TESLA™ S2050
GPU COMPUTING SYSTEM
SUPERCOMPUTING AT 1/10TH THE COST

Based on the new NVIDIA CUDA™ GPU architecture codenamed "Fermi," the Tesla™ S2050 1U Computing Systems are designed from the ground up for high performance computing.

The Tesla S2050 Computing System delivers "must have" features for the technical and enterprise computing space including ECC memory for uncompromised accuracy and scalability, and 7X the double precision performance compared Tesla 10-series GPU computing products. Compared to typical quad-core CPUs, Tesla 20-series computing systems deliver equivalent performance at 1/10th the cost and1/20th the power consumption.

Designed with four Fermi-based Tesla computing processors in a standard 1U chassis, the Tesla S2050 computing system scales to solve the world's most important computing challenges – more quickly and accurately.

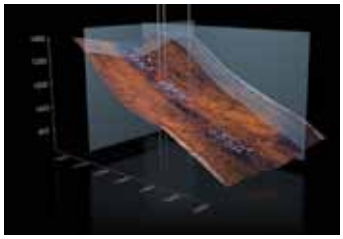

OIL & GAS

SCIENCE

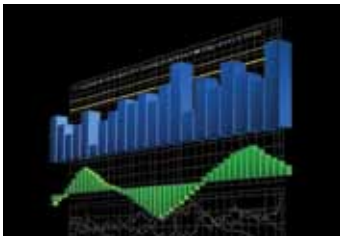

FINANCE

TECHNICAL SPECIFICATIONS

FORM FACTOR
> 1U

OF TESLA GPUS
> 4

GPU MEMORY SPEED
> 1.55 GHz

GPU MEMORY INTERFACE
> 384-bit

GPU MEMORY BANDWIDTH
> 148 GB/sec

DOUBLE PRECISION FLOATING POINT PERFORMANCE (PEAK)
> 2 Tflops

SINGLE PRECISION FLOATING POINT PERFORMANCE (PEAK)
> 4.13 Tflops

TOTAL DEDICATED MEMORY*
> 12GB GDDR5

POWER CONSUMPTION (TYPICAL)
> 900W TDP

SYSTEM INTERFACE
> PCIe x16 Gen2

SOFTWARE DEVELOPMENT TOOLS
> CUDA C/C++/Fortran, OpenCL, DirectCompute Toolkits. NVIDIA Parallel Nsight™ for Visual Studio

*Note: With ECC on, a portion of the dedicated memory is used for ECC bits, so the available user memory is reduced by 12.5%. (e.g. 3 GB total memory yields 2.625 GB of user available memory.)

FEATURES AND BENEFITS

GPUs POWERED BY THE FERMI-GENERATION OF THE CUDA ARCHITECTURE	Delivers cluster performance at 1/10th the cost and 1/20th the power of CPU-only systems based on the latest quad core CPUs.
448 CUDA CORES	Delivers up to 515 Gigaflops of double-precision peak performance in each GPU, enabling 2 Teraflops of double precision performance in a 1U of space. Single precision peak performance is over a Teraflop per GPU.
ECC MEMORY	Meets a critical requirement for mission critical applications with uncompromised computing accuracy and reliability. Offers protection of data in memory to enhance data integrity and reliability for applications. Register files, L1/L2 caches, shared memory, and DRAM all are ECC protected.
SYSTEM MONITORING FEATURES	Simplifies management and remote monitoring post-installation via NVSMI. Status lights on the front and rear of the unit ensures IT staff can see the status whether they are on the either side of the rack.
UP TO 6GB OF GDDR5 MEMORY PER GPU	Maximizes performance and reduces data transfers by keeping larger data sets in local memory that is attached directly to the GPU. Tesla S2050 includes 3GB/GPU.
NVIDIA PARALLEL DATACACHE™	Accelerates algorithms such as physics solvers, ray-tracing, and sparse matrix multiplication where data addresses are not known beforehand. This includes a configurable L1 cache per Streaming Multiprocessor block and a unified L2 cache for all of the processor cores.
NVIDIA GIGATHREAD™ ENGINE	Maximizes the throughput by faster context switching that is 10X faster than previous architecture, concurrent kernel execution, and improved thread block scheduling.
ASYNCHRONOUS TRANSFER	Turbocharges system performance by transferring data over the PCIe bus while the computing cores are crunching other data. Even applications with heavy data-transfer requirements, such as seismic processing, can maximize the computing efficiency by transferring data to local memory before it is needed.
CUDA PROGRAMMING ENVIRONMENT WITH BROAD SUPPORT OF PROGRAMMING LANGUAGES AND APIs	Choose C, C++, OpenCL, DirectCompute, or Fortran to express application parallelism and take advantage of the "Fermi" GPU's innovative architecture.
HIGH SPEED, PCIE GEN 2.0 DATA TRANSFER	Maximizes bandwidth between the host system and the Tesla processors. Enables Tesla systems to work with virtually any PCIe-compliant host system with an open PCI-E slot (x8 or x16).

DRIVERS AND DOWNLOADS

- Tesla S2050 is supported under
  - Linux 32-bit and 64-bit
  - Windows XP, Windows Vista, Windows 7 (32-bit and 64-bit)
- Vertical Solutions and Software page: www.nvidia.com/object/vertical_solutions.html

SUPPORT

- HARDWARE SUPPORT
  For RMA requests, replacements and warranty issues regarding your NVIDIA based product, please contact the reseller that you purchased it from.
- KNOWLEDGEBASE
  Our knowledgebase is available online 24x7x365 and contains answers to the most common questions and issues: www.nvidia.custhelp.com/cgi-bin/nvidia.cfg/php/enduser/std_alp.php
- USER FORUMS
  Discuss Tesla products, talk about CUDA development, and share interesting issues, tips and solutions with your fellow NVIDIA Tesla users on the CUDA discussion forums: www.forums.nvidia.com

To learn more about NVIDIA Tesla, go to **www.nvidia.com/tesla**

NVIDIA