



Intel Accelerator Engines

Accelerating workloads on 4th Gen Intel® Xeon® Scalable processors and Intel® Xeon® CPU Max Series

Jan. 10, 2023 — 4th Gen Intel® Xeon® Scalable processors and Intel® Xeon® CPU Max Series, launched today by Intel, feature the latest Intel® Accelerator Engines for AI, high performance computing (HPC), security, network, analytics and storage. Built-in acceleration is an alternative, more efficient way to achieve higher performance than growing the CPU core count. The latest Intel Accelerator Engines, high bandwidth memory and software optimizations help improve performance and power efficiency across targeted workloads and can lead to cost savings through better CPU utilization.

Built-in Accelerators

Intel® Advanced Matrix Extensions (Intel® AMX)

With new Intel AMX, AI performance on the CPU expands to include fine-tuning and small and medium deep learning training models. Intel AMX is a built-in accelerator that improves the performance of deep learning training and inference. It is ideal for workloads like natural language processing, recommendation systems and image recognition.

Intel® QuickAssist Technology (Intel® QAT)

By offloading encryption, decryption and compression, Intel QAT— now integrated as a built-in accelerator — helps free up processor cores so systems can serve a larger number of clients or use less power. With Intel QAT, 4th Gen Intel Xeon Scalable processors are the highest-performance CPUs that can compress and encrypt in a single data flow.

Intel® Data Streaming Accelerator (Intel® DSA)

Intel DSA drives high performance for storage, networking and data-intensive workloads by improving streaming data movement and transformation operations. Designed to offload the most common data movement tasks that cause overhead in data center-scale deployments, Intel DSA helps speed up data movement across the CPU, memory and caches, as well as all attached memory, storage and network devices.

Intel® Dynamic Load Balancer (Intel® DLB)

Intel DLB helps improve system performance related to handling network data on multicore Intel® Xeon® Scalable processors. It enables the efficient distribution of network processing across multiple CPU cores/threads and dynamically distributes network data across multiple CPU cores for processing as the system load varies. Intel DLB also restores the order of networking data packets processed simultaneously on CPU cores.

Intel® In-Memory Analytics Accelerator (Intel® IAA)

Intel IAA helps run database and analytics workloads faster, with potentially greater power efficiency. This built-in accelerator increases query throughput and decreases the memory footprint for in-memory database and big data analytics workloads. Intel IAA is ideal for in-memory databases, open source databases, and data stores like RocksDB and ClickHouse.

**Intel® Advanced Vector Extensions 512 (Intel® AVX-512)**

Intel AVX-512 is the latest x86 vector instruction set, with up to two fused-multiply add (FMA) units and other optimizations to accelerate performance for demanding computational tasks, including scientific simulations, financial analytics, and 3D modeling and analysis.

Intel® Advanced Vector Extensions 512 (Intel® AVX-512) for vRAN

Intel AVX-512 for virtualized radio access network (vRAN) is designed to deliver greater capacity at the same power envelope for vRAN workloads. This helps communications service providers increase their performance per watt to meet critical performance, scaling and energy efficiency requirements.

Intel® Crypto Acceleration

Intel Crypto Acceleration reduces the impact of implementing pervasive data encryption and increases the performance of encryption-sensitive workloads, like secure sockets layer (SSL) web servers, 5G infrastructure and VPNs/firewalls.

Intel® Speed Select Technology (Intel® SST)

Intel SST is designed to grant more active and expansive control over CPU performance. Intel SST enables improved server utilization and reduced qualification costs by allowing customers to configure a single server to match fluctuating workloads. The result is one flexible server with multiple configurations, leading to improved TCO.

Intel® Data Direct I/O Technology (Intel® DDIO)

Intel DDIO helps remove inefficiencies by enabling direct communication between Intel® Ethernet controllers and adapters and host processor cache. Eliminating frequent visits to main memory can help reduce power consumption, provide greater I/O bandwidth scalability and reduce latency.

Intel® Security Engines**Intel® Software Guard Extensions (Intel® SGX)**

With Intel SGX, organizations can unlock new opportunities for business collaboration and insights — even with sensitive or regulated data. Intel SGX is the most researched, updated and deployed confidential computing technology in data centers on the market today, with the smallest trust boundary. Confidential computing improves the isolation of sensitive data with enhanced hardware-based memory protections. Support for Intel SGX on Intel Xeon CPU Max Series is on DDR flat mode only.

Intel® Trust Domain Extension (Intel® TDX)

Intel TDX is a new capability available through select cloud providers in 2023 that offers increased confidentiality at the virtual machine (VM) level, enhancing privacy and control over data. Within an Intel TDX confidential VM, the guest OS and VM applications are isolated from access by the cloud host, hypervisor and other VMs on the platform.

Intel® Control-Flow Enforcement Technology (Intel® CET)

Intel CET provides enhanced hardware-based protections against return-oriented and jump/call-oriented programming attacks, two of the most common software-based attack techniques. Using this technology helps shut down an entire class of system memory attacks that long evaded software-only solutions.

Availability of accelerators varies depending on SKU. Visit the [Intel product specifications page](#) for product details and acceleration support for the Intel Xeon CPU Max Series.

About Intel

Intel (Nasdaq: INTC) is an industry leader, creating world-changing technology that enables global progress and enriches lives. Inspired by Moore's Law, we continuously work to advance the design and manufacturing of semiconductors to help address our customers' greatest challenges. By embedding intelligence in the cloud, network, edge and every kind of computing device, we unleash the potential of data to transform business and society for the better. To learn more about Intel's innovations, go to newsroom.intel.com and intel.com.

Performance varies by use, configuration and other factors. Learn more at intel.com/performanceindex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration 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. Intel contributes to the development of benchmarks by participating in, sponsoring, and/or contributing technical support to various benchmarking groups, including the BenchmarkXPRT Development Community administered by Principled Technologies.