Product Brief

5th Gen Intel® Xeon® Scalable Processors for Edge

Accelerate Demanding and Evolving Edge Workloads with Built-in AI and Security

5th Gen Intel® Xeon® Scalable processors for edge deliver impressive compute, faster memory, and integrated accelerators to boost AI performance and energy efficiency.

Achieve incredible performance per watt[1] for your demanding and emerging edge workloads, including AI. 5th Gen Intel® Xeon® Scalable processors feature built-in accelerators, optimized software, and enhanced telemetry capabilities to enable more-manageable, more-efficient deployments. Deploy high-value, specialized workloads with enhanced security at the edge while lowering TCO and increasing energy efficiency with five integrated accelerators[2] and Optimized Power Mode.

Advance sustainability goals with select SKUs that support Intel® Speed Select Technology (Intel® SST), which allows you to scale compute power as low as 105W to meet edge workload requirements. Intel® SST provides you with flexible manageability and enhanced control over CPU performance to help you optimize TCO. Feel more confident in the protection and privacy of sensitive and regulated workloads with built-in security features that enable cryptographic isolation for VMs, integrity protection, and memory confidentiality. Build on existing investments with a solution that's fully compatible with 4th Gen Intel® Xeon® Scalable processors, requiring only minimal testing and validation.

Boost AI performance and energy efficiency with integrated accelerators

Stay ahead as business evolves with a platform that delivers better AI inferencing performance, power efficiency, and lower TCO with integrated accelerators and the Optimized Power Mode feature. Built-in accelerators[2] free up CPU cycles for other important workloads. This allows you to push performance further for targeted AI use cases across retail, healthcare, industrial, and energy sectors.

Built-in accelerators in 5th Gen Intel® Xeon® Scalable processors[2]

- **Intel® Advanced Matrix Extensions (Intel® AMX):** Exceptional AI performance for deep-learning training and inference
- **Intel® Data Streaming Accelerator (Intel® DSA):** Optimized streaming data movement and transformation operations
- **Intel® QuickAssist Technology (Intel® QAT):** Fast data compression and encryption to help reduce system resource consumption and improve TCO
- **Intel® In-Memory Analytics Accelerator (Intel® IAA):** Accelerates database and analytics workloads, ideal for in-memory databases
- **Intel® Volume Management Device (Intel® VMD):** Integrated control and management for NVMe storage acceleration; supports Intel® Virtual RAID on CPU (Intel® VROC) to boost performance and reliability

What's new

- Up to 16 percent memory bandwidth improvement with faster DDR5 memory[3]
- Up to 3x increased last-level cache (LLC) vs. 4th Gen Intel Xeon processors[4]
- Intel® Trust Domain Extensions (Intel® TDX) for confidential computing in all SKUs
- Intel® Seamless Firmware Update for improved uptime

5th Gen Intel® Xeon® Scalable processors

PERFORMANCE

Up to
1.59x
average performance gain[5]
vs. 3rd Gen Intel Xeon Gold 6348 processor

Up to
1.29x
average performance per watt gain[1]
vs. 3rd Gen Intel Xeon Gold 6348 processor

ARTIFICIAL INTELLIGENCE

Up to
2.81x
higher real-time inference performance for image classification[6]
vs. 3rd Gen Intel Xeon Gold 6348 processor

Up to
5.28x
higher real-time inference performance for object detection[7]
vs. 3rd Gen Intel Xeon Gold 6348 processor

Up to
3.13x
higher real-time inference performance for image classification[8]
vs. AMD EPYC 9334

For workloads and configurations, visit intel.com/processorclaims: 5th Gen Intel® Xeon® Scalable processors. Results may vary.

Drive more multitasking with faster DDR5 speeds and a larger shared cache

Help overcome data bottlenecks with up to 16 percent memory bandwidth improvement[3] and up to 3x increased LLC over previous-generation processors.[4] 5th Gen Intel Xeon Scalable processors deliver eight channels of DDR5 and up to 16 DIMMs per socket and support up to 5600MT/s at one DIMM per channel (DPC) or 4400 MT/s at two DPC.

More manageability, visibility, and control for better system performance and efficiency

Unlock system performance and operational efficiencies with Intel® Resource Director Technology (Intel® RDT) and Intel® Platform Monitoring Technology (Intel® PMT) on 5th Gen Intel Xeon Scalable processors. Intel RDT brings your business more visibility and control over how shared resources such as LLC and memory bandwidth are used by applications, VMs, and containers. With Intel RDT integrated into every processor, you get workload consolidation density, performance consistency, and dynamic service delivery, helping to drive efficiency and flexibility while reducing overall TCO.

Intel PMT is a unified solution and standard method for discovering, collecting, and managing enhanced telemetry data across components. It offers a more comprehensive and consistent view of infrastructure performance, utilization, and events in real time to enable automation and easy orchestration of workloads across server, storage, and network resources. With Intel PMT, businesses can leverage standardized telemetry to unlock new insights to drive enhanced performance and resource management.

Built-in trust for confidential computing in edge deployments

Embrace confidential computing to address your complex security needs and meet regulatory requirements with built-in security features on 5th Gen Intel Xeon Scalable processors. Featuring Intel TDX and Intel® Software Guard Extensions (Intel® SGX) security technologies, these processors enhance your control over your most sensitive IP and regulated workload data at the edge.

Intel SGX is designed to enhance data protection at rest, in motion, and in use. It enables application isolation and allows users to create access-restricted enclaves to help keep data confidential and secure. Intel SGX also:

- Provides users unique granular-level control and protection at both the enclave and application level
- Adds another layer of defense against software and physical integrity attacks by reducing the attack surface
- Allows users to request verification that applications have not been compromised and the processor is running on the latest security update

Intel TDX provides additional data protection by enabling isolation and confidentiality at the VM level. Within an Intel TDX-enabled confidential VM, the guest OS and VM applications and microservices are isolated from access by the cloud host, hypervisor, and other VMs on the platform. Intel TDX also:

- Extends security capabilities across the whole VM to create a larger trust boundary
- Protects against a broad range of software attacks
- Simplifies the process of migrating existing VMs to a trusted execution environment (TEE)

Add more devices and increase bandwidth with faster UPI speeds and 80 lanes of PCIe 5.0 and PCIe 4.0

Get the flexibility to customize your platform and connect additional devices, accelerators, or faster memory. The 5th Gen platform supports your I/O needs of today and tomorrow with up to 80 lanes of PCIe 5.0 and PCIe 4.0 support and up to four Intel® Ultra Path Interconnect (Intel® UPI) 2.0 links with speeds of up to 20 GT/s.

Maximize the value of your investments

With 5th Gen Intel Xeon Scalable processors, upgrading from the last generation no longer means you lose your previous technology investments. Designed to save time, money, and energy, 5th Gen Intel Xeon processors are software and pin compatible with 4th Gen Intel Xeon processors for faster and easier drop-in upgrades, minimal validation and testing, and shorter development time frames. Intel® Seamless Firmware Updates and enhanced telemetry capabilities mean you experience less downtime and your systems run more efficiently.

Drive ROI with long-life availability[9] and long-use-life SKUs[10]

Customers in the healthcare, public sector, industrial, and energy industries benefit from a steady supply of inventory and equipment designed for continuous operation. 5th Gen processors meet those needs by offering extended availability of up to 10 years for edge SKUs[9] and long-life reliability for up to 10 years on select SKUs.[10] These features not only help customers replenish their inventory but also help prolong the life of systems, extend device life cycles, and drive more value from investments and lengthy certification cycles. With industry-leading support, customers can be confident their systems will endure in the field for many years to come.

Key features

Performance and operational efficiency

- Built-in accelerators[2]: Intel AMX, Intel DSA, Intel QAT, Intel IAA, and Intel VMD
- Larger shared cache[2]
- Processor base power between 105W and 270W[11]
- Support for Intel® Distribution of OpenVINO™ toolkit to accelerate AI inferencing workloads

Memory and I/O

- DDR5 memory speeds up to 5600 MT/s (1DPC)/ 4400MT/s (2DPC)[2]
- 24 GB DRAM for efficient memory use
- 8x DDR5 memory channels per CPU
- Intel® UPI speeds up to 20 GT/s
- CXL Types 1, 2, and 3
- Up to 80 lanes: PCIe 5.0 (x16, x8, x4), PCIe 4.0 (x2)
- Intel® Ethernet network controllers and adapters

Security

- Intel SGX for data protection at rest, in motion, and in use, as well as application isolation and access-restricted enclaves to help keep data confidential and secure
- Intel TDX for VM-level confidentiality and isolation that removes access by the cloud host, hypervisor, and other VMs on the platform

Energy efficiency

- Optimized Power Mode, which is user configurable in the platform BIOS to help achieve incremental power savings for select workloads
- Lower-power options to help lower TCO and advance sustainability goals

Discrete graphics

- Support for Intel® Data Center GPU Flex Series

Software and OS support

- Intel® oneAPI Toolkits
- Linux: Red Hat Enterprise Linux 8.4 LTS, 8.5, and later;[12] SUSE Enterprise Linux SLE 15 SP3 and later;[12] Ubuntu 21.10, 22.04 LTS, and later;[12] Alibaba Cloud Linux[12]
- VMware ESXi[13]
- Microsoft full compatibility:[14] Long-Term Servicing Channel (LTSC) Windows Server 2022, Azure Stack (HCI) v22H2 and later
- Microsoft limited compatibility:[14,15] Long-Term Servicing Channel (LTSC) Windows Server 2019, Windows 10 (CSME or SPS FW), Azure Stack HCI v20H2

Longevity, investment, and value

- Select SKUs that support extended availability of up to 10 years for edge SKUs[9] and long-life reliability for up to 10 years on select SKUs[10]
- Server reliability of up to five years
- Drop-in compatibility with 4th Gen Intel Xeon Scalable processors
- Ecosystem support to minimize required validation for end customers
- Intel Resource Director Technology (Intel RDT), Intel Platform Monitoring Technology (Intel PMT), and Intel Seamless Firmware Update

Use cases

Healthcare and life sciences: Improve medical imaging accuracy, enable faster scan results, and enhance the security of patient data

Applications: Medical imaging, genomics, biomedical manufacturing, digital pathology, health IT

- A larger shared cache[2] and integrated AI accelerator engines[2] deliver incredible performance per watt for AI in hospitals, laboratories, and medical manufacturing.
- Intel TDX and Intel SGX for VM and application isolation help protect sensitive patient data.
- Long-life availability[9] extends device life cycles to drive more value from lengthy certification cycles.

Retail: Provide frictionless shopping journeys and power-efficient AI video analytics in a variety of store environments

Applications: Frictionless retail solutions, vision checkout, transactional back-end servers, VDI servers, and micro fulfillment centers

- The integrated Intel AMX accelerator enables faster training and inference performance without a discrete GPU.
- The portfolio enables you to target a scalable range of solutions from micro markets to big box retailers without a discrete GPU and premium in-store experiences with a discrete GPU.
- A larger shared cache[2] and higher UPI speeds deliver incredible performance per watt for AI in retail workloads.

Industrial and energy: Run compute-intensive automation in less time while boosting factory output, deploy machine vision AI solutions for edge applications, drive consolidation of automation, support the integration of renewables, and build a data-driven power grid

Applications: Digital twins, test and measurement devices, PLC virtualization in process and discrete manufacturing, virtual protection relay for energy substations

- A larger shared cache[2] and faster DDR5 memory speeds help deliver incredible performance per watt for critical workloads vs. 4th Gen Intel Xeon Scalable processors.
- Intel RDT helps optimize control loop timing with the ability to allocate cache and memory to high-priority tasks. When used with cache allocation technology (CAT), Intel RDT supports near-real-time deterministic workloads by providing more visibility and control over how LLC and memory bandwidth are used by applications, VMs, and containers.
- Integrated accelerator engines[2] offload critical workloads for encryption, compression, and AI training/inference from the processor to free up CPU cycles for other important workloads.
- Long-life availability[9] and long-use SKUs[10] drive more value from hard-to-reach deployments in the field and on the factory floor.

Processor block diagram

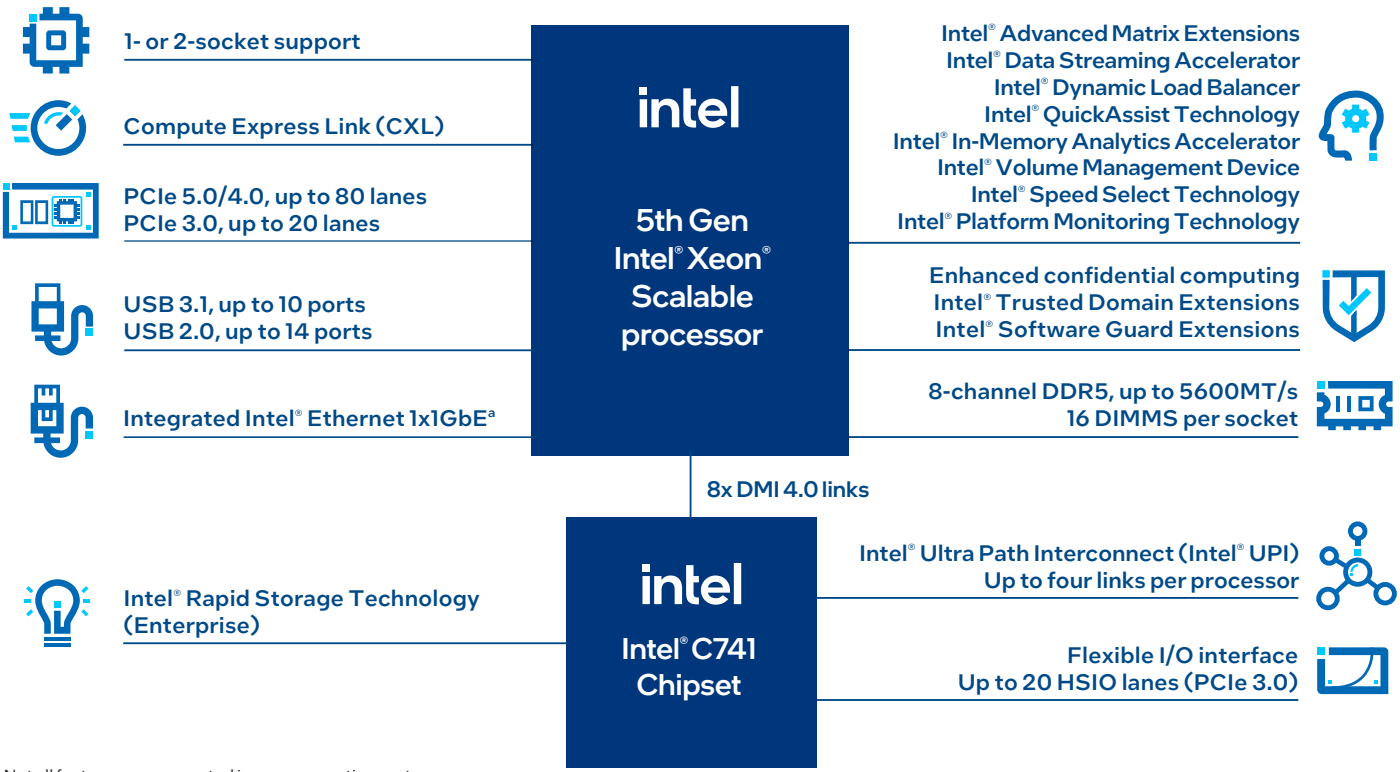

Not all features are supported in every operating system.
Not all features are available on all SKUs.
a. Enabled for workstation and Intel vPro® SKUs only.

Software overview

CATEGORY	OPERATING SYSTEMS/SDKs/BOOT LOADERS	IMPLEMENTATION	DISTRIBUTION AND SUPPORT
Operating systems	Red Hat Enterprise Linux 8.4 LTS, 8.5, and later[a]	Red Hat	Red Hat
	SUSE Enterprise Linux SLE 15 SP3 and later[a]	SUSE	SUSE
	Ubuntu 21.10, 22.04 LTS, and later[a]	Canonical, open source	Canonical
	Alibaba Cloud Linux[a]		
	VMware ESXi[b]	VMware	VMware
	Windows Server 2022 Long-Term Servicing Channel (LTSC) – Full compatibility[c] Windows Server 2019 LTSC/Windows 10 (CSME or SPS FW) – Limited compatibility[c,d]	Intel	Intel, Microsoft
	Azure Stack Hyperconverged Infrastructure (HCI) v22H2 and later – Full compatibility[c] Azure Stack HCI v20H2 – Limited compatibility[c,d]	Intel	Intel, Microsoft
SDK	Intel® oneAPI Video Processing Library (Intel® oneVPL)	Intel	Intel
	Intel® Distribution of OpenVINO™ toolkit	Intel	Intel
	Intel® oneAPI Toolkit	Intel	Intel

Not all features are supported in every operating system. Refer to Intel's IoT Solutions Community for partner contact information.
a. Please contact OS vendor for the exact support and version information.
b. Please contact your VMware rep for compatible ESXi versions.
c. Existing Windows implementation can support HBM Flat assuming it shows up as a NUMA node.
d. Legacy features only. No 4th Gen platform features enabled. Versions will be dropped if the OS vendor discontinues commercial support.

5th Gen Intel® Xeon® Scalable processors for edge

SKU[A,B]	Processor Cores	Base Frequency (GHz)[D]	All-Core Turbo (GHz)	Max Turbo (GHz)	Cache (MB)	TDP (Watts)	Maximum Scalability	DDR5 Memory Speed	Intel® UPI Links Enabled	Default Intel® DSA Devices	Default Intel® IAA Devices	Default Intel® QAT Devices	Default Intel® DLB Devices	Intel® SGX Enclave Capacity (per Processor)	Long-Life Availability[C]	Intel® On-Demand Capable
2S Performance General Purpose																
Intel® Xeon® Gold 6548Y+ processor	32	2.5	3.5	4.1	60	250	2S	5200	3	1	1	1	1	128 GB	Yes	Yes
Intel® Xeon® Gold 6526Y processor	16	2.8	3.5	3.9	37.5	195	2S	5200	3	1	0	0	0	128 GB	Yes	Yes
Intel® Xeon® Gold 5515+ processor	8	3.2	3.6	4.1	22.5	165	2S	4800	3	1	1	1	1	128 GB	Yes	Yes
2S Mainline General Purpose																
Intel® Xeon® Gold 6530 processor	32	2.1	2.7	4.0	160	270	2S	4800	3	1	0	0	0	128 GB	Yes	Yes
Intel® Xeon® Gold 5520+ processor	28	2.2	3.0	4.0	52.5	205	2S	4800	3	1	1	1	1	128 GB	Yes	Yes
Intel® Xeon® Silver 4516Y+ processor	24	2.2	2.9	3.7	45	185	2S	4400	2	1	1	1	1	64 GB	Yes	Yes
Intel® Xeon® Silver 4514Y processor	16	2.0	2.6	3.4	30	150	2S	4400	2	1	0	0	0	64 GB	Yes	Yes
Intel® Xeon® Silver 4510 processor	12	2.4	3.3	4.1	30	150	2S	4400	2	1	0	0	0	64 GB	Yes	Yes
Intel® Xeon® Silver 4509Y processor	8	2.6	3.6	4.1	22.5	125	2S	4400	2	1	0	0	0	64 GB	Yes	Yes
Single-Socket General Purpose																
Intel® Xeon® Bronze 3508U processor	8	2.1	2.2	2.2	22.5	125	1S	4400	0	1	0	0	0	64 GB	Yes	No
Long-Life Use (IoT) General Purpose[D]																
Intel® Xeon® Silver 4510T processor	12	2.0	2.8	3.7	30	115	2S	4400	2	1	0	0	0	64 GB	Yes	Yes

Intel® UPI: Intel® Ultra Path Interconnect
Intel® DSA: Intel® Data Streaming Accelerator
Intel® QAT: Intel® QuickAssist Technology
Intel® DLB: Intel® Dynamic Load Balancer
Intel® IAA: Intel® In-Memory Analytics Accelerator

For more information about Intel® On Demand, visit intel.com/ondemand.

For product specifications, please refer to ark.intel.com.

Intel® processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families.

All processors are lead free (per EU RoHS directive July 2006) and halogen free (residual amounts of halogens are below November 2007 proposed IPC/JEDEC J-STD-709 standards). All processors support Intel® Virtualization Technology (Intel® VT-x, VT-d).

A. Y-SKU: Supports Intel® Speed Select Technology Performance Profile (Intel® SST-PP) 2.0.
B. The frequency of cores and core types varies by workload, power consumption, and other factors. Visit intel.com/content/www/us/en/architecture-and-technology/turbo-boost/intel-turbo-boost-technology.html for more information.
C. Intel does not commit or guarantee product availability or software support by way of road map guidance. Intel reserves the right to change road maps or discontinue products, software, and software support services through standard EOL/PDN processes. Contact your Intel account rep for additional information.
D. High-reliability industrial use supports always-on conditions for up to 10 years.

Get incredible performance, energy efficiency, confidential computing, and support for your important AI workloads at the edge.

Learn more about 5th Gen Intel Xeon Scalable processors at intel.com/5thgenxeon-edge.

Notices and disclaimers

1. See [E2] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.
2. Available on select SKUs.
3. See [G12] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.
4. See [G11] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.
5. See [E1] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.
6. See [E3] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.
7. See [E4] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.
8. See [E5] at intel.com/processorclaims: 5th Gen Intel Xeon Scalable processors. Results may vary.
9. Intel does not commit or guarantee product availability or software support by way of road map guidance. Intel reserves the right to change road maps or discontinue products, software, and software support services through standard EOL/PDN processes. Contact your Intel account rep for additional information.
10. SKUs with long product use life (up to 10 years, up to 100 percent active, no turbo) – Industrial-commercial temperature use condition.
11. Higher-power SKUs are available on the data center road map.
12. Contact OS vendor for exact support and version information.
13. Contact your VMware rep for compatible ESXi versions.
14. Existing Windows implementation can support HBM Flat assuming it shows up as a NUMA node. Currently, it doesn't know how to differentiate HBM vs. DDR5, which means it just treats it as DDR5.
15. Legacy features only; no SPR/EGS features enabled. Versions will be dropped if the OS vendor discontinues commercial support.

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

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 is committed to respecting human rights and avoiding complicity in human rights abuses. See Intel Global Human Rights Principles. Intel® products and software are intended only to be used in applications that do not cause or contribute to a violation of an internationally recognized human right.

0124/LH/CMD/PDF