


PRODUCT BRIEF

Intel® Xeon® Scalable Platform



SECOND GENERATION

Intel® Xeon® Scalable Processors

The Foundation for Data-Centric Innovation



Empowering Transformation in a Data-Centric Era

Across an evolving digital world, disruptive and emerging technology trends in business, industry, science, and entertainment increasingly impact the world's economies. Future workloads will necessitate infrastructures that can seamlessly scale to support immediate responsiveness and diverse performance requirements. The exponential growth of data generation and consumption, the rapid expansion of cloud-scale computing, emerging 5G networks, and the extension of high-performance computing (HPC) and artificial intelligence (AI) into new usages require that today's data centers and networks urgently evolve—or be left behind in a highly competitive environment. These demands are driving the architecture of modernized, future-ready data centers and networks that can quickly flex and scale.

The Intel® Xeon® Scalable platform provides the foundation for a powerful data center platform that creates an evolutionary leap in agility and scalability. Disruptive by design, this innovative processor sets a new level of platform convergence and capabilities across compute, storage, memory, network, and security. Enterprises and cloud and communications service providers can now drive forward their most ambitious digital initiatives with a feature-rich, highly versatile platform.

Enabling Greater Efficiencies and Lower TCO

Across infrastructures, from enterprise to technical computing applications, the Intel Xeon Scalable platform is designed for data center modernization to drive operational efficiencies that lead to improved total cost of ownership (TCO) and higher productivity for users. Systems built on the Intel Xeon Scalable platform are designed to deliver agile services with enhanced performance and groundbreaking capabilities compared to the prior generation.

Performance to Propel Insights

Intel's industry leading, workload-optimized platform with built-in AI acceleration provides the seamless performance foundation for the data-centric era from the multi-cloud to intelligent edge and back. The Intel Xeon Scalable platform with 2nd gen Intel® Xeon® Scalable processors enables a new level of consistent, pervasive, and breakthrough performance.

INTEL® XEON® PLATINUM 9200 PROCESSORS

2X
AVERAGE PERFORMANCE IMPROVEMENT¹
Compared to Intel® Xeon® Platinum 8180 Processor

UP TO 30X
AI PERFORMANCE WITH INTEL® DL BOOST²
Compared to Intel® Xeon® Platinum 8180 Processor

2ND GEN INTEL® XEON® SCALABLE PROCESSORS

UP TO 3.50X
5-YEAR REFRESH PERFORMANCE IMPROVEMENT³
VM density compared to Intel® Xeon® E5-2600 V2 Processor

UP TO 14X
AI PERFORMANCE WITH INTEL® DL BOOST⁴
Compared to Intel® Xeon® Platinum 8180 Processor

2nd Gen Intel® Xeon® Processor provides consistent performance for diverse workloads

	WHAT MATTERS?	INTEL® XEON® SCALABLE PLATFORM ADVANTAGES
Enterprise and Cloud	<ul style="list-style-type: none"> Minimize complexity with compatible virtualization infrastructure Meet stringent customer SLAs 	<ul style="list-style-type: none"> Deploy quickly. Intel VMs coexist w/ other Intel® technology-based servers. Fast response times.
HPC	<ul style="list-style-type: none"> Maximize vector floating performance and efficiency 	<ul style="list-style-type: none"> High performance with fewer servers.
Storage	<ul style="list-style-type: none"> Ensure deterministic storage response 	<ul style="list-style-type: none"> Deterministic performance. Cores, cache, memory, I/O in a single die.
Communications	<ul style="list-style-type: none"> Deliver diverse services efficiently 	<ul style="list-style-type: none"> Incredible efficiency and hardware acceleration from a platform with features required for application, control, packet, and signal processing.

Support for Breakthrough Memory Innovation

A new foundation for performance begins with support for Intel's breakthrough Intel® Optane™ persistent memory, a new class of memory and storage innovation architected for data centric class environments. With individual module capacities up to 512 GB, Intel Optane persistent memory can deliver up-to 36 TB of system-level memory capacity when combine with traditional DRAM. Intel Optane persistent memory complements DRAM memory by affordably enabling unprecedented system memory capacity to accelerate workload processing and service delivery.

Learn more about Intel Optane persistent memory intel.com/optanedcpersistentmemory



2ND GEN INTEL® XEON® SCALABLE PROCESSORS WITH INTEL® OPTANE™ PERSISTENT MEMORY

SAP **13X FASTER RESTART⁵**

Microsoft

UP TO 36% MORE VMs⁶

UP TO 25% LOWER COST PER VM⁶

Intel Optane persistent memory pricing & DRAM pricing referenced in TCO calculations is provided for guidance and planning purposes only and does not constitute a final offer. Pricing guidance is subject to change and may revise up or down based on market dynamics. Please contact your OEM/distributor for actual pricing.

Foundational Enhancements

- **Higher Per-Core Performance:** Up to 56 cores (9200 series) and up to 28 cores (8200 series), delivering high performance and scalability for compute-intensive workloads across compute, storage, and network usages.
- **Greater Memory Bandwidth/Capacity:** Support for Intel Optane persistent memory, supporting up-to 36 TB of system-level memory capacity when combine with traditional DRAM. 50 percent increased memory bandwidth and capacity. Six memory channels versus four memory channels of previous generation for memory-intensive workloads. Support for up-to 4 TB of DDR4 memory, per socket, with speeds up-to 2933 MT/s (1 DPC).
- **Expanded I/O:** 48 lanes of PCIe* 3.0 bandwidth and throughput for demanding I/O-intensive workloads.
- **Intel® Ultra Path Interconnect (Intel® UPI):** Four Intel® UPI (9200 series) and up to three Intel UPI (8200 series) channels increase scalability of the platform to as many as two sockets (9200 series) and up to eight sockets (8200 series), as well as improves inter-CPU bandwidth for I/O-intensive workloads versus previous generation (with Intel® Quick Path Interconnect). Intel UPI offers the perfect balance between improved throughput and energy efficiency.
- **Intel® Deep Learning Boost (Intel® DL Boost) with VNNI:** New Intel® Deep Learning Boost with Vector Neural Network Instructions (VNNI) bring enhanced artificial intelligence inference performance, with up to 30X performance improvement over the previous generation⁴, 2nd gen Intel Xeon Scalable processors help to deliver AI readiness across the data center, to the edge and back.
- **Intel® Infrastructure Management Technologies (Intel® IMT):** A framework for resource management, Intel® Infrastructure Management Technologies (Intel® IMT), combines multiple Intel capabilities that support platform-level detection, reporting, and configuration. This hardware-enhanced monitoring, management, and control of resources can help enable greater data center resource efficiency and utilization.
- **Intel® Security Libraries^o for Data Center (Intel® ISeCL-DC):** A set of software libraries and components, Intel SecL-DC enables Intel hardware-based security features. The open-source libraries are modular and have a consistent interface. They can be used by customers and software developers to more easily develop solutions that help secure platforms and help protect data using Intel hardware-enhanced security features at cloud scale.

- **Intel® Advanced Vector Extensions 512 (Intel® AVX-512):** With double the flops per clock cycle compared to previous generation Intel® AVX2, 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.
- **Security^o without compromise:** Limiting encryption overhead and performance on all secure data transactions.

Innovative Integrations

Platform integrations deliver improvements in performance and latency across the infrastructure:

- **Integrated Intel® QuickAssist Technology (Intel® QAT):** Chipset-based hardware acceleration for growing compression and cryptographic workloads for greater efficiency while delivering enhanced data transport and protection across server, storage, and network infrastructure.
- **Integrated Intel® Ethernet with scalable iWARP* RDMA*:** Provides up to four 10 Gbps high-speed Ethernet ports for high data throughput and low-latency workloads. Ideal for software-defined storage solutions, NVM Express over Fabric solutions, and virtual machine migrations. Integrated in the chipset.

Industry-Leading Memory and Storage Support

Storage innovations can drive significant improvements in efficiency and performance of data-hungry workloads.

- **Support for Intel Optane persistent memory:** Breakthrough memory and storage memory innovation offering groundbreaking capabilities for fast storage solutions. Can be combined with Intel® Optane™ SSDs for the ultimate in storage and data performance.
- **Support for Intel Optane SSDs and Intel® QLC 3D NAND Solid State Drives:** Delivers industry-leading combination of high throughput, low latency, high QoS, and ultra-high endurance to break through data access bottlenecks.
- **Deploy next generation storage with confidence with Intel® Volume Management Device (Intel® VMD):** Enables hot swapping of NVMe SSDs from the PCIe bus without shutting down the system, while standardized LED management helps provide quicker identification of SSD status. This commonality brings enterprise reliability, availability, and serviceability (RAS) features to NVMe SSDs, enabling deployment of next-generation storage with confidence.
- **Intel® Intelligent Storage Acceleration Library (Intel® ISA-L):** Optimizes storage operations, such as encryption, for increased storage performance.

Complementary Offerings for Even Greater Performance, Scalability

Intel offers a broad hardware and software portfolio that complements this new processor.

- Intel® Ethernet 800 series products support up to 100 GbE port speed with Application Device queues (ADQ), which addresses latency-sensitive workloads for higher speed data communication. Data Plane Developer Kit (DPDK) is supported across Intel Ethernet 800 Series products for NFV acceleration, advanced packet forwarding and highly efficient packet processing.

[Learn more at intel.com/ethernet](https://www.intel.com/ethernet)

- Intel® FPGAs offer flexible, programmable acceleration, for low-latency applications, such as virtual switching, network services, data analytics, and AI.

[Learn more at intel.com/fpga](https://www.intel.com/fpga)

- A range of software tools and libraries for general and highly parallel computing help developers optimize applications for Intel® Architecture.

[Learn more at software.intel.com](https://www.software.intel.com)

Enhanced Platform Trust

Data and platform reliability and protection are key concerns for enterprises dealing with increasing concerns and scrutiny regarding data security and privacy. Intel Xeon Scalable platform helps build highly trusted infrastructures with platform data protection, resiliency, and uptime.

Increased Data Protection and Reliability Across Every Workload

- **Enhanced Intel® Run Sure Technology:** New enhancements deliver advanced Reliability, Availability, and Serviceability (RAS) and server uptime for a company's most critical workloads. Hardware-assisted capabilities, including enhanced MCA and recovery and adaptive multi-device error correction, diagnose and recover from previously fatal errors. And, they help ensure data integrity within the memory subsystem.
- **Intel® Key Protection Technology (Intel® KPT) with Integrated Intel QAT and Intel® Platform Trust Technology (Intel® PTT):** Deliver hardware-enhanced platform security[◊] by providing efficient key and data protection at rest, in-use, and in-flight.
- **Intel® Trusted Execution Technology (Intel® TXT) with One-Touch Activation:** Enhanced platform security[◊] while providing simplified and scalable deployment for Intel TXT.

As more data-rich workloads flow through the data center, this comprehensive suite of hardware-enhanced features brings better data- and platform-level protection mechanisms for trusted services in enterprise and cloud environments.[◊]

Dynamic and Highly Efficient Service Delivery

The convergence of enhanced compute, memory, network, and storage performance, combined with software ecosystem optimizations, make Intel Xeon Scalable platform the ideal platform for fully virtualized, software-defined data centers that dynamically self-provision resources—on-premise, through the network, and in the public cloud—based on workload needs.

Powerful Tools and Technologies for an Agile Data Center

Intel® Virtualization Technology (Intel® VT-x) features:

- **Mode based execution control (MBE) virtualization:** Provides an extra layer of protection from malware attacks in a virtualized environment by enabling hypervisors to more reliably verify and enforce the integrity of kernel level code.
- **Timestamp counter scaling (TSC) virtualization:** Provides workload optimization in hybrid cloud environments by allowing virtual machines to move across CPUs operating at different base frequencies.

Intel® Node Manager 4.0: Helps IT intelligently manage and optimize power, cooling, and compute resources in the data center, maximizing efficiency, while reducing the chances of costly overheats.

Faster Time to Value with Intel® Select Solutions

In today's complex data center, hardware and software infrastructure is not "one size fits all." Intel® Select Solutions eliminates guesswork with rigorously benchmark tested and verified solutions optimized for real-world performance. These solutions accelerate infrastructure deployment on Intel® Xeon® processors for today's critical workloads in advanced analytics, hybrid cloud, storage, and networking.

Enterprise and Government – Primed for Business

For enterprise data centers modernizing to take advantage of the era of advanced analytics, the hybrid cloud and future-ready storage, Intel Select Solutions can speed up your data-fueled, IT-driven business transformation.

Communications Service Providers – Tuned Network Enhancements

For Communication Service Providers transforming their network for a 5G-enabled future, Intel Select Solutions offer a faster and more efficient deployment path of tested, reliable infrastructure with verified configurations that take full advantage of virtual network enhancements that support new and emerging customer workload demands.

High Performance Computing – Accelerated Time to Insight

For research in academia and government as well as the enterprises, high performance computing (HPC) capabilities with Intel Select Solutions help push the limits of mainstream data today with deeper insights and more complex problem solving.

Learn more about Intel Select Solutions featuring new 2nd gen Intel Xeon Scalable processors with Intel Optane persistent memory at intel.com/selectsolutions



ARTIFICIAL INTELLIGENCE	ANALYTICS	CLOUD SERVICES	HCI / STORAGE	NETWORK / EDGE	HPC
BigDL on Apache Spark	Microsoft SQL Server* (Windows Server*, Linux)*	Microsoft Azure Stack*	VMware vSAN*	Universal Customer Premises Equipment (Ubuntu*, Adva* Ensemble, Cent OS* + RT)	Simulation & Modeling
AI Inferencing	SAP HANA	Red Hat OpenShift* Container Platform	Microsoft Azure* Stack HCI	NFVI (Red Hat, Ubuntu, FusionSphere*)	Simulation & Visualization
SECURITY	PingCAP* TiDB	Google Cloud's Anthos*	Huawei FusionStorage*	NFVI Forwarding Platform	Genomics Analytics
	GBASE	QingCloud*	Nutanix HCI* <small>COMING SOON</small>	Visual cloud (Delivery network, Media analytics)	HPC & AI Converged Clusters (Magpie, Univa)
	Blockchain: Hyperledger Fabric	Easystack* Cloud Service	xSky		
Hardened Security With Lockheed Martin		Open Cloud			

Strong, Capable Platforms for the Data-Fueled Enterprise

Enterprises are keen to extract value from the exploding data streams being presented to them for rapid insights that can shape their business initiatives. Traditional and emerging applications in the enterprise, including predictive analytics, machine learning, and HPC, require new levels of powerful compute capabilities and massive tiered data storage volumes. The modernized data center is being architected using a converged and holistic approach that can flexibly deliver new services and improve TCO across infrastructure assets today, while providing the most seamless and scalable on-ramp to a self-governing, hybrid data center.

Yet, organizations running their foundational business workloads, such as OLTP and web infrastructure, seek to reduce TCO with higher performing infrastructures.

The Intel Xeon Scalable platform delivers next-generation enterprise capabilities to businesses through a future-ready platform that can serve the hybrid-cloud, data-fueled era, plus it helps improve day-to-day operations. This versatile platform brings disruptive levels of compute performance, coupled with memory and I/O advances, to compute-hungry and latency-sensitive applications. Combined with innovative Intel Optane SSDs and Intel QLC 3D NAND SSD Data Center Family to manage large data volumes across storage, caching, and memory, platforms built on the Intel Xeon Scalable platform are ready to handle the intense demands of the data and cloud era.

With a scalable portfolio of packaging options to suit diverse workload requirements, the Intel Xeon Scalable platform is a performance workhorse designed for deploying highly efficient, virtualized infrastructures for compute, storage, and networking.

Highlights for Enterprise Innovation

- 2nd gen Intel Xeon Scalable Processors
- Intel Optane Persistent Memory
- Intel Deep Learning Boost
- Intel® Speed Select Technology
- Intel Ethernet 800 Series
- Intel Optane SSDs and Intel QLC 3D NAND SSDs
- Intel Infrastructure Management Technologies

2ND GEN INTEL® XEON® SCALABLE PROCESSORS WITH INTEL® OPTANE™ PERSISTENT MEMORY

OLTP DATABASE

Hammer DB



2nd Gen
Intel® Xeon® Platinum 8280

UP TO **3.7X**

HIGHER PERFORMANCE VS
5-YEAR SYSTEM⁷

Intel® Xeon® Processor E5-2697 V2

BIG DATA

Big Bench (3 TB)



2nd Gen
Intel® Xeon® Gold 6248

UP TO **2.3X**

FASTER TIME TO INSIGHTS VS
5-YEAR SYSTEM⁸

Intel® Xeon® Processor E5-2697 V2

Hi Bench



2nd Gen
Intel® Xeon® Gold 6248

UP TO **4.3X**

HIGHER PERFORMANCE VS
5-YEAR SYSTEM⁹

Intel® Xeon® Processor E5-2697 V2



Next-Generation Platform for Cloud-Optimized, 5G-Ready Networks, and Next-Generation Virtual Networks

The coming era of 5G will enable entirely new ecosystems and classes of consumer and enterprise services along with media applications on wireless and wireline networks. These data-rich, innovative use cases, driven by the new Internet of Things (IoT), visual computing, and analytics, represent significant future opportunities for communications service providers (CommSPs) to grow revenue.

The transition from purpose-built, fixed function infrastructure to a new generation of open networks is the essential first step to prepare for a 5G-enabled world. Software-defined networking with Network Functions Virtualization (NFV) is enabling new service opportunities and operations efficiencies for both communication services providers and enterprises alike. Using flexible, optimized, industry-standard servers and virtualized, orchestrated network functions will allow future-ready infrastructures to be able to deliver innovative services with efficiency and ease.

Such distributed communications networks can support extreme levels of scalability, agility, programmability, and security across an ever-growing volume and variety of networking workloads—from the network core to the edge.

The Intel Xeon Scalable platform is the basis for next-generation platforms to build virtualized, cloud-optimized, 5G-ready networks. It offers an architecture that scales and adapts with ease to handle the demands of emerging applications and the convergence of key workloads, such as applications and services, control plane processing, high-performance packet processing, and signal processing. This new processor provides a foundation for agile networks that can operate with cloud economics, be highly automated and responsive, and support rapid and more secure delivery of new and enhanced services enabled by 5G.

Highlights for Communication Service Provider Innovation

- 2nd gen Intel Xeon Scalable processor “N” SKUs, specialized for Networking/NFV

- Intel Optane persistent memory
- Hardware-based acceleration of encryption and compression using integrated Intel QAT
- Intel Ethernet 800 series
- Intel FPGAs maximize versatility in communications infrastructure
- Intel Infrastructure Management Technologies

Additional Resources Optimized for Communication Service Providers

The open source Data Plane Development Kit (DPDK) enables optimized communications operations on Intel Architecture. DPDK has demonstrated ability to scale performance as processor core count and performance increase; workloads, such as Vector Packet Processing (VPP) IPsec, benefit from this enhanced performance. Additionally, these libraries provide pre-optimized mechanisms to allow new processor capabilities (such as Intel AVX-512 and memory and I/O enhancements) to be able to utilize the new functionality for improved packet processing performance with less direct development effort.

Intel offers programs, such as Intel® Network Builders University, ideal for network evolution in the 5G era. With solution guidance and training from these programs, CommSPs can drive their network transformation initiatives forward with increased confidence.

2ND GEN INTEL® XEON® SCALABLE PROCESSORS SPECIALIZED FOR NETWORKING/NFV “N” SKUS

UP TO **1.25-1.58X** NFV WORKLOAD PERFORMANCE IMPROVEMENT¹⁰

Comparing Intel® Xeon® Gold 6230 N processor to Intel® Xeon® Gold 6130 processor

Additional processing headroom delivers enhanced VM/VNF capacity and density

Handle more subscriber capacity and mobile services

Reduced bottlenecks for fixed and mobile 5G networks

Featuring Intel® Select Speed Technology

New in 2019 “N”-Based NFV solutions



Breakthrough HPC and High-Performance Data Analytics Innovation

Today's scientific discoveries are fueled by innovative algorithms, new sources and volumes of data, and advances in compute and storage. Benefitting from exponentially expanding volumes and variety of data, HPC clusters are also the engine for running evolving High-Performance Data Analytics (HPDA) workloads, leading to incredible discoveries and insight for business and human understanding. Machine learning, deep learning, and AI converge the capabilities of massive compute with the flood of data to drive next-generation applications, such as autonomous systems and self-driving vehicles.

Intel Xeon Scalable platform offers a common platform for AI with high throughput for both inference and training—up to 30x higher inference² using 9200 series and up to 14x higher inference⁴ throughput using 8200 series, compared to Intel Xeon Scalable processors introduced in July 2017.

HPC is no longer just the domain of large scientific institutions. Enterprises are increasingly consuming a massive number of HPC compute cycles; some of the world's largest HPC clusters are in private oil and gas companies. Research in personalized medicine applies HPC for highly focused treatment plans. New HPC installations are engaging innovative, converged architectures for non-traditional usages that combine simulation, AI, visualization, and analytics in a single supercomputer.

HPC platforms—from the smallest clusters to largest supercomputers—demand a balance across compute, memory, storage, and network. The Intel Xeon Scalable platform was designed to deliver and enable such balance with massive scalability—to tens of thousands of cores. From its improved core count and mesh architecture to newly integrated technologies and support for Intel Optane persistent memory and storage devices, the Intel Xeon Scalable platform enables the ultimate goals of HPC—to maximize performance across compute, memory, storage, and network without inducing bottlenecks at any intersection of resources.

The integration of Intel® Omni-Path Architecture, an end-to-end high-performance fabric, into the Intel Xeon Scalable platform delivers both increased performance and scaling to distributed, parallel computing clusters. Near linear scaling up to 32 nodes enables building large HPC solutions that are not inhibited by the interconnect. The Intel Xeon Scalable platform and Intel Omni-Path Architecture can enable new discoveries and faster solutions for highly parallel workloads in many data centers.

- Intel® Xeon® Platinum 9200 Processors
- Intel Optane Persistent Memory
- Intel Ultra Path Interconnect
- Intel Deep Learning Boost
- Intel Advanced Vector Extensions 512
- Intel Omni-Path Architecture Host Fabric Interface
- Intel Optane SSDs

Additional Technologies for HPC, HPDA, and AI

- A range of high-productivity software tools, optimized libraries, foundational building blocks, and flexible frameworks for general and highly parallel computing help simplify workflows and assist developers to create codes that maximize the capabilities of IA for HPC and AI.
- Optimizations for popular deep learning frameworks for IA, including Neon,* Caffe,* Theano,* Torch,* and TensorFlow* offer increased value and performance for data scientists.
- Intel® Parallel Studio XE 2017 includes performance libraries, such as Intel® Math Kernel Library for Deep Neural Networks (Intel® MKL-DNN) to accelerate deep learning frameworks on IA, and Intel® Data Analytics Acceleration Library (Intel® DAAL) to speed big data analytics.

Resources Optimized for HPC

To continue to advance discovery through HPC into the Exascale era, the Intel® Modern Code Developer Program offers developers and data scientists easily accessible online and face-to-face code modernization technical sessions on techniques, such as vectorization, memory and data layout, multi-threading, and multi-node programming.

2ND GEN INTEL® XEON® SCALABLE PROCESSORS FOR HPC

UP TO **1.7X** BETTER FLOATING POINT PER CORE¹¹

UP TO **2X** SYSTEM MEMORY CAPACITY
Compared to Intel® Xeon® Platinum 8180 processors

Overview of 2nd Gen Intel® Xeon® Scalable Processors

Intel® Xeon® Scalable Platform

INTEL® XEON® PLATINUM 9200 PROCESSORS



Designed for high performance computing, advanced artificial intelligence and analytics, the Intel® Xeon® Platinum 9000 processors deliver breakthrough levels of performance with the highest Intel® Architecture FLOPS per rack, along with the highest DDR native memory bandwidth support of any Intel® Xeon® processor platform.

- Up to 56 Intel® Xeon® Scalable processing cores per processor
 - Two processors per 2U platform (Intel® Server System S9200WK Data Center Block)
- 12 memory channels per processors, 24 memory channels per platform
- Features new Intel® Deep Learning Boost instructions for enhanced AI inference acceleration and performance
- Enhanced multi-chip package optimized for density and performance

INTEL® XEON® PLATINUM 8200 PROCESSORS



Second generation Intel® Xeon® Platinum 8200 processors are the foundation for secure, agile, hybrid-cloud data centers. With enhanced hardware-based security and exceptional two, four, and eight+ socket processing performance, these processors are built for mission-critical, real-time analytics, machine learning, artificial intelligence and multi-cloud workloads. With trusted, hardware-enhanced data service delivery, this processor family delivers monumental leaps in I/O, memory, storage, and network technologies to harness actionable insights from our increasingly data-fueled world.

INTEL® XEON® GOLD 6200 AND INTEL® XEON® GOLD 5200 PROCESSORS



With support for the higher memory speeds, enhanced memory capacity, and four-socket scalability, the Intel® Xeon® Gold 6200 processors deliver significant improvement in performance, advanced reliability, and hardware-enhanced security. It is optimized for demanding mainstream data center, multi-cloud compute, and network and storage workloads. The Intel® Xeon® Gold 5200 processors deliver improved performance with affordable advanced reliability and hardware-enhanced security. With up-to four-socket scalability, it is suitable for an expanded range of workloads.

INTEL® XEON® SILVER 4200 PROCESSORS



Intel® Xeon® Silver processors deliver essential performance, improved memory speed, and power efficiency. Hardware-enhanced performance required for entry data center computes, network, and storage.

ENTRY-LEVEL PERFORMANCE AND HW-ENHANCED SECURITY



The Intel® Xeon® Bronze processors delivers entry performance for small business and basic storage servers. Hardware-enhanced reliability, availability, and serviceability features designed to meet the needs of these entry solutions.

Delivering Enhanced Performance and Value

The newest 2nd gen Intel Xeon Scalable processors deliver continued innovation, performance, and value to the data center with refreshed product design and production. The new selection of processors offers peak frequencies for high-performance workloads, enhanced performance for mainstream usages, and Increased value and capability for single-socket, entry, edge, networking, and IoT servers—all delivered at similar or lower pricing.[†]

	Intel® Xeon® Bronze Processor (3200 Series)	Intel® Xeon® Silver Processor (4200 Series)	Intel® Xeon® Gold Processor (5200 Series)	Intel® Xeon® Gold Processor (6200 Series)	Intel® Xeon® Platinum Processor (8200 Series)	Intel® Xeon® Platinum Processor (9200 Series)
PERVASIVE PERFORMANCE AND SECURITY						
Highest Core Count Supported	8 cores	16 cores	24 cores	28 cores	28 cores	56 cores
Highest Supported Frequency	1.9 GHz	3.5 GHz	4.0 GHz	4.5 GHz	4.0 GHz	3.8 GHz
Number of CPU Sockets Supported	Up to 2	Up to 2	Up to 4	Up to 4	Up to 8	Up to 2
Intel® Ultra Path Interconnect (UPI)	2	2	2	3	3	4
Intel® UPI Speed	9.6 GT/s	9.6 GT/s	10.4 GT/s	10.4 GT/s	10.4 GT/s	10.4 GT/s
Intel® Advanced Vector Extensions 512 (AVX-512)	1 FMA	1 FMA	1 FMA	2 FMA	2 FMA	2 FMA
Highest Memory Speed Support (DDR4)	2133 MT/s	2400 MT/s	2666 MT/s	2933 MT/s	2933 MT/s	2933 MT/s
Highest Memory Capacity Supported Per Socket ^o	1 TB	1 TB	1 TB, 4.5 TB	1 TB, 4.5 TB	1 TB, 4.5 TB	1.5 TB
16 Gb DDR4 DIMM Support	•	•	•	•	•	•
Intel® Deep Learning Boost (Intel® DL Boost) with Vector Neural Network Instructions (VNNI)	•	•	•	•	•	•
Intel® Optane™ Persistent Memory Module Support ^o		•	•	•	•	
Intel® Omni-Path Architecture (Discrete PCIe* card)	•	•	•	•	•	•
Intel® QuickAssist Technology (Integrated in chipset)	•	•	•	•	•	
Intel® QuickAssist Technology (Discrete PCIe card)	•	•	•	•	•	•
Intel® Optane™ SSDs	•	•	•	•	•	•
Intel® DC SSD Data Center Family (3D NAND)	•	•	•	•	•	•
PCIe 3.0 (48 lanes)	•	•	•	•	•	•
Intel® QuickData Technology (CBDMA)	•	•	•	•	•	•
Non-Transparent Bridge (NTB)	•	•	•	•	•	•
Intel® Turbo Boost Technology 2.0		•	•	•	•	•
Intel® Hyper-Threading Technology (Intel® HT Technology)		•	•	•	•	•
Node Controller Support				•	•	•
^o Supported on select processors only.						
HIGH RELIABILITY						
Reliability, Availability, and Serviceability (RAS) Capability	Standard	Standard	Standard	Advanced	Advanced	Advanced
Intel® Run Sure Technology				•	•	•

[†] Comparing all Intel Xeon Gold 62XXR processors to existing Intel Xeon Gold 62XX processors.


	Intel® Xeon® Bronze Processor (3200 Series)	Intel® Xeon® Silver Processor (4200 Series)	Intel® Xeon® Gold Processor (5200 Series)	Intel® Xeon® Gold Processor (6200 Series)	Intel® Xeon® Platinum Processor (8200 Series)	Intel® Xeon® Platinum Processor (9200 Series)
AGILITY & EFFICIENCY						
Intel® Speed Select Technology (Intel® SST) ^o		•	•	•	•	
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	•	•	•	•	•	•
Mode-Based Execute Control	•	•	•	•	•	•
Timestamp Counter Scaling (TSC) for Virtualization	•	•	•	•	•	•
^o Supported on select processors only.						
SECURITY						
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)	•	•	•	•	•	•

2nd Gen Intel Xeon Scalable Processors

SKU Numbering

Processor numbers for the Intel Xeon Scalable platform use an alphanumeric scheme based on performance, features, processor generation, and any options, following the brand and its class.

Intel® Xeon® Platinum	9	2	#	#	α	α	processor
Intel® Xeon® Platinum	8	2	#	#	α	α	processor
Intel® Xeon® Gold	6	2	#	#	α	α	processor
Intel® Xeon® Gold	5	2	#	#	α	α	processor
Intel® Xeon® Silver	4	2	#	#	α	α	processor
Intel® Xeon® Bronze	3	2	#	#	α	α	processor



2nd Gen Intel® Xeon® Scalable Processor SKUs

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

Processor Identifier	Core Count	Cache (MB)	TDP (Watts)	Processor Base Frequency (GHz)	Max Turbo Frequency Rate (GHz)	Max Memory Speed (MT/s)	Memory Capacities
PLATINUM 9282	56	77	400	2.6	3.8	2933	3 TB
PLATINUM 9242	48	71.5	350	2.3	3.8	2933	3 TB
PLATINUM 9222	32	71.5	250	2.3	3.7	2933	3 TB
PLATINUM 8280	28	38.5	205	2.7	4.0	2933	1 TB / 4.5 TB
PLATINUM 8276	28	38.5	165	2.2	4.0	2933	1 TB / 4.5 TB
PLATINUM 8270	26	35.75	205	2.7	4.0	2933	1 TB
PLATINUM 8268	24	35.75	205	2.9	3.9	2933	1 TB
PLATINUM 8260	24	35.75	165	2.4	3.9	2933	1 TB / 4.5 TB
PLATINUM 8256	4	16.5	105	3.8	3.9	2933	1 TB
PLATINUM 8253	16	22	125	2.2	3.0	2933	1 TB
GOLD 6258R	28	38.5	205	2.7	4.0	2933	1 TB
GOLD 6256	12	33	205	3.6	4.5	2933	1 TB
GOLD 6254	18	24.75	200	3.1	4.0	2933	1 TB
GOLD 6252	24	35.75	150	2.1	3.7	2933	1 TB
GOLD 6250	8	35.75	185	3.9	4.5	2933	1 TB / 4.5 TB ^o
GOLD 6248R	24	35.75	205	3.0	4.0	2933	1 TB
GOLD 6248	20	27.5	150	2.5	4.0	2933	1 TB
GOLD 6246R	16	35.75	205	3.4	4.1	2933	1 TB
GOLD 6246	12	24.75	165	3.3	4.2	2933	1 TB
GOLD 6244	8	24.75	150	3.6	4.2	2933	1 TB
GOLD 6242R	20	35.75	205	3.1	4.2	2933	1 TB
GOLD 6242	16	22	150	2.8	3.9	2933	1 TB
GOLD 6240R	24	35.75	165	2.4	4.0	2933	1 TB / 4.5 TB
GOLD 6240	18	24.75	150	2.6	3.9	2933	1 TB / 4.5 TB
GOLD 6238R	28	38.5	165	2.2	4.0	2933	1 TB / 4.5 TB
GOLD 6238	22	30.25	140	2.1	3.7	2933	1 TB / 4.5 TB
GOLD 6234	8	24.75	130	3.3	4.0	2933	1 TB
GOLD 6230R	26	35.75	150	2.1	4.0	2933	1 TB
GOLD 6230	20	27.5	125	2.1	3.0	2933	1 TB
GOLD 6226R	16	22	150	2.9	3.9	2933	1 TB
GOLD 6226	12	19.25	125	2.9	3.9	2933	1 TB
GOLD 5222	4	16.5	105	3.8	3.9	2933	1 TB
GOLD 5220R	24	35.75	150	2.2	4.0	2667	1 TB
GOLD 5220	18	24.75	125	2.2	3.9	2667	1 TB
GOLD 5218R	20	27.5	125	2.1	4.0	2667	1 TB
GOLD 5218	16	22	125	2.3	3.9	2667	1 TB
GOLD 5217	8	11	115	3.0	3.7	2667	1 TB
GOLD 5215	10	13.75	85	2.5	3.4	2667	1 TB / 4.5 TB
SILVER 4216	8	11	130	3.2	3.0	2400	1 TB
SILVER 4215R	8	11	130	3.2	4.0	2400	1 TB
SILVER 4215	8	11	85	2.5	3.5	2400	1 TB
SILVER 4214R	12	16.5	100	2.4	3.5	2400	1 TB
SILVER 4214	12	16.5	85	2.2	3.2	2400	1 TB
SILVER 4210R	10	13.75	100	2.4	3.2	2400	1 TB
SILVER 4210	10	13.75	85	2.2	3.2	2400	1 TB
SILVER 4208	8	11	85	2.1	3.2	2400	1 TB

^o Additional large-memory SKU available on 6250L only.

2nd Gen Intel® Xeon® Scalable Processor SKUs, continued

Processor Identifier	Core Count	Cache (MB)	TDP (Watts)	Processor Base Frequency (GHz)	Max Turbo Frequency Rate (GHz)	Max Memory Speed (MT/s)	Memory Capacities
BRONZE 3206R	8	11	85	1.9	1.9	2133	1 TB
BRONZE 3204	6	8.25	85	1.9	1.9	2133	1 TB
NETWORKING/NFV SPECIALIZED FEATURING INTEL® SPEED SELECT TECHNOLOGY SKUS							
6252N	24	35.75	150	2.3	3.6	2933	1 TB
6230N	20	27.5	125	2.3	3.9	2933	1 TB
5218N	16	22	105	2.3	3.7	2667	1 TB
SEARCH APPLICATION VALUE SPECIALIZED SKU							
5220S	18	24.75	125	2.7	3.9	2667	1 TB
LONG-LIFE CYCLE AND NEBS-THERMAL FRIENDLY SKUS							
6238T	22	30.25	125	1.9	3.7	2933	1 TB
6230T	20	27.5	125	2.1	3.9	2933	1 TB
5220T	18	24.75	105	1.9	3.9	2667	1 TB
5218T	16	22	105	2.1	3.8	2667	1 TB
4210T	10	13.75	95	2.3	3.2	2400	1 TB
4209T	8	11	70	2.2	3.2	2400	1 TB
SINGLE-SOCKET VALUE SPECIALIZED SKUS							
6212U	24	35.75	164	2.4	3.9	2933	1 TB
6210U	20	27.5	150	2.5	3.9	2933	1 TB
6209U	20	27.5	125	2.1	3.9	2933	1 TB
6208U	16	22	150	2.9	3.9	2933	1 TB
VM DENSITY VALUE SPECIALIZED SKUS							
6262V	24	33	135	1.9	3.6	2400	1 TB
6222V	20	27.5	115	1.8	3.6	2400	1 TB
SKUS FEATURING INTEL® SPEED SELECT TECHNOLOGY							
8260Y	24	35.75	165	2.4	3.9	2933	1 TB
6240Y	18	24.75	150	2.6	3.9	2933	1 TB
4214Y	12	16.5	85	2.2	3.2	2400	1 TB

See [intel.com/products/processor_number](https://www.intel.com/products/processor_number) for details.

Processor details, features, cost, and availability are subject to change without notice.

Please visit [intel.com/xeon](https://www.intel.com/xeon) for the latest product information.

Product Name	SKU	10 Gb/1 Gb Ethernet Ports	Compression	Encryption	RSA
Intel® QuickAssist Technology					
Intel® C621 Chipset	LBG-1G	0/4	N/A	N/A	N/A
Intel® C622 Chipset	LBG-2	2/4	N/A	N/A	N/A
Intel® C624 Chipset	LBG-4	4/4	N/A	N/A	N/A
Intel® C625 Chipset	LBG-E	4/4	20 Gb/s	20 Gb/s	20K Ops/s
Intel® C626 Chipset	LBG-M	4/4	40 Gb/s	40 Gb/s	40K Op/s
Intel® C627 Chipset	LBG-T	4/4	100 Gb/s	100 Gb/s	100K Ops/s
Intel® C628 Chipset	LBG-L	4/4	100 Gb/s	100 Gb/s	100K Ops/s

2nd Gen Intel® Xeon® Scalable Processor SKUs, continued

INTEL® SERVER SYSTEM S9200WK DATA CENTER BLOCK (DCB) TECHNICAL SPECIFICATIONS SUPPORTING THE INTEL® XEON® PLATINUM 9200 PROCESSOR

Form Factor	2U rack enclosure; up to 4 independent warm-swap compute modules
CPU	Intel® Xeon® Platinum 9200 Processors up to 56 cores
Memory	<ul style="list-style-type: none">• DDR4-2933 MT/s DIMMS, 24 DIMMs per platform @ 1DPC• Supports 8 GB to 128 GB DIMM options, number, and capacity configurable
Storage	<ul style="list-style-type: none">• Up to 8x M.2 SSDs per DCB (2x per compute module with 1U compute module; up to 4x M.2 SSDs and 4x hot-swap U.2 NVMe SSDs with 2U Compute Modules)• M.2 and U.2 number and capacity configurable
Ethernet	Integrated 1Gbase-T RJ45 (two ports per compute module), optional shared 1Gbase-T RJ45 management port chassis card
Cooling	Available with high flow rate air-cooling or integrated liquid-cooling options
I/O	2 x16 Gen3 PCIe slots per 1U compute module; 4 x16 Gen3 PCIe slots per 2U compute module for high-speed networking support
Manageability	Dedicated, consolidated Management Module
Security & Serviceability	TPM 2.0 (optional); hot-swap/redundant fans, and PSUs; light path diagnostic LEDs

Not a comprehensive list of all features and capabilities. Learn more at [intel.com/serverproducts](https://www.intel.com/serverproducts)

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration.

No product or component can be absolutely secure.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. For more complete information about performance and benchmark results, visit <http://www.intel.com/benchmarks>.

Performance results are based on testing as of dates shown in configuration and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Pricing guidance as of February 24, 2020. Intel does not guarantee any costs or cost reduction. You should consult other information and performance tests to assist you in your purchase decision.

Intel® Advanced Vector Extensions (Intel® AVX)* provides higher throughput to certain processor operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you can learn more at <http://www.intel.com/go/turbo>.

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

¹ 2x Average Performance Improvement compared with Intel® Xeon® Platinum 8180 processor. Geomean of est SPECrate2017_int_base, est SPECrate2017_fp_base, Stream Triad, Intel Distribution of Linpack, server side Java. Platinum 92xx vs Platinum 8180: 1-node, 2x Intel® Xeon® Platinum 9282 cpu on Walker Pass with 768 GB (24x 32GB 2933) total memory, ucode 0x400000A on RHEL7.6, 3.10.0-957.el7.x86_65, IC19u1, AVX512, HT on all (off Stream, Linpack), Turbo on all (off Stream, Linpack), result: est int throughput=635, est fp throughput=526, Stream Triad=407, Linpack=6411, server side java=332913, test by Intel on 2/16/2019. vs. 1-node, 2x Intel® Xeon® Platinum 8180 cpu on Wolf Pass with 384 GB (12 X 32GB 2666) total memory, ucode 0x200004D on RHEL7.6, 3.10.0-957.el7.x86_65, IC19u1, AVX512, HT on all (off Stream, Linpack), Turbo on all (off Stream, Linpack), result: est int throughput=307, est fp throughput=251, Stream Triad=204, Linpack=3238, server side java=165724, test by Intel on 1/29/2019.

² Up to 30X AI performance with Intel® DL Boost compared to Intel® Xeon® Platinum 8180 processor (July 2017). Tested by Intel as of 2/26/2019. Platform: Dragon rock 2 socket Intel® Xeon® Platinum 9282(56 cores per socket), HT ON, turbo ON, Total Memory 768 GB (24 slots/ 32 GB/ 2933 MHz), BIOS:SE5C620.86B.0D.01.0241.112020180249, Centos 7 Kernel 3.10.0-957.5.1.el7.x86_64, Deep Learning Framework: Intel® Optimization for Caffe version: <https://github.com/intel/caffe/d554cbf1>, ICC 2019.2.187, MKL DNN version: v0.17 (commit hash: 830a10059a018cd-2634d94195140cf2d8790a75a), model: https://github.com/intel/caffe/blob/master/models/intel_optimized_models/int8/resnet50_int8_full_conv.prototxt, BS=64, No datalayer DummyData:3x224x224, 56 instance/2 socket, Datatype:INT8 vs Tested by Intel as of July 11th 2017: 2S Intel® Xeon® Platinum 8180 CPU @ 2.50GHz (28 cores), HT disabled, turbo disabled, scaling governor set to "performance" via intel_pstate driver, 384GB DDR4-2666 ECC RAM. CentOS Linux release 7.3.1611 (Core), Linux kernel 3.10.0-514.10.2.el7.x86_64. SSD: Intel® SSD DC S3700 Series (800GB, 2.5in SATA 6Gb/s, 25nm, MLC). Performance measured with: Environment variables: KMP_AFFINITY=granularity=fine,compact; OMP_NUM_THREADS=56, CPU Freq set with cpupower frequency-set -d 2.5G -u 3.8G -g performance. Caffe: (<http://github.com/intel/caffe/>), revision f96b759f71b2281835f690af267158b82b150b5c. Inference measured with "caffe time --forward_only" command, training measured with "caffe time" command. For "ConvNet" topologies, dummy dataset was used. For other topologies, data was stored on local storage and cached in memory before training. Topology specs from https://github.com/intel/caffe/tree/master/models/intel_optimized_models (ResNet-50). Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with "numactl -l".

³ Up to 3.50X 5-Year Refresh Performance Improvement VM density compared to Intel® Xeon® E5-2600 v2 processor: 1-node, 2x E5-2697 v2 on Canon Pass with 256 GB (16 slots / 16GB / 1600) total memory, ucode 0x42c on RHEL7.6, 3.10.0-957.el7.x86_65, 1x Intel 400GB SSD OS Drive, 2x P4500 4TB PCIe, 2*82599 dual port Ethernet, Virtualization Benchmark, VM kernel 4.19, HT on, Turbo on, score: VM density=74, test by Intel on 1/15/2019. vs. 1-node, 2x 8280 on Wolf Pass with 768 GB (24 slots / 32GB / 2666) total memory, ucode 0x2000056 on RHEL7.6, 3.10.0-957.el7.x86_65, 1x Intel 400GB SSD OS Drive, 2x P4500 4TB PCIe, 2*82599 dual port Ethernet, Virtualization Benchmark, VM kernel 4.19, HT on, Turbo on, score: VM density=21, test by Intel on 1/15/2019.

⁴ Up to 14X AI Performance Improvement with Intel® DL Boost compared to Intel® Xeon® Platinum 8180 Processor (July 2017). Tested by Intel as of 2/20/2019. 2 socket Intel® Xeon® Platinum 8280 Processor, 28 cores HT On Turbo ON Total Memory 384 GB (12 slots/ 32GB/ 2933 MHz), BIOS: SE5C620.86B.0D.01.0271.120720180605 (ucode: 0x200004d), Ubuntu 18.04.1 LTS, kernel 4.15.0-45-generic, SSD 1x sda INTEL SSDSC2BA80 SSD 745.2GB, nvme1n1 INTEL SSDPE2KX040T7 SSD 3.7TB, Deep Learning Framework: Intel® Optimization for Caffe version: 1.1.3 (commit hash: 7010334f159da247db3fe3a9d96a3116ca06b09a), ICC version 18.0.1, MKL DNN version: v0.17 (commit hash: 830a10059a018cd2634d94195140cf2d8790a75a, model: https://github.com/intel/caffe/blob/master/models/intel_optimized_models/int8/resnet50_int8_full_conv.prototxt, BS=64, DummyData, 4 instance/2 socket, Datatype: INT8 vs Tested by Intel as of July 11th 2017: 2S Intel® Xeon® Platinum 8180 CPU @ 2.50GHz (28 cores), HT disabled, turbo disabled, scaling governor set to "performance" via intel_pstate driver, 384GB DDR4-2666 ECC RAM, CentOS Linux release 7.3.1611 (Core), Linux kernel 3.10.0-514.10.2.el7.x86_64, SSD: Intel® SSD DC S3700 Series (800GB, 2.5in SATA 6Gb/s, 25nm, MLC). Performance measured with: Environment variables: KMP_AFFINITY="granularity=fine,compact", OMP_NUM_THREADS=56, CPU Freq set with cpupower frequency-set -d 2.5G -u 3.8G -g performance. Caffe: (http://github.com/intel/caffe/), revision f96b759f71b2281835f690af267158b82b150b5c. Inference measured with "caffe time --forward_only" command, training measured with "caffe time" command. For "ConvNet" topologies, dummy dataset was used. For other topologies, data was stored on local storage and cached in memory before training. Topology specs from https://github.com/intel/caffe/tree/master/models/intel_optimized_models (ResNet-50), Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with "numactl -l"

⁵ 13x faster restart time.

	Baseline Config (DRAM)	AD 2-2-2 Config
System	Lightning Ridge (4S)	Lightning Ridge (4S)
CPU	Intel® Xeon® 8280M	Intel® Xeon® 8280L
CPUs per node	4-socket @ 28 core/socket	4-socket @ 28 core/socket
Memory	6TB 48x 128 GB DDR4 @ 2666 MT/s	9 TB 24x 256 GB Intel® Optane™ PMEM 24x 128 GB DDR4 @ 2666 MT/s
Network	10 GbE Intel X520 NIC	10 GbE Intel X520 NIC
Storage	60x Intel SSD DC S4600 SATA 480GB TB	90x Intel SSD DC S4600 SATA 480GB TB
BIOS	WW48'18	WW48'18
OS or VM version	SUSE 15	SUSE 15
WL Version	Intel IT workload	Intel IT workload
SAP HANA* database size	3TB	6TB
Security mitigations	6TB	SUSE 15
Date costs projected	March 1, 2019	March 1, 2019

⁶ 36% more VMs per node & 25% lower estimated cost per VM configurations.

	Config1-DDR4 (Similar Cost)	Config2-Intel Optane persistent memory (Similar Cost)
Test by	Intel	Intel
Test date	01/31/2019	01/31/2019
Platform	Confidential – Refer to M. Strassmaier if a need to know exists	Confidential – Refer to M. Strassmaier if a need to know exists
# Nodes	1	1
# Sockets	2	2
CPU	Cascade Lake B0 8272L	Cascade Lake B0 8272L
Cores/socket, Threads/socket	26/52	26/52
HT	ON	ON
Turbo	ON	ON
BKC version – E.g. ww47	WW42	WW42
Intel Optane persistent memory FW version	5253	5253
System DDR Mem Config: slots/cap/run-speed	24 slots/32 GB/2666	12 slots/16 GB /2666
System DCPMM Config: slots/cap/run-speed		8 slots/128 GB/ 2666
Total Memory/Node (DDR, DCPMM)	768 GB, 0	192 GB, 1 TB
Storage – boot	1x Samsung PM963 M.2 960 GB	1x Samsung PM963 M.2 960 GB
Storage – application drives	7 x Samsung PM963 M.2 960 GB, 4x Intel SSDs S4600 (1.92 TB)	7x Samsung PM963 M.2 960 GB, 4x Intel SSDs S4600 (1.92 TB)
NIC	1x Intel X520 SR2 (10Gb)	1x Intel X520 SR2 (10 Gb)
PCH	LBG QS/PRQ – T – B2	LBG QS/PRQ – T – B2
Other HW (Accelerator)		
OS	Windows Server 2019 R55-17763	Windows Server 2019 R55-17763
Kernel		
Workload & version	OLTP Cloud Benchmark	OLTP Cloud Benchmark
Compiler		
Libraries		
Other SW (Frameworks, Topologies...)		

⁷ 1. OLTP Warehouse claim of up to 3.7X: 1-node, 2x Intel® Xeon® CPU E5-2697 v2 on Canoe Pass with 256 GB (16 slots / 16 GB / 1866) total memory, ucode 0x42d on RHEL7.6, 3.10.0-957.el7.x86_65, 2 x Intel DC P3700 PCI-E SSD for DATA, 2 x Intel DC P3700 PCI-E SSD for REDO, HammerDB 3.1, HT on, Turbo on, result: transactions per minute=2242024, test by Intel on 2/1/2019. vs. 1-node, 2x Intel® Xeon® Platinum 8280 CPU on Wolf Pass with 384 GB (12 slots / 32 GB / 2933) total memory, ucode 0x4000013 on RHEL7.6, 3.10.0-957.el7.x86_65, 2x Intel® SSD DC P4610 for DATA, 2x Intel SSD DC P4610 for REDO, HammerDB 3.1, HT on, Turbo on, result: transactions per minute=8459206, test by Intel on 2/1/2019. 8^{BigBench} claim of 2.3X: 1+4-node, 2x Intel® Xeon® processor E5-2697 v2 on S2600JF with 128 GB (8 slots / 16GB / 1866) total memory, ucode 0x42d on CentOS-7.6.1810, 4.20.0-1.el7.x86_64, 1x 180GB SATA3 SSD, 3 x Seagate ST4000NM0033 (4TB), 1x Intel I350, TPCx-BB v1.2 (not for publication) / 3TB/ 2 Streams, Mllib, Oracle Hot-Spot 1.8.0_191, python-2.7.5, Apache Hadoop-2.9.2, Apache Spark-2.0.2, Hive 2.2 + CustomCommit, HT on, Turbo on, result: queries per min=265, test by Intel on 1/24/2019. 1+4-node, 2x Intel® Xeon® Gold 6148 processor on S2600WF with 768 GB (384 GB used) (12 slots* / 64 GB / 2400 (384GB used)) total memory, ucode 0x400000A on CentOS-7.6.1810, 4.20.0-1.el7.x86_64, Intel SSD DC S3710, 6 x Seagate ST2000NX0253 (2TB), 1x Intel X722, TPCx-BB v1.2 (not for publication) / 3TB/ 2 Streams, Mllib, Oracle Hot-Spot 1.8.0_191, python-2.7.5, Apache Hadoop-2.9.2, Apache Spark-2.0.2, Hive 2.2 + CustomCommit, HT on, Turbo on, result: queries per min=622, test by Intel on 1/12/2019.

⁸ HiBench claim of 4.3X: 1+4-node, 2x Intel® Xeon® processor E5-2697 v2 on S2600JF with 128 GB (8 slots / 16GB / 1866) total memory, ucode 0x42d on CentOS-7.6.1810, 4.20.0-1.el7.x86_64, 1x 180GB SATA3 SSD, 3 x Seagate ST4000NM0033 (4TB), 1x Intel I350, HiBench v7.1 / bigdata, Mllib, OpenJDK-1.8.0_191, python-2.7.5, Apache Hadoop-2.9.1, Apache Spark-2.2.2, HT on, Turbo on, result: SparkKmeans=119.5M, HadoopKmeans=49.6M, SparkSort=121.4M, HadoopSort=103M, SparkTerasort=107.4M, HadoopTerasort=109M, test by Intel on 1/23/2019. 1+4-node, 2x Intel® Xeon® Gold 6248 processor on S2600WF with 768 GB (384 GB used) (12 slots* / 64 GB / 2400 (384GB used)) total memory, ucode 0x400000A on CentOS-7.6.1810, 4.20.0-1.el7.x86_64, Intel SSD DC S3710, 6 x Seagate ST2000NX0253 (2TB), 1x Intel X722, HiBench v7.1 / bigdata, Mllib, OpenJDK-1.8.0_191, python-2.7.5, Apache Hadoop-2.9.1, Apache Spark-2.2.2, HT on, Turbo on, result: SparkKmeans=1235.8M, HadoopKmeans=92.8M, SparkSort=518.4M, HadoopSort=363.5M, SparkTerasort=589.3M, HadoopTerasort=457.3M, test by Intel on 1/23/2019.

Type	DataSetSize (B)	Overall* Duration (s)		Throughput (B/s)		Throughput Speedup
		Intel® Xeon® E5-2697 v2	Intel® Xeon® Gold 6248	Intel® Xeon® E5-2697 v2	Intel® Xeon® Gold 6248	
SparkKmeans	240,981,849,494	2015	195	119,593,969	1,235,804,356	10.33
SparkSort	307,960,500,694	2535	594	121,483,432	518,452,021	4.27
SparkTerasort	600,000,000,000	5586	1018	107,411,385	589,390,962	5.49
					Spark Geomean	6.23
HadoopKmeans	240,981,849,494	4854	2596	49,646,034	92,828,139	1.87
HadoopSort	307,960,500,694	2990	847	103,002,660	363,589,729	3.53
HadoopTerasort	600,000,000,000	5504	1312	109,011,627	457,317,073	4.20
					Hadoop Geomean	3.03
					Overall Geomean	4.34

¹⁰ Up to 1.25 to 1.58X NVF Workload Performance Improvement comparing Intel® Xeon® Gold 6230N processor to Intel® Xeon® Gold 6130 processor.

VPP IP Security: Tested by Intel on 1/17/2019 1-Node, 2x Intel® Xeon® Gold 6130 Processor on Neon City platform with 12x 16GB DDR4 2666MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 6x Intel XXV710-DA2, Bios: PLYDCRB1.86B.0155.R08.1806130538, ucode: 0x200004d (HT= ON, Turbo= OFF), OS: Ubuntu* 18.04 with kernel: 4.15.0-42-generic, Benchmark: VPP IPsec w/AESNI (AES-GCM-128) (Max Gbits/s (1420B)), Workload version: VPP v17.10, Compiler: gcc7.3.0, Results: 179. Tested by Intel on 1/17/2019 1-Node, 2x Intel® Xeon® Gold 6230N Processor on Neon City platform with 12x 16GB DDR4 2999MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 6x Intel XXV710-DA2, Bios: PLYXCRB1.PFT.0569.D08.1901141837, ucode: 0x4000019 (HT= ON, Turbo= OFF), OS: Ubuntu* 18.04 with kernel: 4.20.0-042000rc6-generic, Benchmark: VPP IPsec w/AESNI (AES-GCM-128) (Max Gbits/s (1420B)), Workload version: VPP v17.10, Compiler: gcc7.3.0, Results: 225

VPP FIB: Tested by Intel on 1/17/2019 1-Node, 2x Intel® Xeon® Gold 6130 Processor on Neon City platform with 12x 16GB DDR4 2666MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 6x Intel XXV710-DA2, Bios: PLYDCRB1.86B.0155.R08.1806130538, ucode: 0x200004d (HT= ON, Turbo= OFF), OS: Ubuntu* 18.04 with kernel: 4.15.0-42-generic, Benchmark: VPP FIB (Max Mpackets/s (64B)), Workload version: VPP v17.10 in ipv4fib configuration, Compiler: gcc7.3.0, Results: 160. Tested by Intel on 1/17/2019 1-Node, 2x Intel® Xeon® Gold 6230N Processor on Neon City platform with 12x 16GB DDR4 2999MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 6x Intel XXV710-DA2, Bios: PLYXCRB1.PFT.0569.D08.1901141837, ucode: 0x4000019 (HT= ON, Turbo= OFF), OS: Ubuntu* 18.04 with kernel: 4.20.0-042000rc6-generic, Benchmark: VPP FIB (Max Mpackets/s (64B)), Workload version: VPP v17.10 in ipv4fib configuration, Compiler: gcc7.3.0, Results: 212.9

Virtual Firewall: Tested by Intel on 10/26/2018 1-Node, 2x Intel® Xeon® Gold 6130 Processor on Neon City platform with 12x 16GB DDR4 2666MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 4x Intel XXV710-DA4, Bios: PLYDCRB1.86B.0155.R08.1806130538, ucode: 0x200004d (HT= ON, Turbo= OFF), OS: Ubuntu* 18.04 with kernel: 4.15.0-42-generic, Benchmark: Virtual Firewall (64B Mpps), Workload version: opnfv 6.2.0, Compiler: gcc7.3.0, Results: 38.9. Tested by Intel on 2/04/2019 1-Node, 2x Intel® Xeon® Gold 6230N Processor on Neon City platform with 12x 16GB DDR4 2999MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 6x Intel XXV710-DA2, Bios: PLYXCRB1.PFT.0569.D08.1901141837, ucode: 0x4000019 (HT= ON, Turbo= OFF), OS: Ubuntu* 18.04 with kernel: 4.20.0-042000rc6-generic, Benchmark: Virtual Firewall (64B Mpps), Workload version: opnfv 6.2.0, Compiler: gcc7.3.0, Results: 52.3

Virtual Broadband Network Gateway: Tested by Intel on 11/06/2018 1-Node, 2x Intel® Xeon® Gold 6130 Processor on Neon City platform with 12x 16GB DDR4 2666MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 6x Intel XXV710-DA2, Bios: PLYDCRB1.86B.0155.R08.1806130538, ucode: 0x200004d (HT= ON, Turbo= OFF), OS: Ubuntu* 18.04 with kernel: 4.15.0-42-generic, Benchmark: Virtual Broadband Network Gateway (88B Mpps), Workload version: DPK v18.08 ip_pipeline application, Compiler: gcc7.3.0, Results: 56.5. Tested by Intel on 1/2/2019 1-Node, 2x Intel® Xeon® Gold 6230N Processor on Neon City platform with 12x 16GB DDR4 2999MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 6x Intel XXV710-DA2, Bios: PLYXCRB1.PFT.0569.D08.1901141837, ucode: 0x4000019 (HT= ON, Turbo= OFF), OS: Ubuntu* 18.04 with kernel: 4.20.0-042000rc6-generic, Benchmark: Virtual Broadband Network Gateway (88B Mpps), Workload version: DPK v18.08 ip_pipeline application, Compiler: gcc7.3.0, Results: 57.7

VCMTS: Tested by Intel on 1/22/2019 1-Node, 2x Intel® Xeon® Gold 6130 Processor on Supermicro* X11DPH-Tq platform with 12x 16GB DDR4 2666MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 4x Intel XXV710-DA2, Bios: American Megatrends Inc.* version: '2.1', ucode: 0x200004d (HT= ON, Turbo= OFF), OS: Ubuntu* 18.04 with kernel: 4.20.0-042000rc6-generic, Benchmark: Virtual Converged Cable Access Platform (iMIX Gbps), Workload version: vcmts 18.10, Compiler: gcc7.3.0, Other software: Kubernetes* 1.11, Docker* 18.06, DPK 18.11, Results: 54.8. Tested by Intel on 1/22/2019 1-Node, 2x Intel® Xeon® Gold 6230N Processor on Neon City platform with 12x 16GB DDR4 2999MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 6x Intel XXV710-DA2, Bios: PLYXCRB1.PFT.0569.D08.1901141837, ucode: 0x4000019 (HT= ON, Turbo= OFF), OS: Ubuntu* 18.04 with kernel: 4.20.0-042000rc6-generic, Benchmark: Virtual Converged Cable Access Platform (iMIX Gbps), Workload version: vcmts 18.10, Compiler: gcc7.3.0, Other software: Kubernetes* 1.11, Docker* 18.06, DPK 18.11, Results: 83.7

OVS DPK: Tested by Intel on 1/21/2019 1-Node, 2x Intel® Xeon® Gold 6130 Processor on Neon City platform with 12x 16GB DDR4 2666MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 4x Intel XXV710-DA2, Bios: PLYXCRB1.86B.0568.D10.1901032132, ucode: 0x200004d (HT= ON, Turbo= OFF), OS: Ubuntu* 18.04 with kernel: 4.15.0-42-generic, Benchmark: Open Virtual Switch (on 4C/4P/8T 64B Mpacket/s), Workload version: OVS 2.10.1, DPK-17.11.4, Compiler: gcc7.3.0, Other software: QEMU-2.12.1, VPP v18.10, Results: 9.6. Tested by Intel on 1/18/2019 1-Node, 2x Intel® Xeon® Gold 6230N Processor on Neon City platform with 12x 16GB DDR4 2999MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 6x Intel XXV710-DA2, Bios: PLYXCRB1.86B.0568.D10.1901032132, ucode: 0x4000019 (HT= ON, Turbo= OFF), OS: Ubuntu* 18.04 with kernel: 4.20.0-042000rc6-generic, Benchmark: Open Virtual Switch (on 6P/6C/12T 64B Mpacket/s), Workload version: OVS 2.10.1, DPK-17.11.4, Compiler: gcc7.3.0, Other software: QEMU-2.12.1, VPP v18.10, Results: 15.2. Tested by Intel on 1/18/2019 1-Node, 2x Intel® Xeon® Gold 6230N Processor with SST-BF enabled on Neon City platform with 12x 16GB DDR4 2999MHz (384GB total memory), Storage: 1x Intel® 240GB SSD, Network: 6x Intel XXV710-DA2, Bios: PLYXCRB1.86B.0568.D10.1901032132, ucode: 0x4000019 (HT= ON, Turbo= ON (SST-BF)), OS: Ubuntu* 18.04 with kernel: 4.20.0-042000rc6-generic, Benchmark: Open Virtual Switch (on 6P/6C/12T 64B Mpacket/s), Workload version: OVS 2.10.1, DPK-17.11.4, Compiler: gcc7.3.0, Other software: QEMU-2.12.1, VPP v18.10, Results: 16.9

¹¹ Up to 1.7x better floating point perf/core using one copy SPECrate2017_fp_base* 2 socket Intel 8280 vs 2 socket AMD EPYC 7601. Xeon-SP 8280, Intel Xeon-based Reference Platform with 2 Intel® Xeon® 8280 processors (2.7GHz, 28 core), BIOS ver SE5C620.86B.0D.01.0348.011820191451, 01/18/2019, microcode: 0x5000017, HT OFF, Turbo ON, 12x32GB DDR4-2933, 1 SSD, Red Hat EL 7.6 (3.10.0-957.1.3.el7.x86_64), 1-copy SPECrate2017_fp_rate base benchmark compiled with Intel Compiler 19.0.1.144, -xCORE-AVX512 -ipo -O, executed on 1 core using taskset and numactl on core 0. Estimated score = 9.6, as of 2/6/2019 tested by Intel with security mitigations for variants 1,2,3,3a, and L1TF. AMD EPYC 7601, Supermicro AS-2023US-TR4 with 25 AMD EPYC 7601 with 2 AMD EPYC 7601 (2.2GHz, 32 core) processors, BIOS ver 1.1c, 10/4/2018, SMT OFF, Turbo ON, 16x32GB DDR4-2666, 1 SSD, Red Hat EL 7.6 (3.10.0-957.5.1.el7.x86_64), 1-copy SPECrate2017_fp_rate base benchmark compiled with AOCC ver 1.0 -Ofast, -march=znver1, executed on 1 core using taskset and numactl on core 0. Estimated score = 5.56, as of 2/8/2019 tested by Intel. Platinum 8280 vs Platinum 8180: 1-node, 2x Intel® Xeon® Platinum 8280M cpu on Wolf Pass with 384 GB (12 X 32GB 2933) total memory, ucode 0x400000A on RHEL7.6, 3.10.0-957.el7.x86_65, IC19u1, AVX512, HT on all (off Stream, Linpack), Turbo on all (off Stream, Linpack), result: est int throughput=317, est fp throughput=264, Stream Triad=217, Linpack=3462, server side java=177561, AIXPRT OpenVino/RN50=2324, test by Intel on 1/30/2019. vs. 1-node, 2x Intel® Xeon® Platinum 8180M cpu on Wolf Pass with 384 GB (12 X 32GB 2666) total memory, ucode 0x200004D on RHEL7.6, 3.10.0-957.el7.x86_65, IC19u1, AVX512, HT on all (off Stream, Linpack), Turbo on all (off Stream, Linpack), result: est int throughput=307, est fp throughput=251, Stream Triad=204, Linpack=3238, server side java=165724, AIXPRT OpenVino/RN50=1170, test by Intel on 1/29/2019.



© 2020 Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

* Other names and brands may be claimed as the property of others.

Printed in USA Q220/EK/8BG/PDF Please Recycle 336198-002