

U.S. to Build Two Flagship Supercomputers for National Labs

Expected to Be 3X Faster Than Current Top Supercomputer, Systems to Couple IBM POWER CPUs With NVIDIA GPU Accelerators via High-Speed NVLink Interface

The U.S. Department of Energy today unveiled plans to build [two GPU-accelerated supercomputers](#) -- expected to deliver at least three-times greater performance than today's most powerful system -- which will move the world closer to the long-held goal of exascale computing.

The supercomputers, to be installed in 2017, will be based on next-generation IBM POWER servers with [NVIDIA® Tesla® GPU accelerators](#) and [NVIDIA NVLink™ high-speed GPU interconnect technology](#).

The "Summit" system at Oak Ridge National Laboratory will be capable of delivering 150 to 300 peak petaflops, and be used for open science. Delivering performance well in excess of 100 peak petaflops, the "Sierra" system will be a key element of Lawrence Livermore National Laboratory's national nuclear security mission.

They will be considerably faster than the U.S.'s current speed champ, Oak Ridge's "Titan," which delivers 27 peak petaflops, as well as the world's fastest, Tianhe-2 at China's National Super Computer Center, in Guangzhou, which delivers 55 peak petaflops.

Visit the [NVIDIA website](#) for more information and a [video](#) about these new systems.

"Today's science is tomorrow's technology," said Jen-Hsun Huang, CEO and co-founder of NVIDIA. "Scientists are tackling massive challenges from quantum to global to galactic scales. Their work relies on increasingly more powerful supercomputers. Through the invention of GPU acceleration, we have paved the path to exascale supercomputing -- giving scientists a tool for unimaginable discoveries."

The U.S. is investing in Summit and Sierra to achieve breakthroughs that lead to greater U.S. energy independence, new approaches to curbing climate change, dramatic improvements in fuel efficiency, natural disaster prediction, safer nuclear material storage, economic competitiveness, and more.

The systems represent the first major milestone in the ongoing partnership between IBM and NVIDIA. They build upon the work of the [OpenPOWER Foundation](#), an open development community formed to develop next-generation computing solutions for high performance computing and enterprise data center customers.

New Technologies Hasten Exascale Computing

The supercomputing community has for many years worked toward building exascale systems, which can perform a quintillion -- a billion billion or 10^{18} -- floating point calculations per second, known as FLOPS. A FLOP is equivalent to a single mathematical calculation, like multiplying two numbers together.

Summit and Sierra will be the next major step on the path to reaching exascale computing levels by virtue of a number of breakthrough technologies.

One is the NVIDIA NVLink high-speed GPU interconnect, which will be integrated into NVIDIA GPUs and IBM POWER CPUs powering the new systems. NVLink allows GPUs and CPUs to share data five to 12 times faster than today, and is designed to ultimately enable supercomputers that are 50 to 100 times faster than today's fastest systems.

The systems will also feature NVIDIA's [future generation GPU architecture, Volta™](#), which will deliver considerably higher performance than the company's current [Maxwell™ architecture](#) and subsequent Pascal™ design. Delivering significantly higher levels of computational performance than anything available today, NVIDIA GPUs will provide Summit and Sierra with more than 90 percent of the peak floating point processing capability.

"Our users have the most complex scientific problems and need exceptionally powerful computers to meet national goals," said Buddy Bland, project director of the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory. "The projected performance of Summit would not have been possible without the combination of these technologies, which will give our users an exceptionally powerful tool to accomplish these goals."

Summit to Accelerate Open Science Computing

Summit, like Titan, will be dedicated to open science, meaning that researchers worldwide will have the opportunity to apply for time on the system. It will also raise the bar for energy-efficient computing, providing five to 10 times higher performance than Titan, while using only 10 percent more power.

Sierra to Bolster National Security

Sierra will deliver five to 10 times higher compute performance than Lawrence Livermore's current fastest system, "Sequoia," which delivers 20 petaflops. It will be used for the National Nuclear Security Administration's program to ensure the safety, security and effectiveness of the nation's nuclear deterrent without testing, and nonproliferation efforts to prevent the spread of weapons of mass destruction worldwide.

Exascale Computing's Potential and Challenge

Exascale supercomputers are anticipated to further discovery into broad areas of science, engineering and industry, such as enabling work on cures for disease, providing insights into the human brain, helping to mitigate the effects of climate change, and increasing our understanding of the origins of the universe.

However, a fundamental challenge in attaining exascale systems is achieving higher levels of performance while minimizing energy consumption, a task GPU accelerators are particularly well suited for. An exascale system built with NVIDIA's latest GPU accelerators would consume five times less power than an x86 CPU-based system, enabling system designers to deliver extremely higher levels of energy-efficient performance.

About the NVIDIA Tesla Accelerated Computing Platform

The Tesla Accelerated Computing Platform is [designed from the ground up](#) for power-efficient, high performance computing, computational science, supercomputing, enterprise, complex data analytics and machine learning applications. It delivers dramatically higher performance and energy efficiency than a CPU-only approach. The platform deeply integrates the world's fastest GPU accelerators, advanced system management features, accelerated communication technology and [NVIDIA CUDA®](#), the world's most pervasive parallel computing model.

To Keep Current on NVIDIA:

- Keep up with the [NVIDIA Blog](#), and follow us on [Facebook](#), [Google+](#), [Twitter](#), [LinkedIn](#) and [Instagram](#).
- View NVIDIA videos on [YouTube](#) and images on [Flickr](#).
- Use the [Pulse news reader](#) to subscribe to the NVIDIA Daily News feed.

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