



TESLA K20 GPU ACTIVE ACCELERATOR

BD-06499-001_v04 | January 2013

Board Specification



DOCUMENT CHANGE HISTORY

BD-06499-001_v04

Version	Date	Authors	Description of Change
01	October 9, 2012	MC, GG, SM	Preliminary Information (Information contained within this document is subject to change)
02	November 9, 2012	GG, SM	<ul style="list-style-type: none">•Removed "Preliminary Information" as this board specification is no longer preliminary•Removed "NVIDIA confidential"•Updated Table 1•Updated "Power Specifications" section
03	December 7, 2012	GG, SM	Updated 6-pin auxiliary power connector section
04	January 14, 2013	GG, SM	Added MTBF data to Table 1

TABLE OF CONTENTS

Overview	1
Key Features	2
Configuration	3
Mechanical Specifications	4
PCI Express System	4
Tesla K20 Active Bracket	5
Power Connectors	6
Auxiliary 6-Pin PCI Express Power Connector	6
Extenders	7
Power Specifications	10
Power by Rail	10
Energy Star Report	11
Thermal Specifications	12
Cooling Solution	12
Support Information	14
Certificates and Agencies	14
Agencies	14
Languages	15

LIST OF FIGURES

Figure 1.	Tesla K20 Active GPU Accelerator Board	1
Figure 2.	Tesla K20 GPU Active	4
Figure 3.	Tesla K20 Active Bracket	5
Figure 4.	Auxiliary 6-Pin PCI Express Power Connector	6
Figure 5.	Long Extender with Offset	8
Figure 6.	Flat Extender	8
Figure 7.	Short Extender with Offset	9
Figure 8.	Active Fan Sink	12

LIST OF TABLES

Table 1.	Board Configuration	3
Table 2.	6-Pin PCI Express Power Connector Pinout	7
Table 3.	PCI Express Power Connectors	10
Table 4.	Power by Rail	10
Table 5.	Energy Star Report	11
Table 6.	Fan Environmental Specifications and Conditions	13
Table 7.	Languages Supported	15

OVERVIEW

The NVIDIA® Tesla® K20 graphics processing unit (GPU) active accelerator board is a PCI Express, dual-slot full height (4.376 inches by 10.5 inches) form factor computing module comprising of a single GK110 GPU. The Tesla K20 active accelerator is designed for workstations and offers a total of 5 GB of GDDR5 on-board memory and supports PCI Express Gen2.

The Tesla K20 can be configured by the OEM or by the end user to enable or disable ECC 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 K20, the register files, cache, and DRAM are ECC protected.

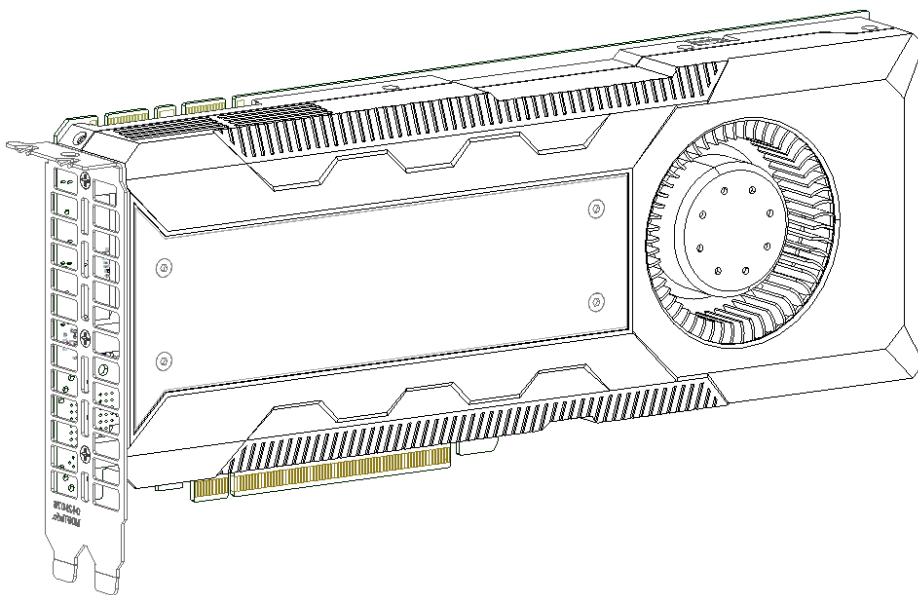


Figure 1. Tesla K20 Active GPU Accelerator Board

KEY FEATURES

GPU

- ▶ Number of processor cores: 2496
- ▶ Processor core clock: 706 MHz
- ▶ Package size: 45 mm × 45 mm 2397-pin flip chip 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

- ▶ Two 6-pin PCI Express power connectors

Memory

- ▶ Memory clock: 2.6 GHz
- ▶ Memory bandwidth: 208 GB/sec
- ▶ Interface: 320-bit
 - Total board memory: 5 GB
 - 20 pieces of 128M × 16 GDDR5, SDRAM

BIOS

- ▶ 2Mbit Serial ROM

CONFIGURATION

The Tesla K20 GPU active accelerator board is available in the following configuration (Table 1).

Table 1. Board Configuration

Specifications	Tesla K20
Generic SKU reference	699-22081-0204-xxx
Chip	GK110
Package size GPU	45 mm × 45 mm 2397-pin S-FCBGA
Processor clock	706 MHz
Memory clock	2.6 GHz
Memory size	5 GB
Memory I/O	320-bit GDDR5
Memory configuration	20 pieces of 128M ×16 GDDR5 SDRAM
Display connectors	None
Power connectors	Two 6-pin PCI Express power connectors
Board power	~225 W
Thermal cooling solution	Active fan sink
Form factor	Tesla Form Factor B (10.5 in length)
Weight without extender	826 grams
Available extenders	<ul style="list-style-type: none"> • Long Extender with Offset • Short Extender with Offset • Flat Extender (short)
Meantime between failures (MTBF)	<ul style="list-style-type: none"> • Uncontrolled environment: 138732 hours at 35 °C • Controlled environment: 222269 hours at 35 °C

MECHANICAL SPECIFICATIONS

PCI EXPRESS SYSTEM

The Tesla K20 GPU active accelerator board (Figure 1) conforms to the PCI Express full height (4.376 inches by 10.5 inches) form factor.




Figure 2. Tesla K20 GPU Active

TESLA K20 ACTIVE BRACKET

As shown in Figure 3, the Tesla K20 includes a vented bracket.



Figure 3. Tesla K20 Active Bracket

POWER CONNECTORS

The Tesla K20 GPU active accelerator board supports two 6-pin PCI Express power connectors on the "east end" of the board.

Auxiliary 6-Pin PCI Express Power Connector

The Tesla K20 GPU active accelerator board is a high powered board solution and it utilizes power from the PCI Express connector as well as the external power (two 6-pin PCI Express power connectors are required for the Tesla K20). Figure 4 shows the specification, and Table 5 shows the pinout for the 6-pin external PCI Express power connector.

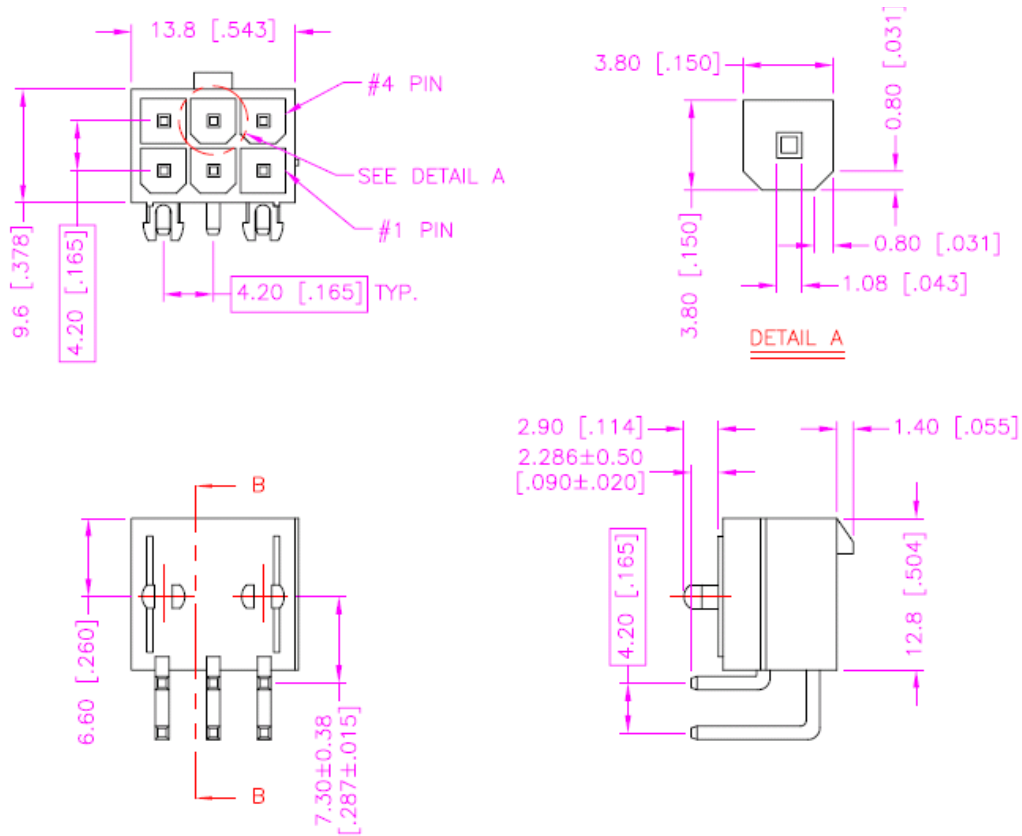


Figure 4. Auxiliary 6-Pin PCI Express Power Connector

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

EXTENDERS

The Tesla K20 GPU active accelerator board provides three extender options as shown in the following figures.

- ▶ NVPN: 320-0866-000 – Long Extender with Offset (Figure 5)
 - Card + Extender = 13.338 inches
- ▶ NVPN: 320-0867-000 – Flat Extender (Figure 6)
 - Card + Extender = 12.283 inches (312 mm) full length
- ▶ NVPN: 320-0868-000 – Short Extender with Offset (Figure 7)
 - Card + Extender = 12.283 inches (312 mm) full length



Figure 5. Long Extender with Offset



Figure 6. Flat Extender




Figure 7. Short Extender with Offset

- ▶ Ensure the system can accommodate standard mounting of “Type B” form factor boards using the NVIDIA extenders.
- ▶ 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 Tesla products, so designing for flexibility is recommended.

POWER SPECIFICATIONS

The Tesla K20 GPU active accelerator board requires power from the PCI Express connector as well as one or two auxiliary power connectors.

Table 3. PCI Express Power Connectors

6-Pin PCIe Power Connector #1	6-Pin PCIe Power Connector #2	Notes
6-pin cable connected	6-pin cable connected	Supported
6-pin cable connected	Not connected	Not supported. Both 6-pin power connectors must always be connected.
Not connected	6-pin cable connected	Not supported. Both 6-pin power connectors must always be connected.
Not connected	Not connected	Not supported. Both 6-pin power connectors must always be connected.

POWER BY RAIL

Table 4 lists the average power (in watts) by rail numbers for the Tesla K20 GPU active graphics while running DGEMM.

Table 4. Power by Rail

PEX12V	PEX3V3	EXT12V	EXT12V	Total Board Power
45 W	2 W	84 W	81 W	212 W

ENERGY STAR REPORT

Table 5 lists the energy star report for the Tesla K20 GPU active graphics board. The Energy Star measurements have been done using NBody simulation.

Table 5. Energy Star Report

Specification	Description
Brand	Tesla K20
GPU	GK110
Board	P2081
SKU	204
Clocks	<ul style="list-style-type: none"> •706 MHz (core) •2.6 GHz (memory)
Frame buffer	320-bit
Memory	GDDR5 SGRAM
PEX	×16
Idle power: ASPM OFF	13.9 W
Idle power: ASPM ON	14.1 W
Maximum power: ASPM OFF	143.55 W
Maximum power: ASPM ON	144.2 W

THERMAL SPECIFICATIONS

COOLING SOLUTION

NVIDIA has designed an active fan sink (Figure 8) to cool the GPU, memories and power components. For fan and environmental specifications refer to Table 6.




Figure 8. Active Fan Sink

Table 6. Fan Environmental Specifications and Conditions

Specifications	Conditions
Rated voltage	12 V
Operating voltage	10.8 to 12.6 V DC
Rated current	0.86 Amp (maximum 1.80 Amp)
Rated power	10.32 W (maximum 21.60 W)
Speed	5000 RPM
Life expectancy	50,000 hours continuous operation at 55 °C with 15 to 65% RH
Operating temperature	-10 °C to 70 °C
Storage temperature	-40 °C to 75 °C
Operating humidity	5% to 90% RH
Storage humidity	5% to 95% RH

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 7. Languages Supported

	Windows 7 (64-Bit), 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, 2013 NVIDIA Corporation. All rights reserved.