PENGUIN™ SOLUTIONS

NEWS RELEASE

Penguin Computing Supplies High Performance Computer to University of Pittsburgh

2011-07-20

Penguin Computing, experts in high performance computing (HPC), today announced the installation of a new heterogeneous cluster at the University of Pittsburgh. This new cluster, used by researchers in multiple departments, is operated by the University of Pittsburgh's Center for Simulation and Modeling (SAM); the hardware and operating systems are maintained and monitored for reliability, performance and security by Computing Services and Systems Development (CSSD), the University's central IT organization. “When our research efforts called for a new cluster, there was no question that we would go to Penguin for support,” said Kenneth D. Jordan, Distinguished Professor in the School of Arts & Sciences at the University of Pittsburgh. “We have purchased several computer clusters from Penguin and have been very pleased with their reliability. There is a major advantage in terms of maintaining the various clusters to use the same operating system and management and queue software on all systems.” The new supercomputing cluster is heterogeneous with three types of nodes, those with Intel Westmere CPUs, those with AMD Magny-Cours CPUs, and those with both Westmere CPUs and NVIDIA Fermi GPUs. The Westmere nodes have 12 cores and 48 GB of memory each and the Magny-Cours nodes have 48 cores and 128 GB of memory each. The choice of hardware was dictated by the mix of calculations carried out by users and networking makes use of QDR Infiniband to enable excellent performance on parallel applications. Similar to earlier Penguin Computing clusters installed at the University of Pittsburg, this new heterogeneous cluster is managed by Penguin’s Scyld ClusterWare software system. “The University of Pittsburgh expands Penguin’s footprint in the research market,” said Charles Wuischpard, CEO, Penguin Computing. “We are pleased that the SAM researchers chose to work with us again when their research needs grew.” Research conducted at the University of Pittsburgh using the new Penguin cluster ranges from understanding mechanisms of protein misfolding to understanding the mechanisms of CO_2 capture and sequestration. Applications running on the cluster make use of codes developed at the University of Pittsburgh as well as from other groups around the world. Popular codes running on the cluster include:

- Amber (http://www.ambermd.org) – A flexible molecular dynamics simulation package for the simulation of biomolecular systems.
- LAMMPS (http://lammps.sandia.gov) – A highly parallel molecular dynamics program for modeling complex materials.
- CASINO (http://www.tcm.phy.cam.ac.uk/~mdt26/casino2.html) – This code solves the electronic Schrodinger Equation using varational and diffusion Monte Carlo methods. It scales well over thousands of CPU cores.
- Molpro (http://www.molpro.net) – This package is designed for highly accurate electronic structure calculations on molecular systems.

The Penguin clusters at the University of Pittsburgh provide a total of about 3000 CPU cores as well as several GPUs.