

# CPU Core Type and Network Workload Affinity

## Maximize Network Efficiency and Scalability: Power Diverse Modern Workloads with Enhanced Performance Using Intel® Xeon® 6 Processors

### Authors

**Brian Keating**

Principal Engineer, Intel

**Georgii Tkachuk**

Performance Engineer, Intel

**Mrittika Ganguli**

Principal Engineer, Intel

**Abhirupa Layek**

Network Hardware Engineer, Intel

**Arpit Menaria**

Enterprise Edge Segment Manager, Intel

**Edwin Verplanke**

Senior Principal Engineer, Intel

**Subhiksha Ravisundar**

Platform Architect, Intel

**Max Coates**

ISB Manager, Intel

### Core Selection Strategies for Intel® Xeon® 6 Processors in Network Workloads

Intel Xeon 6 processors offer a choice between two distinct core types: Efficient cores (E-cores) and Performance cores (P-cores), E-cores provide high core density and exceptional performance per watt, making them ideal for power-sensitive environments. P-cores are designed for a broad range of workloads, excelling in AI and compute-intensive tasks with higher overall performance. Selecting the right core type for a specific application requires a clear understanding of the workload and how it interacts with each core type.

This paper provides a comprehensive analysis of networking workloads, especially those utilizing user mode data planes, and examines their performance on processors with different core configurations. We classify network applications into two main categories: those that benefit from scalability and those that demand high single-threaded performance. By exploring the architectural differences between E-cores and P-cores, we highlight how these variations can impact application performance.

To accurately assess the capabilities of each core, we employ a robust benchmarking methodology designed for controlled and repeatable performance evaluations. The study introduces Key Performance Indicators (KPIs) such as performance per watt and overall performance, adjusted for variations in frequency and core count. Our findings are presented through a series of benchmarks, offering insights into the optimal use of CPU cores for network applications.

This whitepaper aims to guide system architects and network engineers in making informed decisions when deploying user-mode data planes on modern CPUs, optimizing performance by balancing power consumption and computational throughput.

### Table of Contents

- Core Selection Strategies .....1
- E-cores and P-cores ..... 2
- Analysis Methodology..... 2
- Key Hardware Features.....3
- Workload Analysis ..... 3
  - NGFW Benchmark.....4
  - VPP-FIB Benchmark..... 5
  - VPP-IPSEC Benchmark..... 5
  - NGINX Benchmark ..... 6
  - L4 Load Balancer - HDSL.B..... 7
  - Istio-Envoy Benchmark ..... 8
- Conclusions .....9

### CPU Optimization for DPDK and VPP Processing Applications

User mode packet processing applications, built with the Data Plane Development Kit (DPDK), require a highly optimized runtime environment to meet their stringent performance criteria. DPDK, a set of libraries and drivers for fast packet processing, demands low-latency interaction with network interfaces, efficient use of CPU cycles, and the ability to scale across multiple processor cores. The CPU architecture should therefore provide rapid access to I/O, many execution units (cores), and a high-frequency clock rate to maximize packet throughput and minimize processing time.

Similarly, Vector Packet Processing (VPP) applications enhance data plane performance by processing packets in batches. These workloads need CPU architectures with high parallelism and fast context switching to manage vectorized packet processing efficiently. This requires a robust cache hierarchy to reduce memory access latency and a CPU design capable of handling increased vector processing workloads.




Figure 1. Intel® Xeon® 6 processors

### Efficient cores and Performance cores

The choice between Efficient cores (E-cores) and Performance cores (P-cores) allows better alignment with workload needs. Data plane-intensive workloads with high throughput and low latency may perform best on CPUs with fewer, higher-frequency P-cores due to limited scalability across threads. In contrast, application layer workloads involving deep inspection of reassembled traffic streams often scale well across many cores, benefiting from CPUs with numerous lower-frequency E-cores-

#### Performance Cores (P-cores):

- Tuned for high turbo frequencies and high IPC (instructions per cycle).
- Ideal for handling heavy single-threaded workloads, such as those in packet processing applications.
- Capable of hyper-threading, allowing them to run two software threads simultaneously.
- Superior performance/core.
- Wider/deeper microarchitecture, including bigger caches in the un-core, higher memory BW between cores and wider vectors.
- Optimized for AI/ML/Analytics workloads, data acceleration, and vectorized code, benefiting from an extended instruction set, including Intel® AMX.

#### Efficient Cores (E-cores):

- Physically smaller, allowing multiple E-cores to fit into the space of one P-core in a CPU package.
- Designed to maximize CPU and power efficiency, focusing on performance per watt.
- Have a higher core density, enabling more cores within the same power envelope.
- Optimized for performance per unit area.
- Provide superior integer performance per watt and per socket..

### Analysis Methodology

This whitepaper outlines a comprehensive methodology for evaluating CPU performance, focusing on Key Performance Indicators (KPIs) such as performance per watt for energy efficiency and absolute performance in terms of core frequency, core count, cache size, and I/O capabilities. By examining these metrics, we offer insights into optimizing CPU configurations for user-mode packet processing, assisting system architects and network engineers in making informed deployment decisions.

Workload	Data Plane Implementation
Next Generation Firewall (NGFW)	VPP + Snort
IP Forwarding	VPP
Ipssec	VPP
TLS Web Server	Nginx (Snort)
L4 Load Balancer	HDSLB (VPP)
Service Mesh	Istio + Envoy (Linux Kernel)

Table 1. Tested Workloads

We identified a range of networking workloads, as detailed in Table 1, and established benchmarks along with key performance indicators (KPIs) to assess each scenario:

- **Absolute Performance:** Includes metrics like Throughput (Gbps), Connection Rate (cps), and Request Rate (rps).
- **Performance per Watt:** Calculated by dividing the absolute performance by the CPU’s power consumption during the workload.

For benchmarking, we selected processors tailored to each workload type:

- For workloads that scale across a full CPU socket, we compared P-core and E-core processors with similar Thermal Design Power (TDP) to determine the best fit within the thermal envelope.
- For single-threaded performance, we compared performance per core or thread. It is important to note that while using SMT/HT on P-cores can boost core performance by 1.2x, it may reduce per-thread performance by half.

## Key Hardware Features

### Intel® Ethernet Network Adapter

Intel Ethernet Network Adapter E810-2CQDA2 delivers up to 200 Gbps of total bandwidth in systems that are PCIe 4.0 compliant. Each QSFP28 port supports up to 100 Gbps, providing the functionality and throughput of two 100 Gbps adapters in a single bifurcated PCIe 4.0 x16 slot. It is designed for optimizing networking workloads including network functions virtualization (NFV) and features technologies:

- Intelligent Flow Direction: Receiver Side Scaling (RSS)
- Comprehensive Network Virtualization Overlay Protocols Support
- vSwitch Assist
- QoS: Priority-based Flow Control (802.1Qbb)
- Enhanced Transmission Selection (802.1Qaz)
- Differentiated Services Code Point (DSCP)
- Dynamic Device Personalization (DDP)

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

Intel AVX-512 is a powerful SIMD instruction set. It debuted in server products with the 1st Gen Intel® Xeon® Scalable processor, in which the microarchitecture expanded load, store, and execution port widths to accommodate 512-bit wide instructions. Two execution ports are available to retire Intel AVX-512 instructions, enabling up to two Intel AVX-512 instructions to retire, concurrently yielding a total of 1024 bits of data. It's important to note that Intel AVX-512 is only available on the Performance-cores (P-Cores) of Intel processors. This distinction ensures that Intel AVX-512's intensive workloads leverage the high-performance capabilities of the P-Cores, rather than the more efficiency-oriented E-Cores, providing optimal performance in compute-intensive applications.

### Intel® QuickAssist Technology (Intel® QAT)

Intel QAT is a hardware accelerator that enhances system performance by offloading compute-intensive operations such as encryption, decryption, compression, and decompression from the CPU. This offloading frees up CPU resources, allowing them to perform other tasks more efficiently, thereby improving overall system performance, power efficiency, and the ability to handle larger workloads.

Intel QAT supports both symmetric and asymmetric cryptographic operations, including popular algorithms like AES, RSA, and Elliptic Curve Cryptography, as well as advanced data compression techniques.

Integrated into Intel Xeon Scalable processors and other Intel platforms, Intel QAT significantly accelerates tasks such as secure web transactions, VPNs, and data center operations, making it ideal for use in high-performance computing, cloud services, enterprise data centers, and network edge applications. By accelerating these workloads, Intel QAT helps reduce total cost of ownership (TCO) and improves performance per watt, which is critical for data-intensive and network-dependent environments.

Processor SKU	Core Type	Core Count	Base Frequency	TDP (W)
Intel® Xeon® Platinum 8592+ processor	P-Core	64 Cores; 128 Threads	1.9	350
Intel® Xeon® 6780E processor	E-Core	144 Cores; 144 Threads	2.2	330

Table 2. Processors Benchmarked

## Vector Packet Processing (VPP)

VPP is a highly optimized software, it processes packets in large bursts (vectors) and takes advantage of Single Instruction/Multiple Data (SIMD) instructions in Intel architecture. VPP includes specific implementations for each available SIMD bitwidth option in Intel architecture such as Intel® SSE, Intel® AVX2, and Intel AVX-512 for many of its functions. Generally, the widest bitwidth variants offer the best performance.

Each function block treats a packet vector, up to a maximum of 256 packets, as an input and processes them in an identical manner. This helps maximize highly efficient utilization of CPU instruction cache (I-cache). In addition, VPP innovatively adopts a Packet Processing Graph as part of its core design, in which case each function block is abstracted as a graph node. The graph nodes are organized as a tree shape graph by registering the “next” output nodes either initially or at runtime. The packet vectors flow from the Ethernet Adapter Receive (RX) nodes all the way to Transmit (TX) nodes (or dropped) based on the processed destinations in each graph node within.

The DPDK plugin for VPP allows VPP to take advantage of all optimizations and hardware enablement features like NIC and hardware accelerator drivers included in DPDK. This includes the poll mode driver for the Intel Ethernet Network Adapter E810 optimized with Intel AVX-512 intrinsics for vectorizing packet receive and transmit operations.

## Workload Analysis

This section provides a comprehensive analysis of VPP-based and Linux kernel networking workloads, as detailed below in Table 2. To assess both performance and power consumption, we utilized Intel EMON, a specialized command-line (CLI) tool for profiling application and system performance. Intel EMON uses hardware Performance Monitoring Units (PMUs) to gather performance monitoring events, allowing us to capture detailed metrics on CPU activity during workload execution.

By leveraging PMUs, EMON collects critical data such as socket power usage and workload performance, enabling an in-depth examination of how different CPU configurations affect efficiency and throughput. This approach ensures that the performance analysis closely mirrors real-world conditions, offering valuable insights into the interaction between workloads and hardware. These insights help guide the optimization of CPU selection and deployment strategies, ensuring that networking tasks are handled with maximum efficiency and performance.

## NGFW (Next Generation Firewall)

NGFW is a converged network security workload that combines multiple security functions traditionally performed by separate appliances or network security services, including:

- Basic stateful L4 firewall capabilities
- Deep Packet Inspection (DPI), Intrusion Prevention, and other L7 content inspection features
- VPN security gateway (e.g., IPsec)
- TLS Intercepting Proxy for decryption, inspection, and re-encryption of end-to-end TLS encrypted content

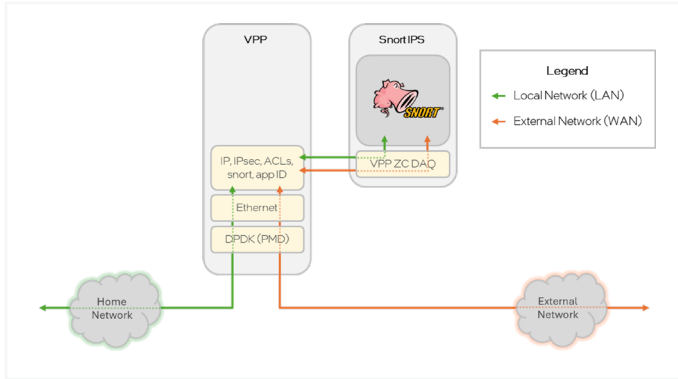


Figure 2. NGFW Sample Architecture - Logical View

Intel developed a reference NGFW application using open-source components for benchmarking, implementing a subset of NGFW functionality: a stateful firewall using VPP and an Intrusion Prevention System (IPS) using Snort. Benchmarks were run using the HTTP 64KB traffic profile with one GET request per connection, Snort lightspd rules, and security policies. Note that the traffic was unencrypted; results may vary with IPsec and TLS enabled.

Benchmarks were conducted on processors with similar Thermal Design Power (TDP), though the TDP of the E-core CPU (Intel® Xeon® 6780E processor) is 20W lower than the P-core equivalent (Intel® Xeon® Platinum 8592+ processor). As shown below in Table 3, the Intel Xeon 6780E processor (E-core) achieves up to 1.38x higher performance and 1.8x better performance per watt compared to the Intel Xeon Platinum 8592+ processor (P-core). The active power consumption during benchmarks was lower than the TDP for both processors, indicating efficient power use.

Performance is driven by Snort, which uses 58 of 64 cores on the Intel Xeon Platinum 5982+ processor and 128 of 144 cores on the Intel Xeon 6780E processor, with the remaining cores for VPP and overhead tasks. As shown in Figure 3, performance scales almost linearly with the number of Snort cores, demonstrating consistent throughput per core. This scaling highlights the architecture's efficiency in distributing workloads across multiple cores, making it beneficial for parallel processing tasks.

Processor SKU	Active Power (W)	Performance per Watt (Mbps/W)
Intel® Xeon® Platinum 8592+ processor	131	416
Intel® Xeon® 6780E processor	168	691

Table 3. Processors Power Consumption Metrics

While a single hyperthreaded P-core delivers up to 1.68x the performance of an E-core, the Intel Xeon 6780E processor's

higher core count (128 vs. 58) compensates for the lower per-core performance, resulting in about 1.28x greater overall throughput. The larger number of E-cores supports extensive parallel processing, making the Intel Xeon 6780E processor highly effective for workloads that scale well with increased core counts, such as network security functions and other multi-threaded applications. This balance between core performance and count offers a powerful tool for optimizing performance and power efficiency in demanding environments.

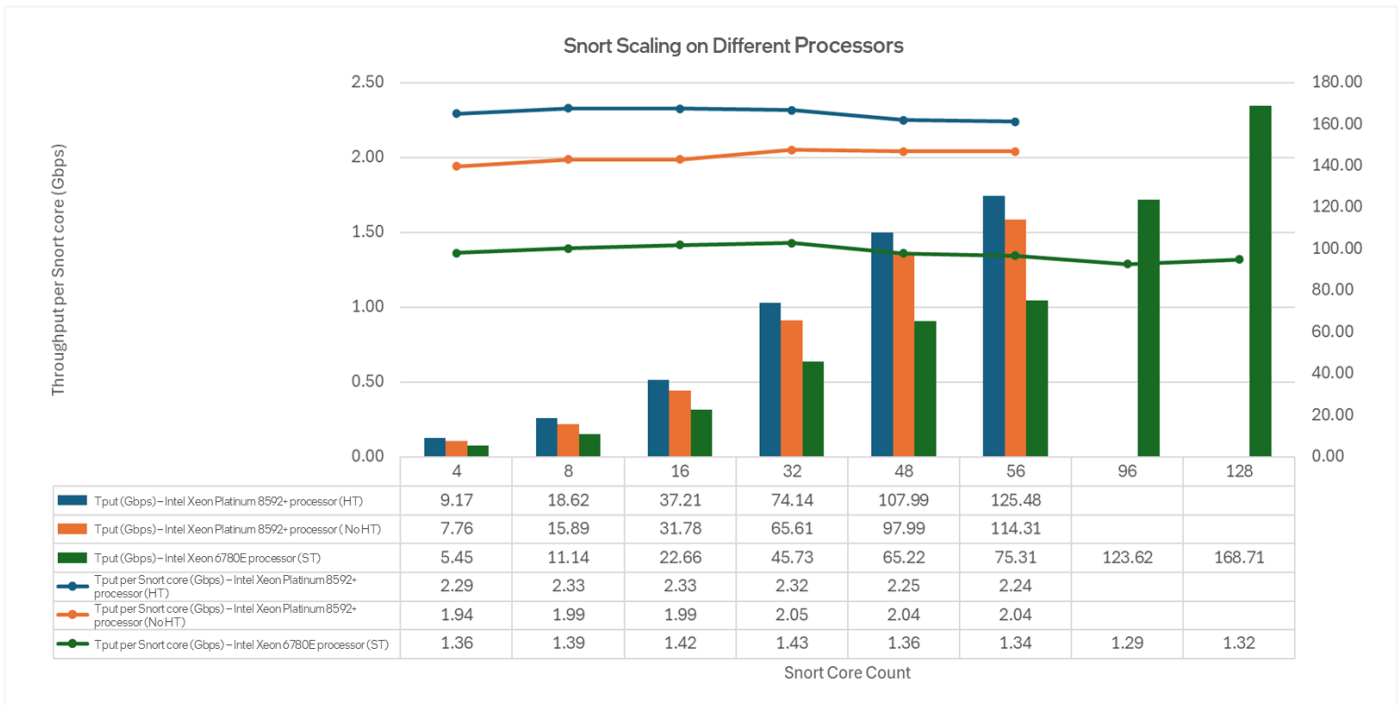


Chart 1. Snort Scaling on Different Processors

## VPP-FIB

The Vector Packer Processor’s (VPP) Forward Information Base (FIB) configuration is one of the most straightforward applications within VPP. It routes network packets based on a user-defined routing table, or FIB.

By default, the packet processing within VPP FIB is performed in a run-to-completion model, meaning that each packet is processed entirely by on VPP thread without needing to communicate with other threads. This absence of interthread communication and locking significantly contributes to the application’s ability to scale with number of CPU cores.

In configurations that utilize many CPU cores and a large FIB, the application is required to move a substantial amount of data. This scenario makes the application increasingly sensitive to the CPU cache size, particularly concerning the L2 and L3 caches. The fastest route lookup times are achieved when the application’s packet buffers, and FIB entries are available in the caches. Additionally, the aggregate data generated by the NIC device, and the CPU puts significant pressure on the CPU’s coherent fabric which may require that it maintains high-frequency to accommodate the high data transfer rates.




Chart 2. VPP - FIB Absolute Performance




Chart 3. VPP - FIB Power Efficiency Scaling

## VPP-IPSEC

VPP IPsec is a core component of Vector Packet Processing (VPP) that supports secure, reliable, and high-performance networking. It offers command-line (CLI) and virtual API (vAPI) commands for configuring key security elements, such as the Security Policy Database (SPD), Security Associations (SA), and cryptographic algorithms. These tools streamline the deployment and management of IPsec protocols, facilitating secure data transmission across networks. VPP IPsec is essential for applications like VPNs, ensuring data integrity and confidentiality with high throughput and low latency, making it ideal for modern secure networking.

VPP IPsec supports:

- Major cipher, authentication, and AEAD cryptographic algorithms
- ESP tunnel and transport mode, optional over UDP or GRE
- Authentication header
- IKEv2 initiator and responder

The most resource-intensive procedure within IPsec is cryptographic operation. To ensure both performance and flexibility, VPP IPsec leverages the underlying crypto infrastructure, optimizing cryptographic operations for maximum efficiency.




Chart 4. VPP - FIB IPsec Performance

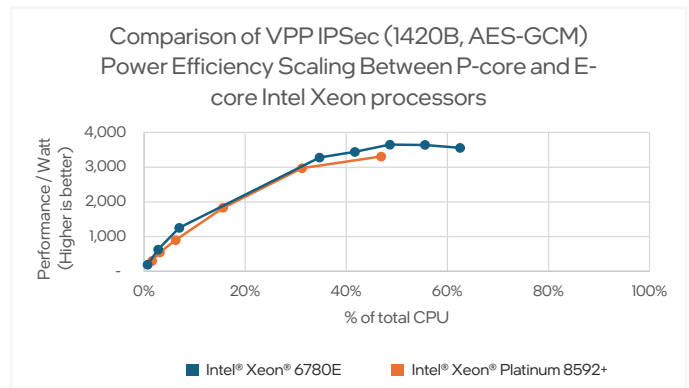


Chart 5. VPP - FIB IPsec Power Efficiency Scaling

## NGINX\*

NGINX is a versatile, high-performance web server that has evolved beyond its original role as a web server to also function as a reverse proxy, load balancer, and HTTP cache. NGINX employs an asynchronous, event-driven architecture that allows it to efficiently manage a large number of simultaneous connections with minimal memory usage. This makes it particularly well-suited for high-traffic websites and applications that require fast and reliable performance.

NGINX’s flexibility comes from its modular design, allowing it to handle various functions like serving static and dynamic content, SSL/TLS termination, and HTTP/2 and HTTP/3 support. It also integrates seamlessly with FastCGI, uWSGI, and SCGI for dynamic content, and provides robust load balancing with in-band health checks. NGINX’s reverse proxy capabilities include caching, compression, and WebSocket support, making it ideal for modern web infrastructures that demand high concurrency and low resource consumption.

Here we focus on NGINX’s secure web server configuration, evaluating the performance of Intel® Xeon® Processors when handling tasks such as downloading large encrypted files over HTTPS and establishing secure TLS connections using ECDSA or RSA 2048 encryption ciphers.

### NGINX Webserver HTTPS Throughput

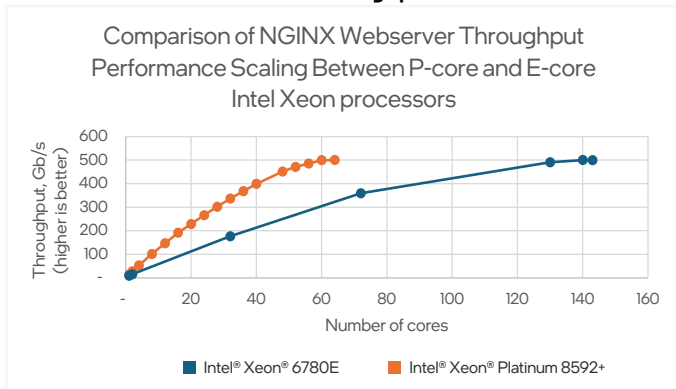


Chart 6. Nginx Webserver Throughput Performance Scaling




Chart 7. Nginx Power Efficiency Scaling

### NGINX Webserver TLS connections per second with ECDSA

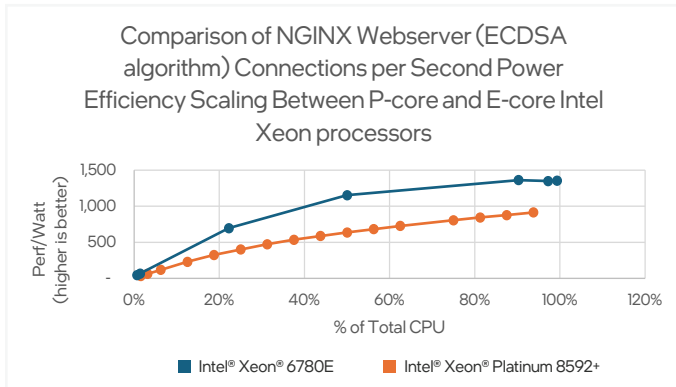


Chart 8. Nginx Webserver Throughput Performance Scaling




Chart 9. Nginx Power Efficiency Scaling

### NGINX Webserver TLS connections per second with RSA2K

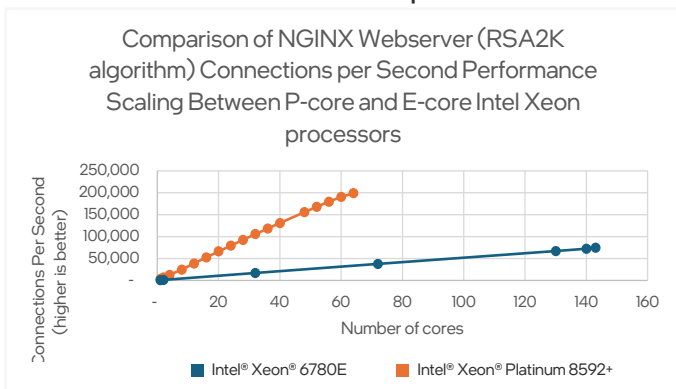


Chart 10. Nginx Webserver Throughput Performance Scaling




Chart 11. Nginx Power Efficiency Scaling

## HDSL B

High Density Scalable Load Balancer (HDSL B) is a Layer 4 load balancer built on the Vector Packet Processing (VPP) framework. It is engineered to achieve industry-leading performance, targeting up to 150 million packets per second (Mpps) throughput, 1 billion concurrent connections, 10 million connections per second (CPS) per node, and exceptional linear scalability.

HDSL B-VPP customizes the lightweight infrastructure from Fd.io/VPP and incorporates advanced load balancing features leveraging Intel Architecture (IA) technologies. Key optimizations include support for Intel AVX-512 instructions, Fast Data I/O Routing (FDIR), and Dynamic Load Balancing (DLB), ensuring optimal performance on Intel platforms. Delivered as a customer reference, HDSL B exemplifies best practices for deploying load balancing solutions on Intel platforms.

For performance evaluations, HDSL B operates in DNAT mode using a weighted round-robin scheduling algorithm with UDP traffic. The Device Under Test (DUT), equipped with either 5th Gen Intel Xeon Scalable processors or Intel Xeon 6 processors, connects to an Ixia traffic generator via one receive (Rx) and one transmit (Tx) port. Traffic sent to the Virtual IP (VIP) of HDSL B undergoes DNAT and is then forwarded and scheduled between two real servers.

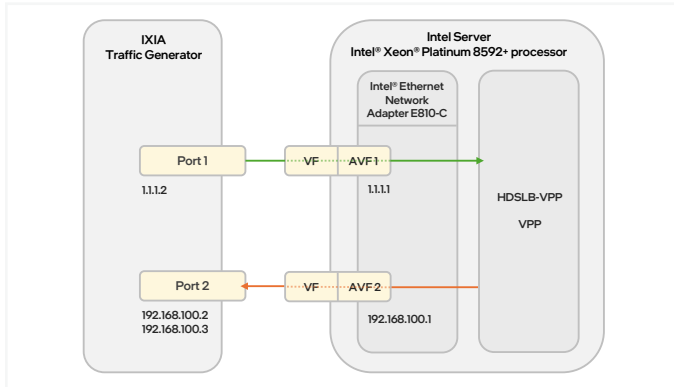


Figure 3. L4 Load Balancer HDSL B Test Setup




Chart 12. HDSL B 8c Performance and Perf-Per-Watt Scaling

Chart 12 above compares the performance of HDSL B on 5th Gen Intel Xeon Scalable processors and Intel Xeon 6 processors across NAT, NAT-SW, FNAT, and Direct Server Return modes. Scaling data indicates that both processors scale effectively with increasing cores and queues, achieving up to 100 Mpps on Intel Xeon 6 and 115 Mpps on 5th Gen Intel Xeon Scalable processors in a 16-core/16-queue configuration. Additionally, the 5th Gen Intel Xeon Scalable processors demonstrate superior performance per watt compared to the Intel Xeon 6 processors, making them an optimal choice for power-sensitive applications.

## Istio - Envoy

Istio is a widely adopted industry-standard microservices framework that leverages worker threads for scalability, particularly emphasizing secure transactions that involve frequent context switches and memory operations. Within Istio, Envoy functions as the network data plane and is utilized as a cloud-native sidecar, ingress gateway proxy, as well as in reverse and forward proxy modes for cloud and Secure Access Service Edge (SASE) environments.

The Istio-Envoy benchmark evaluates the peak performance achievable on single-socket platforms using various combinations of Efficient (E) cores and Performance (P) cores. This benchmark encompasses Intel Xeon 6 processors, 4th Gen and 5th Gen Intel Xeon Scalable processors within a service mesh environment consisting of 23 microservices. Since the workload does not fully utilize the Thermal Design Power (TDP), differences in TDP across these platforms do not impact the performance outcomes.

In this setup, a Nighthawk client generates Layer 7 network traffic, sending it to an Envoy acting as the Ingress proxy. Multiple replicas of server pods, each containing a paired Nighthawk server and a sidecar container, handle the incoming requests.

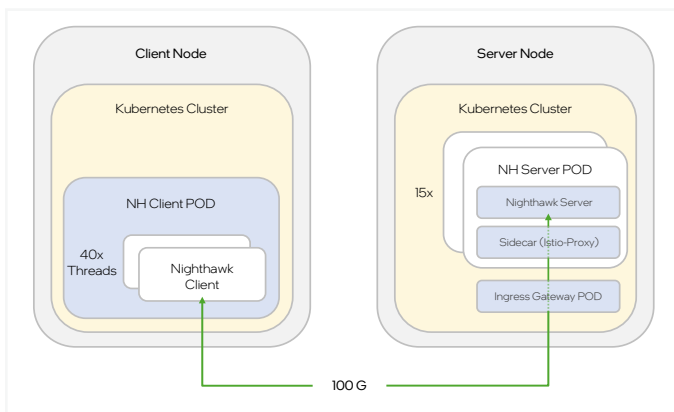


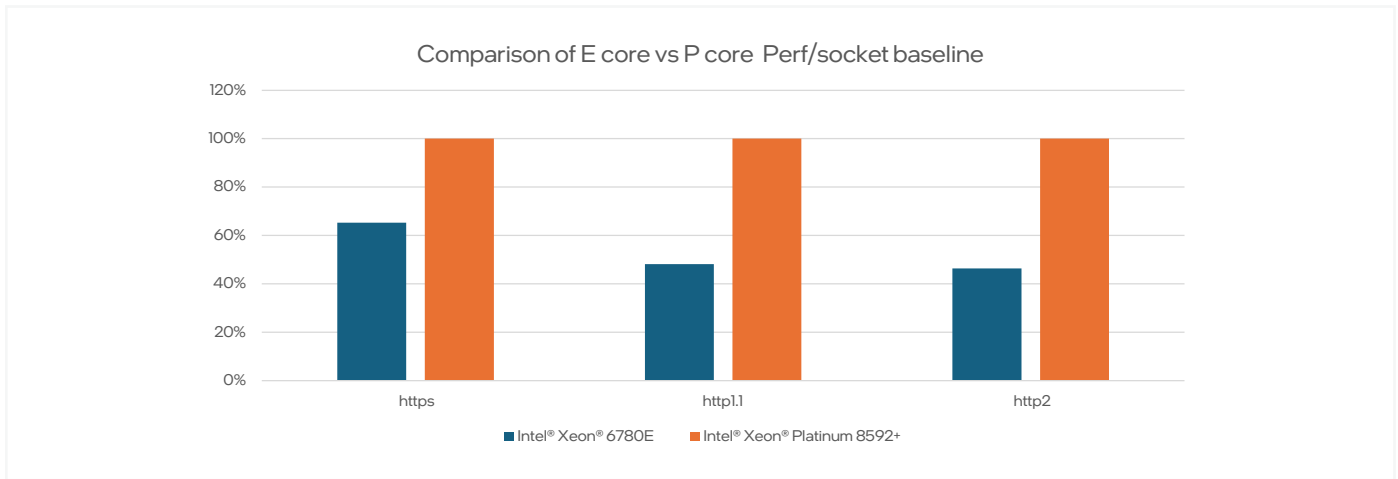
Figure 4. Test Setup for Istio-Envoy 2 Node Cluster

Memory Details	4th Gen Intel® Xeon® Scalable	Intel® Xeon® 6
Installed Memory	256 GB (16x16GB DDR5 5600 MT/s)	256 GB (8x32 GB DDR5 5600 MT/s)
L1d Cache/Core	32 KB	48 KiB
L1i Cache/Core	64 KB	32 KiB
L2 Cache/Core	1 MiB	2 MiB
L3 Cache/Core	96 MiB	320 MiB

Table 4. System Memory Configuration

**White Paper | CPU Core Type and Network Workload Affinity**

The Intel Xeon 6 processor has a higher core count overall—144 cores compared to 112 threads on the 5th Gen Intel Xeon Scalable processor. The RPS scales well on both the 5th Gen Intel Xeon Scalable processor and Intel Xeon 6 platforms, with lower socket power consumption (package power in Watts) across the entire service mesh. E cores in the Intel Xeon 6 platform excel in power efficiency, especially for HTTPS workloads.



**Chart 13. Istio - Envoy Power Efficiency Scaling**

**Glossary**

Abbreviation	Description
Intel AVX-512	Intel Advanced Vector Extensions 512
CPU	Central Processing Unit
DPDK	Data Plane Development Kit
DPI	Deep Packet Inspection
DR	Direct Server Return
E-core	Efficient Core
ECDSA	Elliptic Curve Digital Signature Algorithm
FNAT	Full NAT
FIB	Forwarding Information Base
Gbps	Gigabits per second
HTTPS	Hypertext Transfer Protocol Secure
I-Cache	Instruction Cache
IKEv2	Internet Key Exchange Version 2
I/O	Input/Output
IPC	Instructions Per Cycle
Ipsec	Internet Protocol Security
KPI	Key Performance Indicator
L4	Layer 4
L7	Layer 7
Mpps	Million packets per second
MT/s	Megatransfers per second
NGFW	Next generation Firewall
NIC	Network Interface Card
PCIe	Peripheral Component Interconnect Express

Abbreviation	Description
P-core	Performance Core
PMU	Performance Monitoring Unit
QoS	Quality of Service
Intel QAT	Intel QuickAssist Technology
RAM	Random Access Memory
RPS	Requests Per Second
RSA	Rivest Shamir Adleman
SASE	Secure Access Service Edge
SIMD	Single Instruction/Multiple Data
SMT/HT	Simultaneous Multithreading/Hyper-Threading
SPD	Security Policy Database
SRF	Scalable Radio Frequency
TDP	Thermal Design Power
TLS	Transport Layer Security
UDP	User Datagram Protocol
VPP	Vector Packet Processor
vAPI	Virtual Application Programming Interface
VIP	Virtual IP Address





## Conclusion

Intel's E-core and P-core architectures each offer distinct performance benefits that cater to different networking workloads. E-core platforms, with their focus on energy efficiency excel in performance per watt making them ideal for VPP-based applications and tasks that scale effectively across multiple cores. In contrast, P-core platforms, such as those in the Intel Xeon 6 processors provide higher overall throughput, making them well-suited for compute-intensive workloads that demand strong single-threaded performance.

Applications like NGINX and Envoy take advantage of the larger cache capacities found in Intel Xeon 6 processors, enhancing performance for smaller message sizes. The inclusion of Intel AVX-512 instructions in P-core platforms significantly boosts the performance of cryptographic operations like RSA-2048. For E-core platforms, offloading cryptographic tasks to Intel QuickAssist Technology (Intel QAT) helps maintain strong performance while operating within a lower power envelope, which is particularly beneficial for proxy and service mesh applications.

## Learn More

[Intel® Xeon® 6 Processors](#)

[Intel® QuickAssist Technology \(Intel® QAT\)](#)

[FD.io VPP-SSwan and Linux-CP](#)

[Intel® Quick Assist Technology \(Intel® QAT\) - NGINX\\* Performance](#)

[Service Mesh – Istio and Envoy Optimizations for Intel® Xeon® Scalable Processors Solution Brief](#)

[Next Generation Firewall – Optimizations with 4th Gen Intel® Xeon® Scalable Processor](#)

[Intel® AVX-512 - Packet Processing with Intel® AVX-512 Instruction Set](#)

[Accelerate Cryptographic Operations and Data Compression Workloads with Intel QAT](#)



## Notices & Disclaimers

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 [www.intel.com/PerformanceIndex](http://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.