




TESLA K8 GPU ACTIVE ACCELERATOR

BD-07228-001_v03 | September 2014

Board Specification



DOCUMENT CHANGE HISTORY

BD-07228-001_v03

Version	Date	Authors	Description of Change
01	May 1, 2014	HV, SM	Preliminary Information (Information contained within this board specification is subject to change)
02	July 2, 2014	GG, SM	<ul style="list-style-type: none">•Updated "Overview"•Updated Figure 2
03	September 16, 2014	GG, SM	<ul style="list-style-type: none">•Removed all preliminary and confidential markings. This doc is no longer preliminary•Updated GPU minimum clock speed

TABLE OF CONTENTS

- Overview 1**
 - Key Features 1
 - NVIDIA GPU Boost 2
 - NVIDIA GPU Boost on Tesla K8 Active 2
 - API for NVIDIA GPU Boost on Tesla 3
 - Tesla K8 Block Diagram 5
 - Configuration 6
- Mechanical Specifications 7**
 - PCI Express System 7
 - Standard I/O Connector Placement 8
 - Internal Connectors and Headers 8
 - External PCI Express Power Connectors 8
- Power Specifications 10**
- Support Information 11**
 - Certificates and Agencies 11
 - Agencies 11
 - Languages 12

LIST OF FIGURES

Figure 1.	Tesla K8 Block Diagram.....	5
Figure 2.	Tesla K8 GPU Active Accelerator.....	7
Figure 3.	Tesla K8 Bracket.....	8
Figure 4.	6-Pin PCI Express Power Connector.....	9

LIST OF TABLES

Table 1.	nvidia-smi Commands.....	4
Table 2.	Board Configuration.....	6
Table 3.	6-Pin PCI Express Power Connector Pinout.....	9
Table 4.	External PCI Express Power Connector Configuration.....	10
Table 5.	Languages Supported.....	12

OVERVIEW

The NVIDIA® Tesla® K8 graphics processing unit (GPU) accelerator is a PCI Express, single-slot, full height (4.376 inches by 9.5 inches by 0.737 inches), active cooled computing module. It offers the benefits, reliability, and quality of Tesla GPU accelerators in a small form factor and at a low power budget. The Tesla K8 provides GPGPU acceleration in servers and embedded devices that are power or space constrained at the node or rack level. The Tesla K8 comprises of an NVIDIA GK104 GPU, offers 8 GB of GDDR5 on-board memory, and supports PCI Express Gen2.

The Tesla K8 can be configured to enable or disable error correcting codes (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 K8 the ECC protection is for DRAM only.



Note: This active cooler (fan) in the Tesla Kit is designed for use inside servers and not meant for workstations. The acoustics are 49 dBA at 100% fan speed.

KEY FEATURES

GPU

- ▶ Number of processor cores: 1536
- ▶ Core clocks:
 - Base clock: 693 MHz
 - Boost clock: 811 MHz
 - Minimum clock: 445 MHz
- ▶ Package size: 40 mm × 40 mm 1745-pin ball grid array (BGA)

Board

- ▶ PCI Express Gen2 ×16 system interface
- ▶ Physical dimensions: 4.376 inches × 9.5 inches, single-slot
- ▶ Board power: 100 W

Display Connectors

- ▶ None

Internal Connectors and Headers

- ▶ One 6-pin PCI Express power connector

Memory

- ▶ Memory clock: 2.5 GHz
- ▶ Peak memory bandwidth: 160 GB/s
- ▶ Interface: 256-bit
 - Total board memory: 8 GB
 - 16 pieces of 256M × 16 GDDR5 SDRAM

BIOS

- ▶ 2Mbit serial ROM
- ▶ BAR1 size: 256 MB



Note: The Tesla K8 core clock is set to the base clock by default. Boost clocks can be selected using NVML or NVSMI. Refer to NVML/NVSMI documentation for details.

NVIDIA GPU BOOST

NVIDIA GPU Boost™ makes use of available power headroom to run the core clock to a higher frequency. Application workloads that have power headroom can run at higher GPU clocks to boost application performance. Refer to the *NVIDIA GPU Boost for NVIDIA Professional Products Application Note (DA-06317-001)* for more a more detailed description.

NVIDIA GPU Boost on Tesla K8 Active

NVIDIA GPU Boost for the Tesla K8 Active is optimized to deliver robust and deterministic performance for a wide range of workloads. The Tesla K8 Active gives

end-users full control to manage core clock frequency and power for their workload and system.

The workload or system may have one or more of the following characteristics:

- ▶ Problem set is spread across multiple GPUs and requires periodic synchronization.
- ▶ Problem set spread across multiple GPUs running independently of each other.
- ▶ Workload has “compute spikes.” For example: some portions of the workload are extremely compute intensive, pushing the power higher, and some portions are moderately compute intensive.
- ▶ Workload is compute intensive through-out without any spikes.
- ▶ Workload requires fixed clocks and is sensitive to clocks fluctuating during the execution.
- ▶ Workload runs in a cluster where all GPUs need to start, finish, and run at the same clocks.
- ▶ Peak workload or end user requires predictable performance and repeatable results.
- ▶ Datacenter is used to run different types of workload at different hours in a day to better manage the power consumption.
- ▶ Some boards in cluster have access to better cooling than others.
- ▶ Nodes in the cluster have a limited power budget and need to restrict the power limit of the Tesla K8 boards.

By default the Tesla K8 Active ships with the core clock set to the base clock and the power limit set to 100 W. Some workloads in image and signal processing, etc., may have power headroom and can take advantage of the boost clock or can run at the minimum clock with the power limit set to 70 W.



Note: The memory clock remains constant at 2.5 GHz. Effective memory bandwidth utilization may change with core clock frequency.

When selecting one of the supported clocks, a good strategy is to characterize the workload with the available clocks. For example, SGEMM is extremely demanding on power. Therefore, the minimum clock may be the correct choice when running SGEMM or when a lower power limit is required.

API for NVIDIA GPU Boost on Tesla

The Tesla K8 Active gives end-users full control over the core clock frequency and power limit via NVML or `nvidia-smi`. NVML is a C-based API for monitoring and managing the various states of Tesla products. `nvidia-smi` is a command line utility on top of NVML. NVML documentation is available at

<https://developer.nvidia.com/nvidia-management-library-nvml>. `nvidia-smi` documentation is available at <https://developer.nvidia.com/nvidia-system-management-interface>.

Table 1 gives a summary of the `nvidia-smi` commands for using NVIDIA GPU Boost on the Tesla K8 Active.

Table 1. `nvidia-smi` Commands

Usage	Command
View the clocks the Tesla board supports	<code>nvidia-smi -q -d SUPPORTED_CLOCKS</code>
Set one of the supported clocks	<code>nvidia-smi -ac <MEM clock, Graphics Clock></code>
Enable persistent mode (make the clock settings persistent across driver unload)	<code>nvidia-smi -pm 1</code>
Disable persistent mode (Revert clock settings to base clocks after driver unloads)	<code>nvidia-smi -pm 0</code>
View current clock	<code>nvidia-smi -q -d CLOCK</code>
Reset clocks back to the default base clock (as specified in the board specification)	<code>nvidia-smi -rac</code>

When using non-default applications docks, driver persistence mode should be enabled. Persistence mode ensures that the driver stays loaded even when no NVIDIA® CUDA® or X applications are running on the CPU. This maintains current state, including requested applications clocks. If persistence mode is not enabled, and no applications are using the CPU, the driver will unload and any current user settings will revert back to default for the next application. To enable persistence mode run `'sudo nvidia-smi -pm 1'`

TESLA K8 BLOCK DIAGRAM

Figure 1 is the block diagram for Tesla K8 Active GPU accelerator. It comprises of a GK104 GPU with 8 GB of GDDR5 memory, power supply, and VBIOS ROM. The board supports PCI Express Gen2.

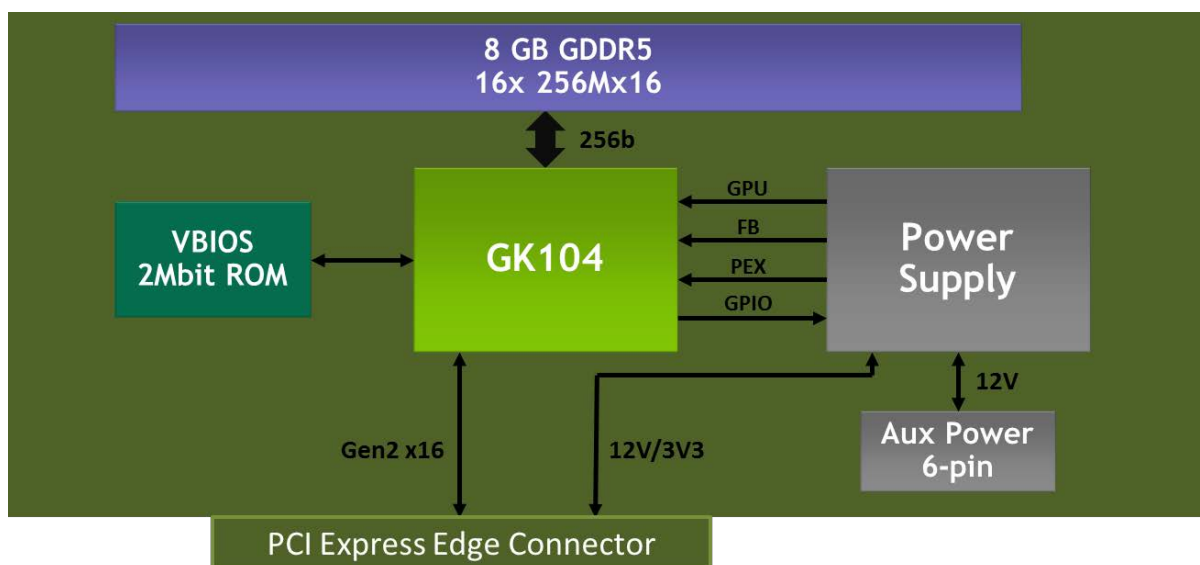


Figure 1. Tesla K8 Block Diagram

CONFIGURATION

The Tesla K8 Active board is available in the following configuration.

Table 2. Board Configuration

Specifications	Tesla K8
Generic SKU reference	699-22004-0200-xxx
Chip	GK104
Package size GPU	40 mm × 40 mm 1745-pin ball grid array (BGA)
Processor clock	<ul style="list-style-type: none"> •Base clock: 693 MHz •Boost clock: 811 MHz •Minimum clock: 445 MHz
Memory clock	2.5 GHz
Memory size	8 GB
Memory I/O	256-bit GDDR5
Memory configuration	16 pieces of 256M × 16 GDDR5 SDRAM
Display connectors	None
Power connectors	6-pin PCI Express power connector
Board power	100 W
Idle power	16 W
Thermal Cooling Solution	Active fan sink
Dimensions (LxHxW)	4.376 inches by 9.5 inches by 0.737 inches
Meantime between failures (MTBF)	<ul style="list-style-type: none"> •Controlled environment: 1116405.1002 hours at 35 °C •Uncontrolled environment: 949848.9651 hours at 35 °C
ASPM	Off
BAR1	256 MB

MECHANICAL SPECIFICATIONS

PCI EXPRESS SYSTEM

The Tesla K8 boards (Figure 2) conform to the PCI Express full height (4.376 inches by 9.5 inches) form factor.



Figure 2. Tesla K8 GPU Active Accelerator

STANDARD I/O CONNECTOR PLACEMENT

As shown in Figure 3, the Tesla K8 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 to be added at a later revision

Figure 3. Tesla K8 Bracket

INTERNAL CONNECTORS AND HEADERS

The Tesla K8 Active supports one 6-pin PCI Express power connector.

External PCI Express Power Connectors

The Tesla K8 Active GPU accelerator is a performance optimized product and uses power from the PCI Express connector as well as an external power connector.

Figure 4 shows the specifications and Table 3 shows the pinout for the 6-pin PCI Express power connector.




Figure 4. 6-Pin PCI Express Power Connector

Table 3. 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

POWER SPECIFICATIONS

The Tesla K8 Active GPU accelerator requires power from the PCI Express connector as well as an auxiliary power connector.

Table 4. External PCI Express Power Connector Configuration

6-Pin Power Connector	Supported	Notes
6-pin connected	Yes	
No cable installed	No	6-pin connector should always be connected



Note: Detailed information about power draw by rail is available to authorized system partners in the Tesla K8 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)		
Swedish	X	
Thai	X	

Note: CUDA software is only supported in English (U.S.)

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, 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

© 2014 NVIDIA Corporation. All rights reserved.