



TESLA K20X GPU ACCELERATOR

BD-06397-001_v05 | November 2012

Board Specification



DOCUMENT CHANGE HISTORY

BD-06397-001_v05

Version	Date	Authors	Description of Change
01	July 31, 2012	GG, SM	Preliminary Information (Information contained within this document is subject to change)
02	September 6, 2012	GG, SM	<ul style="list-style-type: none">•Updated “Key Features” section•Updated Table 1•Updated “Standard I/O Connector Placement” section
03	September 7, 2012	GG, SM	<ul style="list-style-type: none">•Updated Table 4•Updated Table 5•General edits through specification
04	October 9, 2012	GG, SM	<ul style="list-style-type: none">•Added memory bandwidth•Updated idle power to 25 W
05	November 9, 2012	GG, SM	<ul style="list-style-type: none">•Removed “Preliminary Information” as this board specification is now final•Removed “NVIDIA Confidential from document•Updated PCI Express to Gen2•Updated package size•Updated board power (Table 1)•Updated block diagram (Figure 1)

TABLE OF CONTENTS

Overview	1
Key Features	1
Tesla K20X Block Diagram	2
Configuration	3
Mechanical Specifications	4
PCI Express System	4
Tesla K20X Bracket	5
Power Connectors	6
Power Specifications	9
Support Information	10
Certificates and Agencies	10
Agencies	10
Languages	11

LIST OF FIGURES

Figure 1.	Tesla K20X Block Diagram	2
Figure 2.	Tesla K20X GPU Accelerator	4
Figure 3.	Tesla K20X Bracket	5
Figure 4.	6-Pin PCI Express Power Connector	6
Figure 5.	8-Pin PCI Express Power Connector	7

LIST OF TABLES

Table 1.	Board Configuration	3
Table 2.	6-Pin PCI Express Power Connector Pinout	8
Table 3.	8-Pin PCI Express Power Connector Pinout	8
Table 4.	Auxiliary Power Connectors	9
Table 5.	Languages Supported	11

OVERVIEW

The NVIDIA® Tesla® K20X graphics processing (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 of a single GK110 GPU. The Tesla K20X is designed for servers and offers a total of 6 GB of GDDR5 on-board memory and supports PCI Express Gen2.

The Tesla K20X 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 K20X, the register files, cache and DRAM are ECC protected.

KEY FEATURES

GPU

- ▶ Number of processor cores: 2688
- ▶ Processor core clock: 732 MHz
- ▶ Package size: 45 mm × 45 mm 2397-pin ball grid array (S-FCBGA)

Board

- ▶ PCI Express Gen2 ×16 system interface
- ▶ Physical dimensions: 4.376 inches × 10.5 inches, dual-slot

Display Connectors

- ▶ None

Power Connectors

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

Memory

- ▶ Memory clock: 2.6 GHz
- ▶ Memory bandwidth: 250 GB/sec
- ▶ Interface: 384-bit
 - Total board memory: 6 GB
 - 24 pieces of 64M × 16 GDDR5, SDRAM

BIOS

- ▶ 2Mbit Serial ROM

TESLA K20X BLOCK DIAGRAM

Figure 1 is the block diagram for the Tesla K20X GPU dual-slot computing processor module.

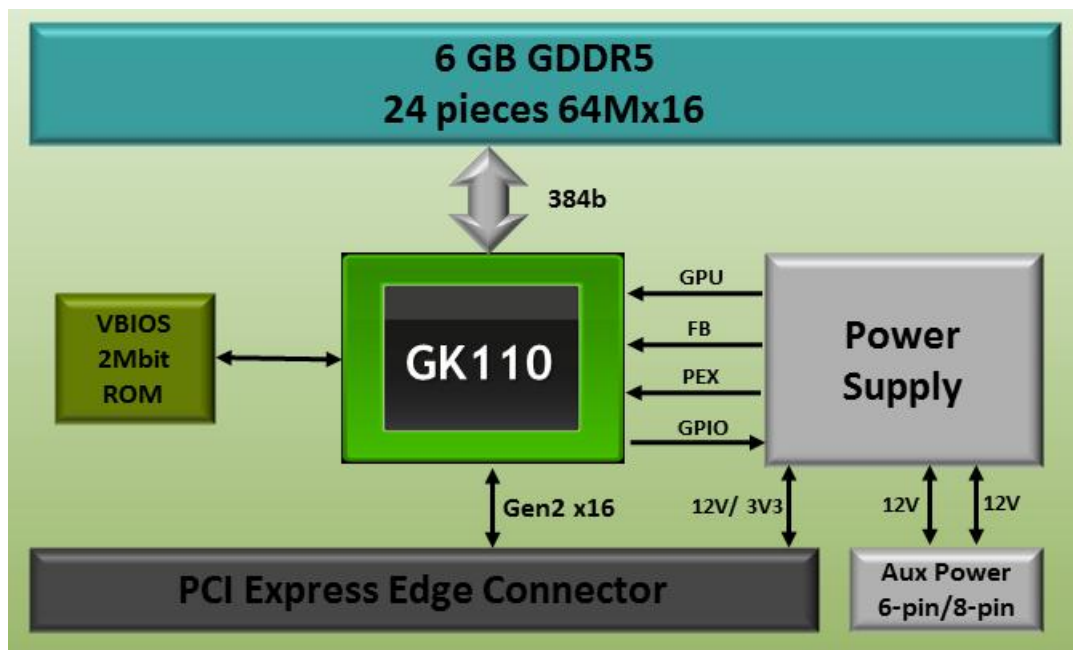


Figure 1. Tesla K20X Block Diagram

CONFIGURATION

The Tesla K20X board is available in the following configuration (Table 1).

Table 1. Board Configuration

Specifications	Tesla K20X
Generic SKU reference	699-22081-0200-xxx
Chip	GK110
Package size GPU	45 mm × 45 mm 2397-pin S-FCBGA
Processor clock	732 MHz
Memory clock	2.6 GHz
Memory size	6 GB
Memory I/O	384-bit GDDR5
Memory configuration	24 pieces of 64M ×16 GDDR5 SDRAM
Display connectors	None
Power connectors	<ul style="list-style-type: none"> •8-pin PCI Express power connector •6-pin PCI Express power connector
Board power	235 W
Idle power	25 W
Thermal cooling solution	Passive heat sink

MECHANICAL SPECIFICATIONS

PCI EXPRESS SYSTEM

The Tesla K20X board (Figure 2) conforms to the PCI Express full height (4.376 inches by 10.5 inches) form factor.




Figure 2. Tesla K20X GPU Accelerator



Note: The final product will ship with shroud.

TESLA K20X BRACKET

As shown in Figure 3, the Tesla K20X includes a vented bracket. If you are an OEM who qualifies for bracket modifications, you have the option of receiving your module with no bracket installed.



Figure 3. Tesla K20X Bracket

POWER CONNECTORS

The Tesla K20X GPU accelerator is a performance optimized, high-end product and uses power from the PCI Express connector as well as external power connectors.

Figure 4 and Figure 5 show the specifications and Table 2 and Table 3 show the pinouts for the 6-pin and 8-pin PCI Express power connectors.




Figure 4. 6-Pin PCI Express Power Connector




Figure 5. 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 K20X GPU accelerator requires power from the PCI Express connector as well as one or two auxiliary power connectors.

Table 4. Auxiliary Power Connectors

8-Pin Header	6-Pin Header	Support	Notes
Connect 8-pin cable	Connect 6-pin cable	Yes	
Connect 8-pin cable	No cable installed	Yes	8-pin cable must supply 150 W
Connect 6-pin cable	Connect 6-pin cable	No	8-pin connector should always be connected



Note: Detailed information about power draw by rail will be available to authorized system partners in the *Tesla K20X Board System Design Guide*.

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)	X	
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.

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.