86_ENABLE=0
```

4. Verify that all ports show EXP_ROM_UEFI_x86_ENABLE as False (0).

```
~$ sudo mlxconfig query|grep -i -e "device\|EXP_ROM_UEFI_x86_ENABLE"
Device #1:
Device type:    ConnectX5
Device:         /dev/mst/mt4119_pciconf8
                EXP_ROM_UEFI_x86_ENABLE                False(0)
Device #2:
Device type:    ConnectX5
Device:         /dev/mst/mt4119_pciconf7
                EXP_ROM_UEFI_x86_ENABLE                False(0)
Device #3:
Device type:    ConnectX5
Device:         /dev/mst/mt4119_pciconf6
                EXP_ROM_UEFI_x86_ENABLE                False(0)
Device #4:
Device type:    ConnectX5
```

```

Device:          /dev/mst/mt4119_pciconf5
                EXP_ROM_UEFI_x86_ENABLE                False(0)
Device #5:
Device type:     ConnectX5
Device:          /dev/mst/mt4119_pciconf4
                EXP_ROM_UEFI_x86_ENABLE                False(0)
Device #6:
Device type:     ConnectX5
Device:          /dev/mst/mt4119_pciconf3
                EXP_ROM_UEFI_x86_ENABLE                False(0)
Device #7:
Device type:     ConnectX5
Device:          /dev/mst/mt4119_pciconf2
                EXP_ROM_UEFI_x86_ENABLE                False(0)
Device #8:
Device type:     ConnectX5
Device:          /dev/mst/mt4119_pciconf1
                EXP_ROM_UEFI_x86_ENABLE                False(0)
Device #9:
Device type:     ConnectX5
Device:          /dev/mst/mt4119_pciconf0
                EXP_ROM_UEFI_x86_ENABLE                False(0)

```

KNOWN DGX-1 SYSTEM ISSUES

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

- ▶ [NVSM APIs Return Duplicate Serial Numbers for PSUs](#)
- ▶ [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](#)

NVSM APIs Return Duplicate Serial Numbers for PSUs

Issue

Utilities or scripts that use NVSM APIs will display duplicate PSU serial numbers instead of unique ones for each PSU.

Cause

This is the result of a race condition in the BMC when getting PSU information.

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. This will be fixed in a future release.

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.

- ▶ [Long Boot Time When KVM Guest is Configured with Private Network](#)
- ▶ [System May Slow Down When Using `mpirun`](#)
- ▶ [PKCS Errors Appear When the System Boots](#)

Long Boot Time When KVM Guest is Configured with Private Network

Issue

When the KVM guest VM is configured for an isolated (`--privateIP`) network, booting to the VM takes a long time due to a `system-networkd-wait-online.service` failure.

Explanation

This is due to a networking issue in the Linux `systemd` service. This is resolved in a Bionic release which will be captured in the next release of the DGX OS software.

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.

UPDATING TO VERSION 4.0.6

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 QA'd versions of the entire DGX software stack, including the drivers.

Perform the updates using commands on the DGX server console.

UPDATE PATH INSTRUCTIONS

The update instructions depend on the software version currently installed on the DGX server. Follow the instructions corresponding to your current DGX OS server software.

Updating from Release 2.x

As of 8/15/2019 with the release of the 4.1.1 patch update, you cannot update to Release 4.0 from Release 2.x.

Updating from Release 3.1

As of 8/15/2019 with the release of the 4.1.1 patch update, you cannot update to Release 4.0 from Release 3.1.

Updating from Version 4.0.1 or later

Follow instructions at [Updating from 4.0.1 \(or Later\) to 4.0.6](#).

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 version 4.0.4 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.

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.

UPDATING FROM 4.0.1 (OR LATER) TO 4.0.6

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. 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 4.0.6.

```
$ 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.

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 \(or Later\) to 4.0.6](#) update instructions

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

© 2018 NVIDIA Corporation. All rights reserved.