

The Intel logo is displayed in white lowercase letters with a registered trademark symbol (®) to the upper right of the 'l'.

intel®

4th Gen Intel® Xeon® Scalable processors – why choose between a CPU and accelerator when you can have both?

This whitepaper will cover what built-in accelerators are, why you should use them, and which Intel® Accelerator Engines are right for your business.

The logo for Exertis Enterprise, featuring the word 'exertis' in lowercase and 'ENTERPRISE' in uppercase, separated by a vertical line.

exertis | ENTERPRISE

a **DCC** business

Overview of Xeon processors

New capabilities in 4th Gen Intel® Xeon® Scalable processors

4th Gen Intel Xeon Scalable processors feature a new architecture with higher per-core performance than the previous generation. Boasting up to 60 cores per CPU.

PCI Express Gen5 (PCIe 5.0)

Unlock new I/O speeds with opportunities to enable the highest possible throughput between the CPU and connected devices. 4th Gen Intel® Xeon® Scalable processors have up to 80 lanes of PCIe 5.0—ideal for fast networking, high-bandwidth accelerators, and high-performance storage devices. PCIe 5.0 doubles the I/O bandwidth from PCIe 4.0,9 maintains backward compatibility and provides foundational slots for CXL.

DDR5

Improve compute performance by overcoming data bottlenecks with higher memory bandwidth. DDR5 offers up to 1.5x bandwidth improvement over DDR4,10 enabling opportunities to improve performance, capacity, power efficiency, and cost. 4th Gen Intel® Xeon® Scalable processors offer up to 4,800 MT/s (1 DPC) or 4,400 MT/s (2 DPC) with DDR5.

CXL


Reduce compute latency in the data centre and help lower total cost of ownership (TCO) with CXL 1.1 for next-generation workloads. CXL is an alternate protocol that runs across the standard PCIe physical layer and can support both standard PCIe devices and CXL devices on the same link. CXL provides a critical capability to create a unified, coherent memory space between CPUs and accelerators, and it will revolutionize how data centre server architectures will be built for years to come.



Up to 8-socket scalability
Four Intel UPI ports at 16 GT/s
80 lanes of PCIe 5.0 with CXL



Up to 4-socket scalability
Three Intel UPI ports at 16 GT/s
80 lanes of PCIe 5.0 with CXL



Up to 2-socket scalability
Two Intel UPI ports at 16 GT/s
80 lanes of PCIe 5.0 with CXL

exertis | ENTERPRISE

a **DCC** business

Speak to an expert
whitebox@exertisenterprise.com



What are built-in accelerators?

Instead of building customised workload solutions every time you need new capabilities, why not rely on the technology already built into your CPU?

4th Gen Intel® Xeon® Scalable processors give you the ability to do this! You likely already depend on these trusted CPUs to run all of your workloads, and they are supporting growing data centre needs with Intel® Accelerator Engines, delivering enhanced performance across various emerging workloads.

exertis | ENTERPRISE

a **DCC** business

Speak to an expert
whitebox@exertisenterprise.com



The real-world benefits of Intel® Accelerator Engines

Whether you're using 4th Gen Intel® Xeon® Scalable processors for on-premises workloads, in the cloud or at the edge, Intel® Accelerator Engines can help your business reach new heights. These accelerators are built-in and designed to provide a range of benefits.



Performance

Intel® Accelerator Engines are designed to help deliver higher performance on targeted workloads.



Cost Savings

Intel® Accelerator Engines allow you to improve performance without having to purchase additional specialised hardware.



Power Savings

Intel® Accelerator Engines can help improve power efficiency by offloading common tasks, boosting performance while minimising power usage.

Which Intel® Accelerator Engines are right for your business?

4th Gen Intel® Xeon® Scalable processors have the most built-in accelerators of any CPU on the market, designed to improve performance and help protect data across today's fast-growing workloads. To help you find the best Intel® Accelerator Engines for your business, let's look at the top offerings across three key categories.



Intel AI Engines

Intel® AMX and Intel® AVX-512

Intel® AMX is Intel's next-generation advancement for deep-learning training and inference on 4th Gen Intel® Xeon® Scalable processors. Ideal for workloads like natural language processing, recommendation systems and image recognition, this new feature extends the built-in AI acceleration capabilities of previous gen Intel® Xeon® Scalable processors while also offering significant performance gains.

Intel® AVX-512 is a continuing feature of 4th Gen Intel® Xeon® Scalable processors, a general-purpose performance-enhancing accelerator with a wide range of uses. When it comes to AI, Intel® AVX-512 can accelerate machine learning workloads for training and inferencing. 4th Gen Intel® Xeon® Scalable processors with Intel® AVX-512 are designed to speed up data pre-processing.

exertis | ENTERPRISE

a **DCC** business

Speak to an expert
whitebox@exertisenterprise.com



Intel® Storage Engines

Intel® DSA, Intel® QAT and Intel® VMD

Intel® DSA works on the CPU – between DRAM, caches and processor cores – and extends across I/Os to attached memory and storage, plus networked resources. As Intel's next-generation direct memory access (DMA) engine, it speeds transfers between volatile and persistent memory and supports virtualised memory and I/Os.

Intel® QAT increases performance of storage workloads and applications by accelerating cryptography and data compression/decompression. For instance, using Intel® QAT as an offload engine provides a significant throughput improvement on compression compared to running the same compression algorithm on CPU cores without acceleration. Intel® QAT also functions as a network engine, boosting performance to meet the demands of today's networking workloads to help systems serve more clients.

Intel® VMD is a legacy feature that enables direct control and management of NVMe SSDs from the PCIe bus without the need for additional hardware adapters. It allows for a smoother and lower cost transition to fast NVMe storage while limiting the downtime of critical infrastructure. With benefits like bootable RAID, robust surprise hot-plug and blink status LED, Intel VMD increases serviceability and enables you to deploy next-generation storage with confidence.

exertis | ENTERPRISE

a **DCC** business

Speak to an expert
whitebox@exertisenterprise.com



Intel® Network Engines

Intel® vRAN Boost and Intel® DLB

Intel® vRAN Boost is a new feature that eliminates the need for an external accelerator card by integrating vRAN acceleration directly into the 4th Gen Intel® Xeon® Scalable processor. By reducing vRAN component requirements, it reduces overall vRAN solution complexity and provides power savings for customers.

Intel® DLB is a new feature that enables the efficient distribution of network processing across multiple CPU cores. It also restores the order of networking data packets processed simultaneously on CPU cores. With Intel® DLB, customers can gain higher performance on packet forwarding compared to software queue management on cores without acceleration. Additionally, applications can achieve higher packet processing performance than the previous generation.

exertis | ENTERPRISE

a **DCC** business

Speak to an expert
whitebox@exertisenterprise.com



Intel® Analytics Engines

Intel® In-Memory Analytics Accelerator

Intel IAA is designed to accelerate database and analytics performance. By increasing query throughput and decreasing the memory footprint for in-memory databases and advanced analytics workloads, this new feature can provide faster data movement and improve CPU core utilization by reducing dependency on CPU cores.

Intel IAA is ideal for in-memory databases, open-source databases and data stores like RocksDB, Redis, Cassandra and MySQL. With Intel IAA, customers can gain higher throughput for data decompression when running the open-source RocksDB database engine compared to using software compression on CPU cores without acceleration.

Intel IAA is ideal for in-memory databases, open-source databases and data stores like RocksDB, Redis, Cassandra and MySQL.

exertis | ENTERPRISE

a **DCC** business

Speak to an expert
whitebox@exertisenterprise.com



Why Exertis?

We use our markets and applications knowledge and combine it with portfolio expertise to supply the ideal Intel® product.

We're a certified Intel Premier Support Provider (Intel PSP). We're qualified on Intel® Server Products, with access to Intel's® technical engineering support teams. We also have a dedicated QA department for you questions on the build, to keep your hardware solutions running and minimize impact.

We provide
bespoke support
packages tailored
to your customer
needs.

So, what next?

There are countless workloads for which Intel® Accelerator Engines will deliver greater business value to customers, including targeting more areas than the ones mentioned above!

Speak to an expert today to find out more about Intel's newest offerings

whitebox@exertisenterprise.com

exertis | ENTERPRISE

a **DCC** business

Speak to an expert
whitebox@exertisenterprise.com

