




TESLA M10

PB-08118-001_v04 | February 2017

Product Brief



DOCUMENT CHANGE HISTORY

PB-08118-001_v04

Version	Date	Authors	Description of Change
01	June 17, 2016	MD, SM	Initial Release
02	October 4, 2016	MD, SM	Updated “Support Information”
03	December 9, 2016	MD, SM	Updated boost specification (Table 1) and certifications
04	February 13, 2017	VK, SM	Updated 8-pin connector figure (Figure 5)

TABLE OF CONTENTS

- Overview 1**
 - Airflow Direction Support 2
- Specifications..... 4**
 - Product Specification..... 4
 - Thermal Specifications 6
- Design Discussion 7**
 - Form Factor 7
 - Power Connector Placement 8
 - Extenders 9
- Support Information..... 11**
 - Certificates and Agencies..... 11
 - Certifications 11
 - Agencies 11
 - Languages 12

LIST OF FIGURES

Figure 1.	Tesla M10 Board (With Optional I/O Bracket)	2
Figure 2.	Tesla M10: Left-to-Right Airflow	2
Figure 3.	Tesla M10: Right-to-Left Airflow	3
Figure 4.	Tesla M10 Board Dimensions (With Optional I/O Bracket).....	7
Figure 5.	PCIe 8-Pin Power Connector	8
Figure 6.	Straight Extender	9
Figure 7.	Long Offset Extender	9

LIST OF TABLES

Table 1.	Product Specifications	4
Table 2.	Memory Specifications	5
Table 3.	Software Specifications.....	5
Table 4.	Board Environmental and Reliability Specifications	5
Table 5.	Thermal Specifications	6
Table 6.	12V External Power Configuration	8
Table 7.	Languages Supported	12

OVERVIEW

The NVIDIA® Tesla® M10 is a dual-slot 10.5 inch PCI Express Gen3 graphics card with four mid-range NVIDIA Maxwell™ graphics processing units (GPUs). The Tesla M10 has 32 GB GDDR5 memory (8 GB per GPU) and a 225 W maximum power limit. The board is passively cooled and supports both airflow directions.

The Tesla M10 is designed to accelerate graphics in virtual desktop and application environments. The main feature of the Tesla M10 board is the support of the NVIDIA GRID™ software which includes NVIDIA GRID vGPU™. This technology enables the virtualization of physical GPUs into full-featured virtual GPUs providing maximum performance and scalability of up to 64 users per Tesla M10.

In addition, NVIDIA Maxwell-based Tesla doubles the H.264 encoding capabilities and improves encoding quality, which enables richer colors, preserve more details after video encoding, and results in a high-quality user experience.

Combined, these new capabilities result in an increased number of concurrent users in the same server space and a lower cost per user.




Figure 1. Tesla M10 Board (With Optional I/O Bracket)

AIRFLOW DIRECTION SUPPORT

The Tesla M10 board is passively cooled and supports airflows in both directions (left-to-right and right-to-left) with a single SKU.

- ▶ Tesla M10 with left-to-right airflow (Figure 2)
- ▶ Tesla M10 with right-to-left airflow (Figure 3)




Figure 2. Tesla M10: Left-to-Right Airflow




Figure 3. Tesla M10: Right-to-Left Airflow

SPECIFICATIONS

PRODUCT SPECIFICATION

Table 1 provides the product specifications for the Tesla M10 board.

Table 1. Product Specifications

Specification		Description
Product SKUs		P2405 SKU 70 (passive bidirectional airflow)
Total board power		225 W
GPU SKU		GM107-570-A2 ×4
NVIDIA® CUDA® cores		2560 (640 per GPU)
GPU clocks	Base	1033 MHz
	Maximum Boost	1306 MHz
	Idle	405 MHz
VBIOS	EEPROM size	4 Mbit
	UEFI	Supported
PCI Express interface		PCI Express Gen3 x16 system interface
Power connectors and headers		8-pin PCI Express power connector
Weight without extender and fans		963 grams (without bracket) 992 grams (with bracket)

Table 2 provides the memory specifications for the Tesla M10 board.

Table 2. Memory Specifications

Specification		Description
Memory clocks	Performance	2600 MHz
	Idle	405 MHz
Memory size	32 GB (8 GB per GPU)	
Memory I/O	128-bit ×4	
Memory configuration	32 pcs 512M ×16 GDDR5	
Memory bandwidth	83 GB/s ×4	

Table 3 provides the software specifications for the Tesla M10 board.

Table 3. Software Specifications

Specification	Description
Compatibility mode	Graphics
Base Address Registers (per GPU)	BAR0: 16 MB BAR1: 256 MB BAR2: 32 MB I/O BAR: 4 KB
PCI class code	0x03 - Display controller
PCI sub class codes	0x00 - VGA-compatible controller
ECC support	Not supported

Table 4 provides the environment conditions specifications for the Tesla M10 graphics board.

Table 4. Board Environmental and Reliability Specifications

Specification	Condition
Operating temperature	0 °C to 45 °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: TBD hours at 35 °C Controlled environment: TBD hours at 35 °C

THERMAL SPECIFICATIONS

Table 5 provides the thermal specifications for the Tesla M10 board.

Table 5. Thermal Specifications

Parameter	Value	Units
Total board power	225	W
GPU shutdown temperature	96	°C
GPU slowdown temperature	91	°C
GPU maximum operating temperature	90	°C
GPU hardware slowdown amount	50	%

DESIGN DISCUSSION

FORM FACTOR

The Tesla M10 board conforms to the NVIDIA Form Factor 2.0 specification. For details about the NVIDIA Form Factor 2.0 specification consult the *System Design Guide for NVIDIA Enterprise GPU Products (DG-07562-001)*.




Figure 4. Tesla M10 Board Dimensions (With Optional I/O Bracket)

POWER CONNECTOR PLACEMENT

The board provides an 8-pin PCI Express power connector on the East edge of the board.




Figure 5. PCIe 8-Pin Power Connector

Table 6. 12V External Power Configuration

Cable Attachment Configuration	Support	Notes
PCIe 8-pin cable	Supported; required	PCIe 8-pin cable must be attached. 8-pin cable must carry up to 150 W.
CPU 8-pin cable	Not supported	CPU 8-pin cable is not compatible with PCIe 8-pin board connector. Plugging in forcibly can cause permanent damage to the board and the system.
PCIe 6-pin cable	Not supported	A 6-pin connector carries insufficient current for the performance needs of this product. A 6-pin cable attached to an 8-pin board connector is an invalid configuration and is not supported.
No auxiliary power cable attached	Not supported	The auxiliary power cable must always be installed.

EXTENDERS

The Tesla M10 board provides two extender options as shown in the following figures.

- ▶ NVPN: 320-0867-003 – Straight extender (Figure 6)
 - Card + extender = 312 mm
- ▶ NVPN: 320-0866-003 – Long offset extender (Figure 7)
 - Card + extender = 339 mm




Figure 6. Straight Extender




Figure 7. Long Offset Extender

- ▶ 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 heat sink baseplate. The extender mounting holes will vary among NVIDIA products, so designing for flexibility is recommended.

SUPPORT INFORMATION

CERTIFICATES AND AGENCIES

Certifications

- ▶ Windows Hardware Quality Lab (WHQL):
 - Certified Windows 7, Windows 8.1, and Windows 10
 - Certified Windows Server 2008 R2, Windows Server 2012 R2, and Windows Server 2016
- ▶ 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:

¹Windows 7, Windows 8, Windows 8.1, Windows 10, Windows Server 2008 R2, Windows Server 2012 R2, and Windows Server 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 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, GRID vGPU, NVIDIA GRID, NVIDIA Maxwell, 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

© 2016, 2017 NVIDIA Corporation. All rights reserved.