



New Intel XPU Innovations Target HPC and AI

Sapphire Rapids with built-in HBM raises performance bar; Intel's GPU, networking and storage capabilities enhance HPC toolbox.

NEWS HIGHLIGHTS

- Latest 3rd Gen Intel® Xeon® Scalable processors will power the next generation of supercomputers and high-performance computing systems.
- The next generation of Intel Xeon Scalable processors (code-named "Sapphire Rapids") will integrate High Bandwidth Memory (HBM).
- Intel's X^e-HPC-based Ponte Vecchio GPU has powered-on, is in system validation, and will include OAM form factor and subsystems.
- Intel announces High Performance Networking (HPN) with Ethernet, extending Intel Ethernet into HPC applications.
- Intel introduces commercial support for DAOS (distributed application object storage).

SANTA CLARA, Calif.--(BUSINESS WIRE)-- At the 2021 International Supercomputing Conference (ISC) Intel is showcasing how the company is extending its lead in high performance computing (HPC) with a range of technology disclosures, partnerships and customer adoptions. Intel® processors are the most widely deployed compute architecture in the world's supercomputers, enabling global medical discoveries and scientific breakthroughs. Intel is announcing advances in its Xeon processor for HPC and AI as well as innovations in memory, software, exascale-class storage, and networking technologies for a range of HPC use cases.

More: [Intel Data Center News](#) | [Intel's HPC GM Trish Damkroger Keynotes 2021 ISC](#) (Keynote Replay) | ["Accelerating the Possibilities with HPC"](#) (Keynote Presentation)

"To maximize HPC performance we must leverage all the computer resources and technology advancements available to us," said Trish Damkroger, vice president and general manager of High Performance Computing at Intel. "Intel is the driving force behind the industry's move toward exascale computing, and the advancements we're delivering with our CPUs, XPUs, oneAPI Toolkits, exascale-class DAOS storage, and high-speed networking are pushing us closer toward that realization."

Advancing HPC Performance Leadership

Earlier this year, Intel extended its leadership position in HPC with the launch of 3rd Gen Intel Xeon Scalable processors. The latest processor delivers up to 53% higher performance across a range of HPC workloads, including life sciences, financial services and manufacturing, as compared to the previous generation processor.

Compared to its closest x86 competitor, the 3rd Gen Intel Xeon Scalable processor delivers better performance across a range of popular HPC workloads. For example, when comparing a Xeon Scalable 8358 processor to an AMD EPYC 7543 processor, NAMD performs 62% better, LAMMPS performs 57% better, RELION performs 68% better, and Binomial Options performs 37% better. In addition, Monte Carlo simulations run more than two times faster, allowing financial firms to achieve pricing results in half the time. Xeon Scalable 8380 processors also outperform AMD EPYC 7763 processors on key AI workloads, with 50% better performance across 20 common benchmarks. HPC labs, supercomputing centers, universities and original equipment manufacturers who have adopted Intel's latest compute platform include Dell Technologies, HPE, Korea Meteorological Administration, Lenovo, Max Planck Computing and Data Facility, Oracle, Osaka University and the University of Tokyo.

Integration of High Bandwidth Memory within Next-Gen Intel Xeon Scalable Processors

Workloads such as modeling and simulation (e.g., computational fluid dynamics, climate and weather forecasting, quantum chromodynamics), artificial intelligence (e.g., deep learning training and inferencing), analytics (e.g., big data analytics), in-memory databases, storage and others power humanity's scientific breakthroughs. The next-generation of Intel Xeon Scalable processors (code-named "Sapphire Rapids") will offer integrated High Bandwidth Memory (HBM), providing a dramatic boost in memory bandwidth and a significant performance improvement for HPC applications that operate memory bandwidth-sensitive workloads. Users can power through workloads using just High Bandwidth Memory or in combination with DDR5.

Customer momentum is strong for Sapphire Rapids processors with integrated HBM, with early leading wins such as the U.S. Department of Energy's Aurora supercomputer at Argonne National Laboratory and the Crossroads supercomputer at Los Alamos National Laboratory.

"Achieving results at exascale requires the rapid access and processing of massive amounts of data," said Rick Stevens, associate laboratory director of Computing, Environment and Life Sciences at Argonne National Laboratory. "Integrating high-bandwidth memory into Intel Xeon Scalable processors will significantly boost Aurora's memory bandwidth and enable us to leverage the power of artificial intelligence and data analytics to perform advanced simulations and 3D modeling."

Charlie Nakhleh, associate laboratory director for Weapons Physics at Los Alamos National Laboratory, said: "The Crossroads supercomputer at Los Alamos National Labs is designed to advance the study of complex physical systems for science and national security. Intel's next-generation Xeon processor Sapphire Rapids, coupled with High Bandwidth Memory, will significantly improve the performance of memory-intensive workloads in our Crossroads system. The [Sapphire Rapids with HBM] product accelerates the largest complex physics and engineering calculations, enabling us to complete major research and development responsibilities in global security, energy technologies and economic competitiveness."

The Sapphire Rapids-based platform will provide unique capabilities to accelerate HPC, including increased I/O bandwidth with PCI express 5.0 (compared to PCI express 4.0) and Compute Express Link (CXL) 1.1 support, enabling advanced use cases across compute,

networking and storage.

In addition to memory and I/O advancements, Sapphire Rapids is optimized for HPC and artificial intelligence (AI) workloads, with a new built-in AI acceleration engine called Intel® Advanced Matrix Extensions (AMX). Intel AMX is designed to deliver significant performance increase for deep learning inference and training. Customers already working with Sapphire Rapids include CINECA, Leibniz Supercomputing Centre (LRZ) and Argonne National Lab, as well as the Crossroads system teams at Los Alamos National Lab and Sandia National Lab.

Intel X^e-HPC GPU (Ponte Vecchio) Powered On

Earlier this year, Intel powered on its X^e-HPC-based GPU (code-named “Ponte Vecchio”) and is in the process of system validation. Ponte Vecchio is an X^e architecture-based GPU optimized for HPC and AI workloads. It will leverage Intel’s Foveros 3D packaging technology to integrate multiple IPs in-package, including HBM memory and other intellectual property. The GPU is architected with compute, memory, and fabric to meet the evolving needs of the world’s most advanced supercomputers, like Aurora. Ponte Vecchio will be available in an OCP Accelerator Module (OAM) form factor and subsystems, serving the scale-up and scale-out capabilities required for HPC applications.

Extending Intel Ethernet For HPC

At ISC 2021, Intel is also announcing its new High Performance Networking with Ethernet (HPN) solution, which extends Ethernet technology capabilities for smaller clusters in the HPC segment by using standard Intel Ethernet 800 Series Network Adapters and Controllers, switches based on Intel® Tofino™ P4-programmable Ethernet switch ASICs and the Intel® Ethernet Fabric suite software. HPN enables application performance comparable to InfiniBand at a lower cost while taking advantage of the ease of use offered by Ethernet.

Commercial Support for DAOS

Intel is introducing commercial support for DAOS (distributed application object storage), an open-source software-defined object store built to optimize data exchange across Intel HPC architectures. DAOS is at the foundation of the Intel® Exascale storage stack, previously announced by Argonne National Laboratory, and is being used by Intel customers such as LRZ and JINR (Joint Institute for Nuclear Research).

DAOS support is now available to partners as an L3 support offering, which enables partners to provide a complete turnkey storage solution by combining it with their services. In addition to Intel’s own data center building blocks, early partners for this new commercial support includes HPE, Lenovo, Supermicro, Brightskies, Croit, Nettrix, Quanta, and RSC Group.

More information about Intel’s participation at ISC 2021, including a full list of talks and demos, can be found at <https://hpcevents.intel.com>.

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 and intel.com.

For performance claims, see [43, 47, 108] at www.intel.com/3gen-xeon-config.
Results may vary.

© 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.

View source version on businesswire.com:
<https://www.businesswire.com/news/home/20210628005138/en/>

Tim Schulte
1-408-609-6845
tim.schulte@intel.com

Source: Intel Corporation