



**PNY TESLA® K10**  
GPU COMPUTING ACCELERATOR

**PART NUMBER:**

Contact us

**TESLA K10 GPU COMPUTING ACCELERATOR**

Optimized for single precision applications, the Tesla K10 is a throughput monster based on the ultra-efficient GK104 Kepler GPU. The accelerator board features two GK104 GPUs and delivers up to 2x the performance for single precision applications compared to the previous generation Fermi-based Tesla M2090 in the same power envelope. With an aggregate performance of 4.58 teraflop peak single precision and 320 gigabytes per second memory bandwidth for both GPUs put together, the Tesla K10 is optimized for computations in seismic, signal, image processing, and video analytics.

With teraflops of single and double precision performance, NVIDIA® Kepler GPU Computing Accelerators are the world's fastest and most efficient high performance computing (HPC) companion processors. Based on the Kepler compute architecture, which is 3 times higher performance per watt than the previous "Fermi" compute architecture, the Tesla Kepler GPU Computing Accelerators make hybrid computing dramatically easier, and applicable to a broader set of computing applications. NVIDIA Tesla GPUs deliver the best performance and power efficiency for seismic processing, biochemistry simulations, weather and climate modeling, image, video and signal processing, computational finance, computational physics, CAE, CFD, and data analytics.

**The innovative design of the Kepler compute architecture includes:**

>> SMX (streaming multiprocessor) design that delivers up to 3x more performance per watt compared to the SM in Fermi. It also delivers 1 petaflop of computing in just 10 server racks.

>> Dynamic Parallelism capability that enables GPU threads to automatically spawn new threads. By adapting to the data without going back to the CPU, it greatly simplifies parallel programming and enables GPU acceleration of a broader set of popular algorithms, like adaptive mesh refinement (AMR), fast multipole method (FMM), and multigrid methods.

>> Hyper-Q feature that enables multiple CPU cores to simultaneously utilize the CUDA cores on a single Kepler GPU, dramatically increasing GPU utilization, slashing CPU idle times, and advancing programmability. Ideal for cluster applications that use MPI.

**TESLA K10 - PRODUCT SPECIFICATIONS**

CUDA PARALLEL PROCESSING CORES	3072 (2 x 1536)
PROCESSOR CORE CLOCK	745 MHz
PEAK DOUBLE PRECISION FLOATING POINT PERFORMANCE (BOARD)	0.19 teraflops
PEAK SINGLE PRECISION FLOATING POINT PERFORMANCE (BOARD)	4.58 teraflops
NUMBER OF GPUS	2 x GK104s
PROCESSOR CORE CLOCK	706 MHz
FRAME BUFFER MEMORY	8 GB
MEMORY BANDWIDTH	320 GBytes/sec
EXTERNAL CONNECTORS	None
INTERNAL CONNECTORS AND HEADERS	(1) 6-pin PCI Express power connector (1) 8-pin PCI Express power connector
GRAPHICS BUS	PCI Express 2.0 x16
FORM FACTOR	110 mm (H) x 265 mm (L) Dual Slot
BOARD POWER	225 W
THERMAL SOLUTION	Passive



Tesla GPU Computing Accelerator Common Features

<b>ECC MEMORY ERROR PROTECTION</b>	Meets a critical requirement for computing accuracy and reliability in datacenters and supercomputing centers. External Both external and internal memories are ECC protected in Tesla K20.
<b>SYSTEM MONITORING FEATURES</b>	Integrates the GPU subsystem with the host system's monitoring and management capabilities such as IPMI or OEM-proprietary tools. IT staff can thus manage the GPU processors in the computing system using widely used cluster/grid management solutions.
<b>L1 AND L2 CACHES</b>	Accelerates algorithms such as physics solvers, ray-tracing, and sparse matrix multiplication where data addresses are not known beforehand
<b>ASYNCHRONOUS TRANSFER WITH DUAL DMA ENGINES</b>	Turbocharges system performance by transferring data over the PCIe bus while the computing cores are crunching other data.
<b>FLEXIBLE PROGRAMMING ENVIRONMENT WITH BROAD SUPPORT OF PROGRAMMING LANGUAGES AND APIS</b>	Choose OpenACC, CUDA toolkits for C, C++, or Fortran to express application parallelism and take advantage of the innovative Kepler architecture.

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

**BOARD**

- >> PCI Express Gen3 x16 system interface
- >> Physical dimensions: 4.376 inches x 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 x 16 GDDR5, SDRAM (per GPU)

**BIOS**

- >> 2Mbit Serial ROM

Software and Drivers

>> **Software applications page:**

[http://www.nvidia.com/object/vertical\\_solutions.html](http://www.nvidia.com/object/vertical_solutions.html)

>> Tesla GPU computing accelerators are supported for both Linux and Windows.

Server modules are only supported on 64-bit OSes and workstation / desktop modules are supported for 32-bit as well.

>> **Drivers:**

Latest drivers can be downloaded from

<http://www.nvidia.com/drivers>

- Learn more about Tesla data center management tools at <http://www.nvidia.com/object/softwarefor-tesla-products.html>

>> Software development tools are available at <http://developer.nvidia.com/getting-started-parallelcomputing>

