NVIDIA Tesla GPUs Now Easier Than Ever to Deploy With New GPU Starter Kit From HP

Ready-to-Use, Affordable NVIDIA GPU-Enabled Cluster From HP Addresses Increased Market Demand for Accelerated Computing Solutions

SANTA CLARA, CA -- NVIDIA today announced the availability of a limited-edition GPU Starter Kit from HP, a pre-configured system that provides researchers with a ready-to-use GPU computing cluster, straight out of the box. It consists of eight HP ProLiant SL390 G7 servers, containing 24 NVIDIA® Tesla™ M2070 GPUs, 16 CPUs, and is pre-configured with the latest NVIDIA CUDA® 4.0 parallel computing software.

Designed for ease of deployment and to address the growing availability of major GPU-enabled scientific applications, the GPU Starter Kit enables every university, government research department and enterprise customer to quickly deploy a rack of GPU-accelerated servers. The system delivers 13.5 teraflops of peak performance and is priced at $99,000, roughly 50 percent below the typical list price. The kit also includes a broad set of third-party development tools at discounted prices to facilitate application optimization.

"Growing demand for GPU computing has fueled the need for a fully integrated, robust and affordable development platform that enables developers to easily create new, accelerated applications," said Sumit Gupta, manager of Tesla products at NVIDIA. "As a result, NVIDIA and HP developed the GPU Starter Kit to remove some of the last hurdles to the mass adoption of GPU computing -- namely, the cost and time of system implementation."

HP ProLiant SL390 G7 servers equipped with Tesla GPUs are the computing platform of choice for leading supercomputers, such as the world's greenest production supercomputer[(i)], Tokyo Institute of Technology's Tsubame 2.0. They also power the Keeneland system -- managed by Georgia Tech, Oak Ridge National Lab and the National Institute for Computational Sciences -- enabling researchers and students to advance scientific discovery in various disciplines. Some of these efforts include:

- To better understand protein-DNA interactions with applications in future drug design, researchers Bo Hong, at Georgia Tech, and Jun-tao Guo, at University of North Carolina, Charlotte, are relying on GPUs to accelerate protein-DNA docking simulations. With 32 GPUs in the Keeneland cluster, they are able to achieve higher performance than with 1,200 CPUs alone.
- In search of advanced materials for wide ranging industrial uses, Jacek Jakowski, researcher from the University of Tennessee, is accelerating a quantum chemical molecular dynamics method by up to 20x using GPUs, reducing typical simulation times from two weeks to hours.
- Researchers Michela Taufer and Sandeep Patel at the University of Delaware are using GPUs to accelerate discoveries in human pathological conditions. With 10x faster simulations, researchers are running much larger size simulations, leading to more accurate insights.

Parallel programming has now become much easier with the recent release of CUDA 4.0 toolkit for C, C++ and Fortran. CUDA 4.0 represents a major leap forward for ease of GPU programmability and offers GPU-accelerated libraries such as FFT, BLAS, LAPACK, RNG, and SPARSE. CUDA C++ support has been enhanced by the powerful Thrust C++ template library, accelerating popular functions such as sorting and reductions.

For more information or to purchase the HP GPU Starter Kit, please contact your HP sales representative, or send an email to hpstarterkit@nvidia.com. For more information on NVIDIA Tesla GPUs for high performance computing, please go here. For more information on CUDA, please go here.

About NVIDIA
NVIDIA (NASDAQ: NVDA) awakened the world to the power of computer graphics when it invented the GPU in 1999. Since then, it has consistently set new standards in visual computing with breathtaking, interactive graphics available on devices ranging from tablets and mobile phones to notebooks and workstations. NVIDIA's expertise in programmable GPUs has led to breakthroughs in parallel processing which make supercomputing inexpensive and widely accessible. The Company holds more than 1,900 issued patents worldwide, including ones covering designs and insights that are essential to modern computing. For more information, see www.nvidia.com.

Certain statements in this press release including, but not limited to statements as to: the performance and pricing of NVIDIA Tesla GPUs, NVIDIA CUDA software and the GPU Starter Kit; the growing availability of major GPU-enabled scientific applications; the interest and demand for GPU computing; the effects of the company's patents on modern computing are forward-looking statements that are subject to risks and uncertainties that could cause results to be materially different than expectations. Important factors that could cause actual results to differ materially include: global economic conditions; our

reliance on third parties to manufacture, assemble, package and test our products; the impact of technological development and competition; development of new products and technologies or enhancements to our existing product and technologies; market acceptance of our products or our partners products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of our products or technologies when integrated into systems; as well as other factors detailed from time to time in the reports NVIDIA files with the Securities and Exchange Commission, or SEC, including its Form 10-Q for the fiscal period ended May 1, 2011. Copies of reports filed with the SEC are posted on the company's website and are available from NVIDIA without charge. These forward-looking statements are not guarantees of future performance and speak only as of the date hereof, and, except as required by law, NVIDIA disclaims any obligation to update these forward-looking statements to reflect future events or circumstances.

© 2011 NVIDIA Corporation. All rights reserved. NVIDIA and the NVIDIA logo, CUDA, and Tesla, are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated. Features, pricing, availability, and specifications are subject to change without notice.

(i) Based upon the Green500 List, November 2010 -- http://www.green500.org/

Ken Brown
Corporate Communications
+1-408-486-2626
kebrown@nvidia.com