intel XEON Accelerate with Xeon

4th Gen Intel® Xeon® Scalable processors Overview

Madhu Bramharouthu
Xeon SP Platform Tech Sales Lead | SMG

intel

Accelerating Data Center Growth

Delivering Leading Platforms for our Customers and Partners

Innovating for the Future of the Data Center

Continuing to Advance Products and Services

Intel® Xeon® Platform Momentum

85 million
Intel Xeon Scalable processors shipped to date

3rd Gen Intel Xeon Scalable processors
15 million Units shipped in less than two years

4th Gen Xeon Scalable processors
450 Design wins complete or in-progress

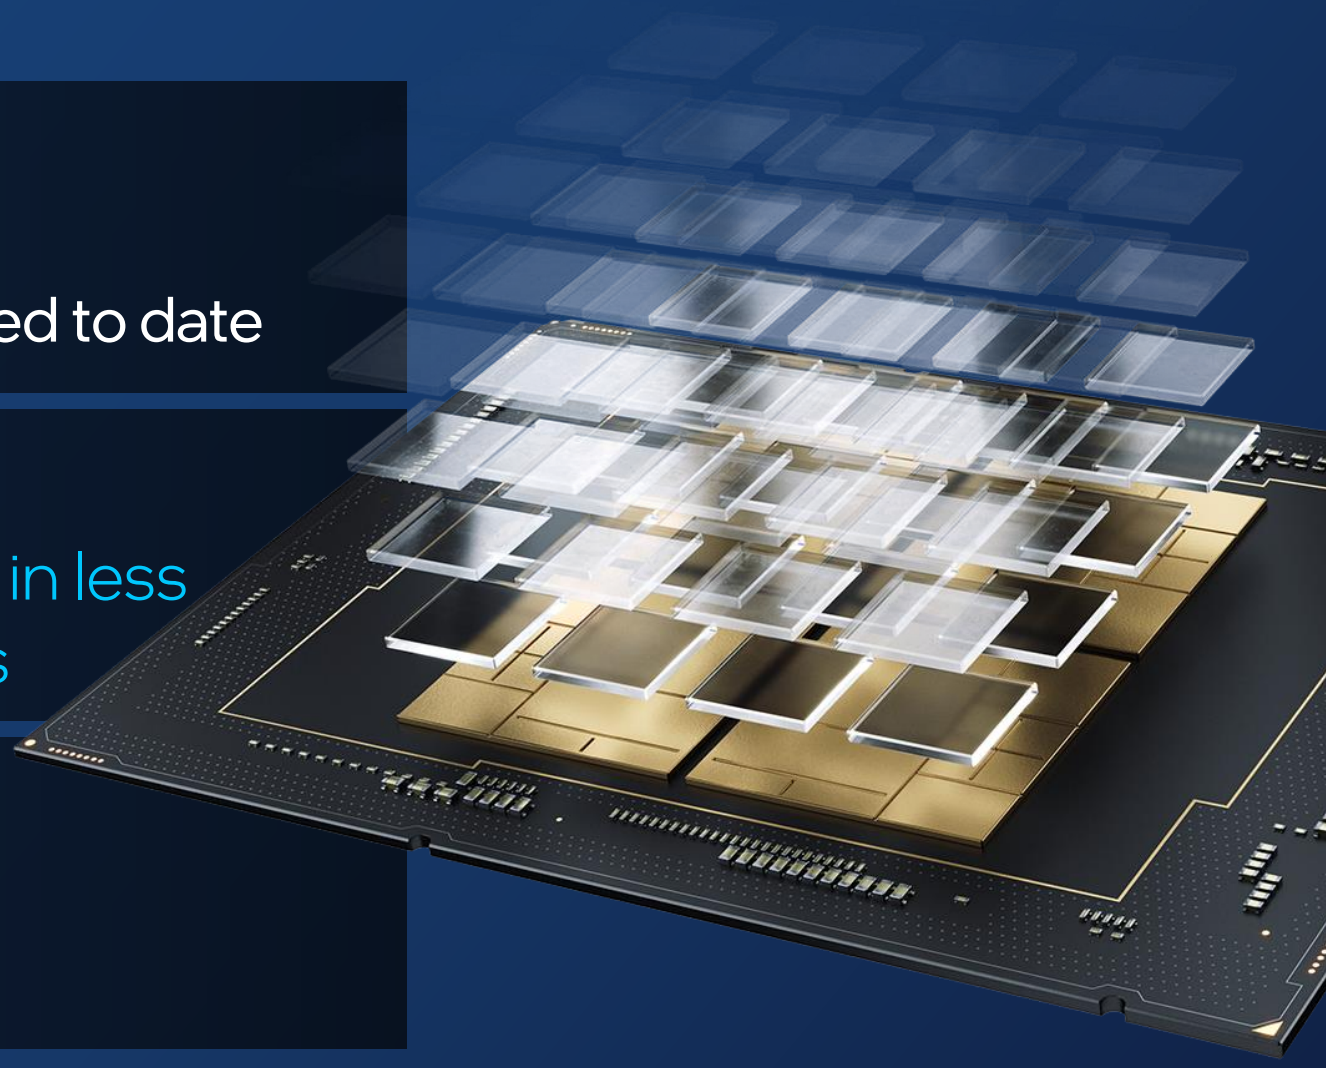

intel XEON Accelerate with Xeon

Compute demands in the data center,
in the network and at the edge
are bigger than ever before

Focus on Customer Real World Workloads

Artificial Intelligence

Networking 5G

Storage

HPC

Data Analytics

Intel's Differentiated Approach

Workload-First

CPU Cores + Built-In Accelerators Wins

Open Software Ecosystem + oneAPI & AI Tools

Higher Performance

Increased Efficiency

Optimal TCO

Intel's Data Center Evolution

Intel® Xeon® Scalable

Most Built-In Accelerators in the Market

Intel® Max CPU + GPU

Breakthrough Memory Bandwidth and Performance

Unrivaled Software Ecosystem

90% of developers are using software developed or optimized by Intel[1]

[1] Global Development Survey conducted by Evans Data Corp., 2021.

INTRODUCING

Intel's Most Feature-Rich Server Platform

4th Gen Intel® Xeon® Scalable Processors and Intel® Xeon® CPU Max Series Processors

New Integrated IP Accelerators

Increased I/O Bandwidth with PCIe 5.0

Compute Express Link (CXL) 1.1

4th Gen Intel Xeon Scalable Processor

Intel Xeon CPU Max Series

Intel® In-Memory Analytics Accelerator (Intel® IAA)

Intel® Data Streaming Accelerator (Intel® DSA)

Intel® Dynamic Load Balancer (Intel® DLB)

Intel® QuickAssist Technology (Intel® QAT)

PCIe 5.0/4.0 80 Lanes

1 to 8 Socket Support

Intel® In-field Scan **

Intel® Seamless Firmware Update

Intel® Virtual RAID on CPU (Intel® VROC)

High Bandwidth Memory Support

8 DMI 4.0 LINKS

Intel® C741 Chipset

Intel® Advanced Matrix Extensions (AMX)

Accelerating Interfacing Architecture (AiA)

Intel® Advanced Vector Extensions for vRAN

Enhanced Intel® Software Guard Extensions

8 Channel DDR5
Up To 4800 MT/s (1DPC)
Up To 4400 MT/s (2DPC)
16 DIMMS per Socket

New RAS features (Enhanced ECC, Error Check & Scrub)

Intel® Ultra Path Interconnect (Up to 4 Links @ 16 GT/s)

Optimized Power Mode

New Integrated ISA Accelerators

HW-Enhanced Security

Increased Memory Bandwidth with DDR5

Increased Inter-Socket Bandwidth with UPI 2.0

Intel® Optane™ SSDs

Intel® Agilex™ FPGAs

Intel® Data Center GPU Max Series
Intel® Data Center GPU Flex Series

Intel® Infrastructure Processing Unit (Intel® IPU)

intel XEON Accelerate with Xeon

** Intel In-Filed Scan is a new capability available through select providers in 2023.

Intel® Xeon® CPU Max Series

Designed for HPC, AI, Analytics and other memory bound workloads

1st x86 CPU to integrate high bandwidth memory and accelerators onto the processor package

Leading performance and efficiency for our customers

4th Gen Intel® Xeon® Scalable Processors

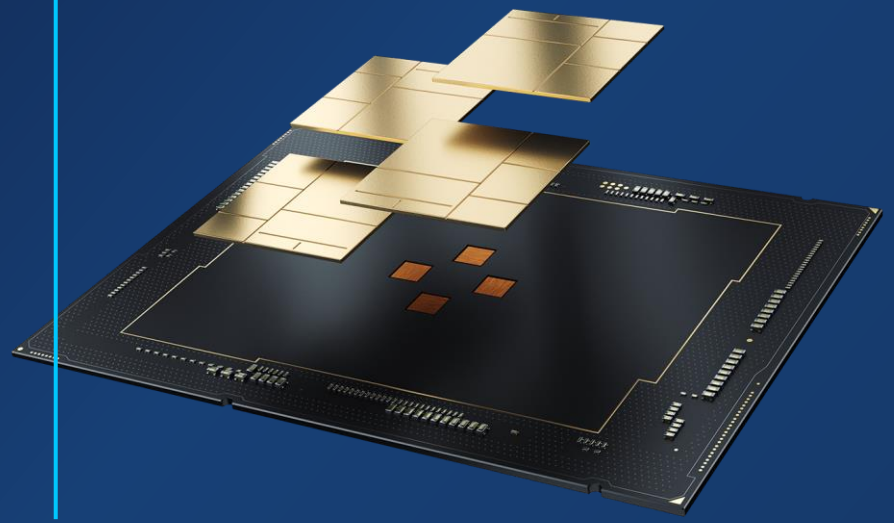

Workload-first approach to innovation, design, and delivery

Most built-in accelerators of any CPU on the market

Leading performance and efficiency for our customers

Industry's most comprehensive Confidential Computing portfolio

4th Gen Intel® Xeon® Scalable Processors

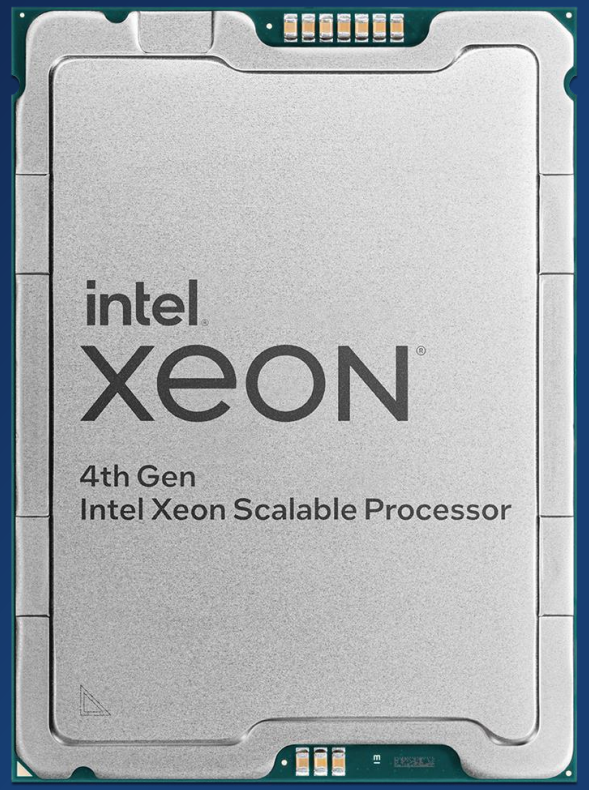

- 1 to 8 socket scalability
- Up to 60 cores
  per processor
- Most built-in accelerators of any CPU
- Increased memory bandwidth with DDR5
- Increased I/O bandwidth with PCIe 5
  80 lanes
- Increased inter-socket bandwidth with UPI 2.0
- Compute Express Link (CXL) 1.1
- Hardware-enhanced security

Maximize the Effectiveness of Every Core

New Integrated IP Acceleration Engines

Intel® acceleration engines help free up cores for more general-purpose compute tasks, increasing overall workload performance and power efficiency

Integrated IP

- Intel® QuickAssist Technology (Intel® QAT)
- Intel® Dynamic Load Balancer (Intel® DLB)
- Intel® Data Streaming Accelerator (Intel® DSA)
- Intel® In-Memory Analytics Accelerator (Intel® IAA)

New Instruction Set Architecture (ISA)

- Intel® Advanced Matrix (AMX)
- Intel® Advanced Vector Extensions for vRAN

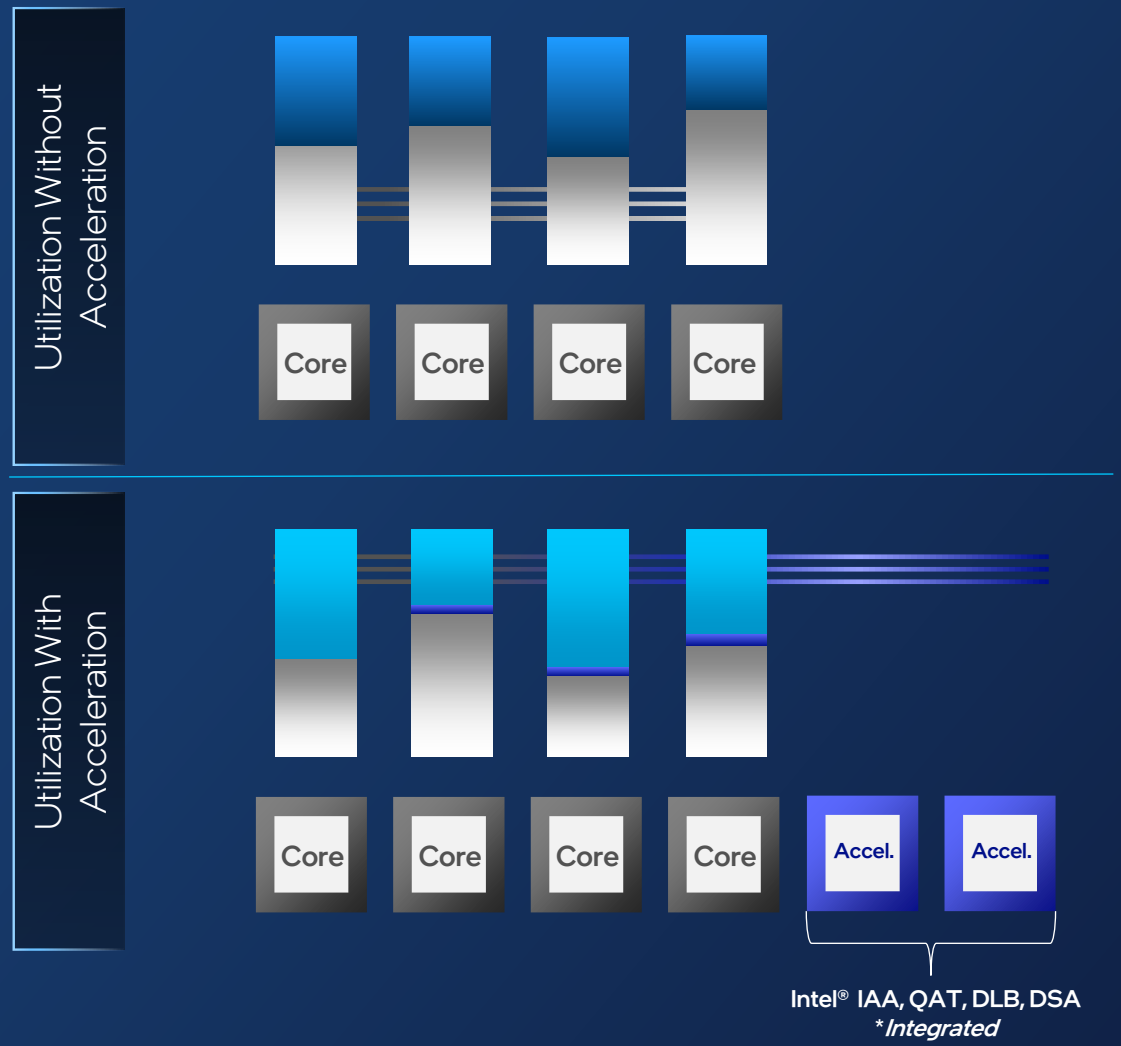

Intel® Accelerator Engines

Most Built-in Accelerators of any CPU on the market providing customers with increased **performance, costs savings** and **sustainability** advantages for the biggest and fastest-growing workloads

Intel® AI Engines

- Intel® Advanced Matrix Extensions (Intel® AMX)
- Intel® Advanced Vector Extensions 512 (Intel® AVX-512)
- Intel® Deep Learning Boost (Intel® DL Boost)

Intel® Security Engines

- Intel® Control-Flow Enforcement Technology (Intel® CET)
- Intel® Crypto Acceleration
- Intel® Software Guard Extensions (Intel® SGX)
- Intel® Trust Domain Extensions (Intel® TDX)
- Intel® QuickAssist Technology (Intel® QAT)

Intel® HPC Engines

- Intel® Advanced Vector Extensions 512 (Intel® AVX-512)
- Intel® Advanced Matrix Extensions (Intel® AMX)
- Intel® Data Streaming Accelerator (Intel® DSA)
- Intel® QuickAssist Technology (Intel® QAT)

Intel® Network Engines

- Intel® QuickAssist Technology (Intel® QAT)
- Intel® Dynamic Load Balancer (Intel® DLB)
- Intel® Data Streaming Accelerator (Intel® DSA)
- Intel® Advanced Vector Extensions (Intel® AVX) for vRAN
- Intel® Speed Select Technology (Intel® SST)

Intel® Analytics Engines

- Intel® In-memory Analytics Accelerator (Intel® IAA)
- Intel® Data Streaming Accelerator (Intel® DSA)
- Intel® Advanced Vector Extensions 512 (Intel® AVX-512)
- Intel® QuickAssist Technology (Intel® QAT)

Intel® Storage Engines

- Intel® Data Streaming Accelerator (Intel® DSA)
- Intel® QuickAssist Technology (Intel® QAT)
- Intel® In-memory Analytics Accelerator (Intel® IAA)
- Intel® Data Direct I/O (Intel® DDIO)
- Intel® Advanced Vector Extensions 512 (Intel® AVX-512)
- Intel® Crypto Acceleration

intel XEON Accelerate with Xeon

Supported in 4th Gen Intel® Xeon® Scalable processors and Intel® Xeon® CPU Max Series processors.
Many features are enabled through Intel® oneAPI Tools and Intel-optimized AI frameworks.

Developer Tools for 4th Gen Intel® Xeon® Scalable Processors

Intel ® oneAPI, AI tools and optimized AI frameworks help developers maximize application performance by activating advanced capabilities of 4th Gen Intel® Xeon® Scalable processors and Intel® Max Series processors. In multiarchitecture systems with Intel Xeon processors and Intel GPUs, using a single codebase through oneAPI delivers productivity and performance.

Compilers, libraries & analysis tools support built-in accelerators to unleash performance, and fast training and inference for AI workloads.

- **Intel® oneAPI Math Kernel Library**
  for HPC and technical compute
- **Intel® oneAPI Deep Neural Network Library**
  for deep learning training + inference
- **Intel® Query Processing & Intel® Data Mover Library***
  for query processing, compression and data movement
- **Intel® VTune™ Profiler**
  helps locate time-consuming parts of code and identify significant issues affecting application performance

Learn more: Software for 4th Gen Intel Xeon & Max Series Processors

*Intel® QPL is open source. Open source Intel® DML in beta, v1 coming soon.

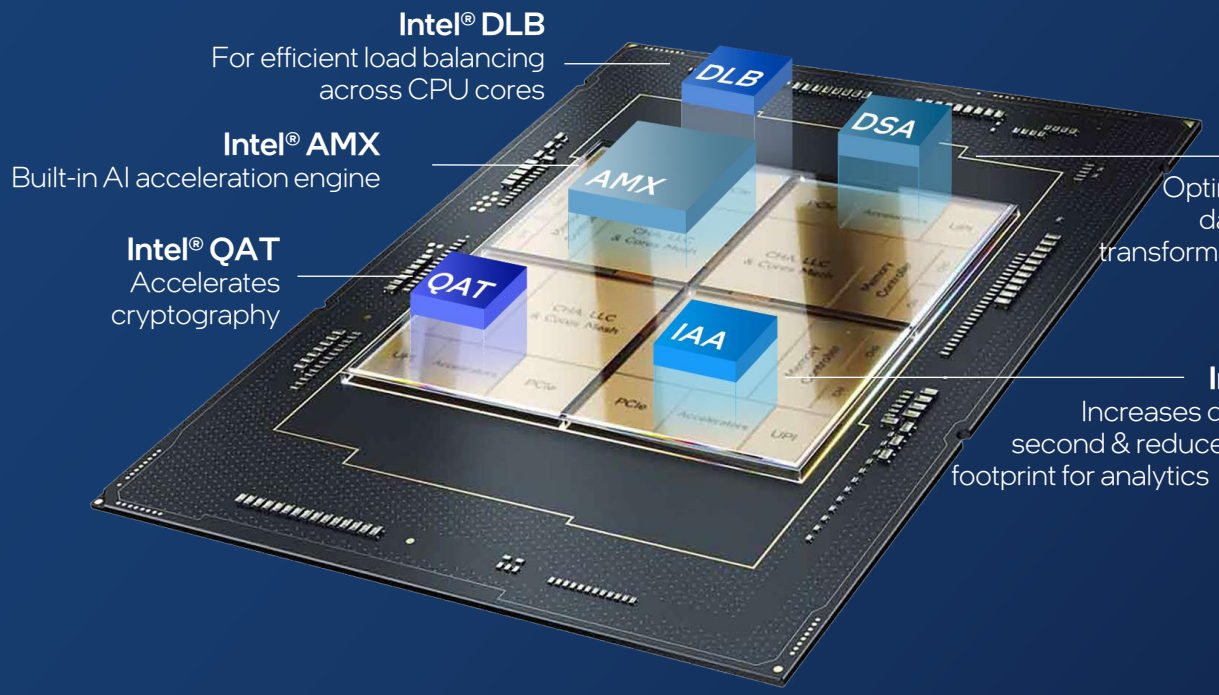

intel oneAPI BASE TOOLKIT | intel oneAPI HPC TOOLKIT | intel oneAPI RENDERING TOOLKIT | intel oneAPI IoT TOOLKIT | intel AI ANALYTICS TOOLKIT | OpenVINO

Powered by oneAPI

intel XEON Accelerate with Xeon

Benefits of Intel® Accelerator Engines

A Higher Performance Server Architecture

Intel® Advanced Matrix Extensions (Intel® AMX)

Up to

8.6x

higher speech recognition inference performance with built-in AMX BF16 vs. FP32

Intel® Dynamic Load Balancer (Intel® DLB)

Up to

96%

lower latency at the same throughput for Istio-Envoy Ingress with Intel® DLB vs. software for Istio Ingress gateway

Intel® Data Streaming Accelerator (Intel® DSA)

Up to

1.7x

higher IOPs for SPDK-NVMe with built-in Intel® DSA vs. ISA-L software

Intel® In-Memory Analytics Accelerator (Intel® IAA)

Up to

2.1x

higher RocksDB performance with Intel® IAA vs Ztsd software

Intel® QuickAssist Technology (Intel® QAT)

Up to

84%

fewer cores to achieve same connections/s on NGINX with built-in QAT vs. out-of-box software

Accelerators Enable Step Function Performance Beyond Base Architecture

intel XEON Accelerate with Xeon

See [A26, W6, N18, D1. N15] at intel.com/processorclaims: 4th Gen Intel Xeon Scalable processors. Results may vary.

CPU + Accelerators: Differentiated Performance On Real Workloads

4th Gen Intel® Xeon® Scalable processors					Intel® Xeon® CPU Max Series
General Purpose Compute	Artificial Intelligence	Network 5G vRAN	Networking & Storage	Data Analytics	HPC
53%	Up to 10x	Up to 2x	Up to 2x	Up to 3x	Up to 3.7x
average performance gain*	higher inference and training performance*	capacity for vRAN workloads at same power envelope*	higher data compression with 95% fewer cores*	higher performance*	on memory-bound workloads**

intel XEON Accelerate with Xeon

See [G1, A17, N10, N16, D1] at intel.com/processorclaims: 4th Gen Intel Xeon Scalable processors. Results may vary.

*4th Gen Intel Scalable Processor vs. 3rd Gen Intel Xeon Scalable processors

** Intel Xeon CPU Max Series vs. Intel Xeon 8380

Architected to Accelerate Real World Workloads

Cloud

Up to 89% performance increase with Intel® QAT vs. prior gen.[11]

inspur

"We were pleased to observe a 20% increase in performance over the current generation C2 VMs from Google Cloud in testing with one of our key workloads."[12]

Google Cloud

Security

Intel® SGX performs up to 4.6x higher vs. prior gen.[13]

Fortanix

AI

Up to 2.48x performance improvement with Intel® AMX vs. prior gen.[14]

Alibaba

Up to 4x performance gain with Intel® AMX vs. prior gen.[15]

BRAINPOOL.AI FUJITSU

"Intel's [4th Gen Xeon processor] provides unprecedented levels of performance for critical graph intelligence tasks."

KATANA GRAPH

5G

"It is not just a software, it is not interfaces, it is not only radio. It is how we can build all the pieces in our architecture."

Telefónica

HPC

Up to 4.3x performance improvement with Intel AMX® on Intel Xeon Max Series vs. prior gen.[16]

CERN openlab SURF

Up to 8.57x performance improvement on Intel Xeon Max Series vs. Intel E5V4.[17]

Los Alamos NATIONAL LABORATORY Hewlett Packard Enterprise

"The reason we use the 4th Gen Intel® Xeon® processor as the building block for immersion born systems is really because of its unrivaled power and efficiency."

Hypertec

The Universal AI Platform

Intel® AI Platform

Any AI code, every workload

4th Gen Intel® Xeon® Scalable Processor

The flexibility of Xeon with the built-in DL performance of an AI accelerator

Up to 10x gen-to-gen model performance
Up to 7.7x performance/watt

Build and deploy everywhere

Intel oneAPI & AI tools + optimized frameworks

Accelerate E2E workflows, increase productivity, speed AI time to results

Up to 6.7x E2E application performance
400+ models validated

Implement pre-built solutions

Extensive Intel AI products and partnership

Accelerate end customer time to market

Example: Meet 10ms SLA with an order-of-magnitude throughout gains

See [A2, A16, A17, A33] at intel.com/processorclaims: 4th Gen Intel Xeon Scalable processors. Results may vary.

Intel's Most Sustainable Data Center Processor Ever

Perf/watt improvements
from the most built-in accelerators ever offered in an Intel® processor

New Optimized Power Mode
delivers up to 20 percent power savings with negligible performance impact on select workloads

Built-in advanced telemetry
enables monitoring and control of electricity consumption and carbon emissions

Available immersion cooling warranty rider for Intel® Xeon® processors

Scope 3 GHG emissions benefits
due to manufacturing with 90-100 percent renewable electricity

Manufactured at sites with state-of-the-art water reclaim
facilities that in 2021 recycled 2.8 billion gallons of water

intel XEON

intel XEON Accelerate with Xeon

CPU + Accelerators: Groundbreaking Efficiency

Higher Performance per Watt

2.9x

average improvement of perf/watt with built-in accelerators*

Lower Power Bills

up to **70W**

power savings per CPU with Optimized Power Mode

Lower TCO More Sustainable

55%

lower TCO and power consumption while reducing 524K kg of CO_2 emissions*

AI Real Time Inferencing workload, ResNet50

intel XEON Accelerate with Xeon

See [E1, E6, E7] at intel.com/processorclaims: 4th Gen Intel Xeon Scalable processors. Results may vary
*For selected workloads.
*4th Gen Intel® Xeon® Scalable Processor vs. 3rd Gen Intel Xeon Scalable processors.

Acceleration Delivers TCO Value

AI real time inferencing	Database	High Performance Computing
55%	52%	66%
vs. prior gen	vs. prior gen	vs. prior gen
lower TCO	lower TCO	lower TCO

intel XEON

intel XEON Accelerate with Xeon

See [E7, E8, E9] at intel.com/processorclaims: 4th Gen Intel Xeon Scalable processors. Results may vary.

Sustainable Compute: Optimized Power Mode

up to 20%
lower processor power savings

<5%
performance impact
for select workloads

Lower Power Bills

up to 70W
power savings per CPU with Optimized Power Mode

intel XEON

intel XEON Accelerate with Xeon

See [E6] at intel.com/processorclaims: 4th Gen Intel Xeon Scalable processors. Results may vary.

Flexibility & Choice for Customers

Most Workload Optimized SKUs on the Market

>56%

Intel® Xeon® Processor Volume supports customer specific or workload specific demand*

Expanded Options for Workload Optimized SKUs

Cloud (-P, -V, -M)	Network (-N)	Storage (-S)
1-Socket (-U)	Long-Life Use (IOT) (-T)	IMDB Analytics (-H)
HPC (w/HBM)	Liquid Cooled (-Q)	CSP Custom

intel XEON Accelerate with Xeon

* Source: Intel Xeon CPU billings on 3rd Gen Intel Xeon Processors, 2022 YTD

Refresh and consolidate Intel® Xeon® processor-based servers

4th Gen Intel® Xeon® processors can significantly lower your total cost of ownership

Up to **5:1** Server consolidation

Free up rack space and reduce data center footprint

Up to **80%** Reduction in server count

Simplify and lower data center operations cost

Up to **60%** Reduction in CO_2 emissions & power

Lower power bills and reach sustainability goals

Up to **75%** Reduction in TCO

Lower overall cost to reinvest for growth

Recover your cost in **4** months

intel XEON Accelerate with Xeon

Comparing benefits transitioning from Intel Xeon 4110 to Intel Xeon 5420+
Performance varies by use, configuration and other factors.
See backup for configurations. Results may vary.

Refresh and consolidate Intel® Xeon® processor-based servers

4th Gen Intel® Xeon® processors can significantly lower your total cost of ownership, cost recovery in 4 months

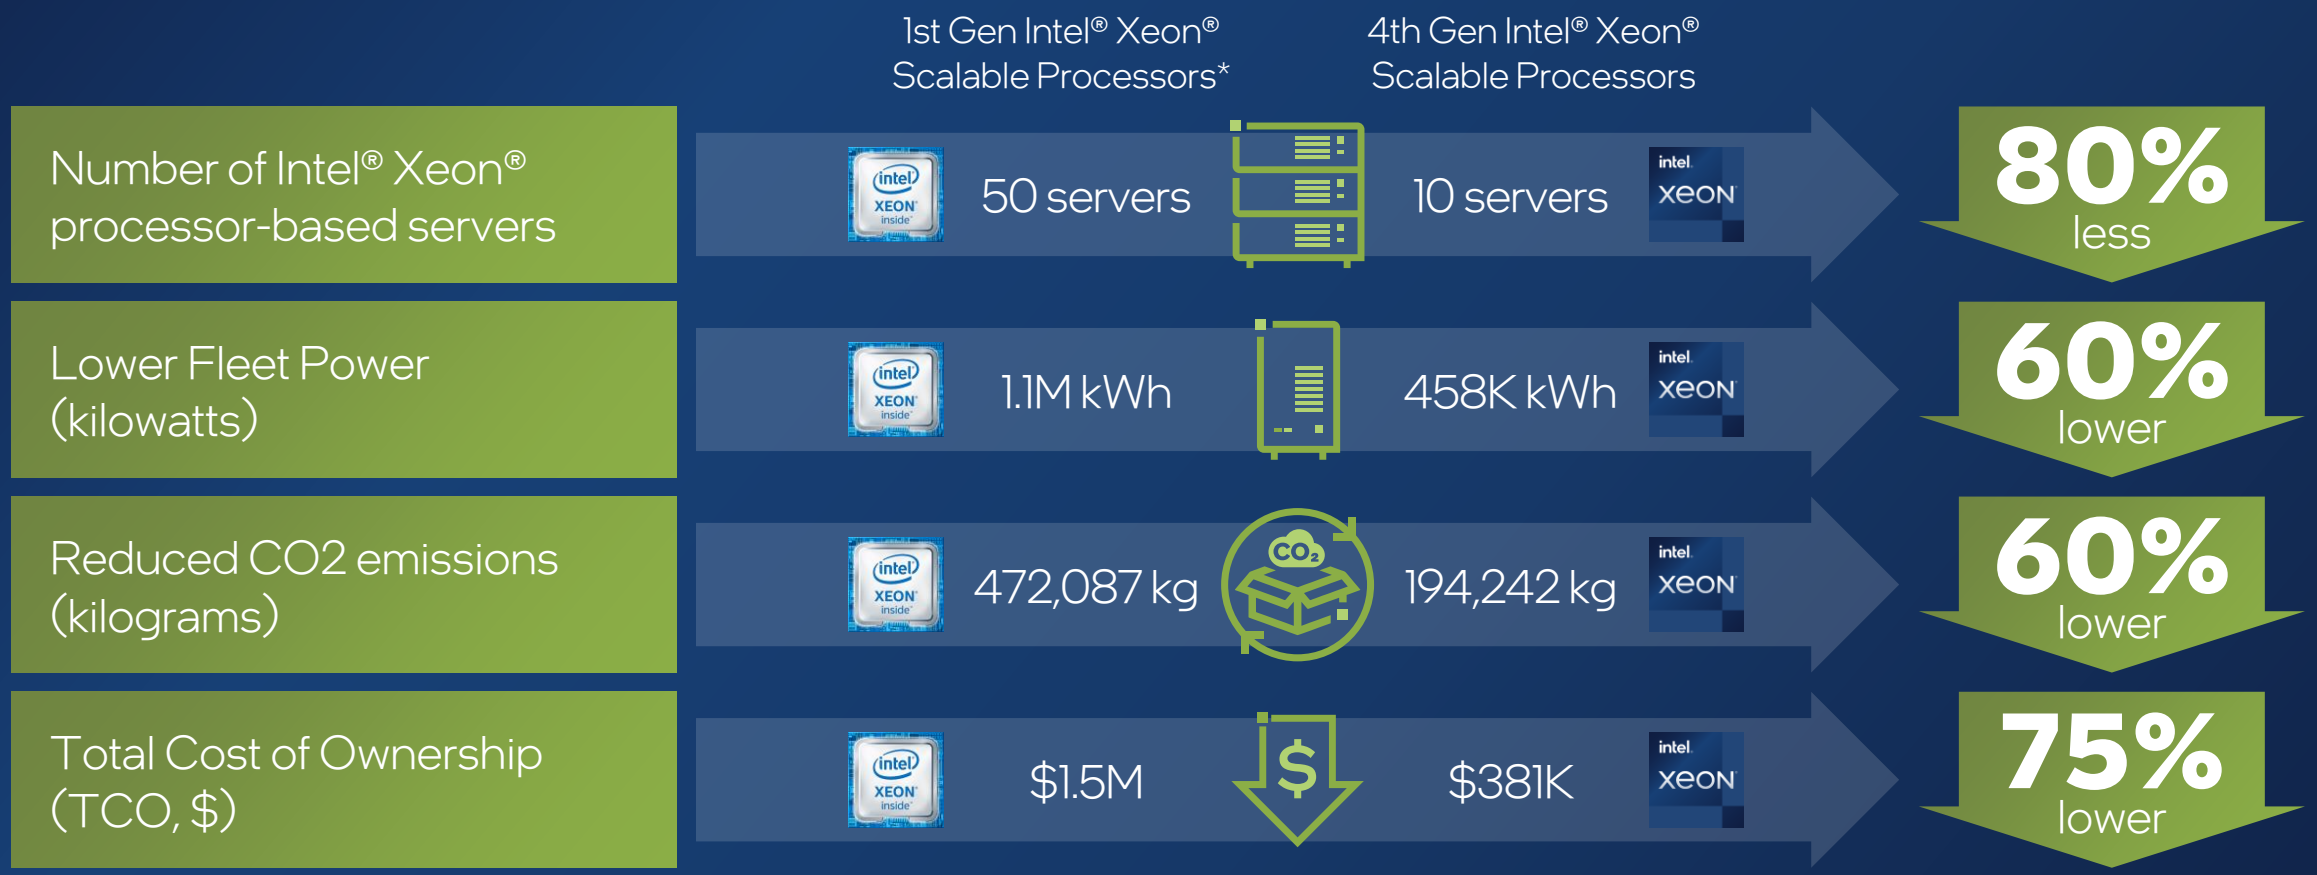

intel XEON Accelerate with Xeon

*Comparing benefits transitioning from Intel Xeon 4110 to Intel Xeon 5420+
Performance varies by use, configuration and other factors.
See backup for configurations. Results may vary.

Refresh and consolidate Intel® Xeon® processor-based servers

4th Gen Intel® Xeon® processors can significantly lower your total cost of ownership

From 1st to 4th Gen Intel Xeon processor	Reduce servers	Reduce energy and CO2	Reduce TCO	Recover costs (months)
8160 → 8460Y+	64%	34%	43%	20
6130 → 6430	62%	26%	49%	11
5120 → 5420+	70%	48%	61%	7
4110 → 5420+	80%	59%	74%	4

An Architecture Influenced by Customers

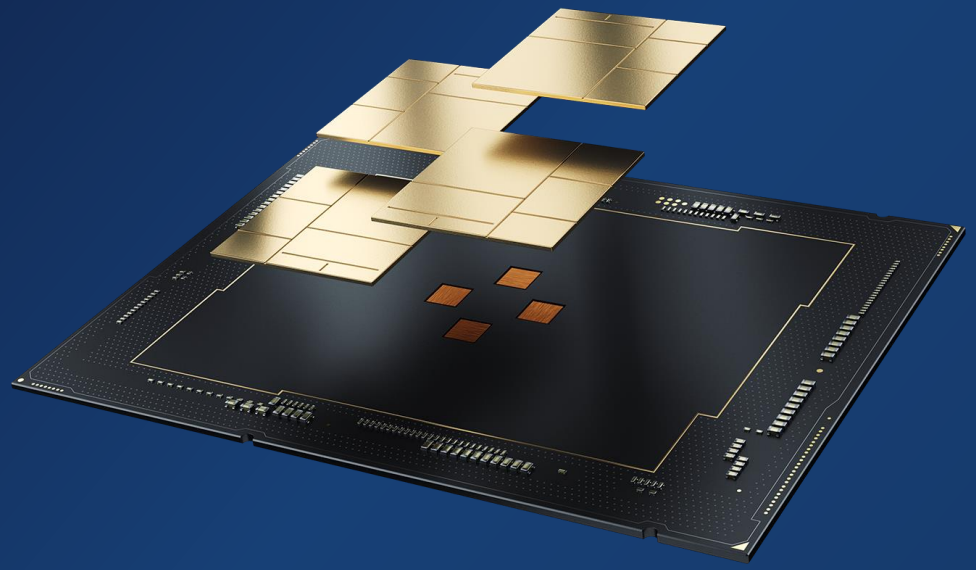

Workload-first approach to innovation, design, and delivery

- Unique Die, SKU, features for Unique Market Needs
- DDR5, PCIe5 and CXL here today
- Cores + Accelerators deliver better value

Most built-in accelerators of any CPU on the market

- Increased performance, power and cost efficiency
- Accelerating AI, Analytics, Networking, Storage, HPC
- 2.9x higher avg perf per watt gains

Leading performance and efficiency for our customers

- Only x86 CPU to offer 4S and 8S scalability & HBM
- Up to 10X higher AI inference and training
- Lower TCO and power consumption

Industry's most comprehensive Confidential Computing portfolio

- Application isolation with Intel® SGX
- Virtual machine isolation with Intel® TDX
- Trust verification services with Project Amber

See backup for workloads and configurations. Results may vary.

intel XEON Accelerate with Xeon

Bringing the Architecture to Life

4th Gen Intel® Xeon® Scalable Processors

Life Sciences

Get up to 53% faster results for life and material sciences for more effective research.

Digital Consumer Web Services

Run social network microservices up to 88% faster for better user experiences.

Financial Services

Meet tight timelines with up to 45% faster results for options pricing.

Retail

Offer personalized product recommendations up to 6.3x faster for smoother e-commerce.

See backup for workloads and configurations. Results may vary.

Thank you!

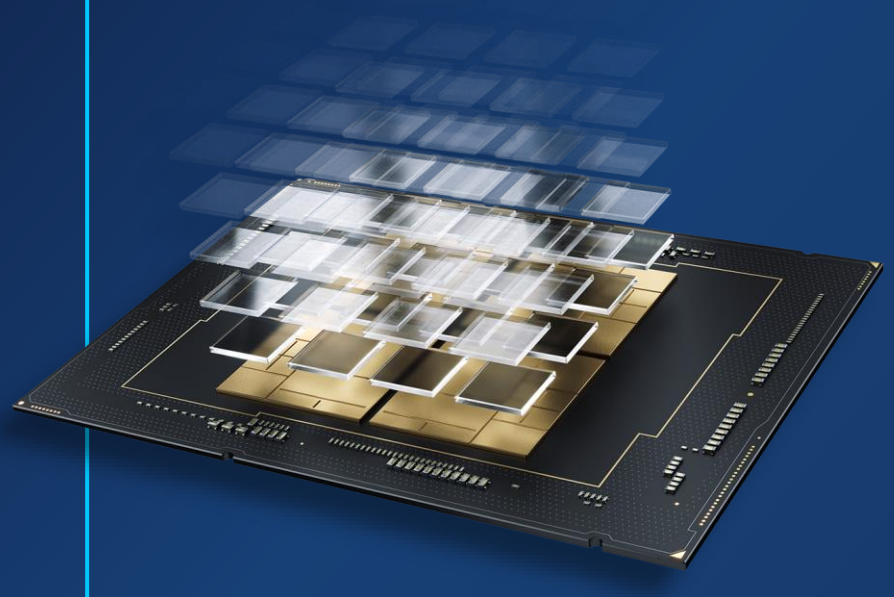

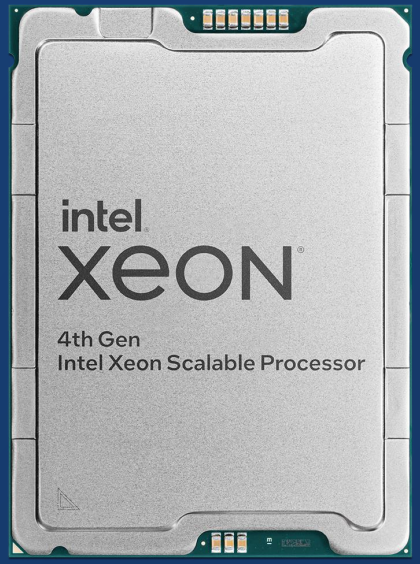

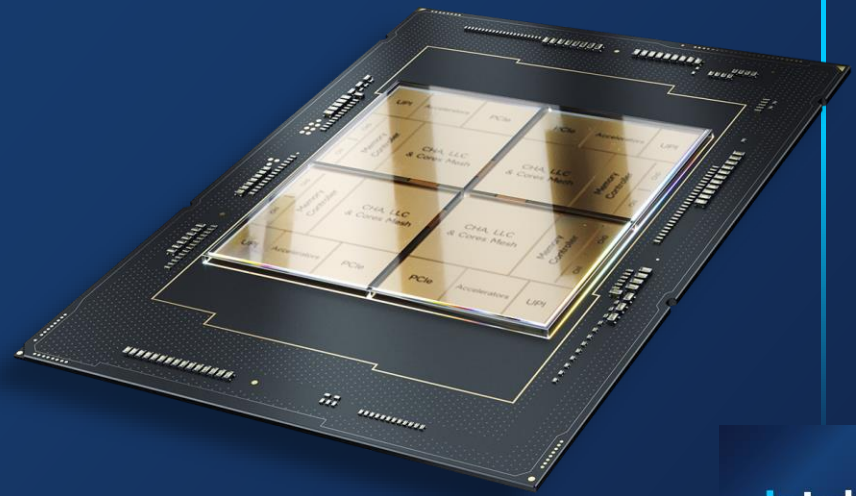

intel XEON

Learn more

- 4th Gen Intel® Xeon® Scalable processor Launch Essentials at: https://partneruniversity.intel.com/learn/learning_plan/view/242/4th-gen-intelr-xeonr-scalable-processor-launch-essentials)


intel

intel XEON Accelerate with Xeon

Resources & Configurations

intel

Resources and Configurations

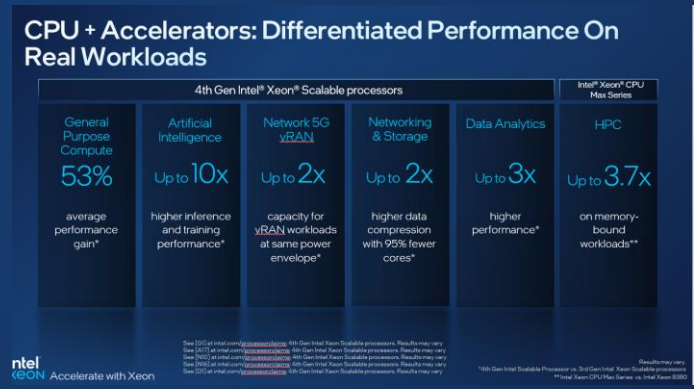

Architecting to Accelerate Customer Workloads

Leading Performance with the most built – in accelerators

- Up to 3.7x on memory-bound workloads - Intel® Xeon® 8380: Test by Intel as of 10/7/2022. 1-node, 2x Intel® Xeon® 8380 CPU, HT On, Turbo On, Total Memory 256 GB (16x16GB 3200MT/s DDR4), BIOS Version SE5C620.86B.01.01.0006.2207150335, ucode revision=0xd000375, Rocky Linux 8.6, Linux version 4.18.0-372.26.1.el8_6.crt1.x86_64, Stream v5.10; Intel® Xeon® CPU Max Series: Test by Intel as of 9/2/2022. 1-node, 2x Intel® Xeon® CPU Max Series, HT On, Turbo On, SNC4, Total Memory 128 GB (8x16GB HBM2 3200MT/s), BIOS Version SE5C7411.86B.8424.D03.2208100444, ucode revision=0x2c000020, CentOS Stream 8, Linux version 5.19.0-rc6.0712.intel_next.1.x86_64+server, Stream v5.10

Resources and Configurations

Bringing the Architecture to Life (1 of 3)

Get up to 53% faster results for life and material sciences for more effective research and Meet tight timelines with up to 45% faster results for options pricing

- DeePMD (Multi-Instance Training)
  8480+: Test by Intel as of 10/12/2022. 1-node, 2x Intel Xeon Platinum 8480+, Total Memory 512 GB, kernel 4.18.0-365.el8_3x86_64, compiler gcc (GCC) 8.5.0 20210514 (Red Hat 8.5.0-10), https://github.com/deepmodeling/deepmd-kit, Tensorflow 2.9, Horovod 0.24.0, oneCCL-2021.5.2, Python 3.9
  8380: Test by Intel as of 10/20/2022. 1-node, 2x Intel Xeon Platinum 8380 processor, Total Memory 256 GB, kernel 4.18.0-372.26.1.el8_6.crt1.x86_64, compiler gcc (GCC) 8.5.0 20210514 (Red Hat 8.5.0-10), https://github.com/deepmodeling/deepmd-kit, Tensorflow 2.9, Horovod 0.24.0, oneCCL-2021.5.2, Python 3.9

- LAMMPS
  8480+: Test by Intel as of 9/29/2022. 1-node, 2x Intel Xeon Platinum 8480+, HT On, Turbo On, SNC4, Total Memory 512 GB (16x32GB 4800MT/s, DDR5), BIOS Version SE5C7411.86B.8713.D03.2209091345, ucode revision=0x2b000070, Rocky Linux 8.6, Linux version 4.18.0-372.26.1.el8_6.crt1.x86_64, LAMMPS v2021-09-29 cmkl:2022.1.0, icc:2021.6.0, impi:2021.6.0, tbb:2021.6.0; threads/core:; Turbo:off; BuildKnobs:-O3 -ip -xCORE-AVX512 -g -debug inline-debug-info -qopt-zmm-usage=high;
  8380: Test by Intel as of 10/11/2022. 1-node, 2x Intel Xeon Platinum 8380 CPU, HT On, Turbo On, NUMA configuration SNC2, Total Memory 256 GB (16x16GB 3200MT/s, Dual-Rank), BIOS Version SE5C620.86B.01.01.0006.2207150335, ucode revision=0xd000375, Rocky Linux 8.6, Linux version 4.18.0-372.26.1.el8_6.crt1.x86_64, LAMMPS v2021-09-29 cmkl:2022.1.0, icc:2021.6.0, impi:2021.6.0, tbb:2021.6.0; threads/core:; Turbo:on; BuildKnobs:-O3 -ip -xCORE-AVX512 -g -debug inline-debug-info -qopt-zmm-usage=high;
  LAMMPS (Atomic Fluid, Copper, DPD, Liquid_crystal, Polyethylene, Protein, Stillinger-Weber, Tersoff, Water)

- Quantum Espresso (AUSURF112, Water_EXX)
  8480+: Test by Intel as of 9/2/2022. 1-node, 2x Intel Xeon Platinum 8480+, HT On, Turbo On, Total Memory 512 GB (16x32GB 4800MT/s, Dual-Rank), ucode revision= 0x90000c0, Rocky Linux 8.6, Linux version 4.18.0-372.26.1.el8_6.crt1.x86_64, Quantum Espresso 7.0, AUSURF112, Water_EXX
  8380: Test by Intel as of 9/30/2022. 1-node, 2x Intel Xeon Platinum 8380 CPU, HT On, Turbo On, Total Memory 256 GB (16x16GB 3200MT/s, Dual-Rank), ucode revision=0xd000375, Rocky Linux 8.6, Linux version 4.18.0-372.26.1.el8_6.crt1.x86_64, Quantum Espresso 7.0, AUSURF112, Water_EXX

- VASP(Geomean: CuC, Si, PdO4, PdO4_k221)
  8480+: Test by Intel as of 10/7/2022. 1-node, 2x 4th Gen Intel® Xeon® Platinum 8480+, HT On, Turbo On, SNC4, Total Memory 512 GB (16x32GB 4800MT/s, DDR5), BIOS Version SE5C7411.86B.8713.D03.2209091345, ucode revision=0x2b000070, Rocky Linux 8.6, Linux version 4.18.0-372.26.1.el8_6.crt1.x86_64, VASP6.3.2
  8380: Test by Intel as of 10/7/2022. 1-node, 2x Intel® Xeon® 8380 CPU, HT On, Turbo On, NUMA configuration SNC2, Total Memory 256 GB (16x16GB 3200MT/s, Dual-Rank), BIOS Version SE5C620.86B.01.01.0006.2207150335, ucode revision=0xd000375, Rocky Linux 8.6, Linux version 4.18.0-372.26.1.el8_6.crt1.x86_64, VASP6.3.2

- GROMACS (geomean: benchMEM, benchPEP, benchPEP-h, benchRIB, hecbiosim-3m, hecbiosim-465k, hecbiosim-61k, ion_channel_pme_large, lignocellulose_rf_large, rnase_cubic, stmv, water1.5M_pme_large, water1.5M_rf_large)
  8480+: Test by Intel as of 10/7/2022. 1-node, 2x 4th Gen Intel® Xeon® Scalable Processor, HT On, Turbo On, SNC4, Total Memory 512 GB (16x32GB 4800MT/s, DDR5), BIOS Version SE5C7411.86B.8713.D03.2209091345, ucode revision=0x2b000070, Rocky Linux 8.6, Linux version 4.18.0-372.26.1.el8_6.crt1.x86_64, GROMACS v2021.4_SP
  8380: Test by Intel as of 10/7/2022. 1-node, 2x Intel® Xeon® 8380 CPU, HT On, Turbo On, NUMA configuration SNC2, Total Memory 256 GB (16x16GB 3200MT/s, Dual-Rank), BIOS Version SE5C620.86B.01.01.0006.2207150335, ucode revision=0xd000375, Rocky Linux 8.6, Linux version 4.18.0-372.26.1.el8_6.crt1.x86_64, Converge GROMACS v2021.4_SP

Resources and Configurations

Bringing the Architecture to Life (2 of 3)

Meet tight timelines with up to 45% faster results for options pricing

- Binomial Options, Black Scholes, Monte Carlo
  8480+: Test by Intel as of 10/7/2022. 1-node, 2x Intel Xeon Platinum 8480+, HT On, Turbo On, SNC4, Total Memory 512 GB (16x32GB 4800MT/s, DDR5), BIOS Version SE5C7411.86B.8713.D03.2209091345, ucode revision=0x2b000070, Rocky Linux 8.6, Linux version 4.18.0-372.26.1.el8_6.crt1.x86_64, Binomial Options v1.1, Black Scholes v1.4, Monte Carlo v1.2
  8380: Test by Intel as of 10/7/2022. 1-node, 2x Intel Xeon Platinum 8380 CPU, HT On, Turbo On, Total Memory 256 GB (16x16GB 3200MT/s DDR4), BIOS Version SE5C620.86B.01.01.0006.2207150335, ucode revision=0xd000375, Rocky Linux 8.6, Linux version 4.18.0-372.26.1.el8_6.crt1.x86_64, Binomial Options v1.1, Black Scholes v1.4, Monte Carlo v1.2

Resources and Configurations

Bringing the Architecture to Life (3 of 3)

Run social network microservices up to 88% faster for better user experiences.

- 8480+:4 (1master, 3worker)-node, each-node, pre-production platform with 2x Intel(R) Xeon(R) Platinum 8480+ on QuantaGrid D54Q-2U with GB (16 slots/ 64GB/ DDR5 4800) total memory, ucode 0x2b000081 , HT on, Turbo on, CentOS Linux release 8.4.2105, 6.0.6, 1x 2.9T INTEL SSDPE2KE032T7, 1x 893.8G AVAGO JBOD, 2x Ethernet Controller X710 for 10GBASE-T, 2x Ethernet Controller E810-C for QSFP, DeathStarBench Social Network, wrk2 - load generator, ICE driver (CVL): 6.0.6, Cilium CNI - 1.11.4, Kubernetes - 1.21.14, ContainerD - 1.4.12, deathstarbench/social-network-microservices:0.0.8, nginx-thrift: yg397/openresty-thrift:xenial, memcached:1.6.7, mongo:4.4.6, redis 7.0.5, dataset: DeathStarBench/socialNetwork/datasets/social-graph/socfb-Reed98/ , test by Intel on 11/2/2022. \
- 8360Y:4 (1master, 3worker)-node, each-node, 2x Intel(R) Xeon(R) Platinum 8360Y on Intel Whitley with GB (16 slots/ 32GB/ DDR4 3200) total memory, ucode 0xd000375, HT on, Turbo on, CentOS Linux release 8.4.2105, 6.0.6, 1x 894.3G INTEL SSDSC2KG96, 2x Ethernet Controller X710 for 10GBASE-T, 1x Ethernet Controller E810-C for QSFP, DeathStarBench Social Network, wrk2 - load generator, ICE driver (CVL): 6.0.6, Cilium CNI - 1.11.4, Kubernetes - 1.21.14, ContainerD - 1.4.12, deathstarbench/social-network-microservices:0.0.8, nginx-thrift: yg397/openresty-thrift:xenial, memcached:1.6.7, mongo:4.4.6, redis 7.0.5, dataset: DeathStarBench/socialNetwork/datasets/social-graph/socfb-Reed98/ , test by Intel on 11/2/2022.
- https://github.com/delimitrou/DeathStarBench#publications

Offer personalized product recommendations up to 6.3x faster for smoother e-commerce.

- 8480+: 1-node, pre-production platform with 2x Intel Xeon Platinum 8480+ on Archer City with 1024 GB (16 slots/ 64GB/ DDR5-4800) total memory, ucode 0x2b0000a1, HT on, Turbo on, CentOS Stream 8, 5.15.0, 1x INTEL SSDSC2KW256G8 (PT)/Samsung SSD 860 EVO 1TB (TF), DLRM, Inf: bs=n [1socket/instance], Inference: bs: fp32=128, amx bf16=128, amx int8=128, Training bs:fp32/amx bf16=32k [1 instance, 1socket], Criteo Terabyte Dataset, Framework: https://github.com/intel-innersource/frameworks.ai.pytorch.private-cpu/tree/d7607bdd983093396a70713344828a989b766a66; Modelzoo: https://github.com/IntelAI/models/tree/spr-launch-public, PT:1.13, IPEX: 1.13, OneDNN: v2.7, test by Intel on 10/24/2022.
- 8380: 1-node, 2x Intel Xeon Platinum 8380 on M50CYP2SBSTD with 1024 GB (16 slots/ 64GB/ DDR4-3200) total memory, ucode 0xd000375, HT on, Turbo on, Ubuntu 22.04 LTS, 5.15.0-27-generic, 1x INTEL SSDSC2KG960G8, DLRM, Inf: bs=n [1socket/instance], Inference: bs: fp32=128, int8=128, Training bs:fp32=32k [1 instance, 1socket], Criteo Terabyte Dataset, Framework: https://github.com/intel-innersource/frameworks.ai.pytorch.private-cpu/tree/d7607bdd983093396a70713344828a989b766a66; Modelzoo: https://github.com/IntelAI/models/tree/spr-launch-public, PT:1.13, IPEX: 1.13, OneDNN: v2.7, test by Intel on 10/24/2022.

Resources and Configurations

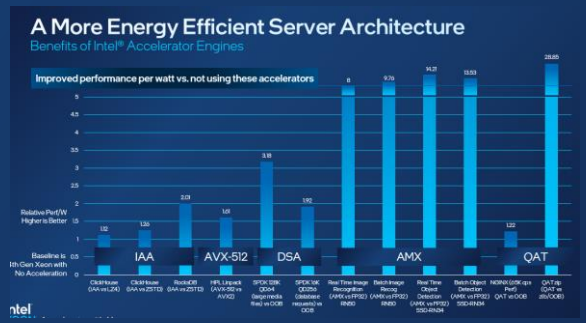

A More Energy Efficient Server Architecture

Up to 1.12x and 1.26x higher performance/W using 4th Gen Xeon Scalable w/Intel Analytics Accelerator vs LZ4 and Zstd on ClickHouse
1-node, 2x pre-production 4th Gen Intel Xeon Scalable processor (60 cores) with integrated Intel In-Memory Analytics Accelerator (Intel IAA), Number of IAA device utilized=8(2 sockets active), on pre-production Intel platform and software, HT On, Turbo On, SNC off, Total Memory 1024GB (16x64GB DDR5 4800), microcode 0x2b0000a1, 1x3.84TB P5510 NVMe, 10GbE x540-AT2, Ubuntu 22.04.1 LTS, 5.18.12-051812-generic, QPL v0.1.21, accel-config-v3.4.6.4, gcc 11.2, Clickhouse 21.12, Star Schema Benchmark, tested by Intel November 2022.

Up to 2.01x higher performance/W using 4th Gen Xeon Scalable w/Intel Analytics Accelerator vs Zstd on RocksDB
1-node, 2x pre-production 4th Gen Intel Xeon Scalable Processor (60 cores) with integrated Intel In-Memory Analytics Accelerator (Intel IAA), on pre-production Intel platform and software, HT On, Turbo On, Total Memory 1024GB (16x64GB DDR5 4800), microcode 0x2b0000a1, 1x3.84TB P5510 NVMe, 10GbE x540-AT2, Ubuntu 22.04.1 LTS, 5.18.12-051812-generic, QPL v0.2.1,accel-config-v3.4.6.4, ZSTD v1.5.2, RocksDB v6.4.6 (db_bench), tested by Intel November 2022.

Up to 1.61 higher performance/W using 4th Gen Xeon Scalable w/AVX-512 vs AVX2 on Linpack
1-node, 2x pre-production 4th Gen Intel® Xeon® Scalable processor (60 core), on pre-production Supermicro SYS-221H-TNR and software with 1024GB DDR5 memory (16x64 GB), microcode 0x2b0000c0, HT On, Turbo On, SNC 4, Ubuntu 22.04.1 LTS, 5.15.0-52-generic, 1x3.84TB P5510 NVMe, 10GbE x540-AT2, One API BaseKit 2022.2.0.262, One API HPC 2022.2.0.191, Linpack ver 2.3, tested by Intel November 2022.

Up to 3.18x and 1.92x higher performance/W using 4th Gen Xeon Scalable w/Data Streaming Accelerator vs out-of-box OS software on SPDK NVMe TCP
1-node, 2x pre-production 4th Gen Intel Xeon Scalable processor (60 core) with integrated Intel Data Streaming Accelerator (Intel DSA), DSA device utilized=1(1 active socket), on pre-production Intel platform and software with 1024GB DDR5 memory (16x64 GB), microcode 0x2b0000a1, 1x3.84TB P5510 NVMe, 10GbE x540-AT2, Ubuntu 22.04.1 LTS, 5.15.0-52-generic, 1x 1.92TB Intel® SSDSC2KG01, 4x 1.92TB Samsung PM1733, 1x Intel® Ethernet Network Adapter E810-2CQDA2, 2x100GbE, FIO v3.30, SPDK 22.05, tested by Intel November 2022.

Up to 8x and 9.76x higher performance/W using 4th Gen Xeon Scalable w/Advanced Matrix Extensions using AMX vs VNNI instructions on ResNet50 Image Processing
1-node, 2x pre-production 4th Gen Intel® Xeon® Scalable processor (60 core) with Intel® Advanced Matrix Extensions (Intel AMX), on pre-production Supermicro SYS-221H-TNR with 1024GB DDR5 memory (16x64 GB), microcode 0x2b0000c0, HT On, Turbo On, SNC Off, CentOS Stream 8, 5.19.16-301.fc37.x86_64, 1x3.84TB P5510 NVMe, 10GbE x540-AT2, Intel TF 2.10, AI Model=Resnet 50 v1_5, best scores achieved: BS1 FP32 8 cores/instance (max. 15ms SLA), BS1 INT8 2 cores/instance (max. 15ms SLA), BS1 AMX 1 core/instance (max. 15ms SLA), BS16 FP32 5 cores/instance, BS16 INT8 5 cores/instance, BS16 AMX 5 cores/instance, using physical cores, tested by Intel November 2022.

Up to 14.21x and 13.53x higher performance/W using 4th Gen Intel Xeon Scalable w/Advanced Matrix Extensions using AMX vs VNNI instructions on SSD-ResNet34 on Object Detection
1-node, 2x pre-production 4th Gen Intel® Xeon® Scalable processor (60 core) with Intel® Advanced Matrix Extensions (Intel AMX), Intel platform with 1024GB DDR5 memory (16x64 GB), microcode 0x2b0000a1, HT On, Turbo On, SNC Off, CentOS Stream 8, 5.19.16-301.fc37.x86_64, 1x3.84TB P5510 NVMe, 10GbE x540-AT2, Intel TF 2.10, AI Model=SSD-ResNet34, best scores achieved: BS1 FP32 60 cores/instance (max. 100ms SLA), BS1 INT8 4 cores/instance (max. 100ms SLA), BS1 AMX 4 core/instance (max. 100ms SLA), BS8 FP32 8 cores/instance, BS2 INT8 1 cores/instance, BS2 AMX 1 cores/instance, using physical cores, tested by Intel November 2022.

Up to 1.22x higher performance/W using 4th Gen Intel Xeon Scalable w/QuickAssist Accelerator vs out-of-box software on NGINX TLS Handshake.
QAT Accelerator: 1-node, 2x pre-production 4th Gen Intel Xeon Scalable Processor (60 cores) with integrated Intel QuickAssist Accelerator (Intel QAT), Number of QAT device utilized=4(1 socket active), on pre-production Intel platform and software with DDR5 memory total 1024GB (16x64 GB), microcode 0x2b0000a1, HT On, Turbo Off, SNC Off, Ubuntu 22.04.1 LTS, 5.15.0-52-generic, 1x3.84TB P5510 NVMe, 10GbE x540-AT2, 1x Intel® Ethernet Network Adapter E810-2CQDA2, 1x100GbE, QAT engine v0.6.14, QAT v20.l.0.9.1, NGINX 1.20.1, OpenSSL 1.1.1l, IPP crypto v2021_5, IPSec v1.1 , TLS 1.3 AES_128_GCM_SHA256, ECDHE-X25519-RSA2K, 65K CPS target SLA, tested by Intel November 2022. Out of box configuration: 1-node, 2x pre-production 4th Gen Intel Xeon Scalable Processor (60 cores) with integrated Intel QuickAssist Accelerator (Intel QAT), Number of QAT device utilized=0, on pre-production Intel platform and software with DDR5 memory total 1024GB (16x64 GB), microcode 0x2b0000a1, HT On, Turbo Off, SNC Off, Ubuntu 22.04.1 LTS, 5.15.0-52-generic, 1x3.84TB P5510 NVMe, 10GbE x540-AT2, 1x Intel® Ethernet Network Adapter E810-2CQDA2, 1x100GbE, NGINX 1.20.1, OpenSSL 1.1.1l, TLS 1.3 AES_128_GCM_SHA256, ECDHE-X25519-RSA2K, 65K CPS target SLA, tested by Intel November 2022.

Up to 28.85x higher performance/W using 4th Gen Intel Xeon Scalable w/QuickAssist Accelerator vs out-of-box zlib on QATzip compression
1-node, 2x pre-production 4th Gen Intel® Xeon Scalable Processor (60 core) with integrated Intel QuickAssist Accelerator (Intel QAT), QAT device utilized=8(2 sockets active), on pre-production Intel platform and software with DDR5 memory Total 1024GB (16x64 GB), microcode 0x2b0000a1, HT On, Turbo Off, SNC Off, Ubuntu 22.04.1 LTS, 5.15.0-52-generic, 1x3.84TB P5510 NVMe, 10GbE x540-AT2, QAT v20.l.0.9.1 , QATzip v1.0.9, tested by Intel November 2022.

Resources and Configurations

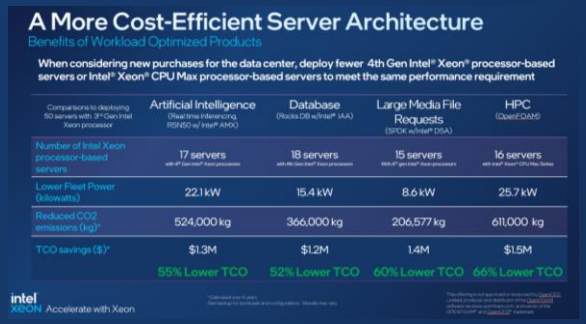

A More Cost-Efficient Server Architecture (1 of 4)

ResNet50 Image Classification

- New Configuration: 1-node, 2x pre-production 4th Gen Intel® Xeon® Scalable 8490H processor (60 core) with Intel® Advanced Matrix Extensions (Intel AMX), on pre-production SuperMicro SYS-221H-TNR with 1024GB DDR5 memory (16x64 GB), microcode 0x2b0000c0, HT On, Turbo On, SNC Off, CentOS Stream 8, 5.19.16-301.fc37.x86_64, 1x3.84TB P5510 NVMe, 10GbE x540-AT2, Intel TF 2.10, AI Model=Resnet 50 v1_5, best scores achieved: BS1 AMX 1 core/instance (max. 15ms SLA), using physical cores, tested by Intel November 2022. Baseline: 1-node, 2x production 3rd Gen Intel Xeon Scalable 8380 Processor (40 cores) on SuperMicro SYS-220U-TNR , DDR4 memory total 1024GB (16x64 GB), microcode 0xd000375, HT On, Turbo On, SNC Off, CentOS Stream 8, 5.19.16-301.fc37.x86_64, 1x3.84TB P5510 NVMe, 10GbE x540-AT2, Intel TF 2.10, AI Model=Resnet 50 v1_5, best scores achieved: BS1 INT8 2 cores/instance (max. 15ms SLA), using physical cores, tested by Intel November 2022.
- For a 50 server fleet of 3rd Gen Xeon 8380 (RN50 w/DLBoost), estimated as of November 2022:
  - CapEx costs: $1.64M
  - OpEx costs (4 year, includes power and cooling utility costs, infrastructure and hardware maintenance costs): $739.9K
  - Energy use in kWh (4 year, per server): 44627, PUE 1.6
  - Other assumptions: utility cost $0.1/kWh, kWh to kg CO2 factor 0.42394
- For a 17 server fleet of 4th Gen Xeon 8490H (RN50 w/AMX), estimated as of November 2022:
  - CapEx costs: $799.4K
  - OpEx costs (4 year, includes power and cooling utility costs, infrastructure and hardware maintenance costs): $275.3K
  - Energy use in kWh (4 year, per server): 58581, PUE 1.6
  - Other assumptions: utility cost $0.1/kWh, kWh to kg CO2 factor 0.42394

Resources and Configurations

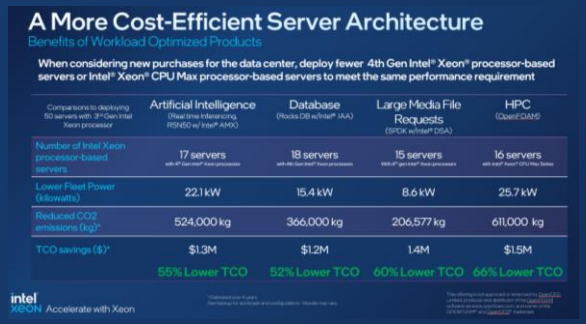

A More Cost-Efficient Server Architecture (2 of 4)

RocksDB

- New Configuration: 1-node, 2x pre-production 4th Gen Intel Xeon Scalable 8490H Processor (60 cores) with integrated Intel In-Memory Analytics Accelerator (Intel IAA), on pre-production Intel platform and software, HT On, Turbo On, Total Memory 1024GB (16x64GB DDR5 4800), microcode 0x2b0000a1, 1x3.84TB P5510 NVMe, 10GbE x540-AT2, Ubuntu 22.04.1 LTS, 5.18.12-051812-generic, QPL v0.2.1,accel-config-v3.4.6.4, ZSTD v1.5.2, RocksDB v6.4.6 (db_bench), tested by Intel November 2022. Baseline: 1-node, 2x production 3rd Gen Intel Xeon Scalable 8380 Processor (40 cores) on SuperMicro SYS-220U-TNR , HT On, Turbo On, SNC Off, Total Memory 1024GB (16x64GB DDR4 3200), microcode 0xd000375, 1x3.84TB P5510 NVMe, 10GbE x540-AT2, Ubuntu 22.04.1 LTS, 5.18.12-051812-generic, ZSTD v1.5.2, RocksDB v6.4.6 (db_bench), tested by Intel November 2022.
- For a 50 server fleet of 3rd Gen Xeon 8380 (RocksDB), estimated as of November 2022:
  - CapEx costs: $1.64M
  - OpEx costs (4 year, includes power and cooling utility costs, infrastructure and hardware maintenance costs): $677.7K
  - Energy use in kWh (4 year, per server): 32181, PUE 1.6
  - Other assumptions: utility cost $0.1/kWh, kWh to kg CO2 factor 0.42394
- For a 18 server fleet of 4th Gen Xeon 8490H (RockDB w/IAA), estimated as of November 2022:
  - CapEx costs: $846.4K
  - OpEx costs (4 year, includes power and cooling utility costs, infrastructure and hardware maintenance costs): $260.6K
  - Energy use in kWh (4 year, per server): 41444, PUE 1.6
  - Other assumptions: utility cost $0.1/kWh, kWh to kg CO2 factor 0.42394

Resources and Configurations

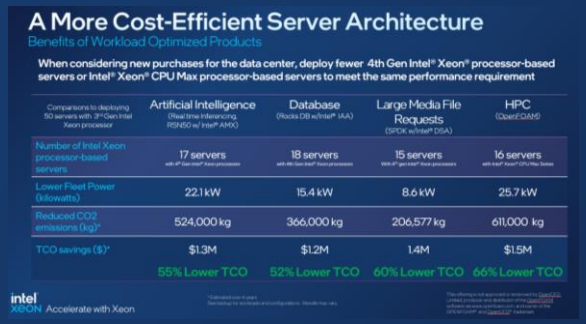

A More Cost-Efficient Server Architecture (3 of 4)

OpenFOAM

- New Configuration: 1-node, 2x pre-production 4th Gen Intel Xeon CPU Max Series (56 cores) on pre-production Intel platform and software, HT On, Turbo On, SNC4 mode, Total Memory 128 GB (8x16GB HBM2 3200MT/s), microcode 0x2c000020, 1x3.5TB INTEL SSDPF2KX038TZ NVMe, CentOS Stream 8, 5.19.0-rc6.0712.intel_next.1.x86_64+server, OpenFOAM 8, Motorbike 20M @ 250 iterations, Motorbike 42M @ 250 iterations, Tools: ifort:2021.6.0, icc:2021.6.0, impi:2021.6.0, tested by Intel December 2022. Baseline: 1-node, 2x production 3rd Gen Intel Xeon Scalable 8380 Processor (40 cores) on SuperMicro SYS-220U-TNR, HT On, Turbo On, 512GB (16x32GB DDR4 3200 MT/s), microcode 0xd000375, 1x2.9TB INTEL SSDPE2KE032T8 NVMe, CentOS Stream 8, 4.18.0-408.el8.x86_64, OpenFOAM 8, Motorbike 20M @ 250 iterations, Motorbike 42M @ 250 iterations, Tools: ifort:2021.6.0, icc:2021.6.0, impi:2021.6.0, tested by Intel December 2022
- For a 50 server fleet of 3rd Gen Xeon 8380 (OpenFOAM), estimated as of December 2022:
  - CapEx costs: $1.50M
  - OpEx costs (4 year, includes power and cooling utility costs, infrastructure and hardware maintenance costs): $780.3K
  - Energy use in kWh (4 year, per server): 52700, PUE 1.6
  - Other assumptions: utility cost $0.1/kWh, kWh to kg CO2 factor 0.42394
- For a 16 server fleet of Intel Xeon CPU Max Series 56 core, estimated as of December 2022:
  - CapEx costs: $507.2K
  - OpEx costs (4 year, includes power and cooling utility costs, infrastructure and hardware maintenance costs): $274.9K
  - Energy use in kWh (4 year, per server): 74621, PUE 1.6
  - Other assumptions: utility cost $0.1/kWh, kWh to kg CO2 factor 0.42394

Resources and Configurations

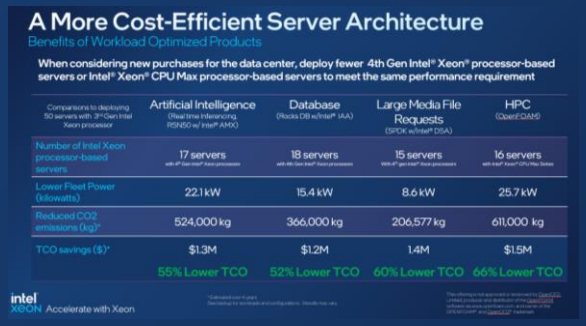

A More Cost-Efficient Server Architecture (4 of 4)

SPDK

- New Configuration: 1-node, 2x pre-production 4th Gen Intel Xeon Scalable processor (60 core) with integrated Intel Data Streaming Accelerator (Intel DSA), DSA device utilized=1(1 active socket), on pre-production Intel platform and software with 1024GB DDR5 memory (16x64 GB), microcode 0x2b0000a1, 10GbE x540-AT2, Ubuntu 22.04.1 LTS, 5.15.0-52-generic, 1x 1.92TB Intel® SSDSC2KG01, 4x 1.92TB Samsung PM1733, 1x Intel® Ethernet Network Adapter E810-2CQDA2, 2x100GbE, FIO v3.30, SPDK 22.05, tested by Intel November 2022. Baseline: 1-node, 2x production 3rd Gen Intel Xeon Scalable Processors(40 cores) on Supermicro SYS-220U-TNR , DDR4 memory total 1024GB (16x64 GB), HT On, Turbo On, SNC Off, microcode 0xd000375, 10GbE x540-AT2, Ubuntu 22.04.1 LTS, 5.15.0-52-generic, 1x 1.92TB Intel SSDSC2KG01, 4x 1.92TB Samsung PM1733, 1x Intel Ethernet Network Adapter E810-2CQDA2, 2x100GbE, FIO v3.30, SPDK 22.05, tested by Intel November 2022.
- For a 50 server fleet of 3rd Gen Xeon 8380 (SPDK), estimated as of November 2022:
  - CapEx costs: $1.77M
  - OpEx costs (4 year, includes power and cooling utility costs, infrastructure and hardware maintenance costs): $630.6K
  - Energy use in kWh (4 year, per server): 22762, PUE 1.6
  - Other assumptions: utility cost $0.1/kWh, kWh to kg CO2 factor 0.42394
- For a 15 server fleet of 4th Gen Xeon 8490H (SPDK w/DSA), estimated as of November 2022:
  - CapEx costs: $743.8K
  - OpEx costs (4 year, includes power and cooling utility costs, infrastructure and hardware maintenance costs): $220.1K
  - Energy use in kWh (4 year, per server): 43387, PUE 1.6
  - Other assumptions: utility cost $0.1/kWh, kWh to kg CO2 factor 0.42394

Leadership Performance with 4th Gen Intel® Xeon® Processors
Disclaimers

- 53% average performance gain over the prior generation1 See [G1] at intel.com/processorclaims: 4th Gen Intel® Xeon® Scalable processors. Results may vary.
- Up to 10x higher PyTorch real-time inference performance with built-in Intel® Advanced Matrix Extensions (Intel® AMX) (BF16) vs. the prior generation (FP32)4 See [A17] at intel.com/processorclaims: 4th Gen Intel® Xeon® Scalable processors. Results may vary.
- Up to 10x higher PyTorch training performance with built-in Intel® Advanced Matrix Extensions (Intel® AMX) (BF16) vs. the prior generation (FP32)5 See [A16] at intel.com/processorclaims: 4th Gen Intel® Xeon® Scalable processors. Results may vary.
- Up to 5:1 consolidation and 75% TCO savings with 4th Gen Intel Xeon processors: Calculations as of March 28, 2023 based on the Intel® Node TCO & Power Calculator using default cost, power and TCO assumptions over a 5-year TCO horizon comparing replacing 50 older servers with Intel Xeon 4110 processors with new servers using new Intel Xeon 5420+ processors. Results may vary. Performance measurements based on published SPECrate®2017_int_base on spec.org as of March 28, 2023 [4110: https://www.spec.org/cpu2017/results/res2020q4/cpu2017-20201015-24218.html 5420+: https://www.spec.org/cpu2017/results/res2023q1/cpu2017-20230130-33925.html]
- 2.9x average performance per watt efficiency improvement for targeted workloads utilizing built-in accelerators compared to the previous generation2 See [E1] at intel.com/processorclaims: 4th Gen Intel® Xeon® Scalable processors. Results may vary.
- Intel® SGX is the most researched, updated, and deployed confidential computing technology in data centers on the market today. With Intel® Security Engines, 4th Gen Intel® Xeon® Scalable processors help bring a zero-trust security strategy to life while unlocking new opportunities for business collaboration and insights—even with sensitive or regulated data. Intel® Software Guard Extensions (Intel® SGX) is designed to enhance data protection at rest, in motion, and in use. Intel SGX is the most researched, updated, and deployed confidential computing technology in data centers on the market today. Intel SGX provides the smallest trust boundary of any confidential computing technology in the data center today.
- Built-in accelerators for encryption help keep data protected while preserving performance. Intel® Crypto Acceleration reduces the impact of implementing pervasive data encryption and increases the performance of encryption-sensitive workloads, such as for Secure Sockets Layer (SSL) web servers, 5G infrastructure, and VPNs/firewalls.. Networking Encryption: Up to 47% fewer cores to achieve the same connections/second using integrated Intel® QuickAssist Technology (Intel® QAT) vs. the prior generation on NGINX key handshake.4 See [N15] at intel.com/processorclaims: 4th Gen Intel® Xeon® Scalable processors. Results may vary.
- Most deployed platform, backed by extensive testing and validation: With more deployments than any other data center CPU in the market, Intel® Xeon® Scalable processors are widely trusted to run critical workloads at scale. From next-gen memory and I/O to software optimizations, 4th Gen Intel Xeon Scalable processors have been extensively tested and validated to deliver the high performance and reliability organizations demand.
- Businesses can speed up time to deployment with the largest ecosystem of partners they know and use—hardware and software vendors and solution integrators around the world build their products on Intel® Xeon® Scalable processors, offering maximum choice and interoperability with the reassurance of thousands of real-world implementations

Refresh and consolidate Intel® Xeon® processor-based servers
Disclaimers

Up to 5:1 consolidation with 75% TCO reduction with 4th Gen Intel Xeon processors

Calculations as of March 28, 2023, based on the Intel® Node TCO & Power Calculator using default cost, power and TCO assumptions over a 5-year TCO horizon comparing replacing 50 older servers with Intel Xeon 4110 processors with new servers using new Intel Xeon 5420+ processors. Results may vary. Performance measurements based on published SPECrate®2017_int_base on spec.org as of March 28, 2023.
4110: https://www.spec.org/cpu2017/results/res2020q4/cpu2017-20201015-24218.html
5420+: https://www.spec.org/cpu2017/results/res2023q1/cpu2017-20230130-33925.html]

4th Gen Intel® Xeon® processors can significantly lower your total cost of ownership

Calculations as of March 28, 2023, based on the Intel® Node TCO & Power Calculator using default cost, power and TCO assumptions over a 5-year TCO horizon comparing replacing 50 older servers with Intel Xeon 4110 processors with new servers using new Intel Xeon 5420+ processors. Results may vary. Performance measurements based on published SPECrate®2017_int_base on spec.org as of March 28, 2023.

8160	https://www.spec.org/cpu2017/results/res2018q4/cpu2017-20181112-09655.html
8460Y	https://www.spec.org/cpu2017/results/res2023q1/cpu2017-20221223-33229.html
6130	https://www.spec.org/cpu2017/results/res2019q2/cpu2017-20190506-13570.html
6430	https://www.spec.org/cpu2017/results/res2023q1/cpu2017-20221223-33187.html
5120	https://www.spec.org/cpu2017/results/res2018q4/cpu2017-20181015-09160.html
5420+	https://www.spec.org/cpu2017/results/res2023q1/cpu2017-20230130-33925.html
4110	https://www.spec.org/cpu2017/results/res2020q4/cpu2017-20201015-24218.html

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.

Intel contributes to the development of benchmarks by participating in, sponsoring, and/or contributing technical support to various benchmarking groups, including the BenchmarkXPRT Development Community administered by Principled Technologies.

Your costs and results may vary.

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

Some results may have been estimated or simulated.

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

All product plans and roadmaps are subject to change without notice.

Statements in this document that refer to future plans or expectations are forward-looking statements. These statements are based on current expectations and involve many risks and uncertainties that could cause actual results to differ materially from those expressed or implied in such statements. For more information on the factors that could cause actual results to differ materially, see our most recent earnings release and SEC filings at www.intc.com.