



NVIDIA Unveils World's Fastest, Most Efficient Accelerators, Powers World's No. 1 Supercomputer

New Tesla K20X GPU Accelerators Deliver 90 Percent of Performance of Titan, Crowned Fastest Supercomputer Today

SALT LAKE CITY, UT -- **SC12** -- NVIDIA today unveiled the NVIDIA® Tesla® K20 family of GPU accelerators, the highest performance, most efficient accelerators ever built, and the technology powering Titan, the world's fastest supercomputer according to the [TOP500 list](#) released this morning at the [SC12 supercomputing conference](#).

Armed with 18,688 NVIDIA Tesla K20X GPU accelerators, the [Titan supercomputer](#) at Oak Ridge National Laboratory in Oak Ridge, Tenn. seized the No. 1 supercomputer ranking in the world from Lawrence Livermore National Laboratory's Sequoia system with a performance record of 17.59 petaflops as measured by the LINPACK benchmark.⁽¹⁾

Tesla K20 - Performance, Energy-Efficiency Leadership

Based on the revolutionary NVIDIA [Kepler™ compute architecture](#), the new Tesla K20 family features the Tesla K20X accelerator, the flagship of NVIDIA's Tesla accelerated computing product line.

Providing the highest computing performance ever available in a single processor, the K20X provides tenfold application acceleration when paired with leading CPUs.⁽²⁾ It surpasses all other processors on two common measures of computational performance -- 3.95 teraflops single-precision and 1.31 teraflops double-precision peak floating point performance.

The new family also includes the Tesla K20 accelerator, which provides 3.52 teraflops of single-precision and 1.17 teraflops of double-precision peak performance. Tesla K20X and K20 GPU accelerators representing more than 30 petaflops of performance have already been delivered in the last 30 days. This is equivalent to the computational performance of last year's 10 fastest supercomputers combined.

"We are taking advantage of NVIDIA GPU architectures to significantly accelerate simulations in such diverse areas as climate and meteorology, seismology, astrophysics, fluid mechanics, materials science, and molecular biophysics," said Dr. Thomas Schulthess, professor of computational physics at [ETH Zurich](#) and director of the [Swiss National Supercomputing Center](#). "The K20 family of accelerators represents a leap forward in computing compared to NVIDIA's prior Fermi architecture, enhancing productivity and enabling us potentially to achieve new insights that previously were impossible."

Additional early customers include: Clemson University, Indiana University, Thomas Jefferson National Accelerator Facility (Jefferson Lab), King Abdullah University of Science and Technology (KAUST), National Center for Supercomputing Applications (NCSA), National Oceanic and Atmospheric Administration (NOAA), Oak Ridge National Laboratory (ORNL), University of Southern California (USC), and Shanghai Jiao Tong University (SJTU).

Energy-Efficiency for "Greener" Data Centers

The Tesla K20X GPU accelerator delivers three times higher energy efficiency than previous-generation GPU accelerators and widens the efficiency advantage compared to CPUs.

Using Tesla K20X accelerators, Oak Ridge's Titan achieved 2,142.77 megaflops of performance per watt, which surpasses the energy efficiency of the No. 1 system on the most recent Green500 list of the world's most energy-efficient supercomputers.⁽³⁾

Fastest on Broadest Range of Data Center Applications

The Tesla K20 family accelerates the broadest range of scientific, engineering and commercial [high performance computing](#) and data center applications. Today, [more than 200 software applications](#) take advantage of GPU-acceleration, representing a 60 percent increase in less than a year.

When Tesla K20X GPU accelerators are added to servers with Intel Sandy Bridge CPUs, many applications are accelerated up to 10x or more, including:⁽⁴⁾

- MATLAB (engineering) - 18.1 times faster
- Chroma (physics) - 17.9 times faster
- SPECfem3d (earth science) - 10.5 times faster
- AMBER (molecular dynamics) - 8.2 times faster

More information about the Tesla K20 GPU accelerators is available at NVIDIA booth 2217 at SC12, Nov. 12-15, and on the NVIDIA [high performance computing website](#). Users can also try the Tesla K20 accelerator for free on remotely hosted clusters. Visit the [GPU Test Drive website](#) for more information.

Availability

The NVIDIA Tesla K20 family of GPU accelerators is shipping today and available for order from leading server manufacturers, including Appro, ASUS, Cray, Eurotech, Fujitsu, HP, IBM, Quanta Computer, SGI, Supermicro, T-Platforms and Tyan, as well as from NVIDIA reseller partners.

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.

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.

¹ Compared to Sequoia system at Lawrence Livermore National Laboratory in Livermore at 16.3 petaflops of performance; Source: Top500.org - www.top500.org/lists/2012/06.

² CPU results: Dual socket E5-2687w, 3.10 GHz, GPU results: Dual socket E5-2687w + 2 Tesla K20X GPUs.

³ Source: Green500.org: www.green500.org/lists/green201206.

⁴ System configuration - CPU results: Dual socket E5-2687w, 3.10 GHz CPUs; GPU results: Dual socket E5-2687w CPUs + 2 Tesla K20X GPUs (MATLAB results running FFT compare one i7-2600K CPU 3.4 GHz vs. Tesla K20 GPU).

Certain statements in this press release including, but not limited to, statements as to: the impact, performance 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.

© 2012 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