



NVIDIA[®] TESLA[™] ACCELERATE SCIENCE WITH SOLUTIONS FROM HP

GAIN SPEEDUP, THROUGHPUT, ACCURACY AND DO MORE ADVANCED SCIENCE WITH THE HP STARTER KIT ACCELERATED BY NVIDIA TESLA GPUs.

GPU computing is a disruptive technology that is transforming life science research. NVIDIA Tesla-based solutions, such as the HP ProLiant SL390s, typically accelerate applications by 10x, enabling researchers to address the world's most challenging computational problems faster than ever before.

Large, realistic biological systems are now simulated in hours or days, rather than weeks or months. And the accuracy and confidence levels obtained from these simulations is reducing the need to run costly, time consuming experiments.


Leading computational chemistry applications including AMBER, NAMD, GROMACS, and LAMMPS are already optimized for the GPU. By simply adding a GPU to your system, research on these codes is typically accelerated by several

orders of magnitude. Researchers may also accelerate proprietary code by leveraging NVIDIA CUDA[™], NVIDIA's powerful parallel computing architecture and programming model.

WHY 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 a x86 CPU system to accelerate applications. GPUs are parallel processors and excel at tackling large amounts of similar data because problems can be split into hundreds or thousands of pieces and calculated simultaneously. When combined with CPUs, which are adept at serial based tasks such as operating systems and data organization, the end result is a tremendous speed up in overall system performance.

AMBER 11 JAC NVE BENCHMARK



(CPU) Dual Socket Intel x5650
(GPU) DualSocket Intel x5650 + 2 Tesla M2070 GPUs



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
- > RHEL
- > CMU
- > Linux Value Pack
- > Latest CUDA 4.0 Release
- > Rack and infrastructure
- > Hardware/Software Integration

NVIDIA CUDA™ ACCELERATED APPLICATIONS FOR LIFE SCIENCES

Computational Chemistry	Computational Fluid Dynamics	Astrophysics, Earth Science	Engineering
AMBER	ALTAIR ACUSOLVE	ASUCA	MATLAB
NAMD	OpenFOAM	GADGET2	Mathematica
GROMACS	PROMETECH PARTICLEWORKS	NBODY	Jacket
LAMMPS	S3D (Sandia Lab)	NOAA NIM	ANSYS
GAMESS-US	TURBOSTREAM	WRF	SIMULIA ABAQUS

ABOUT CUDA

CUDA is NVIDIA's parallel computing model. Because CUDA offers a simple, intuitive method of programming in parallel, it is the easiest and most widely adopted parallel programming model in HPC today. The rich ecosystem of CUDA developer tools include:

- > CUDA C/C++ Compiler
- > Math Libraries
- > Profilers
- > Debuggers
- > Code samples
- > CUDA Programming Guide
- > Supports C, C++, Fortran, OpenCL, DirectCompute, JAVA, Python

CUDA information & free CUDA toolkit www.nvidia.com/cudazone

HP IS THE LEADER IN GPU COMPUTING SOLUTIONS

Since the 1990's, 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 scalability	Built on the HP SL6500 chassis, it is purpose-built to provide a scalable infrastructure for HPC
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	Users to resolve bigger problems at an affordable price

PLEASE CONTACT YOUR HP REPRESENTATIVE FOR MORE INFORMATION.

Marc Hamilton | marc.hamilton@hp.com
Paul Riviere | paul.riviere@hp.com

Len Amadeo | len.amadeo@hp.com
Dave Korf | david.korf@hp.com

For starter kit quotes: 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, Quadro, 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

CASE 1:

► 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

- > 201 TFLOPS in 7 racks
- > Already oversubscribed with scientific use
- > Early success stories in the fields of
 - > Material science
 - > Signal/Image processing
 - > Particle flow
 - > Biochemistry
- > Expansion with more Tesla GPUs vplanned

CASE 2:

► 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

- > World's Greenest Petascale Supercomputer
- > 2.4 PFLOPS
- > Accelerating apps in
 - > Plasma physics
 - > Tsunami prediction
 - > weather modeling
 - > Astrophysics
 - > And many more...

"A 10 nanosecond simulation was run in 2.6 hrs with 2 GPUs versus 6.7 hours using 128 INTEL 2.6 GHz processors. It's amazing how fast it was!"

Researcher from Universidad Autonoma de Mexico

