




TESLA P4 GPU ACCELERATOR

PB-08449-001_v01 | February 2017

Product Brief



DOCUMENT CHANGE HISTORY

PB-08449-001_v01

Version	Date	Authors	Description of Change
01	February 15, 2017	HV, SM	Initial Release

TABLE OF CONTENTS

- Overview 1
- Specifications..... 2
 - Product Specifications 2
 - Form Factor 4
 - Airflow Direction Support..... 5
- Support Information..... 6
 - Certificates and Agencies..... 6
 - Certifications 6
 - Agencies 6
 - Languages 7

LIST OF FIGURES

Figure 1.	Tesla P4 GPU Accelerator	1
Figure 2.	Tesla P4 Board Dimensions.....	4

LIST OF TABLES

Table 1.	Product Specifications	2
Table 2.	Memory Specifications	3
Table 3.	Software Specifications.....	3
Table 4.	Board Environmental and Reliability Specifications	4
Table 5.	Languages Supported	7

OVERVIEW

The NVIDIA® Tesla® P4 is a single-slot, low profile, 6.6 inch PCI Express Gen3 GPU Accelerator with an NVIDIA® Pascal™ graphics processing unit (GPU). The Tesla P4 has 8 GB GDDR5 memory and a 75 W maximum power limit. The Tesla P4 is offered as a 75 W or 50 W passively cooled board that requires system air flow to properly operate the card within thermal limits.

The NVIDIA Tesla P4 features optimized INT8 instructions aimed at deep learning inference computations. As a result, the NVIDIA Tesla P4 delivers 21 TOPs (Tera-Operations per second) of inference performance, enabling smart responsive artificial intelligence (AI)-based services.

For performance optimization this board utilizes NVIDIA GPU Boost™, which will dynamically adjust the GPU clock to maximize performance within thermal limits.




Figure 1. Tesla P4 GPU Accelerator

SPECIFICATIONS

PRODUCT SPECIFICATIONS

Table 1 provides the product specifications for the two configurations of the Tesla P4 GPU Accelerator.

Table 1. Product Specifications

Specification		Description
Product SKUs		PG414 SKU 200; NVPN: 699-2G414-0200-100 PG414 SKU 201; NVPN: 699-2G414-0201-100
Total board power		PG414 SKU 200: 75 Watts PG414 SKU 201: 50 Watts
GPU SKU		GP104-895-A1
PCI Device ID		0x1BB3
SSID		SKU 200: 0x11D8 SKU 201: 0x11E0
Board ID		0x10DE
NVIDIA® CUDA® cores		2560
GPU clocks	Base	885 MHz
	Maximum boost	1531 MHz (1113 MHz default)
	Idle	405 MHz
VBIOS	EEPROM size	4 Mbit
	UEFI	Supported
PCI Express interface		PCI Express 3.0 ×16 Lane and polarity reversal supported
Weight	Board	240 Grams
	Bracket with screws	9.7 Grams (low-profile bracket)
		15.6 Grams (ATX bracket)

Table 2 provides the memory specifications for the Tesla P4 GPU Accelerator.

Table 2. Memory Specifications

Specification		Description
Memory clocks	Performance	2.8 GHz
	Idle	324 MHz
Memory size		8 GB
Memory bus width		256-bit
Memory configuration		8 pcs 256M × 32 GDDR5
Peak memory bandwidth		Up to 192 GBytes/s

Table 3 provides the software specifications.

Table 3. Software Specifications

Specification	Description
Compatibility mode supported	Compute (default) Graphics (optional)
Base address	BAR0: 16 MB BAR1: 256 MB BAR3: 32 MB
PCI class code	0x03 - Display Controller
PCI sub-class code	0x02 - 3D Controller (Compute)
ECC support	Supported (Enabled by default)
SMBus (8-bit address)	0x9E (write), 0x9F (read)
SMBus direct access	Supported
SMBPBI (SMBus Post Box Interface)	Supported
Operating systems	Microsoft Windows 7, Windows 8, Windows 8.1, and Windows 10+ Linux

Table 4 provides the environment conditions specifications for the Tesla P4 GPU Accelerator.

Table 4. Board Environmental and Reliability Specifications

Specification	Condition
Operating temperature	0 °C to 55 °C
Storage temperature	-40 °C to 75 °C
Operating humidity	5% to 90% RH
Storage humidity	5% to 95% RH
Mean time between failures (MTBF)	Uncontrolled environment: 1677428 hours at 35 °C Controlled environment: 2643325 hours at 35 °C

FORM FACTOR

The Tesla P4 board is a low profile passive board. Nominal dimensions are shown. For tolerances see the 2D mechanical drawings attached to this product specification. The dimensions are 68.58 mm, (2.7 inches) × 167.64 mm, (6.6 inches).




Figure 2. Tesla P4 Board Dimensions

AIRFLOW DIRECTION SUPPORT

The Tesla P4 board supports bidirectional airflow.

SUPPORT INFORMATION

CERTIFICATES AND AGENCIES

Certifications

- ▶ Windows Hardware Quality Lab (WHQL):
 - Certified Windows 7 and Windows 8
- ▶ 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 5. 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:

¹Windows 7, Windows 8, and Windows 8.1 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 regard 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 GPU Boost, Pascal, and Tesla 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

© 2017 NVIDIA Corporation. All rights reserved.