



NVIDIA DGX-1 SOFTWARE VERSION 3.1.1

DA-08260-311_v01 | August 2017

Release Notes and Installation Guide



TABLE OF CONTENTS

NVIDIA DGX-1 Software Release Notes for Version 3.1.1	3
Container Update Advisement	3
Highlights and Changes since Version 2.x	4
Known Issues	4
Version Matrix	5
DGX-1 Software Version 3.1.1 Installation Guide	6
About the Re-Imaging Process	6
Obtaining the DGX-1 Software Image	7
Connecting to the DGX-1 Console.....	7
Re-imaging the DGX-1 Remotely	8
Re-imaging the System from a USB Flash Drive	10
First Time Boot Setup	11
Setting Your Time Zone.....	12
Appendix A: Retaining the RAID Partition While Installing the OS	13
Appendix B. Third Party License Notice	15

NVIDIA DGX-1 SOFTWARE RELEASE NOTES FOR VERSION 3.1.1

This document describes the version 3.1.1 release of the NVIDIA® DGX-1™ software. The DGX-1 Software Version 3.1.1 is installed at the factory, and is also provided by Enterprise Support as an ISO image in the event a system needs to be re-imaged.

CONTAINER UPDATE ADVISEMENT

In conjunction with the DGX-1 3.1.1 software, customers should update their NVIDIA Docker containers to the latest container release¹.

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 instructions for how to access them.

HIGHLIGHTS AND CHANGES SINCE VERSION 2.X

- ▶ Ubuntu 16.04 LTS
 - Includes select Ubuntu updates released as of mid-June 2017 time frame.
 - Initialization daemon changed from **Upstart** to **systemd**.
 - Updated network interface naming policy.
Policy now uses predictable names, rather than the native naming scheme used in previous releases. The first and second Ethernet interfaces, enumerated as em1 and em2 in previous releases, will now enumerate as enp1s0f0 and enp1s0f1 respectively
- ▶ Mellanox drivers updated to 4.0.
- ▶ NVIDIA GPU Driver 384.66
- ▶ Docker CE and nvidia-docker are pre-installed, and the docker daemon automatically launched.

KNOWN ISSUES

- ▶ The GPUs cannot be reset while the system is running. If an issue occurs which causes the GPUs to hang or if they need to be reset, you must reboot the system.
- ▶ During the initial setup, the English language is set as the default and only US time zones are available for selection.
You can configure other time zones after the initial setup is completed. See the section [Setting Your Time Zone](#).
- ▶ Apparmor is enabled in the Base OS software, with Docker generating a default profile.
The default profile may or may not work with your containers. 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.

VERSION MATRIX

The following table provides version information for software included in the base OS ISO, the latest firmware and BIOS versions for the DGX-1 hardware, and tested Docker and nvidia-docker versions.

Component	Version	Notes
DGX-1 Base OS Software	3.1.1	Software versions included in this DGX-1 base OS ISO
Graphics Driver	384.66	
Ubuntu	16.04 LTS	
Docker CE	17.03.1	
nvidia-docker	1.0.1	
DGX-1 BMC	3.12.30	Released versions for DGX-1 hardware (Pascal) at the time of this software release. Information provided for reference purposes.
DGX-1 SBIOS	S2W_3A01	
DGX-1 VBIOS	86.00.41.00.05	
DGX-1 BMC	3.16.30	Released versions for DGX-1 hardware (Volta) at the time of this software release. Information provided for reference purposes.
DGX-1 SBIOS	S2W_3A03	
DGX-1 VBIOS	88.00.18.00.00	

DGX-1 SOFTWARE VERSION 3.1.1 INSTALLATION GUIDE

This chapter provides general instructions for re-imaging the DGX-1 with software version 3.1.1 from the ISO file and then performing the first time DGX-1 setup. Refer to the *DGX-1 User Guide* (<http://docs.nvidia.com/dgx/dgx1-user-guide/index.html>) for additional instructions where indicated.

ABOUT THE RE-IMAGING PROCESS



WARNING: This process destroys all data and software customizations that you have made on the DGX-1. Be sure to back up any data that you want to preserve, and push any Docker images that you want to keep to a trusted registry.

The re-imaging process creates a fresh installation of the DGX-1 base OS software. The standard software installation overwrites any data or file systems that may exist on the OS disk as well as the RAID disks. All the disks are re-partitioned, and the RAID array is mounted as /raid.

Since the RAID array on the DGX-1 is intended to be used as a cache and not for long-term data storage, this should not be disruptive. However, if you have set up the disks for a non-cache purpose and want to keep the data on those drives, then select the **Install DGX Server without formatting RAID** option at the boot menu during the boot installation. When the installation is completed, you can repeat any configurations steps needed to use the RAID disks as other than cache disks. See [Appendix A: Choosing not to Reformat the RAID Disks](#) for details on what this option does.

OBTAINING THE DGX-1 SOFTWARE IMAGE

Obtain the correct ISO image file from NVIDIA Support Enterprise Services.

1. Log on to the NVIDIA Enterprise Services (<https://nvid.nvidia.com/enterpriselogin/>) site.
2. Click the Announcements tab to locate the download links for the DGX-1 software image.
3. Download the ISO image and save it to your local disk.
4. Restore the DGX-1 software image from this file either remotely through the BMC or locally from a bootable USB flash drive as explained in the following sections.

CONNECTING TO THE DGX-1 CONSOLE

Connect to the DGX-1 console using either a direct connection or a remote connection through the BMC.

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.

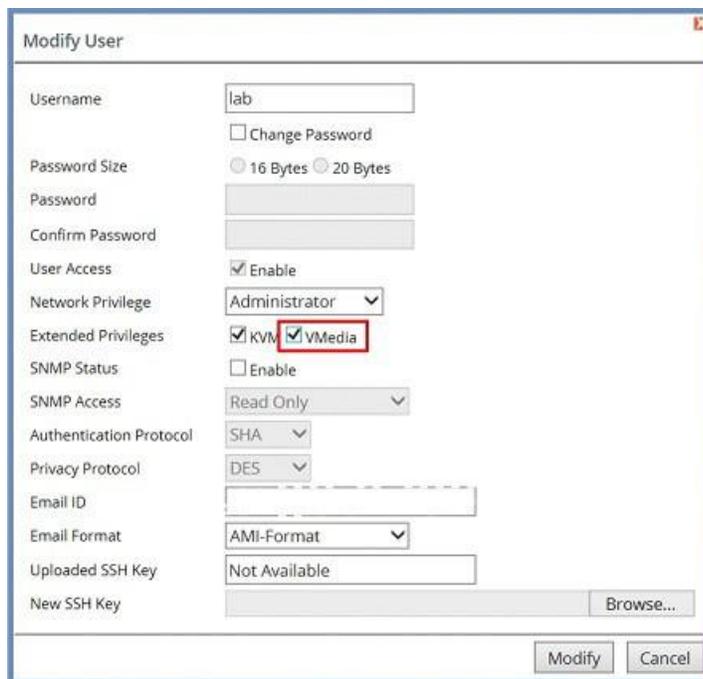
1. Make sure you have connected the IPMI port on the DGX-1 to your LAN.
2. Open a Java-enabled browser within your LAN and go to `http://<IPMI-IP-address>/`. Make sure popups are allowed for the BMC address.
3. Log in.

If the administrator has not manually created a password, then the username that was created during the initial DGX-1 setup is used for both the BMC username and BMC password.

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.

RE-IMAGING THE DGX-1 REMOTELY

1. Connect to the BMC and change user privileges.
 - a) Open a Java-enabled web browser within your LAN and go to `http://<IPMI-IP-address>/`, then log in.
Use Firefox or Internet Explorer. Google Chrome is not officially supported by the BMC.
 - b) From the top menu, click **Configuration** and then select **User Management**.
 - c) Select the user name that you created for the BMC, then click **Modify User**.
 - d) In the **Modify User** dialog, select the **VMedia** check box to add it to the extended privileges for the user, then click **Modify**.



2. Set up the ISO image as virtual media.
 - a) From the top menu, click **Remote Control** and select **Console Redirection**.



- b) Click **Java Console** to open the remote JViewer window.
Make sure pop-up blockers are disabled for this site.
 - c) From the JViewer top menu bar, click **Media** and then select **Virtual Media Wizard**.

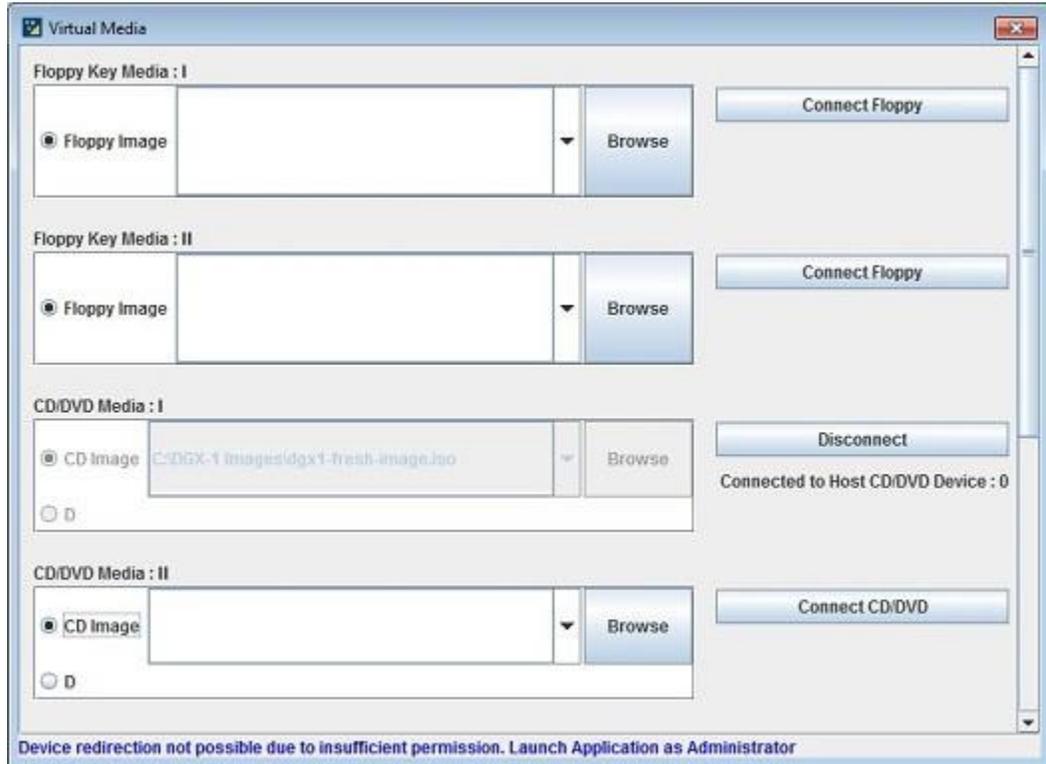


- d) From the **CD/DVD Media: I** section of the **Virtual Media** dialog, click **Browse** and then locate the re-image ISO file and click Open.

You can ignore the device redirection warning at the bottom of the **Virtual Media** wizard as it does not affect the ability to re-image the system.

- e) Click **Connect CD/DVD**, then click **OK** at the Information dialog.

The **Virtual Media** window shows that the ISO image is connected.



- f) Close the window.

The CD ROM icon in the menu bar turns green to indicate that the ISO image is attached.



3. Reboot, install the image, and complete the DGX-1 setup.

- a) From the top menu, click **Power** and then select **Reset Server**.



- b) Click **Yes** and then **OK** at the **Power Control** dialogs, then wait for the system to power down and then come back online.
- c) At the boot selection screen, select **Install DGX Server**.

If you are not using the RAID disks as cache and want to keep data on the RAID disks, then select **Install DGX Server without formatting RAID**.

d) Press **Enter**.

The DGX-1 will reboot from CDR0M0 1.00, and proceed to install the image. This can take approximately 15 minutes.

 **NOTE:** The Mellanox InfiniBand driver installation may take up to 10 minutes.

After the installation is completed, the system ejects the virtual CD and then reboots into the OS.

Refer to the section [First Time Boot Setup](#) for the steps to take when booting up the DGX-1 for the first time after a fresh installation.

RE-IMAGING THE SYSTEM FROM A USB FLASH DRIVE

Before re-imaging the system from a USB flash drive, ensure that you have a bootable USB flash drive that contains the current DGX-1 software image. Refer to the DGX-1 User Guide for instructions on creating a bootable USB flash drive that contains the ISO image.

1. Plug the USB flash drive containing the OS image into the DGX-1.
2. Connect a monitor and keyboard directly into the DGX-1.
3. Boot the system and press **F11** when the NVIDIA logo appears in order to get to the boot menu.
4. Select the USB volume name that corresponds to the inserted USB flash drive, and then boot the system from it.

5. When the system boots up, select **Install DGX Server** from the boot menu.

If you are not using the RAID disks as cache and want to keep data on the RAID disks, then select **Install DGX Server without formatting RAID**.

6. Press **Enter**.

The DGX-1 will reboot and proceed to install the image. This can take more than 15 minutes.

 **NOTE:** The Mellanox InfiniBand driver installation may take up to 10 minutes.

After the installation is completed, the system then reboots into the OS.

Refer to the next section for the steps to take when booting up the DGX-1 for the first time after a fresh installation.

FIRST TIME BOOT SETUP

When you power on the DGX-1 for the first time, you are presented with the system setup screens.

1. Perform the steps to configure the DGX-1 software.

- Select your time zone and keyboard layout.
- Create a user account with your name, username, and password.

You will need these credentials to log in to the DGX-1 as well as to log in to the BMC remotely. When logging in to the BMC, enter your username for both the User ID as well as the password. Be sure to create a unique BMC password at the first opportunity.

- Choose a primary network interface for the DGX-1.



NOTE: After you select the primary network interface, the system attempts to configure the interface for DHCP and then asks you to enter a hostname for the system. If DHCP is not available, you will have the option to configure the network manually. If you need to configure a static IP address on a network interface connected to a DHCP network, then select **Cancel** at the *Network configuration - Please enter the hostname for the system* screen. The system will then present a screen with the option to configure the network manually.

- Choose a hostname for the DGX-1.
- Choose to install predefined software.

Toggle the space bar to select or deselect the software to install.

By default, the DGX-1 installs only minimal software packages necessary to ensure system functionality. During installation, you can deselect the OpenSSH package; however, NVIDIA recommends that you keep this package selected, and uninstall it only if required by your IT security policy.

- Select **OK** to continue.

You are presented with end user license agreement (EULAs) for the NVIDIA software.

2. Accept all EULAs.

The system completes the installation, reboots, then presents the system login prompt:

```
<hostname> login:
Password:
```

3. Log in.

Refer to the DGX-1 User Guide for additional information on configuring your DGX-1 system.

SETTING YOUR TIME ZONE

During the initial software setup, only US time zones were available for selection. You can set a different time zone as follows:

1. Check the current time zone.

```
$ more /etc/timezone
```

```
US/Los Angeles
```

2. Configure a new time zone.

```
$ sudo dpkg-reconfigure tzdata
```

Follow the instructions to set a specific time zone.

3. Verify the new time zone.

```
$ more /etc/timezone
```

Example if Berlin was selected:

```
Europe/Berlin
```

APPENDIX A: RETAINING THE RAID PARTITION WHILE INSTALLING THE OS

As part of the OS installation or re-imaging process, you are presented with a boot menu when booting the installer image. The default selection is **Install DGX Software**. The installation process then repartitions all the SSDs, including the OS SSD as well as the RAID SSDs.

If you have not been using the RAID SSDs as cache drives and do not want them reformatted, then select the **Install DGX Server without formatting RAID** option at the boot menu. This option retains data on the RAID disks and performs the following:

- ▶ Installs the cache daemon but leaves it disabled by commenting out the `RUN=yes` line in `/etc/default/cachefile`.
- ▶ Creates a `/raid` directory, leaves it out of the file system table by commenting out the entry containing `"/raid"` in `/etc/fstab`.
- ▶ Does not format the RAID disks.

You can always choose to use the RAID disks as cache disks at a later time by enabling `cachefile` and adding `/raid` to the file system table as follows:

1. Uncomment the `#RUN=yes` line in `/etc/default/cachefile`.
2. Uncomment the `/raid` line in `etc/fstab`.
3. Run the following:

Mount `/raid`.

```
sudo mount /raid
```

Reload the `systemd` manager configuration.

```
systemctl daemon-reload
```

Start the cache daemon.

```
systemctl start cachefilesd.server
```

These changes are preserved across system reboots.

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/>) 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

© 2017 NVIDIA Corporation. All rights reserved.