

Penguin Computing Deploys World's First HPC Cluster Based on AMD Accelerated Processing Units at Sandia National Labs

2011-11-03

Penguin Computing, experts in high performance computing (HPC) solutions, today announced that the company has successfully installed the world's first HPC cluster powered by AMD (NYSE: AMD) Accelerated Processing Units (APUs) at Sandia National Labs in Albuquerque, New Mexico. The system comprises 104 servers that are interconnected through a QDR Infiniband fabric. It delivers a theoretical peak performance of 59.6 TFLOPs.

Penguin's new Altus 2A00 is the compute platform for this deployment. The Altus 2A00 was specifically designed by Penguin Computing, in partnership with AMD, to support the AMD Fusion APU architecture. It is the world's first Fusion APU system in a rack mountable chassis in a 2U form factor. The unique AMD APU processor design combines multi-core x86 processing, memory controllers, a PCI-E interface and massively parallel GPU computing on a single piece of silicon. The APU includes 400 parallel processing cores that can be leveraged for HPC applications through the OpenCL programming framework. Unlike conventional GPU server architectures, APU parallel multiprocessors share the same physical memory space with CPU cores. As a result, the programming model for APUs is simpler, bottlenecks for data movement between GPU and main memory are avoided and data duplication is eliminated. These capabilities offer particularly compelling benefits when deployed in conjunction with low-latency RDMA interconnects such as Infiniband, as they allow for building efficient distributed GPU applications. "We are interested in research on next generation computer architectures and look forward to collaborating with Penguin and AMD to advance power-efficient computing strategies," said James Ang, Manager of the Scalable Computer Architectures department at Sandia National Labs. "This first of a kind cluster of Altus 2A00 servers will support our exploration of advanced programming models like OpenCL, which seamlessly map MPI applications to the CPU and GPU cores, and research into system software support for advanced data movement capabilities." "With the Altus 2A00, Penguin is the first to bring AMD's unique APU capabilities to the HPC community," says Phil Pokorny, CTO Penguin Computing. "We are extremely proud of our successful deployment of

this platform on such a large scale. We believe that the high level of integration and the resulting benefits for HPC users will further accelerate the adoption of the GPU processing model in HPC. The APU architecture has the potential to become a key component of future exascale systems." "As we progress closer to the exascale era, it's clear that the traditional paradigms of supercomputing are continuing to evolve and require new technologies to keep pace with the rapid levels of innovation that HPC customers demand," said John Byrne, corporate vice president and general manager for Americas Mega Region, AMD. "With industry-leading CPU and GPU technology, AMD has the pieces to assemble a wide range of solutions for HPC deployments, and now HPC customers can leverage the power and efficiency of AMD's APU technology with Penguin Computing's Altus 2A00 server."