NEWS RELEASE

Penguin Computing Announces New Magna and Relion Servers with NVIDIA Tesla P100 GPU Accelerators for High Performance Computing

2016-06-21

Collaboration to Support Customers' Intensive HPC Requirements

FREMONT, CA - June 20, 2016... Penguin Computing, a provider of high performance, enterprise data center and cloud solutions, today announced two new servers powered by NVIDIA® Tesla® P100 GPU accelerators with NVIDIA NVLink™ high-speed interconnect technology. "With support for the Tesla Accelerated Computing Platform, Penguin Computing's Magna and Relion servers offer exceptional solutions to power computationally intensive applications in HPC, deep learning, machine learning and cognitive intelligence," said Roy Kim, group product manager of Accelerated Computing at NVIDIA. Penguin Computing's new servers are the Magna 2002S and the Relion X904GT. The Magna 2002S uses OpenPOWER CPUs and includes the NVIDIA NVLink™ technology supporting four Tesla P100 GPUs in SXM2 sockets in a 2U, 19" EIA form factor. Delivering a 5x speed-up over PCIe Gen3 technology, NVLink™ provides high-bandwidth and ultra-fast communication between GPUs and CPUs, as well as between GPUs. This makes the Magna 2002S highly suitable for deep learning workloads that scale well to multiple GPUs. The Relion X904GT adds an accelerated computing platform to Penguin's Tundra product line. Penguin's design innovation allows customers to choose from a variety of topologies, incorporating NVIDIA NVLink™, to best accelerate their applications. "As the market for deep learning continues to grow, we will enable customers such as AI researchers and data scientists to choose the topology that best accelerates their particular applications," said William Wu, senior product manager at Penguin Computing. "With our Relion X904GT, Penguin Computing can deliver a solution that is tailored to their needs, rather than a one size fits all approach." Tesla P100 with NVIDIA NVLink™ technology enables traditional HPC servers to substantially accelerate time-to-solution for data-intensive applications. A server node with NVLink™ can interconnect up to eight Tesla P100 GPUs at 5x the bandwidth of PCIe, helping customers solve critical technical challenges that arise from scaling HPC and deep learning applications. All existing Relion GT servers - Relion 1904GT, 2904GT and 2908GT - will be P100 ready. Penguin Computing is a sponsor of this week's

ISC High Performance Conference. Visit the company's booth #921 at the event in Frankfurt, Germany.