



NVIDIA Tesla K10 GPU Hits New Performance Milestones for Scientific Simulation

Editorial Note: Update has been made to list of OEMs that are shipping Tesla K10 GPU accelerators.

ISC'12 - NVIDIA® Tesla® K10 GPUs offer performance breakthroughs on popular [high performance computing](#) (HPC) applications -- ranging from seismic processing to life sciences to video processing -- according to new benchmarks NVIDIA released today.

Based on the new [NVIDIA Kepler™ computing architecture](#), the Tesla K10 GPU delivers the industry's highest single precision performance (4.58 teraflops) and highest memory bandwidth (320 GB/sec) in a single accelerator. This is 12 times higher single precision flops and 6.4 times higher memory bandwidth than the latest-generation Intel Sandy Bridge CPUs⁽¹⁾.

The [Tesla K10 GPU](#) outperforms CPUs and previous-generation GPUs across the board on the most popular, compute-intensive applications for four key market segments, including:

- [Defense](#): video analytics, video stabilization, orthorectification, computer vision
- [Life and material sciences](#): molecular dynamics
- [Oil and gas](#): seismic processing, reverse time migration
- [Media and entertainment](#): video editing, video rendering/transcoding, ray tracing

"A distinct advantage of the Tesla K10 GPUs is that it excels in two key areas that have a dramatic impact on overall application performance: floating point operation and memory bandwidth," said Sumit Gupta, senior director of Tesla business at NVIDIA. "Together, these enable the K10 GPU to deliver substantial out-of-the-box performance increases for the top science, engineering and commercial applications with little or no effort on the part of the developer."

New Performance Records on AMBER and LAMMPS

On AMBER, a leading [biomolecular simulation software application](#), four Tesla K10 GPUs achieved world record performance, delivering far superior results than what was available on multiple racks of servers just a few years ago.⁽²⁾

The Tesla system achieved performance of 76 nanoseconds of computer simulation time in a day for a 23,558 atom molecule, outstripping the previous record set with four Tesla M2090s last year, providing [supercomputing](#) performance to thousands of individual researchers to fuel further innovation in such areas as new drug discovery and more effective materials.

"In biomolecular science, adding a few more nanoseconds of simulation time can make a world of difference in the ability of researchers to study and better understand the behavior of complex biological systems," said Ross Walker, assistant research professor, San Diego Supercomputing Center. "It still blows my mind that a single Tesla K10 outperforms some of the largest CPU clusters. The benefit it offers researchers is tremendous, enabling them to accelerate the search for new and better treatments for a host of diseases and disorders."

The Tesla K10 GPU also delivers the highest performance on LAMMPS, another application widely used by the life sciences research community. Running the LAMMPS Lennard Jones Liquid Benchmark, a single Tesla K10 GPU outperforms a Tesla M2090 GPU by 80 percent, delivering the equivalent performance of a cluster with 64 x86 CPUs.⁽³⁾

Accelerating the Search for Energy

NVIDIA Tesla GPUs continue to deliver the highest performance on reverse time migration (RTM) applications for seismic processing in the [oil and gas exploration](#) industry, and for image processing in the computer vision industry. Petrobras, the national oil and gas company of Brazil, achieved an 1.8x speed up on its RTM application on the Tesla K10 GPU, as compared to a Tesla M2090 GPU within the same power envelope.

NVIDIA Tesla K10 GPUs are available from leading OEMs, including Appro Supercomputer Solutions, HP, IBM, SGI and Supermicro, as well as through NVIDIA distribution partners. More information about the Tesla K10 is available on the NVIDIA [Tesla website](#).

About NVIDIA Tesla GPUs

NVIDIA Tesla GPUs are massively parallel accelerators based on the NVIDIA CUDA® parallel computing platform. 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](#). You can also follow us on [Twitter \(@NVIDIATesla\)](#).

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 benefits and impact 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 April 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, Kepler, Tesla and CUDA 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.

(1) Compared to Intel Xeon Processor E5-2690

(2) Server system is 4 nodes, each node configuration with a single Tesla K10 GPUs, Dual Intel Xeon X5670, 72 GB DDR3 memory. For large cluster performance benchmark:

http://ambermd.org/amber10_bench_files/jac_nve_kraken_ranger_large.png

(3) http://lammps.sandia.gov/bench/lj_xt5.html

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