

Penguin Computing Announces Open Compute Project-Driven Tundra Cluster Ecosystem for Increased Customer Choice in High Performance Computing

2015-03-10

SAN JOSE, Calif. – March 9, 2015 - Penguin Computing, a provider of high performance, enterprise data center and cloud solutions, today gave customers a peek at the future of increased customization and choice in high performance computing at the Open Compute Project Summit. “Facebook had the forethought to create the Open Compute Foundation and share IP from designing a highly efficient computing infrastructure at an extremely low cost,” said Phil Pokorny, chief technology officer, Penguin Computing. “We are now building on that collaborative development model to bring expanded flexibility with regard to form factors, processors and configurations for a broad range of customer requirements.” The Penguin Tundra™ cluster platform, based on Open Compute Project rack level infrastructure, delivers the highest density and lowest total cost of ownership for high performance technical computing clusters. Large-scale HPC deployments will benefit from Tundra, which is designed to accommodate future exascale HPC components such as coprocessors and fabrics. Being an active member of the Open Compute Project community is a natural step for Penguin Computing as an early Linux pioneer that understands the benefits of community driven solutions. “Penguin’s Tundra computing platform brings intelligent modularity and flexibility to the Open Compute Project for applications needing the optimum combination of high performance and high density components,” said Steven Hill, senior analyst, Data Center Solutions, Current Analysis. Penguin Computing’s Tundra product line includes compute sled, storage sled and now an Intel® Xeon Phi™ processor-based motherboard. The Intel Xeon Phi processor, code-named Knights Landing, delivers optimized performance for highly parallel applications. Each accelerator sled can be implemented interchangeably for similar or distinct operations. “Working closely with Penguin Computing, we’re pleased to show the first Intel Xeon Phi processor-based motherboard in an HPC platform based on a standard OCP rack design,” said Hugo Saleh, director of Marketing and Industry Development, Technical Computing Group at Intel. “Penguin Computing’s OCP design

provides a compelling implementation for the Intel Xeon Phi processor-based motherboard.” Open Compute Project-based design allows third-party motherboards to fit in a Tundra sled so any branded motherboards can take advantage of the efficient space utilization in OCP’s Open Rack standard. While the previous density was two nodes in a 1U configuration, restricted by the 19” EIA standard traditional server width, Penguin now enables three nodes in a 1U via a 21” wide server. Penguin Computing’s Tundra supports liquid cooling water blocks that allow highly dense racks, air conditioning cost savings and risk minimization for power overloads. Customer segments for Penguin Tundra include manufacturing, financial and many other sectors. Please visit <https://www.penguinolutions.com/products/ocp-hpc-ai-systems/ocp-servers-storage/> for more information.