
IXPUG HPC Asia 2021
Advancing HPC Together
John K. Lee, Sr. Director of High Performance Computing Group
intel

Time to Insight

- Personalized Medicine
- Cleaner Energy
- Better Forecasts Sooner
- Expanded Research
- Smarter, Safer Designs

Accelerating the Growth of HPC

Race to Exascale

HPC-AI Convergence

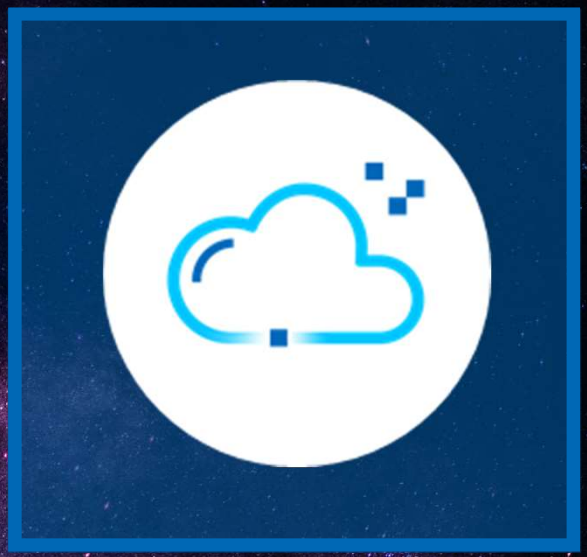

Migration to Cloud

> The pace of change has never been this fast … yet it will never be this slow again.
>
> **Justin Trudeau**

HPC and AI

Converging to Accelerate Insights

AI integrated in HPC workflows

AI accelerating HPC simulations

AI replacing HPC simulations

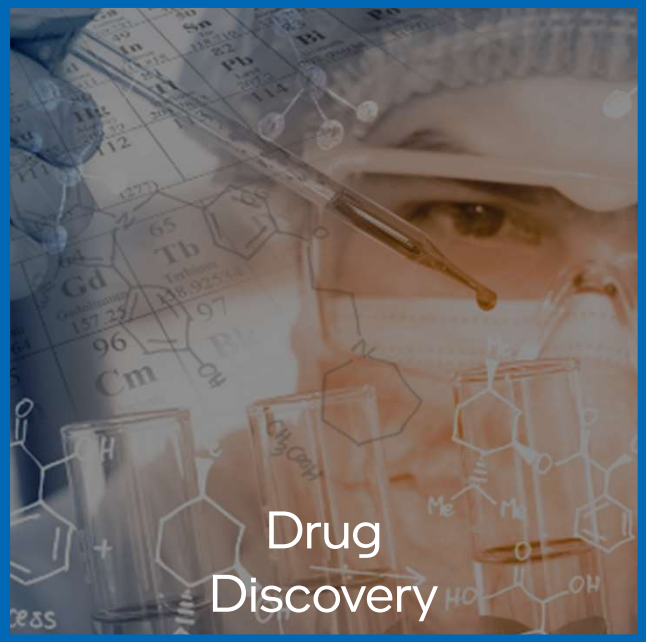

AI determining next series of MD simulations

AI predicting response of drug properties on tumors

AI reading & encoding reports for analysis

Accelerating Convergence

The Intel XPU Strategy

The Foundation for HPC and AI

General Purpose to Dedicated Acceleration

GPU for HPC & AI

FPGA

Data Center Deep Learning

Edge Deep Learning

Automated Driving

Building a Common Platform to Enable Seamless Heterogeneous Acceleration

Intel® Xeon® Scalable Processor

The ONLY Mainstream CPU Optimized for Convergence

Intel® Advanced Vector Extensions 512

Intel® Deep Learning Boost

Intel® Optane™ persistent memory

Cascade Lake

14nm

New AI acceleration built-in

(Intel® DL Boost with VNNI)

New memory storage hierarchy

Ice Lake

10nm

New microarchitecture

Increased memory bandwidth

Sapphire Rapids

10nm Enhanced SuperFin

Next gen Intel® DL Boost

(Intel® Advanced Matrix Extensions)

Accelerating Innovation

from Edge to Cloud to Supercomputing

3rd Gen Intel Xeon Scalable Processor (Ice Lake)

Optimized for an outstanding HPC and AI experience

- Higher memory bandwidth: 8 DDR4 channels & 3200MT/s
- Better performance per core via new architecture
- Faster I / O with PCIe Gen 4
- Supporting exascale storage with up to 6TB memory / socket and Intel® Optane™ PMem
- Security Innovations Intel® SGX & crypto acceleration

Leadership performance at half the cores for key Life Sciences and FSI workloads :

LAMMPS	NAMD STMV	MONTE CARLO
UP TO 1.2X	UP TO 1.2X	UP TO 1.3x

32 core Ice Lake vs 64 core AMD EPYC 7742

Results estimated based on testing on pre-production parts. See backup for configuration details. Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Welcoming New Ice Lake Customers

Announced at SC20:

Korea Meteorological Administration

Powering weather and climate studies to deliver more accurate, reliable and actionable forecasting

50 PFLOPS theoretical peak performance

8x faster than current system

MPCDF

Max Planck Computing & Data Facility

Rechenzentrum Garching der Max-Planck-Gesellschaft

Enabling groundbreaking research in physics, bioscience, theoretical chemistry and beyond

9 PFLOPS theoretical peak performance

Source: Max Planck Compute and Data Facility 10/29/20

ORACLE®

Offering next generation HPC cloud compute instances

30% expected performance increase for workloads like crash, CFD & EDA

For KMA details: www.lenovoxperience.com/newsDetail/283yl044hzgcdv7snkrmmx9oyeo323d8odvnw1h199h4hdoc. For Max Planck details: www.mpcdf.mpg.de/services/computing/raven/about-the-system For Oracle details: https://blogs.oracle.com/cloud-infrastructure/oracle-cloud-infrastructure-compute-and-high-performance-computing-roadmap-update. See backup for details. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

GPU Architecture Strategy

One Architecture and 4 Micro Architectures

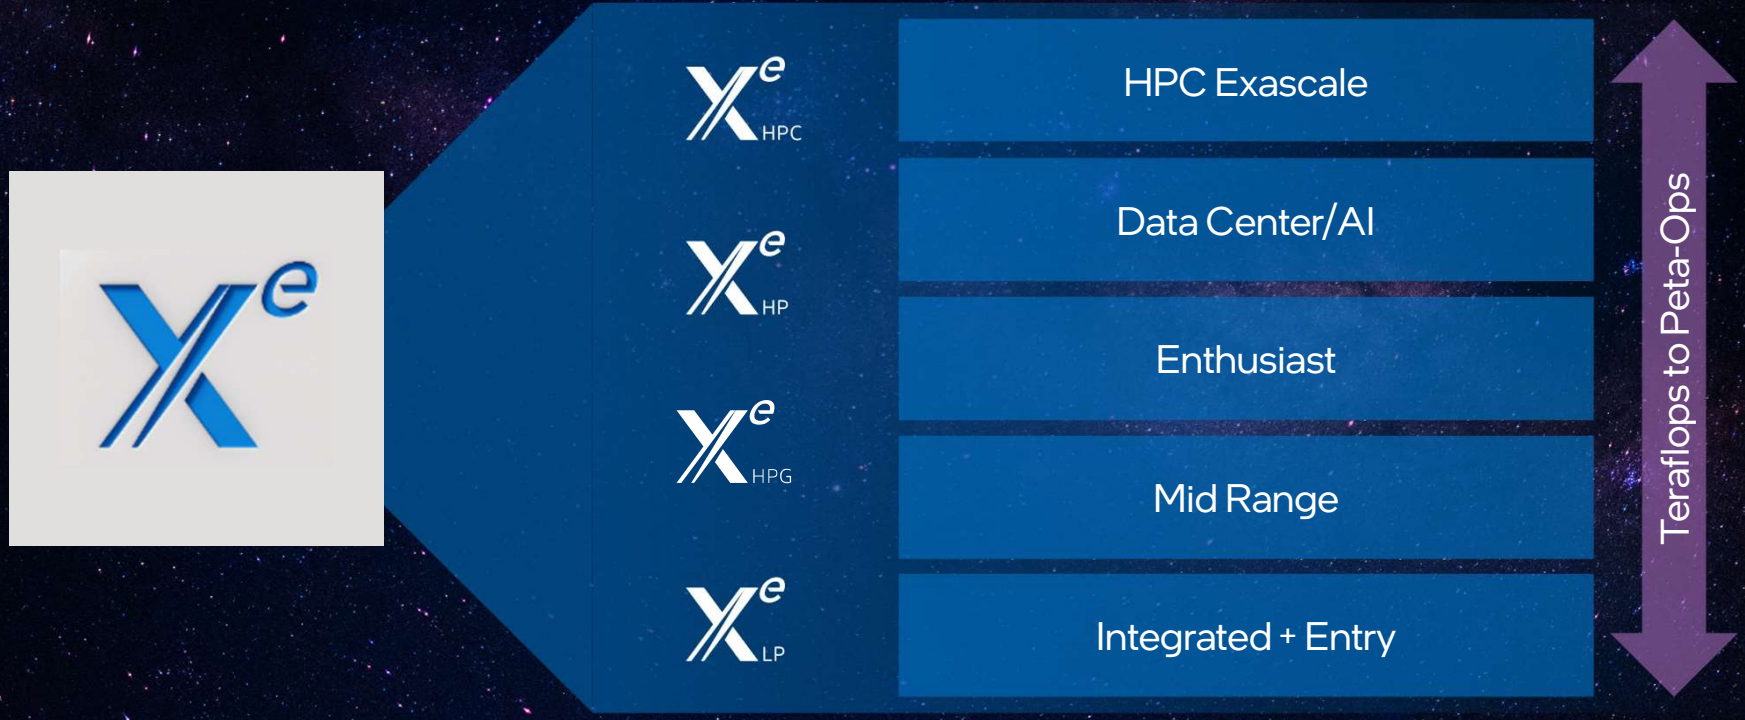

oneAPI

- Cross-architecture programming that delivers **freedom of choice**
- Expose and exploit **cutting-edge features of latest hardware**
- Builds upon **existing HPC programming models: OpenMP, Fortran, MPI**
- Based on **industry standards** and **open specifications**

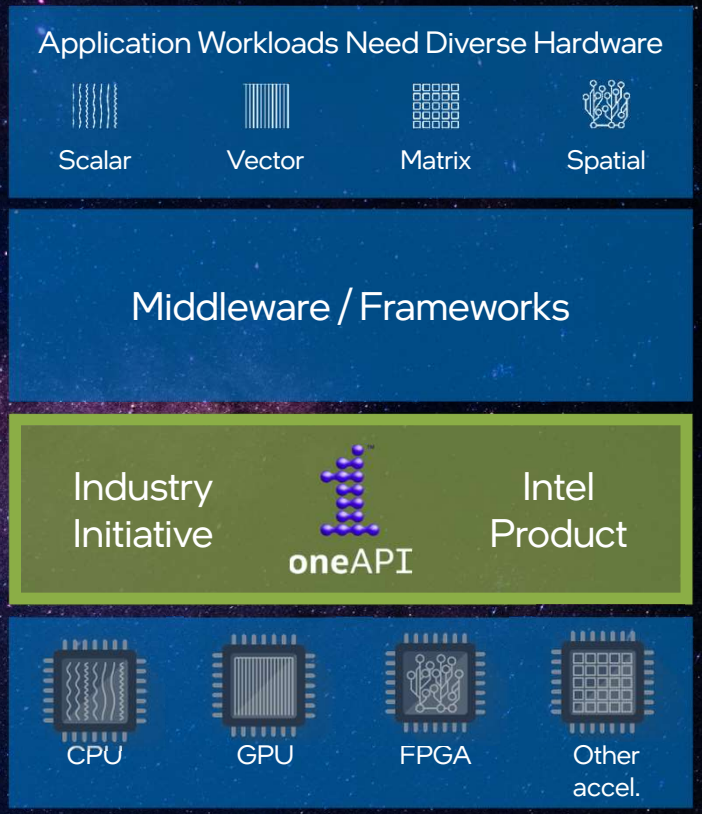

oneAPI Journey

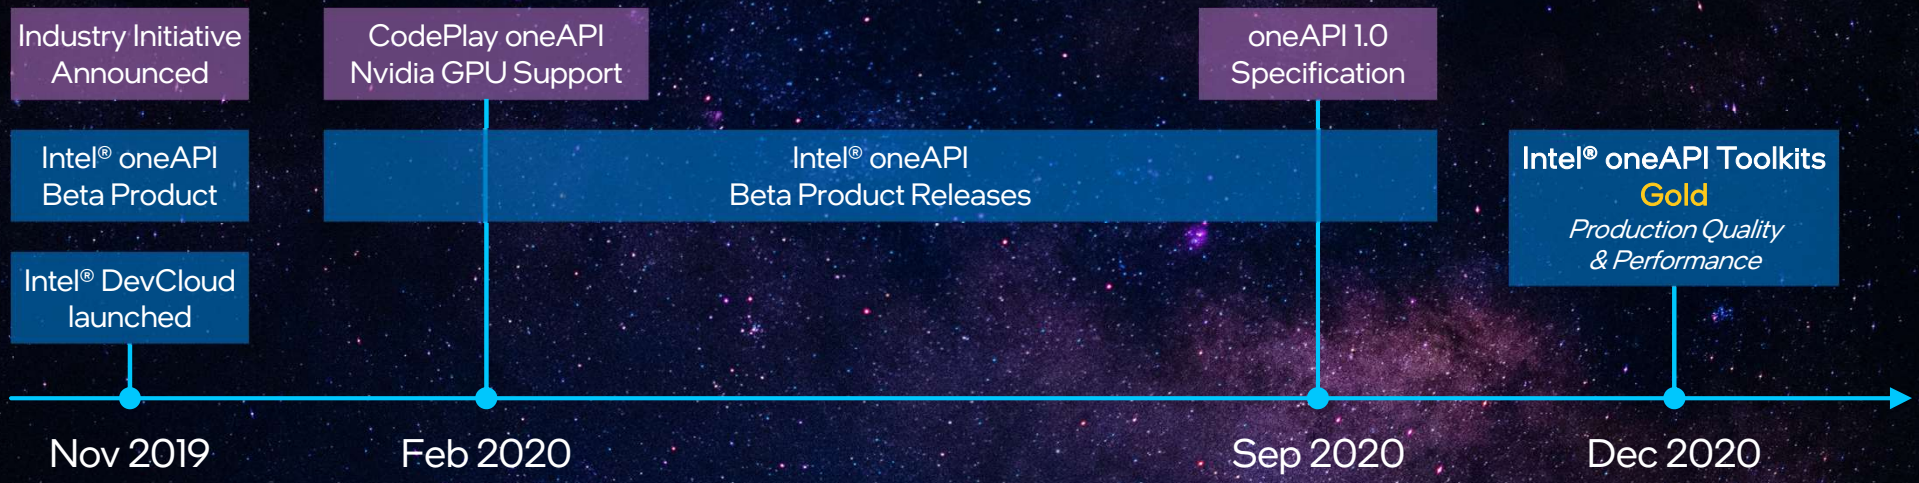

Builds on heritage of proven HPC tools, expanded to XPUs
Next generation of Intel software development products
Use it for FREE: Run the tools locally or on Intel® DevCloud

Intel DevCloud

No downloads

No hardware acquisition

No installation

No set-up & configuration

CPU

FPGA

dGPU
Intel® Iris® Xe MAX
Xe-HP*

*For Early Access Developers

oneAPI

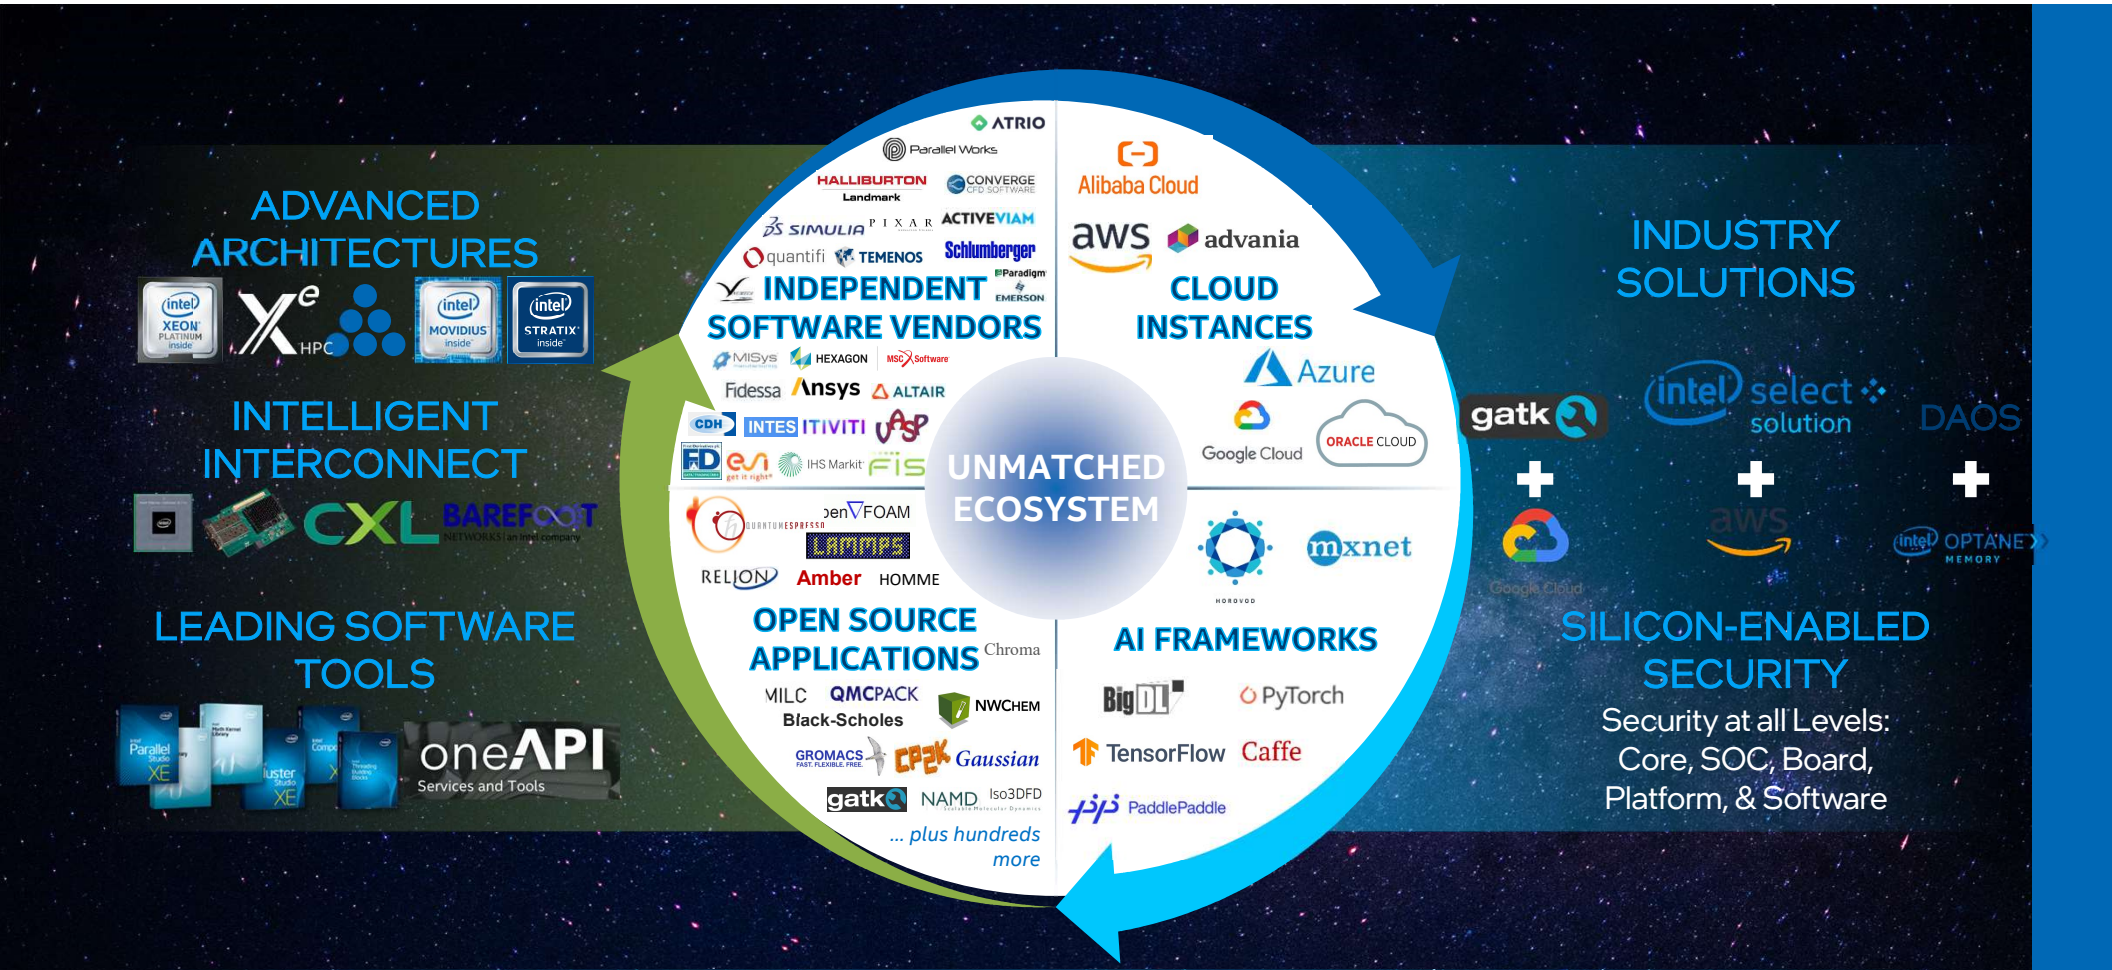

DECADES OF ECOSYSTEM ENGAGEMENT MEANS BETTER APPLICATION PERFORMANCE

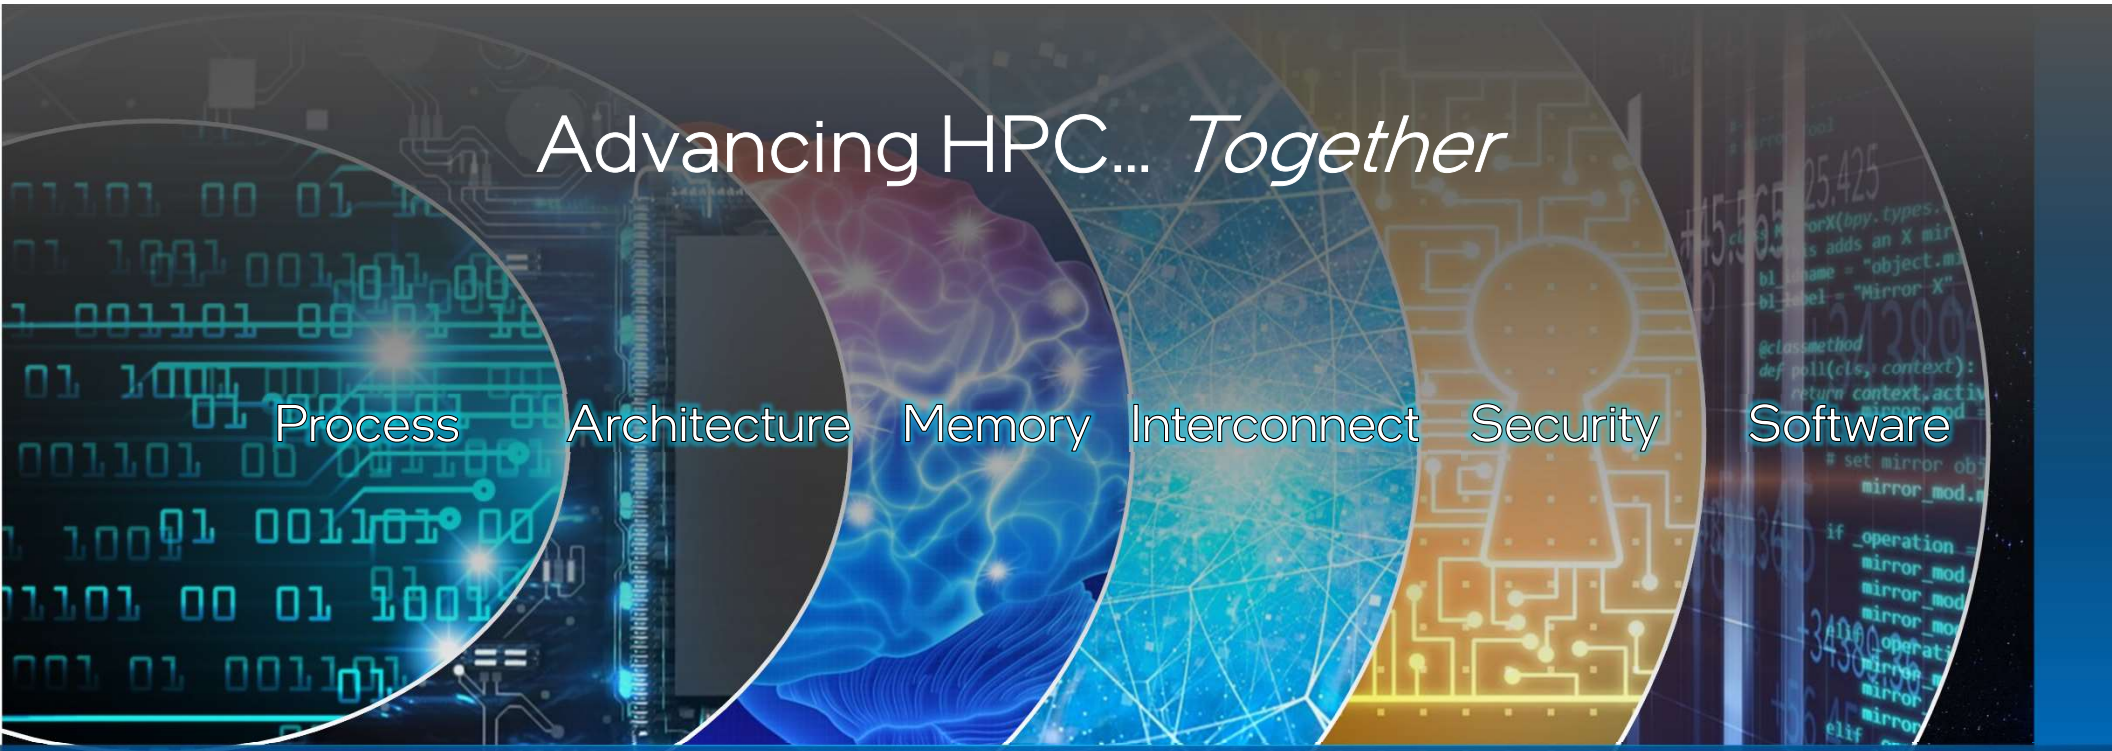

INTEL ARCHITECTURE VISION

Intel will deliver a diverse mix of **Scalar, Vector, Matrix and Spatial Architectures** designed with state of the art **Process technology**, fed by disruptive **Memory** hierarchies, integrated into systems with advanced **Packaging**, deployed at hyperscale with lightspeed **Interconnect** links, unified by a **single Software** abstraction, with benchmark defining Security features


intel
Thank You

Notices and Disclaimers

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

Third party data disclaimer: Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Configuration Details

Higher memory bandwidth: when comparing 8channels on 3rd Gen Intel Xeon Scalable processor with 6channels on 2nd Gen Intel Xeon Scalable Processor

Better performance per core via new architecture: Per core performance estimated based on tests with pre-production parts at iso frequency and memory BW per core vs. prior generation; as of Oct 2020.

NAMD STMV:

2S 3rd Gen Intel Xeon Scalable processor (Ice Lake): 1-node, 2x pre-production 3rd Gen Intel Xeon Scalable processor (Ice Lake - 2.2GHz, 32cores per socket), Intel reference platform, 256GB, 16x16GB 3200MHz DDR4, HT=on, TURBO=on, SNC=disabled, SSDSC2KG96 960GB, BIOS SE5C6200.86B.0017.D92.2007150417, microcode 0x8c000140, CentOS Linux 7.8, 3.10.0-1127.18.2.el7.crt1.x86_64, compiled with Intel C Compiler 2020u2, Intel MKL, NAMD: 2_15-Alpha1, tested by Intel on 9-17-2020. 2S AMD EPYC 7742: 1-node 2x AMD EPYC 7742 (2.25GHz, 64cores per socket), Supermicro platform, 16x16GB 3200MHz DDR4, SMT on, Boost on, NPS=4, SSDSC2KG96 960GB, BIOS2.0b dt 11/15/2019, microcode 0x8301025, CentOS Linux 7.7.1908, 3.10.0-1127.13.1.el7.crt1.x86_64, compiled with AOCC 2.2, Intel MKL, NAMD: 2_15-Alpha1, tested by Intel on 9-10-2020.

Monte Carlo:

2S 3rd Gen Intel Xeon Scalable processor (Ice Lake): 1-node, 2x pre-production 3rd Gen Intel Xeon Scalable processor (Ice Lake - 2.2GHz, 32cores per socket), Intel reference platform, 256GB, 16x16GB 3200MHz DDR4, HT=on, TURBO=on, SNC=disabled, SSDSC2KG96 960GB, BIOS SE5C6200.86B.0017.D92.2007150417, microcode 0x8c000140, CentOS Linux 7.8, 3.10.0-1127.18.2.el7.crt1.x86_64, compiled with Intel C Compiler 2020u2, Intel MKL 2020u2, Monte Carlo FSI Kernel workload developed by Intel, tested by Intel on 10-9-2020. 2S AMD EPYC 7742: 1-node 2x AMD EPYC 7742 (2.25GHz, 64cores per socket), Supermicro platform, 16x16GB 3200MHz DDR4, SMT on, Boost on, NPS=4, SSDSC2KG96 960GB, BIOS2.0b dt 11/15/2019, microcode 0x8301025, CentOS Linux 7.7.1908, 3.10.0-1127.13.1.el7.crt1.x86_64, compiled with Intel C Compiler 2020u2, Intel MKL 2020u2, Monte Carlo FSI Kernel workload developed by Intel, tested by Intel on 7-17-2020.

LAMMPS (Geomean of Atomic Fluid, Copper, Liquid Crystal, Polyethylene, Protein, Stillinger-Weber, Tersoff, and Water):

2S 3rd Gen Intel Xeon Scalable processor (Ice Lake): 1-node, 2x pre-production 3rd Gen Intel Xeon Scalable processor (Ice Lake - 2.2GHz, 32cores per socket), Intel reference platform, 256GB, 16x16GB 3200MHz DDR4, HT=on, TURBO=on, SNC=disabled, SSDSC2KG96 960GB, BIOS SE5C6200.86B.0017.D92.2007150417, microcode 0x8c000140, CentOS Linux 7.8, 3.10.0-1127.18.2.el7.crt1.x86_64, compiled with Intel C Compiler 2020u2, Intel MKL 2020u2, LAMMPS 03/03/2020, tested by Intel on 10-9-20202S AMD EPYC 7742: 1-node 2x AMD EPYC 7742 (2.25GHz, 64cores per socket), Supermicro platform, 16x16GB 3200MHz DDR4, SMT on, Boost on, NPS=4, SSDSC2KG96 960GB, BIOS2.0b dt 11/15/2019, microcode 0x8301025, CentOS Linux 7.7.1908, 3.10.0-1127.13.1.el7.crt1.x86_64, compiled with AOCC 2.2, LAMMPS 07/21/2020, tested by Intel on 8-19-2020

For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

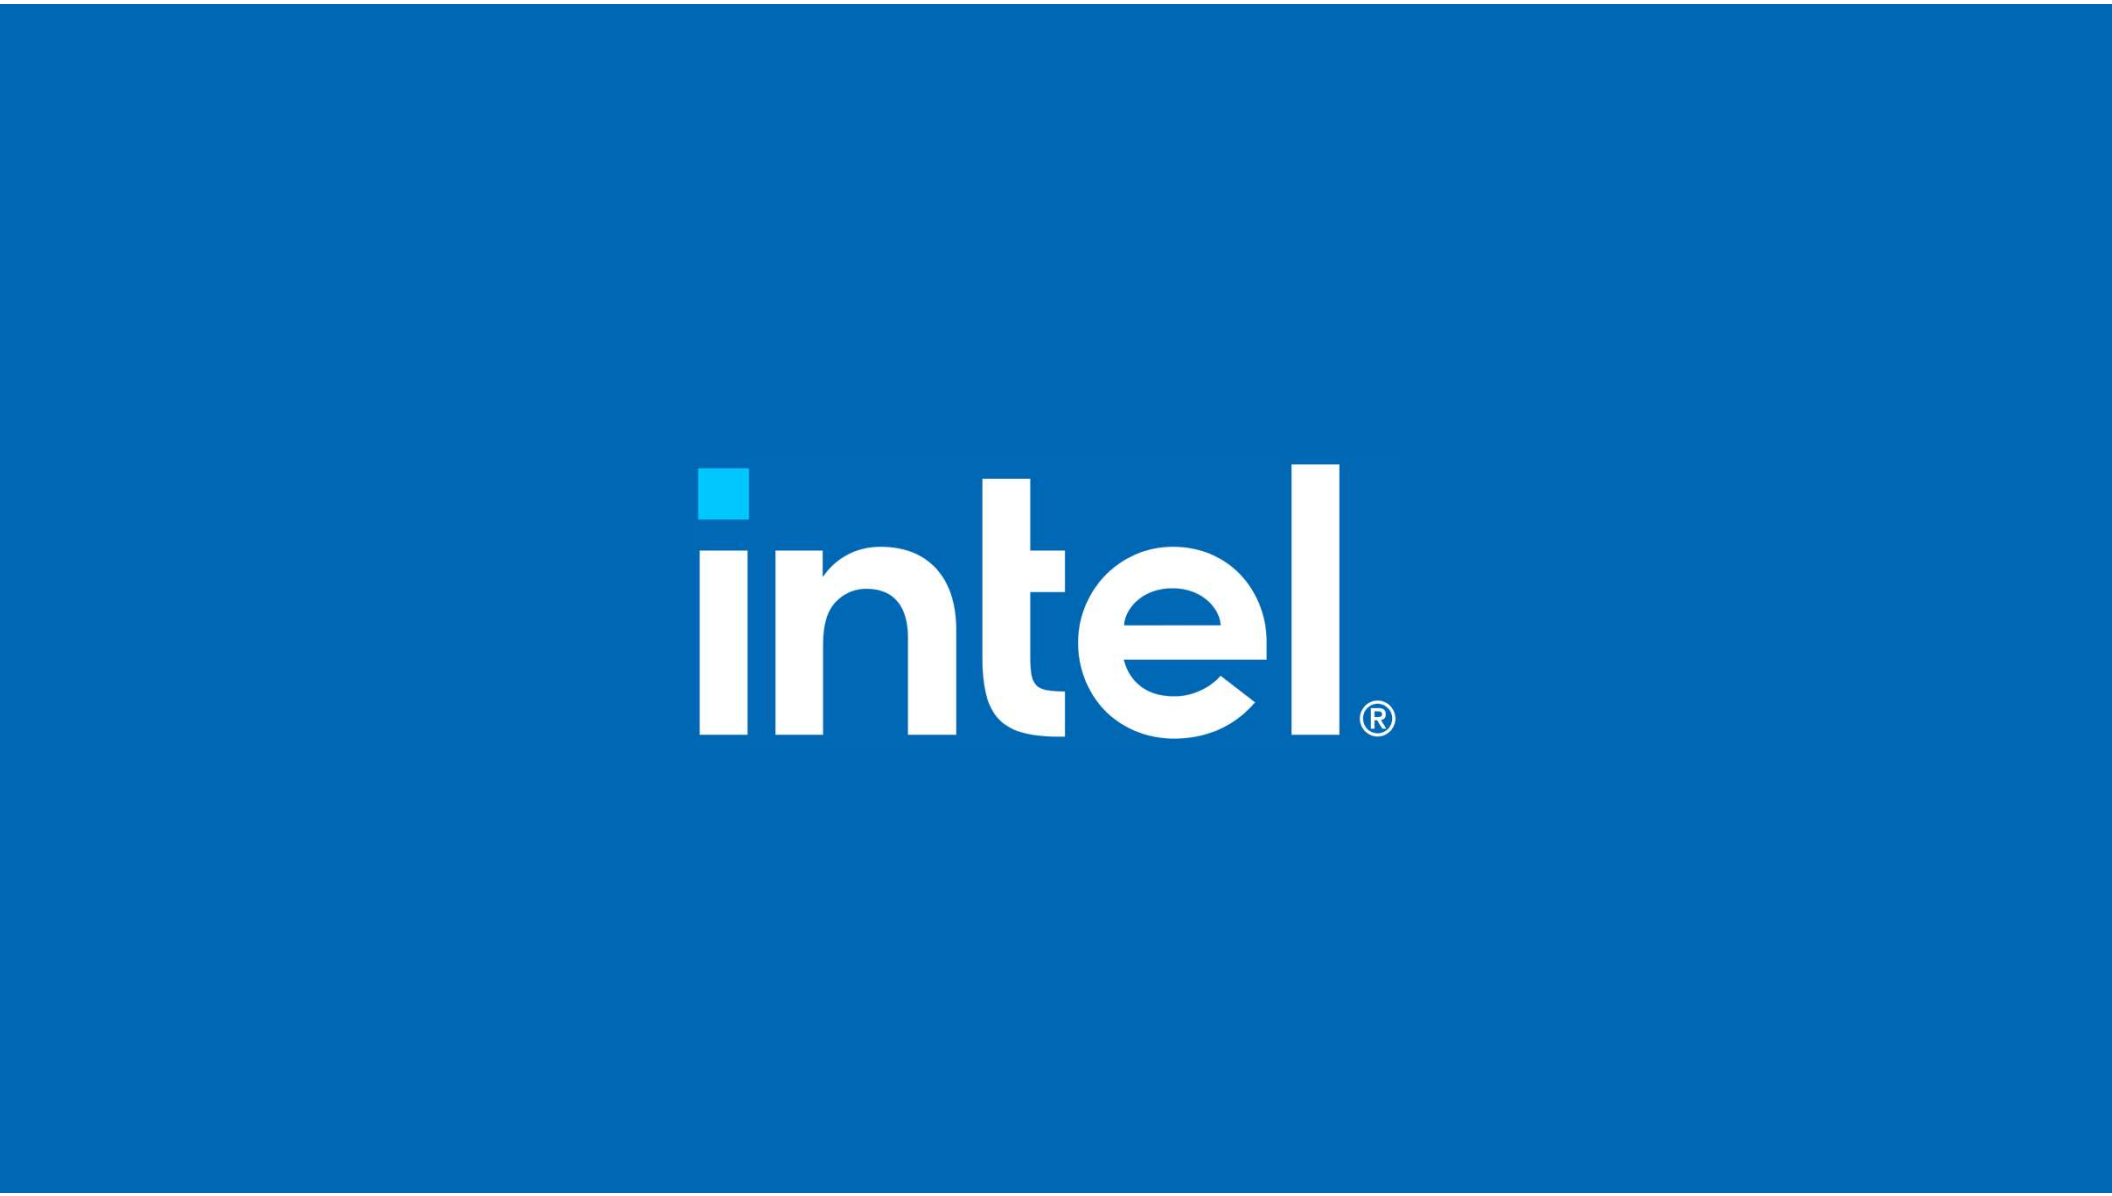
intel