



TESLA K10 GPU ACCELERATOR

BD-06280-001_v07 | November 2012

Board Specification



DOCUMENT CHANGE HISTORY

BD-06280-001_v07

Version	Date	Authors	Description of Change
01	April 10, 2012	GG, SM	Preliminary Information (Information contained within this board specification is subject to change)
02	April 11, 2012	GG, SM	<ul style="list-style-type: none">•Updated board length•Updated "Overview" section•Updated Table 4
03	May 14, 2012	GG, SM	<ul style="list-style-type: none">•Removed "Preliminary Information" statement as this board specification is no longer considered preliminary•Updated title to <i>Tesla K10 GPU Accelerator Board Specification</i>•Updated number of core processors•Updated Figure 2 and Figure 3
04	May 21, 2012	GG, SM	Removed NDA and confidential statements from document
05	June 22, 2012	GG, SM	Updated the "Standard I/O Connector Placement" section with steps on removing the bracket
06	September 18, 2012	GG, SM	<ul style="list-style-type: none">•Added note to "Overview" section•Added peak memory bandwidth to "Key Features" section•Updated Table 1
07	November 7, 2012	GG, SM	Added MTBF information (Table 1)

TABLE OF CONTENTS

Overview	1
Key Features	1
Tesla K10 Block Diagram	3
Configuration	4
Mechanical Specifications	5
PCI Express System	5
Standard I/O Connector Placement	6
Internal Connectors and Headers	8
External PCI Express Power Connectors	8
Power Specifications	12
Support Information	13
Certificates and Agencies	13
Agencies	13
Languages	14

LIST OF FIGURES

Figure 1.	Tesla K10 Block Diagram	3
Figure 2.	Tesla K10 GPU Accelerator.....	5
Figure 3.	Tesla K10 Bracket	6
Figure 4.	Tesla K10 with Bracket	7
Figure 5.	Tesla K10 without Bracket	7
Figure 6.	6-Pin PCI Express Power Connector	9
Figure 7.	8-Pin PCI Express Power Connector	10

LIST OF TABLES

Table 1.	Board Configurations.....	4
Table 2.	6-Pin PCI Express Power Connector Pinout	11
Table 3.	8-Pin PCI Express Power Connector Pinout	11
Table 4.	Configuration with External PCI Express Connectors.....	12
Table 5.	Languages Supported	14

OVERVIEW

The NVIDIA® Tesla® K10 graphics processing unit (GPU) accelerator is a PCI Express, dual-slot, full height (4.376 inches by 10.5 inches by 1.52 inches) form factor computing module comprising two NVIDIA GK104 GPUs. The Tesla K10 offers a total of 8 GB of GDDR5 on-board memory (4 GB per GPU) and supports PCI Express Gen3.

The Tesla K10 can be configured by the OEM or by the end user to enable or disable ECC or error correcting codes that can fix single-bit errors and detect double-bit errors. Enabling ECC will cause some of the memory to be used for the ECC bits, so the user available memory will decrease by 10%. On the Tesla K10 the ECC protection is for DRAM only.



Note: The Tesla K10 can be installed and works in a PCI Express Gen2 system as well. In such configurations the Tesla K10 will run at PCI Express Gen2 speeds.

KEY FEATURES

GPU

The Tesla K10 GPU accelerator has two GK104 GPUs. Characteristics for both GPUs are as follows:

- ▶ Number of processor cores: 1536 per GPU
- ▶ Processor core clock: 745 MHz
- ▶ Package size: 40 mm × 40 mm 1745-pin ball grid array (BGA)

Board

- ▶ PCI Express Gen3 ×16 system interface
- ▶ Physical dimensions: 4.376 inches × 10.5 inches, dual-slot
- ▶ Board power: 225 W

External Connectors

- ▶ None

Internal Connectors and Headers

- ▶ One 6-pin PCI Express power connector
- ▶ One 8-pin PCI Express power connector

Memory

- ▶ Memory clock: 2.5 GHz
- ▶ Peak memory bandwidth: Cumulative 320 GB/s (160 GB/s per GPU)
- ▶ Interface: 256-bit
 - Total board memory: 8 GB (4 GB per GPU)
 - 32 pieces of 128M × 16 GDDR5, SDRAM (per GPU)

BIOS

- ▶ 2Mbit Serial ROM

TESLA K10 BLOCK DIAGRAM

Figure 1 is the block diagram for Tesla K10 GPU accelerator. It comprises two identical GK104 GPUs, each with 4 GB of GDDR5 memory. The GPUs are connected via a PCI Express switch. The board supports PCI Express Gen3.




Figure 1. Tesla K10 Block Diagram

CONFIGURATION

The Tesla K10 boards are available in the following configurations (Table 1) based on the orientation of the airflow inside the system.

Table 1. Board Configurations

Specifications	Tesla K10
Generic SKU reference	<ul style="list-style-type: none"> 900-22055-0010-000: Airflow left to right 900-22055-0020-000: Airflow right to left
Chip	2× GK104
Package size GPU	42.5 mm × 42.5 mm 1981-pin ball grid array (BGA)
Processor clock	745 MHz
Memory clock	2.5 GHz
Memory size	4 GB per GPU (8 GB per board)
Memory I/O	256-bit GDDR5
Memory configuration	32 pieces of 128M × 16 GDDR5 SDRAM
External connectors	None
Internal connectors and headers	<ul style="list-style-type: none"> 8-pin PCI Express power connector 6-pin PCI Express power connector
Board power	225 W
Meantime between failures (MTBF)	<ul style="list-style-type: none"> Controlled environment: 70506.47 hours at 35 °C Uncontrolled environment: 48076.4923 hours at 35 °C

MECHANICAL SPECIFICATIONS

PCI EXPRESS SYSTEM

The Tesla K10 boards (Figure 2) conform to the PCI Express full height (4.376 inches by 10.5 inches) form factor. Figure 2 is shown without the bracket.




Figure 2. Tesla K10 GPU Accelerator



Note: The final product will ship with an 8 and 6-pin connector and shroud.

STANDARD I/O CONNECTOR PLACEMENT

As shown in Figure 3, the Tesla K10 includes a vented bracket. If you are an OEM who qualifies for bracket modifications, you have the option of receiving your modules with no bracket installed.




Figure 3. Tesla K10 Bracket

If you need to remove the standard bracket follow these simple steps:

1. Remove the two shoulder screws on the back side of the PCB.
2. Remove the two flat head screws on the bracket exhaust face.
3. Remove the bracket.
4. Slide the washer in between the PCB and the backplate to maintain clearance between the PCB and the backplate.
5. Attach the shoulder screws.




Figure 4. Tesla K10 with Bracket



Figure 5. Tesla K10 without Bracket

INTERNAL CONNECTORS AND HEADERS

The Tesla K10 GPU accelerator supports the following internal connectors and headers.

- ▶ 8-pin PCI Express power connector (can be used with a 6-pin power cable)
- ▶ 6-pin PCI Express power connector

External PCI Express Power Connectors

The Tesla K10 GPU accelerator is a performance optimized, high-end product and uses power from the PCI Express connector as well as external power connectors. The board can be used in two different ways.

- ▶ One 8-pin PCI Express power connector or
- ▶ Two 6-pin PCI Express power connectors

Figure 6 and Figure 7 show the specifications and Table 2 and Table 3 show the pinouts for the 6-pin and 8-pin PCI Express power connectors.




Figure 6. 6-Pin PCI Express Power Connector




Figure 7. 8-Pin PCI Express Power Connector

Table 2. 6-Pin PCI Express Power Connector Pinout

Pin Number	Description
1	+12 V
2	+12 V
3	+12 V
4	GND
5	Sense
6	GND

Table 3. 8-Pin PCI Express Power Connector Pinout

Pin Number	Description
1	+12 V
2	+12 V
3	+12 V
4	Sense1
5	GND
6	Sense0
7	GND
8	GND

POWER SPECIFICATIONS

The Tesla K10 GPU accelerator requires power from the PCI Express connector as well as one or two auxiliary power connectors.

Table 4. Configuration with External PCI Express Connectors

Connector Type	6-Pin Power Connector	Supported	Notes
8-pin connected	6-pin connected	Yes	
8-pin connected	No cable installed	Yes	8-pin cable must supply 175 W
6-pin connected	N/A	No	6-pin cable in the 8-pin connector is not supported.
Not installed	N/A	No	8-pin connector should always be connected.



Note: Detailed information about power draw by rail is available to authorized system partners in the *Tesla K10 Board System Design Guide* (DG-06105-001)

SUPPORT INFORMATION

CERTIFICATES AND AGENCIES

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)
- ▶ Voluntary Control Council for Interference (VCCI)

LANGUAGES

Table 5. Languages Supported

	Windows Server 2008 and Windows Server 2008 R2	Linux
English (US)	X	X
English (UK)	X	
Arabic	X	
Chinese, Simplified	X	
Chinese, Traditional	X	
Danish	X	
Dutch	X	
Finnish	X	
French	X	
French (Canada)	X	
German	X	
Italian	X	
Japanese	X	
Korean	X	
Norwegian	x	
Portuguese (Brazil)	X	
Russian	X	
Spanish	X	
Spanish (Latin America)		
Swedish	X	
Thai	X	

Note: NVIDIA's CUDA™ software is only supported in English (U.S.)

Notice

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.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication of otherwise under any patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all other information previously supplied. NVIDIA Corporation products are not authorized as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

HDMI

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

ROVI Compliance Statement

NVIDIA Products that support Rovi Corporation's Revision 7.1.L1 Anti-Copy Process (ACP) encoding technology can only be sold or distributed to buyers with a valid and existing authorization from ROVI to purchase and incorporate the device into buyer's products.

This device is protected by U.S. patent numbers 6,516,132; 5,583,936; 6,836,549; 7,050,698; and 7,492,896 and other intellectual property rights. The use of ROVI Corporation's copy protection technology in the device must be authorized by ROVI Corporation and is intended for home and other limited pay-per-view uses only, unless otherwise authorized in writing by ROVI Corporation. Reverse engineering or disassembly is prohibited.

OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

Trademarks

NVIDIA, the NVIDIA logo, CUDA, 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

© 2012 NVIDIA Corporation. All rights reserved.