TESLA K40 GPU ACCELERATOR

BD-06902-001_v06 | October 2014

Board Specification

DOCUMENT CHANGE HISTORY

BD-06902-001_v06

Version	Date	Authors	Description of Change
01	July 3, 2013	GG, SM	Preliminary Information (Information contained within this board specification is subject to change)
02	August 1, 2013	GG, SM	•Updated product name •Updated core clocks speeds •Updated block diagram in Figure 1
03	September 19, 2013	GG, SM	•Added new section: "NVIDIA GPU Boost on Tesla K40" •Updated "Power Specifications" section •General updates throughout the board specification •Added board environmental conditions
04	September 25, 2013	GG, SM	Updated boost clock
05	November 11, 2013	GG, SM, DV	•Updated MTBF numbers •Added Figure 1: K40 Passive Board
06	October 15, 2014	GG, SM	Updated Figure 3

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

OVERVIEW

The NVIDIA® Tesla® K40 graphics processing unit (GPU) is a PCI Express, dual-slot computing module in the Tesla (267 mm length) form factor comprised of a single GK110B GPU. The Tesla K40 is designed for servers and offers a total of 12 GB of GDDR5 on-board memory and supports PCI Express Gen3. The Tesla K40 uses a passive heat sink for cooling.

Tesla K40 boards ship with ECC enabled by default protecting the register files, cache and DRAM. With ECC enabled, some of the memory is used for the ECC bits, so the user available memory is reduced by ~6.25%. On the Tesla K40 the total available memory with ECC turned on will be ~11.25 GB.

Figure 1. Tesla K40 Passive Board

KEY FEATURES

GPU

- Number of processor cores: 2880
- Core clocks
  - Base clock: 745 MHz
  - Boost clocks: 810 MHz and 875 MHz
- Package size: 45 mm × 45 mm 2397-pin ball grid array (S-FCBGA)

Note: All boards ship with core clock set to the base clock value. Boost clocks can be selected using NVML or NVSMI. Refer to the NVML/NVSMI documentation for more details.

Board

- PCI Express Gen3 ×16 system interface
- Physical dimensions: 111.15 mm (height) × 267 mm (length), dual-slot

Thermal Solution

- Passive heat sink

Display Connectors

- None

Power Connectors

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

Memory

- Memory clock: 3.0 GHz
- Memory bandwidth 288 GB/sec
- Interface: 384-bit
  - Total board memory: 12 GB
  - 24 pieces of 256M × 16 GDDR5, SDRAM

BIOS

- 2Mbit serial ROM
- BAR1 size: 16 GB

NVIDIA GPU BOOST ON TESLA K40

NVIDIA GPU Boost™ is a feature available on Tesla K40. It makes use of any power headroom to run the core clock to a higher frequency. Application workloads that have power headroom can run at high GPU clocks to boost application performance.

Note: The memory clock remains constant at 3 GHz. It's likely that the effective memory bandwidth utilization will change depending on the core clock frequency.

NVIDIA GPU Boost for HPC Workloads

NVIDIA GPU Boost for Tesla K40 is optimized to deliver a robust and deterministic boost behavior for a wide range of HPC workloads.

Tesla K40 gives full control to end-users to select the core clock frequency that fits their workload the best. The workload 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 and runs independent 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 moderate.
- 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.
- 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 a cluster have access to better cooling than others.

By default the Tesla K40 ships with the core clock set to the base clock. HPC workloads can have one or more characteristics as described. When selecting one of the supported boost clocks a good strategy is to characterize the workload with the available boost clocks. For example, DGEMM/Linpack are extremely demanding on power. Therefore, the “base clock” may be the correct choice when running Linpack. Some workloads in life sciences, manufacturing, CFD, CAD, etc., may have power headroom and can take advantage of one of the boost clocks.

API FOR NVIDIA GPU BOOST ON TESLA

Tesla K40 gives full control to end-users to select the core clock frequency via NVML or `nvidia-smi`. NVML is a C-based API for monitoring and managing the various states of Tesla products. It provides a direct access to submit queries and commands via `nvidia-smi`. NVML documentation is available at https://developer.nvidia.com/nvidia-management-library-nvml

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

Table 1. nvidia-smi Commands

Usage	Command
View the clocks the Tesla board supports	nvidia-smi –q –d SUPPORTED_CLOCKS
Set one of the supported clocks	nvidia-smi -ac <MEM clock, Graphics clock>
Make the clock settings persistent across driver unload	nvidia-smi -pm 1
Make the clock settings revert to base clocks after driver unloads (or turn off the persistent mode)	nvidia-smi -pm 0
To view the clock in use, use the command	nvidia-smi -q –d CLOCK
To reset clocks back to the base clock (as specified in the board specification)	nvidia-smi –rac
To allow “non-root” access to change graphics clock	nvidia-smi -acp 0

When using non-default applications clocks, 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 GPU. This maintains current state, including requested applications clocks. If persistence mode is not enabled, and no applications are using the GPU, 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`'.

The driver will attempt to maintain requested applications clocks whenever a CUDA context is running on the GPU. However, if no contexts are running the GPU will revert back to idle clocks to save power and will stay there until the next context is created. Thus, if the GPU is not busy, you may see idle current clocks even though requested applications clocks are much higher.

Note: By default changing the application clocks requires root access. If the user does not have root access, the user can request his or her cluster manager to allow non-root control over application clocks. Once changed, this setting will persist for the life of the driver before reverting back to root-only defaults. Persistence mode should always be enabled whenever changing application clocks, or enabling non-root permissions to do so.

TESLA K40 BLOCK DIAGRAM

Figure 2 is the block diagram for the Tesla K40 GPU dual-slot computing processor module.

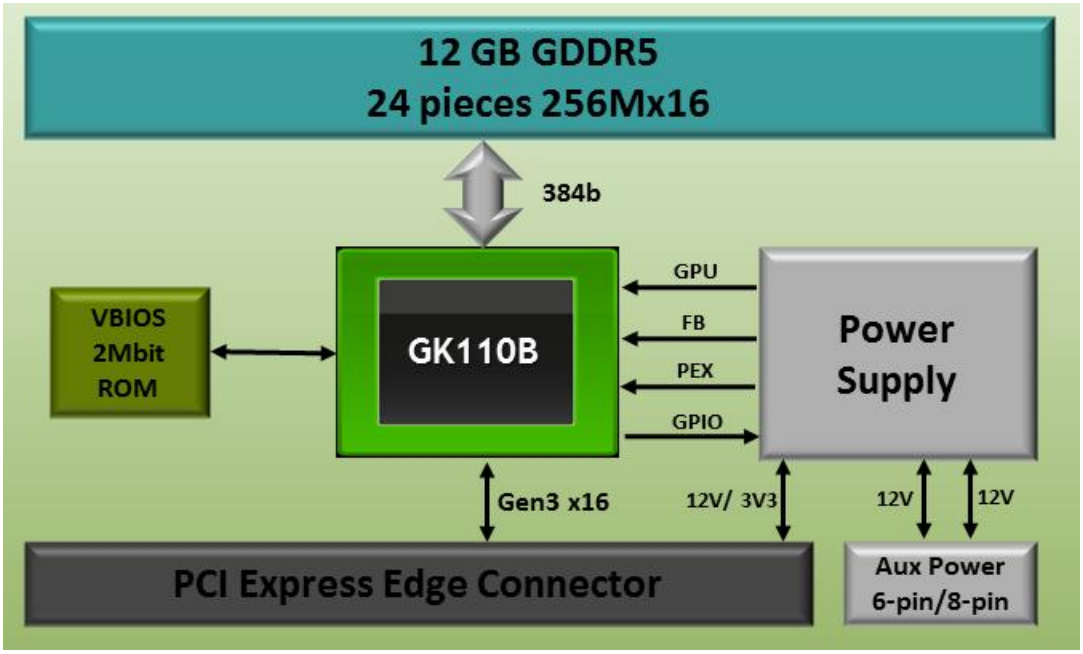

Figure 2. Tesla K40 Block Diagram

ENVIRONMENTAL CONDITIONS

Table 2 lists the environmental operating and storage conditions for the Tesla K40 board.

Table 2. Board Environmental Conditions

Specifications	Conditions
Operating temperature	0 °C to 45 °C
Storage temperature	-40 °C to 75 °C
Operating humidity	5% to 90% RH
Storage humidity	5% to 95% RH

CONFIGURATION

The Tesla K40 board is available in the following configuration.

Table 3. Board Configuration

Specifications	Tesla K40
Generic SKU reference	699-22081-0202-xxx
Chip	GK110B
Core clocks	•Base clock: 745 MHz •Boost clocks: 810 MHz and 875 MHz
Memory clock	3.0 GHz
Memory size	12 GB
Memory I/O	384-bit GDDR5
Memory configuration	24 pieces of 256M × 16 GDDR5 SDRAM
Display connectors	None
Power connectors	•8-pin PCI Express power connector •6-pin PCI Express power connector
Board power	235 W
Power cap level	235 W
BAR1 size	16 GB
Extender support	Straight extender is the default and the long offset extender is available as an option.
Hockey stick defeat	Not supported
Idle power	16 W
Thermal cooling solution	Passive heat sink
Mean time between failures (MTBF)	•Controlled environment: 282,847 hours at 35 °C •Uncontrolled environment: 252,222 hours at 35 °C
ASPM	Off

MECHANICAL SPECIFICATIONS

PCI EXPRESS SYSTEM

The Tesla K40 board (Figure 3) conforms to the PCI Express full height form factor.

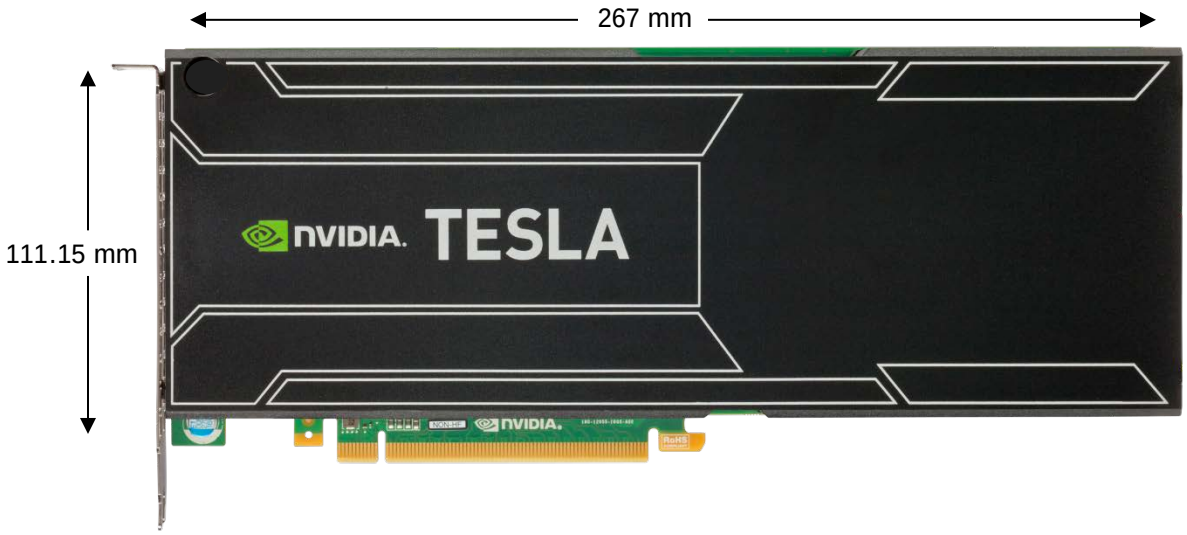

Figure 3. Tesla K40 GPU Accelerator

TESLA K40 BRACKET

As shown in Figure 4, the Tesla K40 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 4. Tesla K40 Bracket

POWER CONNECTORS

The Tesla K40 GPU accelerator is a performance optimized, high-end product and uses power from the PCI Express connector as well as external power connectors.

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

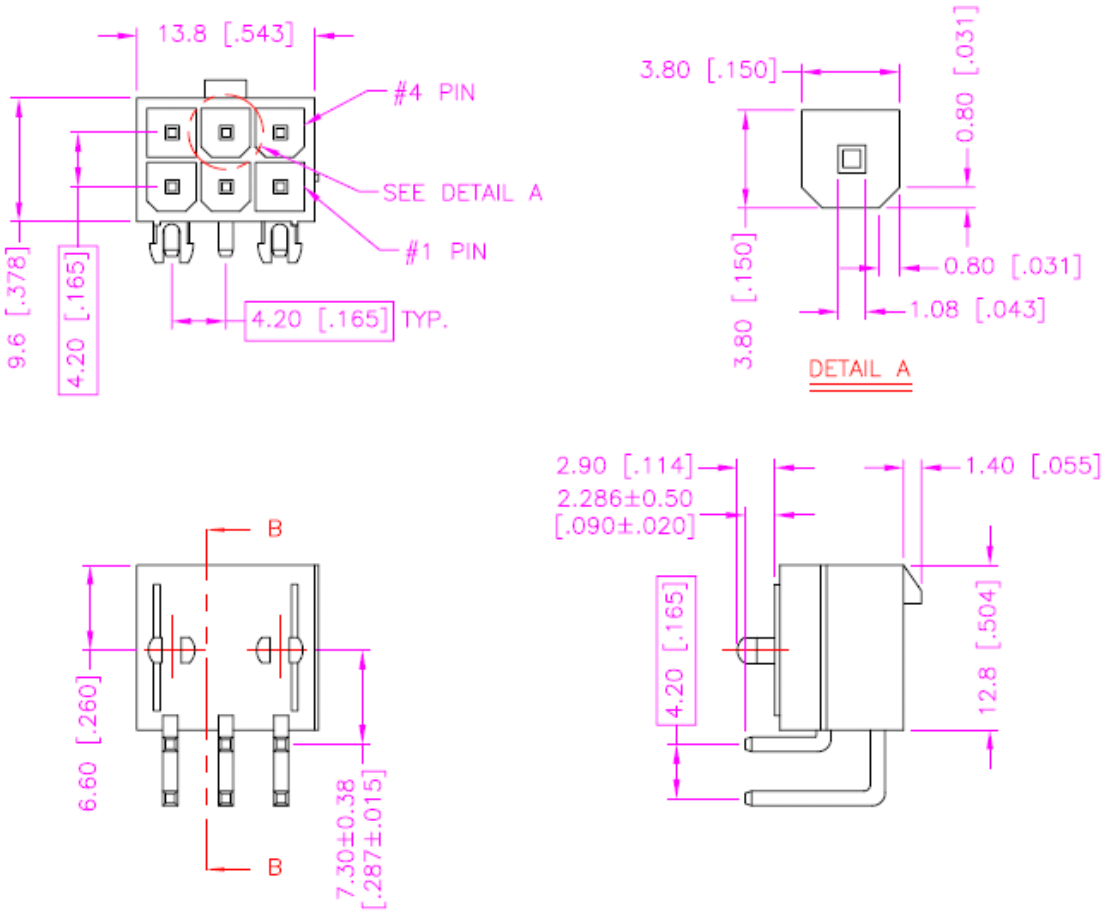

Figure 5. 6-Pin PCI Express Power Connector

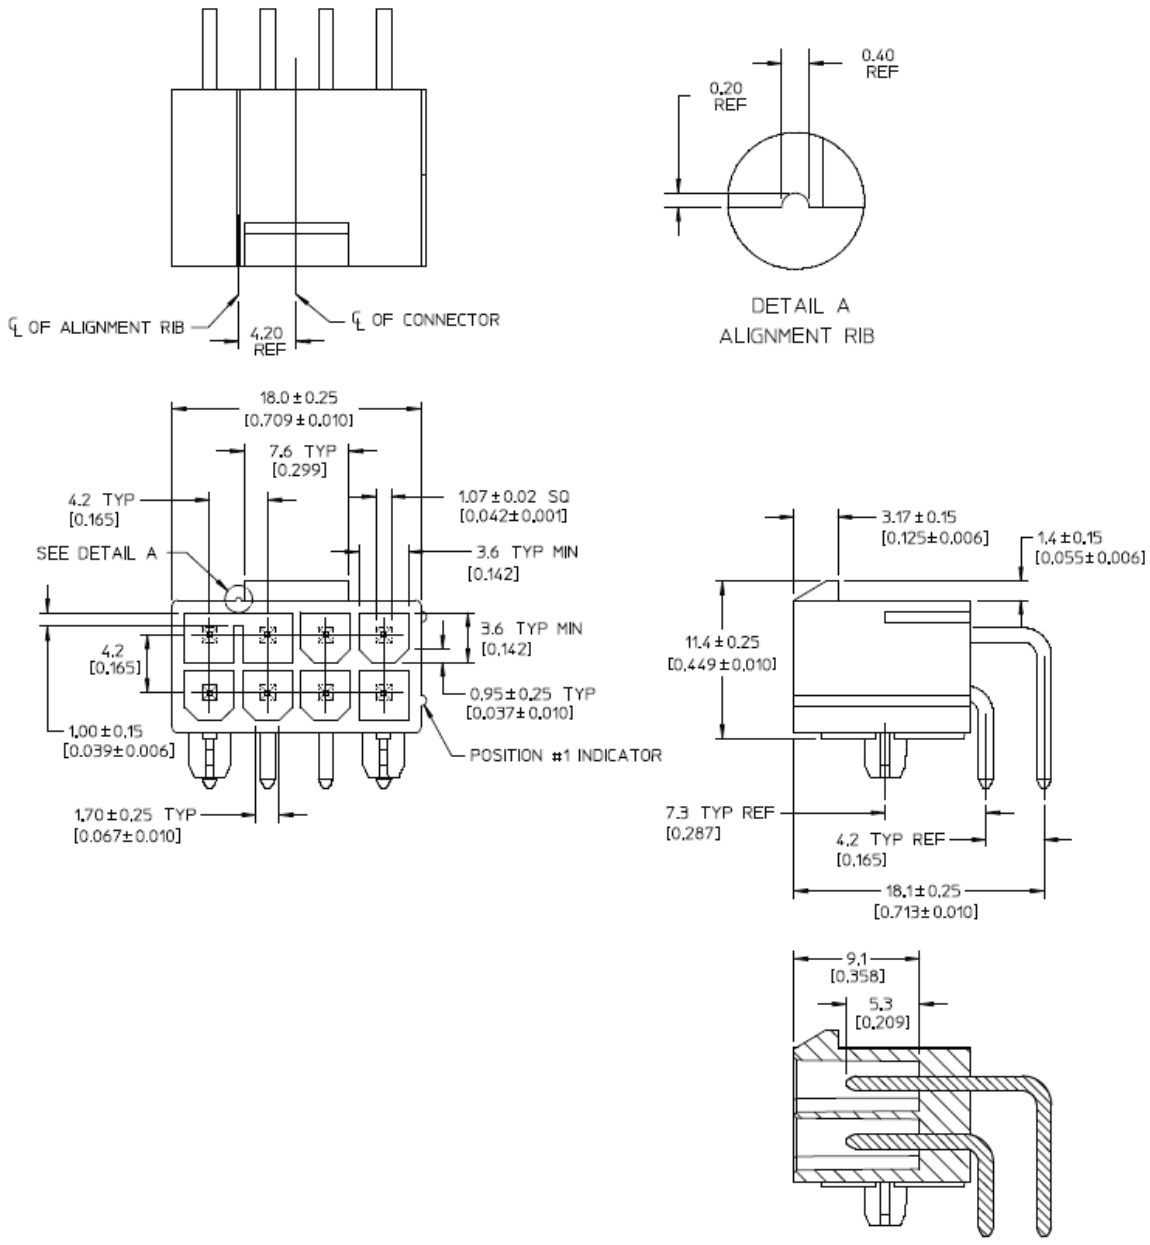

Figure 6. 8-Pin PCI Express Power Connector

Table 4. 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 5. 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 K40 GPU accelerator requires power from the PCI Express connector as well as one or two auxiliary power connectors.

Table 6. 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 175 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 K40 system design guide.

Table 7 provides the power requirements used in thermal and power measurements for the Tesla K40.

Table 7. Power Requirements

Voltage Rail (Volts)	Voltage Tolerance (Minimum)	Voltage Tolerance (Maximum)	Maximum Currents (Amps)
3.3	-8%	+8%	1.0
12	-8%	+8%	19.6

Note:
System power qualification with the Tesla cards should be done with the Thermal Design Power (TDP) application provided by NVIDIA.
The peak current values are characterized over a 1 ms time interval, with 5-sigma confidence. These are values based on characterization data using the TDP application under TDP test conditions. Peak current values may be higher with applications that consume more power than the TDP application.

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 8. 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: 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