PRODUCT BRIEF

Intel® Xeon® Scalable Platform

The Future-Forward Platform Foundation for Agile Digital Services

Empowering Transformation in a Digital World

Across an evolving digital world, disruptive and emerging technology trends in business, industry, science, and entertainment increasingly impact the world's economies. By 2020, the success of half of the world's Global 2000 companies will depend on their abilities to create digitally enhanced products, services, and experiences,[1] and large organizations expect to see an 80 percent increase in their digital revenues,[2] all driven by advancements in technology and usage models they enable.

This global transformation is rapidly scaling the demands for flexible compute, networking, and storage. Future workloads will necessitate infrastructures that can seamlessly scale to support immediate responsiveness and widely 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.[3] 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, 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 and reduce TCO up to 65 percent[4] due to lower software and OS licensing fees, and acquisition, maintenance, and infrastructure costs.

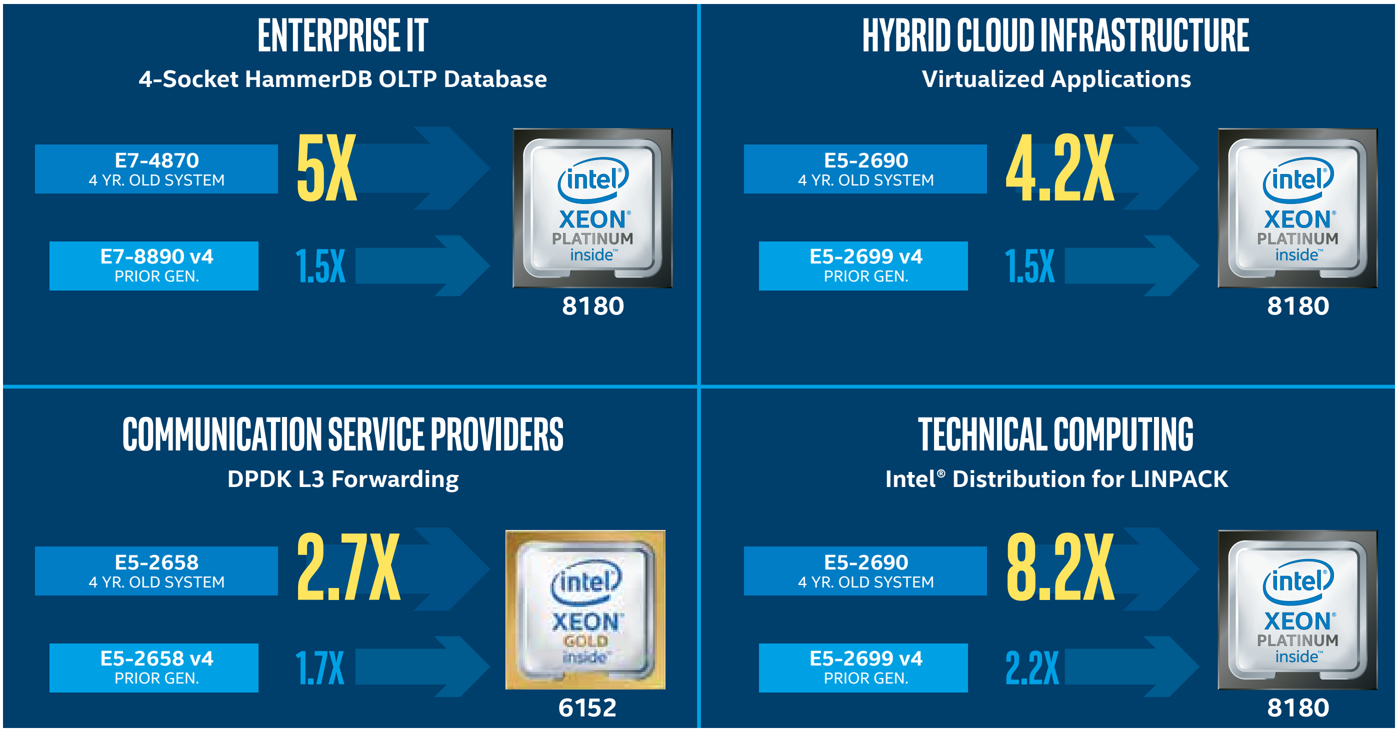

Comparing Intel® Xeon® Scalable platform to both four-year-old and last-generation Intel Xeon processor-based systems, the latest generation Intel CPU delivers better performance and capabilities for enterprise, cloud, communications, and HPC.

Virtualization in the enterprise has been steadily on the rise for a decade. Most organizations have taken advantage of some form of virtualization, creating greater need for servers to run virtual machines (VMs) in the data center. Intel Xeon Scalable platform enables 4.2x[5] more VMs per server compared to earlier generations, allowing IT to increase their consolidation of more services on less hardware.

Reduce costs and deliver more services per server with higher consolidation of VMs per server with Intel® Xeon® Scalable platform.

Pervasive, Breakthrough Performance

From its new Intel® Mesh Architecture and widely expanded resources to its hardware-accelerating and newly integrated technologies, the Intel Xeon Scalable platform enables a new level of consistent, pervasive, and breakthrough performance.

Intel® Xeon® Processor provides consistent performance for diverse workloads

	WHAT MATTERS?	INTEL® XEON® SCALABLE PLATFORM ADVANTAGES
Enterprise and Cloud	Minimize complexity with compatible virtualization infrastructure	Deploy quickly. Intel VMs coexist w/ other Intel® technology-based servers.
	Meet stringent customer SLAs	Fast response times.
HPC	Maximize vector floating performance and efficiency	High performance with fewer servers.
Storage	Ensure deterministic storage response	Deterministic performance. Cores, cache, memory, I/O in a single die.
Communications	Deliver diverse services efficiently	Incredible efficiency and hardware acceleration from a platform with features required for application, control, packet, and signal processing.

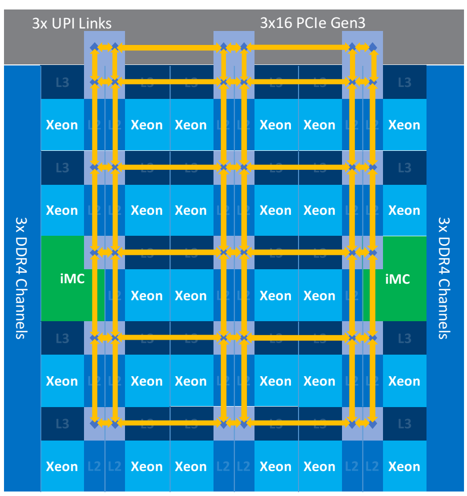

In the Intel® Xeon® Scalable platform Intel® Mesh Architecture with up to 28 cores, the Last Level Cache (LLC), six memory channels, and 48 PCIe* channels are shared among all the cores, giving access to large resources across the entire die and creating dynamic scalability without sacrificing performance for a variety of deployments, such as virtualization. VMs scale easily withfull access to all the shared onboard resources that they need.

Foundational Enhancements

- **Higher Per-Core Performance:** Up to 28 cores, delivering high performance and scalability for compute-intensive workloads across compute, storage, and network usages.
- **Greater Memory Bandwidth/Capacity:** 50 percent increased memory bandwidth and capacity. Six memory channels versus four memory channels of previous generation for memory-intensive workloads.
- **Expanded I/O:** 48 lanes of PCIe* 3.0 bandwidth and throughput for demanding I/O-intensive workloads.
- **Intel® Ultra Path Interconnect (Intel® UPI):** Up to three Intel UPI channels increase scalability of the platform to as many as eight sockets, as well as improves inter-CPU bandwidth for I/O-intensive workloads versus previous generation[3] (with Intel® Quick Path Interconnect). Intel UPI offers the perfect balance between improved throughput and energy efficiency.
- **Intel® Advanced Vector Extensions 512 (Intel® AVX-512):** With double the flops per clock cycle compared to previous generation Intel® AVX2,[6] 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 without compromise:** Near-zero encryption overhead[7] enables higher performance on all secure data transactions.

Innovative Integrations

First-time platform integrations deliver game-changing improvements in performance[8] and latency across the infrastructure:

- **Integrated Intel® Omni-Path Architecture (Intel® OPA) Host Fabric Interface:** End-to-end high-bandwidth, low-latency fabric that optimizes performance and eases deployment of HPC clusters by eliminating the need for a discrete host fabric interface card. Integrated in the CPU package.

- **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 Storage Support

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

- **Support for Intel® Optane™ SSDs and Intel® 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.[9]
- **Deploy next generation storage with confidence with Intel® Volume Management Devices (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.[10]
  - Up to 3.1x faster cryptographic hashing of SHA algorithms using Intel AVX-512.
  - Up to 1.2x algorithm speedup for AES-128-GCM.
  - Up to 2x faster Reed Solomon Erasure Code performance using Intel AVX-512.

Complementary Offerings for Even Greater Performance, Scalability

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

- The Intel® Xeon Phi™ Processor is an ideal foundation for highly parallel applications, such as machine learning training, simulation, and visualization.
- Intel® FPGAs offer efficient acceleration[11] along with the flexibility that comes with using programmable hardware for low-latency applications, such as virtual switching, network services, data analytics, and AI.
- A range of software tools and libraries for general and highly parallel computing help developers optimize applications for Intel® Architecture.

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 Optane SSDs and Intel® 3D NAND SSDs:** Bring larger data volumes closer to the CPU for faster backup and storage performance. Increases server efficiency and agility, while minimizing disruptions, with mission-critical reliability and advanced manageability and serviceability features.
- **Intel Volume Management Device:** Unlocks the power of NVMe-based storage (Intel Optane SSDs and Intel 3D NAND SSDs) with new management capabilities that include hot-plug storage and error containment functions.
- **Intel® VT-x new 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.

Strong, Capable Platforms for the Hybrid-Cloud, 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 on their foundational business workloads, such as OLTP and web infrastructure, seek to reduced 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 with up to 58 percent more requests/second on runtime workloads.[12] 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 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

- Up to 28 high-performance cores, six memory channels, and 48 lanes of PCIe 3.0
- Up to three Intel UPI channels
- Support for Intel Optane SSDs and Intel 3D NAND Solid State Drives
- Enhanced Intel Run Sure Technology

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.

Accelerate Network Traffic with Intel® Xeon® Scalable Platform, DPDK, and Intel® QAT

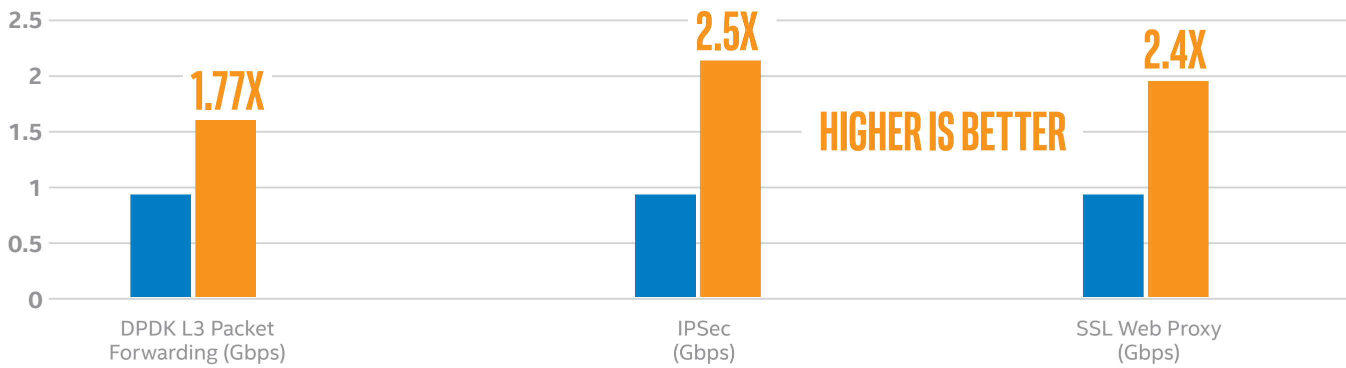

DPDK optimizations[13]: Utilize new cache hierarchy with large mid-level cache and balanced bandwidth improvements to address data, audio, and video growth.

IPSec[14]: Provide data authentication, integrity, and confidentiality with accelerated cryptography and compression/decompression.

SSL Web Proxy[15]: Accelerate data encryption/decryption for secure transmissions.

Previous
Intel® Xeon® Scalable Platform

Combined with the Data Plane Development Kit (DPDK) and Intel QAT, this new processor can boost network performance, enabling service providers to accommodate greater traffic to grow services—and revenue—and helping them to prepare for 5G.

Highlights for CommSP Innovation

- Integrated Intel QAT
- Hardware-based acceleration of encryption and compression using integrated Intel® QAT.
- Integrated Intel Ethernet
- Intel FPGAs maximize versatility in communications infrastructure
- Intel KPT with Integrated Intel QAT and Intel PTT
- Intel TXT with One-Touch Activation

Additional Resources Optimized for CommSPs

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.

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 18x higher inference[16] and up to 19x higher training[17] throughput, compared to four-year-old systems.

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.

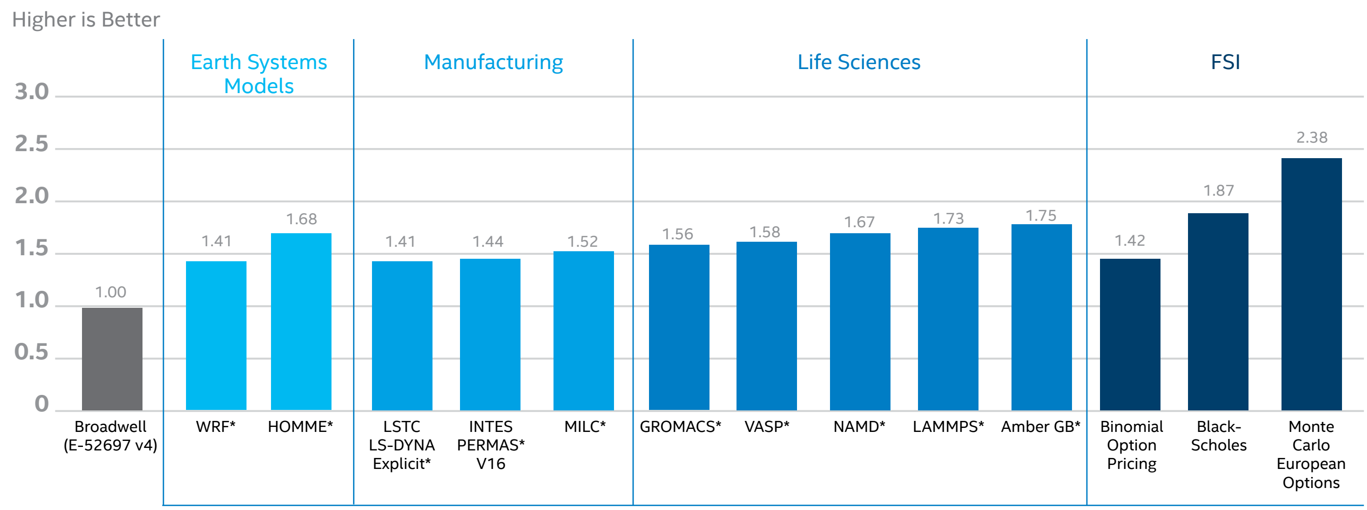

Across 13 common HPC workloads, Intel® Xeon® Scalable platform delivers improved performance.

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 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[18] up to 32 nodes enables building large HPC solutions that are not inhibited by the interconnect. And, the Intel Xeon Scalable platform, compared to previous generation Intel Xeon processors, delivers 1.63x better overall performance across 13 common HPC workloads used throughout the sciences and financial services.[19] 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.

Highlights for HPC Innovation

- Intel Ultra Path Interconnect
- Intel Advanced Vector Extensions 512
- Integrated Intel Omni-Path Architecture Host Fabric Interface
- Support for Intel Optane SSDs and Intel 3D NAND Solid State Drives

Additional Technologies for HPC, HPDA, and AI

- The bootable Intel® Xeon Phi™ Processor 7200 Series, with its Many Integrated Core (MIC) architecture, is an ideal foundation for highly parallel workloads, such as machine learning training, simulation, and visualization.
- 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.

Intel Xeon Scalable Platform at a Glance

Intel Xeon Scalable platform offers a range of performance, scalability, and feature options to meet a wide variety of workloads in the data center, from the most advanced (Intel® Xeon® Platinum processor 8XXX family) to entry-level (Intel® Xeon® Bronze processor 3XXX family).

BEST PERFORMANCE, MOST SCALABLE, BEST BUSINESS AGILITY

Intel® Xeon® Platinum Processor 8XXX Family

- Best choice for mission-critical, real-time analytics, machine learning, and artificial intelligence workloads
- Best workload-optimized performance for general purpose compute and hybrid-cloud deployments
- Best performance for the most demanding storage and networking workloads
- Best memory bandwidth and 2, 4, 8+ socket scalability

GREAT PERFORMANCE, FAST MEMORY, AND MORE INTERCONNECT/ACCELERATOR ENGINES

Intel® Xeon® Gold Processor 6XXX Family

- Significant workload-optimized performance improvements for general purpose compute
- Significant improvements for demanding storage and networking workloads
- Highest memory speed, highest memory capacity, and enhanced Intel AVX-512
- Enhanced 2-4 socket scalability and performance

BETTER PERFORMANCE, ADVANCED RELIABILITY

Intel® Xeon® Gold Processor 5XXX Family

- Improved performance for compute-bound workloads
- Affordable Advanced RAS and 4-socket scalability
- Suitable for a wider range of workloads

EFFICIENT PERFORMANCE AT LOW POWER

Intel® Xeon® Silver Processor 4XXX Family

- Solid compute capability (Hyper-Threading, Turbo Boost)
- Improved memory speed, energy efficiency
- Suitable for a moderate range workloads

ENTRY-LEVEL PERFORMANCE AND HW-ENHANCED SECURITY

Intel® Xeon® Bronze Processor 3XXX Family

- Affordable, entry-level 2-socket support suitable for light-range workloads
- Reliable upgrade versus Intel® Xeon® processor E3 product family

Family Features

	Intel® Xeon® Bronze Processor (3000 Series)	Intel® Xeon® Silver Processor (4000 Series)	Intel® Xeon® Gold Processor (5000 Series)	Intel® Xeon® Gold Processor (6000 Series)	Intel® Xeon® Platinum Processor (8000 Series)
PERVASIVE PERFORMANCE AND SECURITY					
Highest Core Count Supported	8 cores	12 cores	14 cores	22 cores	28 cores
Highest Supported Frequency	1.7 GHz (8C/85W)	2.2 GHz (10C/85W)	3.6 GHz (4C/105W)	3.4 GHz (6C/115W)	3.6 GHz (4C/105W)
Number of CPU Sockets Supported	Up to 2	Up to 2	Up to 4	Up to 4	Up to 8
Intel® Ultra Path Interconnect (UPI)	2	2	2	3	3
Intel® UPI Speed	9.6 GT/s	9.6 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
Memory Speed Support (DDR4)	2133 MHz	2400 MHz	2400 MHz	2666 MHz	2666 MHz
Highest Memory Capacity Supported Per Socket	768 GB	768 GB	768 GB	768 GB, 1.5	768 GB, 1.5 TB
Intel® Omni-Path Architecture (Discrete PCIe* card)	•	•	•	•	•
Intel® QuickAssist Technology (Integrated in chipset)	•	•	•	•	•
Intel® QuickAssist Technology (Discrete PCIe card)	•	•	•	•	•
Intel® Optane™ Technology-based SSDs (3D XPoint™)	•	•	•	•	•
Intel® 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		•	•	•	•
Node Controller Support				•	•
Intel® Omni-Path Architecture (Integrated)				•	•
HIGH RELIABILITY					
Reliability, Accessibility, and Serviceability (RAS) Capability	Standard	Standard	Standard	Advanced	Advanced
Intel® Run Sure Technology				•	•
AGILITY & EFFICIENCY					
Intel® Volume Management Device (Intel VMD)	•	•	•	•	•
Intel® VT-x	•	•	•	•	•
Intel® Speed Shift Technology	•	•	•	•	•
Intel® Node Manager 4.0	•	•	•	•	•
SECURITY					
Intel® Memory Protection Extensions (MPX)	•	•	•	•	•
Mode-based Execute Control	•	•	•	•	•
Intel® Key Protection Technology (KPT) w/Integrated Intel® QAT	•	•	•	•	•
Intel® Platform Trust Technology (PTT)	•	•	•	•	•
Intel® TXT w/ One-Touch Activation (OTA)	•	•	•	•	•

Intel Xeon Scalable Platform

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	**8**	**1**	**#**	**#**	**α**	**α**	**processor**
Intel® Xeon® Gold	**6**	**1**	**#**	**#**	**α**	**α**	**processor**
Intel® Xeon® Gold	**5**	**1**	**#**	**#**	**α**	**α**	**processor**
Intel® Xeon® Silver	**4**	**1**	**#**	**#**	**α**	**α**	**processor**
Intel® Xeon® Bronze	**3**	**1**	**#**	**#**	**α**	**α**	**processor**

SKU Level

8 = Platinum

6, 5 = Gold

4 = Silver

3 = Bronze

Processor Generation

1 = 1st Gen (Skylake)

Processor SKU

(ex. 20, 34...)

Integrations and Optimizations
(if applicable)

F = Fabric
T = High Tcase/
Extended Reliability

Memory Capacity

No Suffix = 768 GB per socket

M = 1.5 TB per socket

Intel Xeon Scalable Platform SKUs

SKUS OPTIMIZED FOR HIGHEST PER-CORE PERFORMANCE

Class	SKU	Cores	Base Non-AVX Speed (GHz)	TDP (W)
Platinum	8180[a]	28	2.5	205
Platinum	8168	24	2.7	205
Platinum	8158[b]	12	3.0	150
Platinum	8156[c]	4	3.6	105
Gold	6148	20	2.4	150
Gold	6154	18	3.0	200
Gold	6150	18	2.7	165
Gold	6142[a]	16	2.6	150
Gold	6132	14	2.66	140
Gold	6146[b]	12	3.2	165
Gold	6136[b]	12	3.0	150
Gold	6126[d]	12	2.6	125
Gold	6144[b]	8	3.5	150
Gold	6134[a,b]	8	3.2	130
Gold	6128[d]	6	3.4	115
Gold	5122[c]	4	3.6	105

SKUS OPTIMIZED FOR BALANCED ENERGY EFFICIENT PERFORMANCE/WATT

Class	SKU	Cores	Base Non-AVX Speed (GHz)	TDP (W)
Platinum	8176[a]	28	2.1	165
Platinum	8170[a]	26	2.1	165
Platinum	8164	26	2.0	150
Platinum	8160[a]	24	2.1	150
Platinum	8153[e]	16	2.0	125
Gold	6152	22	2.1	140
Gold	6138	20	2.0	125
Gold	6140[a]	18	2.3	140
Gold	6130	16	2.1	125
Gold	5120	14	2.2	105
Gold	5118	12	2.3	105
Gold	5115	10	2.4	85
Silver	4116	12	2.1	85
Silver	4114	10	2.2	85
Silver	4112	4	2.6	85
Silver	4110	8	2.1	85
Silver	4108	8	1.8	85
Bronze	3106	8	1.7	85
Bronze	3104	6	1.7	85

Intel Xeon Scalable Platform SKUs, continued

SKUS WITH EXTENDED LIFE (10-YEAR USE) AND NEBS-FRIENDLY THERMAL SPECIFICATION				
Class	**SKU**	**Cores**	**Base Non-AVX Speed (GHz)**	**TDP (W)**
Platinum	8160T	24	2.1	150
Gold	6138T	20	2.0	125
Gold	6130T	16	2.1	125
Gold	6126T[d]	12	2.6	125
Gold	5120T	14	2.2	105
Gold	5119T	14	1.9	85
Silver	4116T	12	2.1	85
Silver	4114T	10	2.2	85
Silver	4109T	8	2.0	70

SKUS WITH INTEGRATED INTEL® OMNI-PATH ARCHITECTURE FABRIC				
Class	**SKU**	**Cores**	**Base Non-AVX Speed (GHz)**	**TDP (W)**
Platinum	8176F	28	2.1	173
Platinum	8160F	24	2.1	160
Gold	6148F	20	2.4	160
Gold	6142F	16	2.6	160
Gold	6138F	20	2.0	135
Gold	6130F	16	2.1	135
Gold	6126F[d]	12	2.6	135

PRODUCT NAME	SKU	10Gb/1Gb 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

[a] "M" SKU available supporting 1.5 TB/socket
[b] Non-default 24.75 MB cache
[c] Non-default 16.5 MB cache
[d] Non-default 19.25 MB cache
[e] Compatible for 8-socket builds

[1] Source: IDC, http://www.idc.com/getdoc.jsp?containerId=prUS41888916.

[2] Source: Gartner Group, http://www.gartner.com/newsroom/id/3142917.

[3] Requires platform-to-platform comparison between Intel® Xeon® Processor Scalable Family and Intel® Xeon® E5 V1/V2. Up to 2.27x claim based on LINPACK*: 1-Node, 2 x Intel® Xeon® Processor E5-2699 v4 on Grantley-EP (Wellsburg) with 64 GB Total Memory on Red Hat Enterprise Linux* 7.0 kernel 3.10.0-123 using MP_LINPACK 11.3.1 (Composer XE 2016 U1). Data Source: Request Number: 1636, Benchmark: Intel® Distribution of LINPACK, Score: 1446.4 Higher is better vs. 1-Node, 2 x Intel® Xeon® Platinum 8180 Processor on Wolf Pass SKX with 384 GB Total Memory on Red Hat Enterprise Linux* 7.3 using mp_linpack_2017.1.013. Data Source: Request Number: 3753, Benchmark: Intel® Distribution of LINPACK, Score: 3295.57 Higher is better.

[4] Up to 65% lower 4-year TCO estimate example based on equivalent rack performance using VMware ESXi* virtualized consolidation workload comparing 20 installed 2-socket servers with Intel Xeon processor E5-2690 (formerly "Sandy Bridge-EP") running VMware ESXi* 6.0 GA using Guest OS RHEL6.4 compared at a total cost of $919,362 to 5 new Intel® Xeon® Platinum 8180 (Skylake) running VMware ESXi6.0 U3 GA using Guest OS RHEL 6 64bit at a total cost of $320,879 including basic acquisition. Server pricing assumptions based on current OEM retail published pricing for Broadwell based servers – subject to change based on actual pricing of systems offered.

Total Cost of Ownership Estimated		**Server A: 2S Xeon® E5-2690 (2.9 GHz, 8C, 20 MB)**	**Server B: 2S Xeon® Platinum 8180 (2.5 GHz, 28C, 38.5 M)**
Number of Servers		20	5
Server Acquisition		$273,220	$150,995
Operating System & Software	OS License	$123,100	$30,775
	OS Maintenance	$241,360	$60,340
	Software License	$39,800	$9,950
	Software Maintenance	$51,680	$12,920
Server Maintenance		$134,080	$33,520
Infrastructure & Utilities	Power and Cooling	$42,522	$18,979
	Rack/Floor Space	$12,400	$3,100
	Networking	$1,200	$300
Total		**$919,362**	**$320,879**

[5] Up to 4.2x more VMs based on server virtualization consolidation workload: Based on Intel® internal estimates 1-Node, 2 x Intel® Xeon® Processor E5-2690 on Romley-EP with 256 GB Total Memory on VMware ESXi* 6.0 GA using Guest OS RHEL6.4, glassfish3.1.2.2, postgresql9.2. Data Source: Request Number: 1718, Benchmark: server virtualization consolidation, Score: 377.6 @ 21 VMs vs. 1-Node, 2 x Intel® Xeon® Platinum 8180 Processor on Wolf Pass SKX with 768 GB Total Memory on VMware ESXi6.0 U3 GA using Guest OS RHEL 6 64bit. Data Source: Request Number: 2563, Benchmark: server virtualization consolidation, Score: 1580 @ 90 VMs. Higher is better.

[6] As measured by Intel® Xeon® Processor Scalable Family with Intel® AVX-512 compared to an Intel® Xeon® E5 v4 with Intel® AVX2.

[7] BigBench, Near Zero encryption overhead: BigBench query Runtime/second. Testing done by Intel. BASELINE: Platform 8168, NODES 1 Mgmt + 6 Workers, Make Intel Corporation, Model S2600WFD, Form Factor 2U, Processor Intel® Xeon® Platinum 8168, Base Clock 2.70 GHz, Cores per socket 24, Hyper-Threading Enabled, NUMA mode Enabled, RAM 384 GB DDR4, RAM Type 12x 32 GB DDR4, OS Drive Intel® SSD DC S3710 Series (800 GB, 2.5 in SATA 6Gb/s, 20nm, MLC), Data Drives 8x – Seagate Enterprise 2.5 HDD ST2000NX0403 2 TB, Intel® SSD DC P3520 Series (2.0TB), Temp Drive DC 3520 2 TB, NIC Intel X722 10 GBE – Dual Port, Hadoop Cloudera 5.11, Benchmark TPCx-BB 1.2, Operating System CentOS Linux release 7.3.1611 (Core); HDFS encryption turned OFF. vs. NEW: Platform 8168, NODES 1 Mgmt + 6 Workers, Make Intel Corporation, Model S2600WFD, Form Factor 2U, Processor Intel® Xeon® Platinum 8168, Base Clock 2.70 GHz, Cores per socket 24, Hyper-Threading Enabled, NUMA mode Enabled, RAM 384 GB DDR4, RAM Type 12x 32 GB DDR4, OS Drive Intel® SSD DC S3710 Series (800 GB, 2.5 in SATA 6 Gb/s, 20nm, MLC), Data Drives 8x – Seagate Enterprise 2.5 HDD ST2000NX0403 2 TB, Intel® SSD DC P3520 Series (2.0 TB), Temp Drive DC 3520 2 TB, NIC Intel X722 10 GBE – Dual Port, Hadoop Cloudera 5.11, Benchmark TPCx-BB 1.2, Operating System CentOS Linux release 7.3.1611 (Core); HDFS encryption turned ON.

[8] Up to 5x claim based on OLTP Warehouse workload: 1-Node, 4 x Intel® Xeon® Processor E7-4870 on Emerald Ridge with 512 GB Total Memory on Oracle Linux* 6.4 using Oracle 12c* running 800 warehouses. Data Source: Request Number: 56, Benchmark: HammerDB, Score: 2.46322e+006 Higher is better vs. 1-Node, 4 x Intel® Xeon® Platinum 8180 Processor on Lightning Ridge SKX with 768 GB Total Memory on Red Hat Enterprise Linux* 7.3 using Oracle 12.2.0.1 (including database and grid) with 800 warehouses. Data Source: Request Number: 2542, Benchmark: HammerDB, Score: 1.2423e+007 Higher is better.

Up to 1.73x claim based on HammerDB:1-Node, 2 x Intel® Xeon® Processor E5-2699 v4 on Grantley-EP (Wellsburg) with 384 GB Total Memory on Red Hat Enterprise Linux* 7.1 kernel 3.10.0-229 using Oracle 12.1.0.2.0 (including database and grid) with 800 warehouses, HammerDB 2.18. Data Source: Request Number: 1645, Benchmark: HammerDB, Score: 4.13568e+006 Higher is better vs. 1-Node, 2 x Intel® Xeon® Platinum 8180 Processor on Purley-EP (Lewisburg) with 768 GB Total Memory on Oracle Linux* 7.2 using Oracle 12.1.0.2.0, HammerDB 2.18. Data Source: Request Number: 2510, Benchmark: HammerDB, Score: 7.18049e+006 Higher is better.

Up to 4.2x more VMs based on virtualization consolidation workload: Based on Intel® internal estimates 1-Node, 2 x Intel® Xeon® Processor E5-2690 on Romley-EP with 256 GB Total Memory on VMware ESXi* 6.0 GA using Guest OS RHEL6.4, glassfish3.1.2.2, postgresql9.2. Data Source: Request Number: 1718, Benchmark: server virtualization workload, Score: 377.6 @ 21 VMs Higher is better vs. 1-Node, 2 x Intel® Xeon® Platinum 8180 Processor on Wolf Pass SKX with 768 GB Total Memory on VMware ESXi6.0 U3 GA using Guest VM's utilize RHEL 6 64bit OS. Data Source: Request Number: 2563, Benchmark: server virtualization workload, Score: 1580 @ 90 VMs Higher is better.

Up to 1.5x more VMs based on virtualization consolidation workload : 1-Node, 2 x Intel® Xeon® Processor E5-2699 v4 on Grantley-EP (Wellsburg) with 512 GB Total Memory on VMware ESXi* 6.0 Update 1 using Guest VM's utilize RHEL 6 64bit OS. Data Source: Request Number: 1637, Benchmark: server virtualization workload, Score: 1034 @ 58 vs. 1-Node, 2 x Intel® Xeon® Platinum 8180 Processor on Wolf Pass SKX with 768 GB Total Memory on VMware ESXi6.0 U3 GA using Guest VM's utilize RHEL 6 64bit OS. Data Source: Request Number: 2563, Benchmark: server virtualization workload, Score: 1580 @ 90 VMs Higher is better.

Up to 2.7x claim based on DPDK L3 Packet Forwarding: Intel® Xeon® E5-2650 processors 2.00GHz, 8 GT/s QPI, 20MB L3 cache, Patsburg Chipset (C0 stepping), Grizzly Pass Platform (R2216G-ZBPP), DDR3 1333MHz, 8 x dual rank 4GB (total 32GB), 4 memory channels per socket Configuration, 1 DIMM per channel, 6 x Intel® 82599 dual-port PCI-Express Gen2 x8 10 Gb Ethernet NIC, 1 x Intel® 82599 dual-port Gen2 x8 I/O expansion module10 Gb Ethernet NIC, BIOS version S5500.86B.01.00.0048, Operating system: Fedora Core 15, Kernel version: 2.6.38.4, IxNetwork* 6.0.400.22, DPDK version: FD5_1 Score: 102Gbits/s packet forwarding at 256B using cores vs. Gold 6152: Estimates based on Intel internal testing on Intel Xeon 6152 2.1 GHz, 2x Intel®, FM10420(RRC) Gen Dual Port 100GbE Ethernet controller (100Gbit/card) 2x Intel® XXV710 PCI Express Gen Dual Port 25GbE Ethernet controller (2x25G/card), DPDK 17.02. Score: 281 Gbits/s packet forwarding at 256B packet using cores, IO and memory on a single socket.

Up to 1.7x claim based on DPDK L3 Packet Forwarding: E5-2658 v4: 5 x Intel® XL710-QDA2, DPDK 16.04. Benchmark: DPDK l3fwd sample application Score: 158 Gbits/s packet forwarding at 256B packet using cores. Gold 6152: Estimates based on Intel internal testing on Intel Xeon 6152 2.1 GHz, 2x Intel®, FM10420(RRC) Gen Dual Port 100GbE Ethernet controller (100Gbit/card) 2x Intel® XXV710 PCI Express Gen Dual Port 25GbE Ethernet controller (2x25G/card), DPDK 17.02. Score: 281 Gbits/s packet forwarding at 256B packet using cores, IO and memory on a single socket.

[9] System Configuration: 4 Node vSAN* Cluster. Per Node configuration: Supermicro* SuperServer 2028U-TN24R4T+ Dual Intel® Xeon® E5-2687Wv4 (12 Core @ 3.0 Ghz), Supermicro* Server Board, 256 GB DDR4 RAM, Boot Drive, 1x Intel® SSD DC S3710 Series (200 GB, 2.5"), vSAN Intel 3D NAND Cluster: Virtual SAN SSDs - 2 Disk Groups comprised of 2x Intel® SSD DC P4600 Series (1.6TB, 2.5" SFF), 8x Intel® SSD DC P4500 Series (4 TB, 2.5" SFF), vSAN Intel 2D NAND Cluster: Virtual SAN SSDs - 2 Disk Groups comprised of 2x Intel® SSD DC P3700 Series (800GB, 2.5" SFF), 8x Intel® SSD DC P3500 Series (2 TB, 2.5" SFF), Intel® Ethernet Server Adapter X540-DA2.

[10] 3.1x, 1.2x, 2x ISA-L configuration. Intel Xeon® Processor Scalable Family: Platinum 8180 Processor, 28C, 2.5 GHz, H0, Neon City CRB, 12x16 GB DDR4 2666 MT/s ECC RDIMM, BIOS PLYCRB1.86B.0128.R08.1703242666.

Intel® Xeon® E5-2600v4 Series Processor, E5-2650 v4, 12C, 2.2 GHz, Aztec City CRB, 4x8 GB DDR4 2400 MT/s ECC RDIMM, BIOS GRRFCRB1.86B.0276.R02.1606020546

Operating System: Redhat Enterprise Linux 7.3, Kernel 4.2.3, ISA-L 2.18, BIOS Configuration, P-States: Disabled, Turbo: Disabled, Speed Step: Disabled, C-States: Disabled, ENERGY_PERF_BIAS_CFG: PERF.

[11] As measured by Intel® Xeon® Processor Scalable Family with Intel® FPGA optimized workload and Intel® Xeon® Processor Scalable Family without FPGA optimized workload.

[12] Up to 1.58x claim based on Ghost-NodeJS workload: 1-Node, 2 x Intel® Xeon® Processor E5-2699 v4 on Wildcat Pass with 384 GB Total Memory on Ubuntu 16.04 LTS using Node.js version 6.9.2, MySQL Maria DB version 15.1 Distrib 10.0.30. Data Source: Request Number: 2687, Benchmark: Ghost-NodeJS, Score: 2308 Higher is better vs. 1-Node, 2 x Intel® Xeon® Platinum 8180 Processor on Wolf Pass SKX with 384 GB Total Memory on Ubuntu 16.10 using Node.js version 6.9.2, MySQL Maria DB version 15.1 Distrib 10.0.30. Data Source: Request Number: 2687, Benchmark: Ghost-NodeJS, Score: 3647 Higher is better.

[13] Up to 1.77x claim based on DPDK L3 Forwarding: E5-2658 v4: 5 x Intel® XL710-QDA2, DPDK 16.04. Benchmark: DPDK l3fwd sample application. Vs. Gold 6152: Estimates based on Intel internal testing on Intel Xeon 6152 2.1 GHz, 2x Intel®, FM10420(RRC) Gen Dual Port 100GbE Ethernet controller (100Gbit/card) 2x Intel® XXV710 PCI Express Gen Dual Port 25GbE Ethernet controller (2x25G/card), DPDK 17.02. Score: 281 Gbits/s packet forwarding at 256B packet using cores, IO and memory on a single socket.

[14] Up to 2.5x claim based on DPDK IPSec Seg-gw benchmark: Intel® Xeon® E5-2658 v4, Intel® PCH C612, DDR4-2400 Intel 895XCC based Quick Assist Accelerator Adapter PCIe Gen3 x8 links, DPDK 16.11 IPSec-secgw,1420 B packet). Intel® DPDK 16.11 IPsec-secgw sample application. Cores, IO, packet buffer memory, and processing cores are on a single socket. 6 cores used on one Socket, Crypto algorithm: AES-128-CBC-HMAC-SHA1 vs. Intel Xeon 6152 2.1 GHz, 3x Intel® Corporation, Ethernet Controller X710 (4x10 Gbe ports per card), Lewisburg B0 Quick Assist Accelerator with PCIe Gen3 x24 links, Intel® DPDK 17.02 IPsec-secgw, Intel® QAT1.7.Upstream.L.1.0.0-15, 6 cores used. Cores, IO, packet buffer memory, and processing cores are on a single socket. 6 cores used on one Socket, Crypto algorithm: AES-128-CBC-HMAC-SHA1.

[15] Up to 2.4x claim based on TLS Web Proxy using NGINX®: Intel® Xeon® E5-2658 v4, DDR4-2133, Intel® PCH C612, Intel® 895XCC based QuickAssist Accelerator Adapter PCIe Gen3 x8 links, OpenSSL-Async (0.4.9-009) + NGINX-1.6.2 (0.1.0-008), QAT1.6.L.2.6.0-60. Cores, IO, packet buffer memory, and processing cores are on a single socket. 6 cores used on one Socket 12Cores are used, Crypto algorithm: AES-128-CBC-HMAC-SHA1 vs. Intel Xeon 6152 2.10 GHz, DDR4-2400 3x Intel® Corporation Ethernet Controller X710 (4 x10 Gbe ports per card), 1x Intel® Corporation Ethernet Controller X710 (2 x10 Gbe ports per card), PCIe x16 to 2 x8 PCIe bifurcation plugin card, Lewisburg-L B1 QuickAssist Accelerator with PCIe Gen3 x24 links, Intel®OpenSSL-1.0.1u + NGINX-1.9.6, Intel® QAT1.7.Upstream.L.1.0.0-15. Cores, IO, packet buffer memory, and processing cores are on a single socket. 6 cores used on one Socket, 20Core are used. Crypto algorithm: AES-128-CBC-HMAC-SHA1.

[16] Inference: Platform: 2S Intel® Xeon® Platinum 8180 CPU @ 2.50 GHz (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.

Deep Learning Frameworks: 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 (GoogLeNet, AlexNet, and ResNet-50), https://github.com/intel/caffe/tree/master/models/default_vgg_19 (VGG-19), and https://github.com/soumith/convnet-benchmarks/tree/master/caffe/imagenet_winners (ConvNet benchmarks; files were updated to use newer Caffe prototxt format but are functionally equivalent). Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with "numactl l".

[17] Training: Platform: 2S Intel® Xeon® CPU E5-2697 v2 @ 2.70GHz (12 cores), HT enabled, turbo enabled, scaling governor set to "performance" via intel_pstate driver, 256GB DDR3-1600 ECC RAM. CentOS Linux release 7.3.1611 (Core), Linux kernel 3.10.0-514.21.1.el7.x86_64. SSD: Intel® SSD 520 Series 240GB, 2.5in SATA 6Gb/s, 25nm, MLC.

Performance measured with: Environment variables: KMP_AFFINITY='granularity=fine, compact,1,0', OMP_NUM_THREADS=24, CPU Freq set with cpupower frequency-set -d 2.7G -u 3.5G -g performance.

Deep Learning Frameworks: Caffe: (http://github.com/intel/caffe/), revision b0ef3236528a2c7d2988f249d347d5fdae831236. 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 (GoogLeNet, AlexNet, and ResNet-50), https://github.com/intel/caffe/tree/master/models/default_vgg_19 (VGG-19), and https://github.com/soumith/convnet-benchmarks/tree/master/caffe/imagenet_winners (ConvNet benchmarks; files were updated to use newer Caffe prototxt format but are functionally equivalent). GCC 4.8.5, Intel MKL small libraries version 2017.0.2.20170110.

[18] LINPACK Scaling: Intel® Xeon® Platinum 8170 processor, 2.10 GHz 26 cores , Dual socket servers connected to a single Intel® OPA Eldorado Forest edge switch. 64 GB 2666 MHz DDR4 per node (2x16GB DDR4 per socket). Intel® Turbo Boost Technology enabled, Intel® Hyper-Threading Technology enabled. Intel(R) MPI Library for Linux* OS, Version 2017 Update 1 Build 20161016, xhpl_intel64_dynamic binary as packaged in Intel® Parallel Studio XE 2017.1.043. Intel Fabric Suite (IFS) 10.3.1.0.22. RHEL* 7.3, 3.10.0-514.el7.x86_64 kernel.

[19] Up to 1.63x Gains based on Geomean of Weather Research Forecasting - Conus 12Km, HOMME, LSTCLS-DYNA Explicit, INTES PERMAS V16, MILC, GROMACS water 1.5M_pme, VASPSi256, NAMDstmv, LAMMPS, Amber GB Nucleosome, Binomial option pricing, Black-Scholes, Monte Carlo European options. Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance. 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 information go to http://www.intel.com/performance/datacenter.

Printed in USA 0617/SC/HBD Please Recycle