IDC LINK
Real-time IDC Research opinion on industry news, trends and events

DOE Goes All in on IBM-NVIDIA for 300 Petaflops System

November 15, 2014 - IDC Link

By: Earl Joseph; Robert Sorensen

The U.S. Department of Energy has signed a contract with IBM to build two new flagship supercomputers capable of performance ranging as high as 300 petaflops for delivery in 2017. The larger system, named Summit, will be installed at the DOE's Oak Ridge National Laboratory, and with an initial target of 150 petaflops — with the ability to scale up to 300 petaflops if an upgrade budget permits — it will provide five times the performance of the current-generation, Cray-built Titan supercomputer there, currently the most powerful HPC in the United States. Sierra will support the DOE's Office of Science in its broad science and energy mission, addressing the most challenging and impactful science problems for government, academia, and industry including combustion science, climate change, energy storage, and nuclear power.

The second, smaller system, named Sierra, will be installed at the Lawrence Livermore National Laboratory where 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. Sierra will have performance well in excess of 100 peak petaflops.

Both systems will be built by IBM and include major component suppliers NVIDIA and Mellanox. The larger Summit will consist of more than 3,400 nodes, each with IBM POWER9 processors and NVIDIA Volta GPUs. CPUs and GPUs will be connected with 80–200GBps NVLink with a large coherent memory, more than 512GB of combined DDR4, and high-bandwidth memory — all directly addressable from the CPUs and GPUs and 800GB of NVRAM, which can be configured as either a burst buffer or as extended memory. Each node will have 40 teraflops of peak performance: Just four nodes of Summit would make the current Top 500 list. The Sierra system will require about the same amount of total power as the current Titan system — about 10MW — but will have a footprint only about one-fifth its size.

The systems represent the first major milestone in the ongoing partnership between IBM and NVIDIA, which was announced at last year's SC conference and builds upon the work of the OpenPOWER Foundation, an open development community formed to develop next-generation computing solutions for high-performance computing and enterprise datacenter customers. Mellanox is also a member of the OpenPOWER Foundation.

This is a significant win for IBM, NVIDIA, and the OpenPOWER Foundation in general. IBM currently has three HPC systems in the June 2014 list of the top 10 supercomputers, but all are based on the Blue Gene chip. This procurement signals that IBM is looking to the Power chip to drive its next-generation HPC products, especially now that it has sold off its x86 series servers — formerly a major component of its overall HPC product line — to China's Lenovo. In addition, NVIDIA stands to gain by even more clearly drawing the battle lines with other accelerator options — notably the Intel Phi — by aligning so closely with IBM and its Power series.

Finally, these systems are part of the "Coral" DOE joint procurement process that will also include a 100 + petaflop supercomputer with a different architecture for Argonne National Laboratory. The winner of the Argonne contract is expected to be announced soon. The next procurement round after Coral, projected to occur sometime between 2020 and 2024, should introduce supercomputers with peak performance of

an exaflop or more. The two Coral systems announced today should nicely fill the performance slot between other DOE procurements such as the Trinity and NERSC-8 systems slated for acceptance in the next two years and the first exascale systems. These new procurements should also shed more light on what the DOE thinks an exascale system may ultimately look like: smaller numbers of more powerful nodes, large and fast CPUs, lots of GPU counts to achieve high flops counts, nimble ties between GPU and CPU, smaller footprints, and at least some cap on total power consumption.

Subscriptions Covered:

Technical Computing, High-Performance Data Analysis