NVIDIA DGX OS Server Release 4.99.11

Release Notes and Update Guide

DA-08260-49911_v02 | August 2020

Table of Contents

NVIDIA DGX OS Server 4.99.11 Release Notes

This document describes the NVIDIA® DGX™ OS Server Release 4.99 software package.

About Release 4.99

The following are the primary new features of DGX OS Server Release 4.99 since Release 4.4:

- Support only for the NVIDIA DGX A100 system
- Updated NVIDIA GPU driver to Release 450.

Delivery and Update Mechanisms

Initial 4.99 Release

DGX OS Server Release 4.99, version 4.99.11, is provided as an "over-the-network" update, which requires an internet connection and ability to access the NVIDIA public repositories. It is also provided as an ISO image for re-imaging the system.

Refer to the *DGX-A100 User Guide* (https://docs.nvidia.com/dgx/dgxa100-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

- IMPORTANT: 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. Failure to do so will result in the GPUs not getting recognized.
- NGC Containers

  In conjunction with DGX OS Server v4.99, customers should update their NGC containers to the latest container release[1].
- 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 4.99. See DGX OS Server Software Content for the software component list and versions.

Version 4.99.11

Changes since Version 4.99.10:

- CUDA 11 toolkit (if previously installed) updated to CUDA 11 Update 1
- Updated NVSM to 20.05.17
  - Supports the latest SBIOS and Broadcom FW.
  - Implemented bug fixes
- Upgraded the `nvidia-peer-memory` module to 1.1-0.

Version 4.99.10

Changes since Version 4.99.9:

- Updated NVIDIA GPU driver to 450.51.06
- Updated DCGM to 2.0.10

Version 4.99.9

Changes since Version 4.99.8:

- Updated NVIDIA GPU driver to 450.51.05
- Updated NVSM software to version 20.05.9
- Updated DCGM to version 2.0.8

Version 4.99.8

Changes since Version 4.99.6:

- Updated NVIDIA GPU driver to 450.36.06
- Updated NVSM software to version 20.05.3
- Updated Docker-CE to 19.03.8
- Updated the NVIDIA Container Toolkit stack to enable use of MIG with containers

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 4.99.11 Update

The following table shows the version information for software included in the DGX OS Server update 4.99.10 for DGX A100.

Component	Version
GPU Driver	450.51.06
NVIDIA Container Toolkit	libnvidia-container1 1.1.0-1 libnvidia-container-tools 1.1.0-1 nvidia-container-runtime 3.1.4-1 nvidia-container-toolkit 1.0.6-1 nvidia-docker2 2.2.2-1
Ubuntu	18.04.4 LTS
Ubuntu kernel	5.4.0-42[2]
Docker Engine	19.03.8
NVIDIA System Health Monitor (NVSM)	**NVSM 20.05.17**
Data Center GPU Management (DCGM)	2.0.10
Mellanox OFED	MLNX 5.0-2.1.8.0

Package Versions After Getting the 4.99.10 Update

The following table shows the version information for software included in the DGX OS Server update 4.99.10 for DGX A100.

Component	Version
GPU Driver	**450.51.06**
NVIDIA Container Toolkit	libnvidia-container1 1.1.0-1 libnvidia-container-tools 1.1.0-1 nvidia-container-runtime 3.1.4-1 nvidia-container-toolkit 1.0.6-1 nvidia-docker2 2.2.2-1
Ubuntu	18.04.4 LTS

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

Component	Version
Ubuntu kernel	5.3.0-59[3]
Docker Engine	19.03.8
NVIDIA System Health Monitor (NVSM)	NVSM 20.05.9
Data Center GPU Management (DCGM)	**2.0.10**
Mellanox OFED	MLNX 5.0-2.1.8.0

Package Versions After Getting the 4.99.9 Update

The following table shows the version information for software included in the DGX OS Server update 4.99.9 for DGX A100.

Component	Version
GPU Driver	**450.51.05**
NVIDIA Container Toolkit	libnvidia-container1 1.1.0-1 libnvidia-container-tools 1.1.0-1 nvidia-container-runtime 3.1.4-1 nvidia-container-toolkit 1.0.6-1 nvidia-docker2 2.2.2-1
Ubuntu	18.04.4 LTS
Ubuntu kernel	5.3.0-59[4]
Docker Engine	19.03.8
NVIDIA System Health Monitor (NVSM)	**NVSM 20.05.9**
Data Center GPU Management (DCGM)	**2.0.8**
Mellanox OFED	MLNX 5.0-2.1.8.0

Package Versions After Getting the 4.99.8 Update

The following table shows the version information for software included in the DGX OS Server update 4.99.8 for DGX A100.

Component	Version
GPU Driver	**450.36.06**
NVIDIA Container Toolkit	libnvidia-container1 1.1.0-1 libnvidia-container-tools 1.1.0-1 nvidia-container-runtime 3.1.4-1 nvidia-container-toolkit 1.0.6-1

[3] The resulting kernel version on your system may be a later version depending on when the update is performed.

[4] The resulting kernel version on your system may be a later version depending on when the update is performed.

Component	Version
	nvidia-docker2 2.2.2-1
Ubuntu	18.04.4 LTS
Ubuntu kernel	5.3.0-53[5]
Docker Engine	19.03.8
NVIDIA System Health Monitor (NVSM)	**NVSM 20.05.3**
Data Center GPU Management (DCGM)	2.0.4
Mellanox OFED	MLNX 5.0-2.1.8.0

DGX Server Firmware Version Reference

See the DGX A100 System Firmware Update Container Version 20.05.12.3 release notes for the firmware versions available at the time of this DGX OS release.

[5] The resulting kernel version on your system may be a later version depending on when the update is performed.

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-Server 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 DGX A100 user guide 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-international
http://international.download.nvidia.com/dgx/repos/bionic/dists/bionic/
Release
$ wget -O f6-international
http://international.download.nvidia.com/dgx/repos/bionic/dists/bionic-
4.99/Release
```

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

Performing the Updates

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.

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 the latest version.

```
$ 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 A100, you will not be able to perform the over-the-network update. You will need to re-image the DGX A100 with the latest image (see the DGX A100 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 update instructions

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 4.99.11.

- NVSM raises alerts against the PCIe links to the NVSwitches, stating that the link width is degraded due to the link width being 2 when it is expected to be 4. (Reported in 4.99.10)
- nvsm show health may report "Unhealthy" for some NVMe and IB checks.
- NVSM raises multiple PCIe alerts after the DGX OS is installed.
- NVSM reports "An unexpected Network port is missing" when the optional dual port network card is not installed.
- NVSM reports the installed IB controller is unhealthy if the default dual-port network card is configured as InfiniBand and the optional dual-port card is not installed,
- On systems with encrypted rootfs, the RAID 1 array cannot be rebuilt using NVSM.(Originally reported in 4.99.9)

  The workaround was to manually rebuild the RAID 1 array by adding the partition of the restored drive to the RAID array.

  ```
  $ sudo mdadm --manage /dev/md1 --add /dev/<device-name>
  ```
- After running the /usr/sbin/update-pciids script, NVSM will report the PCIe switches as unhealthy in response to nvsm show health. (Reported in 4.99.10)

  The workaround was to restore the pci.ids file from /usr/share/misc/pci.ids.old.

The following issues that existed in version 4.99.9 are fixed in version 4.99.10.

- Out-of-memory error/memory leak occurs when memory compression is enabled.

  Refer to the GPU driver notes for details.

- GPU row-remapper error: Incorrect address translation causes erroneous row_remapping data to be stored in InfoROM.

  Refer to the GPU driver notes for details.
- nv_peer_mem module is not installed or updated when performing apt full-upgrade.

The following issues that existed in version 4.99.8 are fixed in version 4.99.9.

- If the ESP becomes corrupted on one of the NVMe M.2 drives, you cannot rebuild it using the `nvsm start rebuild` command.
- Up to 25% GPU utilization is reported by `nvml/nvidia-smi` even when the GPU is idle.

Known Software Issues

The following are currently open issues reported in version 4.99.11.

- With Eight NVMe drives installed, nvsm-plugin-pcie generates "ERROR Device not found in mapping table" error.

The following are currently open issues originally reported in version 4.99.10.

- `nvsm show alerts` Reports nvswitch PCIe Link Width Degraded
- cuMemFree CUDA API Performance Regression
- nvsm stress-test Does not Stress the System if MIG is Enabled

The following are currently open issues originally reported in version 4.99.9.

- 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
- Syslog Contains Numerous "SM LID is 0, maybe no SM is running" Error Messages
- BMC is not Detectable After Restoring BMC to Default
- NVSM Enumerates NVSwitches as 8-13 Instead of 0-5
- The System May Fail to Boot if one of the M.2 drives is Corrupted
- NVSwitch non-fatal Error 10003 Occurs When Fabric Manager Service Started

The following are currently open issues originally reported in version 4.99.8.

- Mellanox Firmware is not Updated Automatically
- NVSM Reports "Unknown" for Number of logical CPU cores on non-English system

With Eight NVMe drives installed, nvsm-plugin-pcie generates "ERROR Device not found in mapping table" error.

Issue

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*).

Explanation

This is an issue with the NVSM plugin PCIe service, which is not detecting the additional four drives. nvsm show health and nvsm dump health function normally and no false alerts are raised in connection with this issue.

`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.

Broadcom firmware fixes this issue and will be released in revision ef_3, to be provided at a later date.

cuMemFree CUDA API Performance Regression

Issue

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 stress-test Does not Stress the System if MIG is Enabled

Issue

On a system with MIG enabled, the GPU load and system power draw do not increase when running `nvsm stress-test`.

Explanation

MIG must be disabled before running `nvsm stress-test`.

AMD Crypto Co-processor is not Supported

Issue

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

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

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.

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
```

BMC is not Detectable After Restoring BMC to Default

Issue

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

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.

The System May Fail to Boot if one of the M.2 drives is Corrupted

Issue

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
```

NVSwitch non-fatal Error 10003 Occurs When Fabric Manager Service Started

Issue

When starting the Fabric Manager service, the following error is reported:

```
detected NVSwitch non-fatal error 10003 on NVSwitch pci
```

Explanation

This error is not fatal and no functionality is affected. This issue will be resolved in a future driver release.

Mellanox Firmware is not Updated Automatically

Issue

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.

If you intend to run connect the ConnectX-6 cards to an FDR network, the firmware needs to be updated to version 20.27.6008.

Explanation and Workaround

To work around, update the Mellanox firmware manually as follows.

1. Download the firmware for the Mellanox ConnectX-6 cards from the Mellanox website: https://www.mellanox.com/support/firmware/connectx6ib.
   Download firmware for the following products (OPN):
   - MCX653105A-HDAT (PSID MT_0000000223)
   - MCX653106A-HDAT (PSID MT_0000000225)
2. Copy and extract the files into a directory on the DGX A100.

```
:~/temp-mlnx-download$ ls
fw-ConnectX6-rel-20_27_2008-MCX653105A-HDA_Ax-UEFI-14.20.22-
FlexBoot-3.5.901.bin
fw-ConnectX6-rel-20_27_2008-MCX653106A-HDA_Ax-UEFI-14.20.22-
FlexBoot-3.5.901.bin
```

3. Run the "`sudo mlxfwmanager -u -y`" command from that directory.
   The "`-u`" option indicates that you want to perform an upgrade, and the "`-y`" option indicates you don't want to be prompted.

```
:~/temp-mlnx-download$ sudo mlxfwmanager -u -y

Querying Mellanox devices firmware ...

Device #1:
----------

Device Type:    ConnectX6
Part Number:    MCX653106A-HDA_Ax
Description:    ConnectX-6 VPI adapter card; HDR IB (200Gb/s) and
200GbE; dual-port QSFP56; PCIe4.0 x16; tall bracket; ROHS R6
PSID:           MT_0000000225
PCI Device Name: /dev/mst/mt4123_pciconf9
Base MAC:       1c34da4d72f6
Versions:       Current     Available
    FW          20.27.1016  20.27.2008
    PXE         3.5.0901    3.5.0901
```

```
    UEFI        14.20.0019  14.20.0022
Status:   Update required
...
...
...
---------
Found 10 device(s) requiring firmware update...
Device #1: Updating FW ...
Initializing image partition - OK
Writing Boot image component - 26%
...
...
```

If all ten ConnectX-6 cards need a firmware update, the total update time will be approximately 30 minutes.

4. Restart the nvidia-mnlx-config service.

```
$ sudo systemctl restart nvidia-mlnx-config
```

5. Reboot once the firmware update completes to load the new firmware.

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 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

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.

Explanation and 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
  ```

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 list known limitations and other issues that will not be fixed.

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.

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 Unites States and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright