

Powering Next Gen Edge CDN with 4th Gen Intel® Xeon® Scalable Processors

Network-optimized 4th Gen Intel® Xeon® processors with higher core counts, faster I/O and built-in acceleration for load-balancing, copy-and-forward data operations and crypto processing deliver up to 1.46X performance improvements for CDN over previous generations.¹ This latest processor is breaking through existing limits on throughput performance and the number of simultaneous subscriber connections supported by a single server.


High-throughput CDNs that can cost-effectively scale to meet growing performance requirements are a strategic necessity for content providers and broadcasters. And while CDNs move physically closer to the end users, they are also evolving as low-latency infrastructure for the transformation to edge compute. Server platforms that can meet high-throughput, low-latency and cost-efficiency requirements are vital to this evolution.

Content providers continue to evolve their infrastructures to meet the rapidly expanding demand for high-bandwidth streaming content such as HD and UHD video. The efficiency of CDNs is an important focus of that innovation, in terms of metrics such as the density of sessions supported per node and energy requirements per stream. Leaner, more efficient mechanisms are vital to meeting emerging service requirements while maintaining relatively steady CapEx and OpEx. Growth of this market segment is projected to grow at a CAGR of more than 27% to reach a value of nearly \$50 billion by 2026.²

As CDNs move physically closer to the end users they serve, operators are adding low-latency compute infrastructure for the evolution of edge compute beyond core CDN workloads. Server platforms that can meet those low-latency requirements are vital to this evolution.

Network-optimized 4th Gen Intel® Xeon® Scalable processors dramatically improve performance, reduce latency and reduce costs for CDN providers, relative to predecessors. To provide a balanced platform, 4th Gen Intel Xeon Scalable processors incorporate advances across execution, memory and I/O:

- **Enhanced execution resources** benefit encryption and other CDN packet processing with better performance per core than predecessors and an unprecedented array of hardware accelerators built into the processor.
- **Memory bandwidth and speed increase of up to 1.5x**, with up to eight DDR5 memory channels and a transfer rate of up to 4800 MT/s, enabling each node to hold more content in DRAM, to support more streams per node.
- **Faster, higher capacity I/O**, with up to 80 lanes of PCIe 5.0 per socket compared to 64 lanes of PCIe 4.0 in the previous generation, for greater capacity of storage per node and faster networking. These I/O advances help drive exceptional numbers of HD and UHD sessions out to more subscribers and with fewer cache-miss events.



Hardware accelerators built into the processor dramatically improve performance by offloading tasks critical to CDN operation from the processor cores to dedicated hardware resources, which also frees up processor cores for other work. In addition to its throughput benefits, this approach improves latency and energy efficiency by avoiding data transfer to off-processor hardware through the PCIe bus.

Load balancing without occupying processor cores

Dispatching tens of thousands of simultaneous HTTPS connections across many processor cores is typically a software task that occupies dedicated cores. Intel® Dynamic Load Balancer (Intel® DLB) implements load balancing logic and rules in a silicon engine that operates independently of the processor cores. It consumes less energy than software approaches and frees up resources to provide additional connection capacity. That improved efficiency manifests as optimized density of streams per node and performance per watt, for lower equipment requirements and operating costs.

Accelerated movement of streaming data

HTTPS CDN delivery involves large numbers of copy-and-forward data operations. Requests come in from a network interface to the CPU, which fetches content from storage, copies it to DRAM, encrypts it, copies it back to DRAM and passes it to an outgoing NIC port. Each of these operations conventionally requires a software-driven mapping and copy operation between virtual and physical memory address space, which occupies processor core resources and incurs latency that can negatively impact quality of experience for streaming media. Intel® Data Streaming Accelerator (Intel® DSA) handles the memory copy operations using silicon resources that do not occupy processor cores. The Intel DSA silicon engine is much faster than software, which improves latency, and requires less energy, which helps reduce total cost of ownership (TCO).

Reduced performance overhead from encrypt/decrypt operations

Delivering high-bandwidth content over HTTPS requires TLS encryption at scale, with deterministic high throughput to safeguard quality of experience. Intel® QuickAssist Technology (Intel® QAT) built into 4th Gen Intel Xeon Scalable processors accelerates crypto processing, improving new connection rates CDN nodes can support. Offloading execution of TLS handshakes from processor cores frees up precious compute resources for other important edge functions, thus adding to the efficiency of CDN platforms. The latest generation of Intel QAT hardware offload can handle more handshakes per second than predecessor platforms, thus improving scalability of the platform in terms of new connection rates. Intel QAT also improves throughput for content payload encryption, adding further to overall capacity.

View the latest performance data at www.intel.com/PerformanceIndex

Conclusion

4th Gen Intel Xeon Scalable processors offer a significant step forward in the cost-effective delivery of CDN services, with higher densities of simultaneous subscriber connections per CDN node and higher performance per watt. The balanced platform provides core count, memory bandwidth and I/O expandability improvements as the foundation for an unprecedented array of accelerators built into the silicon. The accelerators improve performance and efficiency of load balancing, streaming data movement and encryption, to increase throughput and reduce latency while freeing up compute resources and reducing energy usage.

Learn More

www.intel.com/4thgenxeon-network
networkbuilders.intel.com



¹ See [N60] at intel.com/processorclaims: 4th Gen Intel Xeon Scalable processors. Results may vary.

² Mordor Intelligence. "Content Delivery Network (CDN) Market — Growth, Trends, Covid-19 Impact and Forecasts (2022-2027)." <https://www.mordorintelligence.com/industry-reports/content-delivery-market#:~:text=Market%20Overview,%25%2C%20over%20the%20forecast%20period>.

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 <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 configuration details. 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. 1222/FS/MESH/350505-001US