

Penguin Computing Updates Densest GPU Compute Solution with NVIDIA Tesla K40 GPU Accelerators

2013-11-19

Penguin Computing, a provider of high performance and enterprise computing solutions, today announced that it will offer NVIDIA® Tesla® K40 GPU accelerators across its range of GPU computing platforms, including its flagship Relion 2808GT servers. The Relion 2808GT was designed to deliver maximum compute power with a minimal datacenter footprint and maximum efficiency. The system supports eight GPU accelerators or coprocessors in two rack units and provides a higher compute density than any other server on the market. Configured with NVIDIA Tesla K40 GPU accelerators, a single Relion 2808GT can achieve over 33TFLOPs of peak performance. "Our GPU computing strategy is to offer the widest range of GPGPU ratios and densities in the industry in order to support the widest range of customer requirements," said Jussi Kukkonen, Penguin Computing Director of Product Management. "The Relion 2808GT platform in conjunction with NVIDIA Tesla K40 GPU accelerators delivers unprecedented levels of performance and provides the most cost-effective solution for GPU-centric applications." The NVIDIA Tesla K40 GPU accelerator is the first and highest-performance accelerator optimized for big data analytics and large-scale scientific workloads. It delivers double the memory and up to 40 percent higher performance than its predecessor, the NVIDIA Tesla K20X GPU accelerator, and 10 times higher performance than today's fastest CPU. It features intelligent NVIDIA GPU Boost technology, which converts power headroom into a user-controlled performance boost, enabling users to unlock the untapped performance of a broad range of applications. The Tesla K40 GPU accelerator is based on the **NVIDIA Kepler™ GPU computing architecture** and powered by NVIDIA CUDA®, the world's most pervasive parallel computing model. "Our goal is to always bring the very latest technology to our HPC and enterprise customers in a well managed system backed by world-class support," said Charles Wuischpard, CEO Penguin Computing. "The Relion 2808GT server paired with NVIDIA Tesla K40 GPU accelerators is truly the next step in highperformance computing." The Relion 2808GT is the ideal platform for running scientific and engineering applications that support GPUs and is the highest density server that uses the NVIDIA Tesla K40 GPU accelerator at launch. The Relion 2808 GT is available in a 2U rack-mount, supporting eight full width GPUs. "Tomorrow's high-end computing systems will need to be optimized for both performance and

energy efficiency,” said Sumit Gupta, general manager of Tesla Accelerated Computing products at NVIDIA. “By incorporating Tesla K40 GPU accelerators, Penguin’s high-density servers give HPC customers new levels of performance for their most demanding HPC workloads while minimizing costs.” In addition to our Relion 2808GT supporting highest density 8x NVIDIA K40 GPUs in 2U, our Relion 1800GT - 3x NVIDIA K40 GPUs in 1U, and Relion 2800GT - 4x NVIDIA K40 GPUs in 2U, will also be supported 11/18/2013 for the launch of the NVIDIA K40 GPU accelerator. For more information, please visit www.penguincomputing.com.