

Scale Up Network and Edge Performance with a Power-Optimized and Highly Integrated Server Platform

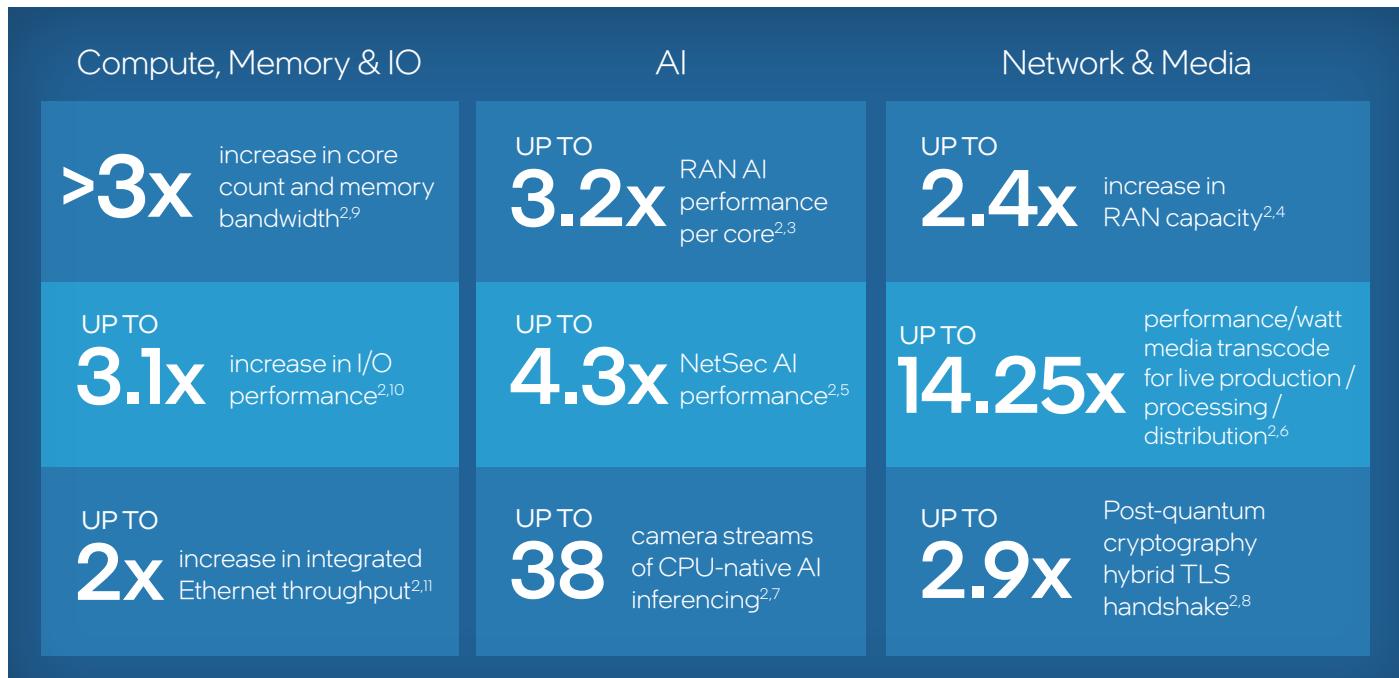
The Intel® Xeon® 6 SoC with Performance-cores (P-cores) drives up throughput for network and edge use cases with high per-core performance and several edge optimized accelerators that target a range of market segments.¹ The platform's core count scalability helps make it an excellent fit across a wide range of workloads.



Enterprises and network operators continue to realize new business value by converging networking, AI, security and other workloads at the edge. Processing data closer to where it is generated and consumed avoids the cost, delay and cybersecurity exposure associated with backhauling data to a cloud or data center. Maximum advantage for network and edge deployments requires density-optimized servers that help reduce TCO with low equipment and operating costs, while delivering high performance and advanced packet processing.

As the next-generation platform for network and edge servers, the Intel® Xeon® 6 SoC delivers the performance and flexibility to scale across a range of high-throughput, low-latency usage models. A power-efficient microarchitecture and several market-specific edge accelerators enhance performance per watt to help reduce operating costs. Availability across a wide range of core counts and feature sets with software compatibility including previous hardware generations enables solutions to deploy and scale effectively in various environments and implementations.

An optimized software ecosystem compounds the Intel Xeon 6 SoC's performance and efficiency advantages and streamlines solution time-to-production. Intel validates and fine-tunes solutions to help solve the complexities of network and edge deployments, helping customers overcome obstacles. This engineering leadership helps transform data into split-second insights, bridging the gaps between technology and business objectives.



Architecture for improved network and edge outcomes

The Intel Xeon 6 SoC with P-cores helps enterprises and network operators pursue a competitive advantage with sophisticated network and edge implementations. Solutions can scale up performance for AI and analytics workloads to help support novel functionality with larger datasets, more streams per server and more connected devices such as cameras and industrial controllers. To ensure high responsiveness and excellent user experiences, the low latency offered by the Intel Xeon 6 SoC enables sensitive workloads such as synchronous communications and near-real-time analytics. Outstanding performance per watt helps optimize TCO. The Intel Xeon 6 SoC delivers advances across the platform:

- **Expanded compute resources.** Up to 72 Performance-cores built on Intel 3 process technology for energy-efficient high throughput, complemented by enhancements to built-in accelerators and instruction sets.
- **Enhanced memory subsystem.** Up to eight channels of DDR5 memory operating at up to 6400 MT/s and an expanded L2 cache to keep more warm data close to processing resources.
- **Increased I/O throughput with lower latency.** Integrated Intel® Ethernet supporting up to 200 Gbps through up to eight ports, as well as up to 32 lanes PCIe 5.0 with CXL 2.0 and up to 16 lanes PCIe 4.0, for 48 lanes total.

To enable dense, durable designs and short-depth servers in challenging edge environments, the BGA package of the Intel Xeon 6 SoC provides options for extended temperatures.¹² Hardware-enabled security and data isolation enhance protection of intellectual property including confidential AI models.

Intel software libraries and other development tools simplify development of software optimized for Intel platform features. OpenVINO™ toolkit streamlines deployment of AI solutions and enables cross-architectural programming for flexibility that speeds up time to production and protects infrastructure investments. It enables network and edge AI solutions built to run on the Intel Xeon 6 SoC to interoperate seamlessly not only on other Intel CPUs, but also on discrete GPUs and other accelerators.

AI acceleration for fast inference on the CPU

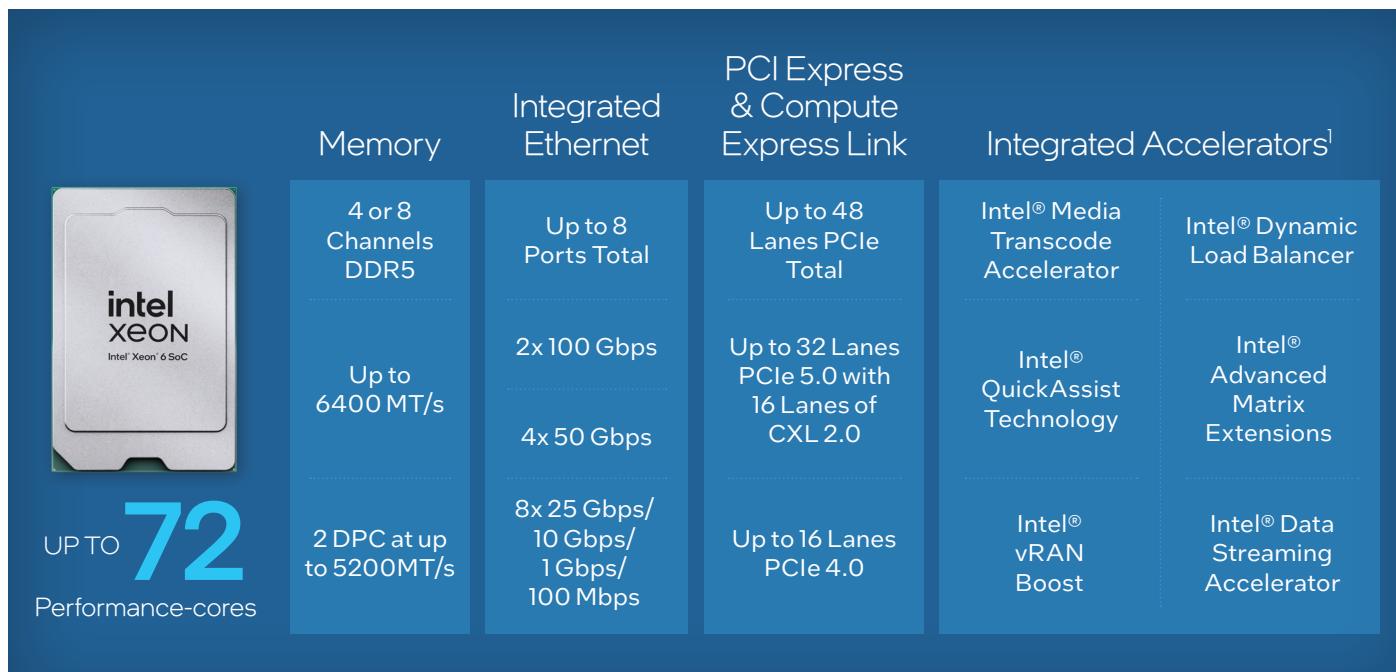
Growing AI deployments will be a key means of improving automation and efficiency in new and existing use cases for the foreseeable future. The Intel Xeon 6 SoC delivers cost efficiencies with high-performance inference on the CPU, without add-on hardware. Intel® Advanced Matrix Extensions (Intel® AMX) is a hardware technology that accelerates matrix multiplication, a core task in AI workloads.

Intel AMX uses two-dimensional registers called tiles to hold more data than conventional one-dimensional registers and performs Tile Matrix Multiplication (TMUL) on them to increase throughput. With the Intel Xeon 6 SoC, Intel AMX gains support for FP16 datatypes, an addition that accelerates matrix multiplication further in many pipelines and enables faster inference on many GPU-trained models. Support for Intel AMX is built into popular AI frameworks, including PyTorch, TensorFlow and ONNX. Accelerating these functions on the CPU enables enterprises and network operators to take advantage of existing server investments as they evolve their solutions for new capabilities and opportunities.

Cost, latency and power advantages from component integration

The Intel Xeon 6 SoC incorporates an unprecedented level of component integration, including built-in Ethernet networking and accelerators based on dedicated built-in hardware. That integration can reduce system cost by eliminating the need for discrete components, streamlining the server bill of materials.

Avoiding the need to move data across the PCIe bus to add-in cards, the Intel Xeon 6 SoC reduces system latency and power consumption. Hardware-based accelerators are generally more power-efficient than performing the same workloads on the processor cores, and where tasks are offloaded from the cores, they also free those resources for other work.



Integrated components in the Intel® Xeon® 6 SoC

Component	Benefit
Intel® Advanced Matrix Extensions (Intel® AMX)	Accelerates matrix multiplication, a key mathematical operation for AI and deep learning, by dramatically increasing the amount of data handled per computational cycle.
Integrated Intel® Ethernet	Provides flexible network connectivity for high throughput and advanced capabilities such as profile-specific performance tuning and timing for network synchronization.
Intel® Dynamic Load Balancer (Intel® DLB)	Increases network data-handling performance with the efficient dynamic distribution of processing across multiple CPU cores as the system load varies.
Intel® QuickAssist Technology (Intel® QAT)	Accelerates compression/decompression, public key cryptography and cryptographic ciphers including AES GCM/CCM/CTR, ChaChaPoly and wireless algorithms SNOW 3G and ZUC.
Intel® Media Transcode Accelerator	Built-in hardware accelerator increases performance of software-based video encoders alongside the CPU cores, using shared system memory to avoid workload overhead from memory copies.
Intel® Data Streaming Accelerator (Intel® DSA)	Improves streaming data movement and transformation operations by offloading them from the CPU cores, for more efficient use of storage and network resources.
Intel® vRAN Boost	Integrates vRAN acceleration into the CPU, helping provide higher capacity and density per server than with prior-generation processors.

Protecting data in use with confidential computing

Strong encryption protects data at rest or in transit, but data must generally be unencrypted while it is in use. Confidential computing mitigates the associated risk with hardware-based measures to isolate sensitive active data. These protections are critical to provide continuous protection of workloads in distributed networks. Confidential computing enables protected data sharing with third parties, including training AI models without exposing the training data itself. The Intel Xeon 6 SoC offers the choice of two built-in confidential computing technologies for flexibility in meeting security and regulatory requirements:

- **Intel® Software Guard Extensions (Intel® SGX)** is the most researched, updated and deployed confidential computing technology in data centers on the market today, with the smallest trust boundary of any confidential computing technology in the data center.
- **Intel® Trust Domain Extensions (Intel® TDX)** offers confidentiality at the virtual machine (VM) level, isolating the guest OS and all VM applications from the cloud host, hypervisor and other VMs on the platform. Intel TDX is designed to streamline the deployment and management of confidential VMs at scale.

Ecosystem Support Through Intel® Industry Solution Builders

The [Intel Industry Solution Builders](#) program helps partners innovate and adapt to evolving business, technology and end-user needs, effectively and cost-efficiently. The program provides members with a variety of technical enablement options such as hands-on support from subject matter experts, access to virtual testing and optimization labs, training, tools and other resources.

To help advance the state of the art in network engineering, Intel Industry Solution Builders provides partner resources such as developer kits, verified reference implementations and tools. It offers expertise and co-engineering guidance with Intel's full suite of open standards-based networking software to help network and edge operators build, deploy, manage and secure solutions more effectively.

Sample deployment use cases

The performance, efficiency and flexibility of the Intel Xeon 6 SoC make it a compelling foundation for network and edge solutions, including those evolving to make greater use of AI. In deployment, it scales up performance, drives efficiency and consolidates workloads to boost density.

Virtual radio access networks (vRANs)

 The Intel® Xeon® 6 SoC's initial and ongoing cost advantages include giving mobile network operators (MNOs) the ability to shrink infrastructure requirements for most vRAN sites to a single server. Intel® vRAN Boost improves quality of service and bandwidth efficiency dramatically by improving transmission reliability, even over unreliable communication channels. Hardware acceleration for AI inference can optimize vRAN deployments with intelligent traffic steering, cybersecurity and other functions that deliver increased value from server investments.

Secure access service edge (SASE) and security appliances

 The Intel Xeon 6 SoC enhances performance and TCO as SASE services move from public cloud infrastructure to private points of presence (PoPs). Accelerated AI inference enhances network traffic steering, adjusting data paths dynamically to improve throughput and reliability. Security solutions can use deep learning to establish expected network behaviors and then detect anomalies from that baseline that suggest malicious activity. The Intel Xeon 6 SoC is also the foundation for the next generation of security appliances, to deliver AI-assisted threat detection and response.

Media transcoding for live video services

 The Intel Xeon 6 SoC with built-in Intel Media Transcode Accelerator provides the features and capabilities to support use cases that include low-latency live streaming, high-quality broadcast, video-on-demand encoding, high-density Virtual Desktop Infrastructure (VDI), responsive cloud gaming and other media-intensive services. It provides high-quality, low-latency transcoding of AVC, HEVC and AV1 video content.

Edge AI

 The Intel Xeon 6 SoC provides higher CPU-native inference performance than predecessors, enabling critical functionality at the edge that can reduce latency and enhance performance for end-user applications. Business and security functions benefit from AI acceleration in hardware, as well as optimizations for popular frameworks. The Intel Xeon 6 SoC handles demanding AI requirements without the added complexity and expense of discrete accelerators or backhauling data to a data center or cloud, enabling better TCO.

Federal and aerospace



The Intel Xeon 6 SoC helps achieve extended mission goals to replace proprietary, expensive systems with open technologies, including commercial off the shelf (COTS) servers. The highly integrated SoC comes in a BGA package and is engineered specifically for durable, reliable, space-constrained equipment designs, targeting federal and aerospace usages.

Retail



The Intel Xeon 6 SoC powers fast, efficient, scalable applications at retail locations, with high density to help support more cameras and enterprise workloads on the same enterprise network server. Built-in AI acceleration drives up inference performance without discrete GPUs or other accelerators, to help improve TCO, better understand customer needs and deliver better shopping experiences to build sales and loyalty.

Specifications

SKU	Cores	TDP (Watts)	Base Frequency (GHz)	All-Core Turbo (GHz)	Uncore	DDR5 Memory Speed	Integrated Ethernet	Intel® QAT Support	Intel® vRAN Boost	Intel® Media Transcode Accelerator	Intel® AMX	Extended Temp
6726P-B	42	235	2.3	2.9	1.6	6400	200 Gbps	Yes	Yes	—	Yes	—
6716P-B	40 SST	235	2.3	2.9	1.6	6400	200 Gbps	Yes	Yes	—	Yes	—
6706P-B	40	235	2.5	2.9	1.6	6400	200 Gbps	Yes	Yes	—	Yes	—
6563P-B	38	235	2.4	3.1	1.8	6400	200 Gbps	Yes	—	Yes	Yes	—
6556P-B	36	215	2.3	2.9	1.6	6400	200 Gbps	Yes	Yes	—	Yes	Optional
6553P-B	36	235	2.6	3.4	1.6	6400	200 Gbps	Yes	—	Yes	Yes	—
6546P-B	32	195	2.3	2.9	1.6	6400	200 Gbps	Yes	Yes	—	Yes	—
6543P-B	32	160	2	2.6	1.4	5600	100 Gbps	Yes	—	—	Yes	—
6533P-B	32	205	2.2	2.6	1.6	5600	None	Yes	—	—	Yes	—
6523P-B	24	175	2.5	3.3	1.8	5600	100 Gbps	Yes	—	—	Yes	—
6516P-B	20	145	2.3	2.9	1.6	4800	200 Gbps	Yes	Yes	—	Yes	—
6513P-B	20	130	2	2.6	1.4	5600	100 Gbps	Yes	—	—	Yes	—
6503P-B	12	110	2	2.6	1.4	4800	100 Gbps	Yes	—	—	Yes	—

Additional higher core count SKUs available in Q4'25

Learn More

www.intel.com/xeon
networkbuilders.intel.com



¹ Availability of accelerators varies depending on SKU. Visit the Intel® Product Specifications page for additional product details.

² Versus prior generation.

³ See [7ND34] at intel.com/processorclaims: Intel® Xeon® 6. Results may vary.

⁴ See [7ND21] at intel.com/processorclaims: Intel Xeon 6. Results may vary.

⁵ See [7ND30] at intel.com/processorclaims: Intel Xeon 6. Results may vary.

⁶ See [7ND32] at intel.com/processorclaims: Intel Xeon 6. Results may vary.

⁷ See [7ND20] at intel.com/processorclaims: Intel Xeon 6. Results may vary.

⁸ See [7ND33] at intel.com/processorclaims: Intel Xeon 6. Results may vary.

⁹ Intel® Xeon® 6 SoC 72 cores, 8 channel DDR5 6400 MT/s versus Intel® Xeon® D-2899NT 22 cores, 4 channel DDR4 3200 MT/s.

¹⁰ See [7ND25] at intel.com/processorclaims: Intel Xeon 6. Results may vary.

¹¹ See [7ND26] at intel.com/processorclaims: Intel Xeon 6. Results may vary.

¹² Commercial temp: 0°C to ~85°C Tcase. Extended temp: -40°C to ~85°C Tcase.

No product or component can be absolutely secure.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Your costs and results may vary.

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

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a nonexclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

© Intel Corporation. Intel, the Intel logo and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others. 0225/DK/MESH/PDF 356903-001US