




# TESLA M40

PB-07859-001\_v03 | January 2016

## Product Brief



## DOCUMENT CHANGE HISTORY

PB-07859-001\_v03

Version	Date	Authors	Description of Change
01	November 6, 2015	WT, SM	Initial Release
02	January 20, 2016	WT, SM	<ul style="list-style-type: none"><li>•Updated PCI Device ID in Table 1</li><li>•Updated Table 2 and Table 5</li><li>•Updated Windows in “Certifications”</li></ul>
03	January 22, 2016	WT, SM	Updated certifications and languages supported

# TABLE OF CONTENTS

- Overview ..... 1**
  - Airflow Direction Support ..... 2
- Specifications..... 3**
  - Product Specification..... 3
  - Power and Electrical Specifications ..... 5
- Design Discussion ..... 7**
  - Form Factor ..... 7
  - Power Brake Header Placement ..... 8
  - Power Connector Placement ..... 8
    - CPU 8-Pin to PCIe 8-Pin Dongle ..... 9
  - Extenders ..... 10
- Support Information..... 12**
  - Certificates and Agencies ..... 12
    - Certifications ..... 12
    - Agencies ..... 12
  - Languages ..... 13

## LIST OF FIGURES

Figure 1.	Tesla M40 Board (With Optional I/O Bracket) .....	2
Figure 2.	Tesla M40 Airflow Directions (With Optional I/O Bracket) .....	2
Figure 3.	Tesla M40 Board Dimensions (With Optional I/O Bracket).....	7
Figure 4.	Power Brake Header (With Optional I/O Bracket) .....	8
Figure 5.	CPU 8-Pin Power Connector.....	8
Figure 6.	CPU 8-Pin to PCIe 8-Pin Dongle.....	9
Figure 7.	Long Offset Extender .....	10
Figure 8.	Straight Extender .....	10

## LIST OF TABLES

Table 1.	Product Specifications .....	3
Table 2.	Memory Specifications .....	4
Table 3.	Software Feature Specifications .....	4
Table 4.	Board Environmental and Reliability Specifications .....	4
Table 5.	Input Power Measurements .....	5
Table 6.	Power Management Specifications.....	5
Table 7.	Peak Input Current .....	5
Table 8.	Hardware Threshold Current Limits .....	6
Table 9.	Supported Auxiliary Power Connections .....	9
Table 10.	Languages Supported .....	13

# OVERVIEW

The NVIDIA® Tesla® M40 is a dual-slot 10.5 inch PCI Express Gen3 card with a single NVIDIA Maxwell™ GM200 graphics processing unit (GPU). The Tesla M40 has 12 GB GDDR5 on-board memory and a 250 W maximum power limit. It uses a passive heat sink for cooling, which requires system air flow to properly operate the card within its thermal limits. The Tesla M40 is designed for scale-out deep learning training deployment.

For performance optimization, NVIDIA GPU Boost™ feature is supported. By adjusting the GPU clock dynamically, maximum performance is achieved within the power cap limit.

Tesla M40 boards are shipped with EDC and ECC enabled by default to protect the GPU's memory interface and the on-board DRAM memories. EDC protects the memory interface by detecting any single, double, and all odd-bit errors. The GPU will replay any memory transaction that have an EDC error until the data transfer is error-free. ECC protects the DRAM content by fixing any single-bit errors and detecting double-bit errors. There is no replay associated with ECC.

With ECC enabled, some of the memory is used for the ECC bits, so the user available memory is reduced by ~6.25%. On the Tesla M40 the total available memory with ECC on will be ~11.25 GB.




Figure 1. Tesla M40 Board (With Optional I/O Bracket)

## AIRFLOW DIRECTION SUPPORT

The Tesla M40 board employs a bidirectional heat sink, which accepts airflow from both left-to-right and right-to-left directions.




Figure 2. Tesla M40 Airflow Directions (With Optional I/O Bracket)

# SPECIFICATIONS

## PRODUCT SPECIFICATION

Table 1 provides the product specifications for the Tesla M40 board.

Table 1. Product Specifications

Specification		Description
Board SKU		PG600 SKU 202
Total board power		250 W (default) Maximum: 250 W
GPU SKU		GM200-895-A1
PCI Device ID		0x17FD
NVIDIA® CUDA® cores		3072
GPU clocks	Base	948 MHz
	Maximum boost	1114 MHz
	Idle	324 MHz
VBIOS	EEPROM size	4 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 One 2-pin power brake header
Weight	Board	898 Grams
	Bracket with screws	20 Grams
	Long offset extender	52 Grams
	Straight extender	42 Grams

Table 2 provides the memory specifications for the Tesla M40 graphics board.

**Table 2. Memory Specifications**

Specification	Description
Memory size	12 GB
Memory I/O	384-bit
Memory configuration	24 pcs 256M × 16 GDDR5
Peak memory bandwidth	Up to 288 GB/s

Table 3 provides the software feature specifications.

**Table 3. Software Feature Specifications**

Specification	Description
Base address	BAR0: 16 MB, non-prefetchable, 32-bit system address space BAR1: 16 GB, prefetchable, 64-bit system address space BAR2: 32 MB, prefetchable, 64-bit system address space I/O: Disabled
PCI class codes	PCI class: 0x03 (Display controller) PCI sub-class: 0x02 (3D controller)
ECC support	Supported (enabled by default)
SMBus (8-bit address)	0x9E (write), 0x9F (read)
SMBus direct access	Supported
SMBPBI (SMBus Post Box Interface)	Supported

Table 4 provides the environment conditions specifications for the Tesla M40 graphics board.

**Table 4. Board Environmental and Reliability Specifications**

Specification	Description
Operating temperature	0 °C to 45 °C
Storage temperature	-40 °C to 75 °C
Operating humidity	5% to 90% relative humidity
Storage humidity	5% to 95% relative humidity
Mean time between failures (MTBF)	Uncontrolled environment: 744126 hours at 35 °C Controlled environment: 961014 hours at 35 °C



## POWER AND ELECTRICAL SPECIFICATIONS

Table 5 provides the average current and power measurements for each input rail under worst-case total graphics power (TGP) test conditions for each configuration.

**Table 5. Input Power Measurements**

Parameter	Specification	Time Scale
PCI Express edge connector (12V ± 8%)	4.3 A	1 s
	52 W	1 s
PCI Express edge connector (3V3 ± 9%)	0.6 A	1 s
	2 W	1 s
CPU 8-pin power connector (12V ± 8%)	16.4 A	1 s
	197 W	1 s

Table 6 provides the power management specifications for the Tesla M40 board.

**Table 6. Power Management Specifications**

Parameter	Value	Units
Power capping threshold	250	W
Power capping sample time	30	ms
Power capping response time (typical)	100	ms

The Tesla M40 uses an on-board current and voltage monitoring device to limit the duration of current spikes that exceed a programmable hardware threshold (HT) limit. In addition, the maximum peak input current is also specified, based on characterization of worst-case transient workloads.

Table 7 provides the peak input current specifications for 12V input rails based on characterization of worst-case transient workloads.

**Table 7. Peak Input Current**

Input Rails	Value	Time Scale
PCI Express edge connector (12V)	7.2 A	1 ms
CPU 8-pin power connector (12V)	24.9 A	1 ms

Table 8 provides the HT limit and the maximum duration for the 12V inputs.

**Table 8. Hardware Threshold Current Limits**

<b>Input Rails</b>	<b>HT Limit</b>	<b>Maximum Duration</b>
PCI Express edge connector (12V)	8 A	1.5 ms
CPU 8-pin power connector (12V)	25 A	1.5 ms

# DESIGN DISCUSSION

## FORM FACTOR

The Tesla M40 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 3. Tesla M40 Board Dimensions (With Optional I/O Bracket)

## POWER BRAKE HEADER PLACEMENT

Figure 4 shows the placement of the power brake header connectors for the Tesla M40 board.




Figure 4. Power Brake Header (With Optional I/O Bracket)

## POWER CONNECTOR PLACEMENT

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




Figure 5. CPU 8-Pin Power Connector

Table 9. 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 <sup>1</sup> 1× PCIe 8-pin cable and 1× PCIe 6-pin cable <sup>1</sup>

**Note:**  
<sup>1</sup>Each PCIe 6-pin cable must be capable of carrying up to 50% of the average powers and tolerate 50% of the peak currents specified in Table 5 and Table 7

## CPU 8-Pin to PCIe 8-Pin Dongle

Figure 6 lists the pin assignments of the dongle.




Figure 6. CPU 8-Pin to PCIe 8-Pin Dongle

**Note:** The NVPN for the CPU 8-pin dongle is 030-0571-000hf.

## EXTENDERS

The Tesla M40 board provides two extender options as shown in the following figures.

- ▶ NVPN: 320-0866-003 – Long offset extender (Figure 7)
  - Card + extender = 339 mm
- ▶ NVPN: 320-0867-003 – Straight extender (Figure 8)
  - Card + extender = 312 mm




Figure 7. Long Offset Extender




Figure 8. Straight 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, Windows 10, Windows Server 2008 R2 and 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 Authority and Radio Spectrum Management Group of New Zealand (C-Tick)
- ▶ 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 10. Languages Supported

Languages	Windows <sup>1</sup>	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:**

<sup>1</sup>Windows 7, Windows 8, Windows 8.1, Windows 10, Windows Server 2008 R2, and Windows Server 2012 R2 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, 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

© 2015, 2016 NVIDIA Corporation. All rights reserved.