



TESLA M6

PB-07865-001_v02 | December 2016

Product Brief



DOCUMENT CHANGE HISTORY

PB-07865-001_v02

Version	Date	Authors	Description of Change
01	November 13, 2015	YI, SM	Initial Release
02	December 9, 2016	MD, SM	Updated default mode and certifications

TABLE OF CONTENTS

- Overview 1
- Specifications..... 3
 - Product Specification..... 3
- Design Discussion 5
 - Form Factor 5
 - MXM PCB Mounting Holes 6
- Compute and Graphics Modes 8
 - Compute Mode 8
 - Graphics Mode 9
- Support Information..... 10
 - Certificates and Agencies..... 10
 - Certifications 10
 - Agencies 10

LIST OF FIGURES

Figure 1.	Tesla M6 Board.....	2
Figure 2.	Tesla M6 Board Outline	5
Figure 3.	Mounting Holes	6

LIST OF TABLES

Table 1.	Product Specifications.....	3
Table 2.	Memory Specifications.....	4
Table 3.	Software Feature Specifications	4
Table 4.	Thermal Specifications	4
Table 5.	Tesla M6 Board Outline Specifications	6
Table 6.	Mounting Holes Specifications	7
Table 7.	Compute Mode Settings	8
Table 8.	Graphics Mode Settings.....	9

OVERVIEW

The NVIDIA® Tesla® M6 is an MXM 3.1 Type B card with a single NVIDIA Maxwell™ GM204 graphics processing unit (GPU). It has 8 GB GDDR5 on-board memory and a 100 W maximum power limit.

Tesla M6 is specifically designed to fit into constrained space available in blade servers. NVIDIA does not ship it with a cooling solution attached. However, it provides thermal specifications that OEMs can use to design their custom heat sinks.

Tesla M6 can function in two modes:

- ▶ **Compute:** Compute mode has large memory BAR and ECC enabled, making it suitable for single precision compute applications. ECC protects the DRAM content by fixing any single-bit errors and detecting double-bit errors.
- ▶ **Graphics:** By using the `gpumodeswitch` utility, Tesla M6 can be switched to graphics mode. This mode allows it to be compatible with the NVIDIA GRID™ software to use in virtualized graphics environments.




Figure 1. Tesla M6 Board

SPECIFICATIONS

PRODUCT SPECIFICATION

Table 1 provides the product specifications for the Tesla M6 board.

Table 1. Product Specifications

Specification		Description
Board SKU		P2754 SKU 200
Total board power		100 W (maximum)
GPU SKU		GM204-995
IDs		DEVID: 0x13F3 SSID: 0x1143
Form Factor		MXM 3.1 Type B
NVIDIA® CUDA® cores		1536
GPU clocks	Base	722 MHz (TGP: 75 W) 950 MHz (TGP: 100 W)
	Boost	886 MHz (TGP: 75 W) 1051 MHz (TGP: 100 W)
PCI Express interface		P0: Gen3 16 lanes, 8.0 Gbps P8: Gen1 16 lanes, 2.5 Gbps

Table 2 provides the memory specifications for the Tesla M6 graphics board.

Table 2. Memory Specifications

Specification	Description
Memory clock	2300 MHz
Memory size	8 GB
Memory I/O	256-bit
Memory configuration	16 pcs 256M × 16 GDDR5
Memory bandwidth	147.2 GB/s

Table 3 provides the software feature specifications.

Table 3. Software Feature Specifications

Specification	Description
EEPROM size	4 Mbit
PCI classcodes	PCI base class: 0x03 PCI sub-class: 0x02
ECC support	Supported
SMBPBI (SMBus Post Box Interface)	Supported

Table 4 provides the thermal specifications for the Tesla M6 graphics board.

Table 4. Thermal Specifications

Action	T _j (°C)
GPU shutdown temperature	91
GPU slowdown temperature	88
GPU maximum operating temperature	86
GPU hardware slowdown amount	50%

DESIGN DISCUSSION

FORM FACTOR

Tesla M6 follows the MXM 3.1 Type B mechanical specifications. For more details on the mechanical specifications, refer to the *MXM Electromechanical Specification Version 3.1*.




Figure 2. Tesla M6 Board Outline

Table 5. Tesla M6 Board Outline Specifications

Symbol	mm			inches		
	Minimum	Nominal	Maximum	Minimum	Nominal	Maximum
A1		82.00			3.228	
A2		70.00			2.756	
A3		105.00			4.134	
A4	3.37	3.50	3.63	0.133	0.138	0.143
A5	78.37	78.50	78.63	3.085	3.091	3.096
A6	3.87	4.00	4.13	0.152	0.157	0.163
A7	65.87	66.00	66.13	2.593	2.598	2.604
A8	38.87	39.00	39.13	1.530	1.535	1.541

MXM PCB MOUNTING HOLES

The module has 6 holes. Two are used to secure the board to the system and the other four to fasten the thermal solution to the module.




Figure 3. Mounting Holes

Table 6. Mounting Holes Specifications

Symbol	mm			inches		
	Minimum	Nominal	Maximum	Minimum	Nominal	Maximum
B1		75.00			2.953	
B2		8.25			0.325	
B3		54.25			2.136	
B4		5.50			0.217	
B5		51.50			2.028	
B6	3.07	3.20	3.33	0.121	0.126	0.131
B7	3.07	3.20	3.33	0.121	0.126	0.131

COMPUTE AND GRAPHICS MODES

Tesla M6 can be configured into compute or graphics mode (default).

COMPUTE MODE

Compute mode is optimized for high-performance compute (HPC) applications. Table 7 provides details of the compute mode settings.

Table 7. Compute Mode Settings

Setting	Value	Notes
Classcode	3D Controller	This classcode indicates to operating systems (OS) that the GPU is not intended for use as a primary display device.
Memory BAR	8 gigabytes	Tesla GPUs expose a large memory base address register (BAR) for direct access to the frame buffer from the CPU, and other PCI Express devices.
I/O base BAR	Disabled	The GPU need not consume any legacy I/O resources when used as a non-display device.
ECC protection	Enabled	Error Correcting Code (ECC) is enabled on the GPU frame buffer to protect against single- and multi-bit memory errors.

GRAPHICS MODE

While compute mode is optimal for HPC usage, it can cause compatibility problems with OS and hypervisors when the GPU is used primarily as a graphics device:

- ▶ Some OS require that the GPU advertise a VGA display controller classcode in order for the GPU to be used as a primary graphics device.
- ▶ Some hypervisors cannot support pass through of GPUs with large memory BARs to guest virtual machines.

To address these problems, Tesla M6 supports graphics mode for compatibility with NVIDIA GRID software. Table 8 provides details of the graphics mode settings.

Table 8. Graphics Mode Settings

Setting	Value	Notes
Classcode	VGA Controller	This classcode indicates to OS that the GPU can function as a primary display device.
Memory BAR	256 megabytes	The GPUs a smaller memory BAR for direct access to the frame buffer.
I/O base BAR	Enabled	The GPU exposes an I/O BAR to claim the resources required to operate as a VGA controller.
ECC protection	Disabled	ECC protection is disabled by default, though it can still be enabled by use of the <code>nvidia-smi</code> management tool

The mode of the GPU is established directly at power-on, from settings stored in the GPU's non-volatile memory. `gpumodeswitch` is used to program the mode of the GPU by updating the GPU's non-volatile memory settings.

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 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)

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 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

© 2015, 2016 NVIDIA Corporation. All rights reserved.