



NVIDIA Quadro RTX 6000 PCIe Server Card

Product Brief

Document History

PB-09762-001_v01

Version	Date	Authors	Description of Change
01	December 13, 2019	DK, SM	Initial Release

Table of Contents

- Overview 1
 - Airflow Direction Support2
- Specifications..... 3
 - Product Specifications3
 - Environmental and Reliability Specifications.....5
- Product Features 6
 - Form Factor6
 - NVLink Connector Placement7
 - Power Connector Placement.....8
- Support Information 10
 - Certifications10
 - Agencies10
 - Languages.....11

List of Figures

Figure 1.	NVIDIA Quadro RTX 6000 Server Card	1
Figure 2.	Quadro RTX 6000 Server Card: Left-to-Right Airflow.....	2
Figure 3.	Quadro RTX 6000 Server Card: Right-to-Left Airflow.....	2
Figure 4.	NVIDIA Quadro RTX 6000 Server Card Dimensions.....	6
Figure 5.	NVLink Connector Placement – Top View	7
Figure 6.	NVLink Bridge Clearance	8
Figure 7.	CPU 8-Pin Power Connector.....	8

List of Tables

Table 1.	Product Specifications.....	3
Table 2.	Memory Specifications	4
Table 3.	Software Specifications	4
Table 4.	Board Environmental and Reliability Specifications	5
Table 5.	Supported Sync Configurations.....	7
Table 6.	Supported Auxiliary Power Connections.....	9
Table 7.	Languages Supported	11


Overview

The NVIDIA® Quadro RTX™ 6000 Server Card is a dual-slot, 10.5-inch PCI Express Gen3 graphics solution based on the state-of-the-art NVIDIA Turing™ architecture. This server card version of the Quadro RTX 6000 is a passively cooled board capable of 250 W maximum board power.

The Quadro RTX 6000 Server Card is the world's first server-specific graphics card delivering hardware-accelerated ray tracing, revolutionary AI features, advanced shading and powerful simulation capabilities to creative professionals. With 24 GB of GDDR6 memory (expandable to 48 GB with NVIDIA® NVLink®), even the most graphics-intensive applications run with the highest level of user experience, including those with very large data sets.

The Quadro RTX 6000 Server Card is intended for deployment in data center or edge servers.

Figure 1. NVIDIA Quadro RTX 6000 Server Card



Airflow Direction Support

The NVIDIA Quadro RTX 6000 passively cooled server card supports bidirectional airflow. The following diagrams illustrate the airflow options supported.

Figure 2. Quadro RTX 6000 Server Card: Left-to-Right Airflow





Figure 3. Quadro RTX 6000 Server Card: Right-to-Left Airflow



Specifications

Product Specifications

Table 1 provides the product specifications for the NVIDIA Quadro RTX 6000 Server Card.

Table 1. Product Specifications

Specification	Description
Product SKU	PG150 SKU 230 NVPN: 699-2G150-0230-xxx
Total board power	250 W
Thermal solution	Passive
Mechanical form factor and dimensions	NVIDIA Form Factor 4.0 for PCIe Card Server Products 4.376" High × 10.5" Long, dual slot
GPU SKU	TU102-895-A1
PCI IDs	Device ID: 0x1E78 Vendor ID: 0x10DE Sub-Vendor ID: 0x10DE Sub-System ID: 0x13D9
NVIDIA® CUDA® cores	4608
GPU clocks	Base: 1275 MHz Boost: 1620 MHz
VBIOS	EEPROM size: 8 Mbit UEFI: Supported
PCI Express interface	PCI Express 3.0 × 16 Lane and polarity reversal supported
Power connectors and headers	One CPU 8-pin auxiliary power connector
Weight	Board: 1200 Grams (excluding bracket and extenders) Bracket with screws: 20 Grams Long offset extender: 64 Grams Straight extender: 39 Grams

Table 2 provides the memory specifications for the NVIDIA Quadro RTX 6000 Server Card.

Table 2. Memory Specifications

Specification	Description
Memory clocks	Performance: 6501 MHz Idle: 405 MHz
Memory size	24 GB
Memory bus width	384-bit
Peak memory bandwidth	Up to 624 GB/s

Table 3 provides the software specifications for the NVIDIA Quadro RTX 6000 Server Card.

Table 3. Software Specifications

Specification	Description
SR-IOV support	Supported -- 24 VF (virtual functions)
BAR address (physical function)	BAR0: 16 MB BAR1: 32 GB BAR3: 32 MB
BAR address (virtual function)	BAR0: 384 MB, 32-bit (24 VF x 16 MB) BAR1: 6 GB, 64-bit (24 VF x 256 MB) BAR3: 768 MB, 64-bit (24 VF x 32 MB)
Message signaled interrupts	MSI-X: Supported MSI: Not supported
PCI class code	0x03 – Display Controller
PCI sub-class code	0x02 – 3D Controller
ECC support	Enabled (by default). Can be disabled via software
SMBus (8-bit address)	0x9E (write), 0x9F (read)
SMBus direct access	Supported
SMBPBI SMBus Post-Box Interface)	Supported

Environmental and Reliability Specifications

Table 4 provides the environment conditions specifications for the Quadro RTX 6000 Server Card.

Table 4. Board Environmental and Reliability Specifications


Specification	Description
Ambient operating temperature	0 °C to 50 °C
Storage temperature	-40 °C to 75 °C
Operating humidity	5% to 95% relative humidity
Storage humidity	5% to 95% relative humidity
Mean time between failures (MTBF)	Uncontrolled environment: ¹ 703379.3 hours at 35 °C Controlled environment: ² 913208.1 hours at 35 °C
Notes: ¹ Some environmental stress with limited maintenance. ² No environmental stress with optimum operation and maintenance.	

Product Features

Form Factor

The NVIDIA Quadro RTX 6000 Server Card conforms to the *NVIDIA Form Factor 4.0 Specification for Server Products*.


Figure 4. NVIDIA Quadro RTX 6000 Server Card Dimensions



NVLink Connector Placement

Figure 5 shows the connectors for the NVLink bridge on the Quadro RTX 6000 Server Card.

Figure 5. NVLink Connector Placement – Top View



The Quadro RTX 6000 Server Card supports the following high bandwidth NVLink bridges:

- ▶ 2-slot span NVLink bridge
- ▶ 3-slot span NVLink bridge


Table 5 presents the supported Sync configurations for the Quadro RTX 6000 Server Card.

Table 5. Supported Sync Configurations

Configuration	Supported
High-bandwidth NVIDIA NVLink bridge (2-slot span) on NVLink connector	Yes
High-bandwidth NVIDIA NVLink bridge (3-slot span) on NVLink connector	Yes

Sufficient clearance must be provided both above the north edge of the card and behind the backside of the card's PCB to accommodate an NVIDIA Quadro NVLink bridge. The clearance above the card's north edge should meet or exceed 18 mm. The backside clearance (from the PCB's rear surface) should meet or exceed 7 mm. Figure 6 illustrates the NVLink bridge clearance.

Figure 6. NVLink Bridge Clearance



Power Connector Placement

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

Figure 7. CPU 8-Pin Power Connector

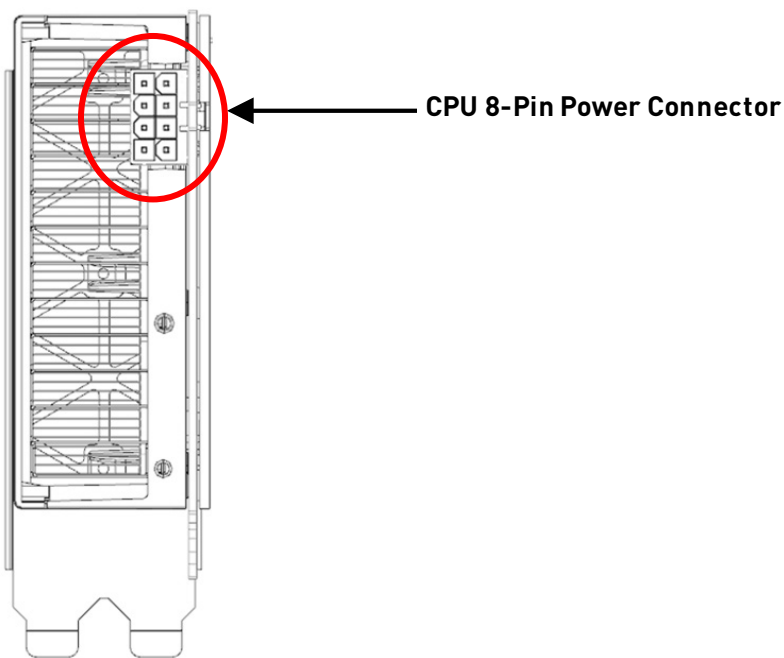


Table 6 lists supported auxiliary power connections for the NVIDIA Quadro RTX 6000 Server Card.

Table 6. Supported Auxiliary Power Connections

Board Connector	PSU Cable
CPU 8-pin	1× CPU 8-pin cable
CPU to PCIe 8-pin dongle	2× PCIe 8-pin cable 2× PCIe 6-pin cable ¹ 1× PCIe 8-pin cable and 1× PCIe 6-pin cable ¹
Note: ¹ Each PCIe 6-pin cable must be capable of carrying up to 50% of the average power and tolerate 50% of the peak currents	

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

The information provided in this specification is believed to be accurate and reliable as of the date provided. However, NVIDIA Corporation ("NVIDIA") does not give any representations or warranties, expressed or implied, as to the accuracy or completeness of such information. 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 publication supersedes and replaces all other specifications for the product that may have been previously supplied.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and other changes to this specification, at any time and/or to discontinue any product or service without notice. Customer should obtain the latest relevant specification 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. NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this specification.

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 these specifications 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 specification. 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 specification, 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 specification. 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 specification 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.

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. 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 NVIDIA terms and conditions of sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, CUDA, NVIDIA Turing, NVLink, and Quadro RTX 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

© 2019 NVIDIA Corporation. All rights reserved.