

## News Byte

June 17, 2019

[Contact Intel PR](#)

Dr. Rajeeb Hazra, Intel corporate vice president and general manager of the Enterprise and Government Group, deliver Intel's opening keynote at the 2019 International Supercomputing Conference on June 17, 2019, in Frankfurt, Germany. He introduces Intel's latest high-performance computing products and technological breakthroughs.

**What's New:** At the 2019 International Supercomputing Conference (ISC'19), Intel provided the latest updates on how its data-centric portfolio is transforming next-generation high-performance computing (HPC) systems and propelling the industry toward exascale computing.

*"In today's data-centric world, Intel continues to drive innovation and performance. Our portfolio of HPC solutions, including Intel® Xeon® Platinum 9200 processors and Intel® Optane® DC persistent memory, enables unprecedented explorations in science and discovery."*

*–Trish Damkroger, Intel vice president and general manager of the Extreme Computing Organization*

**Why Today's News Matters:** Today's advanced HPC systems leverage both traditional HPC data analytics and artificial intelligence (AI) technologies to efficiently address the most complex scientific problems. As the need for more computing performance increases in HPC systems, memory bandwidth is increasingly a bottleneck. Intel's disclosures at ISC'19 demonstrate how its data-centric portfolio addresses the unique challenges of HPC systems by bringing together HPC data analytics and AI acceleration into a single computing environment, while delivering a new memory and storage paradigm that feeds the compute engine.

**What Updates Did Intel Provide at ISC'19:** Intel's data-centric portfolio provides the foundation for system architects to build advanced computing systems that move, store and process massive amounts of data. For systems that demand the highest computing performance, Intel announced today that pre-configured systems based on Intel® Xeon® Platinum 9200 processors are available for purchase from select OEMs, including Atos\*, HPE\*, Lenovo\*, Penguin Computing\*, Megware\* and authorized Intel resellers. An extension of the 2nd generation Intel Xeon Scalable processor family, the Intel Xeon Platinum 9200 processor series is architected to deliver leadership performance for workloads and usages such as scientific simulations, financial analytics, artificial intelligence/deep learning, 3D modeling and analysis, cryptography, and data compression. Intel Xeon Platinum 9200 processors feature integrated Intel® Deep Learning Boost

(Intel DL Boost) technology to accelerate AI performance up to 30x<sup>1</sup> over the previous generation Xeon Scalable processor at launch.

**What Else was Announced at ISC'19:** Also at ISC'19, Intel, together with ECMWF\*, EPCC\*<sup>4</sup>, Fujitsu\*, Arctur\* and the other partners of the NEXTGenIO project<sup>2</sup>, disclosed the latest breakthrough performance results using Intel Optane DC persistent memory across various supercomputing applications.

- The European Centre for Medium-Range Weather Forecasts (ECMWF) achieved 10-times<sup>3</sup> higher bandwidth when its Fields Database, which holds the meteorological data for medium-range weather forecasts, was stored in persistent memory and distributed across multiple computing nodes. Use of compute nodes equipped with Intel Optane DC persistent memory accelerated ECMWF's global weather forecasts and a reduced number of I/O nodes needed to run its models.
- The Arctur HPC center, in partnership with Barcelona Supercomputer Center, achieved a 2x<sup>3</sup> speed-up in simulating 3D models of an electric lightweight aircraft, which reduced its OpenFOAM runtimes by 50%<sup>3</sup> on 16 nodes.
- EPCC achieved 2-times<sup>3</sup> higher throughput on the CASTEP\* material science application when running its code on computing nodes equipped with Intel Optane DC persistent memory, accelerating material science research across multiple domains.

Intel is further accelerating the adoption of Intel Optane DC persistent memory in HPC systems. Intel announced a revolutionary new storage architecture for supercomputing that leverages Intel Optane DC persistent memory and Distributed Asynchronous Object Storage (DAOS). DAOS is an open-source software-defined scale-out object store providing high-bandwidth, low-latency and high I/O operations, and is designed specifically for convergence of HPC and AI workloads. This new software-defined storage engine eliminates several limitations of today's parallel file systems.

Susan Coghlan, ALCF-X\* project director/exascale computing systems deputy director, said: "The Argonne Leadership Computing Facility will be the first major production deployment of the DAOS storage system as part of Aurora, the first U.S. exascale system coming in 2021. The DAOS storage system is designed to provide the levels of metadata operation rates and bandwidth required for I/O extensive workloads on an exascale-level machine."

Intel also disclosed at ISC'19 more information about its One API project, which will deliver a unified programming model to simplify application development across diverse computing architectures. Intel's One API will be based on industry standards and open specifications and will be interoperable with OpenMP\*, MPI\* and Fortran\* among others.

#### More Context: [Data Center News](#)

##### The Small Print:

<sup>1</sup> Up to 30X AI performance with Intel® DL Boost compared to Intel® Xeon® Platinum 8180 processor when launched (July 2017). Tested by Intel as of 2/26/2019. Platform: Dragon rock 2 socket Intel® Xeon® Platinum 9282(56 cores per socket), HT ON, turbo ON, Total Memory 768 GB (24 slots/ 32 GB/ 2933 MHz), BIOS:SE5C620.86B.0D.01.0241.112020180249, Centos 7 Kernel 3.10.0-957.5.1.el7.x86\_64, Deep Learning Framework: Intel® Optimization for Caffe version: <https://github.com/intel/caffe> d554cbf1, ICC 2019.2.187, MKL DNN version: v0.17 (commit hash: 830a10059a018cd2634d94195140cf2d8790a75a), model: [https://github.com/intel/caffe/blob/master/models/intel\\_optimized\\_models/int8/resnet50\\_int8\\_full\\_conv.prototxt](https://github.com/intel/caffe/blob/master/models/intel_optimized_models/int8/resnet50_int8_full_conv.prototxt), BS=64, No datalayer DummyData:3x224x224, 56 instance/2 socket, Datatype: INT8 vs Tested by Intel as of July 11th 2017: 2S Intel® Xeon® Platinum 8180 CPU @ 2.50GHz (28 cores), HT disabled, turbo disabled, scaling governor set to "performance" via intel\_pstate driver, 384GB DDR4-2666 ECC RAM. CentOS Linux release 7.3.1611 (Core), Linux kernel 3.10.0-514.10.2.el7.x86\_64. SSD: Intel® SSD DC S3700 Series (800GB, 2.5in SATA 6Gb/s, 25nm, MLC). Performance measured with: Environment variables: KMP\_AFFINITY='granularity=fine, compact', OMP\_NUM\_THREADS=56, CPU Freq set with cpupower frequency-set -d 2.5G -u 3.8G -g performance. Caffe: (<http://github.com/intel/caffe/>), revision f96b759f71b2281835f690af267158b82b150b5c. Inference measured with "caffe time -forward\_only" command, training measured with "caffe time" command. For "ConvNet" topologies, dummy dataset was used. For other topologies, data was stored on local storage and cached in memory before training. Topology specs from [https://github.com/intel/caffe/tree/master/models/intel\\_optimized\\_models](https://github.com/intel/caffe/tree/master/models/intel_optimized_models) (ResNet-50), Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with "numactl -l".

<sup>2</sup> The NEXTGenIO project is funded by the European Union's Horizon 2020 Research and Innovation program under Grant Agreement no. 671951.

<sup>3</sup> System Configuration details provided by EPCC:

- 34 DP nodes with Intel Xeon Scalable processor 8260M CPUs (A0 stepping), Fujitsu mainboard
- 96 GByte DDR4 DRAM per socket (6x16 GByte DIMMs, 2666 speed grade), plus 1.5 TByte Intel Optane DC Persistent Memory (6x256 GByte DIMMs, QS)
- Dual-Rail Omni-Path networks (2 OPA NICs per node) connected via 2 48-port OPA switches
- Two additional Storage server nodes running Lustre

<sup>4</sup> EPCC is the Advanced Computing Facility located at the University of Edinburgh

Tags: [HPC](#), [Intel Xeon](#), [supercomputing](#)

## Other News



April 8, 2021

[SD Supercomputer Center Selects Habana, Intel for Efficient AI](#)

April 6, 2021

[Intel Xeon Advances Nasdaq's Homomorphic Encryption R&D](#)

April 6, 2021

[3rd Gen Intel Xeon Scalable Launch](#)

### About Intel

Intel (Nasdaq: INTC) is an industry leader, creating world-changing technology that enables global progress and enriches lives. Inspired by Moore's Law, we continuously work to advance the design and manufacturing of semiconductors to help address our customers' greatest challenges. By embedding intelligence in the cloud, network, edge and every kind of computing device, we unleash the potential of data to transform business and society for the better. To learn more about Intel's innovations, go to [newsroom.intel.com](https://newsroom.intel.com) and [intel.com](https://intel.com).


© Intel Corporation. Intel, the Intel logo and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

## Latest News: Artificial Intelligence



April 8, 2021

[SD Supercomputer Center Selects Habana, Intel for Efficient AI](#)



April 6, 2021

[Intel Launches Its Most Advanced Performance Data Center Platform](#)



April 1, 2021

[At John Deere, 'Hard Iron Meets Artificial Intelligence'](#)

[Read More](#)