

# VMware Modernization Partner Enablement Package

How our valued partners can build modernization solutions  
based on latest gen Intel® Xeon® processors

# Contents

- **Intel & VMware Better Together**
  - What's in it for me?
  - 15+ Years of Collaboration
- **Benefits of upgrading to 5th Gen Intel® Xeon® processors**
  - Optimized performance for your workloads
  - Why choose 5th Gen Intel® Xeon® processors for your server refresh?
- **Benefits of 5th Gen Intel® Xeon® processors with VMware vSphere Foundation**
  - Benchmarks on 4th and 5th Gen Intel® Xeon® processors
- **VMware Enterprise AI with Intel**
- **Call to Action**
- **Resources**



# Modernizing the Data Center with VMware and Intel Value Proposition

WATCH NOW



[Modernize with VMware and Intel](#)

## For SIs / SPs

- Future VMware direction and impact is top of mind following Broadcom's acquisition
- Connect with your customers and show them how to save money and modernize
- Server & Software upgrades unleashes new AI & Security capabilities leading to multiple revenue streams from modernization of the data center

## For End Customer

- Achieve lower cost through server consolidation, resulting in energy savings
- Reducing the number of servers -> lower SW costs
- Additional incremental benefits from better security and running AI workloads cost-effectively

# Intel and VMware: 15+ Years of Collaboration

VMware and Intel solutions combine co-engineered building blocks, reference designs, and tools to enable digital transformation



## MODERNIZE FOR PRIVATE CLOUD

- **Unify management** of compute, storage, and network resources with hyperconverged infrastructure
- Optimize **VM density and data performance** with Intel® architecture (IA) and built-in accelerators
- **Future proof for AI workloads**



## INTEGRATE MULTI-CLOUD

- **Deploy software-defined data centers** to leverage hybrid clouds
- Leverage agile building blocks to **optimize total cost of ownership (TCO)**
- **Quickly deploy** best of breed designs, optimized for containers, to public clouds
- **Support VMware vMotion (EVC)** across multi-generations of Intel® Xeon® Processors



## VIRTUALIZE CLOUD NETWORKS

- Protected, **any-to-any** connectivity
- Seamlessly **extend workloads from on-prem to multi-cloud**
- Experience cloud-like elastic **scalability**
- **Accelerate packet processing** and encryption performance with IA
- **Reduce CapEx and OpEx costs** by consolidating network functions on Intel-based servers

---

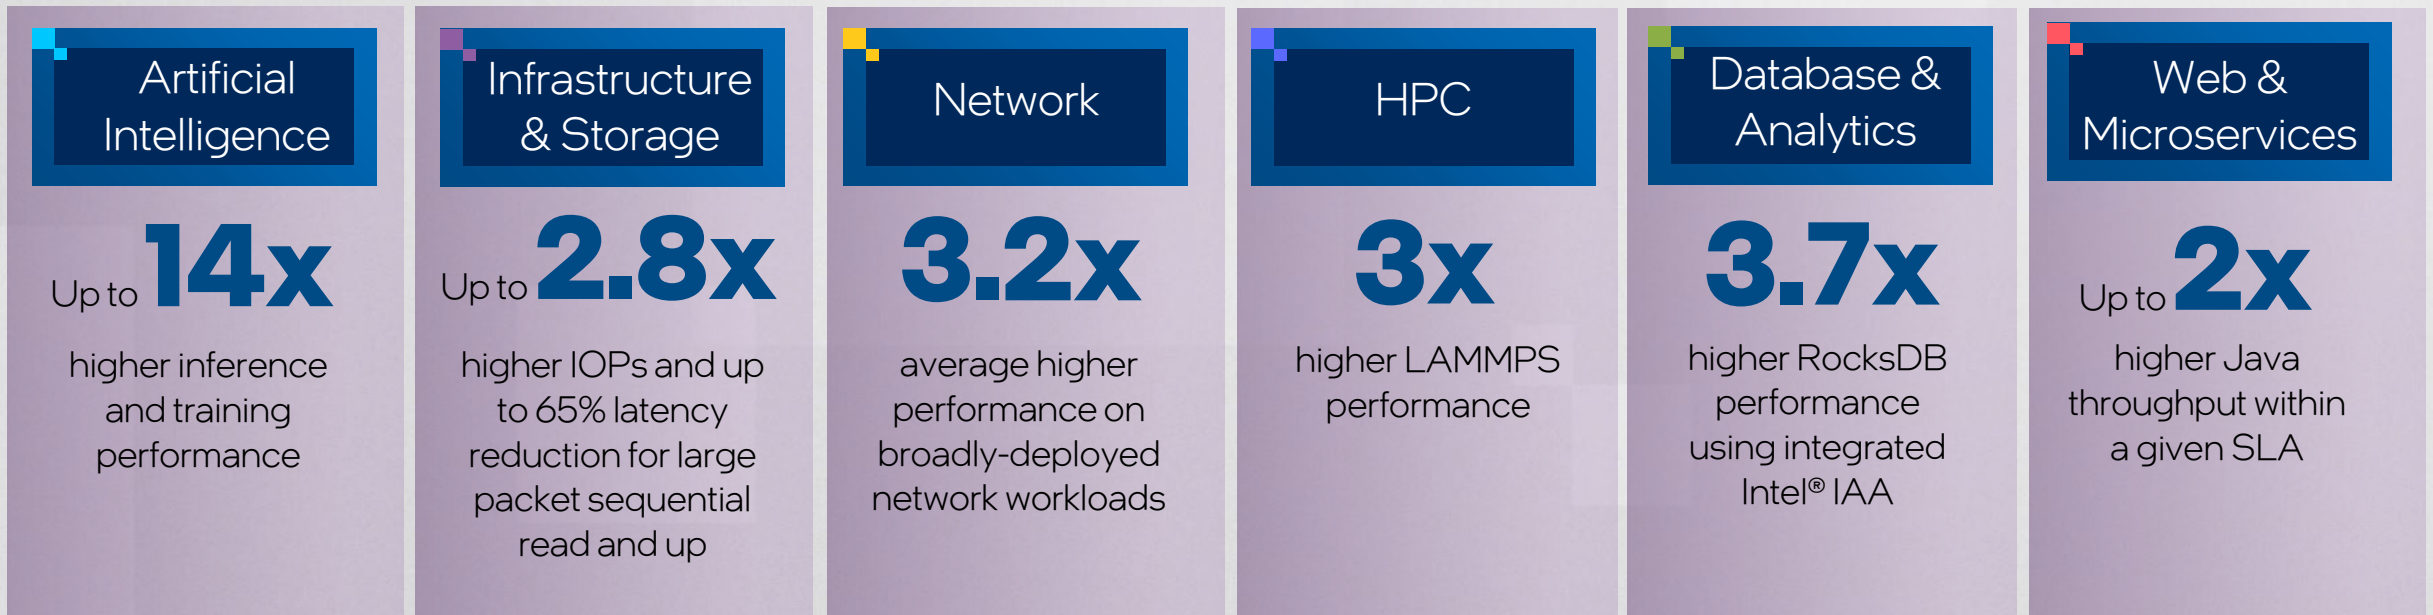
**SECURE FOUNDATION FOR ALL JOINT SOLUTIONS**

# Upgrade to 5th Gen Intel® Xeon® processors

## Optimized Performance for all your Workloads

### Experience Significant Performance Gains

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





# 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<sup>1</sup>

## Improve 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<sup>3</sup>

## 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<sup>2</sup>

## Server Consolidation

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<sup>1</sup>

## Better Security with 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.



<sup>1,2,3</sup>See [T7, G1, T13] at [intel.com/processorclaims](https://www.intel.com/processorclaims): 5th Gen Intel Xeon Scalable processors. Results may vary.

# 5th Gen Intel® Xeon® Processor and vSphere Foundations 8.0

## Keep existing versions of Intel hardware and VMware software

- ❖ Intel® Xeon® Gen 1 out of Support
- ❖ No AI Acceleration
- ❖ Standard Security

Upgrade hardware ONLY

## New Intel hardware and Existing VMware Software

- TCO Savings → Lower Power
- Improved Performance → Server Consolidation
- Enhanced Security and Storage (QAT)

Upgrade BOTH current hardware and software

Upgrade software ONLY

## Existing Intel hardware and new VMware software

- Enhanced Security
- Life Cycle Management

## New Intel hardware and VMware software

- ✓ Better TCO and Performance
- ✓ AI Ready (Intel® AMX)
- ✓ Increased Security
- ✓ Improved Life Cycle Management
- ✓ Significant Storage Performance Improvement (ESA)

[Beyond Savings: How Server Consolidation with VMware vSAN 8 Boosts Performance by more than 7.4x](#)

[How to Consolidate Your Infrastructure with VMware vSphere Foundation and Intel video](#)

READ MORE >

[Is your Hardware and Software Upgrade Strategy Creating a Game of "Upgrade Leapfrog?"](#)

# Optimize Workloads with 5th Gen Intel® Xeon® Processors Accelerators Supported by VMware vSphere Foundation

## Hardware Accelerators

Intel's workload-first approach with **cores + built-in accelerators**, alongside optimized software, **delivers superior performance, higher efficiency, and better TCO.**


**Intel® Advanced Matrix Extensions (Intel® AMX)** Accelerates AI deep learning inference and training workloads

**Intel® QuickAssist Technology (Intel® QAT)** Accelerates cryptography and data de/compression

**Intel® Software Guard Extensions (Intel® SGX)** For increased protection of confidential data

**Memory.** DDR5 and more memory channels increases bandwidth

**Storage.** PCIe Gen 5 doubles the I/O bandwidth



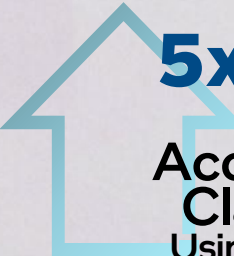
## VMware vSphere Foundation vSphere 8 vSAN 8

[Built-in Accelerators and Why You Should Use Them](#)



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


## BENCHMARKS



Up to  
**5x Faster<sup>1</sup>**  
and Still  
Accurate Image  
Classification  
Using Intel® AMX for  
BF16 compared to Intel®  
AVX-512 for FP32

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

[<sup>1</sup>Solution Brief](#)



**45% Lower  
TCO<sup>2</sup>**  
Lower total cost of  
ownership by more  
than 45% when using  
vSAN vs without

**Boost Performance and  
Lower Latency** with  
VMware vSAN8 and 4<sup>th</sup> Gen  
Intel® Xeon®

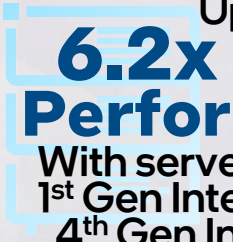
[<sup>2</sup>Solution Brief](#)



Up to  
**5.7x Higher<sup>3</sup>**  
Natural  
Language  
Processing  
INT8 with Intel® AMX vs  
FP32 with Intel® AVX-512

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

[<sup>3</sup>Solution Snapshot](#)



Up to  
**6.2x Better  
Performance<sup>4</sup>**  
With servers featuring  
1<sup>st</sup> Gen Intel® Xeon® vs  
4<sup>th</sup> Gen Intel® Xeon®

VMware vSAN on 4<sup>th</sup> Gen  
Intel® Xeon® for  
**Modernization**

[<sup>4</sup>Solution Snapshot](#)

[\*\*READ MORE >\*\*](#)

[VMware vSAN and 4th Gen Intel® Xeon® Processors Outpace AMD Genoa](#)

# Value of VMware with 4<sup>th</sup> & 5<sup>th</sup> Gen Intel® Xeon® Processors



Up to  
**22.4x** Higher  
Performance


with 5<sup>th</sup> Gen Intel Xeon Processors compared to older servers using 2<sup>nd</sup> Gen Intel Xeon Scalable Processors

Up to  
**12.5:1** Server  
Consolidation

with 5<sup>th</sup> Gen Intel Xeon Processors compared to older servers using 2<sup>nd</sup> Gen Intel Xeon Scalable Processors

## MORE INFO

[Boost VMware Performance While Consolidating Servers and Reducing Cost](#)



Up to  
**39.6%** More LLM Tokens

with Intel® AMX on 5<sup>th</sup> gen Intel Xeon Platinum vs. Gold processors

Up to  
**4.8x** More Images

with Intel® AMX on 5<sup>th</sup> vs 3<sup>rd</sup> gen Intel Xeon Processors

Up to  
**4.6x** More NLP Samples

with Intel AMX on 5<sup>th</sup> vs 3<sup>rd</sup> gen



**1.42x** Faster Backup during ACTIVE database use using Intel® QAT for compression on 4<sup>th</sup> Gen Intel Xeon processors over standard SQL compression.<sup>1</sup>



**1.28x** More VDI Users using 4<sup>th</sup> Gen Intel Xeon processors vs. 3<sup>rd</sup> Gen<sup>1</sup>

## MORE 5TH GEN PROOF POINTS

Power Efficiency + Image Classification and NLP

<sup>1</sup>See backup for workloads and configurations. Results may vary.



# Benefits of Bringing AI Everywhere in the Enterprise

## Intel and VMware Better Together

Enterprise AI

Intel® Xeon®

VMware

Better Together

90%<sup>1</sup>

Enterprise Apps will be Infused with AI by 2025

100M+

Intel® Xeon® Install Base

100%<sup>2</sup>

of Fortune 500 Global Companies use VMware Technologies & Services

4<sup>th</sup> and 5<sup>th</sup> Gen Intel® Xeon®

with Built-in AI

+

VMware Products & Services

Bringing AI to your Mainstream Infrastructure

1: Forbes: <https://www.forbes.com/sites/gilpress/2019/11/22/top-artificial-intelligence-ai-predictions-for-2020-from-idc-and-forrester/#4fef9821315a>

2: VMware: <https://npifinancial.com/blog/how-will-broadcoms-acquisition-of-vmware-affect-vmware-customers/>

# VMware Enterprise AI with Intel

## Unlock AI Everywhere



### Enable Privacy and Security for AI Models

Build and deploy on a robust multi-cloud platform with integrated security and management



### Boost AI Performance

Achieve excellent AI model performance using VCF and Intel processors, hardware accelerators and optimized software



### Deploy AI Everywhere

Get a fully validated AI stack on already deployed clusters

Announcing New Collaborations in  
**VMware Enterprise AI**

[Blog](#)  
[Webinar](#)  
[Website](#)

**Intel® AMX** - The next big step in AI; cost-effective means of running AI workloads without the need for a GPU  
[Technical Brief](#) showing how **VMware Enterprise AI with 4th Gen Intel® Xeon® and built-in Intel® AMX acceleration** will let you run your AI apps anywhere you want



# Enterprise AI with Intel: transformative AI with superior privacy, security, performance, scale and TCO



The combination of **VMware vSphere Foundation** and **Intel's AI software suite and Intel® Xeon® processors with built-in AI accelerators** will deliver a **validated and benchmarked AI stack** for data preparation, model training, fine-tuning and inferencing to accelerate scientific discovery and enrich business and consumer services



VMware Enterprise AI will be supported by servers from Dell Technologies, Hewlett Packard Enterprise and Lenovo running 4th and 5th Gen Intel® Xeon® CPUs with Intel® Advanced Matrix Extensions (Intel® AMX).

# Call to Action for SIs

Your customers have questions on VMware moving to a subscription model

**!** Switching platforms is **costly and time consuming**



Showcase the streamlined migration path to vSphere 8.0 with the latest Xeon generation servers, which will **reduce TCO**



Highlight how to **save more money** by running emerging AI workloads on latest generation Intel® Xeon® processors

**How to Get Started:** Access the [Intel® Xeon® Processor Advisor Suite](#) to calculate your customers' best route to lower TCO and their ROI

[Learn More](#)



# Intel® Xeon® Processor Advisor Suite

## Optimize Performance, TCO, ROI and Power Consumption

Discover the best time to refresh and how to scale your data center:



Access up-to-date Specifications for Benchmarks and Pricing for On-Prem Deployments



Get Product Recommendations for CPUs, PMEM, SSDs and Networking



Get Solutions Recommendations for Systems and Instances by Workload



Calculate TCO and ROI for Node and Rack Based Data Center Solutions

**Instant Expert Advice 24x7 to Optimize your Data Center TCO and Sustainability**

[START NOW](#)

# Additional Resources

Asset Type	Title and Link
Product Briefcase	<a href="#">Accelerate End-to-End AI Pipelines on VMware Powered by 4th Gen Intel Xeon Scalable Processors</a>
Research Paper	<a href="#">The CIO of 2025: Leading Technology-driven Growth Across The Entire Organization - Forbes Insights by Intel and VMware</a>
OEM Solution Brief	<a href="#">HPE expands partnership with VMware and Intel to accelerate AI for all organizations</a>
OEM Solution Brief	<a href="#">New Intel and VMware Technologies Turbocharge Lenovo ThinkAgile VX V3 Systems</a>
Webinar	<a href="#">Multi-cloud Environments with Intel &amp; VMware</a>
Sales Guide	<a href="#">Reasons to Modernize: A Conversation Guide</a>
Blog	<a href="#">Configuration Guidance for VMware Private AI with Intel</a>
Solution Brief	<a href="#">VMware Horizon with 4th Gen Intel® Xeon Processors — Lab Insights</a>
Research Paper	<a href="#">2024 VMware Explore: The Return of the Private Cloud</a>
Blog	<a href="#">Navigating the Clouds: A Story of Transformation and Partnership with VMware Cloud Foundation</a>
VMware Brief	<a href="#">Why Run Modern Apps on VMware Hyperconverged Infrastructure</a>
VMware Solution Overview	<a href="#">Deployment, Management, and Monitoring for Mixed Workloads on VMware vSAN™</a>
Video	<a href="#">Intel vPro® and VMware Workspace ONE: Modern Management and Security from Chip to Cloud</a>
Solution Overview	<a href="#">Intel vPro and VMware Workspace ONE</a>



# Notices and Disclaimers

- Performance varies by use, configuration and other factors. Learn more on the [Performance Index site](#).
- Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.
- Your costs and results may vary.
- Intel technologies may require enabled hardware, software or service activation.
- © 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.

The Intel logo is centered on a dark blue background. It features the word "intel" in a white, lowercase, sans-serif font. A small, bright blue square is positioned above the letter 'i'. To the right of the word "intel" is a registered trademark symbol (®). The background is a solid dark blue with several faint, semi-transparent squares of varying shades of blue scattered across it.

intel®



# 5<sup>th</sup> Generation Intel® Xeon® Processor Proof Points




# Performance Boost Plus Increased Power Efficiency for VMware VCF on 5<sup>th</sup> Gen Intel® Xeon® Processors

Performance of HCI Bench - vSAN 8.U2 ESA, 4-node cluster

Higher is better

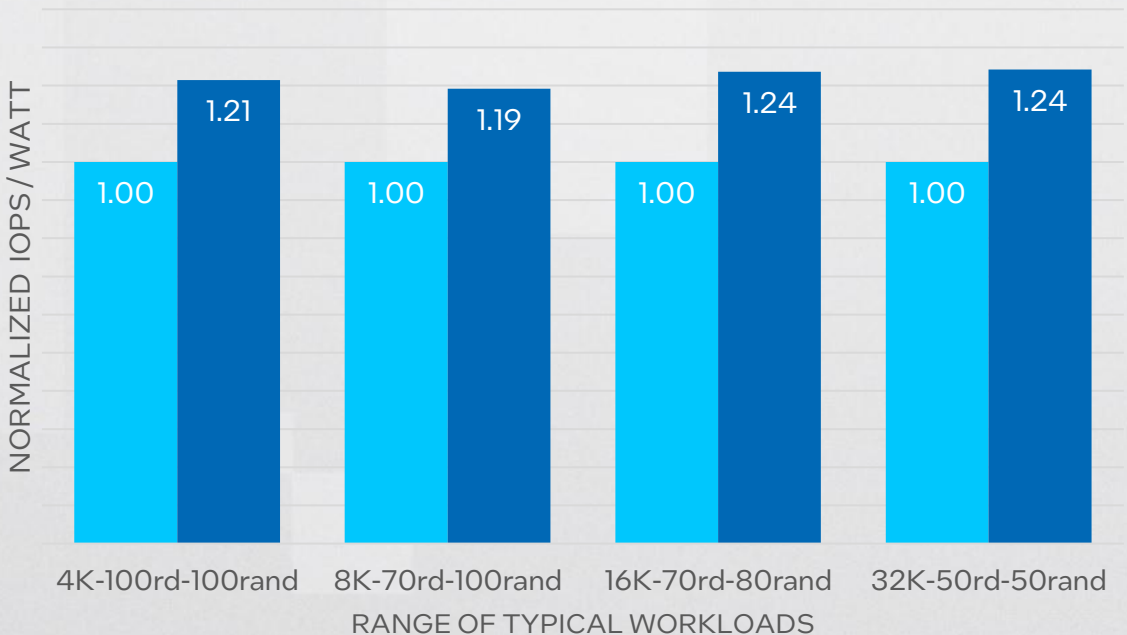
■ Intel® Xeon® Platinum 8490H ■ Intel® Xeon® Platinum 8592+



Performance/Power of HCI Bench - vSAN 8.U2 ESA

Higher is better

■ Intel® Xeon® Platinum 8490H ■ Intel® Xeon® Platinum 8592+



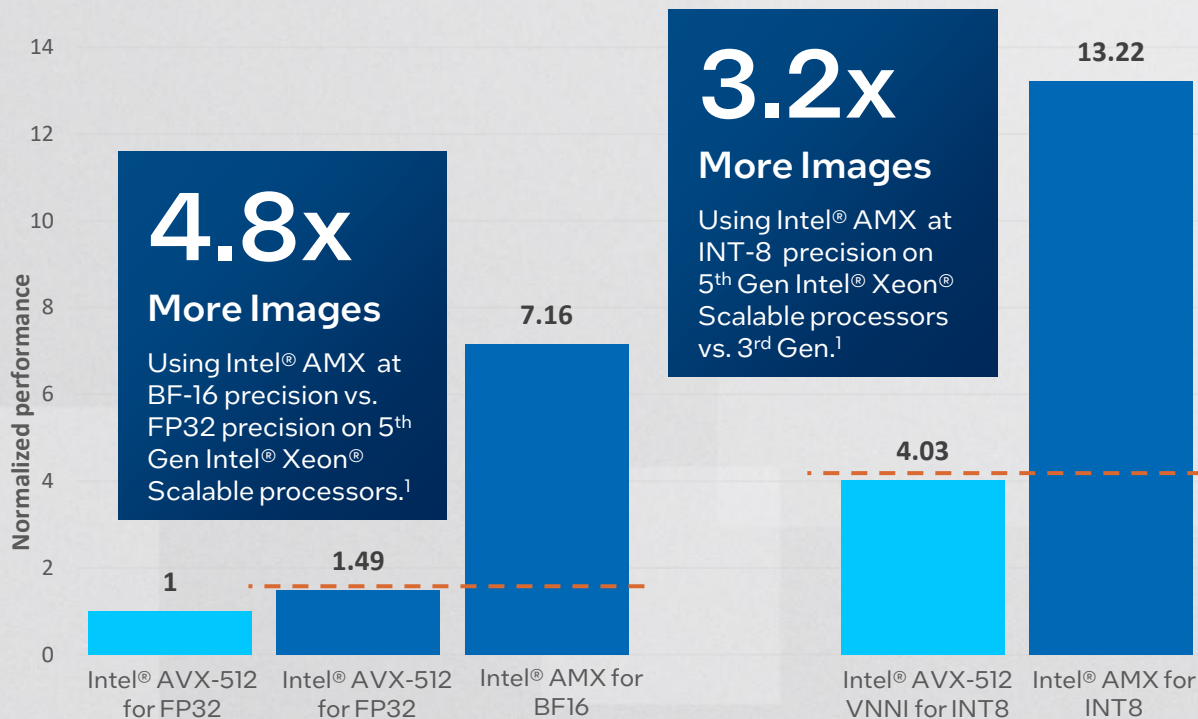
**Up to 7% higher IOPS with up to 24% reduction in performance/watt by using 5<sup>th</sup> Gen Intel® Xeon® processors vs. previous generation**

# AI: Image Classification and NLP on VMware VCF

## 5<sup>th</sup> vs. 3<sup>rd</sup> Generation Intel® Xeon® Processors

### Normalized Performance on TensorFlow 2.14 using ResNet-50


Batch Size = 128, Multi-Instance (4 cores per Instance)  
(Higher is Better)



Intel® Xeon® Gold 6348, 2.6GHz, 28c  
Intel® AVX-512 + Intel® DL Boost

### Normalized Performance on TensorFlow 2.14 using BERT-Large

Batch Size = 128, 28x2 vs. 32x2 Instances  
(Higher is Better)




Intel® Xeon® Platinum 6548Y+, 2.5GHz, 32c  
Intel® AVX-512 + Intel® DL Boost; Intel® AMX for BF16, INT8

# AI: Image Classification and NLP on VMware VCF

## 5<sup>th</sup> vs. 3<sup>rd</sup> Generation Intel® Xeon® Scalable Processors on Dell PowerEdge Servers

### Normalized Performance on TensorFlow 2.14 using ResNet-50

Batch Size = 128, Multi-Instance (4 cores per Instance)  
(Higher is Better)




Intel® Xeon® Platinum 8358, 2.6GHz, 32c  
Intel® AVX-512 + Intel® DL Boost

Intel® Xeon® Platinum 8562Y+, 2.8GHz, 32c  
Intel® AVX-512 + Intel® DL Boost; Intel® AMX for BF16, INT8

### Normalized Performance on TensorFlow 2.14 using BERT-Large

Batch Size = 128, 32x2 Instances  
(Higher is Better)

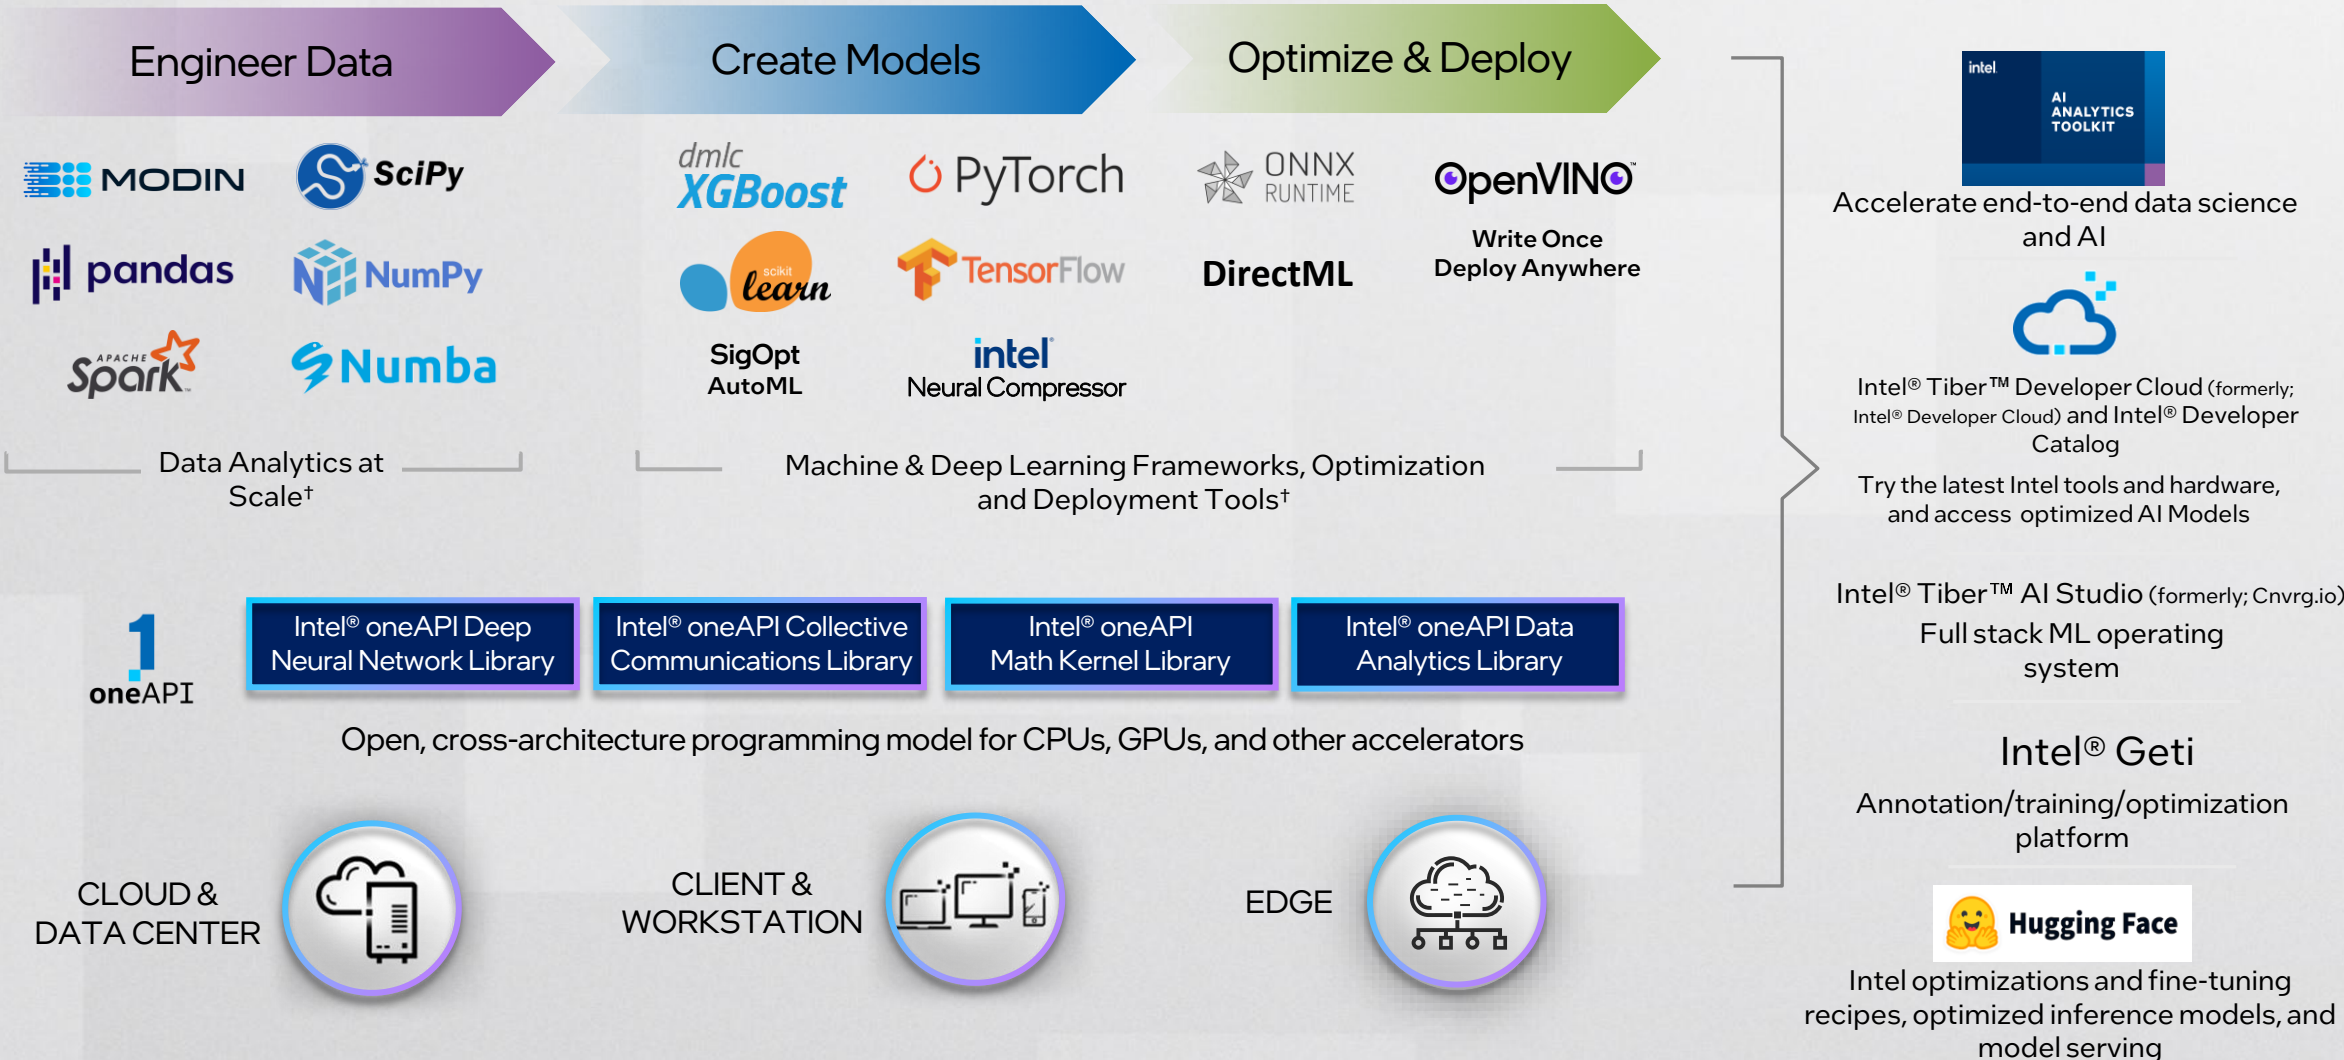


Intel® Xeon® Platinum 8358, 2.6GHz, 32c  
Intel® AVX-512 + Intel® DL Boost

Intel® Xeon® Platinum 8562Y+, 2.8GHz, 32c  
Intel® AVX-512 + Intel® DL Boost; Intel® AMX for BF16, INT8




# Intel® AI Software Portfolio



Note: components at each layer of the stack are optimized for targeted components at other layers based on expected AI usage models, and not every component is utilized by the solutions in the rightmost column

<sup>†</sup> This list includes popular open-source frameworks that are optimized for Intel hardware


# Configuration: VMware vSAN 8 on 5<sup>th</sup> Gen Intel® Xeon® Processors



- 1. Intel® Xeon® Platinum 8490H:** Tested by Intel as of 11/17/23; 4-node cluster; Each node: 2x Intel Xeon Platinum 8490H, 60 cores, 1.9GHz, HT On, Turbo On, NUMA 2, Integrated Accelerators Available [used]: DLB 8 [0], DSA 8 [0], IAA 8 [0], QAT 8 [0]. Total Memory 512GB (16x32GB DDR5 4800 MT/s [4800 MT/s]), BIOS 05.01.00, microcode 0x2b000461, 2x Intel® Ethernet Controller E810-C for QSFP 100G, Drives: 1x 894.3G Intel SSDSC2KG960G8, 8x 3.5TB Intel SSDPF2KX038TZ. OS/Software: VMware vSphere/vSAN 8.0 U2, build 22380479, vSAN ESA default, using HCI Bench 2.82, FIO3.3. Throughput test in IOPS and IOPS/watt with multiple profiles as shown on slide. #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB.
- 2. Intel® Xeon® Platinum 8592+:** Tested by Intel as of 10/12/23. 4-node cluster; Each node: 2x Intel Xeon Platinum 8592+, 64 cores, 1.9GHz, HT On, Turbo On, NUMA 2, Accelerators Available [used]: DLB 2 [0], DSA 2 [0], IAA 2 [0], QAT 2 [0]. Total Memory 512GB (16x32GB DDR5 5600 MT/s [5600 MT/s]), BIOS 3B05.TEL4P1, microcode 0x21000161, 2x I350 Gigabit Network Connection, 2x Intel® Ethernet Controller E810-C for QSFP 100G, Drives: 1x 894.3G Intel SSDSC2KG960G8, 8x 3.5TB Intel SSDPF2KX038TZ. OS/Software: VMware vSphere/vSAN 8.0 U2, build 22380479, vSAN ESA default, using HCI Bench 2.82, FIO3.3. Throughput test in IOPS and IOPS/watt with multiple profiles as shown on slide. #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB.



# 5<sup>th</sup> Gen vs. 2<sup>nd</sup> Gen Intel® Xeon® Processors I/O Performance Workloads and Configurations



- 1. Intel® Xeon® Gold 6242 – RAID-5:** Test by Intel as of 05/7/24. 4-node clusters, 2x Intel Xeon Gold 6242 CPU @ 2.80GHz, 16 cores, HT On, Turbo On, Total Memory 384GB (12x32GB DDR4 2666 MT/s [2666 MT/s]), BIOS 2.20.1, microcode 0x5003604, 2x I350 Gigabit Network Connection, 2x Ethernet Controller X710 for 10GbE SFP+, Boot: 2x 349.3G INTEL MDTPE21K375GA, Storage: 6x 1.8T INTEL SSDPE2KX020T8, OS/Software: VMware 7.0U3, 23794027, vSAN OSA – default policy (RAID5, 2DG), using HCI Bench 2.8.3, FIO3.3. Throughput test 8K profile: I/O size 8K, Read percentage 70%, Random percentage 100%, latency target mode<10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB. Throughput test 16K profile: I/O size 16K, Read percentage 70%, Random percentage 80%, latency target mode<10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB. Throughput test 32K profile: I/O size 32k, Read percentage 50%, Random percentage 50%, latency target mode<10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB.
- 2. Intel Xeon Gold 6242 – RAID-1:** Test by Intel as of 05/7/24. 4- node clusters, 2x Intel Xeon Gold 6242 CPU @ 2.80GHz, 16 cores, HT On, Turbo On, Total Memory 384GB (12x32GB DDR4 2666 MT/s [2666 MT/s]), BIOS 2.20.1, microcode 0x5003604, 2x I350 Gigabit Network Connection, 2x Ethernet Controller X710 for 10GbE SFP+, Boot: 2x 349.3G INTEL MDTPE21K375GA, Storage: 6x 1.8T INTEL SSDPE2KX020T8, OS/Software: VMware 7.0U3, 23794027, vSAN OSA – default policy (RAID-1, 2DG), using HCI Bench 2.8.3, FIO3.3. Throughput test 8K profile: I/O size 8K, Read percentage 70%, Random percentage 100%, latency target mode<10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB. Throughput test 16K profile: I/O size 16K, Read percentage 70%, Random percentage 80%, latency target mode<10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB. Throughput test 32K profile: I/O size 32k, Read percentage 50%, Random percentage 50%, latency target mode<10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB.
- 3. Intel Xeon Gold 6548Y+:** Test by Intel as of 05/7/24. 4-node cluster, 2x Intel Xeon Gold 6548Y+ CPU @ 2.50GHz, 32 cores, HT On, Turbo On, NUMA 2, Integrated Accelerators Available [used]: DLB 0 [0], DSA 2 [0], IAA 0 [0], QAT 0 [0], Total Memory 512GB (16x32GB DDR5 5600 MT/s [5200 MT/s]), BIOS 3B07.TEL2P1, microcode 0x21000200, 2x Ethernet Controller E810-C for QSFP, Boot: 2x 223.6G INTEL SSDSCKKB240GZ, Storage: 8x 2.9T KIOXIA KCD81PUG3T20, OS/Software: VMware ESXi 8.0.2, 23825572, vSAN ESA – Optimal default policy (RAID-5, flat), using HCI Bench 2.8.3, FIO3.3. Throughput test 8K profile: I/O size 8K, Read percentage 70%, Random percentage 100%, latency target mode<10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB. Throughput test 16K profile: I/O size 16K, Read percentage 70%, Random percentage 80%, latency target mode<10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB. Throughput test 32K profile: I/O size 32k, Read percentage 50%, Random percentage 50%, latency target mode<10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB

## Notices and Disclaimers

Performance varies by use, configuration and other factors. Learn more at [www.intel.com/PerformanceIndex](http://www.intel.com/PerformanceIndex).

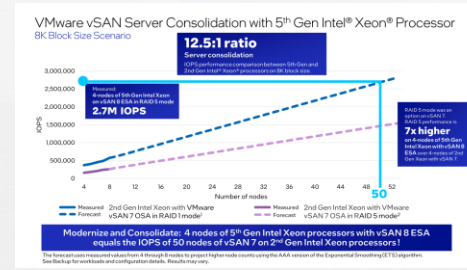
Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

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



# 5<sup>th</sup> Gen vs. 2<sup>nd</sup> Gen Intel® Xeon® Processors Server Consolidation Workloads and Configurations



- 1. Intel® Xeon® Gold 6242 – RAID-5:** Test by Intel as of 05/7/24. 4- to 8-node clusters, 2x Intel Xeon Gold 6242 CPU @ 2.80GHz, 16 cores, HT On, Turbo On, Total Memory 384GB (12x32GB DDR4 2666 MT/s [2666 MT/s]), BIOS 2.20.1, microcode 0x5003604, 2x I350 Gigabit Network Connection, 2x Ethernet Controller X710 for 10GbE SFP+, Boot: 2x 349.3G INTEL MDTPE21K375GA, Storage: 6x 1.8T INTEL SSDPE2KX020T8, OS/Software: VMware 7.0U3, 23794027, vSAN OSA – default policy (RAID5, 2DG), using HCI Bench 2.8.3, FIO3.3. Throughput test 8k profile (I/O size 8k, Read percentage 70%, Random percentage 100%, latency target mode<10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB).
  - 2. Intel Xeon Gold 6242 – RAID-1:** Test by Intel as of 05/7/24. 4- to 8-node clusters, 2x Intel Xeon Gold 6242 CPU @ 2.80GHz, 16 cores, HT On, Turbo On, Total Memory 384GB (12x32GB DDR4 2666 MT/s [2666 MT/s]), BIOS 2.20.1, microcode 0x5003604, 2x I350 Gigabit Network Connection, 2x Ethernet Controller X710 for 10GbE SFP+, Boot: 2x 349.3G INTEL MDTPE21K375GA, Storage: 6x 1.8T INTEL SSDPE2KX020T8, OS/Software: VMware 7.0U3, 23794027, vSAN OSA – default policy (RAID-1, 2DG), using HCI Bench 2.8.3, FIO3.3. Throughput test 8k profile (I/O size 8k, Read percentage 70%, Random percentage 100%, latency target mode<10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB).
- For #1 and #2, performance was measured at 4 to 8 nodes (five data points), and then projected to higher node counts using exponential smoothing forecasting. Exponential smoothing forecasting in Excel is based on the AAA version (additive error, additive trend and additive seasonality) of the Exponential Triple Smoothing (ETS) algorithm, which smooths out minor deviations in past data trends by detecting seasonality patterns and confidence intervals.
- 3. Intel Xeon Gold 6548Y+:** Test by Intel as of 05/7/24. 4-node cluster, 2x Intel Xeon Gold 6548Y+ @ 2.50GHz, 32 cores, HT On, Turbo On, NUMA 2, Integrated Accelerators Available [used]: DLB 0 [0], DSA 2 [0], IAA 0 [0], QAT 0 [0], Total Memory 512GB (16x32GB DDR5 5600 MT/s [5200 MT/s]), BIOS 3B07.TEL2P1, microcode 0x21000200, 2x Ethernet Controller E810-C for QSFP, Boot: 2x 223.6G INTEL SSDSCKKB240GZ, Storage: 8x 2.9T KIOXIA KCD81PUG3T20, OS/Software: VMware ESXi 8.0.2, 23825572, vSAN ESA – Optimal default policy (RAID-5, flat), using HCI Bench 2.8.3, FIO3.3. Throughput test 8k profile (I/O size 8k, Read percentage 70%, Random percentage 100%, latency target mode<10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB).

## Notices and Disclaimers

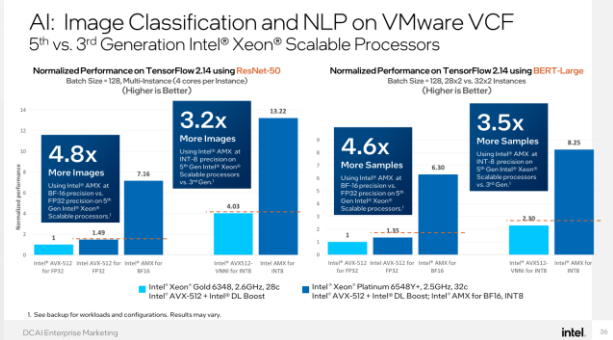
Performance varies by use, configuration and other factors. Learn more at [www.intel.com/PerformanceIndex](http://www.intel.com/PerformanceIndex).

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

Your costs and results may vary.

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

# Configuration Details – AI on 5th Gen Intel® Xeon® Processors

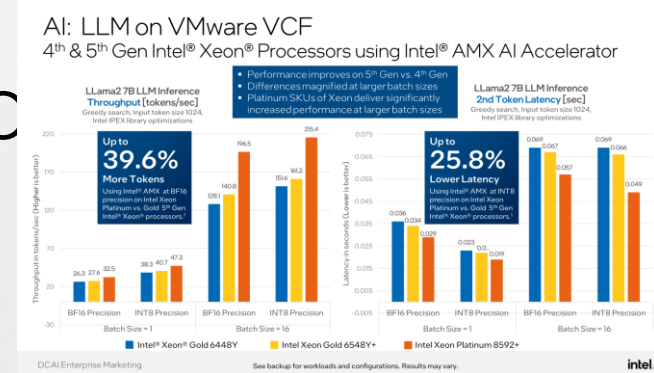


BASELINE on 3<sup>rd</sup> Gen Xeon Config: Intel® Xeon® Gold 6348. 1-node, 2x Intel Xeon Gold 6348 CPU, 28 cores @ 2.60GHz, HT On, Turbo On, NUMA 2, Integrated Accelerators Available [used]: DLB 0 [0], DSA 0 [0], IAA 0 [0], QAT 0 [0], Total Memory 512GB (16x32GB DDR4 3200 MT/s [3200 MT/s]), BIOS SE5C620.86B.01.01.0009.2311021928, microcode 0xd0003b9, 2x Ethernet Controller E810-C for QSFP, Drives: 9x 3.5TB INTEL SSDPF2KX038TZ, 2x 54.9G INTEL SSDPEK1A058GA, VMware vSphere 8.0U2, build 22380479, Ubuntu Server 22.04.3 VM (vHW=21, vmxnet3), Kernel 5.15, intel-optimized-tensorflow:2.14, ResNet50v1.5, Batch size=128, VM=56vCPU+400GB RAM, Multi instance scenario (4 cores per instance), BERT-Large, SQuAD 1.1, Batch size=128, VM=56vCPU+400GB RAM. Test by Intel as of 12/11/23.

5th Gen Xeon Configuration: Intel® Xeon® Gold 6548Y+. 1-node, 2x Intel Xeon Gold 6548Y+, 32 cores @ 2.5GHz, HT On, Turbo On, NUMA 2, Integrated Accelerators Available [used]: DLB 2 [0], DSA 2 [0], IAA 2 [0], QAT 2 [0], Total Memory 512GB (16x32GB DDR5 5600 MT/s [5200 MT/s]), BIOS 3B05.TEL4P1, microcode 0x21000161, 2x Ethernet Controller E810-C for QSFP, Drives: 8x 3.5TB INTEL SSDPF2KX038TZ, 1x 894.3G INTEL SSDSC2KG960G8, VMware vSphere 8.0U2, build 22380479, Ubuntu Server 22.04.3 VM (vHW=21, vmxnet3), Kernel 5.15, intel-optimized-tensorflow:2.14, ResNet50v1.5, Batch size=128, VM=64vCPU+400GB RAM, Multi instance scenario (4 cores per instance), BERT-Large, SQuAD 1.1, Batch size=128, VM=64vCPU+400GB RAM. Test by Intel as of 11/24/23.



# Configuration Details – LLM on VCF on 5th Gen Intel® Xeon® Processors



**4th Gen Xeon Config: Intel® Xeon® Gold 6448Y.** 1-node, 2x Intel Xeon Gold 6448Y CPU, 32 cores @ 2.10GHz, HT On, Turbo On, NUMA 2, Integrated Accelerators Available [used]: DLB 0 [0], DSA 2 [0], IAA 0 [0], QAT 0 [0], Total Memory 512GB (16x32GB DDR5 5600 MT/s [4800 MT/s]), BIOS 3B07.TEL2P1, microcode 0x2b000590, 2x Ethernet Controller E810-C for QSFP, Drives: 1x 894.3G Intel SSDSC2KG960G8, 8x 3.5T Intel SSDPF2KX038TZ, VMware vSphere 8.0U2, build 22380479, VM: 60 vCPU + 256GB vRAM, Ubuntu Server 22.04.4 (vHW=21, vmxnet3), kernel 5.15.0-94, LLM Llama2 inference, Pytorch/IPEX 2.2, score=?UNITS, VM=56vCPU+400GB RAM. Test by Intel as of 6/19/24.

**5th Gen Xeon Configuration: Intel® Xeon® Gold 6548Y+.** 1-node, 2x Intel Xeon Gold 6548Y+, 32 cores @ 2.5GHz, HT On, Turbo On, NUMA 2, Integrated Accelerators Available [used]: DLB 2 [0], DSA 2 [0], IAA 2 [0], QAT 2 [0], Total Memory 512GB (16x32GB DDR5 5600 MT/s [5200 MT/s]), BIOS 3B07.TEL2P1, microcode 0x21000200, 2x Ethernet Controller E810-C for QSFP, Drives: 1x 894.3G Intel SSDSC2KG960G8, 8x 3.5T Intel SSDPF2KX038TZ, VMware vSphere 8.0U2, build 22380479, VM: 60 vCPU + 256GB vRAM, Ubuntu Server 22.04.3 VM (vHW=21, vmxnet3), Kernel 5.15.0-94, LLM Llama2 inference, Pytorch/IPEX 2.2, score=?UNITS. Test by Intel as of 6/19/24.

**5th Gen Xeon Configuration: Intel® Xeon® Platinum 8592+.** 1-node, 2x Intel Xeon Platinum 8592+, 64 cores @ 1.9GHz, HT On, Turbo On, NUMA 2, Integrated Accelerators Available [used]: DLB 2 [0], DSA 2 [0], IAA 2 [0], QAT 2 [0], Total Memory 512GB (16x32GB DDR5 5600 MT/s [5600 MT/s]), BIOS 3B07.TEL2P1, microcode 0x21000200, 2x Ethernet Controller E810-C for QSFP, Drives: 1x 894.3G Intel SSDSC2KG960G8, 8x 3.5T Intel SSDPF2KX038TZ, VMware vSphere 8.0U2, build 22380479, VM: 124 vCPU + 256GB vRAM, OS: Ubuntu Server 22.04.4 (vHW=21, vmxnet3), kernel 5.15.0-94, LLM Llama2 inference, Pytorch/IPEX 2.2, score=?UNITS. Test by Intel as of 6/19/24.



# AI on vSphere/vSAN 8.0 Configuration Details (3<sup>rd</sup> Gen vs. 4<sup>th</sup> Gen)

**3rd Gen Intel® Xeon® Scalable platform configuration:** 4-node cluster, Each node: 2x Intel® Xeon® Gold 6348 Processor, 1x Server Board M50CYP2UR, Total Memory 512 GB (16x 32GB DDR4 3200MHz), HyperThreading: Enable, Turbo: Enabled, NUMA noSNC, Intel VMD: Enabled, BIOS: SE5C620.86B.01.01.0006.2207150335 (ucode:0xd000375), Storage (boot): 2x 80 GB Solidigm SSD P1600X, Storage (cache): 2x 400 GB Intel® Optane™ DC SSD P5800X Series, Storage (capacity): 6x 3.84 TB Solidigm SSD DC P5510 Series PCIe NVMe, Network devices: 1x Intel Ethernet E810CQDA2 E810-CQDA2, FW 4.0, at 100 GbE RoCE, Network speed: 100 GbE, OS/Software: VMware/vSAN 8.0, 20513097, Test by Intel as of 03/08/2023 using Ubuntu Server 22.04 VM (vHW=20, vmxnet3), vSAN default policy (RAID-1, 2DG), Kernel 5.15, Intel-optimized-tensorflow:2.11.0, ResNet50v1.5, Batch size=128, VM=56vCPU+64GBRAM, Multi-instance scenario (4 cores per instance), BERT-Large, SQuAD 1.1, Batch size=128, VM=56vCPU+64GBRAM

**4th Gen Intel® Xeon® Scalable platform configuration:** 4-node cluster, Each node: 2x Intel® Xeon® Gold 6448Y Processor QS pre-production, 1x Server Board M50FCP2SBSTD, Total Memory 512 GB (16x DDR5 32GB 4800MHz), HyperThreading: Enable, Turbo: Enabled, NUMA noSNC, Intel VMD: Enabled, BIOS: SE5C741.86B.01.01.0002.2212220608 (ucode:0x2b000161), Storage (boot): 2x240GB Solidigm S4520, Storage (data): 6x 3.84 TB Solidigm SSD DC P5510 Series PCIe NVMe, Network devices: 1x Intel Ethernet E810CQDA2 E810-CQDA2, FW 4.0, at 100 GbE RoCE, Network speed: 100 GbE, OS/Software: VMware/vSAN 8.0, 20513097, Test by Intel as of 03/13/2023 using Ubuntu Server 22.04 VM (vHW=20, vmxnet3), vSAN ESA – Optimal default policy (RAID-5, flat), Kernel 5.15, intel-optimized-tensorflow:2.11.0, ResNet50v1.5, Batch size=128, VM=64vCPU+64GBRAM, Multi-instance scenario (4 cores per instance), BERT-Large, SQuAD 1.1, Batch size=128, VM=64vCPU+64GBRAM

# MS SQL on vSphere/vSAN Configuration Details

**4th Gen Intel® Xeon® Scalable processor systems:** Tested by Evaluator Group as of 03/17/23

4-nodes, 4x Intel(R) Xeon(R) Platinum 8462Y+, 32 cores each, HT On, Turbo On, Total Memory 2TB Samsung 64GB DIMMS, BIOS 3A11.uh, microcode 0x2b000111, 2x Ethernet Controller 10-Gigabit X540-AT2, 2x P5800x (Gen4) Optane 1.6TB, 6x 3.8T INTEL SSDPF2KX038TZ, Windows Server 2022, HammerDB 4.5, MS SQL 2022-SSEI-Eval/SQL Server Management Studio 18.12.1, ESXi QAT driver: QAT2.0.W.2.0.1, Windows QAT driver: QAT2.0.W.2.0.1, vCenter 8.0.0, 21216066

# “Large Enterprise Relational Database” (Oracle DB) on vSphere/vSAN Configuration Details

## Config 1 – 3rd Gen Intel® Xeon® Platinum


4-node, Each node: Intel Software Development Platform, 2x Intel® Xeon® Platinum 8358 processor (32C, 2.6GHz, 250W TDP), HT On, Turbo ON, SNC OFF, Total Memory: 2 TB (32x64GB DDR4 2DPC 3200 MHz), ucode: 0x0d000375, Intel E810-CQDA2 100GbE, 2 Groups: Per node cache tier: 2x 1.6TB P5800X Gen 4 SD 5800X, Per node capacity tier: 6x 3.84TB D7-P5510 Series, Gen4. ESXi 8.0.0, 21216066, vCenter -8.0.0, 21216066, Oracle 21c, HammerDB 4.7. 32 VM per 4-node vSAN cluster. Per VM 12vCPU, 224GB Memory, 1 OS disk 90GB, 1 data disk 650GB, 1 log disk 150GB, 2000 WH. Tested by Evaluator Group as of July 2023.

## Config 2 – 4<sup>th</sup> Gen Intel® Xeon® Platinum

4-node, Each node: QuantaGrid D54Q-2U, 2x Intel® Xeon® Platinum 8462Y+ processor (32C, 2.8GHz, 300W TDP), HT On, Turbo ON, SNC OFF, Total Memory: 2 TB (32x64GB DDR5 2DPC 4800 MHz), ucode: 0x2b000161, Intel E810-CQDA2 100GbE, 2 Groups: Per node cache tier: 2x 1.6TB P5800X Gen 4 SD 5800X, Per node capacity tier: 6x 3.84TB D7-P5510 Series, Gen4. ESXi 8.0.0, 21216066, vCenter -8.0.0, 21216066, Oracle 21c, HammerDB 4.7. 32 VM per 4-node vSAN cluster. Per VM 12vCPU, 224GB Memory, 1 OS disk 90GB, 1 data disk 650GB, 1 log disk 150GB, 2000 WH. Tested by Evaluator Group as of July 2023.



# VMware Horizon Config Information – 4<sup>th</sup> Gen Intel® Xeon® Processors




## Config 1 – 3rd Gen Intel® Xeon® Platinum, 830 users

4-node, Each node, Intel Software Development Platform, 2x Intel® Xeon® Platinum 8358 processor (32C, 2.6GHz, 250W TDP), HT On, Turbo ON, SNC OFF, Total Memory: 2 TB (32x64GB DDR4 2DPC 3200 MHz), ucode: 0x0d000375, Intel E810-CQDA2 100G, 2 Groups: Per node cache tier: 2x 1.6TB P5800X Gen 4 SD 5800X, Per node capacity tier: 6x 3.84TB D7-P5510 Series, Gen4, ESXi 8.0.0, 21216066, vCenter -8.0.0, 21216066, Horizon 8.9.0 2303 Build 21593375, LoginVSI 5.2.2. Knowledge worker profile 2vCPU/4GB. EUX 7.5. Tested by Evaluator Group as of August 2023.

## Config 2 – 4<sup>th</sup> Gen Intel® Xeon® Platinum, 1060 users

4-node, Each node, QuantaGrid D54Q-2U, 2x Intel® Xeon® Platinum 8462Y+ processor (32C, 2.8GHz, 300W TDP), HT On, Turbo ON, SNC OFF, Total Memory: 2 TB (32x64GB DDR5 2DPC 4800 MHz), ucode: 0x2b000161, Intel E810-CQDA2 100G, 2 Groups: Per node cache tier: 2x 1.6TB P5800X Gen 4 SD 5800X, Per node capacity tier: 6x 3.84TB D7-P5510 Series, Gen4, ESXi 8.0.0, 21216066, vCenter -8.0.0, 21216066, Horizon 8.9.0 2303 Build 21593375, LoginVSI 5.2.2. Knowledge worker profile 2vCPU/4GB. EUX 7.5 Tested by Evaluator Group as of August 2023.

# Configuration: VMware vSAN 8 on 5<sup>th</sup> Gen Intel® Xeon® Processors




- 1. Intel® Xeon® Platinum 8490H:** Tested by Intel as of 11/17/23; 4-node cluster; Each node: 2x Intel Xeon Platinum 8490H, 60 cores, 1.9GHz, HT On, Turbo On, NUMA 2, Integrated Accelerators Available [used]: DLB 8 [0], DSA 8 [0], IAA 8 [0], QAT 8 [0]. Total Memory 512GB (16x32GB DDR5 4800 MT/s [4800 MT/s]), BIOS 05.01.00, microcode 0x2b000461, 2x Intel® Ethernet Controller E810-C for QSFP 100G, Drives: 1x 894.3G Intel SSDSC2KG960G8, 8x 3.5TB Intel SSDPF2KX038TZ. OS/Software: VMware vSphere/vSAN 8.0 U2, build 22380479, vSAN ESA default, using HCI Bench 2.82, FIO3.3. Throughput test in IOPS and IOPS/watt with multiple profiles as shown on slide. #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB.
- 2. Intel® Xeon® Platinum 8592+:** Tested by Intel as of 10/12/23. 4-node cluster; Each node: 2x Intel Xeon Platinum 8592+, 64 cores, 1.9GHz, HT On, Turbo On, NUMA 2, Accelerators Available [used]: DLB 2 [0], DSA 2 [0], IAA 2 [0], QAT 2 [0]. Total Memory 512GB (16x32GB DDR5 5600 MT/s [5600 MT/s]), BIOS 3B05.TEL4P1, microcode 0x21000161, 2x I350 Gigabit Network Connection, 2x Intel® Ethernet Controller E810-C for QSFP 100G, Drives: 1x 894.3G Intel SSDSC2KG960G8, 8x 3.5TB Intel SSDPF2KX038TZ. OS/Software: VMware vSphere/vSAN 8.0 U2, build 22380479, vSAN ESA default, using HCI Bench 2.82, FIO3.3. Throughput test in IOPS and IOPS/watt with multiple profiles as shown on slide. #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50GB.

# Configuration Details – AI on 5th Gen Intel® Xeon® Processors

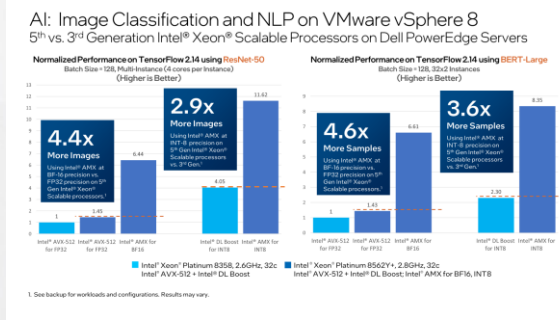
**BASELINE on 3<sup>rd</sup> Gen Intel® Xeon® Config:** Intel® Xeon® Gold 6348. 1-node, 2x Intel Xeon Gold 6348 CPU, 28 cores @ 2.60GHz, HT On, Turbo On, NUMA 2, Integrated Accelerators Available [used]: DLB 0 [0], DSA 0 [0], IAA 0 [0], QAT 0 [0], Total Memory 512GB (16x32GB DDR4 3200 MT/s [3200 MT/s]), BIOS SE5C620.86B.01.01.0009.2311021928, microcode 0xd0003b9, 2x Ethernet Controller E810-C for QSFP, Drives: 9x 3.5TB INTEL SSDPF2KX038TZ, 2x 54.9G INTEL SSDPEK1A058GA, VMware vSphere 8.0U2, build 22380479, Ubuntu Server 22.04.3 VM (vHW=21, vmxnet3), Kernel 5.15, intel-optimized-tensorflow:2.14, ResNet50v1.5, Batch size=128, VM=56vCPU+400GB RAM, Multi instance scenario (4 cores per instance), BERT-Large, SQuAD 1.1, Batch size=128, VM=56vCPU+400GB RAM. Test by Intel as of 12/11/23.

**5th Gen Intel® Xeon® Configuration: Intel® Xeon® Gold 6548Y+.** 1-node, 2x Intel Xeon Gold 6548Y+, 32 cores @ 2.5GHz, HT On, Turbo On, NUMA 2, Integrated Accelerators Available [used]: DLB 2 [0], DSA 2 [0], IAA 2 [0], QAT 2 [0], Total Memory 512GB (16x32GB DDR5 5600 MT/s [5200 MT/s]), BIOS 3B05.TEL4P1, microcode 0x21000161, 2x Ethernet Controller E810-C for QSFP, Drives: 8x 3.5TB INTEL SSDPF2KX038TZ, 1x 894.3G INTEL SSDSC2KG960G8, VMware vSphere 8.0U2, build 22380479, Ubuntu Server 22.04.3 VM (vHW=21, vmxnet3), Kernel 5.15, intel-optimized-tensorflow:2.14, ResNet50v1.5, Batch size=128, VM=64vCPU+400GB RAM, Multi instance scenario (4 cores per instance), BERT-Large, SQuAD 1.1, Batch size=128, VM=64vCPU+400GB RAM. Test by Intel as of 11/24/23.





# Configuration Details – Dell AI on 5th Gen Intel® Xeon® Processors



**BASELINE on 3rd Gen Xeon Config. Intel® Xeon® Platinum 8358:** 1-node, 2x Intel Xeon Platinum 8358 CPU, 32 cores @ 2.60GHz, HT On, Turbo On, NUMA 2, Integrated Accelerators Available [used]: DLB 0 [0], DSA 0 [0], IAA 0 [0], QAT 0 [0], Total Memory 512GB (16x32GB DDR4 3200 MT/s [3200 MT/s]), BIOS 1.12.1, microcode 0xd0003b9, 2x Ethernet Controller E810-C for QSFP, 1x 447.1G DELLBOSS VD, Drives: 2x 7TB Samsung MZQL27T6HBLA-00A07, VMware vSphere 8.0U2, build 22380479, Ubuntu Server 22.04.3 VM (vHW=21, vmxnet3), Kernel 5.15, intel-optimized-tensorflow:2.14, ResNet50v1.5, Batch size=128, VM=64vCPU+64GB RAM, Multi instance scenario (4 cores per instance), BERT-Large, SQuAD 1.1, Batch size=128, VM=64vCPU+64GB RAM. Test by Intel as of 11/17/23.

**5th Gen Xeon Configuration: Intel® Xeon® Platinum 8562Y+:** 1-node, 2x Intel Xeon Platinum 8562Y+ CPU, 32 cores @ 2.80 GHz, HT On, Turbo On, NUMA 2, Integrated Accelerators Available [used]: DLB 2 [0], DSA 2 [0], IAA 2 [0], QAT 2 [0], Total Memory 1536GB (16x96GB DDR5 5600 MT/s [5600 MT/s]), BIOS 1.9.11, microcode 0x210001a0, 2x Ethernet Controller E810-C for QSFP, Drives: 2x Dell Ent NVMe PM1735a MU 6.4TB, 1x Dell Ent NVMe v2 AGN MU U.2 6.4TB, VMware vSphere 8.0U2, build 22380479, Ubuntu Server 22.04.3 VM (vHW=21, vmxnet3), Kernel 5.15, intel-optimized-tensorflow:2.14, ResNet50v1.5, Batch size=128, VM=64vCPU+64GB RAM, Multi instance scenario (4 cores per instance), BERT-Large, SQuAD 1.1, Batch size=128, VM=64vCPU+64GB RAM. Test by Intel as of 11/14/23.