

3rd Gen Intel Xeon Scalable processors



Performance made flexible

We know that you must build for your needs of today while planning for unknown future demands. Flexible infrastructure is key, one that is optimized for multi-cloud and AI, serving many workloads, anywhere, at any time.

That's why we've created a platform that delivers the flexibility you need, bringing AI everywhere, from edge to cloud, helping our customers achieve more. More performance. More efficiency. And more wonderful.

Introducing 3rd Gen Intel Xeon Scalable processors, a balanced architecture with built-in acceleration and advanced security capabilities, designed through decades of innovation for the most in-demand workload requirements.

Through partnership with the world's software leaders and solution providers, 3rd Gen Intel Xeon Scalable processors are optimized for many workload types and performance levels, all with the consistent, open Intel architecture you know and trust.

3rd Gen Intel Xeon Scalable processors are...

Optimized for cloud, enterprise, HPC, network, security, and IoT workloads with **8 to 40 powerful cores** and a wide range of frequency, feature, and power levels.

Infused with Intel Crypto Acceleration, increasing the performance of encryption-intensive workloads including SSL web serving, 5G infrastructure, and VPN/firewalls, while reducing the performance impact of pervasive encryption.



The only data center CPU with **built-in AI acceleration**, end-to-end data science tools, and an ecosystem of smart solutions.

Engineered for the demands of cloud workloads and to support a wide range of XaaS environments.

Powered by Intel SGX, which protects data and application code while in use from the edge to the data center and multi-tenant public cloud.

Outstanding gen-on-gen performance

The new 3rd Gen Xeon Scalable processors are based on a balanced, efficient architecture that increases core performance, memory, and I/O bandwidth to accelerate diverse workloads from the data center to the edge.

- Available with up to 40 powerful cores.
- Built-in workload acceleration features include Intel Deep Learning Boost, Intel Advanced Vector Extensions 512, and Intel Speed Select Technology.

Available multi-socket scale to propel insights

Select 3rd Gen Intel Xeon Scalable processors support up to 28 cores per processor in four-and-eight-socket configurations, driving enhanced performance, throughput, and CPU frequencies compared to previous-gen processors.

- Four-to-eight-socket 3rd Gen Intel Xeon Scalable processors deliver multi-socket core count density, with up to 224 cores per platform in an 8-socket configuration.
- Up to six Intel UltraPath Interconnect (Intel UPI) channels increase platform scalability and improve inter-CPU bandwidth for I/O-intensive workloads versus the previous generation, offering the perfect balance between improved throughput and energy efficiency.

Technology Overview

Built-in workload and service acceleration

Intel Deep Learning Boost (Intel DL Boost) acceleration is built-in specifically for the flexibility to run complex AI workloads on the same hardware as your existing workloads.

- **with int8** – available on all 3rd Gen Intel Xeon Scalable processors, Vector Neural Network Instructions (VNNI) enhance inference workloads by maximizing use of compute resources, improving cache utilization, and reducing potential bandwidth bottlenecks.
- **with bfloat16** – available on select 3rd Gen Intel Xeon Scalable processors, the industry's first x86 support of Brain Floating Point 16-bit (bfloat16)⁵ brings enhanced artificial intelligence inference and training performance with Intel Deep Learning Boost.

Learn more at ai.intel.com

1.46x
average gen-on-gen
performance improvement¹

Up to 1.60x
higher memory bandwidth
vs. prior gen²

Up to 2.66x
higher memory capacity
vs. prior gen³

Up to 1.33x
more PCI Express lanes
per processor
vs. prior gen⁴

Now with PCIe Gen 4

Intel Advanced Vector Extensions 512 (Intel AVX-512)


boosts performance and throughput for the most demanding computational tasks in applications such as modeling and simulation, data analytics and machine learning, data compression, visualization, and digital content creation. Take advantage of 3rd Gen Intel Xeon Scalable processors' increased memory bandwidth, improved frequency management, and 2x the FMA—now across Platinum, Gold, and Silver SKUs. Compared to Intel AVX2, Intel AVX-512 is designed to enable greater performance than ever before.

Learn more at intel.com/avx512

Intel Speed Select Technology (Intel SST) is a powerful collection of capabilities giving you fine-grain control over CPU performance that can help to optimize total cost of ownership. With Intel SST, one server can do more. Lots more.

- 3rd Gen Intel Xeon Scalable processors offer enhanced Intel SST capabilities, featuring **Intel SST – Base Frequency** (Intel SST-BF), **Intel SST – Core Power** (Intel SST-CP), and **Intel SST – Turbo Frequency** (Intel SST-TF) capabilities across most Intel Xeon Platinum and Gold processors.
- 3rd Gen Intel Xeon Scalable processors, Y SKUs, support the new **Intel SST – Performance Profile 2.0** (Intel SST-PP), with additional core count, frequency, profile, and power configuration opportunities.

Learn more at intel.com/speedselect



Customizable performance to power a range of workloads

Intel SST-PP ▪ Intel SST-BF ▪ Intel SST-CP ▪ Intel SST-TF



Security

Intel Software Guard Extensions (Intel SGX) provides fine grain data protection via application isolation in memory, independent of operating system or hardware configuration.^{6,7}

Learn more at intel.com/sgx

Intel Crypto Acceleration brings built-in Vector AES-NI, Vector CLMUL, Intel Secure Hash Algorithm Extensions, VPMADD52 instructions, and RSA/DH encryption protocols.⁵

Learn more at newsroom.intel.com/crypto

Intel Total Memory Encryption delivers full physical memory encryption support to enhance data and VM protection.^{6,7}

Intel Platform Firmware Resilience (Intel PFR) uses an Intel FPGA to protect, detect, and correct, providing NIST SP800-193 firmware resiliency. Platform firmware can be validated before execution, while runtime monitoring and filtering help protect against manipulation. In case of a compromise, Intel PFR provides automated recovery in minutes.



Flexibility

Intel Resource Director Technology (Intel RDT) delivers visibility and control over shared platform resources to optimize performance and help increase resource utilization. The latest 3rd Gen Intel Xeon Scalable processors bring improvements including the second gen of Memory

Bandwidth Allocation (MBA 2.0), now featuring an advanced hardware controller with flexible and efficient bandwidth controls, and higher-resolution counters (32b) for Memory Bandwidth Monitoring.

Learn more at intel.com/intelrdt

Overview of 3rd Gen Intel Xeon Scalable processors

Intel Xeon Platinum 8300 processors are the foundation for secure, agile, hybrid-cloud data centers. With enhanced hardware-based security and exceptional multi-socket processing performance, these processors are built for mission-critical, real-time analytics, machine learning, artificial intelligence, high performance computing (HPC), and multi-cloud workloads.⁶

With trusted, hardware-enhanced data service delivery, these CPUs deliver improvements in I/O, memory, storage, and network technologies to harness actionable insights from our increasingly data-fueled world. Designed for advanced analytics, artificial intelligence, and high-density infrastructure, Intel Xeon Platinum 8300 processors deliver exceptional performance, platform capabilities, and workload acceleration.

Learn more at intel.com/xeonscalable

With support for higher memory speeds, enhanced memory capacity, and up to four-socket scalability, **Intel Xeon Gold 6300 and 5300 processors** deliver improved performance, enhanced memory capabilities, hardware-enhanced security, and workload acceleration.⁶ These processors are optimized for demanding mainstream data center, multi-cloud compute, and network and storage workloads. With up to four-socket scalability, they are suitable for an expanded range of workloads.

Intel Xeon Silver 4300 processors deliver essential performance, improved memory speed, and power efficiency. They offer hardware-enhanced performance required for entry data center compute, network, and storage.



Features Intel SGX



Features built-in AI acceleration



Up to 40 cores
per Intel Xeon Platinum processor

Up to 8 memory channels
at up to 3200 MT/s.

Features built-in AI acceleration (Intel DL Boost) with int8 (VNNI) and bfloat16⁵ quantization for accelerated AI inference and training performance.

Features Intel SGX, **the most deployed trusted execution environment**, with up to 512GB enclave per processor.



Up to 32 cores
per Intel Xeon Gold processor

Up to 8 memory channels
at up to 3200 MT/s.

Featuring Intel SGX, **the most deployed trusted execution environment**, with standard SKUs supporting up to 64GB enclave per processor.



Up to 20 cores
per Intel Xeon Silver processor

Up to 8 memory channels
at up to 2667 MT/s.

Featuring Intel SGX, **the most deployed trusted execution environment**, with up to 8GB enclave per processor.

Portfolio Products & Technologies

Intel Optane persistent memory 200 series

Intel Optane persistent memory is a category of devices that deliver flexibility to the data center. Use it to augment memory capacity beyond what's available on DRAM-only systems or to deploy a new persistent memory tier, delivering high performance and low latency for accelerated data processing.

Intel Optane persistent memory 200 series is Intel's next-gen security-enabled, reliable, persistent memory module. It provides large capacity and native persistence to help extract more value from larger datasets and increase agility with fast access to more data closer to the CPU. Intel Optane persistent memory is supported by leading software providers, including SAP, Oracle, Microsoft, VMware, Nutanix, Citrix, Apache Spark, Aerospike, Redis, MemVerge, and more.

Four-socket 3rd Gen Intel Xeon Scalable platforms support in-memory database deployments fueled by Intel Optane PMem 200 series. Only App Direct Mode is supported on four-socket solutions.

Learn more at intel.com/optanepersistentmemory

The latest Intel Optane persistent memory 200 series delivers⁸:

Average of
32% higher memory bandwidth
vs. prior gen & platform

Up to
6 TB of memory per socket
for faster analysis of the largest datasets

Up to
16% lower power
vs. prior gen

Intel Optane SSD P5800X

Data creation continues to grow at exponential rates, and workloads are becoming ever more intense. For the hottest stored data, legacy storage options are increasingly a performance bottleneck, inhibiting the evolution of new architectures and application performance.

Caching and tiering the hottest data to higher-performance storage can break these bottlenecks. However, reliance on today's NAND SSDs for write-intensive environments risks disk wear-out, which can increase maintenance costs and downtime. The Intel Optane SSD P5800X offers the ideal combination of no-compromise I/O performance and never-before-seen SSD endurance to deliver unprecedented storage value.

Learn more at intel.com/optane

Compared to the previous generation, the Intel Optane SSD P5800X enables unprecedented storage value by delivering up to⁹:

- 60% lower latency**
for faster actionable insights
- 50% better Quality of Service (QoS)**
to enable monetization of improved SLAs
- 4K greater random 4K mixed read/write IOPs**
to better saturate high-speed networks
- 67% higher endurance**
to help extend the life of lower-endurance NAND SSDs



Intel SSD D5-P5316 (NAND)


Architected with the industry's first 144-layer quad-level cell (QLC) NAND and PCIe Gen 4, the Intel SSD D5-P5316 enables significant TCO savings while accelerating warm storage.

Built from the ground up to optimize and accelerate storage, the Intel SSD D5-P5316, with Intel's most advanced QLC NAND technology, delivers industry-leading SSD storage density for the data center and the high bandwidth of PCIe Gen 4. Compared to HDDs, the Intel SSD D5-P5316 can accelerate access to stored data by up to 25x, while performance optimization brings up to 2x higher sequential read performance, up to 38% higher random read, and up to 5x higher endurance compared to previous-gen NAND-based Intel SSDs.¹⁰

The innovative form factor also enables up to 1 petabyte of storage in just 1U of rack space, allowing massive storage consolidation for lower TCO. In fact, it facilitates up to a 20x reduction in storage footprint in the data center.¹⁰ Well-suited for a wide range of read-intensive, low-latency workloads, the Intel SSD D5-P5316 is ideal for content delivery network (CDN), hyperconverged infrastructure (HCI), big data, artificial intelligence (AI), cloud elastic storage (CES), and HPC workloads.

Learn more at intel.com/3dnand

- Unlock the value of stored data
- Firmware enhancements for drive performance, IT efficiency, data security, and manageability
- NAND technology industry leader



Connectivity


Intel Ethernet 800 series adapters now offer port data rates from 10 Gbps to 100 Gbps, supporting both PCIe Gen 3 and Gen 4 across a variety of port counts to meet the needs of almost every workload.

- Our **Application Device Queues (ADQ)** feature prioritizes application traffic to help deliver the performance required for high-priority, network-intensive workloads.
- **Enhanced Dynamic Device Personalization (DDP)** uses the fully programmable pipeline to enable frame classification for advanced and proprietary protocols on the adapter—to increase throughput, lower latency, and reduce host CPU overhead.
- Supports **RDMA over iWARP** or **RoCEv2** protocols and **NVMe over TCP** with **ADQ** for high-throughput, low-latency cloud, storage, and HPC clusters.

The newest member of the adapter family, the **Intel Ethernet Network Adapter E810-2CQDA2** increases network data throughput up to 200 Gbps per adapter for high-performance vRAN, NFV forwarding plane, storage, HPC, cloud, and CDN.

Learn more at intel.com/ethernet

Up to **2x**
increase in resources
for virtualized and containerized networks
compared to Intel Ethernet 700 Series





Customer Segment Benefits

5G performance made flexible

5G use cases come in all shapes and sizes—so should your infrastructure. Workloads are unique and need to be delivered across multiple locations, meaning that as your network evolves, the technology must too.

Network-optimized 3rd Gen Intel Xeon Scalable processors, N SKUs, are designed to support diverse network environments. Optimized for many workloads and performance levels, they are available in a wide range of cores, frequencies, features, and power. For organizations

ready to drive 5G networking to the next level, these CPUs increase 5G user plane function (UPF) performance by up to 1.42x vs. the prior gen.¹

Through partnership with a proven, broad ecosystem, including 400+ Intel Network Builders members, Intel is delivering solution blueprints based on 3rd Gen Intel Xeon Scalable processors, N SKUs—all resulting in accelerated qualification and shortened time-to-deployment for vRAN, NFVI, virtual CDN, and more.

Key uses and workloads

vRAN



5G UPF



VPP IPsec




vBNG



CDN



DPDK/L3F Forwarding



- **Network-optimized N SKUs** deliver low latency and high throughput.
- **Enhanced Intel AVX-512 & Crypto processing.**
- Robust and proven ecosystem to ensure time to deployment.
- **Intel Crypto Acceleration** increases encrypted network workload performance and virtually eliminates the performance impact of full encryption for normal workloads.
- With **up to 36 latest-gen cores**, these CPUs offer **higher base frequency** for greater throughput for virtualized network functions and **lower power** for dense or constrained physical deployments.
- Accelerate networking infrastructure deployment time with Intel Select Solutions, our OEM partners, and leading cloud providers.
- Multiple SKUs available with increased supply availability of up to 7 years.

Learn more at intel.com/networking

Your 5G advantage¹

1.62x
 average performance improvement
 across network and communication workloads
 vs. prior gen

2x
 massive MIMO throughput
 in a similar power envelope for a best-in-class
 3x100MHz 64T64R configuration

Up to **1.76x**
 enhanced DPDK L3 forwarding
 vs. prior gen

Up to **1.63x**
 increased CDN throughput,
 enabling you to serve the same number of users at a higher
 resolution or more subscribers at the same resolution
 vs. prior gen



AI performance made flexible

Whether you're making product recommendations, optimizing your supply chain, or sequencing genomes, AI is integral to your digital future. But delivering AI into production from edge to cloud has been challenging, with infrastructure considerations, software and application compatibility, and fragmented tools top of mind. Until now.

As the only data center CPU with built-in AI acceleration, hardware-enhanced security, and software optimizations, 3rd Gen Intel Xeon Scalable processors can help deliver higher performance per TCO across the most diverse set of smart AI workloads.

Through Intel's optimizations to streamline popular end-to-end data science tools, powered by the open oneAPI standard, 3rd Gen Intel Xeon Scalable processors make it easier and faster for data practitioners to build and widely deploy smarter models—and simpler to move smoothly from PoC to production.

Choose from a broad selection of pre-integrated and verified enterprise solutions for data analytics and AI. Combined with the largest enterprise databases and a broad range of applications ready to unlock your insights, accelerate your time to AI deployment with 3rd Gen Intel Xeon Scalable processors.

Key uses and workloads

Classical machine learning



Natural language processing



Recommender system



High-res computer vision



Federated learning



- **Intel DL Boost** delivers higher performance with built-in AI.
- **Intel SGX** helps improve security with multi-party computing through federated learning.
- Use built-in **Intel Crypto Acceleration** to virtually eliminate the performance impact of full data encryption and dramatically increase the performance of encryption-intensive workloads.
- **Software optimizations** on industry standard framework help boost performance.

Learn more at ai.intel.com

Your AI advantage¹

Up to **1.74x**
more AI inference performance
with enhanced Intel Deep Learning Boost
vs. prior gen



Security made flexible

The connected world can expose valuable data assets to theft or tampering. Privacy concerns and regulations can make some collaboration opportunities impossible with today's solutions, which protect data integrity in storage or on a network, but can leave data exposed while in use.

The latest 3rd Gen Intel Xeon Scalable processors—delivering a **revolutionary step forward** for privacy and security—include Intel SGX, which protects data and application code while in use, enabling new ways to collaborate using shared data without compromising privacy. With the **smallest attack surface** within the system, Intel SGX is already proven in countless research studies and production deployments by many leading security organizations.

- Intel SGX creates secured enclaves in memory where data and application code can be verified, protected, and kept private while being processed.
- Strengthen today's security workloads, such as key management systems, and open new possibilities of trusted multi-party compute and collaboration while keeping each party's data private.
- Keep data and code isolated and better protected from intruders, malicious software, and even your cloud provider.

Key uses and workloads

Trusted foundation for all workloads



Cloud confidential computing



Multi-party compute



Enterprise blockchain



Key management



- **Built-in set of security capabilities** address current and future privacy and security concerns customers face.
- **Intel SGX** delivers the most proven and deployed Trusted Executive Environment (TEE) for the data center, now with huge memory enclaves.
- New **Intel Crypto Acceleration** reduces CPU overhead for better encryption performance and SLA on encryption security.
- New **Intel Total Memory Encryption (Intel TME)** encrypts full system memory for added protection against physical attacks.
- Support for up to 512 GB Intel SGX enclave capacity, per CPU, available on multiple SKUs. See SKU table on Page 13 for more details.

Learn more at intel.com/sgx

Your security advantage¹

Up to **1.48x**

faster encryption performance

with Intel Crypto Acceleration

vs. prior gen



Cloud and enterprise made flexible

You benefit when your cloud infrastructure is flexible enough to serve many workloads. Deploying on an optimized platform tuned to your multi-cloud needs and delivered with broad ecosystem support can help you scale your solution.

Intel enables a simpler path and seamless migration to trusted cloud deployments and speeds the deployment of proven, repeatable cloud use cases in most industries, built from deep, varied industry expertise and collaboration.

The latest 3rd Gen Intel Xeon Scalable processors are built upon years of cloud innovation, tuned to your unique needs with built-in AI and Intel Crypto Acceleration and advanced security capabilities.

Built on open standards and APIs, with fully optimized software and security from hardware up the stack, 3rd Gen Intel Xeon Scalable processors scale up and out with ease. As a result, you have flexibility to evolve your infrastructure for changing needs, streamlining cost and enhancing data governance.

All of the world's largest cloud service providers plan to offer services powered by 3rd Gen Intel Xeon Scalable processors, ranging from IaaS, e-commerce, and content distribution, to social media, confidential computing, and more.

Key uses and workloads

Relational database



Analytics



AI



Web services



Virtual desktop (VDI, WVD)



Security (confidential compute)



- **Built-in acceleration** technologies and advanced security to enable a safer remote workforce.
- **Accelerate deployment** with ISV partnerships.
- Interoperability is tailored to power **mission-critical workloads** with consistency across the ecosystem.
- Intel delivers **disruptive performance efficiency** across workloads, **agile, trusted infrastructure**, and **differentiated services** to grow new markets.
- Intel SGX helps deliver privacy and security in cloud solutions on the market today, like Microsoft Azure Confidential Compute, Alibaba Cloud, Baidu MesaTEE, and IBM Cloud Data Shield.

P SKU: IaaS optimized for orchestration efficiency in high-frequency VM environments

V SKU: SaaS optimized SKU for orchestration efficiency in high-density, lower-power VM environments

Y SKU: Intel Speed Select Technology – Performance Profile

Learn more at intel.com/cloud and intel.com/enterprise

Your cloud & enterprise advantage¹

Up to **1.58x**
higher performance
on cloud-based microservices
vs. prior gen

Up to **1.72x**
higher virtualization performance
vs. prior gen

Up to **1.64x**
more database transactions per minute
vs. prior gen

High performance computing (HPC) made flexible

3rd Gen Intel Xeon Scalable processors are designed to power new business breakthroughs and scientific discoveries through outstanding performance and flexibility in an efficient, future-forward server infrastructure that supports the convergence of HPC with AI in the cloud or on-premises.

Enhanced core architecture, memory bandwidth, and security capabilities deliver faster time to results for HPC customers, while amazing performance, both gen-on-gen and versus competition, powers a wide range of diverse and challenging HPC applications.

New Q SKUs, optimized for liquid-cooled systems, provide HPC customers with max core count and frequencies delivering enhanced top-bin performance. Memory-bound workloads are fueled with faster I/O, featuring increased memory bandwidth, more capacity with 8 channels, and 64 lanes of PCIe Gen 4 per CPU.

Built-in advantages such as Intel DL Boost and enhanced Intel AVX-512 deliver AI and HPC convergence and workload performance, while new Intel Speed Select Technology capabilities empower HPC customers to tailor performance for their specific workload needs.

Key uses and workloads

Life sciences



Manufacturing



Finance



Cloud



AI



Security
(confidential compute)



- Core count and frequency flexibility.
- 8 DDR4 memory channels per CPU at up to 3200 MT/s.
- Up to 64 lanes PCI Express Gen 4 per socket to enable higher I/O bandwidth per core.
- Added improvements for performance per core and performance per watt.
- Accelerate deployment time for HPC workloads with Intel Select Solutions from our OEM partners, and leading cloud providers.

Q SKU: Liquid cooled for high performance

Learn more at intel.com/hpc

Your HPC advantage¹

Up to **1.53x**
higher HPC performance
vs. prior gen



Internet of Things (IoT) made flexible

Compute-intensive workloads such as AI and video analytics are headed to the edge and IoT, where processing can take place closer to the data source or point of action. The need for performance, security, and manageability continues to grow as the diversity, sophistication, and intelligence of edge and IoT deployments accelerates.

The newest 3rd Gen Intel Xeon Scalable processors deliver the performance, security, and operational controls required for powerful AI, complex image or video analytics, and consolidated workloads at the edge, on-prem, or wherever work needs to be done.

This processor family offers enhanced gen-over-gen performance and built-in AI acceleration for faster analytics, processing of more images or video streams, and more powerful AI in edge and IoT deployments. Built-in hardware-based security protects critical code and private data from tampering or interception by malware or hackers, a particular challenge in dispersed edge, industrial, or IoT deployments.

Key uses and workloads

Video



Healthcare



Industrial manufacturing



Public sector



Retail, banking, hospitality, and education



- Greater throughput and built-in AI acceleration for faster analytics, more video streams, and more powerful AI in edge and IoT.
- Increased performance and more cores along with increased memory bandwidth³ allows for faster object recognition analysis on multiple video streams, simultaneously, for vision usages.
- Built-in Intel Deep Learning Boost can accelerate deep learning workloads such as image classification, object detection and more. In healthcare deployment this can augment clinical workflows and help provide diagnosis.
- Enhanced per-core performance enables convergence of workloads and more virtualized desktops onto single or fewer backend servers for retail, banking, and education deployments.
- Long-life availability to meet the specific support needs of the edge, industrial, and embedded segment.

T SKU = Up to 10-year usability on select SKUs, support for higher case temperature (TC_{case})

Learn more at intel.com/iot

Your IoT advantage¹

Up to **1.56x**
improvement in AI inference
for image classification
with enhanced Intel DL Boost
vs. prior gen

Hardened security and long-life availability
to meet the specific needs of edge and embedded systems

3rd Gen Intel Xeon Scalable processor SKUs

For the most up-to-date information, please visit intel.com/xeon or ark.intel.com

SKU	Cores	Base (GHz)	Single Core Turbo (GHz)	All Core Turbo (GHz)	Cache (MB)	TDP (Watts)	Support for Intel Optane Persistent Memory 200 Series	Intel SGX Enclave Capacity (Per Processor)	Recommended Customer Pricing (RCP) in \$ US Dollars
OPTIMIZED FOR HIGHEST PER-CORE SCALABLE PERFORMANCE									
8380	40	2.3	3.4	3.0	60	270	Yes	512 GB	\$8099
8368	38	2.4	3.4	3.2	57	270	Yes	512 GB	\$6302
8362	32	2.8	3.6	3.5	48	265	Yes	64 GB	\$5448
8360Y	36	2.4	3.5	3.1	54	250	Yes	64 GB	\$4702
8358	32	2.6	3.4	3.3	48	250	Yes	64 GB	\$3950
6354	18	3.0	3.6	3.6	39	205	Yes	64 GB	\$2445
6348	28	2.6	3.5	3.4	42	235	Yes	64 GB	\$3072
6346	16	3.1	3.6	3.6	36	205	Yes	64 GB	\$2300
6342	24	2.8	3.5	3.3	36	230	Yes	64 GB	\$2529
6334	8	3.6	3.7	3.6	18	165	Yes	64 GB	\$2214
6326	16	2.9	3.5	3.3	24	185	Yes	64 GB	\$1300
5317	12	3.0	3.6	3.4	18	150	Yes	64 GB	\$950
5315Y	8	3.2	3.6	3.5	12	140	Yes	64 GB	\$895
SCALABLE PERFORMANCE									
8352Y	32	2.2	3.4	2.8	48	205	Yes	64 GB	\$3450
6338	32	2.0	3.2	2.6	48	205	Yes	64 GB	\$2612
6336Y	24	2.4	3.6	3.0	36	185	Yes	64 GB	\$1977
6330	28	2.0	3.1	2.6	42	205	Yes	64 GB	\$1894
5320	26	2.2	3.4	2.8	39	185	Yes	64 GB	\$1555
5318Y	24	2.1	3.4	2.6	36	165	Yes	64 GB	\$1273
4316	20	2.3	3.4	2.8	30	150		8 GB	\$1002
4314	16	2.4	3.4	2.9	24	135	Yes	8 GB	\$694
4310	12	2.1	3.3	2.7	18	120		8 GB	\$501
4309Y	8	2.8	3.6	3.4	12	105		8 GB	\$501

Y Supports Intel Speed Select Technology – Performance Profile 2.0 (Intel SST-PP). All 8300, 6300, 5300, and 4300 SKUs, non-H/HL SKUs, are supported on a unique 1 or 2 socket platform. Please contact your hardware provider for a list of system availability supporting your specific SKU configuration.

All 8300, 6300, 5300, and 4300 processors, non-H/HL SKUs, are configured to support up to 6 TB of system memory, per processor. Intel has validated for up to 4 TB of Intel Optane persistent memory 200 series, per processor. Intel has validated for up to 256 GB capacity DRAM memory modules, as of March 2021.

Unless noted, all 8300, 6300, and 5300 processors, non-H/HL SKUs, include support for Intel Speed Select Technology (Intel SST), featuring Intel SST – Base Frequency (SST-BF), Intel SST – Core Power (SST-CP) and Intel SST – Turbo Frequency (SST-TF) capabilities.

8362 does not support Intel Speed Select Technology – Base Frequency (SST-BF).

SKUS SUPPORTING MAXIMUM INTEL SGX ENCLAVE CAPACITY									
8380	40	2.3	3.4	3.0	60	270	Yes	512 GB	\$8099
8368Q	38	2.6	3.7	3.3	57	270	Yes	512 GB	\$6743
8368	38	2.4	3.4	3.2	57	270	Yes	512 GB	\$6302
8352S	32	2.2	3.4	2.8	48	205	Yes	512 GB	\$4046
5318S	24	2.1	3.4	2.6	36	165	Yes	512 GB	\$1667

8352S and 5318S support Intel Speed Select Technology – Performance Profile 2.0 (Intel SST-PP).

Product Brief | 3rd Gen Intel Xeon Scalable processors

SKU	Cores	Base (GHz)	Single Core Turbo (GHz)	All Core Turbo (GHz)	Cache (MB)	TDP (Watts)	Support for Intel Optane Persistent Memory 200 Series	Intel SGX Enclave Capacity (Per Processor)	Recommended Customer Pricing (RCP) in \$ US Dollars
CLOUD OPTIMIZED FOR VM UTILIZATION									
8358P	32	2.6	3.4	3.2	48	240	Yes	8 GB	\$3950
8352V	36	2.1	3.5	2.5	54	195	Yes	8 GB	\$3450

P IaaS Cloud Specialized Processor, V SaaS Cloud Specialized Processor.

8352V supports Intel Speed Select Technology – Performance Profile 2.0 (Intel SST-PP).

LIQUID COOLED									
8368Q	38	2.6	3.7	3.3	57	270	Yes	512 GB	\$6743

8368Q supports up to 512 GB Intel Software Guard Extensions (Intel SGX) enclave capacity.

NETWORKING/NFV OPTIMIZED									
8351N	36	2.4	3.5	3.1	54	225	Yes	64 GB	\$3027
6338N	32	2.2	3.5	2.7	48	185	Yes	64 GB	\$2795
6330N	28	2.2	3.4	2.6	42	165	Yes	64 GB	\$2029
5318N	24	2.1	3.4	2.7	36	150	Yes	64 GB	\$1375

8351N is supported in a one-socket configuration only.

5318N supports Intel Speed Select Technology – Performance Profile 2.0 (Intel SST-PP).

MEDIA PROCESSING OPTIMIZED									
8352M	32	2.3	3.5	2.8	48	185	Yes	64 GB	\$3864

Optimized for processing AI and media workloads and services.

LONG-LIFE USE AND NEBS-THERMAL FRIENDLY									
6338T	24	2.1	3.4	2.7	36	165	Yes	64 GB	\$2742
5320T	20	2.3	3.5	2.9	30	150	Yes	64 GB	\$1727
4310T	10	2.3	3.4	2.9	15	105		8 GB	\$555

Support for up to 10-year reliability, higher Tcase.

SINGLE-SOCKET OPTIMIZED									
8351N	36	2.4	3.5	3.1	54	225	Yes	64 GB	\$3027
6314U	32	2.3	3.4	2.9	48	205	Yes	64 GB	\$2600
6312U	24	2.4	3.6	3.1	36	185	Yes	64 GB	\$1450

Supported in one-socket configurations only.

SKU	Cores	Base (GHz)	Single Core Turbo (GHz)	All Core Turbo (GHz)	Cache (MB)	TDP (Watts)	Support for Intel Optane Persistent Memory 200 Series	Recommended Customer Pricing (RCP) in \$ US Dollars	
FOUR & EIGHT SOCKET SCALABLE PERFORMANCE									
8380HL	28	2.9	4.3	3.8	38.5	250	Yes	\$13012	
8380H	28	2.9	4.3	3.8	38.5	250	Yes	\$10009	
8376HL	28	2.6	4.3	3.5	38.5	205	Yes	\$11772	
8376H	28	2.6	4.3	3.5	38.5	205	Yes	\$8719	
8360HL	24	3.0	4.2	3.8	33	225	Yes	\$7203	
8360H	24	3.0	4.2	3.8	33	225	Yes	\$4200	
8356H	8	3.9	4.4	4.3	35.75	190	Yes	\$3400	
8354H	18	3.1	4.3	4.0	24.75	205	Yes	\$3500	
8353H	18	2.5	3.8	3.3	24.75	150	Yes	\$3003	

SKU	Cores	Base (GHz)	Single Core Turbo (GHz)	All Core Turbo (GHz)	Cache (MB)	TDP (Watts)	Support for Intel Optane Persistent Memory 200 Series	Recommended Customer Pricing (RCP) in \$ US Dollars
FOUR & EIGHT SOCKET SCALABLE PERFORMANCE (CONT.)								
6348H	24	2.3	4.2	3.1	33	165	Yes	\$2700
6330H	24	2.0	3.7	2.8	33	150	Yes	\$1894
6328HL	16	2.8	4.3	3.7	22	165	Yes	\$4779
6328H	16	2.8	4.3	3.7	22	165	Yes	\$1776
5320H	20	2.4	4.2	3.3	27.5	150	Yes	\$1555
5318H	18	2.5	3.8	3.3	24.75	150	Yes	\$1273

H and HL SKUs are only supported on a unique 4 or 8-socket platform. Please contact your hardware provider for a list of system availability supporting your specific SKU configuration.

H SKUs are configured to support up to 1.12 TB of system memory, per processor.

HL SKUs are configured to support up to 4.5 TB of system memory, per processor.

H and HL SKUs are validated for up to 256 GB capacity DRAM memory modules, as of March 2021.

H and HL SKUs support Intel Optane persistent memory 200 series in a 4-socket platform only.

H SKUs are validated for up to 768 GB of Intel Optane persistent memory 200 series, per processor.

HL SKUs are validated for up to 3 TB of Intel Optane persistent memory 200 series, per processor.

6330H, 6328H, 6328HL & 5320H SKUs include Intel Speed Select Technology (Intel SST), supporting Intel SST – Core Power (SST-CP) and Intel SST – Turbo Frequency (SST-TF) capabilities.

Features Supported

One and Two Socket Supported Platforms

Not a comprehensive list of features or capabilities. Subject to change without notice. Visit [intel.com/xeon](https://www.intel.com/xeon) for additional details or contact your Intel representative.

	Intel Xeon Silver 4300 Processors	Intel Xeon Gold 5300 Processors	Intel Xeon Gold 6300 Processors	Intel Xeon Platinum 8300 Processors
Performance				
Highest Core Count Supported	20 cores	24 cores	32 cores	40 cores
Lowest Core Count Supported	8 cores	8 cores	8 cores	8 cores
Highest Supported Turbo Frequency	3.4 GHz	3.4 GHz	3.6 GHz	3.7 GHz
Highest Supported Base Frequency	2.8 GHz	3.2 GHz	3.6 GHz	2.8 GHz
Number of CPU Sockets Supported	Up to 2	Up to 2	Up to 2	Up to 2
Intel Ultra Path Interconnect (Intel UPI) and Intel UPI Speed	2 UPI at 10.4 GT/s	3 UPI at 11.2 GT/s	3 UPI at 11.2 GT/s	3 UPI at 11.2 GT/s
Highest Memory Speed Support (DDR4)	2667 MT/s	2933 MT/s	3200 MT/s	3200 MT/s
Highest Memory Capacity Supported per Socket	6 TB	6 TB	6 TB	6 TB
Intel Optane Persistent Memory 200 Series Module Support	4314 only	•	•	•
PCI Express Gen 4 (64 lanes per socket)	•	•	•	•
Intel Turbo Boost Technology 2.0	•	•	•	•
Intel Hyper-Threading Technology (Intel HT Technology)	•	•	•	•
Intel Deep Learning Boost (Intel DL Boost) with Vector Neural Network Instructions (VNNI)	•	•	•	•
Intel Advanced Vector Extensions 512 (Intel AVX-512)	2 FMA	2 FMA	2 FMA	2 FMA
Intel QuickAssist Technology supported in select Intel C620 series chipsets and discrete PCIe cards	•	•	•	•
Support for Intel Optane SSDs and Intel SSDs (3D NAND)	•	•	•	•

	Intel Xeon Silver 4300 Processors	Intel Xeon Gold 5300 Processors	Intel Xeon Gold 6300 Processors	Intel Xeon Platinum 8300 Processors
Reliability				
Reliability, Availability, and Serviceability (RAS) Capability	Standard	Advanced	Advanced	Advanced
Intel Run Sure Technology		•	•	•
Agility & Efficiency				
Intel Speed Select Technology (Intel SST) featuring Performance Profile 2.0 (Intel SST-PP) supported	4309Y only	5315Y, 5318Y, 5318N, 5318S	6336Y only	8352S, 8352Y, 8352V, 8360Y
Intel Speed Select Technology (Intel SST) featuring Intel SST Base Frequency (SST-BF), Intel SST Core Power (SST-CP)*, and Intel SST Turbo Frequency (SST-TF) capabilities		•	•	•
Intel Infrastructure Management Technologies (Intel IMT)	•	•	•	•
Intel Resource Director Technology (Intel RDT)	•	•	•	•
Intel Volume Management Device (Intel VMD)	•	•	•	•
Intel Virtualization Technology (Intel VT)	•	•	•	•
Intel Speed Shift Technology	•	•	•	•
Intel Node Manager 4.0	•	•	•	•
Security				
Intel Software Guard Extensions (Intel SGX)	•	•	•	•
Standard Intel Software Guard Extensions (Intel SGX) enclave capacity supported	Up to 8GB	Up to 64GB	Up to 64GB	Up to 64GB
Maximum Intel Software Guard Extensions (Intel SGX), up to 512GB, enclave capacity supported		5318S only		8352S, 8368, 8368Q, 8380
Intel Crypto Acceleration	•	•	•	•
Intel Total Memory Encryption (Intel TME)	•	•	•	•
Intel Platform Firmware Resilience (Intel PFR)	•	•	•	•
Intel Advanced Vector Extensions 512 (Intel AVX-512)	•	•	•	•
Intel Key Protection Technology (Intel KPT) with Integrated Intel QAT	•	•	•	•
Intel Platform Trust Technology (Intel PTT)	•	•	•	•
Intel Trusted Execution Technology (Intel TXT) with One-Touch Activation (OTA)	•	•	•	•

*8362 SKU does not support Intel Speed Select Technology Base Frequency (SST-BF).

Features Supported

Four and Eight Socket Supported Platforms

Not a comprehensive list of features or capabilities. Subject to change without notice. Visit intel.com/xeon for additional details or contact your Intel representative.

	Intel Xeon Gold 5300H Processors	Intel Xeon Gold 6300 HL & H Processors	Intel Xeon 8300 HL & H Processors
Performance			
Highest Core Count Supported	20 cores	24 cores	28 cores
Lowest Core Count Supported	12 cores	16 cores	8 cores
Highest Supported Turbo Frequency	4.2 GHz	4.3 GHz	4.4 GHz
Highest Supported Base Frequency	2.5 GHz	2.8 GHz	3.9 GHz
Number of CPU Sockets Supported	Up to 4	Up to 4	Up to 8
Intel Ultra Path Interconnect (Intel UPI) and Intel UPI Speed	6 UPI at 10.4 GT/s	6 UPI at 10.4 GT/s	6 UPI at 10.4 GT/s
Highest Memory Speed Support (DDR4)	2933 MT/s	2933 MT/s	3200 MT/s (1 DPC) 2933 MT/s (2 DPC)
Highest Memory Capacity Supported per Socket	1.12 TB	4.5 TB (HL), 1.12 TB (H)	4.5 TB (HL), 1.12 TB (H)
Intel Optane Persistent Memory 200 Series Module Support (Supported in 4 socket platforms only)	•	•	•
Highest Memory Capacity Supported per Socket (DRAM + Optane PMem 200 Series)	4.5 TB	4.5 TB	4.5 TB
PCI Express 3 (48 lanes per socket)	•	•	•
Intel Turbo Boost Technology 2.0	•	•	•
Intel Hyper-Threading Technology (Intel HT Technology)	•	•	•
Intel Deep Learning Boost (Intel DL Boost) with Vector Neural Network Instructions (VNNI) and Brain Floating Point 16-bit (bfloat16) numeric format processing	•	•	•
Intel Advanced Vector Extensions 512 (Intel AVX-512)	2 FMA	FMA	2 FMA
Intel QuickAssist Technology supported in select Intel C620 series chipsets and discrete PCIe cards	•	•	•
Support for Intel Optane SSDs and Intel SSDs (3D NAND)	•	•	•
Reliability			
Reliability, Availability, and Serviceability (RAS) Capability	Standard	Advanced	Advanced
Intel Run Sure Technology	•	•	•
Agility & Efficiency			
Intel Speed Select Technology (Intel SST) featuring Performance Profile 2.0 (Intel SST-PP) supported	5320H only	6330H, 6328H, 6328HL only	
Intel Speed Select technology (Intel SST) supporting Intel SST Core Power (SST-CP) and Intel SST Turbo Frequency (SST-TF) capabilities	•	•	•
Intel Infrastructure Management Technologies (Intel IMT)	•	•	•
Intel Resource Director Technology (Intel RDT)	•	•	•
Intel Volume Management Device (Intel VMD)	•	•	•
Intel Virtualization Technology (Intel VT)	•	•	•
Intel Speed Shift Technology	•	•	•
Intel Node Manager 4.0	•	•	•

Intel Xeon Gold 5300H Processors

Intel Xeon Gold 6300 HL & H Processors

Intel Xeon 8300 HL & H Processors

Security			
Support for Intel Platform Firmware Resilience (Intel PFR)	•	•	•
Intel Security Libraries for Data Center (Intel ISeCL-DC)	•	•	•
Intel Advanced Vector Extensions 512 (Intel AVX-512)	•	•	•
Intel Key Protection Technology (KPT) with Integrated Intel QAT	•	•	•
Intel Platform Trust Technology (PTT)	•	•	•
Intel Trusted Execution Technology (Intel TXT) with One-Touch Activation (OTA)	•	•	•

Product Name	SKU	Compression	Encryption	RSA	Max PCIe Uplink	Rec. Min Uplink Config	PCIe Uplink x8 Optional Muxed Link	TDP (W)	Est Reduced TDP4 (W)
		Intel QuickAssist Technology							
Intel C621A Chipset	LBG-1G	N/A	N/A	N/A	X1 (NA)	X1 (NA)	disabled	15	10
Intel C627A Chipset	LBG-T	~65 Gb/s	100 Gb/s	100K Ops/s	x16	x16	enabled	28.6	26
Intel C629A Chipset	LBG-C	~75-80 Gb/s	100 Gb/s	N/A	x16	X16	Enabled	28.6	26



Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. 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 optimizations, for Intel compilers or other products, may not optimize to the same degree for non-Intel products.

Intel may make changes to specifications and product descriptions at any time, without notice.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. All processors support Intel Virtualization Technology (Intel VT-x).

Performance varies by use, configuration and other factors.

1. Please visit www.intel.com/3gen-xeon-config and use the corresponding performance number [#] to access full system configuration and performance detail.

- Process up to 1.64x more database transactions per minutes vs. prior gen **[81]**
- Up to 1.72x higher virtualization performance vs. prior gen **[84]**
- 1.62x average performance improvement across network and communications workloads vs. prior gen **[91]**
- 2x Massive MIMO throughput in a similar power envelope for a best-in-class 3x100MHz 64T64R configuration **[91]**
- 1.76x enhanced DPDK L3 forwarding vs. prior gen **[91]**
- Up to 1.63x increased throughput enabling you to serve the same number of users at a higher resolution or a greater number of subscribers at the same resolution vs. prior gen **[91]**
- Increase 5G UPF performance by 1.42x vs. prior gen **[91]**
- Up to 1.48x faster encryption performance with Intel Crypto Acceleration vs. prior gen **[97]**
- 1.58x higher performance on cloud-based microservices vs. prior gen **[98]**
- 1.53x higher HPC performance vs. prior gen **[108]**
- 1.56x improvement in AI inference for image classification with enhanced Intel Deep Learning Boost vs. prior gen **[119]**
- Up to 1.74x more AI inference performance with enhanced Intel Deep Learning Boost vs. prior gen **[120]**
- 1.46x average gen-on-gen performance improvement **[125]**

2. 3rd Gen Intel Xeon Platinum 8380 CPU: 8 channels, 3200 MT/s (2 DPC) vs. 2nd Gen Intel Xeon Platinum 8280 CPU: 6 channels, 2666 MT/s (2 DPC).

3. 3rd Gen Intel Xeon Platinum 8380 CPU: 8 channels, 2 DPC (256GB DDR4) vs. 2nd Gen Intel Xeon Platinum 8280 CPU: 6 channels, 2 DPC (128GB DDR4).

4. 3rd Gen Intel Xeon Platinum 8380 CPU: 64 lanes of PCI Express 4 per processor vs. 2nd Gen Intel Xeon Platinum 8280 CPU: 48 lanes of PCI Express 3 per processor.

5. bfloat16 is supported only on 4S and 8S Intel Xeon Scalable processors.

6. No product or component can be absolutely secure.

7. Intel Optane persistent memory does not work with Intel SGX—see intel.com for additional details on security interoperability with Advanced RAS and Intel Optane PMem.

8. Based on testing by Intel as of April 27, 2020 (Baseline) and March 23, 2021 (New).

Baseline configuration: 1-node, 1 x Intel Xeon Platinum 8280L processor (28 cores at 2.7 GHz) on Neon City with a single Intel Optane PMem module configuration (6 x 32 GB DRAM; 1 x {128 GB, 256 GB, 512 GB} Intel Optane PMem module), ucode rev: 04002F00 running Fedora 29 kernel 5.1.18-200.fc29.x86_64 and Intel Memory Latency Checker (Intel MLC) version 3.8 with App Direct Mode.

New Configuration: 1-node, 1 x Intel Xeon pre-production ICX-XCC processor (38 cores at 2.0 GHz) on Wilson City with a single Intel Optane PMem module configuration (8 x 32 GB DRAM; 1 x {128 GB, 256 GB, 512 GB} Intel Optane PMem module), ucode rev: 8d000270 running RHEL 8.1 kernel 4.18.0-147.el8.x86_64 and Intel MLC version 3.9 with App Direct Mode.

9. Source – Intel; Date tested – March 18, 2021

Workload – FIO rev 3.5, based on random 512B transfer size with total queue depth of 64 (QD=8, workers/jobs=8) workload, 4KB transfer size with total queue depth of 32 (QD=4, workers/jobs=8) workload, 8KB transfer size with total queue depth of 16 (QD=4, workers/jobs=4) workload in most case, except where specified.

System configuration

Intel Optane SSD P5800X 1.6TB: CPU: Intel Xeon Platinum 8380 2.30GHz 270W 40 cores per socket, CPU Sockets: 2, BIOS: SE5C6200.86B.3021.D40.2103160200, UCODE: 0X8D05A260, RAM: 32GB @3200 MT/s DDR4, DIMM Slots Populated: 16 slots, PCIe Attach: CPU (not PCH lane attach), OS: Ubuntu 20.04.2 LTS, Kernel: 5.4.0-67-generic, FIO version: 3.16; NVMe Driver: Inbox, C-states: Disabled, Hyper Threading: Disabled, CPU Governor (through OS): Performance Mode, Intel Turbo Mode, and P-states = Disabled; IRQ Balancing Services (OS) = Off; SMP Affinity, set in the OS; FIO with ioengine=io_uring.

See Intel Optane SSD DC P4800X product specifications at <https://ark.intel.com/content/www/us/en/ark/products/97161/intel-optane-ssd-dc-p4800x-series-375gb-2-5in-pcie-x4-3d-xpoint.html>

10. Test and System Configuration: Mainboard: Intel Server Board S2600WFT, Version: R2208WFTZS, BIOS: SE5C620.86B.00.01.0014.070920180847, Platform architecture: x86_64, CPU: Intel Xeon Gold 6140 CPU @ 2.30GHz, CPU Sockets: 2, RAM Capacity: 32G, RAM Model: DDR4, OS version: centos-release-7-5, Build id: 1804, kernel: 4.14.74, NVMe Driver: Inbox, Fio version: 3.5, G4SAC PCIe Gen4 switch PCIe card (Microsemi). Intel SSD D5-P5316 was tested on ACV10100 firmware.

Up to 25x accelerated access: Sequential read performance of Intel SSD D5-P5316 compared to Seagate Exos X18 (seagate.com/files/www-content/datasheets/pdfs/exos-x18-channel-DS2045-1-2007GB-en_SG.pdf).

Up to 38% higher random read: Comparing 4K random read between Intel SSD D5-P5316 15.36TB (800K IOPS) and Intel SSD DC P4326 15.36TB (580K IOPS).

Up to 2x higher sequential read: Comparing 128K sequential read bandwidth between Intel SSD D5-P5316 15.36TB (7.0 GB/s) and Intel SSD DC P4326 15.36TB (3.2 GB/s).

Up to 5x higher endurance gen over gen: Comparing endurance (64K random write) between Intel SSD D5-P5316 30.72TB (22,930 TBW) and Intel SSD DC P4326 15.36TB (4,400 TBW).

Up to 20x reduction of warm storage footprint: With 4TB HDD drive, it takes 10 (2U) of rack space to fill up 1PB or storage. With Intel SSD D5-P5316 30.72TB E1.L or U.2, it takes 1U of rack space to fill up 1PB of storage—which equals up to 20x greater rack consolidation.

© 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.