

Intel Reduces High Performance Computing Complexity and Fuels Artificial Intelligence Insight

Nov. 15, 2016 — As businesses and industries increasingly adopt High Performance Computing (HPC) to tackle complex challenges – ranging from weather forecasting to product development to diagnosing diseases – there is a growing need to reduce complexity, balance performance and scale to needs. Additionally, organizations are seeking to lower the barriers of using artificial intelligence (AI) across many commercial, government and academic applications. Today at Supercomputing 2016 in Salt Lake City, Utah, Intel announced it is helping solve many of these issues with key updates to its HPC portfolio of products and upcoming technologies.

Intel® HPC Orchestrator

To help optimize performance by balancing compute, memory, storage, energy efficiency and other capabilities, Intel has designed the Intel® Scalable System Framework (Intel® SSF). Intel SSF is an advanced architectural approach for simplifying the procurement, deployment and management of HPC systems. It is here that Intel is leading a broad-based effort to transform the ecosystem, simplifying system management through the delivery of a reliable, modular, integrated and validated HPC software platform.

- Available now, [Intel® HPC Orchestrator](#) is a software platform that HPC system builders, application developers and end users can use to gain efficiencies in the setup and management of systems, enable testing at scale and simplify deployment for users. Intel HPC Orchestrator is a customizable platform that supports traditional HPC workloads, as well as emerging uses such as machine learning, high-performance data analysis and autonomous driving.
- Based on the [OpenHPC](#) community system software stack, Intel HPC Orchestrator includes multiple options of compiler, MPI family, and base OS combinations along with tools to assist with install automation. The integrated offering also includes Intel® Parallel Studio XE Cluster Edition with a 90-day evaluation license, providing ease of installation from initial download.
- As a key component of Intel SSF, Intel HPC Orchestrator includes more than 60 pre-integrated modular components and features ongoing updates and technical support. This allows HPC system developers and users to focus their technical resources on enhancing and differentiating their HPC solutions.
- Intel HPC Orchestrator is available through OEM launch system suppliers Dell* and Fujitsu*, with more availability announcements expected soon.

Intel Scores Big on Top500 List; Sees 2.5 Times More Intel® Omni-Path Architecture Systems

In the nine months since Intel® Omni-Path Architecture (Intel® OPA) began shipping, it has become the standard fabric for 100 gigabit (Gb) systems. Intel OPA is featured in 28 of the top 500 most powerful supercomputers in the world announced at Supercomputing 2016 and now has 66 percent of the 100Gb market¹. Top500 designs include Oakforest-PACS, MIT Lincoln Lab and CINECA.

- Intel OPA has been deployed in 28 clusters in the November 2016 Top500 Supercomputer list, which is twice the number of InfiniBand* EDR systems and now accounts for around 66 percent of all 100GB systems. Additionally, two systems are ranked in the Top 15: Oakforest-PACS is

ranked sixth with 8,208 nodes and CINECA is ranked 12th with 3556 nodes. The Intel OPA systems on the list add up to total floating-point operations per second (FLOPS) of 43.7 petaflops (Rmax), or 2.5 times the FLOPS of all InfiniBand* EDR systems.

- Intel OPA has seen rapid market adoption in the nine months it has been shipping broadly, driven by clear customer benefits such as high performance, price-performance and innovative fabric features, such as error detection and correction without additional latency.
- Intel OPA is an end-to-end fabric solution that improves HPC workloads for clusters of all sizes, achieving up to 9 percent higher application performance and up to 37 percent lower fabric costs on average compared to InfiniBand EDR.²
- Major installations include the University of Tokyo and Tsukuba University (JCAHPC), Texas Tech University, the University of Washington, the University of Colorado Boulder, MIT Lincoln Lab, and Met Malaysia. There are now well over 100 successfully-deployed Intel OPA clusters and most adoption is a result of competitive benchmarking and leadership price-performance.

Accelerate Time to Value for AI with Intel® Deep Learning Inference Accelerator

Intel-based platforms are already powering virtually all HPC and machine learning workloads due to industry-leading performance and a robust software ecosystem. Intel is broadening its AI solution portfolio in 2017 with Intel® Deep Learning Inference Accelerator, a field-programmable gate array (FPGA)-based integrated hardware and software solution for neural network acceleration.

- Intel® Deep Learning Inference Accelerator offers a scalable acceleration solution for image recognition applications, bringing exceptional throughput and power efficiency to data centers so users can deploy deep learning solutions in a matter of minutes.
- The hardware features an Intel® Arria® 10 FPGA add-in card optimized for targeted topologies of convolutional neural networks (CNN). Its FPGA-based architecture enables remote updates so the hardware keeps up with the rapid pace of innovation in AI.
- Intel Deep Learning Inference Accelerator is software-programmable through industry standard frameworks, such as Intel® Distribution for Caffe* and Intel® Math Kernel Library for Deep Neural Networks.
- Intel Deep Learning Inference Accelerator will be available in 2017.

Expanding Intel Scalable System Framework Ecosystem

Intel and ecosystem partners continue to expand the Intel SSF ecosystem with additional large system announcements, optimized solutions, new tools and resources, and a new Intel SSF program.

- Recently, the following companies and organizations have adopted Intel SSF:
 - Leading universities and research centers, including MIT Lincoln Lab, which deployed Intel SSF in its Top500 Supercomputer entry.
 - Hewlett Packard Enterprise* (HPE) integrated Intel SSF into its release of the HPE Next Generation Sequencing Solution to increase the speed and cost-effectiveness of genomic sequencing.
 - Dell integrated Intel SSF into Dell HPC System for Life Sciences to accelerate time-to-insight for a range of fields, including drug design, cancer research, agriculture and forensics.

- Intel released an update to Intel® Cluster Checker 2017. This comprehensive diagnostic tool is distributed in Intel® Parallel Studio XE 2017 Cluster Edition and supports Intel® Xeon™ processors, Intel® Xeon Phi™ processors, Intel OPA, and Intel® Enterprise Edition for Lustre* software.

Intel Xeon Phi Processor Family Updates

- Intel Xeon Phi processor 7210, 7230, 7250 and 7290 are now shipping broadly and processors with integrated Intel OPA are on track to ship broadly this month.
- More than 50 HPC designs on Intel Xeon Phi processors are in progress across the industry.
- The latest Intel Xeon Phi processor was selected for nine new systems on the Top500, including two systems in the Top 10, with the Cori system ranking fifth and Oakforest-PAC system ranking sixth. Additionally, Intel was the prime contractor supporting the Collaboration of Oak Ridge, Argonne and Lawrence Livermore (CORAL) in a Top 20 system, the Theta system which will help to advance U.S. leadership in scientific research.³
- The Developer Access Program for Intel Xeon Phi processor-based platforms continues to gain traction as Intel expands worldwide to include collaborations in Japan, Korea, Europe and the U.S. Learn more [here](#).
- The Intel Xeon Phi coprocessor will be available in early 2017.

Continued Intel® Xeon® Processor Innovation

Intel® Xeon® platforms are already powering virtually all HPC and machine learning workloads and to push that further, Intel is introducing its fastest two-socket processor, the Intel® Xeon® processor E5-2699A v4.

- The Intel® Xeon® processor E5-2699A v4 delivers industry-leading performance for demanding, data-intensive applications ideal for customers who need high performance from the world's most-chosen data center processor family.
- At Supercomputing 2016, Intel will publically demonstrate an HPC cluster featuring an early version of the next-generation Intel Xeon processor utilizing integrated Intel OPA and Intel® Advanced Vector Extensions-512 (Intel® AVX-512), complemented by Intel Xeon Phi processors.
- Intel AVX-512 processor instructions, launched June 2016 with the Intel Xeon Phi processor, accelerates compression and encryption algorithms. It provides hardware-enhanced data protection and fast time-to-insight, and can accelerate floating-point operations used in a wide range of HPC applications.

For more information, visit the Intel [Supercomputing 2016 press kit](#).

CONTACT: Dan Francisco
(916) 812-8814
daniel.francisco@intel.com

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com/ssf.

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 information, go to intel.com/benchmarks.

Cost reduction scenarios described are intended as examples of how a given Intel- based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

© 2016 Intel Corporation.

Intel and the Intel logo, Xeon and Xeon Phi are trademarks of Intel Corporation in the United States and other countries

* Other names and brands may be claimed as the property of others.

¹ <https://www.top500.org/lists/2016/11/>

² Configuration for WIEN2k version 14.2 lapw1c_mpi benchmark, GROMACS version 5.0.4 ion_channel benchmark, NWChem release 6.6 Siosi5 benchmark, LS-DYNA MPP R8.1.0 3cars benchmark, ANSYS Fluent v17.0 rotor_3m benchmark, NAMD 2.11 stmv benchmark, Quantum Espresso version 5.3.0 ausurf112 benchmark, CD-adapco STAR-CCM+® version 11.04.010 lemanx_poly 17m benchmark, LAMMPS Feb 16, 2016 stable version release rhodopsin protein benchmark, WRF version 3.5.1 conus2.5km benchmark, Spec MPI2007 Large suite (Intel Internal measurements marked estimates until published): Intel® Xeon® Processor E5-2697A v4 dual socket servers. 64 GB DDR4 memory per node, 2133 MHz. RHEL 7.2. BIOS settings: Snoop hold-off timer = 9, Early snoop disabled, Cluster on die disabled. IOU Non-posted prefetch disabled. Intel® Omni-Path Architecture (Intel® OPA): Intel Fabric Suite 10.0.1.0.50. Intel Corporation Device 24f0 – Series 100 HFI ASIC (Production silicon). OPA Switch: Series 100 Edge Switch – 48 port (Production silicon). EDR Infiniband: MLNX_OFED_LINUX-3.2-2.0.0.0 (OFED-3.2-2.0.0). Mellanox EDR ConnectX-4 Single Port Rev 3 MCX455A HCA. Mellanox SB7700 - 36 Port EDR Infiniband switch. Configuration for MiniFE 2.0, VASP (developer branch), GaAsBI-64 benchmark: Intel® Xeon® Processor E5-2697 v4 dual socket servers. 128 GB DDR4 memory per node, 2400 MHz. RHEL 6.5. Snoop hold-off timer = 9. Intel® OPA: Intel Fabric Suite 10.0.1.0.50. Intel Corporation Device 24f0 – Series 100 HFI ASIC (Production silicon). OPA Switch: Series 100 Edge Switch – 48 port (Production silicon). IOU Non-posted prefetch disabled. 2). Mellanox EDR based on internal measurements: Mellanox EDR ConnectX-4 Single Port Rev 3 MCX455A HCA. Mellanox SB7700 - 36 Port EDR Infiniband switch. 37% lower fabric cost: Configuration assumes full bisectional bandwidth (FBB) Fat-Tree configurations for all calculated clusters. All cluster configurations, in 12 node increments, are estimated via internal Intel configuration tool. Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction. Intel® and Mellanox component pricing from www.kernelsoftware.com, with prices as of October 20, 2016.

³ <https://www.top500.org/lists/2016/11/>