

NVIDIA CUDA-X AI And HPC Software Stack Now Available On Marvell ThunderX Platforms

Drives Industry-Leading Application Results for Energy-Efficient High Performance Computing

SANTA CLARA, Calif. and DENVER, Nov. 18, 2019 /[PRNewswire](#)/ -- Marvell (NASDAQ: MRVL) today announced the availability of NVIDIA GPU support on its ThunderX® family of Arm®-based server processors. Following NVIDIA's [June announcement](#) to bring CUDA to the Arm architecture, Marvell has collaborated with NVIDIA to port its CUDA-X AI™ and HPC libraries, GPU-accelerated AI frameworks and software development tools to the ThunderX platform. The computational performance and memory bandwidth of ThunderX2®, Marvell's latest 64-bit Armv8-A based server processor, combined with the parallel processing capabilities of NVIDIA GPUs provide a compelling path to energy-efficient exascale computing.

Artificial intelligence (AI) and machine learning (ML) continue to become essential technology components to data center server requirements at the cloud and network edge. To address these evolving AI and ML workloads, as well as the most challenging and complex problems in science and research, supercomputers need processors that are optimized to provide cutting-edge throughput, application latency and power.

With an initial focus on computational science applications including GROMACS, NAMD, MILC and LAMMPS, the ThunderX2 configurations are demonstrating compelling performance with an enhanced ability to drive higher and more efficient combined application results in a GPU-enabled system.

"NVIDIA GPU support for our ThunderX2 server processor brings clear, differentiated value to meet the distinctive performance and power requirements of the exascale computing era," said Gopal Hegde, vice president and general manager, Server Processor Business Unit at Marvell Semiconductor, Inc. "The availability of NVIDIA GPU acceleration on the Arm architecture will further expand the ThunderX2 ecosystem for HPC, cloud computing and edge markets, spurring innovation across low level firmware through system software to commercial ISV applications."

"The availability of CUDA acceleration for ThunderX2 processors marks a significant milestone in bringing the power efficiency and high performance of the Arm architecture to the infrastructure market," said Chris Bergey, senior vice president and general manager, Infrastructure Line of Business at Arm. "The breadth and depth of innovation across the ecosystem enables an easy migration path and robust support for existing and future GPU workloads from the edge to the cloud."

"NVIDIA GPU-accelerated computing on Arm provides customers worldwide with greater choice in building next-gen AI-enabled supercomputers," said Ian Buck, general manager and vice president of Accelerated Computing at NVIDIA. "Combining NVIDIA's unmatched platform for AI and HPC with Marvell's powerful ThunderX2 Arm-based server processors is already delivering impressive application performance."

Ecosystem Support

ThunderX2 is the most widely supported Armv8-A server processor with an ecosystem of over 100 partners across commercial, open source and industry standards engagements. NVIDIA's full software suite support is enabling the acceleration of more than 600 HPC applications and AI frameworks on ThunderX2 systems.

"Our collaboration with Marvell enables us to support servers with the industry-leading performance of ThunderX2 with our SC8000 compute acceleration expansion platform, bringing data center AI capabilities to a host of edge applications," said Steve Cooper, CEO at One Stop Systems. "The SC8000 is the industry's first solution that incorporates NVIDIA Tesla GPUs with NVLink and Arm servers. The addition of Arm-based architecture into our solutions extends the value of use cases for AI on the Fly edge appliances for our customers."

"Red Hat and Marvell have a long history of collaborating in the Arm server ecosystem, helping to bring open, industry-wide standards to enterprise Arm architecture," said Chris Wright, senior vice president and chief technology officer at Red Hat. "Enabling NVIDIA GPUs on ThunderX-based systems paired with the CUDA-X SDK and libraries supports customer choice in terms of architecture for running HPC, AI and ML applications on top of Red Hat platforms."

"SUSE and Marvell have a successful track record of strong HPC collaboration as seen most recently in the Catalyst UK deployment," said Daniel Nelson, vice president, Products & Solutions Management at SUSE. "We are excited to build upon this collaboration with the addition of NVIDIA's GPUs and suite of HPC, AI and ML applications support on Marvell's ThunderX2. SUSE looks forward to driving additional optimizations and software infrastructure support in the Arm ecosystem and will be demonstrating GPU support at the SC19

industry conference in Denver."

The Marvell® ThunderX platform with NVIDIA GPU support will be demonstrated at the Marvell booth (#349) at The International Conference for High Performance Computing, Networking, Storage and Analysis (SC19), running through November 22, 2019 in Denver, Colorado.

About Marvell

Marvell first revolutionized the digital storage industry by moving information at speeds never thought possible. Today, that same breakthrough innovation remains at the heart of the company's storage, processing, networking, security and connectivity solutions. With leading intellectual property and deep system-level knowledge, Marvell's semiconductor solutions continue to transform the enterprise, cloud, automotive, industrial, and consumer markets. To learn more, visit: <https://www.marvell.com/>

Marvell, the M logo, ThunderX and ThunderX2 are registered trademarks of Marvell and/or its affiliates in the US and/or elsewhere. Other names and brands may be claimed as the property of others.

Marvell Media Relations

Kristin Hehir
Senior Manager, Public Relations
408-222-8744
kristinh@marvell.com

Hanna Kang
Senior Manager, Public Relations
408-222-3780
hkang@marvell.com

View original content to download multimedia:<http://www.prnewswire.com/news-releases/nvidia-cuda-x-ai-and-hpc-software-stack-now-available-on-marvell-thunderx-platforms-300960385.html>



SOURCE Marvell

<https://investor.marvell.com/2019-11-18-NVIDIA-CUDA-X-AI-and-HPC-Software-Stack-Now-Available-on-Marvell-ThunderX-Platforms>