NVIDIA A10 GPU Accelerator

Product Brief

PB-10415-001_v04 | March 2022

Document History

PB-10415-001_v04

Version	Date	Authors	Description of Change
01	March 16, 2021	AV, SM	Initial Release
02	April 12, 2021	KD, SM	Removed confidential markings
03	January 12, 2022	AV, SM	• Corrected VF BAR1 information in Table 3 • Updated Table 4
04	March 18, 2022	VK, SM	• Removed Secure Boot from Table 1 • Updated Secure Boot and CEC description in Table 3 • Updated "Root of Trust" section • Added Table 5 "Root of Trust Feature Set"

Table of Contents

List of Figures

List of Tables

Overview

The NVIDIA® A10 Tensor Core graphics processing unit (GPU) delivers a versatile platform for Graphics and Video processing, as well as Deep Learning Inferencing in distributed computing environments. It combines the 2nd generation NVIDIA® TensorRT™ cores, 3rd generation tensor cores with 24 GB of GDDR6 memory in a single-slot 10.5-inch PCI Express Gen4 form factor, with 150 W maximum board power. The card is passively cooled that requires system airflow to operate within its thermal envelope.

Powered by the NVIDIA Ampere architecture, the NVIDIA A10 universal GPU provides revolutionary multi-precision performance to accelerate mixed workloads from a single, GPU accelerated infrastructure. When combined with NVIDIA RTX™ Virtual Workstation (vWS) software, A10 is ideal for running high-performance virtual workstations running professional visualization applications or combine with NVIDIA Virtual PC (vPC) software for multimedia-rich virtual desktops. It can also support deep learning and machine learning training and inference, video transcoding, cloud gaming, AI audio and video effects, rendering, data analytics, and many other workloads.

As part of the NVIDIA AI Platform, the A10 supports all AI frameworks and network models, delivering dramatic performance and efficiency that maximizes the utility of at-scale deployments.

Figure 1. NVIDIA A10 PCIe Card

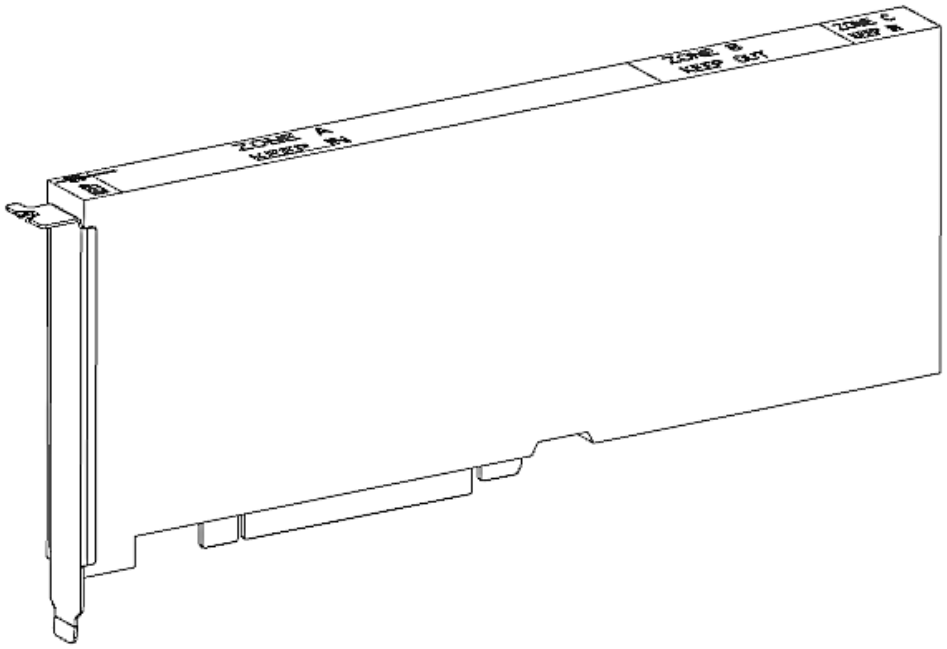

Specifications

Product Specifications

Table 1 through Table 3 provide the product, memory, and software specifications for the NVIDIA A10 PCIe card.

Table 1. Product Specifications

Specification	NVIDIA A10
Product SKU	PG133 SKU 215 NVPN: 699-2G133-0215-C01 or later
Total board power	150 W
Thermal solution	Passive
Mechanical Form Factor	Full-height, full-length (FHFL) 10.5", single slot
GPU SKU	GA102-890
PCI Device IDs	Device ID: 0x2236 Vendor ID: 0x10DE Sub-Vendor ID: 0x10DE Sub-System ID: 0x1482
GPU clocks	Base: 885 MHz Boost: 1695 MHz
Performance States	P0, P8
VBIOS	EEPROM size: 8 Mbit UEFI: Supported
PCI Express interface	PCI Express 4.0 ×16, x8; PCIe 3.0 x16 Lane and polarity reversal supported
Zero Power	Not supported
NEBS readiness	Supported
Power connectors	One PCIe 8-pin auxiliary power connector

Specification	NVIDIA A10
Weight	Board: 550 Grams (excluding bracket and extenders) Bracket with screws: 12 Grams Long offset extender: 64 Grams Straight extender: 39 Grams

Table 2. Memory Specifications

Specification	Description
Memory clock	6251 MHz
Memory type	GDDR6
Memory size	24 GB
Memory bus width	384 bits
Peak memory bandwidth	Up to 600 GB/sec

Table 3. Software Specifications

Specification	Description[1]
SR-IOV support	Supported: 32 VF (virtual functions)
BAR address (physical function)	BAR0: 16 MiB[1] BAR1: 32 GiB[1] BAR3: 32 MiB[1]
BAR address (virtual function)	BAR0: 8 MiB, (256 KiB per VF)[1] BAR1: 64 GiB, 64-bit (2 GiB per VF)[1] BAR3: 1 GiB, 64-bit (32 MiB per VF)[1]
Message signaled interrupts	MSI-X: Supported MSI: Not supported
Multi-Instance GPU (MIG)	Not supported
ARI Forwarding	Supported
Driver support	Linux: R460.21 or later Windows: R460.57 or later
Secure Boot	Supported (See "Root of Trust" section)
CEC Firmware	v5.01 or later (for CEC-enabled cards)
NVIDIA® CUDA® Support	CUDA 11.2 or later
Virtual GPU Software Support	Supports vGPU 12.x or later: • NVIDIA RTX Virtual Workstation (vWS) • NVIDIA Virtual PC (vPC)/Virtual Applications (vApps) • NVIDIA AI Enterprise • NVIDIA Virtual Compute Server (vCS)

Specification	Description[1]
NVIDIA® NGC-Ready™ Test Suite	NGC-Next Certification 2.x or later
PCI class code	0x03 – Display Controller
PCI sub-class code	0x02 – 3D Controller
Primary Boot Device Capability	Not supported
ECC support	Enabled (by default). Can be disabled via software
SMBus (8-bit address)	0x9E (write), 0x9F (read)
Reserved I2C addresses[2]	0xAA, 0xAC
SMBus direct access	Supported
SMBPBI SMBus Post-Box Interface)	Supported

Note:

[1]The KiB, MiB and GiB notation emphasizes the "power of two" nature of the values. Thus,

- 256 KiB = 256 x 1024
- 16 MiB = 16 x 1024^2
- 64 GiB = 64 x 1024^3

[2]See "Root of Trust" section of this product brief.

The operator is given the option to configure this power setting to be persistent across driver reloads or to revert to default power settings upon driver unload.

Environmental and Reliability Specifications

Table 4 provides the environment conditions specifications for the A10 PCIe card.

Table 4. Board Environmental and Reliability Specifications

Specification	Description
Ambient operating temperature	0 °C to 50 °C
Ambient operating temperature (short term)[1]	-5 °C to 55 °C
Storage temperature	-40 °C to 75 °C
Operating humidity (short term)[1]	5% to 93% relative humidity
Operating humidity	5% to 85% relative humidity
Storage humidity	5% to 95% relative humidity
Mean time between failures (MTBF)	Uncontrolled environment:[2] 1,198,847 hours at 35 °C Controlled environment:[3] 1,412,904 hours at 35 °C

Notes: Specifications in this table are applicable up to 6000 feet.

[1]A period not more than 96 hours consecutive, not to exceed 15 days per year.

[2]Some environmental stress with limited maintenance (GF35).

[3]No environmental stress with optimum operation and maintenance (GB35).

Airflow Direction Support

The NVIDIA A10 PCIe card employs a bidirectional heat sink, which accepts airflow either left-to-right or right-to-left directions.

Figure 2. A10 Airflow Directions

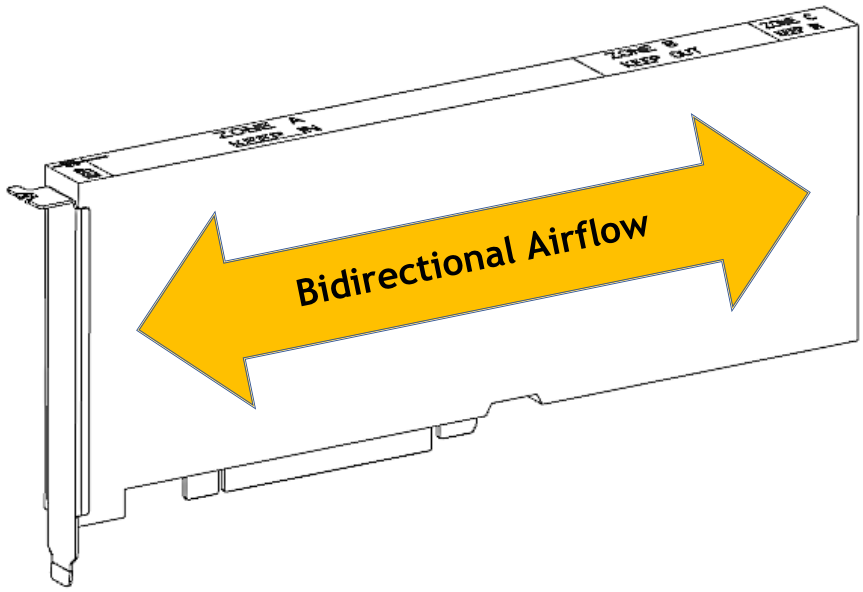

Product Features

PCI Express Interface Specifications

The following sub-sections describe the PCIe interface specifications for the A10 PCIe card.

PCIe Speed Support

The A10 card supports PCIe Gen4.

Polarity Inversion and Lane Reversal Support

Lane Polarity Inversion, as defined in the PCIe specification, is supported on the A10 PCIe card.

Lane Reversal, as defined in the PCIe specification, is supported on the A10 PCIe card. When reversing the order of the PCIe lanes, the order of both the Rx lanes and the Tx lanes must be reversed.

Root of Trust

The NVIDIA A10 GPU has a primary root of trust within the GPU chip that provides the following:

- Secure boot
- Secure firmware upgrade
- Firmware rollback protection
- Ability to disable In-band firmware update (established after each GPU reset)
- Secure application processor recovery

In addition, NVIDIA offers some A10 boards with an onboard CEC1712 chip, acting as a secondary root of trust, extending the security capabilities allowing for firmware attestation, key revocation, and out-of-band firmware updates. The CEC1712 device authenticates the contents of the GPU firmware ROM before permitting the GPU to boot from its ROM. For CEC1712-enabled cards, the root of trust feature occupies up to two I2C addresses (in addition

to the SMBus addresses). I2C addresses 0xAA and 0xAC should therefore be avoided for system use.

Identification of the two variants of A10 boards (with or without CEC1712) can be done using the 900-level part number on the back of the GPU or running the nvidia-smi -q command.

- 900-2G133-XXXX-0XX A10 GPUs without CEC1712
- 900-2G133-XXXX-1XX A10 GPUs with CEC1712
- Exceptions
  - 900-2G133-0020-100 is without CEC1712
  - 900-2G133-0020-000 with CEC1712

The following table shows the features that are available using the primary and secondary root of trust.

Table 5. Root of Trust Feature Set

Features	Primary Root of Trust within GPU Chip	Secondary Root of Trust Using External CEC Chip on Board
Secure Boot	Yes	Yes
Secure Firmware Upgrade	Yes	Yes
Firmware Rollback Protection	Yes	Yes
In-Band Firmware Update Disable	Yes[1]	Yes
Key Revocation	No	Yes
Firmware Attestation	No	Yes

Notes:

[1]"In-Band Firmware Update Disable" feature must be established after every GPU reset.

Form Factor

In this product brief, nominal dimensions are shown in Figure 3.

Figure 3. NVIDIA A10 PCIe Card Dimensions

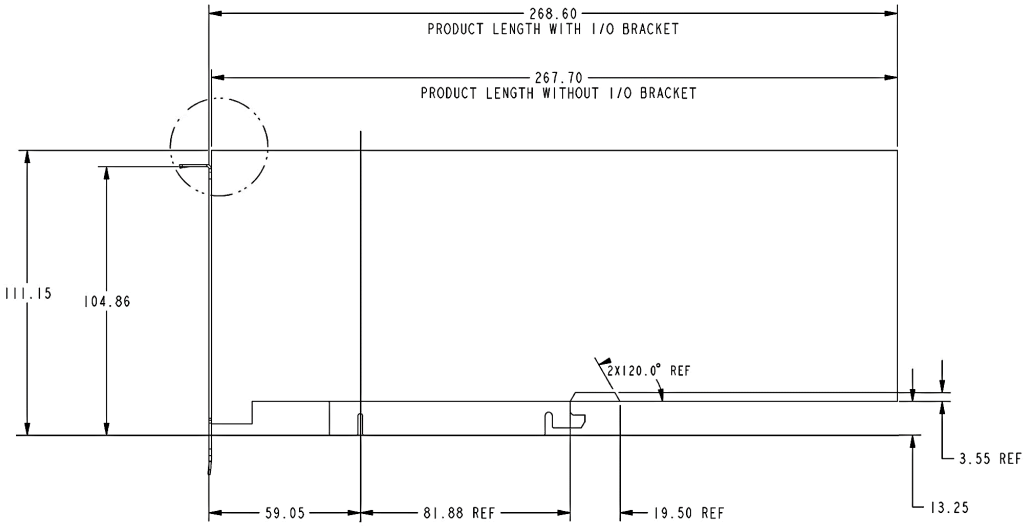

Power Connector Placement

The board provides a PCIe 8-pin power connector on the east edge of the board.

Figure 4. PCIe 8-Pin Power Connector

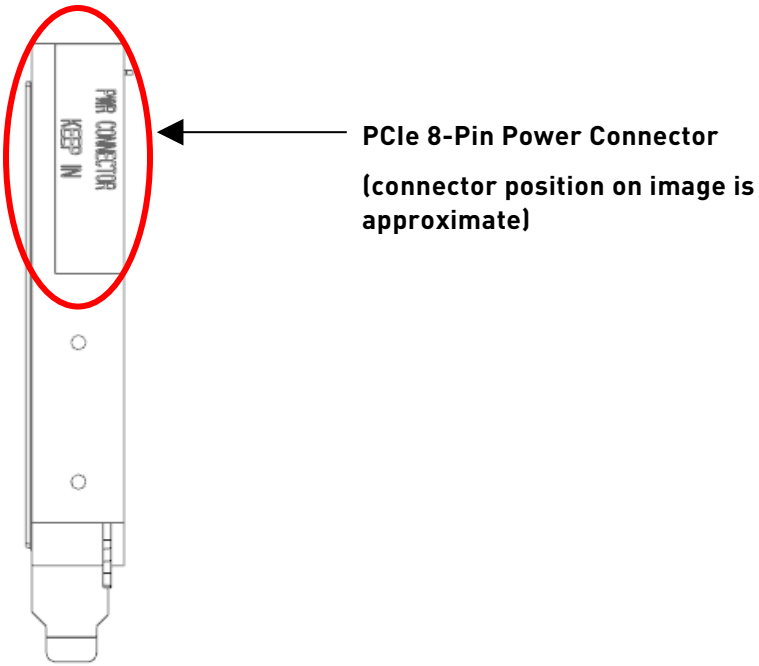

Table 6 lists supported auxiliary power connections for the NVIDIA A10 PCIe card.

Table 6. Supported Auxiliary Power Connections

Board Connector	PSU Cable
PCIe 8-pin	1× PCIe 8-pin cable

Extenders

The A10 PCIe card provides two extender options, shown in Figure 5.

- NVPN: 682-00007-5555-000 – Long offset extender
  - Card + extender = 339 mm
- NVPN: 682-00007-5555-001 – Straight extender
  - Card + extender = 312 mm

Using the standard NVIDIA extender ensures greatest forward compatibility with future NVIDIA product offerings.

If the standard extender will not work, OEMs may design a custom attach method using the extender mounting holes on the east edge of the PCIe card.

Figure 5. Long Offset and Straight Extenders

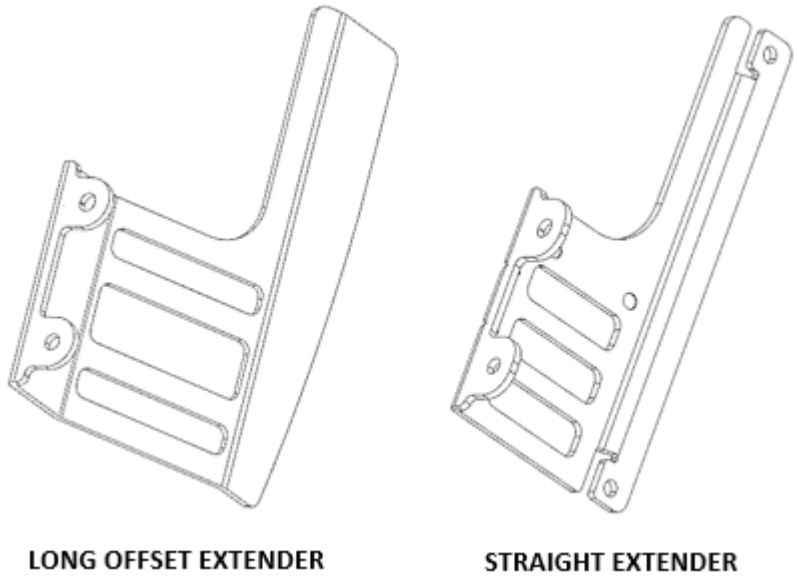

Support Information

Certifications

- Windows Hardware Quality Lab (WHQL):
  - Certified Windows 7, Windows 8.1, Windows 10
  - Certified Windows Server 2008 R2, Windows Server 2012 R2
- Ergonomic requirements for office work W/VDTs (ISO 9241)
- EU Reduction of Hazardous Substances (EU RoHS)
- Joint Industry guide (J-STD) / Registration, Evaluation, Authorization, and Restriction of Chemical Substance (EU) – (JIG / REACH)
- Halogen Free (HF)
- EU Waste Electrical and Electronic Equipment (WEEE)

Agencies

- Australian Communications and Media Authority and New Zealand Radio Spectrum Management (RCM)
- Bureau of Standards, Metrology, and Inspection (BSMI)
- Conformité Européenne (CE)
- Federal Communications Commission (FCC)
- Industry Canada - Interference-Causing Equipment Standard (ICES)
- Korean Communications Commission (KCC)
- Underwriters Laboratories (cUL, UL)
- Voluntary Control Council for Interference (VCCI)

Languages

Table 7. Languages Supported

Languages	Windows[1]	Linux
English (US)	Yes	Yes
English (UK)	Yes	Yes
Arabic	Yes	
Chinese, Simplified	Yes	
Chinese, Traditional	Yes	
Czech	Yes	
Danish	Yes	
Dutch	Yes	
Finnish	Yes	
French (European)	Yes	
German	Yes	
Greek	Yes	
Hebrew	Yes	
Hungarian	Yes	
Italian	Yes	
Japanese	Yes	
Korean	Yes	
Norwegian	Yes	
Polish	Yes	
Portuguese (Brazil)	Yes	
Portuguese (European/Iberian)	Yes	
Russian	Yes	
Slovak	Yes	
Slovenian	Yes	
Spanish (European)	Yes	
Spanish (Latin America)	Yes	
Swedish	Yes	
Thai	Yes	
Turkish	Yes	

Note:

[1]Microsoft Windows 7, Windows 8, Windows 8.1, Windows 10, Windows Server 2008 R2, Windows Server 2012 R2, and Windows 2016 are supported.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform 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 accepts no 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.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, CUDA, NGC-Ready, NVIDIA RTX, and TensorRT are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

NVIDIA Corporation | 2788 San Tomas Expressway, Santa Clara, CA 95051
http://www.nvidia.com