

NVIDIA TESLA ACCELERATE SCIENCE WITH A MINI-SUPERCOMPUTER FROM HP

BREAKTHROUGH SCIENTIFIC BARRIERS WITH THE TESLA-BASED HP STARTER KIT

REDEFINING HIGH PERFORMANCE IN THE DATACENTER

From the largest supercomputing centers to thousands of research institutes worldwide, IT managers are adding NVIDIA[®] Tesla[™] GPUs to their datacenter servers to cost-effectively increase overall system performance.

GPUs are power efficient, massively parallel processors that typically increase the performance of leading parallel applications by up to 10x. Most industry standard scientific and engineering applications are optimized for the GPU, resulting in lower power and cooling costs, smaller footprint, and lower capital cost.

SERVER EFFICIENCY WITH GPUs



1U server with dual-socket CPUs, comparing with vs without GPUs

ENABLING BREAKTHROUGH SCIENCE FOR RESEARCHERS

Many leading researchers in the field of life sciences, engineering, physics, and mathematics are leveraging GPUs to solve their most challenging problems. With accelerations of up to 10x, researchers are now able to run computer simulations which were previously impossible. For example, GPUs enable large, realistic biological systems to be simulated in minutes or hours rather than days and weeks.

Many popular scientific and engineering codes are accelerated on the GPU today, including:

Life Sciences

- > AMBER
- > GROMACS
- > LAMMPS
- > NAMD

Engineering

- > ANSYS Mechanical
- > CST MICROWAVE STUDIO
- > FluiDyna OpenFOAM
- > SIMULIA Abaqus

Math/Physics

- > Chroma
- Mathematica
- > MATLAB
- > MILC



24 GPUs inside with 10,752 compute cores!

HP GPU STARTER KIT

Configuration:

- > 1 DL380 control node w/ E5620 (QC, 2.4 GhZ, 80w) CPUs, 24 GB RAM and 2x 500 GB HDD
- > 2 SL6500 enclosures
- > 8 SL390s 2u server trays w/ X5675 (6C, 3.06 GhZ, 95w) CPUs, 48 GB RAM, 500 GB HDD, 3 Nvidia M2070 GPU modules
- > Voltaire IB 4x QDR 36 port managed switch
- > HPN ProCurve 2910 24 port 10/100/1000 Ethernet switch
- > Latest CUDA® 4.0 release
- > 3rd party development tools

OVER 460 UNIVERSITIES TEACHING & USING GPUs INCLUDING:



ACCELERATE YOUR APPLICATIONS

Dramatically increase the performance of your applications using GPUaccelerated libraries, directive-based solutions, and programming languages, such as the easy-to-use and widely adopted NVIDIA® CUDA® Accelerate the most performance critical areas in your application using the development tools you already know.

Learn more at www.nvidia.com/cudazone

GPU COMPUTING

With the ever-increasing demand for more compute performance, the HPC industry is moving toward a hybrid computing model where GPUs are added to an x86 CPU system to accelerate applications. GPUs are parallel processors that excel at tackling large amounts of similar data that is calculated simultaneously. When combined with CPUs, the end result is a tremendous speedup in overall system performance.

HP IS THE LEADER IN GPU COMPUTING SOLUTIONS

Since the 1990s, HP and NVIDIA have been partners in developing innovative GPU-based solutions, first with powerful workstations for graphic applications and today with GPU-based clusters for high-performance computing.

Business Drivers	HP ProLiant SL390s Resolution
Massive scalabilty	Built on the HP SL6500 chassis, the SL390s is purpose-built to provide a scalable infrastructure for high performance computing.
Increase performance without increased operating costs	Compared to quad-core CPUs, HP ProLiant SL390s servers with NVIDIA Tesla™ M2070 GPUs deliver equivalent performance at 1/20th the power consumption and 1/10th the cost.
Integrated management	Use HP Cluster Management Utility (CMU) to manage clusters of HP SL390s systems all from a single, centralized console.
Computing for less	Resolve bigger problems at an affordable price.

For more information, please contact your HP representative.

For Starter Kit guotes email hpc-sales@hp.com

NVIDIA Corporation | 2701 San Tomas Expressway, Santa Clara, CA 95050 T (408) 486 2000 | F (408) 486 2200 | www.nvidia.com

©2011 NVIDIA Corporation. NVIDIA, the NVIDIA logo, Quadro, CUDA and Tesla are trademarks and/or registered trademarks of NVIDIA Corporation. All rights reserved. All company and product names may be trademarks or registered trademarks of the respective owners with which they are associated. Features, availability, and specifications are subject to change without notice.

SUCCESSFUL CUSTOMERS

Harvard University

"(Harvard has) some of the world's best scientists doing cutting-edge research on some of the great scientific challenges of our times. ...We leverage GPU computing for what I call high-throughput science, ranging from radio astronomy and neuroscience to quantum chemistry and physics."

Hanspeter Pfister

Professor of the Practice of Computer Science Harvard University

Keeneland Project, Georgia Tech

"Georgia Tech has a long history of education and research that depends heavily on the parallel processing capabilities that NVIDIA has introduced with its CUDA architecture."

Jeffrey Vetter

Professor at Georgia Tech, Principal Investigator for Keeneland Group Leader, Oak Ridge National Laboratory

Tsubame 2.0, Tokyo Institute of Technology

"In testing our key applications, the Tesla GPUs delivered speed-ups that we had never seen before, sometimes even orders of magnitude."

Satoshi Matsuoka

Professor at Tokyo Institute of Technology

San Diego Supercomputing Center

"With Tesla GPUs, AMBER users in university departments can obtain application performance that outstrips what is possible even with extensive supercomputer access."

Ross Walker

Assistant Research Professor San Diego Supercomputing Center

