



NEWS RELEASE

Penguin Computing™ Announces Intel Xeon Scalable Processor Availability for Penguin Computing On-Demand HPC Cloud

11/12/2017

POD HPC Cloud Demand and Growth to Continue into 2018

FREMONT, CA – November 13, 2017 – Penguin Computing, provider of high performance computing, enterprise data center and cloud solutions, today announced that more than 11,500 cores of the latest Intel® Xeon® Scalable processor (codenamed: Skylake-SP) will be available in December 2017 on Penguin Computing® On-Demand™ (POD) HPC cloud. The new POD HPC cloud compute resources use Intel® Xeon® Gold 6148 processors, a cluster-wide Intel® Omni-Path Architecture low-latency fabric and are integrated with Penguin Computing® Scyld Cloud Workstation™ for web-based, remote desktop access into the public HPC cloud service. “As an HPC cloud provider, we know that it is critical to provide our customers with the latest processor technologies,” said Victor Gregorio, Senior Vice President, Cloud Services, Penguin Computing. “The latest Intel Xeon Scalable processor expansion will provide an ideal compute environment for MPI workloads that can leverage thousands of cores for computation. We have significant customer demand for POD HPC cloud in applicable areas like high-resolution weather forecasting and computational fluid dynamics, including solutions from software partners like ANSYS, Flow Science and CD-adapco.” “Intel offers a balanced portfolio of HPC optimized components like the Intel Xeon Scalable processor and Intel Omni-Path Architecture, which provides the foundation for researchers and innovators to drive new discoveries and build new products faster than ever before,” said Trish Damkroger, Vice President of technical computing at Intel. “Penguin Computing On-Demand provides an easy and flexible path to access the latest technology so more users can realize the benefits of HPC.” Scientists and engineers at every company are trying to innovate faster while holding down costs. Modeling and simulation are the backbone of these efforts. Customers may wish to run simulations at scale, or many different permutations simultaneously but may require more computing resources than are readily available in-house. The POD HPC cloud offers organizations a flexible, cost effective approach to meeting these requirements. The Intel Xeon Scalable processor provides increased



performance, a unified stack optimized for key workloads including data analytics, and integrated technologies including networking, acceleration and storage. The processor's increased performance is realized through innovations including Intel® AVX-512 extensions that can deliver up to 2x FLOPS per clock cycle, which is especially important for HPC, data analytics and hardware-enhanced security/cryptography workloads. Along with numerous acceleration refinements, the new processor offers integrated 100 Gb/s Intel® Omni-Path Architecture fabric options. With these improvements, the Intel Xeon Scalable Platinum 8180 processor yielded an increase of up to 8.2x more double precision GFLOPS/sec when compared to Intel Xeon processor E5-2690 (codenamed Sandy Bridge) common in the server installed-base, and a 2.27x increase over the previous-generation Intel Xeon processor E5-2699 v4 (codenamed Broadwell) . The doubling of cores in the publicly available POD HPC cloud resources in 2017 was preceded by a 50 percent increase in capacity in 2016. As customer demand continues to increase, POD HPC cloud will continue to grow using the most current technologies to deliver the actionable insights that organizations require. Visit Penguin Computing at Booth 1801 during SC17 in Denver.

About Penguin Computing

Penguin Computing is one of the largest private suppliers of enterprise and high-performance computing solutions in North America and has built and operates the leading specialized public HPC cloud service, Penguin Computing On-Demand (POD). Penguin Computing pioneers the design, engineering, integration and delivery of solutions that are based on open architectures and comprise non-proprietary components from a variety of vendors. Penguin Computing is also one of a limited number of authorized Open Compute Project (OCP) solution providers leveraging this Facebook-led initiative to bring the most efficient open data center solutions to a broader market, and has announced the Tundra product line which applies the benefits of OCP to high performance computing. Penguin Computing has systems installed with more than 2,500 customers in 40 countries across eight major vertical markets.

Visit www.penguincomputing.com to learn more about the company and follow [@PenguinHPC](https://twitter.com/PenguinHPC) on Twitter.

Penguin Computing, Scyld ClusterWare, Scyld Insight, Scyld Cloud Workstation, Scyld Cloud Manager, Relion, Altus, Penguin Computing On-Demand, Tundra, Arctica and FrostByte are trademarks or registered trademarks of Penguin Computing, Inc.

1. Baseline config: 1-Node, 2 x Intel® Xeon® Processor E5-2699 v4 on Red Hat Enterprise Linux* 7.0 kernel 3.10.0-123 using Intel® Distribution for LINPACK Benchmark, score: 1446.4 GFLOPS/s vs. estimates based on Intel internal testing on 1-Node, 2x Intel Xeon Scalable processor (codename Skylake-SP) system. Score: 3295.57

Baseline config: 1-Node, 2 x Intel® Xeon® Processor E5-2690 based system on Red Hat Enterprise Linux* 6.0 kernel version 2.6.32-504.el6.x86_64 using Intel® Distribution for LINPACK Benchmark. Score: 366.0 GFLOPS/s vs. 1-Node,

2 x Intel® Xeon® Scalable Platinum 8180 processor on Ubuntu 17.04 using MKL 2017 Update 2. Score: 3007.8

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

