



FOR IMMEDIATE RELEASE

Contact:

Mike LaPan
Cirrascale Corporation
(858) 874-3800
mike.lapan@cirrascale.com

CIRRASCALE® LAUNCHES EXPANDED MULTI-GPU SOLUTIONS WITH NVIDIA TESLA P100 ACCELERATED COMPUTING PLATFORM

The Cirrascale GX Series Rackmount and GB5600 Blade Server support up to eight NVIDIA Tesla P100 GPU accelerators with 16GB HBM2 memory, the world's most advanced datacenter accelerator

SAN DIEGO, CA -- September 8, 2016 -- Cirrascale Corporation®, a premier developer of blade and rackmount solutions enabling GPU-driven deep learning infrastructure, today announced it will begin shipping the recently announced NVIDIA® Tesla® P100 GPU accelerators with 16GB HBM2 memory throughout its high-performance, deep learning product lines.

Utilizing the company's proprietary 96-lane Gen3 PCIe switch-enabled risers, the GX Series and GB5600 Series product lines can attach up to eight discrete NVIDIA Tesla P100 GPUs on the same PCIe root complex in a single rackmount or blade server chassis, with additional room for other Gen3 PCIe x16 devices such as InfiniBand® or NVMe cards.

"Our customers have been clamoring to get their hands on NVIDIA's Pascal-based architecture to help further advance their HPC and deep learning needs," said PJ Go, president of Cirrascale Corporation. "The Tesla P100 GPU tied with Cirrascale hardware, can deliver one of the world's fastest compute nodes containing with one of the most advanced datacenter accelerators capable of solving today's important challenges in HPC and deep learning."

Extending the capabilities of these accelerators, the Cirrascale SR3615 PCIe switch riser enables up to 10 PCIe Gen3 x16 compatible devices -- such as GPU accelerators, InfiniBand network or NVMe storage cards -- to communicate directly with each other on the same PCI root complex. This eliminates the need for host CPU intervention by allowing the accelerators to share a single memory address space and make use of DMA to control data movement. When used in conjunction with NVIDIA GPUDirect™ technology, compatible PCIe Gen3 x16 devices can directly read and write CUDA™ host and device memory, including memory owned by network and storage devices. By doing so, it eliminates unnecessary memory copies, dramatically lowers CPU overhead, and reduces latency resulting in significant performance improvements in data transfer times.

"Cirrascale is helping us bring the latest technology to advanced datacenters and deliver the world's fastest GPU compute nodes," said Roy Kim, Tesla Product Lead for the Accelerated Computing Group at NVIDIA. "Cirrascale's purpose-built solutions, now shipping with Tesla P100 accelerators, will help researchers and data scientists develop better, more accurate artificial intelligence applications to drive the deep learning revolution forward."

The Cirrascale GX Series rackmount and GB5600 Series blade servers supporting the NVIDIA Tesla P100 GPU accelerators -- as well as the Cirrascale proprietary PCIe switch-enabled riser -- are immediately available to order and are shipping to customers now.

About Cirrascale Corporation

Cirrascale Corporation is a premier developer of hardware and cloud-based solutions enabling GPU-driven deep learning infrastructure. Cirrascale leverages its patented Vertical Cooling Technology and proprietary PCIe switch riser technology to provide the industry's densest rackmount and blade-based peered multi-GPU platforms. The company sells hardware solutions to large-scale deep learning infrastructure operators, hosting and cloud service providers, and HPC users. Cirrascale also licenses its award winning technology to partners globally. To learn more about Cirrascale and its unique multi-GPU infrastructure solutions, please visit <http://www.cirrascale.com> or call (888) 942-3800.

Cirrascale and the Cirrascale logo are trademarks or registered trademarks of Cirrascale Corporation. NVIDIA, the NVIDIA logo, Pascal, GPUDirect, CUDA and Tesla are trademarks or registered trademarks of NVIDIA Corporation. All other names or marks are property of their respective owners.

###