NVIDIA Powers Titan, World's Fastest Supercomputer For Open Scientific Research

Oak Ridge National Lab's 20+ Petaflops System -- Powered by 18,000-Plus NVIDIA Tesla K20 GPUs -- Paves Path to Efficient, Affordable Exascale Computing

SANTA CLARA, CA -- Titan, the world's fastest open-science supercomputer,[1] was completed this month at Oak Ridge National Laboratory in Tennessee, opening new windows of opportunity into the exploration of some of the world's toughest scientific challenges.

Titan's peak performance is more than 20 petaflops -- or 20 *million billion* floating-point operations per second -- about 90 percent of which comes from 18,688 NVIDIA® Tesla® K20 GPU accelerators. These are based on the NVIDIA Kepler™ architecture, the fastest, most efficient, highest-performance computing architecture ever built.

Researchers use ever faster supercomputers to accelerate the pace of discovery and innovation across a range of scientific fields of inquiry -- from developing more efficient engines and higher capacity, lighter weight batteries, to studying climate change and finding cures for disease. Titan is a milestone on the path to exascale computing, which targets building a 1,000 petaflops supercomputer.

Titan is operated by Oak Ridge National Laboratory, part of the U.S. Department of Energy's network of research labs, as an open-science system. This means it is available to researchers from academia, government laboratories, and a broad range of industries who will use Titan to model physical and biological phenomena and seek breakthroughs faster than possible by experimentation alone.

Supported by the energy efficiency and cost-effectiveness of the Tesla K20 GPU, Titan is more than 10 times faster and five times more energy efficient than its predecessor, the 2.3-petaflops[2] Jaguar system, while occupying the same floor space. Had Oak Ridge upgraded Jaguar by simply expanding its CPU-based architecture, the system would be more than four times its current size and consume more than 30 megawatts of power.[3]

"Basing Titan on Tesla GPUs allows Oak Ridge to run phenomenally complex applications at scale, and validates the use of accelerated computing to address our most pressing scientific problems," said Steve Scott, chief technology officer of the GPU Accelerated Computing business at NVIDIA. "You simply can't get these levels of performance, power- and cost-efficiency with conventional CPU-based architectures. Accelerated computing is the best and most realistic approach to enable exascale performance levels within the next decade."

Titan development began three years ago with Oak Ridge's decision to upgrade Jaguar, the previous open science system leader and a former world No. 1 most powerful supercomputer. The upgrade includes the Tesla K20 GPU accelerators, a replacement of the compute modules to convert the system's 200 cabinets to a Cray XK7 supercomputer, and 710 terabytes of memory.

"Science and technology have always been our primary goal, and Titan is a groundbreaking tool that will allow researchers worldwide to leverage GPU-accelerated computing to make unparalleled breakthroughs," said Jeff Nichols, associate laboratory director for computing and computational sciences at Oak Ridge National Laboratory. "The new Tesla GPU accelerators offer the performance and energy efficiency that enable Titan to scale to unprecedented performance levels without consuming the energy equivalent of a small city."

About NVIDIA Tesla GPUs
NVIDIA Tesla GPUs are massively parallel accelerators based on the NVIDIA CUDA® parallel computing platform and programming model. Tesla GPUs are designed from the ground up for power-efficient, high performance computing, computational science and supercomputing, delivering dramatically higher application acceleration for a range of scientific and commercial applications than a CPU-only approach.

More information about NVIDIA Tesla GPUs is available at the Tesla website. To learn more about CUDA or download the latest version, visit the CUDA website. More NVIDIA news, company and product information, videos, images and other information is available at the NVIDIA newsroom. Follow us on Twitter at @NVIDIATesla.

About NVIDIA
NVIDIA (NASDAQ: NVDA) awakened the world to computer graphics when it invented the GPU in 1999. Today, its processors power a broad range of products from smartphones to supercomputers. NVIDIA's mobile processors are used in cell phones, tablets and auto infotainment systems. PC gamers rely on GPUs to enjoy spectacularly immersive worlds. Professionals use them to create 3D graphics and visual effects in movies and to design everything from golf clubs to jumbo jets. And researchers utilize GPUs to advance the frontiers of science with high performance computing. The company has

more than 5,000 patents issued, allowed or filed, including ones covering ideas 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 impact and benefits of NVIDIA Tesla GPUs and 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 July 29, 2012. 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.

(1) Based on June 2012 Top500 list.
(2) November 2011 Top500 List (http://top500.org/list/2011/11/100) - Jaguar = 2.3 petaflops peak performance @ 7 megawatts; Titan's peak performance will be in excess of 20 petaflops, consuming approximately 9 megawatts of electrical power.
(3) Assuming the same peak performance with a system composed of dual-CPU nodes, using the same CPU as in Titan.

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