INTEL NEWS FACT SHEET

Bringing the Power of HPC Mainstream

Nov. 16, 2015 — Intel Corporation today announced several new and expanded high performance computing (HPC) products and programs designed to bring HPC system capabilities and benefits to more industries and workloads.

Holistic Approach to HPC System Designs

- Intel® Scalable System Framework (Intel® SSF) is an advanced architectural approach for simplifying the procurement, deployment and management of HPC systems while broadening the accessibility of HPC to more industries and workloads.
- The framework combines next-generation Intel® Xeon® processors and Intel® Xeon Phi™ processors, Intel® OPA, silicon photonics, innovative memory technologies, and the Intel® Lustre* parallel file system, along with the ability to efficiently integrate them into a broad spectrum of system solutions. The framework also provides a ubiquitous and standards-based programming model, extending the ecosystem's current investments in existing code for future generations.
- Intel will provide reference architectures, reference designs and system specifications, including hardware and software bill of materials, for Intel SSF validated systems.
- Colfax*, Cray*, Dell*, Fujitsu Systems Europe*, HPE*, Inspur*, Lenovo*, Penguin Computing*, SGI*, Sugon* and Supermicro* to announce plans to launch systems based on the Intel Scalable System Framework early next year.
- Leading software vendors expected to announce support of the Intel SSF include Altair*, Ansys*, Dassault Systemes SIMULIA* and MSC Software*.
- To ensure strong alignment and collaboration with the recently announced U.S. National Strategic Computing Initiative (NSCI) to spur the creation and deployment of HPC technology, Intel has created an Extreme Scale Program Office that will fully leverage the Intel SSF capabilities for advancing the NSCI's goals.
- Intel and the Barcelona Supercomputing Center (BSC) agreed to evaluate and demonstrate the many new exascale-class capabilities of Intel SSF. Collaboration areas include Intel's upcoming Intel® Optane™ non-volatile memory technology, Intel Lustre-based scalable storage units and contribution of BSC performance analysis tools to the OpenHPC system software stack.

The Future of High Performance Fabrics

- Launch of Intel® Omni-Path Architecture (Intel® OPA), an end-to-end fabric solution that cost-effectively improves the performance of HPC applications for entry-level to large-scale HPC clusters.
- Intel OPA's 48-port switch enables up to 26 percent more servers than InfiniBand* EDR within the same budget[1] for a more efficient switch and system designs and an accelerated time to discovery and innovation.
- Intel OPA enables up to 17 percent lower Message Passing Interface (MPI) latency[2] and 7 percent higher MPI messaging rate than InfiniBand EDR[3] for greater overall system performance.

- Intel OPA provides up to 60 percent lower power consumption than InfiniBand EDR solutions based on higher-efficiency components and an overall reduction in fabric infrastructure with the 48-port switch chip[4].
- Innovative new fabric features that deliver performance and resiliency without compromise such as packet integrity protection, with no-latency penalty for error detection[5], and traffic flow optimization for significantly lower MPI latency in mixed storage and MPI traffic environments[6].
- Intel OPA, an end-to-end solution, is composed of PCIe host adapters, edge switches, director switches, cabling and open source software tools. It will also be available as an integrated device in Intel Xeon Phi processors (code-named Knights Landing) and future 14nm Intel Xeon processors.
- Intel OPA is currently being used at several large sites, including the Texas Advanced Computing Center and the Pittsburgh Supercomputer Center.
- Colfax, Cray, Dell, Fujitsu Systems Europe, Hitachi*, Huawei*, HPE, Inspur, Lenovo, NEC*, SGI, Sugon, Supermicro and other system vendors starting to announce Intel OPA-based switches and server platforms, with volume shipments ramping in the first quarter of next year.

Intel and HPE Collaboration to Expand HPC Access

- In July, Intel and HPE announced a new HPC Datacenter Specialty designation providing qualified system channel sellers access to exclusive HPC research, technical support, training and tools.
- Intel officially launched the program in North America and Europe in October and 65 channel hardware vendors have achieved the HPC Datacenter Specialty designation. The number of designees continues to grow and more will join when the program is launched in Asia next year.
- HPC Datacenter Specialty designation requires vendors meet several minimum qualifications, including the completion of an extensive training program and demonstrated experience deploying clusters of a minimum size and quantity.

Modern Code Developer Program

- Since the introduction of the Intel® Modern Code Developer Community at International Supercomputing earlier this year, Intel has reached over 500,000 developers with Modern Code tools, trainings and messaging.
- Intel and CERN openlab* created the Intel® Modern Code Developer Challenge to spur advancement in the use of modern coding techniques and the science it supports while encouraging students to pursue careers in the field of high performance computing.
- The challenge attracted more than 17,000 students representing over 130 universities across 19 countries. Over a thousand students downloaded the code and took advantage the associated trainings as part of the challenge. These students had remote access to clusters based on Intel Xeon processors and Intel Xeon Phi coprocessors.
- Intel and CERN openlab jointly announced the winners of the challenge at the annual Intel® HPC Developer Conference, held in conjunction with SC15.
- The grand prize winner of a nine-week internship at CERN was able to speed up a brain development code over 320x, reducing the runtime of a large dataset from 45 hours to less than nine minutes.

INTEL NEWS FACT SHEET

Intel Solutions for Lustre Software

- Launched five new Intel Parallel Computing Centers around the world that will specialize in improving the performance and ease of use of Intel Lustre software.
- Intel expanded its Lustre software capabilities with the release of Intel® Enterprise Edition for Lustre software v2.4 and Intel® Cloud Edition for Lustre software v1.2.

About Intel
Intel (NASDAQ: INTC) is a world leader in computing innovation. The company designs and builds the essential technologies that serve as the foundation for the world's computing devices. As a leader in corporate responsibility and sustainability, Intel also manufactures the world's first commercially available "conflict-free" microprocessors. Additional information about Intel is available at newsroom.intel.com and blogs.intel.com, and about Intel's conflict-free efforts at conflictfree.intel.com.

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 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 http://www.intel.com.

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 **http://www.intel.com/performance**.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice Revision #20110804

[1] Based on Intel projects that assume a 750-node cluster, and number of switch chips required is based on a full bisectional bandwidth (FBB) Fat-Tree configuration. Intel® OPA uses one fully-populated 768-port director switch, and Mellanox EDR solution uses a combination of 648-port director switches and 36-port edge switches. Mellanox component pricing from www.kernelsoftware.com, with prices as of November 3, 2015. Compute node pricing based on Dell PowerEdge R730 server from www.dell.com, with prices as of May 26, 2015. Intel® OPA pricing based on estimated reseller pricing based on expected Intel MSRP pricing, to be available on ark.intel.com in November 2015.
[2] Tests performed by Intel on Intel® Xeon® Processor E5-2697v3 dual-socket servers with 2133 MHz DDR4 memory. 16 servers/448 MPI ranks. Intel® Turbo Boost Technology enabled and Intel® Hyper-Threading Technology disabled. 28 MPI ranks per node. HPCC 1.4.3 Random order ring latency. OPA: Open MPI 1.10.0 with PSM2. Pre-production Intel Corporation Device 24f0 – Series 100 HFI ASIC, Series 100 Edge Switch – 48 port. IOU Non-posted Prefetch disabled in BIOS. Mellanox EDR based on internal measurements. EDR: Open MPI 1.8-mellanox released with hpcx-v1.3.336-icc-MLNX_OFED_LINUX-3.0-1.0.1-redhat6.6-x86_64.tbz. Mellanox EDR ConnectX-4 Single Port Rev 3 MCX455A HCA. Mellanox SB7700 - 36 Port EDR Infiniband switch.

[3] Tests performed by Intel on Intel® Xeon® Processor E5-2697v3 dual-socket servers with 2133 MHz DDR4 memory. Intel® Turbo Boost Technology enabled and Intel® Hyper-Threading Technology disabled. Intel OPA: Open MPI 1.10.0 with PSM2. Pre-production Intel Corporation Device 24f0 – Series 100 HFI ASIC, Series 100 Edge Switch – 48 port. IOU Non-posted Prefetch disabled in BIOS. EDR: Open MPI 1.8-mellanox released with hpcx-v1.3.336-icc-MLNX_OFED_LINUX-3.0-1.0.1-redhat6.6-x86_64.tbz. Mellanox EDR ConnectX-4 Single Port Rev 3 MCX455A HCA. Mellanox SB7700 - 36 Port EDR Infiniband switch. 17% claim: HPCC 1.4.3 Random order ring latency. 16 nodes, 28 MPI ranks per node. 7% message rate claim: Ohio State Micro Benchmarks v. 4.4.1. osu_mbw_mr 28 MPI ranks per node, 8 byte message. osu_mbw_mr modified to use maximum rank pair communication time instead of average rank pair communication time.

[4] Based on Intel projections that assume 750-node cluster, and number of switch chips required is based on a full bisectional bandwidth (FBB) Fat-Tree configuration. Intel® OPA uses one fully-populated 768-port director switch, and Mellanox EDR solution uses a combination of director switches and edge switches. Mellanox power data based on Mellanox CS7500 Director Switch, Mellanox SB7700/SB7790 Edge switch, and Mellanox ConnectX-4 VPI adapter card installation documentation posted on www.mellanox.com as of November 1, 2015. Intel OPA power data based on product briefs posted on www.intel.com as of November 16, 2015. Intel® OPA pricing based on estimated reseller pricing based on expected Intel MSRP pricing, to be available on ark.intel.com in November 2015.

[5] A CRC check is performed on every Link Transfer Packet (LTP, 1056-bits) transmitted through a switch hop as defined by the Intel® OPA wire protocol, so the stated switch latencies always include error detection by definition.

[6] Based on preliminary Intel internal testing using two pre-production Intel® OPA edge switches with one inter-switch link, comparing MPI latency over multiple iterations with varying bandwidth allocations for storage and MPI traffic over multiple virtual lanes, both with Traffic Flow Optimization enabled and disabled.