

Efficient Performance for General-Purpose Workloads

Intel Xeon 6700/6500-series processors with P-cores provide flexible high performance with cost and power efficiency, enabling TCO gains across the broadest range of workloads.



Achieve up to 5:1 average server consolidation at maintained performance with an average total cost of ownership (TCO) gain of 40 percent compared to 2nd Gen Intel Xeon Scalable processors.¹

In today's dynamic business landscape, organizations depend on a variety of general-purpose workloads to keep their operations running smoothly and to provide key services for customers. These workloads, including storage, networking, database, edge, and many others, require servers with efficient processors that offer high performance per core.

Intel Xeon 6700/6500-series processors with Performance-cores (P-cores) provide a cost-effective, high-performance solution. Compared to the previous generation, these efficient processors offer higher socket performance over the same thermal design power (TDP) range. They also provide increased memory bandwidth and higher input/output (I/O), making them ideal for a wide range of workloads. Specifically, Intel Xeon 6700/6500-series processors enable organizations to:^{2,3,4}

- **Accelerate AI with Intel® Advanced Matrix Extensions (Intel® AMX).** Intel AMX speeds up inferencing for INT8 and BF16 data formats, and it offers support for FP16-trained models, providing up to 2,048 floating point operations per cycle per core for INT8 and up to 1,024 floating point operations per cycle per core for BF16/FP16.
- **Improve memory input with fast Gen 3 DDR5 Multiplexed Rank DIMMs (MRDIMMs)** that can deliver more than 25 percent greater memory bandwidth than RDIMMs, with an expected data-transfer rate of up to 8,000 megatransfers per second (MT/s). Intel Xeon 6700/6500-series processors with P-cores also support DDR5-6400 high-speed memory.
- **Efficiently accelerate mainstream workloads.** Intel Xeon 6700/6500-series processors with P-cores provide up to 86 cores per socket and one-to-eight socket scalability, with up to 336 MB L3 cache per processor and exceptionally low latency at large L3 access sizes. Intel® Advanced Vector Extensions 512 (Intel® AVX-512) can be used to boost the speed of vector math, which is common to high-performance computing (HPC) and traditional AI workloads.
- **Enhance privacy and control over data with increased confidentiality at the virtual machine (VM) level.** With Intel® Trust Domain Extensions (Intel® TDX) confidential VMs, the guest operating system (OS) and VM applications are isolated from access by the cloud host, hypervisor, and other VMs on the platform. And with Intel TDX 2.0, support is extended to the trusted execution environment (TEE) for device I/O, enabling encrypted communications with connected PCIe devices via Intel TDX Connect. Intel is actively collaborating with a broad ecosystem of partners to [rapidly expand adoption](#) of this security-enabling technology.

Efficient performance in a single-socket platform

While a two-socket platform is still a popular mainstream solution, many businesses are shifting to single-socket platforms to lower their power and cooling costs and to reduce their total cost of ownership (TCO). But that shift is only possible with systems and processors that can provide high I/O performance.

The single-socket configurations of Intel Xeon 6700/6500-series processors with P-cores offer 50 percent more PCIe lanes/socket, compared to a multi-socket platform. The added PCIe lanes help eliminate power and latency inefficiencies, resulting in significantly greater overall I/O performance for peripherals and storage devices.

These efficient processors also provide significant flexibility for rack-power and density configurations, enabling cost optimizations that can help lower TCO. And the greater core density offered by these single-socket platforms can lead to improved performance/watt for greater overall efficiency that can lower operating expenses (OpEx). As a result, single-socket Intel Xeon 6 processors with P-cores are ideal for a broad array of commonly used, critical business workloads, as shown in Tables 1 and 2.

Table 1. High PCIe lane count and improved I/O per socket make single-socket Intel Xeon 6700/6500-series processors with P-cores ideal for many common workloads that businesses rely on

Data center mainstream use cases	Dependency on improved I/O per socket
Storage	High data transfer rates between devices/storage drives to enable higher capacity storage solutions, including NVMe Express (NVMe) caching and high fan-out storage that only requires modest levels of compute
Scale-out databases*	Increased I/O required to retrieve and update information for read/write operations common to scale-out databases
Virtualization	Multiple VMs running on the same physical hardware accessing shared storage simultaneously for database workloads
Virtual desktop infrastructure (VDI)	Large numbers of virtual desktops running simultaneously to support grocery, retail, and many other use cases
Mixed AI workloads/AI host CPUs	Rapid offloading of tasks to achieve optimal processor and accelerator performance

*E-cores offer another high-efficiency option for scale-out databases. See the [Intel Xeon website](#) for more information and to compare features.

Table 2. Intel Xeon 6700/6500-series processors with P-cores provide single-socket, high I/O solutions that are well-suited to edge use cases

Edge use cases	Dependency on improved I/O per socket
Wireless core/networking	Greater I/O bandwidth for consistent transfer of large datasets
AI cybersecurity	Ability to apply built-in AI capabilities for real-time network traffic pattern recognition
Content delivery networks (CDNs)	Improved balance between memory, CPU, and caching support via NVMe drives
Internet of Things (IoT)	Peripheral and accelerator requirements and the ability to quickly send and receive data

Secure and efficient edge workloads

As the demand for seamless connectivity between data center workloads and the edge grows, businesses need solutions that bridge the gap while maintaining high performance and efficiency. Intel Xeon 6 processors are purpose-built to meet the demanding requirements of edge networking workloads, delivering exceptional performance, scalability, and energy efficiency. With advanced AI capabilities, Intel Xeon processors enable real-time insights and threat detection that is critical for latency-sensitive edge applications. Their support for zero-touch provisioning and automated lifecycle management further simplifies deployment and maintenance, empowering businesses to optimize their edge infrastructures with ease.

Performance data

Up to **1.53x** higher ResNet-50 batch inference performance (INT8) with 33 percent fewer cores compared to an AMD EPYC processor.⁶

Up to **2.45x** better CDN video-on-demand performance and up to 1.88x better performance/watt compared to prior generations.⁸

Up to **40%** average performance gain across broad workloads, compared to the previous generation.¹

Up to **1.76x** higher query throughput compared to prior generations when running an IBM Db2 big data insights (BDI) workload.⁵

Up to **1.54x** higher High Performance Conjugate Gradient (HPCG) performance with MRDIMMs, compared to the previous generation.⁷

Up to **45%** lower TCO for new data center deployments when using Intel Xeon 6 processors with P-cores versus AMD EPYC processors.⁹

Up to **5:1** average server consolidation at maintained performance with an average TCO gain of 40 percent compared to 2nd Gen Intel Xeon Scalable processors.¹

Highlight technologies

The innovative microarchitecture of Intel Xeon 6700/6500-series processors with P-cores delivers the following advanced features and benefits:

- Up to 86 cores in a single socket, enabling ultra-high-density compute performance and scalability.
- Support for up to eight sockets to support in-memory databases and mission-critical workloads.
- Rich one-socket systems support storage systems by providing high I/O lanes and the ability to attach a greater number of drives for a lower TCO.
- Up to eight memory channels with 6,400 MT/s of DDR5 memory and 8,000 MT/s of MRDIMMs. MRDIMMs are capable of providing more than 25 percent additional memory bandwidth compared to standard DDR5 DIMMs, supporting bandwidth-constrained use cases found in AI and HPC.⁴
- Intel® Flat Memory Mode helps expand system memory and improve TCO when using low-cost memory, such as DDR4 with Compute Express Link (CXL) 2.0.
- Up to 64 lanes of CXL 2.0 with data-transfer rates of up to 32 gigatransfers per second (GT/s) per lane and with support for CXL capabilities such as memory expansion and sharing, including Type 3 devices.
- Intel® Ultra Path Interconnect (Intel® UPI) 2.0 provides up to 24 GT/s of inter-socket bandwidth—a 20 percent increase over the prior generation.
- Up to 88 lanes of PCIe 5.0 for two-socket servers, with options of up to 136 lanes for one-socket server designs, to allow for significant I/O add-in components including accelerators, network adapters, storage controllers, and storage.
- Intel AMX provides up to 16x more multiply accumulate (MAC) operations than Intel AVX-512 for BF16- and FP16-based models to enhance AI performance.
- Intel AVX-512 encompasses unique instructions and two 512-bit fused-multiply add (FMA) units per core, boosting the speed of vector mathematics common to AI, HPC, and database workloads.
- Intel® QuickAssist Technology (Intel® QAT) allows offload of bulk cryptography and compression to accelerate networking and storage.
- Intel® Data Streaming Accelerator (Intel® DSA) 2.0 provides double the throughput for a given workload compared to Intel DSA 1.0, which greatly improves data offloading abilities while enabling transform operations such as move, fill, compare, cyclic redundancy checking (CRC), data integrity field (DIF), delta, and flush.
- Intel® In-Memory Analytics Accelerator (Intel® IAA) 2.0 allows offload of memory compression and decompression, scan and filter functions, and CRC, with significantly improved data bandwidth over Intel IAA 1.0.
- Intel® Dynamic Load Balancer (Intel® DLB) enables dynamic distribution of network packet processing and offload of reordering operations.
- Intel TDX version 2, with support for up to 2,048 encryption keys per CPU, extends Intel TDX support to the TEE for device I/O, enabling encrypted communications with connected PCIe devices via Intel TDX Connect.
- Built-in security meets FIPS 140-3 Cryptographic Algorithm Validation Program (CAVP) certification.

The complete list of Intel Xeon 6 processor families is shown in Table 3.

Table 3. The Intel Xeon 6900-, 6700-, 6500-, and 6300-series processors are built to provide ideal performance for a wide range of use cases

Series	Designed for
Intel Xeon 6900-series processors	Maximum performance ideal for the most demanding cloud, AI, and HPC environments
Intel Xeon 6700-series processors	Enhanced performance ideal for a wide array of data center and telco environments
Intel Xeon 6500-series processors	Essential performance ideal for mainstream server and edge environments
Intel Xeon 6300-series processors	Entry-level performance ideal for small/medium business environments

Optimize your data center workloads

Intel Xeon 6700/6500-series processors with P-cores are ideal for handling the common workhorse applications of most organizations. With flexible high performance plus cost and power efficiency, this family of Intel Xeon processors is built to take on the widest range of workloads while bringing significant TCO gains to the data center.

And with 136 PCIe lanes per socket, the single-socket variant of the Intel Xeon 6700/6500-series processors with P-cores offers even greater rack density and configuration flexibility, providing outstanding performance-per-watt efficiency, low latency, and high I/O performance for storage, wireless core, networking, CDNs, or other general-purpose computing needs.

Learn more

- Learn more about [Intel Xeon 6 processors](#).
- Discover the [best data center refresh options](#) available for your organization.
- Explore [SKU details and pricing](#) for Intel Xeon 6 processors with P-cores
- Explore [developer tools](#) for Intel Xeon 6 processors.



¹ Workload geomean. See [intel.com/processorclaims](https://www.intel.com/processorclaims): Intel® Xeon® 6. Results may vary.

² See [9A2] at [intel.com/processorclaims](https://www.intel.com/processorclaims): Intel Xeon 6. Results may vary.

³ See [9H10] at [intel.com/processorclaims](https://www.intel.com/processorclaims): Intel Xeon 6. Results may vary.

⁴ In comparison to DDR5 6,400 RDIMMs.

⁵ Up to 1.76x higher query throughput with an Intel Xeon 6787P processor, compared to a 4th Gen Intel Xeon Platinum 8480+ processor, when running an IBM Db2 BDI workload: **Baseline:** 1 node, 2 x Intel Xeon Platinum 8480+ processor, 56 cores, 350 W thermal design power (TDP), Intel® Hyper-Threading Technology (Intel® HT Technology) on, Intel® Turbo Boost Technology on, NUMA 2, 1,024 GB total memory (16 x 64 GB DDR5 4,800 MT/s), BIOS 3B08.TEL3P1, microcode 0x2b0005d1, 2 x Intel® Ethernet Controller X710 for 10GBASE-T, 1 x 894.3 GB Intel SSDSC2KG96, 4 x 1.8 TB Intel SSDPE2KX020T8, Ubuntu 24.04.1 LTS, 6.8.0-47-generic, running IBM Db2 v11.5.9 and IBM BDI kit v0.9, tested by Intel as of October 22, 2024. **New:** 1 node, 2 x Intel Xeon 6787P processor, 86 cores, 350 W TDP, Intel HT Technology on, Intel Turbo Boost Technology on, NUMA 2, 1,024 GB total memory (16 x 64 GB DDR5 6,400 MT/s), BIOS BHSDCRB1.IPC.3544.P07.2410011104, microcode 0x1000311, 2 x Intel Ethernet Controller X710 for 10GBASE-T, 1 x 28.9 GB USB flash drive, 4 x 3.5 TB Samsung MZWLJ3T8HBL5-00007, 1 x 1.7 TB Samsung MZWLRT9HBJR-00007, Ubuntu 24.04 LTS, 6.8.0-49-generic, running IBM Db2 v11.5.9 and IBM BDI kit v0.9, tested by Intel as of November 27, 2024.

⁶ See [7A28] at [intel.com/processorclaims](https://www.intel.com/processorclaims): Intel® Xeon® 6. Results may vary.

⁷ See [7H2] at [intel.com/processorclaims](https://www.intel.com/processorclaims): Intel® Xeon® 6. Results may vary.

⁸ See [7N24] at [intel.com/processorclaims](https://www.intel.com/processorclaims): Intel® Xeon® 6. Results may vary.

⁹ See [9T22] at [intel.com/processorclaims](https://www.intel.com/processorclaims): Intel® Xeon® 6. Results may vary.

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

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See configuration disclosure for additional details.

No product or component can be absolutely secure.

Your costs and results may vary.

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

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

© 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.