Solution Brief

Network-Optimized 4th Gen Intel® Xeon® Scalable Processors
5G User Plane Function (UPF)

Improved Density and Power Efficiency for 5G Core Deployments

Network-optimized 4th Gen Intel® Xeon® Scalable processors enable high performance for 5G core deployments while reducing total cost of ownership (TCO). Built-in accelerators help increase throughput and decrease latency, while advances in power management enhance both the responsiveness and the efficiency of the platform.

Dramatic growth in global mobile traffic has continued from the beginning of the COVID-19 pandemic, with more than a 5.6x increase projected from 47.6 million TB/month in 2020 to 267.9 million TB/month in 2027.[1] Communication service providers (CoSPs) must have the operational agility to meet those needs cost-effectively while maintaining the high levels of service quality needed to attract and retain customers.

True 5G core implementations are key for CoSPs to deliver innovative and monetizable 5G services and handle increases in traffic. Workload consolidation and cost reduction from replacing proprietary equipment with general-purpose, standards-based servers has been a key contributor to new efficiencies, as has the reduction in data backhaul achieved by distributing 5G workloads.

Network-optimized 4th Gen Intel® Xeon® processors are the next step in that evolution, accelerating 5G core workloads such as 5G UPF while increasing energy efficiency, with a high-throughput, low-latency platform engineered for the core and scalable to the edge. The processor provides a balanced array of advances over predecessors to drive performance and efficiency in UPF deployments:

- **Enhanced execution pipeline** based on improved microarchitecture for high per-core performance as well as up to 52 cores per socket in high core-count SKUs and an industry-leading set of built-in hardware accelerators.
- **Increased memory bandwidth and speed** up to 1.5x, with up to eight DDR5 channels operating at up to 4800 MT/s, enabling the platform to hold more user plane data close to the processor for enhanced throughput.
- **Expanded I/O speed and capacity** provided by up to 80 lanes of PCIe Gen 5.0 per processor compared to 64 lanes of PCIe in the predecessor, for faster movement of user plane data.

The platform further enhances the scalability and performance of software-defined infrastructure with an enhanced instruction set architecture (ISA). Intel Ethernet 800 Network Adapters complement the Intel Xeon processor for UPF deployments, with protocol-specific (e.g., GTP-U) parsing, classification and steering by means of Intel Dynamic Device Personalization (Intel DDP). Advanced architecture and the industry's largest lineup of hardware accelerators built into network-optimized 4th Gen Intel Xeon Scalable processors provide performance and efficiency across UPF workloads, as introduced in the sections below.

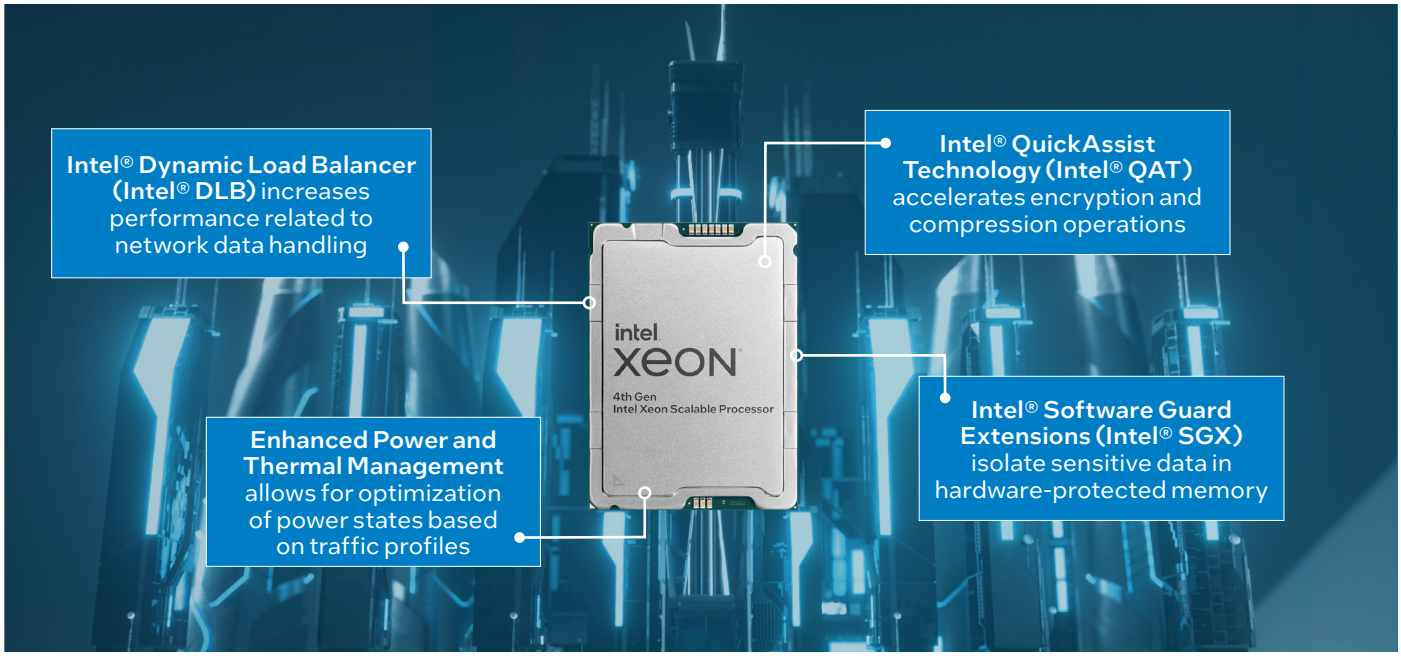

Increased throughput with silicon-based load balancing

Dynamically tuning the placement of shifting workloads across the available processing cores to be as efficient as possible is central to making best use of system resources. Intel® Dynamic Load Balancer (Intel® DLB) implements that orchestration with silicon-based load balancing rules and logic built into network-optimized 4th Gen Intel Xeon Scalable processors. This built-in feature operates without occupying the processor cores that would be needed to execute software-based traffic control.

Performing those functions in hardware is inherently faster than using software logic, which contributes to better hardware utilization, for lower operating expense (OpEx) and increased throughput. Likewise, Intel DLB performs the load balancing function using less energy than software-based methods, which provides both financial and sustainability benefits to network operators.

Power-management improvements for efficiency and responsiveness

The need to reduce and optimize energy consumption in widely distributed UPF deployments is vital to cost efficiency and a point of climate social responsibility. Network-optimized 4th Gen Intel Xeon Scalable processors represent an ideal balance for service providers between power and performance featuring high per-core throughput, performance per watt and core counts that increase the number of connections per host or instance. The platform also introduces new CO.1 and CO.2 light sleep power states, which makes power management more responsive by taking idle cores out of sleep much faster than with predecessor technologies.

The ability to shift more quickly between power states reduces exit latency, defined as the delay in the execution pipeline as processor cores come out of sleep to do work. Lower exit latency helps ensure that processing resources are available when they are needed. In the context of real-time communications, that factor is critical to avoid dropped packets that would adversely affect service quality and user experience. Carriers can therefore save on energy usage and OpEx with more effective power-management policy, while meeting service level agreements (SLAs) and maintaining high quality of experience (QoE).

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

Accelerated encryption with enhanced power efficiency

5G core implementations must protect the privacy and integrity of both control plane and data plane transmissions with extensive encryption and decryption operations, which can consume significant processor resources. 4th Gen Intel Xeon Scalable processors offer hardware acceleration for encryption using built-in Intel® QuickAssist Technology (Intel® QAT), which contributes to accelerating UPF workloads, including back-office services such as OSS and BSS.

Increased crypto performance also enhances and reduces overhead associated with application and device authentication. Because Intel QAT is built into the processor, gaining the encryption and decryption throughput increase does not require sending data over the PCIe bus to communicate with external hardware. Eliminating that requirement reduces power consumption as well as transfer latency, for better overall responsiveness and energy efficiency.

Confidential computing protects private data in public environments

In highly distributed environments, 5G core workloads inter-operate with third-party applications, data, services and infrastructure, creating a larger, low-visibility attack surface. Encryption can protect the privacy of data at rest and in transit, but that data must generally be unencrypted while in use.

Confidential computing protects data while in use by isolating it in hardware-enforced private memory spaces known as "enclaves." Sensitive data is decrypted only within enclaves, where it can only be accessed by trusted code; no user or process that is not explicitly granted that trust can reach the data, regardless of privilege level. Intel® Software Guard Extensions (Intel® SGX) is the most researched, updated and deployed confidential computing technology on the market today, with the smallest trust boundary.

Comprehensive Telemetry

Service assurance, especially in a cloud native infrastructure framework, is another concern for network operators. Existing systems pick up on real-time issues in the network and service layers, but not those within the infrastructure. This visibility gap between the underlying network infrastructure and the service layers often prevents visibility into the root causes of service degradations and other IT Ops incidents.

The combination of 4th Gen Intel Xeon Scalable processors, with built-in accelerators, and Intel® Platform Telemetry Insights offers a granular view into cloud-native infrastructure operations. This visibility into the heart of the foundational platform infrastructure enables assessment of operational metrics, including health, utilization, congestion, power consumption and configuration checks. By adopting CNCF's OpenTelemetry project, Intel provides a standard implementation for the industry to benefit from infrastructure observability.

Conclusion

Network-optimized 4th Gen Intel Xeon Scalable processors help CoSPs handle fast-expanding UPF data requirements, with energy-efficient operation and high density of connections per host or instance. The processor is designed for high throughput and low latency from the core to the edge. Built-in hardware accelerators enhance distribution of work across cores, accelerate encryption and protect data with confidential computing. The energy efficiency of handling these operations in hardware is complemented by enhanced power states and power management. With the ability to handle higher data volumes within strict latency requirements while improving on cost, scalability and sustainability metrics, this platform provides a future-ready hardware foundation for next-generation services.

Learn More

www.intel.com/4thgenxeon-network

networkbuilders.intel.com

intel®

[1] ReportLinker, October 2022. "Global Mobile Data Traffic Industry." https://www.reportlinker.com/p05442636/Global-Mobile-Data-Traffic-Industry.html.

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.

1222/FS/MESH/350506-001US