



NVIDIA DGX-1 SOFTWARE VERSION 2.0.4

DA-08260-204_v02 | April 2017

Release Notes and Upgrade Guide



TABLE OF CONTENTS

NVIDIA DGX-1 Software Release Notes for Version 2.0.4	3
Highlights	4
Known Issues.....	4
DGX-1 Software Version 2.0.4 Upgrade Guide	5
Connecting to the DGX-1 Console.....	5
Performing the Over-the-Network Update	6
Recovering from an Interrupted Update	10
Appendix A: MetaPackage 2.0.5 Contents	12
Appendix B. Third Party License Notice	16

NVIDIA DGX-1 SOFTWARE RELEASE NOTES FOR VERSION 2.0.4

This document describes the version 2.0.4 release of the NVIDIA® DGX-1™ software and the update package for performing the over-the-network update.

Software Update Advisement

All DGX-1 systems should be updated with the update package. See the update instructions below. “Over-the-network” updates require an internet connection and ability to access the NVIDIA public repository using the *apt-get* command.

If you prefer to re-image the system using an ISO image, or if your DGX-1 is not connected to a network with internet access, refer to the DGX-1 User Guide for re-imaging instructions.

Container Update Advisement

In conjunction with the DGX-1 2.0.4 update, customers should update their NVIDIA Docker containers to the latest container release¹.

¹ Containers access instructions are available through the NVIDIA Enterprise Support site.

HIGHLIGHTS

- ▶ Includes the following software:
 - Ubuntu 14.04 LTS
 - Includes Xenial HWE update 4.4.0-45, which fixes vulnerabilities reported in CVE-2016-5196.
 - NVIDIA GPU Driver 375.20
- ▶ Updated the installation process based on Canonical OEM configuration.
- ▶ Initial setup defaults to Base OS mode².
- ▶ Revised over-the-network software update mechanism.

KNOWN ISSUES


- ▶ The installed *ipmitool* cannot manipulate FRU data.

If you need to access FRU information, go to the NVIDIA Enterprise Support site (<https://nvidia-esp.custhelp.com/>) and refer to the Knowledge Base article *DGX-1 Asset Tag Programming* for instructions on using the *ipmi fru* command from the *freeipmi* package.

² Cloud Management mode will be available at a later date. At that time you will be able to switch to Cloud Management mode using an over-the-network update process.


DGX-1 SOFTWARE VERSION 2.0.4 UPGRADE GUIDE

These instructions explain how to update the DGX-1 software through an internet connection to the NVIDIA public repository. Perform the updates using commands on the DGX-1 console.

 **NOTE:** This update uses NVIDIA DGX-1 update script and metapackage version 2.0.5, which corresponds to ISO image version 2.0.4. If your DGX-1 is already installed with ISO image version 2.0.4, you do not need to perform this update.

CONNECTING TO THE DGX-1 CONSOLE

Connect to the DGX-1 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-1 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.

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-1

Remote Connection through the BMC

This method requires that you have the BMC login credentials.

1. Make sure you have connected the IPMI port on the DGX-1 to your LAN.
2. Open a Java-enabled browser within your LAN and go to `http://<IPMI IP Address>/`. Make sure popups are allowed for the BMC address.
3. Log in.
4. From the top menu, click **Remote Control** and then select **Console Redirection**.
5. Click **Java Console** to open the popup window.
6. If necessary, power on the DGX-1 using the power button icon on the upper right corner of the BMC window.

PERFORMING THE OVER-THE-NETWORK UPDATE



NOTE: The DGX-1 will automatically reboot during the update process in order to complete the upgrade of the Linux kernel, CUDA driver, and Mellanox firmware.

See the section [Recovering from an Interrupted Update](#) in the event the script is interrupted for any reason.

1. Verify that networking is enabled so that you can access the DGX-1 public repository.

For example:

```
$ ping www.google.com
```

You may require alternate methods of verifying, if your network is not configured for ping.

2. Verify that the DGX-1 public repository is configured.

```
$ dpkg -l dgx1-repo-ubuntu1404
```

Expected output:

```
ii dgx1-repo-ubuntu1404 1.1-1 amd64 dgx1 repository configuration files
```

If this output appears, then proceed directly to step 3.

If the `dgx1-repo-ubuntu1404` package is not present, then install the package manually as follows:

a) Download the debian repo using *wget*.

```
$ sudo wget
http://international.download.nvidia.com/dgx1/repos/pool/multiverse/
d/dgx1-repo/dgx1-repo-ubuntu1404_1.1-1_amd64.deb -O /tmp/dgx1-repo-
ubuntu1404_1.1-1_amd64.deb
```

b) Install the package.

```
$ sudo dpkg -i /tmp/dgx1-repo-ubuntu1404_1.1-1_amd64.deb
```

c) Verify that the package installed.

```
$ dpkg -l dgx1-repo-ubuntu1404
```

Expected output:

```
ii dgx1-repo-ubuntu1404 1.1-1 amd64 dgx1 repository configuration
files
```

3. Get the new package list.

```
$ sudo apt-get update
```

4. Confirm update script 2.0.5-6 is available.

```
$ apt-cache policy dgx1-ota-update
```

Expected output:

```
dgx1-ota-update:
  Installed: (whatever version was previously installed, or none)
  Candidate: 2.0.5-6
```

5. Install the update script.

```
$ sudo apt-get -y --force-yes install dgx1-ota-update
```

6. Run the update script.

The script automatically performs several steps:

- Removes packages (nvidia-361, ar-mgr, isert-dkms, etc).
- Installs the update meta package (dgx1-ota-update-meta).
- Reboots the DGX-1 upon completion.

To run the script, enter the following:

```
$ sudo /usr/bin/dgx1-ota-update.sh
```

- While the script is running, you may be presented with a configuration choice such as the following:

```
Setting up dgx-limits (1.0-1) ...

Configuration file '/etc/security/limits.d/dgx1- limits.conf'

==> File on system created by you or by a script.
==> File also in package provided by package maintainer.
   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.

*** dgx1-limits.conf (Y/I/N/O/D/Z) [default=N] ?
```

Press **Y** to all such messages in order to “install the package maintainer’s version”.

- You may see the following error message during DKMX configuration for the `nv_peer_mem` module while running the script:

```
DKMS: install completed.
modprobe: ERROR: could not insert 'nv_peer_mem': Exec format error
```

Step 11 of this procedure provides instructions for remedying this error.

The update script will automatically reboot upon completion.

7. Wait for the system to come back up.
8. Confirm that the Linux kernel version is 4.4.0-45.

```
$ uname -a
```

Expected output:

```
Linux dgx-1 4.4.0-45-generic #66~14.04.1-Ubuntu SMP Wed Oct
19 15:05:38 UTC 2016 x86_64 x86_64 x86_64 GNU/Linux
```

9. Confirm CUDA driver version is 375.20.

```
$ nvidia-smi
```

Expected output (first line)

```
NVIDIA-SMI 375.20      Driver Version: 375.20
```

10. If your network is configured for DHCP, then make sure that dynamic DNS updates are enabled.

Check whether `/etc/resolv.conf` is a link to `/run/resolvconf/resolv.conf`. and skip to step 11 if the expected output appears.

```
$ ls -l /etc/resolv.conf
```


Expected output:

```
lrwxrwxrwx 1 root root 29 Dec  1 21:19 /etc/resolv.conf ->
../run/resolvconf/resolv.conf
```

If this does not appear, then enable dynamic DNS updates as follows:

- a) Launch the *Resolvconf Reconfigure* package.

```
$ sudo dpkg-reconfigure resolvconf
```

The *Configuring resolvconf* screen appears.

- b) Select **<Yes>** when asked whether to prepare `/etc/resolv.conf` for dynamic updates.
 c) Select **<No>** when asked whether to append original file to dynamic file.
 d) Select **<OK>** at the *Reboot recommended* screen.

You do not need to reboot.

You are returned to the command line.

- e) Bring down the interface, where `<network interface>` is `em1` or `em2`, whichever you have set up as your primary network interface:

```
$ sudo ifdown <network interface>
```

Expected output:

```
ifdown: interface <network interface> not configured
```

- f) Bring up the interface, where `<network interface>` is `em1` or `em2`, whichever you have set up as your primary network interface:

```
$ sudo ifup <network interface>
```

Expected output (last line):

```
...
bound to <IP address> -- renewal in ...
```

- g) Repeat step 10 to confirm that `/etc/resolv.conf` is a link to `/run/resolvconf/resolv.conf`.

11. Make sure the `nvidia-peer-memory` module is installed.

The `nvidia-peer-memory` module should be installed already if the previously running kernel was 4.4.

```
$ lsmod | grep nv_peer_mem
```

Expected output:

```
nv_peer_mem      16384  0
nvidia           11911168  30 nv_peer_mem,nvidia_modeset,nvidia_uvm
```

```
ib_core          143360  13
rdma_cm,ib_cm,ib_sa,iw_cm,nv_peer_mem,mlx4_ib,mlx5_ib,ib_mad,ib_ucm
,ib_umad,ib_uverbs,rdma_ucm,ib_ipoib
```

- If the expected output appears, then you have completed the DGX-1 setup.
 - If there is no output, then build and install the nvidia-peer-memory module as follows:
- a) Get and install the module.

```
$ sudo apt-get install --reinstall mlnx-ofed-kernel-dkms nvidia-
peer-memory-dkms
```

Expected output:

```
DKMS: install completed.
Processing triggers for initramfs-tools (0.103ubuntu4.2) ...
update-initramfs: Generating /boot/initrd.img-4.4.0-45-generic
```

- b) Add the module to the Linux kernel.

```
$ sudo modprobe nv_peer_mem
```

There is no expected output for this command.

- c) Repeat step 11 to verify that the nvidia-peer-memory module is installed.

RECOVERING FROM AN INTERRUPTED UPDATE

If the script is interrupted during the update, such as from a loss of power or loss of network connection, you can attempt to recover as follows:

1. 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-1, you will not be able to perform the over-the-network update. You will need to re-image the DGX-1 with the latest image to update the software. Refer to the DGX-1 User Guide for instructions.
 - If you are successfully returned to the Linux command line, continue the following steps.
2. Reset the script.

```
$ sudo dpkg --configure -a
```

3. Start the script again.

```
$ sudo /usr/bin/dgx1-ota-update.sh
```

- While the script is running, you may be presented with a configuration choice such as the following:

```
Setting up dgx-limits (1.0-1) ...

Configuration file '/etc/security/limits.d/dgx1- limits.conf'
==> File on system created by you or by a script.
==> File also in package provided by package maintainer.
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.

*** dgx1-limits.conf (Y/I/N/O/D/Z) [default=N] ?
```

Press **Y** to all such messages in order to “install the package maintainer’s version”.

- You may see the following error message during DKMX configuration for the `nv_peer_mem` module while running the script:

```
DKMS: install completed.
modprobe: ERROR: could not insert 'nv_peer_mem': Exec format error
```

Step 11 of this procedure provides instructions for remedying this error.

The update script will automatically reboot upon completion.

4. Continue following the instructions from step 7 in the [Performing the Update](#) instructions.

APPENDIX A: METAPACKAGE 2.0.5 CONTENTS

CUDA

cuda-drivers(=375.20-1)

cuda-drivers-diagnostic(=375.20-1)

DCGM

datacenter-gpu-manager(=375.20-1)

DGX1 Config

dgx1-configurations(=1.0-1)

dgx1-repo-ubuntu1404(=1.1-1)

Kernel

linux-headers-4.4.0-45-generic(=4.4.0-45.66~14.04.1)

linux-headers-4.4.0-45(=4.4.0-45.66~14.04.1)

linux-image-4.4.0-45-generic(=4.4.0-45.66~14.04.1)

linux-image-extra-4.4.0-45-generic(=4.4.0-45.66~14.04.1)

linux-lts-xenial-tools-4.4.0-45(=4.4.0-45.66~14.04.1)

linux-tools-4.4.0-45-generic(=4.4.0-45.66~14.04.1)

linux-signed-image-4.4.0-45-generic(=4.4.0-45.66~14.04.1)

thermald(=1.4.3-5~14.04.4)

Mellanox

dapl2-utils(=2.1.9mlnx-OFED.3.3.0.0.6.34100)
 ibacm(=1.2.1mlnx1-OFED.3.4.0.1.5.34100)
 ibacm-dev(=1.2.1mlnx1-OFED.3.4.0.1.5.34100)
 ibsim(=0.6-0.34100)
 ibsim-doc(=0.6-0.34100)
 ibutils(=1.5.7.1-0.34100)
 ibverbs-utils(=1.2.1mlnx1-OFED.3.4.0.1.4.34100)
 infiniband-diags(=1.6.6.MLNX20160814.999c7b2-0.34100)
 iser-dkms(=1.8.1-OFED.3.4.1.0.0.1.g2ed8a21)
 kernel-mft-dkms(=4.5.0-31)
 knem(=1.1.2.90mlnx1-OFED.3.3.1.5.5.1.ga659703)
 knem-dkms(=1.1.2.90mlnx1-OFED.3.3.1.5.5.1.ga659703)
 libdapl2(=2.1.9mlnx-OFED.3.3.0.0.6.34100)
 libdapl-dev(=2.1.9mlnx-OFED.3.3.0.0.6.34100)
 libibcm1(=1.0.5mlnx2-OFED.3.4.0.0.4.34100)
 libibcm-dev(=1.0.5mlnx2-OFED.3.4.0.0.4.34100)
 libibdm1(=1.5.7.1-0.34100)
 libibmad(=1.3.12.MLNX20160814.4f078cc-0.34100)
 libibmad-devel(=1.3.12.MLNX20160814.4f078cc-0.34100)
 libibmad-static(=1.3.12.MLNX20160814.4f078cc-0.34100)
 libibumad(=1.3.10.2.MLNX20150406.966500d-0.34100)
 libibumad-devel(=1.3.10.2.MLNX20150406.966500d-0.34100)
 libibumad-static(=1.3.10.2.MLNX20150406.966500d-0.34100)
 libibverbs1(=1.2.1mlnx1-OFED.3.4.0.1.4.34100)
 libibverbs1-dbg(=1.2.1mlnx1-OFED.3.4.0.1.4.34100)
 libibverbs-dev(=1.2.1mlnx1-OFED.3.4.0.1.4.34100)
 libmlx4-1(=1.2.1mlnx1-OFED.3.4.0.0.4.34100)
 libmlx4-1-dbg(=1.2.1mlnx1-OFED.3.4.0.0.4.34100)
 libmlx4-dev(=1.2.1mlnx1-OFED.3.4.0.0.4.34100)
 libmlx5-1(=1.2.1mlnx1-OFED.3.4.1.0.0.34100)

libmlx5-1-dbg(=1.2.1mlnx1-OFED.3.4.1.0.0.34100)
 libmlx5-dev(=1.2.1mlnx1-OFED.3.4.1.0.0.34100)
 libopensm(=4.8.0.MLNX20160906.32a95b6-0.34100)
 librdmacm1(=1.1.0mlnx-OFED.3.4.0.0.4.34100)
 librdmacm-dev(=1.1.0mlnx-OFED.3.4.0.0.4.34100)
 librdmacm-utils(=1.1.0mlnx-OFED.3.4.0.0.4.34100)
 libsdp1(=1.1.108-OFED.3.0.8.gfbd01df.34100)
 libsdp-dev(=1.1.108-OFED.3.0.8.gfbd01df.34100)
 mlnx-ethtool(=4.2-OFED.3.1.1.5.2.34100)
 mlnx-fw-updater(=3.4-1.0.0.0)
 mlnx-ofed-kernel-dkms(=3.4-OFED.3.4.1.0.0.1.g2ed8a21)
 mlnx-ofed-kernel-utils(=3.4-OFED.3.4.1.0.0.1.g2ed8a21)
 mlnx-rds-dkms(=3.4-OFED.3.4.1.0.0.1.g2ed8a21)
 mlnx-sdp-dkms(=3.4-OFED.3.4.1.0.0.1.g2ed8a21)
 mpitests(=3.2.18-0.34100)
 mstflint(=4.5.0-1.34100)
 nvidia-peer-memory(=1.0-1)
 nvidia-peer-memory-dkms(=1.0-1)
 ofed-scripts(=3.4-OFED.3.4.1.0.0)
 openmpi(=1.10.5a1-0.34100)
 opensm(=4.8.0.MLNX20160906.32a95b6-0.34100)
 perftest(=3.0-3.1.34100)
 rds-tools(=2.0.7-OFED.2.4.1.g299420c.34100)
 sdpNetstat(=1.60-OFED.3.0.10.g3cf409a.34100)
 srp-dkms(=1.6.1-OFED.3.4.1.0.0.1.g2ed8a21)
 srptools(=1.0.3-4.34100)

Misc

autoconf(>=2.69-6)
 build-essential(>=11.6ubuntu6)
 cachefilesd(>=0.10.5-1)
 cmake(>=2.8.12.2-0ubuntu3)

fail2ban(>=0.8.11-1)

iotop(>=0.6-1)

man

python(>=2.7.5-5ubuntu3)

python-dev(>=2.7.5-5ubuntu3)

ssh(>=1:7.3p1-1libressl)

vim(>=2:7.4.052-1ubuntu3)

APPENDIX B. 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.

libressl (OpenSSL)

LibreSSL <http://www.libressl.org/> version 2.4.1 is provided under the following terms:

```
* Copyright (c) 1998-2011 The OpenSSL Project. 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.
*
* 3. All advertising materials mentioning features or use of this
* software must display the following acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit. (http:// www.openssl.org/)"
*
```



```

* 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used
to * endorse or promote products derived from this software without
* prior written permission. For written permission, please contact
* openssl-core@openssl.org.
*
* 5. Products derived from this software may not be called "OpenSSL"
* nor may "OpenSSL" appear in their names without prior written *
permission of the OpenSSL Project.
*
* 6. Redistributions of any form whatsoever must retain the following
* acknowledgment:
* "This product includes software developed by the OpenSSL Project
* for use in the OpenSSL Toolkit (http:// www.openssl.org/)"
*
* THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY
* EXPRESSED 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 OpenSSL PROJECT OR
* ITS 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.

```

Mellanox (OFED)

MLNX OFED (<http://www.mellanox.com/>) version 3.3-1.0.4.0 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, and DGX-1 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

© 2016 NVIDIA Corporation. All rights reserved.