



NVIDIA Tesla GPUs Again Power World's Greenest Petaflop Supercomputer

Tokyo Institute of Technology's "Tsubame 2.0" System Is Three Times More Energy-Efficient Than Comparable CPU-Only Supercomputers

SANTA CLARA, CA -- NVIDIA today announced that, for the second year in a row, the world's most energy efficient petaflop-class supercomputer is powered by NVIDIA® Tesla™ GPUs.

The Tsubame 2.0 system at the Tokyo Institute of Technology's Global Scientific Information Center (GSIC) ranks as the greenest petaflop-class supercomputer on the recently released [Green500 list](#). Published twice annually, the Green500 list, rates the 500 most energy efficient supercomputers based on performance achieved relative to power consumed.

Tsubame 2.0 is a heterogeneous supercomputer (combining both CPUs and GPUs) used to accelerate a range of scientific and industrial research in Japan. With sustained performance of 1.19 petaflops per second while consuming 1.2 megawatts, Tsubame 2.0 delivers 958 megaflops of processing power per watt of energy. It is 3.4-times more energy efficient than the next-closest x86 CPU-only petaflop system, the Cielo Cray supercomputer at Los Alamos National Laboratory, which delivers 278 megaflops per watt.

In the race to exascale computing, power efficiency has become the defining element of computing performance. Heterogeneous GPU-accelerated systems are inherently more energy efficient than CPU-only systems because applications can take advantage of the different processors for executing different jobs. The sequential parts of the application runs on CPUs, and the data- and compute-intensive parts are accelerated by the massively parallel GPU processor.

Tsubame 2.0 is comprised of HP ProLiant SL390 servers with Intel Xeon CPUs accelerated by NVIDIA Tesla GPUs. The Tesla GPUs provide more than 80 percent of its performance, enabling Tsubame 2.0 to achieve high levels of performance with very low power usage. This year, two of the five finalists for the prestigious [Gordon Bell Prize](#) ran on Tsubame 2.0, including the winner for Special Achievement in Scalability and Time-to Solution.

The latest Green500 list underscores the energy efficiency of heterogeneous computer design. Five of the world's 10 most efficient systems, and 22 of the top 30 most efficient systems, combine GPUs with CPUs.

Tesla GPUs are massively parallel accelerators based on the CUDA® parallel computing architecture. Application developers can accelerate their applications either by using CUDA C, CUDA C++, CUDA Fortran or using the simple, easy-to-use directive-based compilers.

For more information about Tsubame 2.0, visit the [Tokyo Institute of Technology, Global Scientific Information and Computing Center web site](#). To learn more about Tesla GPUs, visit the [Tesla web site](#). To learn more about CUDA, visit the [CUDA web site](#).

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 [smart phones](#) 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 visual effects in movies and design everything from golf clubs to jumbo jets. And researchers utilize GPUs to advance the frontiers of science with [high-performance computing](#). The company holds more than 2,100 patents worldwide, including ones covering ideas essential to modern computing. For more information, see [www.nvidia.com](#).

Tags / Keywords:

NVIDIA, CUDA, Tesla, SC11, Gordon Bell, Green500, Tsubame, Top500, GPU, GPU computing, supercomputing, parallel computing, GPGPU, high performance computing, HPC, programmer, directives, compiler, developers, research, scientific computing

Certain statements in this press release including, but not limited to statements as to: the benefits of Tesla GPUs and their impact on Tsubame 2.0; 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 October 30, 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, 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.

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