

NVIDIA Tesla



NVIDIA Tesla GPU Computing Solutions for HPC

Revolutionary NVIDIA® Tesla™ high performance computing solutions put personal supercomputing into the hands of individual scientists and engineers by expanding the capability of any workstation or system with the power of GPU computing. Scientific and technical professionals now have an incredible opportunity to expand their ability to solve problems previously impossible with current computing approaches.

Parallel Performance

Tesla computing solutions enable users to process large datasets with a massively multi-threaded computing architecture. By developing a parallel architecture from the ground up, NVIDIA has designed its new Tesla computing products to meet the requirements of HPC software. Exclusive computing features

include a Thread Execution Manager to coordinate the concurrent execution of thousands of computing threads and a Parallel Data Cache enabling computing threads to share data easily, delivering results in less time.

C for the GPU

The world's only C-language development environment for the GPU, the NVIDIA CUDA™ software development kit includes a standard C compiler, hardware debugger tools, and a performance profiler for simplified application development.

Developer Community

NVIDIA is the catalyst for the largest GPU computing developer community. NVIDIA's interactive, on-line GPU developer community provides access to forums, educational materials, and additional resources and tools.

Compatible Solutions

As an industry-standard solution, Tesla easily fits into existing HPC environments. Available products include a Tesla C870 GPU computing processor for users to upgrade their existing workstation, a Tesla D870 desktide GPU computing system to add additional performance alongside a workstation, and a Tesla S870 GPU computing system for deployment within an enterprise data center. Used in tandem with multi-core CPU systems, Tesla solutions provide a flexible computing platform that runs on both Microsoft® Windows® and Linux® operating system environments.



Features and Benefits

Massively Multi-threaded Computing Architecture	Executes thousands of concurrent processing threads for high throughput parallel processing of mathematically intensive problems.
NVIDIA GPU Computing Drivers	Management of the GPU resources and an extensive runtime library for enhanced data management and program execution. Offers a high speed data transfer path and streamlined driver for computing, independent of the graphics driver.
Multi-GPU Computing	Multiple Tesla GPUs can be controlled by a single CPU via the GPU computing driver, delivering incredible throughput on computing applications. The power of the GPU to solve large-scale problems can be multiplied by splitting the problem across multiple GPUs.

Technical Specifications

NVIDIA Tesla Architecture

- Massively-parallel computing architecture with 128 multi-threaded processors per GPU
- Scalar thread processor with full integer and floating point operations
- Thread Execution Manager enables thousands of concurrent threads per GPU
- Parallel Data Cache enables processors to collaborate on shared information at local cache performance
- Ultra-fast memory access with 76.8 GB/sec. peak bandwidth per GPU
- IEEE 754 single-precision floating point

Supporting Platforms

- Tesla certified system*
- Microsoft Windows XP (32-bit)
Available on Tesla C870 and Tesla D870 only
- Linux (64-bit and 32-bit)
 - Red Hat Enterprise Linux 4 and 5
 - SUSE 10.1, 10.2 and 10.3 (coming soon)

*For Tesla D870 and Tesla S870

Product Details

Tesla C870 GPU Computing Processor

- One Tesla GPU (128 thread processors)
- 1.5 GB dedicated memory
- Fits in one full-length, dual slot with one open PCI Express x16 slot

Tesla D870 Deskside GPU Computing System

- Two Tesla GPUs (128 thread processors per GPU)
- 3 GB system memory (1.5 GB dedicated memory per GPU)
- Quiet operation (40dB) suitable for office environment
- Connects to host via cabling to a low power PCI Express x8 or x16 adapter card
- Optional rack mount kit

Tesla S870 GPU Computing System

- Four Tesla GPUs (128 thread processors per GPU)
- 6 GB of system memory (1.5 GB dedicated memory per GPU)
- Standard 19", 1U rack-mount chassis
- Connects to host via cabling to a low power PCI Express x8 or x16 adapter card
- Configuration: 2 PCI Express connectors driving 2 GPUs each (4 GPUs total)

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

