Data Center Modernization
Enablement Package

How our valued partners can address customers' business challenges with Intel based solutions

Microsoft Data Center Products Refresh on 4th Generation Intel® Xeon® Scalable Processors

Windows Server 2022

Windows Server 2012

Microsoft ended support for Windows Server 2012 in October 2023

Why Upgrade Your Hardware Now?

Upgrading to Windows Server 2022 on older, out of warranty hardware will **not allow** you to use all the features of Windows Server

Upgrading to Intel® 4th Generation Xeon® processors unlocks the full value of Windows Server 2022 for enhanced TCO and Security

LEARN HOW

Data Center Modernization Infographic

Intel® technologies may require enabled hardware, software, or service activation.
No product or component can be absolutely secure.

4th Gen Intel® Xeon® Scalable Processor and Windows Server 2022

Keep existing versions of Intel hardware and Microsoft software

- Standard security
- More challenging administration
- Microsoft support ended October 2023

Upgrade hardware only

New Intel hardware and existing Microsoft software

- Increased security
- Lower power consumption

Upgrade BOTH current hardware and software

Upgrade software only

Existing Intel hardware and new Microsoft software

- Security enhancements
- Improved container support
- Hybrid capabilities with Azure

New Intel hardware and Microsoft software

- Performance improvements
- Increased security
- Better scalability
- Network improvements

Intel® technologies may require enabled hardware, software, or service activation.
No product or component can be absolutely secure.

Better Together: Windows Server 2022 + 4th Gen Intel® Xeon® Scalable Processor

Upgrade Paths / Value Propositions

Software Modernization ▲ NEW

SW ONLY UPGRADE

1st and 2nd Gen Intel® Xeon® Processors

- ✗ Lesser Platform Security (UEFI, TPM 2.0, Secure Boot)
- ✗ No Shielded VMs (Requires TPM 2.0)
- ✗ No HCI Management
- ✗ No Advanced Flash Support (NVMe, NVDIMM, or Intel® Optane™ SSD)

Windows Server 2022

- ✓ Security Enhancements (VBS, Windows Defender)
- ✓ Performance Improvements (Faster boot, login times)
- ✓ Improved Container support
- ✓ Hybrid Capabilities w/Azure (Arc Enabled; SMB Compression; WAC)
- ✓ Scalability (48TB Memory; 2048 Logical Cores across 64 sockets)

HW + SW UPGRADE

4th Gen Intel® Xeon® Processors

- ✓ Platform Security (UEFI, TPM 2.0, Secure Boot, Secure Core, TME-MK)
- ✓ Support for NVMe, NVDIMM, Intel® Optane™ SSD
- ✓ Up to 60 cores / 120 threads on 2-socket system, 6TB of memory
- ✓ Support for 100Gb NICs
- ✓ Accelerators (AMX, QAT)

Windows Server 2022

- ✓ Security Enhancements (VBS, Windows Defender)
- ✓ Performance Improvements (Faster boot, login times)
- ✓ Improved Container support
- ✓ Hybrid Capabilities w/Azure (Arc Enabled; SMB Compression; WAC)
- ✓ Scalability (48TB Memory; 2048 Logical Cores across 64 sockets)

CURRENT HW & SW – DO NOTHING

1st and 2nd Gen Intel® Xeon® Processors

- ✗ Lesser Platform Security (UEFI, TPM 2.0, Secure Boot)
- ✗ No Shielded VMs (Requires TPM 2.0)
- ✗ No HCI Management
- ✗ No Advanced Flash Support (NVMe, NVDIMM, or Intel® Optane™ SSD

Windows Server 2012

- ✗ Less secure
- ✗ Non-compliance. Cost of Extended Security Updates
- ✗ Lack of hybrid/multi-cloud integration

HW ONLY UPGRADE

4th Gen Intel® Xeon® Processors

- ✓ Platform Security (UEFI, TPM 2.0, Secure Boot, Secure Core, TME-MK)
- ✓ Support for NVMe, NVDIMM, Intel® Optane™ SSD
- ✓ Up to 60 cores / 120 threads on 2-socket system, 6TB of memory
- ✓ Support for 100Gb NICs
- ✓ Accelerators (AMX, QAT)

Windows Server 2012

- ✗ OS incapable of taking advantage of new HW security features (Secure Core, TME-MK, Shielded VMs)
- ✗ Unable to scale to take advantage of HW innovations (Cores, Memory and VM Support)
- ✗ Lack of hybrid / multi-cloud integration
- ✗ Non-compliance. Cost of Extended Security Updates

OLD — Hardware Modernization ▶ NEW

Intel® technologies may require enabled hardware, software, or service activation.
No product or component can be absolutely secure.

4th Gen Intel® Xeon® Scalable Processor Summary

The 4th Gen Intel® Xeon® Scalable platform improves throughput and efficiency gen-over-gen, enabling server consolidation, improving TCO and helping achieve sustainability goals.

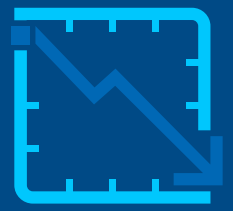

4th Gen Intel® Xeon® Scalable systems run complex AI workloads on the same hardware as your existing workloads, leveraging your data to accelerate insights and innovation.

4th Gen Intel® Xeon® Scalable systems help secure and protect your data across environments without compromising performance.

Microsoft SQL Server 2022

Microsoft SQL Server 2012

Microsoft ended support for Microsoft SQL 2012

Why Upgrade Your Hardware Now?

Upgrading to Microsoft SQL Server 2022 on older, out of warranty hardware will not allow you to use all the features of SQL 2022

Upgrading to Intel® 4th Generation Xeon® processors unlocks the full value of Windows Server 2022 for enhanced TCO and Security

LEARN HOW

Data Center Modernization Infographic

Intel® technologies may require enabled hardware, software, or service activation.
No product or component can be absolutely secure.

4th Gen Intel® Xeon® Scalable Processor and Microsoft SQL Server

Flexible, scalable, cloud-connected database for modern workloads

Keep existing versions of Intel hardware and Microsoft software

- Standard security
- More challenging administration
- Microsoft support no longer available

Upgrade hardware only

New Intel hardware and existing Microsoft software

- Faster database backups
- Increased security
- Lower power consumption

Upgrade BOTH current hardware and software

Upgrade software only

Existing Intel hardware and new Microsoft software

- Unified data management
- Cloud-based analytics
- Access to latest software features

New Intel hardware and Microsoft software

- Faster, real-time insights
- Increased security
- Significantly easier administration

Microsoft® SQL Server®

Intel® technologies may require enabled hardware, software, or service activation.
No product or component can be absolutely secure.

Better Together: Microsoft SQL Server 2022 + 4th Gen Intel® Xeon® Processor

Upgrade Paths / Value Propositions

Microsoft SQL Server

SW ONLY UPGRADE

1st and 2nd Gen Intel® Xeon® Processors

- ✗ Limited platform security
- ✗ No Shielded VMs
- ✗ No HCI Management
- ✗ No Advanced Flash Support (NVMe, NVDIMM, or Intel® Optane™ SSD)

SQL Server 2022

- ✓ Reduced workload times with much faster queries
- ✓ Seamless analytics over on-prem, operational data
- ✓ More secure database
- ✓ High availability, business continuity
- ✓ Resolution of conflicts with different data replicas is automatic
- ✓ Unified data governance and management

HW + SW UPGRADE

4th Gen Intel® Xeon® Processors

- ✓ 114% faster database backups with QAT
- ✓ Up to 34% more NOPM transactions compared to 3rd gen
- ✓ Up to 28% faster query response compared to 3rd Gen
- ✓ MSFT Storage Spaces Direct w/NVME, Intel® Optane™
- ✓ Support for NVMe, NVDIMM, Intel® Optane™ SSD
- ✓ Up to 60 cores / 120 threads on 2- socket system, 6TB of memory
- ✓ Support for 100Gb NICs
- ✓ Platform Security (UEFI, TPM 2.0, Secure Boot)

SQL Server 2022

- ✓ Reduced workload times with much faster queries
- ✓ Seamless analytics over on-prem, operational data
- ✓ More secure database
- ✓ High availability, business continuity
- ✓ Resolution of conflicts with different data replicas is automatic
- ✓ Unified data governance and management

CURRENT HW & SW – DO NOTHING

1st and 2nd Gen Intel® Xeon® Processors

- ✗ Limited platform security
- ✗ No Shielded VMs
- ✗ No HCI Management
- ✗ No Advanced Flash Support (NVMe, NVDIMM, or Intel® Optane™ SSD

SQL Server 2012 / 2016 / 2019

- ✗ More difficulty with real-time analysis
- ✗ Lower availability
- ✗ Less encryption and protection
- ✗ Administration can be more difficult
- ✗ Lose opportunity to migrate to Azure or hybrid cloud at same time
- ✗ Resolution of conflicts with different data replicas is manual

HW ONLY UPGRADE

4th Gen Intel® Xeon® Processors

- ✓ 114% faster database backups with QAT
- ✓ Up to 34% more NOPM transactions compared to 3rd gen
- ✓ Up to 28% faster query response compared to 3rd Gen
- ✓ MSFT Storage Spaces Direct w/NVME, Intel® Optane™
- ✓ Support for NVMe, NVDIMM, Intel® Optane™ SSD
- ✓ Up to 60 cores / 120 threads on 2-socket system, 6TB of memory
- ✓ Support for 100Gb NICs
- ✓ Increased platform security (UEFI, TPM 2.0, Secure Boot)

SQL Server 2012 / 2016 / 2019

- ✗ More difficulty with real-time analysis
- ✗ Lower availability
- ✗ Less encryption and protection
- ✗ Administration can be more difficult
- ✗ Lose opportunity to migrate to Azure or hybrid cloud at same time
- ✗ Resolution of conflicts with different data replicas is manual

Software Modernization ▲ NEW

OLD — Hardware Modernization ▶ NEW

Intel® technologies may require enabled hardware, software, or service activation.
No product or component can be absolutely secure.

See backup for workloads and configurations. Results may vary.

Summary

Performance across Workloads

Experience up to **34% faster** transaction processing[1], **23% faster** query response times, and **53%[2] faster** backup times for Gen-over-Gen performance[3] on SQL Server 2022 when running on 4th Gen Intel® Xeon® Scalable processor.

Highly Secured Data Platform

SQL Server is the **most secure database**[4] and run it on 4th Gen Intel® Xeon® Scalable systems for enabling confidential and protection to the silicon level.

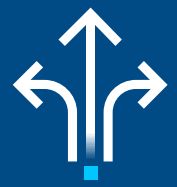

Ready for Modernization

SQL Server 2022 and 4th Gen Intel® Xeon® deliver **cloud-connected, flexible platform to leverage digital transformation opportunities**.

[1, 2, 3], See backup slides in Microsoft SQL Server 2022 on Intel® Technologies for workloads and configuration

[4] According to the National Institute of Standards and Technology Comprehensive Vulnerability Database, as of September 2022

Microsoft Azure Stack HCI

Why Modernize Your Business with Microsoft Azure Stack HCI & Intel?

Portability

Intel® Virtualization Technology is foundational for seamless movement of data and apps between the data center, ever-growing edge, and Azure public cloud

Flexibility

Tailor performance and capacity to your needs with Intel's broad portfolio of products and accelerators, tools, libraries and frameworks

Increased HW-based Security

with Intel® Total Memory Encryption, Intel® Crypto Acceleration, and Secured Core

Proven Solutions

Tested, verified for modernizing existing IT environments; certified by Microsoft, validated by the server vendor, and with performance verified by Intel

Enhance your infrastructure with hybrid cloud, optimized by trusted partners

Intel® technologies may require enabled hardware, software, or service activation. No product or component can be absolutely secure.

Microsoft Azure Stack HCI Using 4th Gen Intel® Xeon® Scalable Processors Delivers Performance for Demanding Workloads

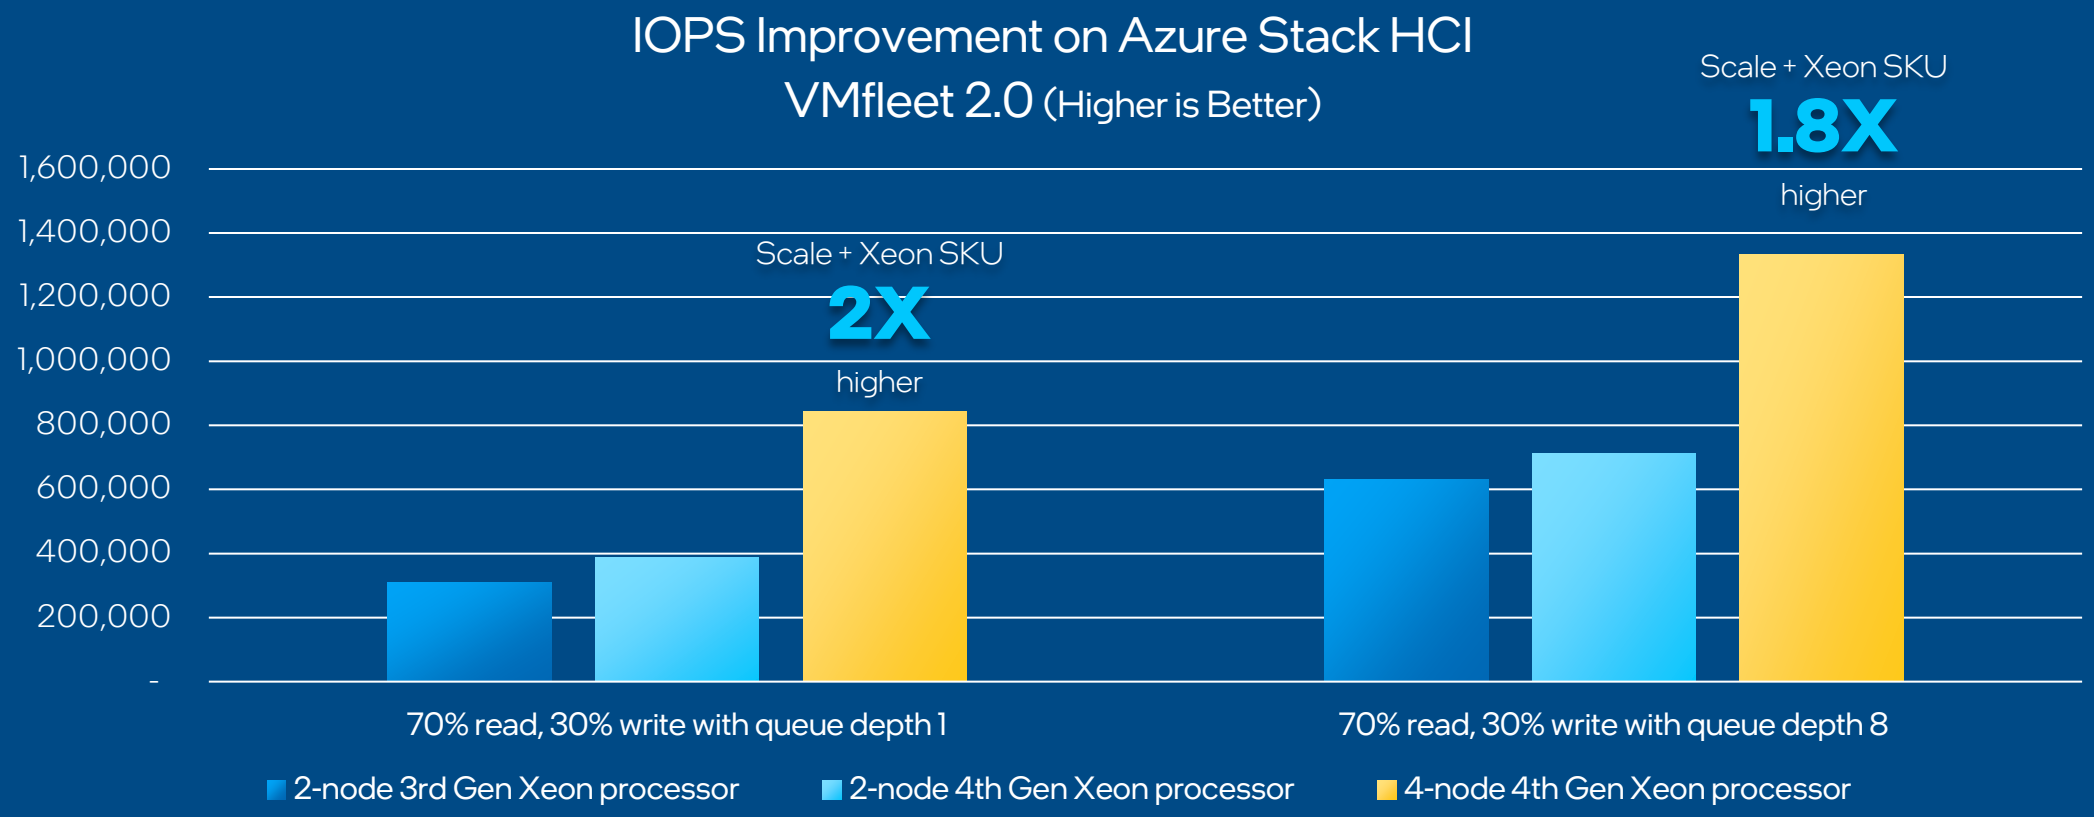

Performance increased gen-to-gen, scales with cluster size, and higher series of Xeon

Results using 4th Intel ®Xeon® Processor - QS processors and silver systems. Performance varies by part, use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex. See backup in Microsoft Azure Stack HCI on 4th Gen Intel® Xeon® Scalable Processors for workloads and configurations. Results may vary.

Accelerate AI - Image Classification on Microsoft Azure Stack HCI using 4th Gen Intel® Xeon® Scalable Processors with Intel® AMX

TensorFlow
Higher is better

Image Classification on Tensorflow 2.11 using ResNet50

BS=128, Multi-instance (16x2 and 40x2 instances)

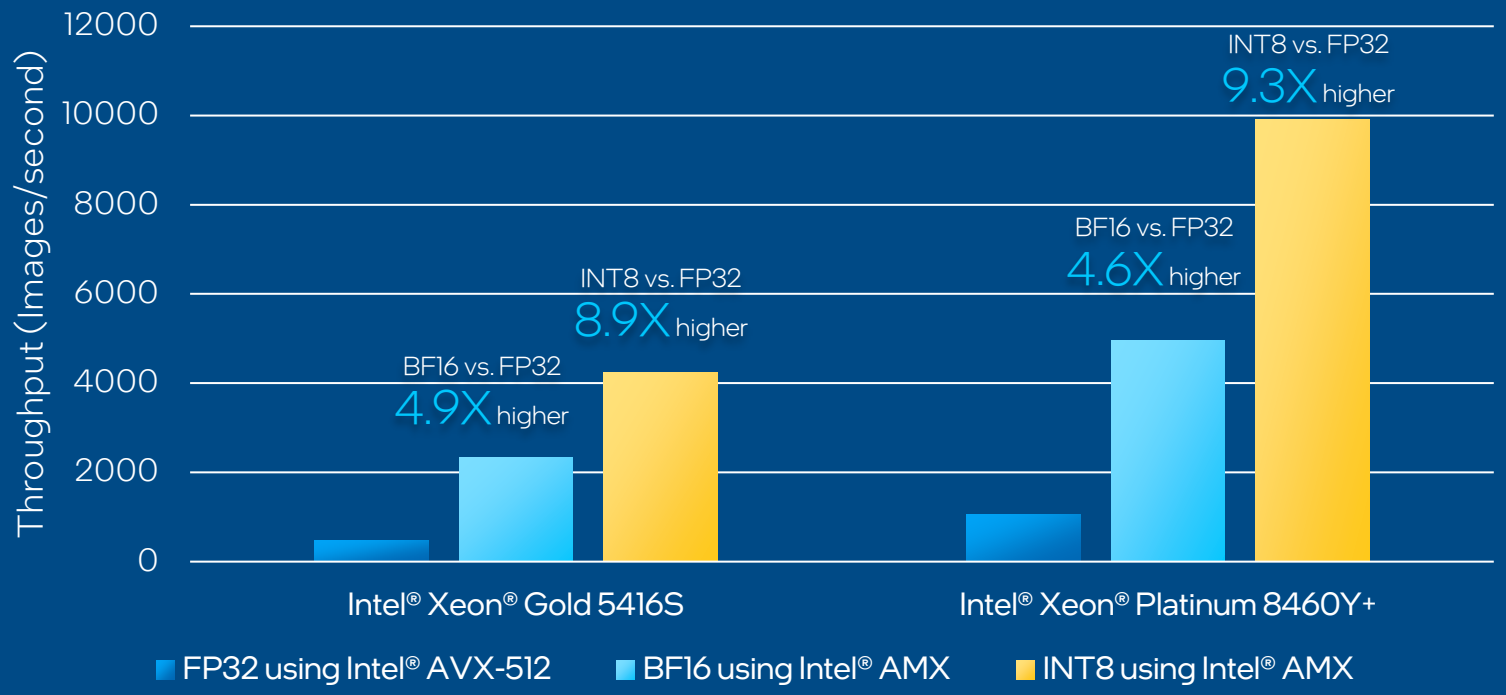

- The ResNet-50 benchmark measures image classification/vision workloads
- FP32 is a standard 32-bit floating point data type used to train deep learning models and for inferencing
- Bfloat16 is a truncated version of 32-bit floating point, used for both training and inference, offering similar accuracy but faster computation
- INT8 offers higher performance and is least computationally demanding for constrained environments, with minimal impact on accuracy
- Many DL workloads are mixed precision and 4th Gen Intel Xeon Scalable processors can seamlessly transition between Intel AMX and Intel AVX-512 to use the most efficient instruction set

Increase performance with higher series of Intel® Xeon® processor or by changing precision

Results using 4th Gen Intel Xeon Processor - pre-production processors and systems. Performance varies by part, use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.
See backup in Microsoft Azure Stack HCI on 4th Gen Intel® Xeon® Scalable Processors for workloads and configurations. Results may vary. Intel AVX-512=Intel Advanced Vector Extensions 512, Intel AMX=Intel Advanced Matrix Extensions

Accelerate Natural Language Processing (NLP) on Microsoft Azure Stack HCI using 4th Gen Intel® Xeon® Scalable Processors with Intel® AMX

TensorFlow
Higher is better

NLP on Tensorflow 2.11 using BERT-Large

BS=128, Multi-instance (16x2 and 40x2 instances)

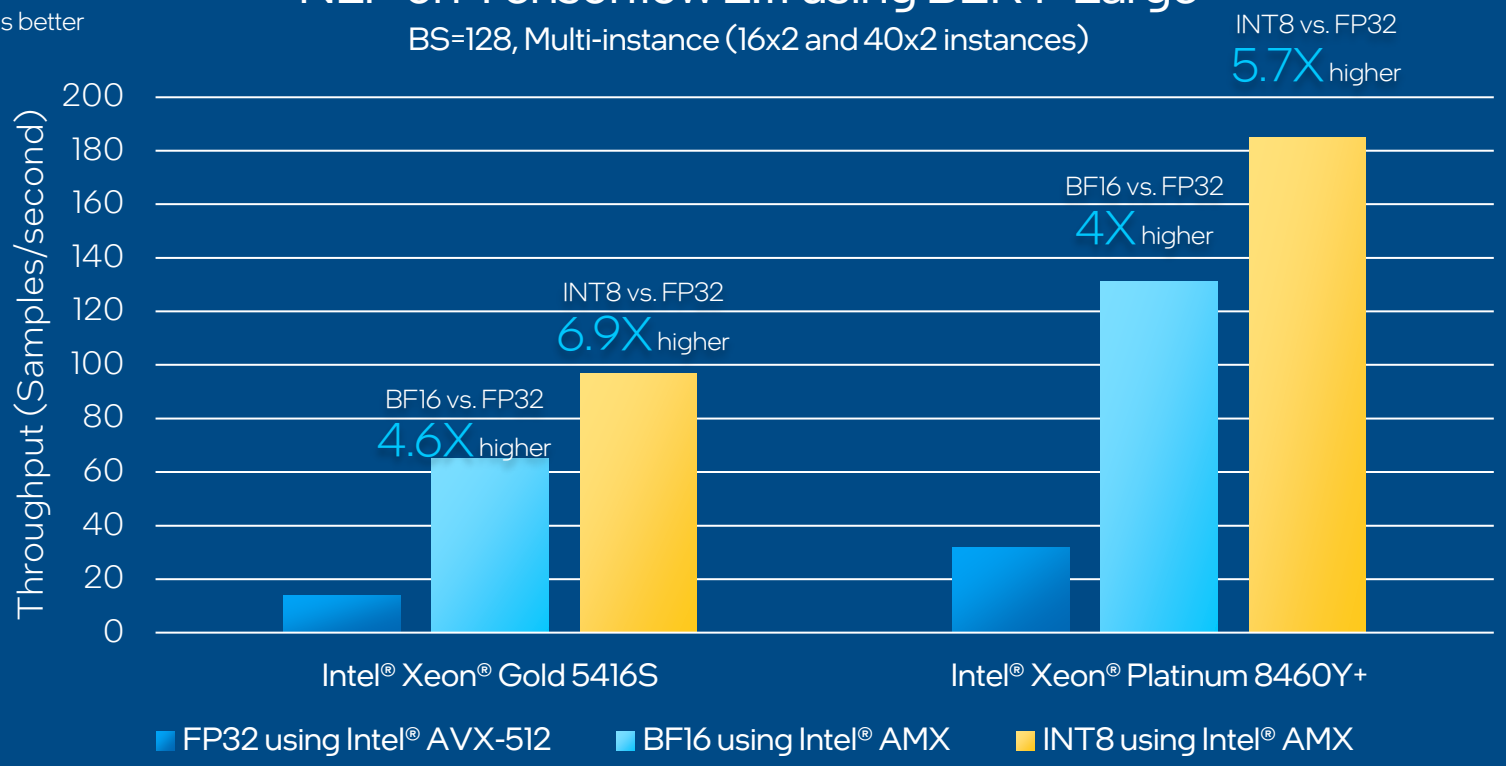

- BERT-Large is a pretrained model used for Natural Language Processing
- FP32 is a standard 32-bit floating point data type used to train deep learning models and for inferencing
- Bfloat16 is a truncated version of 32-bit floating point, used for both training and inference, offering similar accuracy but faster computation
- INT8 offers higher performance and is least computationally demanding for constrained environments, with minimal impact on accuracy
- Many DL workloads are mixed precision and 4th Gen Intel® Xeon® Scalable processors can seamlessly transition between Intel® AMX and Intel® AVX-512 to use the most efficient instruction set

Increase performance with higher series of Intel Xeon processor or by changing precision

Results using 4th Gen Intel® Xeon® Processor - pre-production processors and systems. Performance varies by part, use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.
See backup in Microsoft Azure Stack HCI on 4th Gen Intel® Xeon® Scalable Processors for workloads and configurations. Results may vary. Intel AVX-512=Intel Advanced Vector Extensions 512, Intel AMX=Intel Advanced Matrix Extensions

Data Center Modernization & Optimization

ACCESS NOW

- Reasons to Modernize: A Conversation Guide

Why Modernize Now?

New Applications Demand New Infrastructure

Efficiency

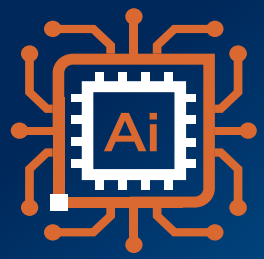

Artificial Intelligence

Security

[5 ways Intel® Xeon® Scalable processors can support your fastest-growing workloads]


Efficiency

Significant Performance Leaps

5th Gen Intel® Xeon® CPUs provides generational improvements on CPU and platform upgrades

	General Purpose Compute	AI	HPC	Networking and Storage
CPU upgrade 4th Gen Intel® Xeon® CPU vs. 5th Gen Intel Xeon CPU	**1.21x** average performance gain	up to **1.42x** higher inference	up to **1.4x** higher HPC performance gain	up to **1.7x** higher throughput
Server platform upgrade 3rd Gen Intel® Xeon® CPU vs. 5th Gen Intel Xeon CPU	**1.84x** average performance gain	up to **14x** higher inference and training performance	up to **2.1x** average performance gain	up to **3.6x** higher throughput

See [G1, A17, A16, H7, H1, N6, N13] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.

intel

Significant Performance Gains

5th Gen Intel® Xeon® Scalable processors vs. 3rd Gen Intel® Xeon® processors

Artificial Intelligence

Up to **14x**

higher inference and training performance

Infrastructure & Storage

Up to **2.8x**

higher IOPs and up to 65% latency reduction for large packet sequential read and up

Network

3.2x

average higher performance on broadly-deployed network workloads

HPC

3x

higher LAMMPS performance

Database & Analytics

3.7x

higher RocksDB performance using integrated Intel® IAA

Web & Microservices

Up to **2x**

higher Java throughput within a given SLA

See [A17, N16, N11, H10, D1, W3] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.

intel

The Sustainable Data Center

Reduce power consumption, even in demanding workloads with Intel's configurable hardware and intelligent software

Improve **efficiency and performance** across targeted workloads

Up to **10x** improved performance/ watt[2] on AI workloads with built-in acceleration*

*in 5th Gen Intel® Xeon® Scalable processors with Intel® AMX

Increase **energy efficiency** on lower-utilization workloads

Up to **14%** performance/watt improvement at ~50% load with Optimized Power Mode enabled[3] *

*on 5th Gen Intel Xeon, versus OPM disabled

Enable AI and ML with carbon- and power-efficient **deep learning training and interference**

79% higher throughput/watt with Intel® Gaudi®2 vs. NVIDIA H100[4]

Save power deploying **fewer new servers** to meet performance goals

Up to **1,482 MWh** fleet energy savings with 5th Gen Intel Xeon* on data storage workloads[5]

*vs. 3rd Gen Intel® Xeon® processor based servers over 4 years

LEARN MORE

The Sustainable Data Center

Sustainability with Intel Technologies

The Sustainable CTO: The Road to Tech Positive

Intel Sustainability: Server Consolidation

WATCH >

Cloud TV: Sustainability and the Cloud

Discover 5 ways you can help your customers practice sustainability with Intel technologies

[2,3,4,5] Performance configurations outlined in The Sustainable Data Center Infographic above. Results may vary.

The Sustainable Data Center

How Intel® Xeon® processor–powered servers, compares to AMD EPYC processor–powered servers

96% Less electricity used in AI workloads[1]

51% Less electricity used in Network Infrastructure[2]

39% Less electricity used in HPC workloads[3]

LEARN MORE

5 Reasons Why Processor Selection Makes a Difference Infographic

[1, 2, 3,] Performance configurations outlined in the 5 Reasons Why Processor Selection Makes a Difference Infographic above. Results may vary.

Liquid Cooling

Improving TCO through Energy Efficiency and Water Reduction

Liquid Cooling Benefits

Power / Performance

- Reduce PUE[1] (Power Usage Effectiveness) from 1.3+ to as low as 1.03[3]
- Lower power consumption by up to 30%[3]
- Extends cooling range for higher system thermals

Density

- More compute volume within same rack footprint[4]
- Less real estate needed per compute output[4]

Water

- Liquid cooling can significantly reduce the billions of gallons of water used in air cooled data centers[2]

[1]PUE= power consumed by the entire data center divided by power consumed by IT equipment in it

[2]LiquidStack, 2022

[3]GRC Cooling

[4]Source: Intel analysis

Liquid Cooling
Liquid Cooling Solution Benefits

Energy

Up to 40%[2]
reduction in TCO

~1.03 Enables PUEs to 1.03[1]
PUE: Power Usage Effectiveness

40% Reduction in cooling CAPEX[1]

95% Reduction in cooling OPEX[1]

~30% Reduction in power consumption[1]

Water

Up to 100%
reduction in water use[1]

Use Heat for District Heating

Use Heat for Urban Farming

If evaporative cooling is used, water reduction can still be significant over conventional rack air cooling

Density

10x
Increase in compute density[2]

Enables compute in dense edge environments

Increases compute density per sqm

Eliminates physical components (fans, chillers)

Forecasted Growth for Liquid Cooling (2022-2027): 36.3%[3]

PUE: Power Usage Effectiveness =
amount of power entering data center/amount of power needed to run IT equipment

[1]Source: Green Revolution Cooling
Water reduction dependent on system and heat reuse

[2]Source: Submer

[3]BIS Research

Why Choose 5th Gen Intel® Xeon® processors for Server Refresh?

Lower Total Cost of Ownership (TCO)

Intel's portfolio of hardware, software, systems, and tools can help advance your data center's overall efficiency, creating energy savings and reducing your carbon footprint, without sacrificing performance, while giving you the TCO and flexibility you need.

Up to **77%** reduction in TCO[1]

Efficiency

Intel® Accelerator Engines boost CPU utilization, reduce electricity consumption resulting in lower impact on the environment.

10x better efficiency (perf/watt) with built-in accelerators[3]

Optimized workload performance

By delivering more performance per core with built-in accelerators, 5th Gen Intel® Xeon® processors help you meet requirements for even the most demanding workloads.

84% Performance gain[2]

Modernization

Intel® Xeon® processors deliver the low-latency, high-bandwidth capabilities required by modern and AI-infused workloads. Replacing aging infrastructure with these speedy and energy efficient processors will help you keep pace with rapidly evolving market needs.

Up to **16:1** server consolidation[1]

intel. XEON

Confidential Computing

With Intel, you can choose from the most deployed confidential computing options in data centers on the market today—now including application or VM-level isolation.

What's the right transition for your customer?

[1, 2, 3] See [T7. G1, T13]at intel.com/processorclaims:5th Gen Intel Xeon Scalable processors. Results may vary

Artificial Intelligence

ACCESS NOW

- Enterprise AI / Generative AI Partner Enablement Package
- AI Partner Enablement Package

AI Continuum

Bringing AI everywhere

In today's hypercompetitive environment, enterprises that embrace AI are pulling ahead.

Intel infrastructure is engineered for enterprise AI, empowering you to maximize your investments and realize your vision at a lower cost. And, with enterprise-ready solutions and open, optimized software, you can go to market fast, even with sensitive and regulated data.

It's time to think differently about enterprise AI.

> Bringing AI Everywhere Infographic

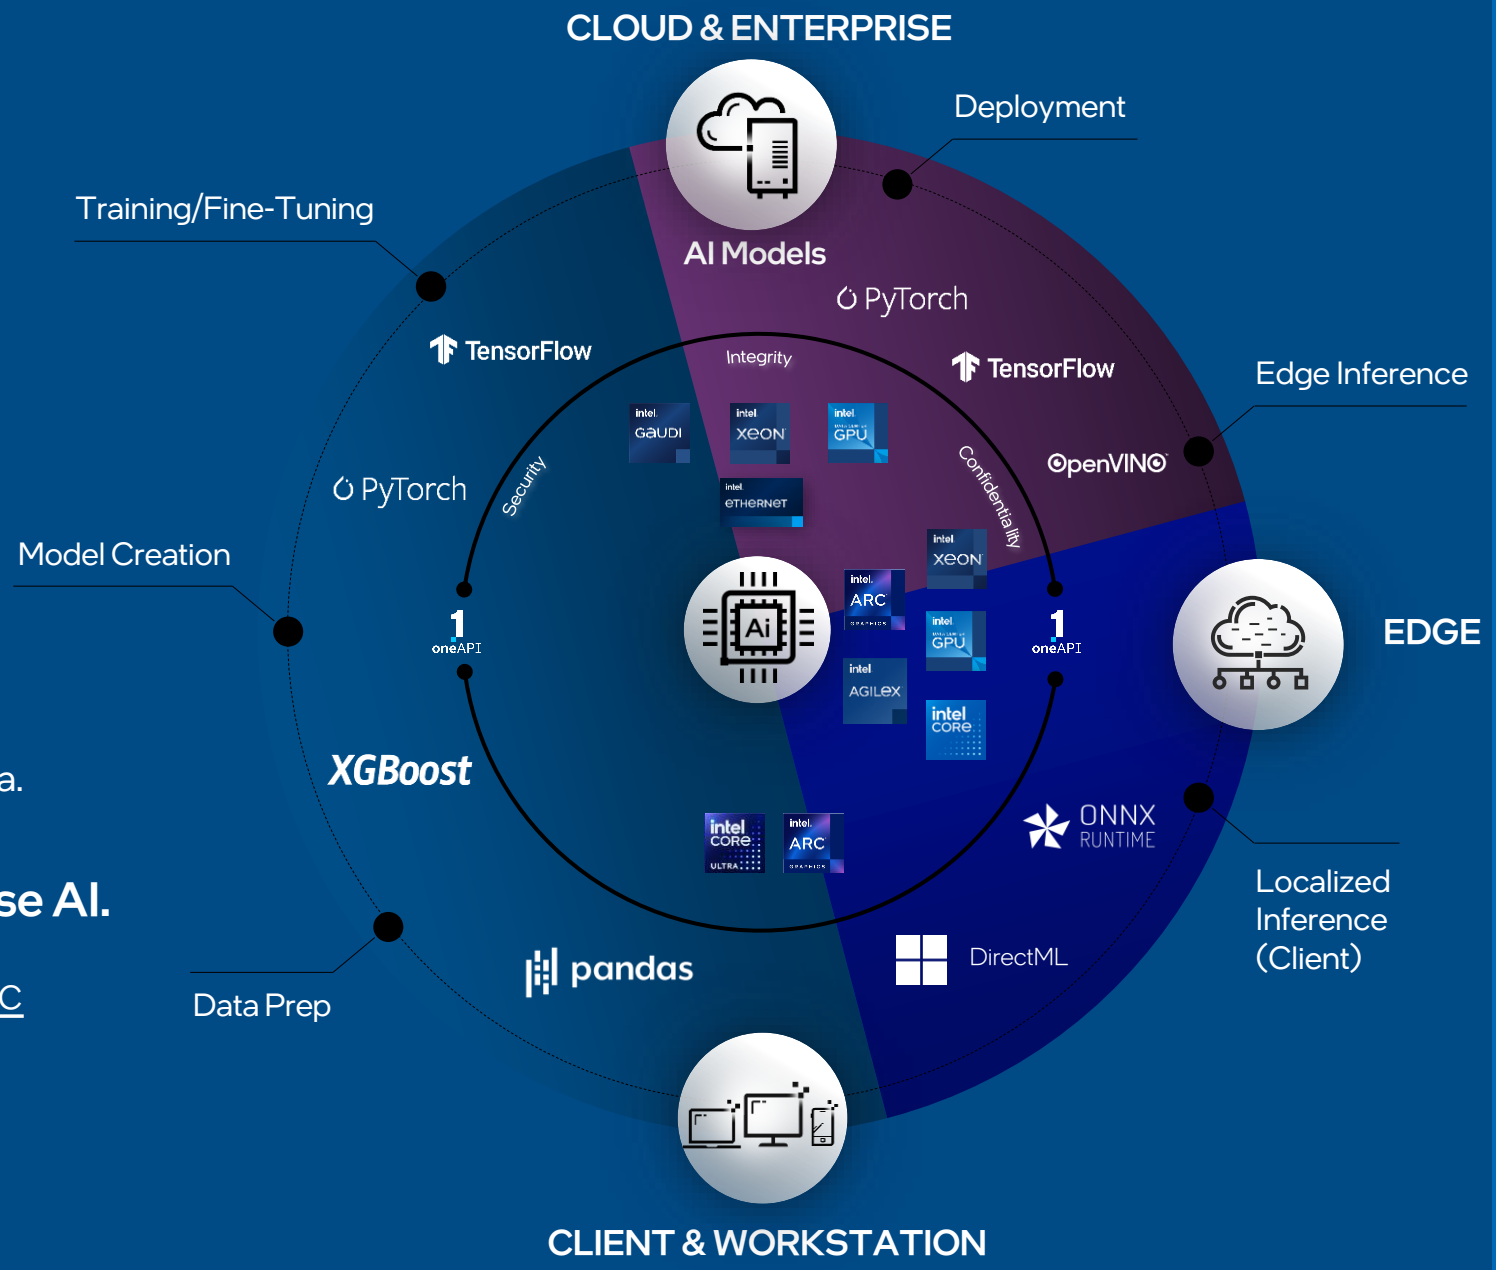

Note: Intel® Core™ Ultra processors integrate NPU low power inference engine from 15th Gen processors onwards.

© 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.

Intel AI Software Enables AI Everywhere Faster

Accelerating development with optimizations of the most popular industry libraries and open source tools, the Intel® AI software suite unleashes the performance of Intel® Xeon® Scalable processors without code changes

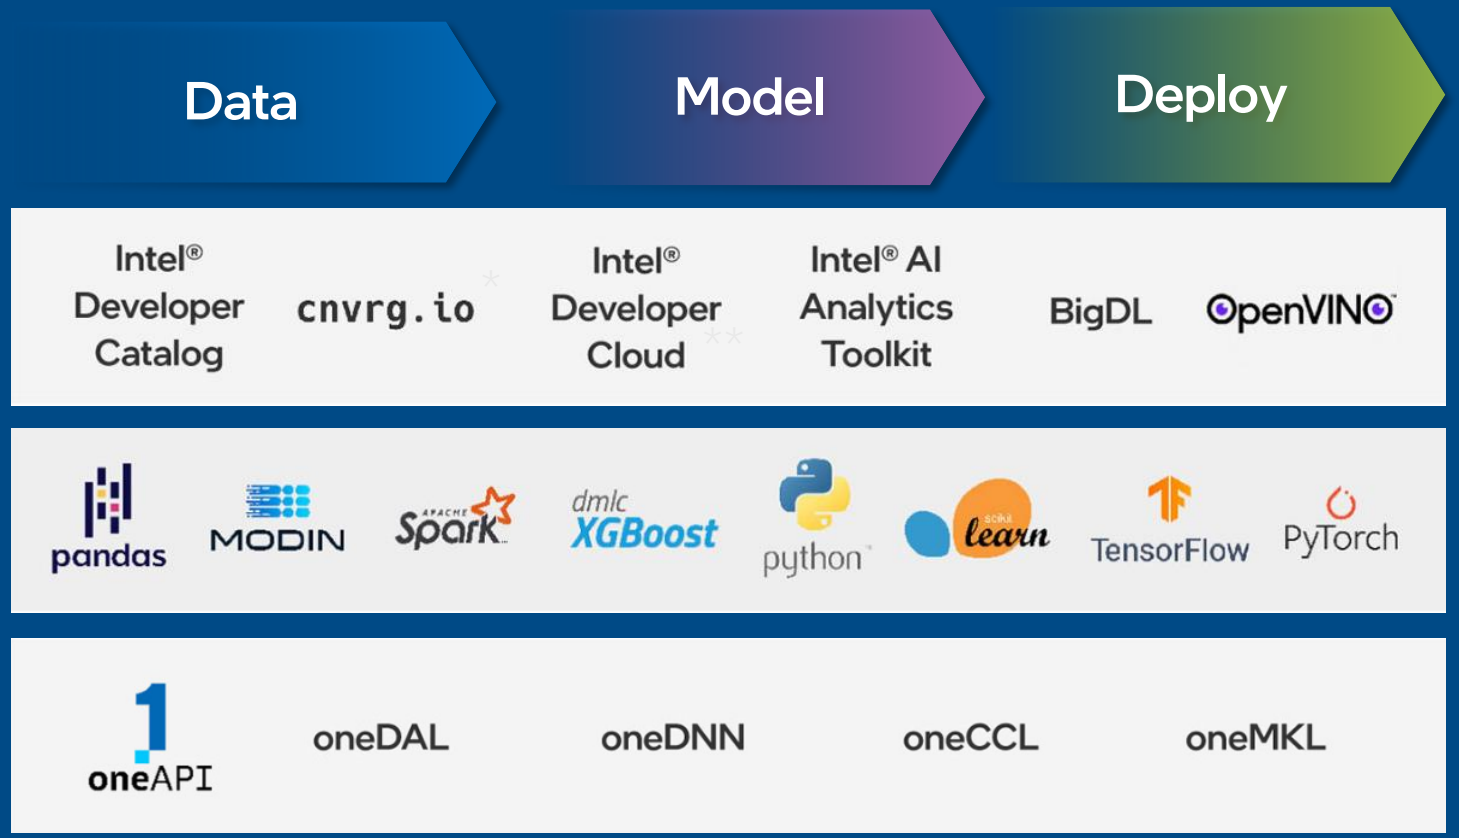

The Intel® AI software suite has been validated on over 400 AI models and use cases to help ensure that you achieve out-of-the-box application performance

* Now known as: Intel® Tiber™ AI Studio

** Now known as: Intel® Tiber™ Developer Cloud

Accelerate AI Development with Reference Kits

Optimized AI reference kits help developers and data scientists innovate faster

Built on the oneAPI open, standards-based, heterogeneous programming model and components of Intel's end-to-end AI software portfolio, such as Intel® AI Analytics Toolkit and the Intel® Distribution of OpenVINO™ toolkit, the reference kits enable AI developers to streamline the process of introducing AI into their applications, enhancing existing intelligent solutions and accelerating deployment.

The result is proven performance improvements with a shorter, more productive workflow versus a traditional model development workflow

Using the **AI reference kit** designed to set up interactions with an enterprise conversational AI chatbot, users can experience inferencing in batch mode **up to 45% faster** with oneAPI optimizations

The **AI reference kit** designed to automate visual quality control inspections for life sciences demonstrated training **up to 20% faster** and inferencing **55% faster** for visual defect detection with oneAPI optimizations

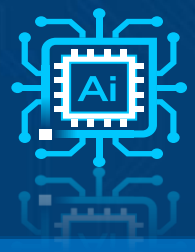

To enable developers to predict utility asset health and deliver higher service reliability, there is an **AI reference kit** that provides **up to a 25% increase** in prediction accuracy

Press Release

AI Reference Kits eBook

4th and 5th Gen Intel® Xeon® Scalable Processors with Accelerators for AI Inference

Accelerators like Intel® AVX-512 and Intel® AMX are designed to improve performance, reduce latency and increase memory bandwidth, making them well suited for running demanding Inference AI workloads

Built-in Accelerators and Why You Should Use Them

Intel® Advanced Matrix Extensions (Intel® AMX)

significantly accelerates deep learning training and inference, ideal for workloads like natural language processing, recommendation systems and image recognition

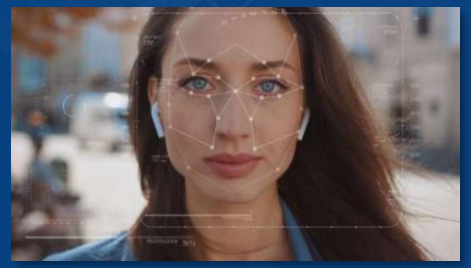

Website Solution Brief
Video User Guide

Intel® Advanced Vector Extensions 512 (Intel® AVX-512)

can accelerate classical machine learning and other workloads in the end-to-end AI workflow, such as data prep

Website Solution Brief
Video User Guide and Downloads

READ THE ARTICLE

Taboola Improves Content Recommendation Engines

Intel AI Optimizations Quick Start Guide

Case Studies

	Challenge	Solution / Results	Intel Products	More info
Tencent Cloud *Search engine for cloud compute service*	How to handle large-scale queries and respond promptly with the search results	Tencent can use the optimized BERT model to deliver better service experiences and to help reduce TCO	4th Gen Xeon® + Intel® AMX	Case Study
美团 Meituan *Leading retail technology company*	Cost effective vision AI services	Meituan increased the overall efficiency of its online resources by over 3x and saved 70% on service costs	4th Gen Xeon® + Intel® AMX + Intel® IPP + Intel® Extension for PyTorch (Intel® IPEX)	Case Study
SIEMENS *Medical Image Processing*	Improving efficiency of radiation therapy professionals	Supporting radiation therapy professionals with AI-based auto contouring technology increases workload efficiency, improve consistency, and help free up staff to focus on value adding work	4th Gen Xeon® + Intel® AMX + OpenVINO™	Case Study Video
Alibaba Cloud *Leading Cloud Computing Provider*	Improve performance of address-purification services	Faster end-to-end performance translates to better business results for Alibaba's customers in logistics, e-commerce, energy, retail, and finance. Using a built-in accelerator helps Alibaba control TCO	4th Gen Xeon® + Intel® AMX + Intel® oneDNN	Case Study

See case study links for workloads and configurations. Results may vary.

Testimonials on Intel's AI Technology

FOX SPORTS

"We've shaved weeks off of setup time"

"For us, Intel® Xeon® processors are a cornerstone of how we deploy technology. We run only on Intel® Xeon® CPUs, and that gives us the ability to run everywhere: in VMs, in dedicated on-premises bare metal, in the cloud."

Case Study

SIEMENS

35x speedup in AI inference time for auto contouring algorithms compared to previous gen[1]

20% reduction in energy consumption compared to previous gen[2]

Case Study Video

[12]See case study links for workloads and configurations. Results may vary.

Security

ACCESS NOW

- Confidential Computing Partner Enablement Package

Intel Offers the Most Comprehensive Security Portfolio for Confidential Computing

Confidential computing with trusted execution environments (TEEs) helps protect data and AI models

With 4th and 5th Gen Intel® Xeon® processors, you can choose from the most researched and updated confidential computing options in data centers on the market today

READ MORE

Intel® Software Guard Extensions (Intel® SGX)

Application isolation

Intel® Trust Domain Extensions (Intel® TDX)

Virtual machine isolation

*Intel® TDX available through select cloud providers

Intel® Tiber™ Trust Services formerly Intel® Trust Authority

Independent trust verification services for multi-cloud & hybrid cloud

PERFORMANCE PROOFPOINT

Up to **4x higher VPP IPSec (1420B) throughput**[1]

with the new 5th Gen Intel Xeon Platinum 8592+ processor compared to the 3rd Gen Intel Xeon Processor

Confidential Computing 1-pager

[1]See [N13] at intel.com/processorclaims: 5th Gen Intel Xeon processors. Results may vary

Intel Trusted Execution Environments

Application-level isolation: **Intel® SGX**

Advantages

- Separation from cloud provider and other tenants
- Smaller trust boundary and potential attack surface
- More amenable to code inspection and monitoring
- Deployable in VMs, cloud-native containers and bare-metal

Considerations

- Apps may require specific development or tailoring
- Frequent calls outside the enclave may impact performance

SOLUTION BRIEF

[Microsoft moves credit card transactions to Azure Cloud Services running Intel® Software Guard Extensions (Intel® SGX)]

Microsoft

VM-level isolation: **Intel® TDX**

Advantages

- Separation from cloud provider and other tenants
- Lowest porting effort for existing applications
- More amenable to enterprise-wide deployment mandates
- Can be a simple instance configurator setting

Considerations

- Larger trust boundary (guest OS, all apps, VM admins)
- Possible re-validation with updated guest OS & hypervisor
- Less granular attestation

INFOGRAPHIC

READ MORE

[Which Intel Trusted Execution Environment is right for you?]

Intel® Tiber™ Trust Services

formerly Intel® Trust Authority

Put Zero Trust Within Reach and Get Public Cloud Flexibility with Private Cloud Security

Intel® Tiber™ Trust Services is a new portfolio of software and services that brings enhanced security and assurance to Confidential Computing with Zero Trust principles

In its first generation, it offers an independent attestation service that attests to Trusted Execution Environments (TEEs) that are based on (Intel® SGX) and (Intel® TDX)

Implement the tenets of Zero Trust without incurring the cost and complexity of building your own attestation service

Independent

Scalable

Easy to Deploy

LEARN MORE

Product Brief

Video

CASE STUDIES

click on logos for more info

noname zscaler THALES

Intel® TDX Availability

Intel® TDX is available on 4th and 5th Gen Intel® Xeon® Scalable instances in public preview through three leading cloud providers

Click on the logos below for more information on each cloud provider's offering

Microsoft Azure

Alibaba Cloud

Google Cloud

IBM Cloud
*public preview TBA

Intel® TDX is enabled on the following guest OS vendors

Red Hat

SUSE

Canonical Ubuntu

5th gen Intel® Xeon®

WHITE PAPER > Alibaba Cloud ApsaraDB Confidential Database Empowered By Intel® TDX

Partner Benchmarking

ACCESS NOW

- VMware Modernization Partner Enablement Package

VMware vSphere 8.0 on 4th and 5th Gen Intel® Xeon® Scalable Processors

VMware vSphere 8

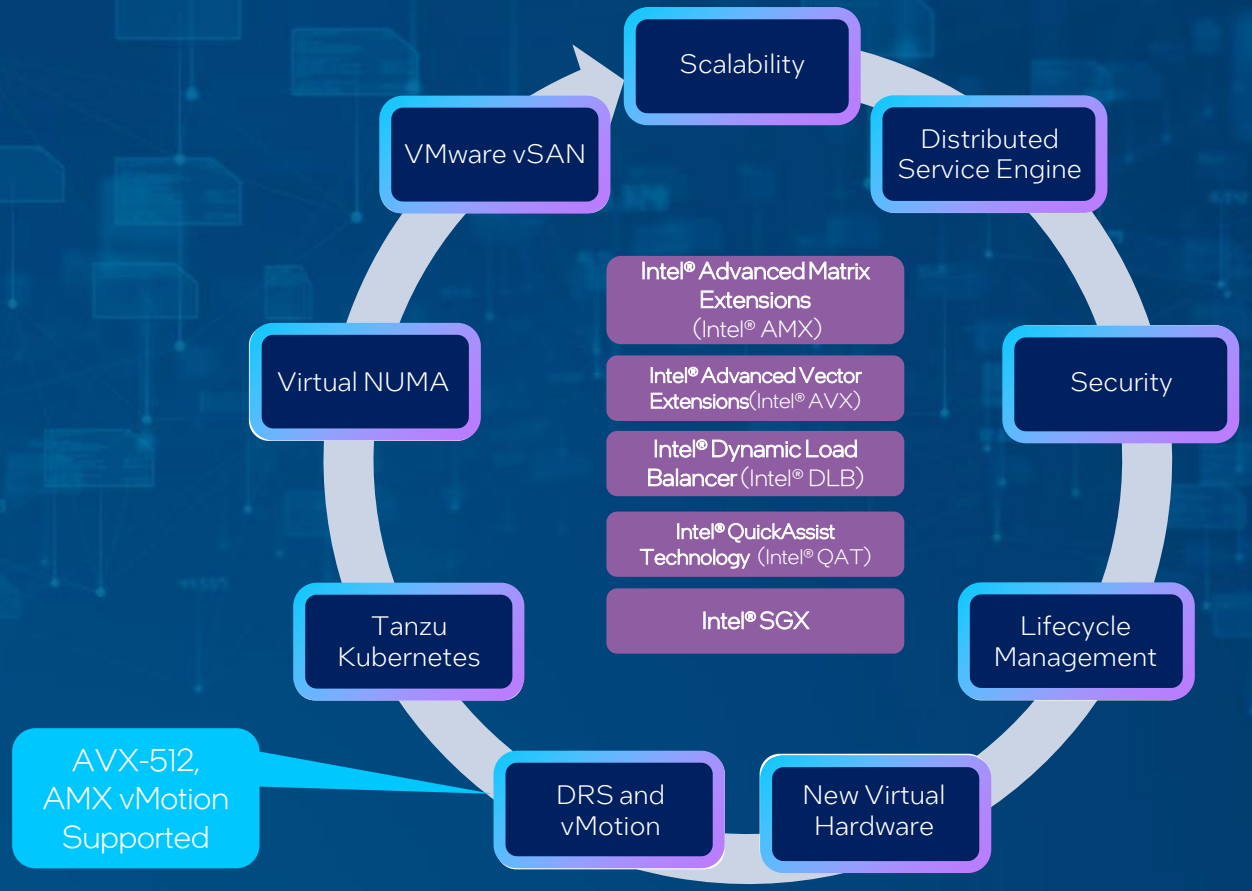

https://docs.vmware.com/en/VMware-vSphere/8.0/rn/vmware-vsphere-80-release-notes/index.html

Unlocking the Value of Accelerators with Software

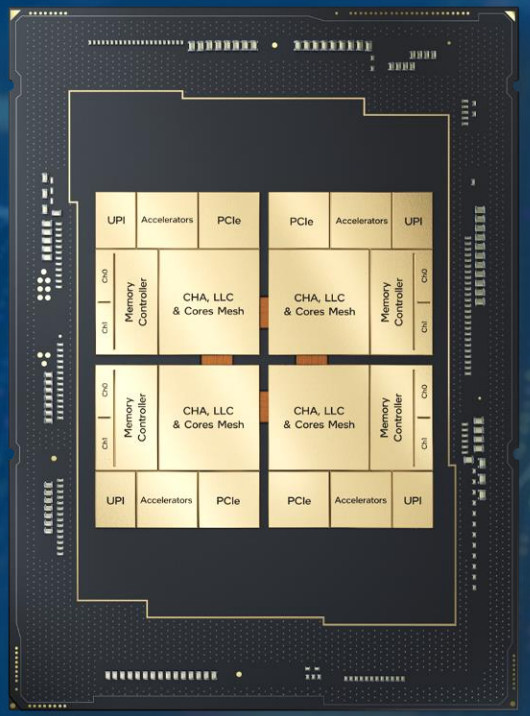

Intel® Advanced Matrix Extensions (Intel® AMX)

- TensorFlow
- PyTorch
- ONNX Runtime
- OpenVINO
- oneDNN (Intel oneAPI)

Intel® Advanced Vector Extensions (Intel® AVX) for vRAN

- FlexRAN
- Data Plane dev Kit (DPDK)*

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

- Intel Query Processing Library

Intel® Data Streaming Accelerator (Intel® DSA)

- Storage Perf Dev Kit (SPDK)*
- Data Plane Dev Kit (DPDK)*

Intel® QuickAssist Technology (Intel® QAT)

- QATzip* (Intel lib)
- OpenSSL**
- Boring SSL

Intel® Dynamic Load Balancer (Intel® DLB)

- VPP IPsec
- Data Plane Dev Kit (DPDK)*

*Intel open-source library (not part of stock SW).
**Difference between Intel version and stock version.
***Intel® QPL and *Intel® DML in open-source beta, v1.0.0 coming shortly.

Benefits of vSphere Foundation on 4th Gen Intel® Xeon®

BENCHMARKS

Up to
5x Faster[1]
and Still Accurate Image Classification
Using Intel® AMX for BF16 compared to Intel® AVX-512 for FP32

VMware vSphere/vSAN8 on 4th Gen Xeon® with Intel® AMX for **Image Classification**

[1]Solution Brief

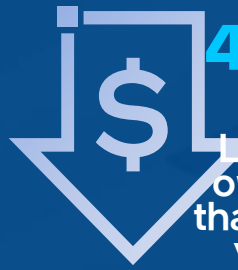

45% Lower TCO[2]
Lower total cost of ownership by more than 45% when using vSAN vs without

Boost Performance and Lower Latency with VMware vSAN8 and 4th Gen Xeon®

[2]Solution Brief

Up to
5.7x Higher[3]
Natural Language Processing
INT8 with Intel® AMX vs FP32 with Intel® AVX-512

vSphere/vSAN with Intel® AMX for **Natural Language Processing**

[3]Solution Snapshot

Up to
6.2x Better Performance[4]
With servers featuring 1st Gen Xeon® vs 4th Gen Xeon®

VMware vSAN on 4th Gen Xeon® for **Modernization**

[4]Solution Snapshot

READ MORE >

VMware vSAN and 4th Gen Intel® Xeon® Processors Outpace AMD Genoa

See relevant Solution Briefs and Snapshots (linked above) for configuration details. Results may vary.

4th Gen Intel® Xeon® Scalable Processor Support Red Hat Q3'23

Feature	RHEL – BareMetal	RHEL – Virtualized	OpenShift (K8S)	Accelerator Getting Started Guide
4th Gen Xeon®	8.6, 9.0	8.6, 9.0	4.11	N/A
DSA (Data Streaming Accelerator)	8.6 / 9.0	TBD	4.13 (Q4'23/Q1'24)	DSA Guide
IAA (In-memory Analytics Accel.)	8.6 / 9.0	TBD	4.13 (Q4'23/Q1'24)	IAA Guide
QAT (Quick Assist Technology)	8.6 & 9.0	TBD - OOT* until Q2'24	4.12	QAT Guide
AMX (Adv. Matrix eXtensions)	8.6 & 9.0	8.6, 9.0	4.11	AMX Guide
AVX (Adv. Vector eXtensions)	8.6, 9.0	8.6, 9.0	4.11	Not Available
DLB (Dynamic Load Balancer)	OOT* until Q2'24	OOT* until Q2'24	TBD	TBD
SGX (SW Guard eXtensions)	8.6, 9.0	8.6, 9.0	4.11	SGX Guide
TDX (Trust Domain eXtensions)	TBD	8.8, 9.2 (VM Guest & TBD on Host)	TBD	TDX Guides
SIOV (Scalable I/O Virtualization)	9.2 (target)	-	-	Not Available
SST (Speed Select Technology)	8.6 / 9.0	N/A	Power Operator (Q1'23)	SST Guide
Intel On-Demand	8.7 / 9.1	N/A	N/A	N/A

* OOT = Out-of-Tree

Accelerators monitoring tool: **pcm-accel** (https://github.com/intel/pcm)

Boosting AI Performance with Red Hat® OpenShift® 4.12 on 4th Gen Intel® Xeon® Scalable Processors

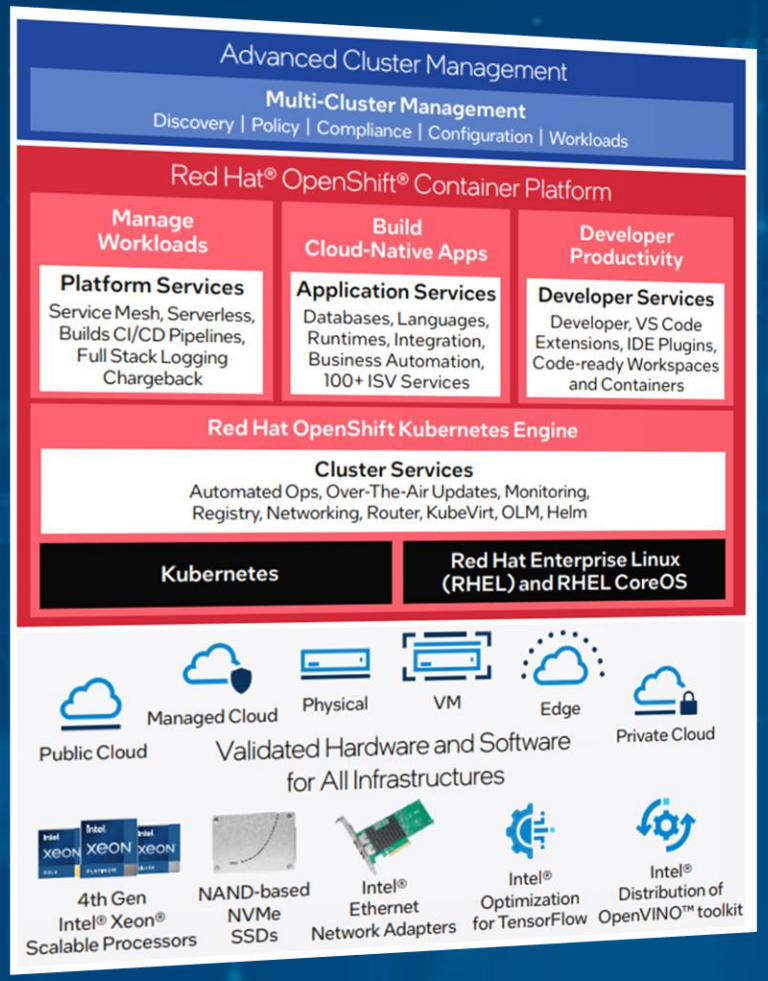

Natural Language Processing:

Smoother Experiences with Faster Responses

Up To **5.7x** higher
End-to-End Real-Time Inference Performance Speedup[2]

Up To **6.2x** higher
Real-Time NLP Inference Performance[3]

Recommendation Systems:

Recommendations in Real Time

Up To **6.3x** higher
Batch Recommendation System Inference Performance[4]

Up To **4x** higher
Recommendation System Training Performance[4]

UPDATE:
5th Gen Xeon®

- Accelerate Red Hat OpenShift AI Workflows using 5th Gen Xeon® Features with Red Hat Validated Patterns

[2,3,4] https://www.intel.com/content/www/us/en/partner/showcase/redhat/openshift-container-4-12-boost-ai-ref-arch.html

Call to Action

Education

Understand the value of modernizing your customers' data centers on 4th and 5th Gen Intel® Xeon® Scalable processors and how it will reduce operating costs and increase your AI & Security capabilities over older infrastructure

Engagement

Connect with your Intel representative to understand how to leverage Intel's technology portfolio to modernize your customers' data centers

Cloud TV

Intel® Cloud TV explores cloud computing news, trends, and strategies to drive your success

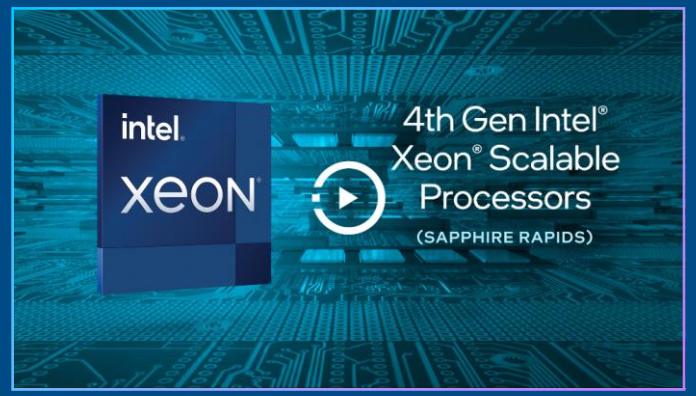

Sapphire Rapids in the Cloud

Sustainability and the Cloud

Modernizing the Hybrid Data Center

5th Gen Intel® Xeon® Scalable Processors Overview

Modernization Opportunities with Microsoft

Intel® Xeon® Processor Advisor Tool Suite

New: Updated advisors for 4th Generation Intel® Xeon®
Scalable processors are now available!

Sign Up!

Then choose your deployment environment to begin

On-Prem Advisors

Find the best solutions for your workload whether it's refreshing existing infrastructure or building a new data center. Get instant recommendations and optimize based on TCO, Sustainability and Power.

Cloud Advisors

Discover the best solutions for refreshing or migrating your workload to the cloud. Get instant recommendations optimized for performance and TCO. Find Intel based CSP Instances and pricing.

Modernization Information and Resources

Public Content

Asset Type	Title and Link
Infographic	Intel Sustainability Use Case - Server Consolidation
Sales Brief	Intel Sustainability Use Case - AI
Solution Brief	Advance Your Energy Initiatives
Whitepaper	Worker Experiences Redefined with 4th Gen Intel® Xeon® Scalable Processors and New Accelerators Innovate Faster with Integrated AI
Video	Sustainability with Intel technologies
Case Study	Gunpowder Cuts Digital Rendering Time and Cost on New Google Cloud Instances
Performance Index	4th Generation Intel® Xeon® Scalable Processors
Live Webinar	Cloud Solution Architect (CSA) Tech Talk: Reduce TCO and Improve Efficiency with 4th Gen Intel® Xeon® Scalable Processors
Recorded Webinar	Cloud Solution Architect (CSA) Tech Talk: Building Sustainability Practices in the Data Center and Cloud
Recorded Webinar	Cloud Solution Architect (CSA) Tech Talk: Accelerating Critical Workloads with 4th Generation Intel® Xeon® Scalable Processors
Intel® Optimization Hub	Optimizations as Code
Training	In-deck links to Online Tutorials

Microsoft Data Center Products Refresh Information and Resources

Public Content

Asset Type	Title and Link
Microsoft SQL Server 2022	
Tuning Guide	Tuning SQL Server for OLTP
Tuning Guide	Tuning SQL Server for OLAP
Solution Brief	Optimizing Microsoft SQL Server 2022 on Lenovo ThinkSystem SR650 V3
Solution Snapshot	Microsoft SQL Server 2022 on 4th Gen Intel® Xeon® Scalable Processors
Solution Design Brief	Microsoft SQL Server 2022 on Intel® Technologies
Whitepaper	Intel QAT Performance on 4th Gen Intel® Xeon® Processors
Microsoft Azure Stack HCI	
Whitepaper	Unify Operations Across Hybrid and Multi-Cloud Environments
Solution Design Brief	Microsoft Azure Stack HCI on 4th Gen Intel® Xeon® Scalable Processors
Article	MSFT Azure HCI & Arc Wall Street Journal - The Path to Greener IT in a Hybrid Cloud World
Article	MSFT Azure HCI & Arc Wall Street Journal - Driving Sustainability for IT Infrastructure
Case Study	Franz Morat Group Gears Up for the Future
Security Assets	Infographic | White paper | Video animation
Windows Server 2022	
Report	Deploying Windows Server 2022 on Dell PowerEdge Servers

Cloud TV - Public

Modernization Opportunities with Microsoft

Learn how to capitalize on two critical modernization opportunities for you and your customers with the end of support for Microsoft Windows Server and SQL Server 2012

How to Access Intel® Partner Alliance Customer Support

Intel Virtual Assistant

This Chat Bot, located in the bottom-right corner of each Partner Alliance webpage, provides self-help to most questions or a quick link to a live support agent.

Get Help "Blade"

Submit an online support request.

This link is found on the footer of most pages within the Partner Alliance website.

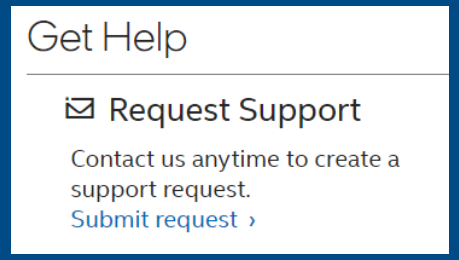

Partner Alliance "Get Help" page

The Get Help page provides detailed self-help guides on most of the tools and benefits available to Partner Alliance members.

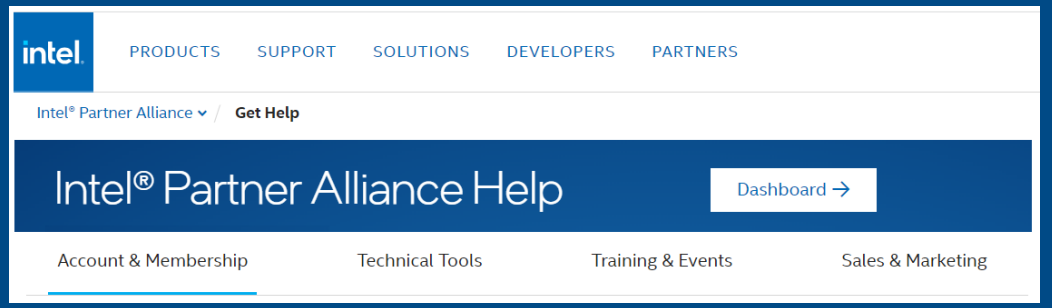

Training

Topic -- Audience
Data Center Sustainability with Intel Data Center Manager DevOps / Cloud Architects
Corporate Sustainability: A Blueprint for Reducing Carbon Emissions C-Suite
One Intel: Introduction to Intel Sustainability Initiative ALL
How to Reduce Data Center Power Cost with Sustainability Regulations C-Suite / Procurement
Intel's Contribution to Cloud Native DevOps / Cloud Architects
Application Architecture and Development in the Cloud DevOps

Topic -- Audience
AWS: Intel Instances and Affinity Cloud Architects
Azure: Intel Instances and Affinity Cloud Architects
Google Cloud Platform: Intel Instances and Affinity Cloud Architects
Workload Placement Cloud Architects


intel

Backup

What's the right transition for your customer?

Guide your customer to the best refresh option

Current install base	Good	Better	Best
Intel Xeon Silver 4000 Series	4514Y 16C/150W	4516+ 24C/185W	5520+ 28C/205W
Intel Xeon Gold 5000 Series	5515Y 8C/185W	5520+ 28C/205W	6530 32C/270W
Intel Xeon Gold 6000 Series	6526Y 16C/195W	6548Y 32C/250W	8558 48C/330W
Intel Xeon Platinum 8000 Series	8562Y+ 32C/300W	8568Y+ 48C/350W	8592+ 64C/350W

Refresh your infrastructure today, to be ready for tomorrow's demands

Refresh from 1st Gen Intel® Xeon® to 5th Gen Intel Xeon

AI
(BertLarge)

up to **42x** higher performance[1]

up to **20x** higher performance/ watt[1]

Refresh from 2nd Gen Intel® Xeon® to 5th Gen Intel Xeon

AI
(DLRM)

up to **10x** higher performance[2]

up to **6x** higher performance/ watt[2]

Intel vs. AMD

AI
(Recommender - DLRM)

up to **2.8x** higher batched inference performance[3]

up to **2.6x** higher performance/ watt[3]

Exceptional Performance & Efficiency

5th Gen Intel® Xeon® processors deliver impressive performance-per-watt gains across all workloads, plus outsized performance and lower TCO

- Optimize AI, HPC, network, data analytics & storage workloads with Intel® Accelerator Engines
- Enhanced platform capabilities
  - 3x increase in shared last-level cache
  - PCIe 5 – Double I/O bandwidth
  - DDR5 – Increase memory bandwidth
  - CXL® 1.1 – Next Gen I/O for low latency and performance
  - Optimized Power Mode enables energy savings with minimal impact of performance*
  - Advanced security technologies to help protect data with Intel Software Guard Extensions (Intel SGX) and Intel® Trust Domain Extensions (Intel TDX)
- To learn more about SKU transitions and refresh opportunities:
  - https://xeonprocessoradvisor.intel.com

TCO Advisor Tool

[1, 2, 3] See [A37, A38, A208] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary
*Enabled in platform BIOS, visit intel.com/processorclaims: 5th Gen Intel Scalable processors for more information

What's the right transition for your customer?

Guide your customer to the best refresh option

Current install base	Good	Better	Best
Intel Xeon Silver 4000 Series	4514Y 16C/150W	4516+ 24C/185W	5520+ 28C/205W
Intel Xeon Gold 5000 Series	5515Y 8C/185W	5520+ 28C/205W	6530 32C/270W
Intel Xeon Gold 6000 Series	6526Y 16C/195W	6548Y 32C/250W	8558 48C/330W
Intel Xeon Platinum 8000 Series	8562Y+ 32C/300W	8568Y+ 48C/350W	8592+ 64C/350W

Intel® Xeon®, the processor designed for AI

Total Cost of Ownership Savings
3rd Gen Intel® Xeon® to 5th Gen Intel Xeon

AI
(Recommender -DLRM)

up to **5:1**
server consolidation[1]

up to **72%**
TCO savings[1]

Performance & Efficiency Gains
3rd Gen Intel® Xeon® to 5th Gen Intel Xeon

AI
(Real Time Inference)

up to **14x**
higher performance[2]

up to **9.5x**
higher performance/ watt[2]

Intel vs. AMD

AI
(Recommender - DLRM)

up to **2.8x**
higher batched inference performance[3]

up to **2.6x**
higher performance/ watt[3]

Exceptional Performance & Efficiency

5th Gen Intel® Xeon® processors deliver impressive performance-per-watt gains across all workloads, plus outsized performance and lower TCO

- Optimize AI, HPC, network, data analytics & storage workloads with Intel® Accelerator Engines
- Enhanced platform capabilities
  - 3x increase in shared last-level cache
  - PCIe 5 – Double I/O bandwidth
  - DDR5 – Increase memory bandwidth
  - CXL® 1.1 – Next Gen I/O for low latency and performance
  - Optimized Power Mode* enables energy savings with minimal impact of performance
  - Advanced security technologies to help protect data with Intel Software Guard Extensions (Intel SGX) and Intel® Trust Domain Extensions (Intel TDX)
- To learn more about SKU transitions and refresh opportunities:
  - https://xeonprocessoradvisor.intel.com

TCO Advisor Tool

[1, 2, 3] See [T12, A16, A208]at intel.com/processorclaims:5th Gen Intel Xeon Scalable processors. Results may vary
*Enabled in platform BIOS

3rd Gen Intel® Xeon® Scalable Processor with SQL Server 2019 vs 4th Gen Intel ® Xeon® Scalable Processor with SQL Server 2022 on Standard Edition

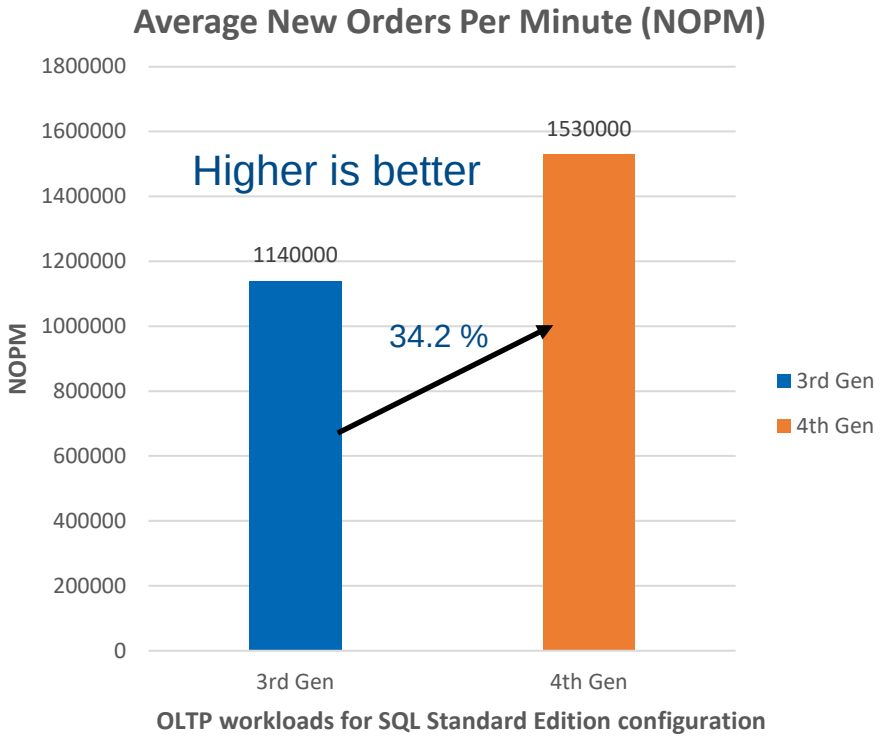

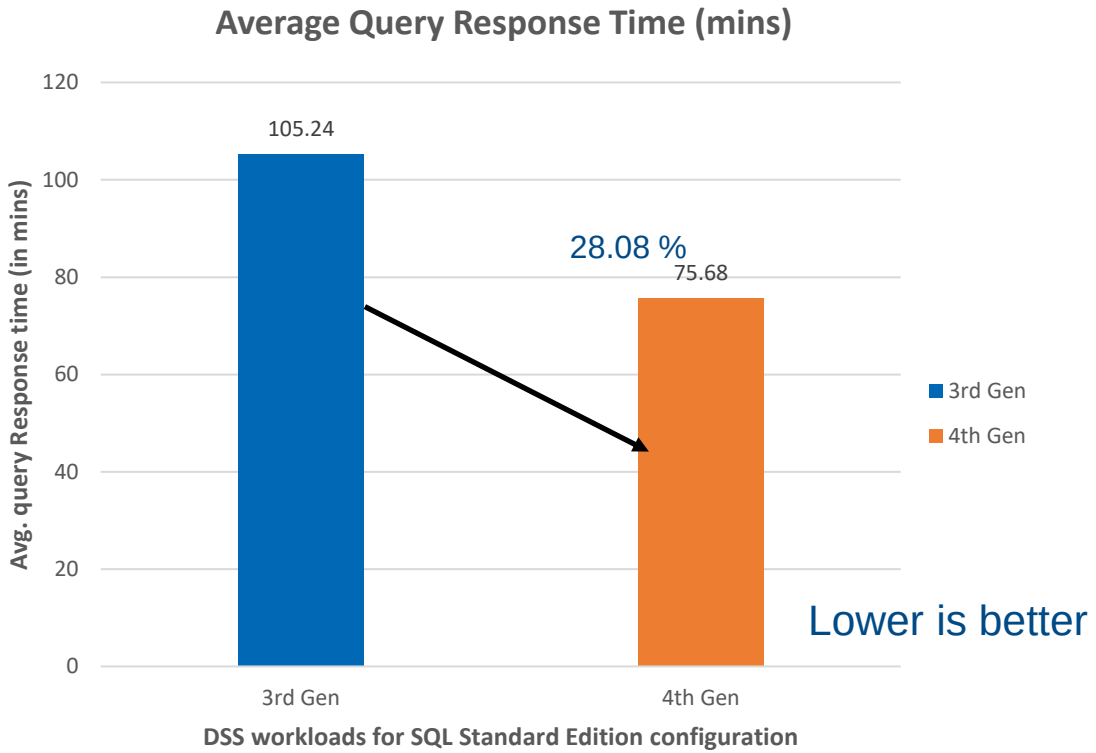

Up to 34% more NOPM transactions with 4th Gen Xeon processors over 3rd Gen Xeon processors

Up to 28 % faster query response time with 4th Gen Xeon processors over 3rd Gen Xeon processors

Tested by Intel as of 04/12/2023. 1 Node, 2x Intel® Xeon® Gold 6444Y+ (12C, 3.6GHz, 225W) CPU, 1x Quanta SDP QuantaGrid D54Q-2U, Total Memory: 256GB (16 x 16 GB 4800MHz DDR5 DIMM), Intel® Hyper-Threading Technology: Enabled, Turbo: Enabled, Storage (boot): 1 x Solidigm DC S4610 , 960 GB , Storage (Data drive): 6x Solidigm® SATA S4500 Series (3.84TB) , Storage (Log drive): 2 x Intel® SSD D7-P5510 3.84TB (NVMe), , Network devices: 1 x 25 GbE Intel(R) Ethernet Network Adapter E810-C-Q2, Network speed: 25 GbE, 1 x 10 GbE Intel(R) Ethernet Converged Network Adapter X550-T2, Network Speed: 1 GbE, OS/Software: Windows 2022 standard Edition with SQL Server 2022 Standard Edition (RTM) – 16.0.1000.6 (x64), HammerDB v4.0

Tested by Intel as of 03/19/2021. 1 Node, 2x Intel® Xeon® Silver 4310 (12C, 2.1GHz, 120W) CPU, 1x Intel® Server Board M50CYP , Total Memory: 256GB (16 x 16 GB 3200MHz DDR4 DIMM), Intel® Hyper-Threading Technology: Enabled, Turbo: Enabled, Storage (boot): 1 x Intel D3-S4510 , 240 GB , Storage (Data drive): 6x Intel® SATA S4610 Series (960GB) , Storage (Log drive): 2 x Intel® SSD DC P4610 1.6TB (NVMe) , Network devices: 1 x 10 GbE Intel(R) Ethernet Converged Network Adapter X550-T2, Network Speed: 1 GbE, OS/Software: Windows 2019 Data Center Edition with Microsoft SQL Server 2019 Std Edition (RTM-CU10) (KB5001090) - 15.0.4123.1 (X64), HammerDB v4.0
Results may vary.

3rd Gen Intel® Xeon® Scalable Processor vs 4th Gen Intel® Xeon® Scalable Processor with SQL Server 2022 on Plus Enterprise Configuration

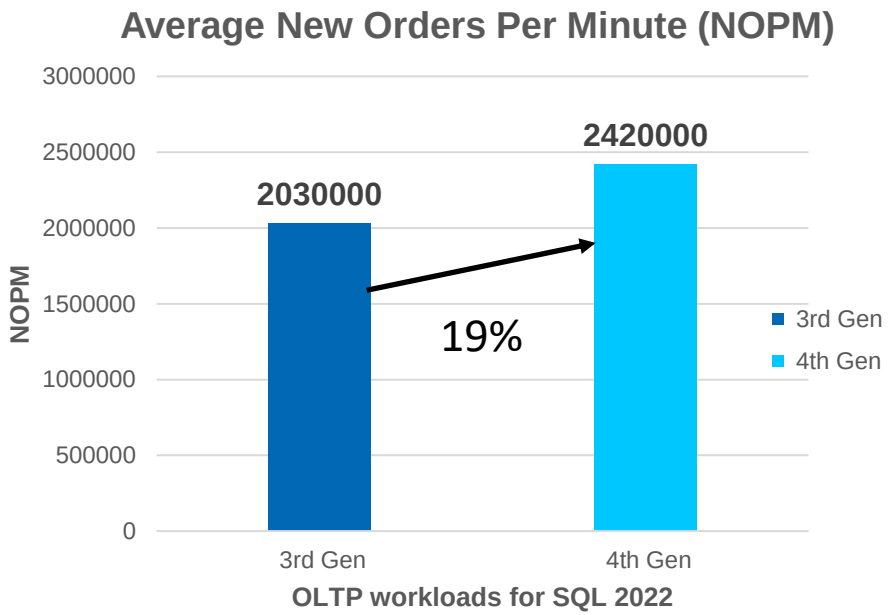

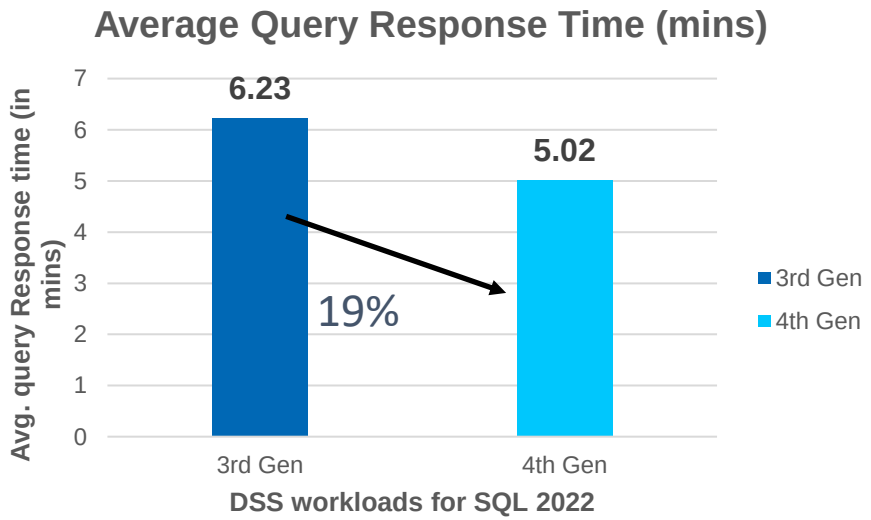

Up to 19% more NOPM transactions with 4th Gen Xeon® processors over 3rd Gen Xeon® processors

Up to 19% faster query response time with 4th Gen Xeon® processors over 3rd Gen Xeon® processors

Tested by Intel as of 03/07/2023. 1 Node, 2x Intel® Xeon® Gold 8460Y+ (32C, 2.3GHz, 300W) CPU, 1x Quanta SDP QuantaGrid D54Q-2U , Total Memory: 512GB (16 x 32 GB 4800MHz DDR5 DIMM), Intel® Hyper-Threading Technology: Enabled, Turbo: Enabled, Storage (boot): 1 x Solidigm DC S4610 , 960 GB , Storage (Data drive): 6x Solidigm® D7 P5510 Series (3.84TB) (NVMe) , Storage (Log drive): 2 x Intel® SSD DC P5800X 400GB (NVMe), , Network devices: 1 x 25 GbE Intel(R) Ethernet Network Adapter E810-C-Q2, Network speed: 25 GbE, 1 x 10 GbE Intel(R) Ethernet Converged Network Adapter X550-T2, Network Speed: 1 GbE, OS/Software: Windows 2022 standard Edition with SQL Server 2022 Enterprise Edition (RTM) – 16.0.1000.6 (x64), HammerDB v4.5

Tested by Intel as of 03/07/2023. 1 Node, 2x Intel® Xeon® Gold 6348 (28C, 2.6GHz, 235W) CPU, 1x M50CYP , Total Memory: 512GB (16 x 32 GB 2933MHz DDR5 DIMM), Intel® Hyper-Threading Technology: Enabled, Turbo: Enabled, Storage (boot): 1 x Solidigm DC P4101 , 512 GB , Storage (Data drive): 6x Solidigm® D7 P5510 Series (3.84TB) (NVMe) , Storage (Log drive): 2 x Intel® SSD DC P5800X 400GB (NVMe), Network devices: 1 x 25 GbE Intel(R) Ethernet Network Adapter E810-C-Q2, Network speed: 25 GbE, 1 x 10 GbE Intel(R) Ethernet Converged Network Adapter X550-T2, Network Speed: 1 GbE, OS/Software: Windows 2022 standard Edition with SQL Server 2022 Enterprise Edition (RTM) – 16.0.1000.6 (x64), HammerDB v4.5

Results may vary.

3rd Gen Intel® Xeon® Scalable Processor with SQL server 2019 vs 4th Gen Intel® Xeon® Scalable Processor with SQL Server 2022 on Base Enterprise Edition

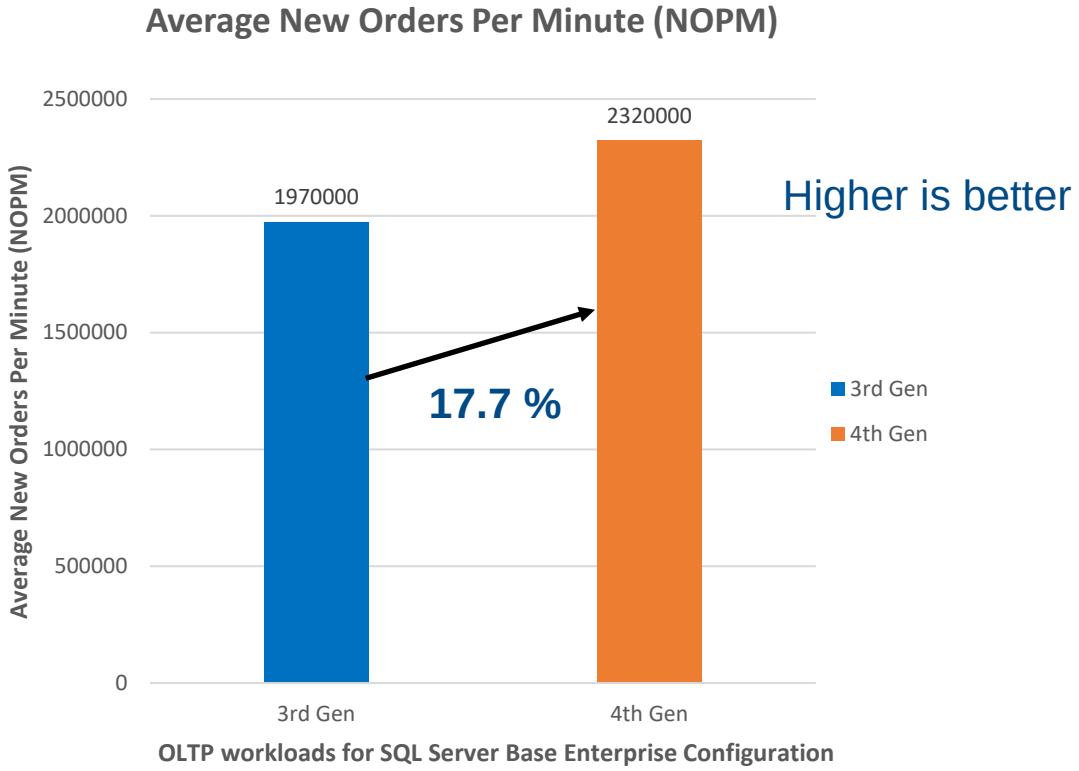

Up to 17.7% more NOPM transactions with 4th Gen Xeon processors over 3rd Gen Xeon processors

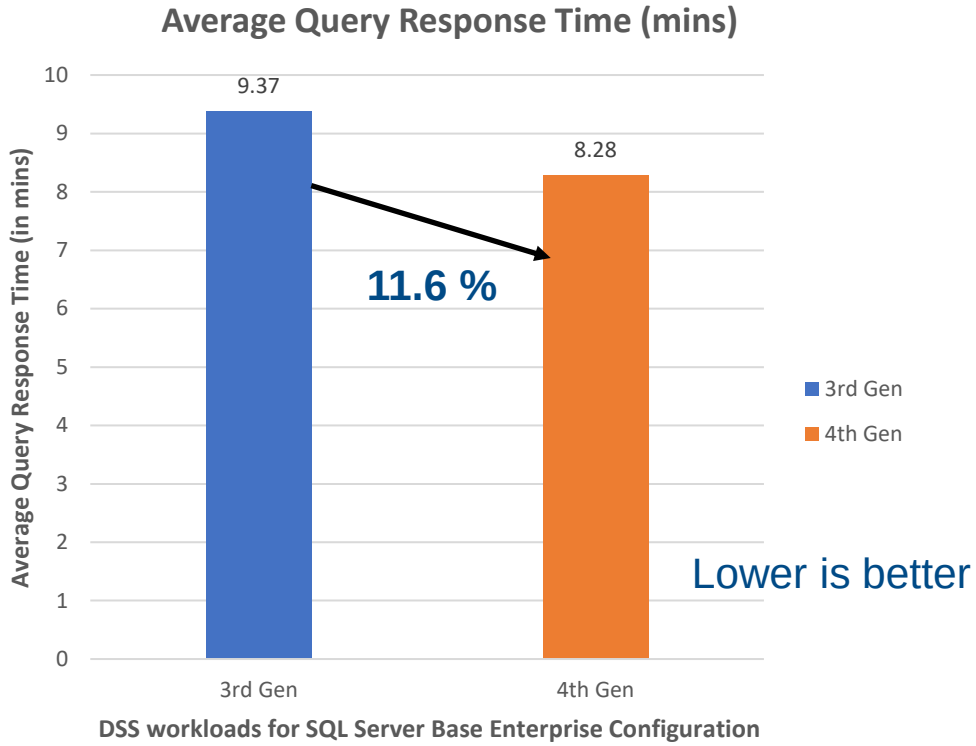

Up to 11.6% faster query response time with 4th Gen Xeon processors over 3rd Gen Xeon processors

Tested by Intel as of 04/27/2023. 1 Node, 2x Intel® Xeon® Gold 6438Y+ (24C, 2.1GHz, 185W) CPU, 1x Quanta SDP QuantaGrid D54Q-2U , Total Memory: 256GB (16 x 16 GB 4800MHz DDR5 DIMM), Intel® Hyper-Threading Technology: Enabled, Turbo: Enabled, Storage (boot): 1 x Solidigm DC S4610 , 960 GB , Storage (Data drive): 6x Solidigm® SSD D7-P5510 3.84TB (NVMe) , Storage (Log drive): 2 x Solidigm® SSD D7-P5510 3.84TB (NVMe), Network devices: 1 x 25 GbE Intel(R) Ethernet Network Adapter E810-C-Q2, Network speed: 25 GbE, 1 x 10 GbE Intel(R) Ethernet Converged Network Adapter X550-T2, Network Speed: 1 GbE, OS/Software: Windows 2022 Standard Edition with SQL Server 2022 Enterprise Edition (RTM) – 16.0.1000.6 (x64), HammerDB v4.0

Tested by Intel as of 05/30/2023. 1 Node, 2x Intel® Xeon® Gold 5318S (24C, 2.1GHz, 165W) CPU, 1x Intel® Server Board M50CYP , Total Memory: 256GB (16 x 16 GB 2933MHz DDR4 DIMM), Intel® Hyper-Threading Technology: Enabled, Turbo: Enabled, Storage (boot): 1 x Solidigm DC P4101 512GB SSD, Storage (Data drive): 6x Intel® P4510 Series (2 TB) , Storage (Log drive): 2 x Intel® SSD DC P4610 1.6TB (NVMe) , Network devices: 1 x 10 GbE Intel(R) Ethernet Converged Network Adapter X550-T2, Network Speed: 1 GbE, OS/Software: Windows Server 2022 Standard Edition with Microsoft SQL Server 2019 Enterprise Edition (RTM-CU20) (KB5024276) - 15.0.4312.2 (X64), HammerDB v4.0
Results may vary.

SQL Server 2022 with QAT Performance Comparisons

Intel® 4th Generation Xeon® Scalable processor with QAT vs Intel® 4th Generation Xeon® Scalable processor without QAT enabled

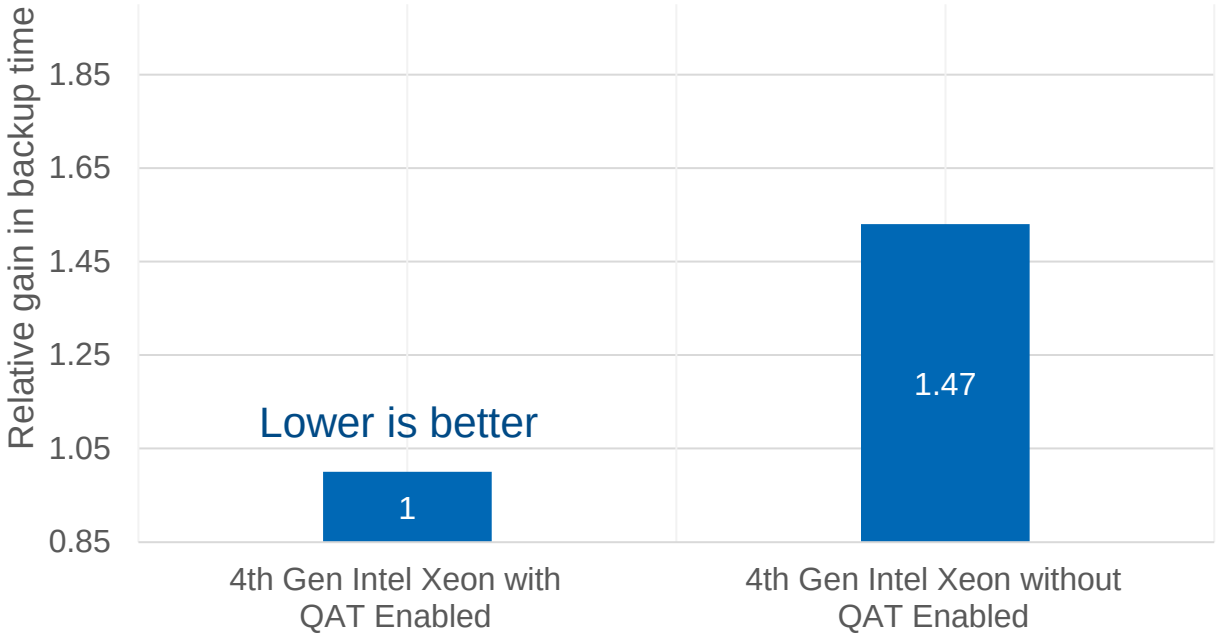

Up to 47% faster backup with 4th Gen Intel® Xeon® Processor and QAT enabled in idle state

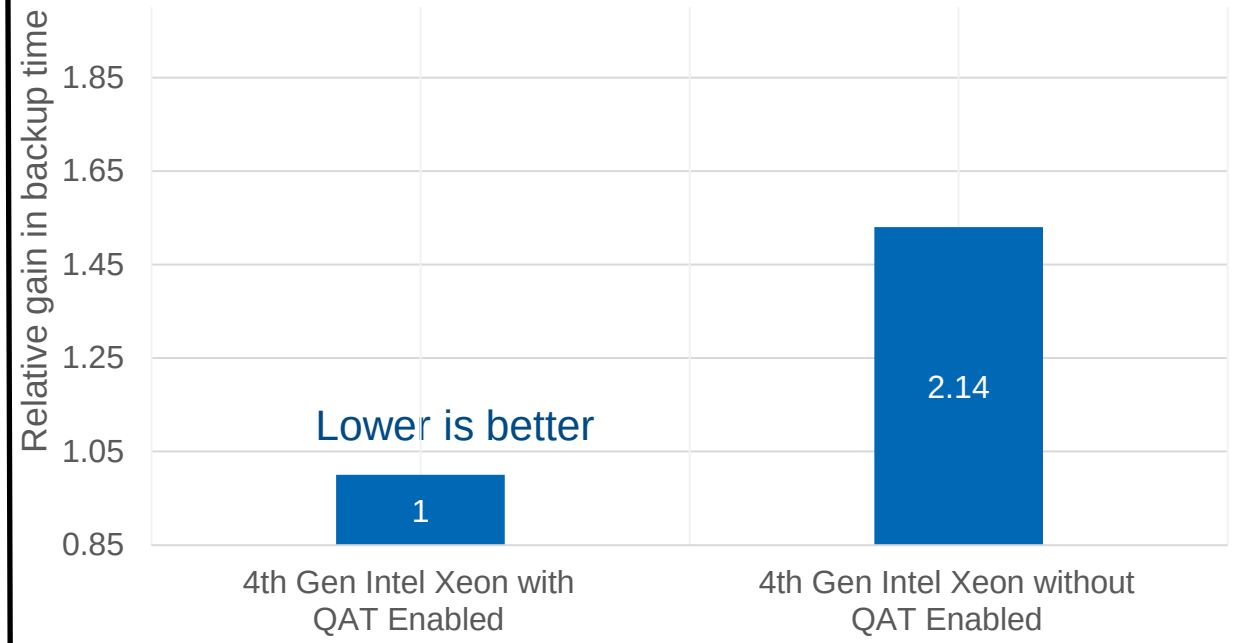

Up to 114% faster backup with 4th Generation Intel® Xeon® Processor and QAT under peak load

Tested by Intel as of 12/12/2022. 1-node, 2x Intel® Xeon® Gold 8460Y+ Processor, 32 cores, HT On, Turbo On, Total Memory 512 GB (16 slots/ 32GB/ 4800 MHz [run @ 4800MHz]) DDR4 memory, one QAT device enabled, ucode 0x2B000081, Windows 2022 Standard Edition 21H2, 10.0.20348, SQL Server 2022, 16.0.1000.6 (X64), database backup without QAT using Xpress software compression.

Results may vary.

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.

© 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.

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