

NVIDIA and Tech Leaders Team to Build GPU-Accelerated Arm Servers for New Era of Diverse HPC Architectures

Arm, Ampere, Cray, Fujitsu, HPE, Marvell to Build NVIDIA GPU-Accelerated Servers for Hyperscale-Cloud to Edge, Simulation to AI, High-Performance Storage to Exascale Supercomputing

SC19 -- NVIDIA today introduced a reference design platform that enables companies to quickly build [GPU-accelerated Arm®-based servers](#), driving a new era of high performance computing for a growing range of applications in science and industry.

Announced by NVIDIA founder and CEO Jensen Huang at the SC19 supercomputing conference, the reference design platform -- consisting of hardware and software building blocks -- responds to growing demand in the HPC community to harness a broader range of CPU architectures. It allows supercomputing centers, hyperscale-cloud operators and enterprises to combine the advantage of NVIDIA's accelerated computing platform with the latest Arm-based server platforms.

To build the reference platform, NVIDIA is teaming with Arm and its ecosystem partners -- including Ampere, Fujitsu and Marvell -- to ensure NVIDIA GPUs can work seamlessly with Arm-based processors. The reference platform also benefits from strong collaboration with Cray, a Hewlett Packard Enterprise company, and HPE, two early providers of Arm-based servers. Additionally, a wide range of HPC software companies have used [NVIDIA CUDA-X™ libraries](#) to build GPU-enabled management and monitoring tools that run on Arm-based servers.

"There is a renaissance in high performance computing," Huang said. "Breakthroughs in machine learning and AI are redefining scientific methods and enabling exciting opportunities for new architectures. Bringing NVIDIA GPUs to Arm opens the floodgates for innovators to create systems for growing new applications from hyperscale-cloud to exascale supercomputing and beyond."

Rene Haas, president of the IP Products Group at Arm, said: "Arm is working with ecosystem partners to deliver unprecedented performance and efficiency for exascale-class Arm-based SoCs. Collaborating with NVIDIA to bring CUDA acceleration to the Arm architecture is a key milestone for the HPC community, which is already deploying Arm technology to address some of the world's most complex research challenges and driving further advances in the embedded, automotive and edge segments."

The reference platform's debut follows NVIDIA's announcement earlier this year that it will bring its [CUDA-X software platform to Arm](#). Fulfilling this promise, NVIDIA is making available as a preview its Arm-compatible software development kit, consisting of NVIDIA CUDA-X libraries and development tools for accelerated computing.

Collaboration with Broader HPC Ecosystem

In addition to making its own software compatible with Arm, NVIDIA is working closely with its broad ecosystem of developers to bring GPU acceleration to Arm for HPC applications such as GROMACS, LAMMPS, MILC, NAMD, Quantum Espresso and Relion. NVIDIA and its HPC-application ecosystem partners have compiled extensive code to bring GPU acceleration to their applications on the Arm platform.

To enable the Arm ecosystem, NVIDIA collaborated with leading Linux distributors Canonical, Red Hat, Inc., and SUSE, as well as the industry's leading providers of essential HPC tools.

World-leading supercomputing centers have begun testing GPU-accelerated Arm-based computing systems. This includes Oak Ridge and Sandia National Laboratories, in the United States; the University of Bristol, in the United Kingdom; and Riken, in Japan.

Ecosystem Support

"Ampere is thrilled to be collaborating with NVIDIA to develop GPU-based acceleration solutions that will seamlessly work with our high-performance and power-efficient Ampere server processors. Our new products will provide our customers with the flexibility to choose the optimal NVIDIA GPU accelerator to efficiently run their most demanding workloads from cloud to edge."

-- Renee James, chairman and CEO, Ampere Computing

"It is great to see the quick progress NVIDIA has made bringing CUDA and OpenACC to Arm servers. We're very keen to work closely with NVIDIA and others to compile, profile and debug accelerated applications on this architecture. Our work to date with our 4,096-core Arm system has shown the strength of this collaborative approach."

-- Professor Mark Parsons, director, EPCC

"NVIDIA is a welcome and important member of the growing Arm ecosystem. Fujitsu believes NVIDIA will expand the Arm ecosystem and help ensure that customers have more choice for high performance computing and data science, particularly for AI, as we move into the new era of computing."

-- Takeshi Horie, corporate executive officer, vice head of Service Platform Business Group, Fujitsu

"Through our close collaboration with NVIDIA, customers deploying Marvell ThunderX2-based servers will now have access to the full suite of NVIDIA's GPU-accelerated software, marking an important milestone in the availability of accelerated computing for the Arm ecosystem. Together, we continue to drive new levels of energy-efficiency while delivering uncompromised performance for a multitude of supercomputing and AI applications for the exascale era."

-- Gopal Hegde, vice president and general manager, Server Processor Business Unit, Marvell Semiconductor, Inc.

"With the help of HPE, Marvell and NVIDIA, we're excited to see how quickly our Oak Ridge National Laboratory-led effort could upgrade our Arm test-bed system, pull together performance testing, and get positive results. In just two weeks we were able to compile and correctly run approximately eight leadership-class applications, three important community libraries, and benchmark suites frequently used for evaluating the Arm-HPC ecosystem. Based on early results, the functionality of this Arm-hosted accelerated computing ecosystem appears to be already on par with the POWER and x86 environments. It's a terrific beginning for an accelerated computing ecosystem within Arm."

-- Jack C. Wells, director of science, National Center for Computational Sciences, Oak Ridge National Laboratory

"We've had a long history of collaboration with NVIDIA. It's exciting to see the company delivering on its promise to bring its leading CUDA-X software stack and ecosystem to the Arm HPC community. We've started evaluating Riken's codes on NVIDIA's GPU-accelerated Arm system and feel it will drive innovation for a new wave of HPC and AI converged workloads in Japan."

-- Satoshi Matsuoka, director, Riken

"With NVIDIA's newly added support for Arm host CPUs, Kokkos and LAMMPS worked out of the box. The results met our expectations and are in line with our experience deploying NVIDIA GPUs in systems with x86-based CPUs."

-- Christian Trott, principal member, Technical Staff, Sandia National Laboratories

"NVIDIA's Arm-based software stack is really solid straight out of the box. In the past we've made extensive use of Arm and NVIDIA platforms separately, but being able to combine them is an exciting prospect for us. Performance of NVIDIA's GPU drivers on Arm is really good, comparable with the x86 platforms we tried, which is an impressive achievement in such a short time period."

-- Simon McIntosh-Smith, professor of high performance computing, University of Bristol

About NVIDIA

[NVIDIA](#)'s (NASDAQ: NVDA) invention of the GPU in 1999 sparked the growth of the PC gaming market, redefined modern computer graphics and revolutionized parallel computing. More recently, GPU deep learning ignited modern AI — the next era of computing — with the GPU acting as the brain of computers, robots and self-driving cars that can perceive and understand the world. More information at <http://nvidianews.nvidia.com/>.

Certain statements in this press release including, but not limited to, statements as to: the benefits and impact of a reference design platform for building GPU-accelerated Arm-based servers, breakthroughs in machine learning and AI, bringing NVIDIA GPUs and GPU-accelerated applications to Arm, and NVIDIA's collaboration with Arm and its ecosystem partners; a renaissance in high performance computing; the HPC community deploying Arm technology to address complex research challenges and driving further advances; NVIDIA bringing its CUDA-X software platform to Arm; and Fujitsu's belief that NVIDIA will expand the Arm ecosystem and help ensure that customers have more choice for high performance computing and data science 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 most recent reports NVIDIA files with the Securities and Exchange Commission, or SEC, including, but not limited to, its annual report on Form 10-K and quarterly reports on Form 10-Q. 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.

© 2019 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo and CUDA-X 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.

Media Contacts

Kristin Uchiyama

+1-408-486-2248

kuchiyama@nvidia.com