Benchmarking Network Performance and Power Efficiency on AMD EPYC

White Paper

Virtualized Networking Software

www.6wind.com

Table of Content

Executive Summary

This white paper evaluates the performance scalability and power-efficiency of the 6WIND Virtual Service Router (VSR) running on AMD EPYC™ processors.

Using a reproducible and transparent benchmarking methodology, the study characterizes packet forwarding and IPsec performance across a wide range of packet sizes, traffic profiles, and CPU core allocations

These results highlight AMD EPYC processors as a high-performance, energy-efficient foundation for software-defined networking High core density, DDR5 memory bandwidth, and PCIe® Gen5 I/O enable 6WIND VSR to scale linearly across CPU cores while delivering predictable throughput and strong performance per watt in data center and telco edge deployments

Key findings include:

- Near-linear processor core scaling for IPv4 and IPv6 forwarding up to the aggregate 400 Gbps bidirectional NIC line rate
- 400 Gbps line-rate forwarding achieved with single-digit core counts (~4 cores) for large packet profiles
- Predictable packet-per-second (PPS) scaling for 64-byte traffic up to 16 CPU cores
- Linear IPsec AES-GCM-256 scalability with no observed efficiency collapse
- Peak IPv4 forwarding efficiency exceeding approximately 4 Gbps/W

These results demonstrate that 6WIND VSR running on AMD EPYC processors, delivers a balanced and energy-efficient platform suitable for high-throughput data center, cloud gateway, and power-sensitive telco edge deployments

Abstract

This white paper presents performance benchmarking results for the 6WIND Virtual Service Router (VSR) running on AMD EPYC server processors.

The objective is to characterize data-plane forwarding, throughput, scalability, and power efficiency under representative service-provider and enterprise edge workloads.

The results show predictable CPU scaling and efficient packet processing across a wide range of packet sizes, up to the 400 Gbps aggregate NIC line rate.

6WIND VSR leverages a high-performance, DPDK-based I/O architecture and an accelerated Fast Path packet-processing engine designed for deterministic, multi-core scalability

Background

As network functions virtualization (NFV) and cloud-native architectures become mainstream, software-based routing solutions must deliver deterministic performance at 100–400 Gbps while preserving deployment flexibility.

Cloud and telco environments require network functions to support both east–west and north–south traffic at scale, integrate with orchestration frameworks, and adapt dynamically to changing workloads.

Use cases such as multi-gigabit edge routing, virtual CPE, cloud interconnect, and secure access services impose strict requirements on throughput, latency stability, and power-efficiency. In this context, predictable CPU scaling and efficient packet processing are critical design goals.

6WIND Virtual Service Router is optimized for these requirements through an accelerated Fast Path data plane that minimizes per-packet processing overhead and reduces latency variability. Its architecture is designed to scale efficiently on modern multi-core processors, making it well suited for cloud, edge, and service-provider deployments

AMD EPYC Processors

AMD EPYC 9004 Series processors span multiple product families addressing both performance-driven data center workloads and power-constrained telco edge environments.

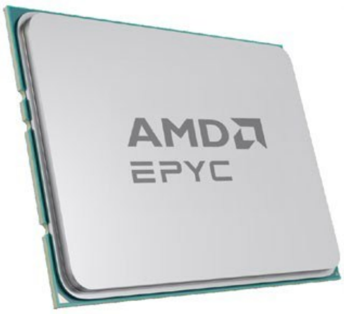

The 4th Gen AMD EPYC processors, built on Zen 4 and Zen 4c architectures optimized for high core density and per-socket performance, targets high-throughput cloud and core data center environments. A balanced memory and I/O subsystem, including support for DDR5 memory and PCIe® Gen5 connectivity, enables efficient scaling for network-intensive workloads.

Support for DDR5 memory increases aggregate memory bandwidth, while PCIe Gen5 connectivity enables attachment of high-speed network adapters required for 100G, 200G, 400G and higher-speed Ethernet deployments commonly found in modern data centers.

Complementing the AMD EPYC 9004 series, the EPYC 8004 (codenamed “Siena”) family is designed for power-efficiency and footprint optimization in telecommunications and edge data center environments.

While offering lower peak core counts and I/O bandwidth, the 8004 series provides sufficient compute capability for routing, packet processing, and security functions with reduced power consumption and simplified cooling requirements

AMD EPYC 8004 processors target scenarios such as distributed telco edge, virtualized RAN, and access and aggregation nodes, where performance per watt, thermal constraints, and system density are critical considerations

Figure 1: 4th Gen AMD EPYC Family

For software-defined networking workloads, these processor characteristics translate into high packet-processing efficiency, predictable multi-core scaling, and strong performance per watt—key requirements for both high-throughput data center gateways and power-constrained telco edge deployments

Solution Overview – 6WIND VSR Architecture

6WIND VSR separates the control plane from an accelerated data plane referred to as the Fast Path.

The Fast Path is responsible for all performance-critical packet processing functions, including packet reception and transmission, forwarding lookups, policy enforcement, and service chaining. It is explicitly designed to scale across multiple CPU cores, with each Fast Path instance operating independently to process traffic in parallel

The 6WIND VSR architecture provides highly deterministic performance characteristics that are essential for both high-throughput data center deployments and latency-sensitive telco cloud use cases.

The number of processor cores dedicated to the Fast Path can be precisely tuned based on target throughput, packet size, and enabled services.

This flexibility allows operators to balance performance, power consumption, and consolidation density across a wide range of deployment scenarios.

Deployment Model

6WIND VSR supports bare-metal, virtual machine, and containerized deployments.

In this benchmarking study, 6WIND VSR was deployed in a virtualized environment using a KVM-based hypervisor. SR-IOV was enabled to provide direct NIC access, CPU cores were explicitly pinned to Fast Path threads, and NUMA locality was carefully managed to ensure predictable and repeatable performance.

Test Objectives

The objectives of this benchmarking effort were to validate forwarding and IPsec scalability, quantify performance-per-core and performance-per-watt behavior, and analyze how throughput evolves as additional CPU cores are allocated to the data plane

Baseline IPv4 and IPv6 forwarding performance was measured across a wide range of packet sizes, from minimum-size packets to full MTU frames. Additional scenarios included IPsec, VLAN, and MPLS to evaluate the impact of incremental processing complexity.

These benchmarks are intended to characterize workload-specific behavior and scaling trends for 6WIND VSR running on AMD EPYC processors, including CPU scalability, throughput, and power-efficiency trends.

No direct vendor-to-vendor performance comparison is implied, and results should be interpreted as a workload-specific evaluation of 6WIND VSR on AMD EPYC processors.

Testbed Configuration

Hardware Configuration

- Server platform: Lenovo ThinkEdge SE455
- Processor: AMD EPYC 8534P @ 2.3 GHz (64 cores)
- Memory: 384 GB (6x64G) DDR5 @ 4800 MHz
- Network interfaces: Nvidia ConnectX-7 (2x200 Gbps) - MCX755106AS-HEAT
- BIOS settings: SMT on, Power profile performance, NPS1, C-states disabled
- Determinism Slider: Power, L1/L2 Stream HW Prefetcher: Enabled, ACPI SRAT L3 Cache as NUMA Domain: Disabled, Global C State Control: Disabled, DF C-States: Disabled.

Software Configuration

- Host operating system: Ubuntu 22.04.5 LTS
- Kernel Version: 6 8 0-90-generic #91~22.04.1
- Ubuntu Virtualization stack: KVM/Qemu v6.2.0, vCPU pinning policy
- Hugepages: 16 pages (1G page size)
- Traffic generator: Trex v3.04
- 6WIND VSR version: v3.10
  - Number of interfaces and VRFs: 2 interfaces / 1 VRF
  - Routing configuration: Static
  - Data plane / Fast Path:
    - Core allocation: Test dependent
    - SMT enabled at the BIOS level
    - Fast Path threads pinned to one hardware thread per core
    - Core/port mapping applied
    - Offloads: Disabled

Bidirectional traffic is generated across both 200 Gbps NIC ports to achieve 400 Gbps aggregate throughput

Figure 2: test setup

Methodology

Metrics

Measured metrics include throughput (Gbps), packet rate (pps), and CPU power consumption Power was measured using the Linux turbostat tool, which reports CPU package-level power. Measurements exclude NIC and system-level power and are intended for relative efficiency comparison rather than absolute platform TCO evaluation.

CPU package-level power measurement was selected to isolate processor efficiency and scaling behavior, independent of NIC or system-level variables. This approach enables clearer analysis of how packet forwarding and security workloads translate CPU resources into usable throughput.

Latency was intentionally excluded, as the focus of this study is throughput scalability and power-efficiency.

Traffic Profiles

The setup uses bidirectional traffic to saturate the 400 Gbps line rate Traffic profiles included fixed 64-byte packets, multiple IMIX profiles with average packet sizes of approximately 350 and 700 bytes, and a large-packet profile using 1400-byte frames Both single-flow and multi-flow scenarios were tested.

Test Procedure

Each test included a warm-up period to reach a steady operational state. This prevented cache, CPU frequency, and I/O initialization transients from affecting the results. Measurements were then collected over a fixed duration to capture stable throughput, packet rate, and power consumption behavior.

All tests were executed using an automated benchmarking framework to ensure consistency. Each scenario was run multiple times under identical conditions, with results reported as the median to reduce the impact of outliers.

System configuration and power management settings—including CPU frequency policies, power profiles, BIOS settings, virtualization parameters, and automation workflows—were kept constant throughout. This ensured that observed performance and power-efficiency trends were attributable only to changes in traffic profiles, enabled services, and Fast Path core allocation.

Benchmark Results

Across all scenarios, throughput increases predictably as additional CPU cores are allocated. Large packet profiles reach the 400 Gbps capacity with relatively few cores, while small-packet profiles remain CPU-bound and continue scaling with additional cores.

Baseline Forwarding Performance

- IPv4 forwarding:

	Bidirectional Bandwidth (Gbps)					
Packet Size	**1C**	**2C**	**4C**	**8C**	**16C**	**32C**
64 (64B)	9.34 Gbps	19.61 Gbps	40.29 Gbps	76.79 Gbps	83.17 Gbps	98.35 Gbps
IMIX 350 (349B)	41.49 Gbps	87.01 Gbps	173.26 Gbps	221.29 Gbps	323.02 Gbps	357.89 Gbps
IMIX 700 (696B)	80.25 Gbps	165.33 Gbps	322.06 Gbps	371.56 Gbps	400.00 Gbps	400.00 Gbps
1400 (1400B)	156.63 Gbps	314.50 Gbps	399.77 Gbps	400.00 Gbps	400.00 Gbps	400.00 Gbps

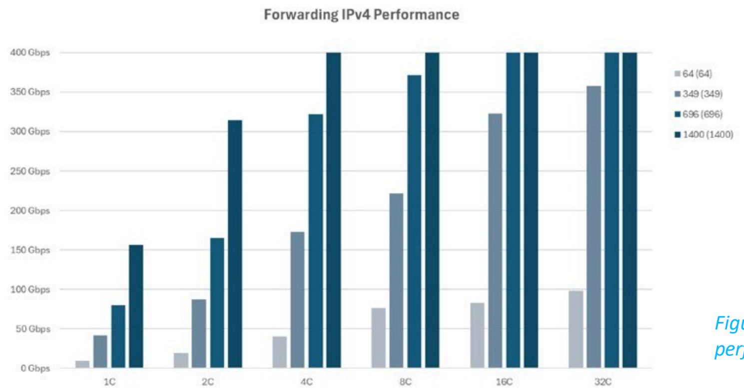

Figure 3: IPv4 forwarding performance

- IPv6 forwarding:

	Bidirectional Bandwidth (Gbps)					
Packet Size	**1C**	**2C**	**4C**	**8C**	**16C**	**32C**
64 (64B)	8.92 Gbps	18.69 Gbps	35.42 Gbps	73.47 Gbps	82.97 Gbps	98.35 Gbps
IMIX 350 (349B)	39.54 Gbps	82.87 Gbps	153.60 Gbps	213.02 Gbps	320.53 Gbps	357.89 Gbps
IMIX 700 (696B)	76.50 Gbps	159.13 Gbps	285.92 Gbps	379.19 Gbps	400.00 Gbps	400.00 Gbps
1400 (1400B)	149.30 Gbps	302.01 Gbps	399.50 Gbps	400.00 Gbps	400.00 Gbps	400.00 Gbps

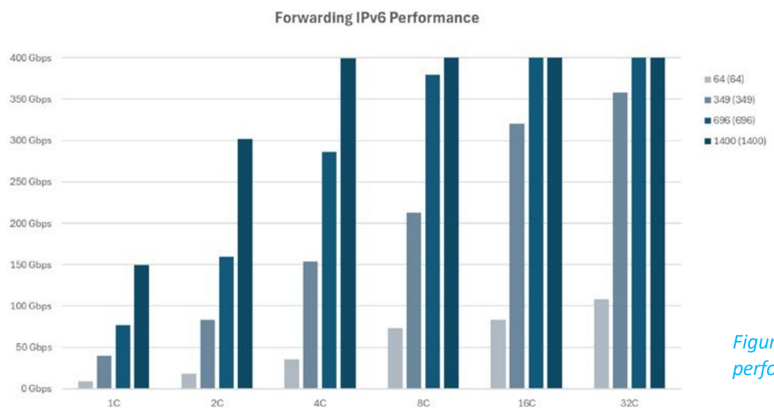

Figure 4: IPv6 forwarding performance

IPsec Performance

- IPsec ESP 4in4 AEAD=AES-GCM key-size=256 Crypto=multibuffer:

	Bidirectional Bandwidth (Gbps)					
Packet Size	**1C**	**2C**	**4C**	**8C**	**16C**	**32C**
64 (64B)	3.34 Gbps	7.05 Gbps	14.05 Gbps	22.20 Gbps	43.64 Gbps	82.17 Gbps
IMIX 350 (349B)	8.13 Gbps	16.97 Gbps	31.29 Gbps	53.34 Gbps	81.32 Gbps	147.77 Gbps
IMIX 700 (696B)	12.10 Gbps	25.09 Gbps	45.49 Gbps	75.00 Gbps	103.90 Gbps	184.70 Gbps
1400 (1400B)	19.69 Gbps	39.38 Gbps	65.75 Gbps	85.24 Gbps	129.05 Gbps	247.45 Gbps

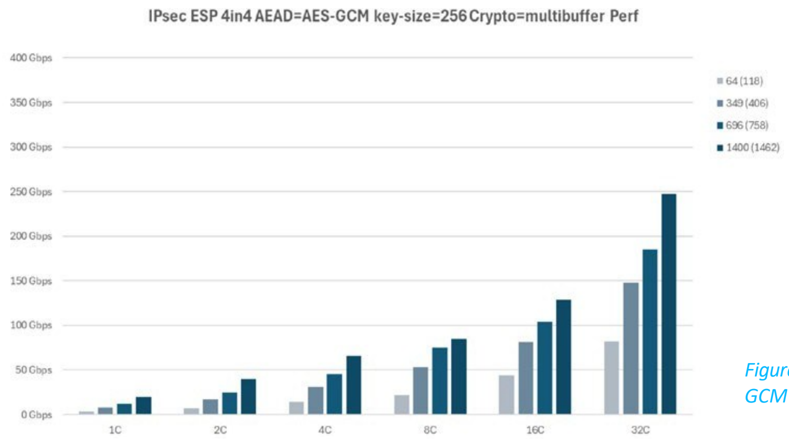

Figure 5: IPsec ESP 4in4 - AES GCM 256

- IPsec ESP 6in6 AEAD=AES-GCM key-size=256 Crypto=multibuffer:

	Bidirectional Bandwidth (Gbps)					
Packet Size	**1C**	**2C**	**4C**	**8C**	**16C**	**32C**
64 (64B)	2.79 Gbps	5.92 Gbps	11.67 Gbps	11.67 Gbps	40.86 Gbps	77.35 Gbps
IMIX 350 (349B)	6.59 Gbps	13.92 Gbps	26.53 Gbps	26.53 Gbps	80.97 Gbps	141.61 Gbps
IMIX 700 (696B)	9.92 Gbps	21.42 Gbps	39.93 Gbps	39.93 Gbps	100.44 Gbps	176.01 Gbps
1400 (1400B)	16.66 Gbps	35.06 Gbps	60.23 Gbps	60.23 Gbps	123.47 Gbps	228.38 Gbps

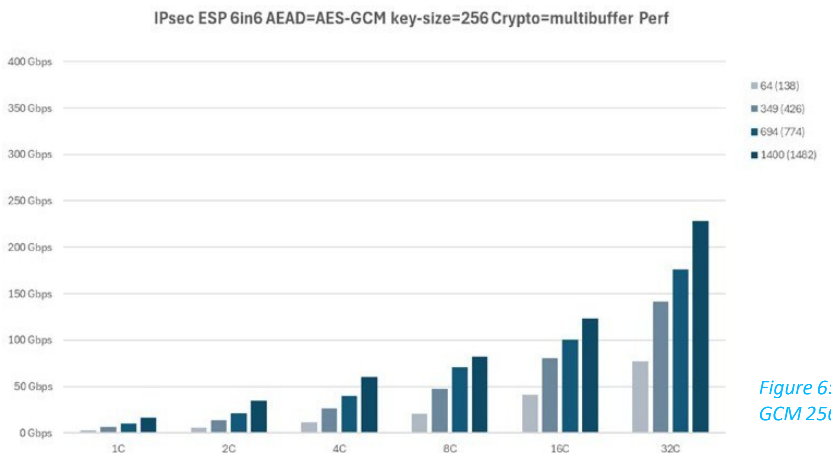

Figure 6: IPsec ESP 6in6 - AES GCM 256

VLAN Performance

	Bidirectional Bandwidth (Gbps)					
Packet Size	**1C**	**2C**	**4C**	**8C**	**16C**	**32C**
64 (64B)	8.46 Gbps	17.91 Gbps	34.16 Gbps	64.86 Gbps	87.38 Gbps	102.38 Gbps
IMIX 350 (349B)	36.72 Gbps	76.24 Gbps	142.79 Gbps	194.90 Gbps	310.73 Gbps	361.79 Gbps
IMIX 700 (696B)	69.90 Gbps	146.20 Gbps	266.22 Gbps	366.00 Gbps	399.86 Gbps	400.00 Gbps
1400 (1400B)	136.30 Gbps	278.50 Gbps	398.80 Gbps	400.00 Gbps	400.00 Gbps	400.00 Gbps

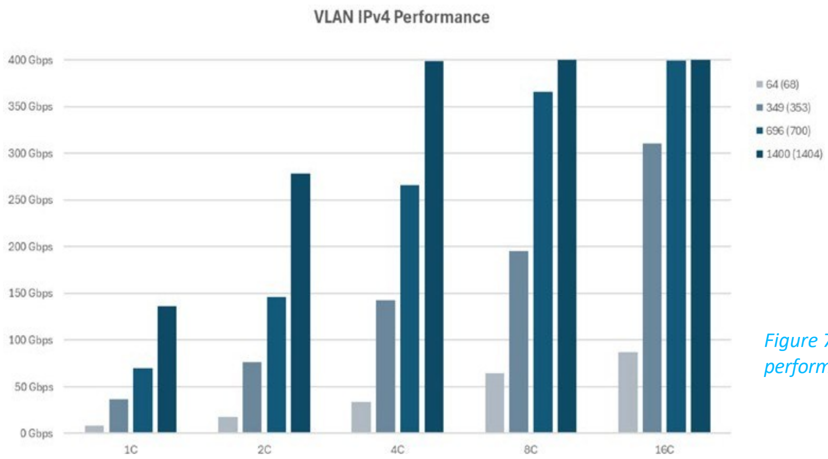

Figure 7: VLAN IPv4 performance

MPLS Performance

	Bidirectional Bandwidth (Gbps)					
Packet Size	**1C**	**2C**	**4C**	**8C**	**16C**	**32C**
64 (64B)	8.35 Gbps	17.61 Gbps	35.32 Gbps	62.99 Gbps	77.59 Gbps	97.59 Gbps
IMIX 350 (349B)	37.12 Gbps	77.68 Gbps	145.86 Gbps	189.68 Gbps	268.91 Gbps	357.65 Gbps
IMIX 700 (696B)	71.94 Gbps	149.81 Gbps	271.38 Gbps	370.48 Gbps	371.93 Gbps	400.00 Gbps
1400 (1400B)	140.20 Gbps	285.52 Gbps	399.04 Gbps	400.00 Gbps	400.00 Gbps	400.00 Gbps

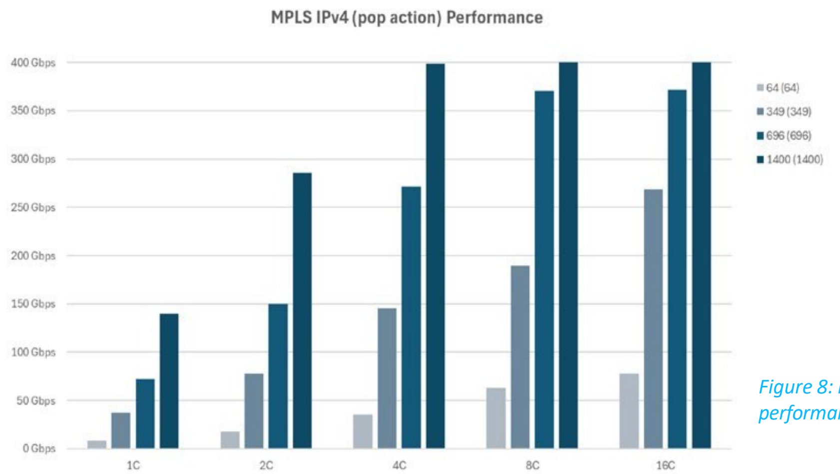

Figure 8: MPLS IPv4 performance

Power Measurements

Power-efficiency analysis correlates achieved throughput with CPU power consumption.

For forwarding workloads, efficiency improves rapidly at low to mid core counts before declining as the system becomes I/O-bound.

For IPsec workloads, efficiency increases monotonically with core count, reflecting a compute-bound workload

Baseline Forwarding Efficiency

	Efficiency (Gbps/W)					
Packet Size	**1C**	**2C**	**4C**	**8C**	**16C**	**32C**
64 (64B)	0.12	0.24	0.45	0.75	0.73	0.67
IMIX 350 (349B)	0.52	1.04	1.94	2.16	2.73	2.45
IMIX 700 (696B)	0.99	1.99	3.54	3.71	3.46	2.74
1400 (1400B)	1.90	3.69	4.37	3.96	3.47	2.77

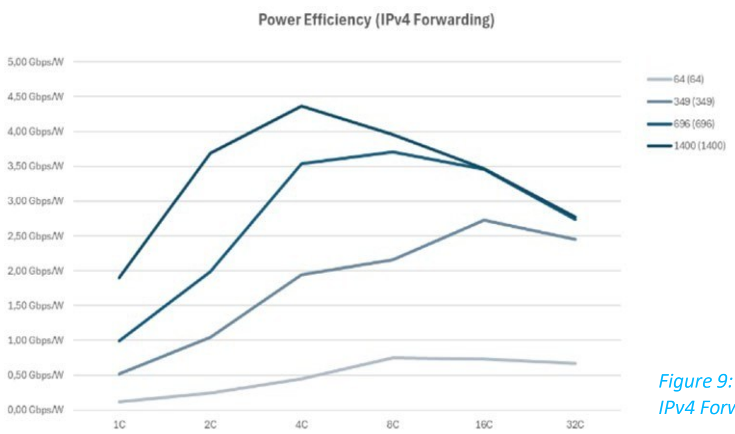

Figure 9: Power Efficiency - IPv4 Forwarding

As shown in the figure, for IPv4 forwarding, power-efficiency improves rapidly as CPU cores scale from 1C to mid-range configurations, driven by the amortization of fixed platform power over increasing throughput.

Larger packet profiles achieve the highest efficiency, peaking at approximately ~4 4 Gbps/W, while smaller packets remain PPS-limited and less energy efficient.

Efficiency peaks around 4C to 8C, after which it gradually declines as the system approaches the 400 Gbps line-rate and additional CPU power no longer translates into proportional throughput gains.

IPsec Efficiency

	Efficiency (Gbps/W)					
Packet Size	**1C**	**2C**	**4C**	**8C**	**16C**	**32C**
64 (64B)	0.04	0.09	0.16	0.23	0.39	0.56
IMIX 350 (349B)	0.10	0.21	0.36	0.54	0.73	1.05
IMIX 700 (696B)	0.15	0.31	0.52	0.76	0.96	1.34
1400 (1400B)	0.25	0.47	0.73	0.89	1.23	1.88

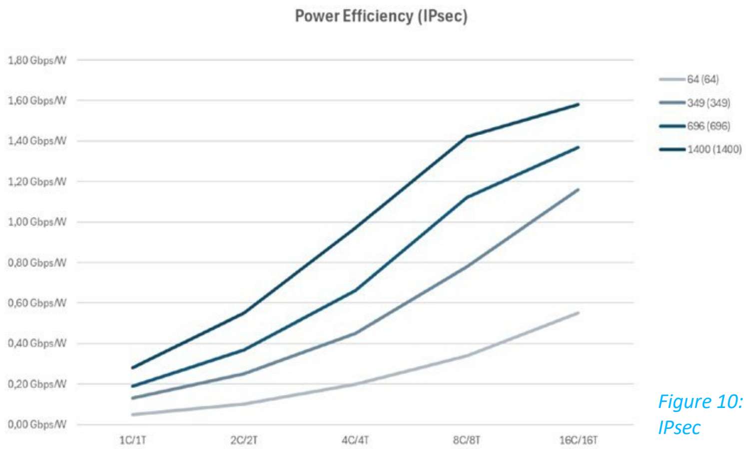

Figure 10: Power Efficiency – IPsec

In contrast to IPv4 forwarding, the previous figure shows that IPsec power-efficiency increases monotonically with core count across all traffic profiles There is no mid-range efficiency peak within the tested configurations:

- Efficiency continues to improve from 1C through 32C
- This reflects a fully CPU-bound workload, where additional cores translate directly into higher encrypted throughput
- The absence of NIC saturation allows power to be converted into useful work with high efficiency

Packet size has a pronounced impact on IPsec efficiency:

- 1400-byte packets achieve the highest efficiency, approaching ~1 9 Gbps/W
- Mid-sized packets (696, 349 bytes) scale similarly with lower absolute efficiency
- 64-byte packets show the lowest Gbps/W due to high per-packet crypto overhead Despite this spread, all profiles show consistent upward efficiency trends as cores increase

Results – Forwarding Performance Analysis

The results show strong forwarding performance even at low core counts, indicating high efficiency per CPU core.

Large-packet traffic profiles achieve substantial throughput with only a small number of cores, demonstrating low per-packet processing overhead and efficient use of CPU pipelines, caches, and memory subsystems.

This shows that each additional AMD EPYC core contributes meaningful forwarding capacity rather than marginal gains.

Across all traffic profiles, throughput increases proportionally with the number of processor cores. From single-core operation through multi-core configurations, the forwarding path scales predictably and consistently, with no signs of contention, lock saturation, or shared-resource limitations. This linear scaling continues until the system reaches the 400 Gbps NIC line rate, validating the effectiveness of the platform's parallel processing design.

Small-packet traffic profiles, which are inherently packet-per-second (PPS) intensive, provide a stringent test of CPU scalability In these scenarios, throughput continues to scale with additional cores well beyond the point where large-packet profiles have already saturated the NIC. This behavior confirms that the forwarding engine remains CPU-bound under PPS stress and that additional cores translate directly into increased packet processing capacity.

The following chart, built based on the IPv4 forwarding results, illustrates the platform's near-linear CPU scalability across multiple traffic profiles, clearly showing how forwarding throughput increases proportionally with the number of CPU cores until the 400 Gbps NIC line rate is reached.

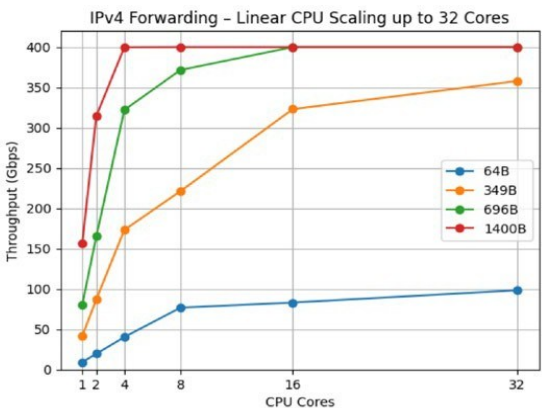

Figure 11: Forwarding performance - linear scalability

For PPS-intensive workloads such as the 64-byte profile, throughput scales almost perfectly linearly from 1 to 8 cores, demonstrating that each additional core contributes directly to packet processing capacity with no observable contention or efficiency loss.

Larger packet profiles reach line rate with fewer cores, flattening at 400 Gbps once the NIC becomes the limiting factor rather than the CPU.

This behavior confirms both high per-core forwarding efficiency and a clean transition from CPU-bound to I/O-bound operation, validating the platform's ability to fully utilize available CPU resources and to scale predictably under increasing parallelism.

The following figure shows throughput per core (Gbps/core) vs CPU cores, normalized from the raw bandwidth numbers.

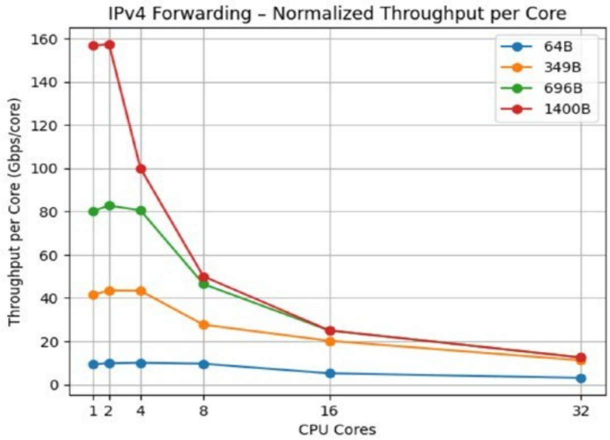

Figure 12: IPv4 Forwarding - Normalized Throughput per Core

The figure highlights the existence of two distinct and well-defined scaling regimes in the platform's forwarding behavior From one to four CPU cores, the normalized throughput per core is stable (i e throughput increases almost linearly), indicating a purely CPU-bound regime in which additional cores translate directly into higher forwarding capacity with no observable contention, synchronization overhead, or software inefficiency.

Beyond four cores, the slope of the curve gradually decreases, marking a controlled transition into bandwidth- and packet-rate-limited operation rather than a loss of CPU or software efficiency. This non-linearity is smooth, packet-size dependent, and aligned with known external limits, demonstrating that the forwarding data path remains scalable and well balanced while efficiently driving the platform toward its physical I/O constraints.

The following figure provides a representation of these two regions based on the IMIX 350B profile:

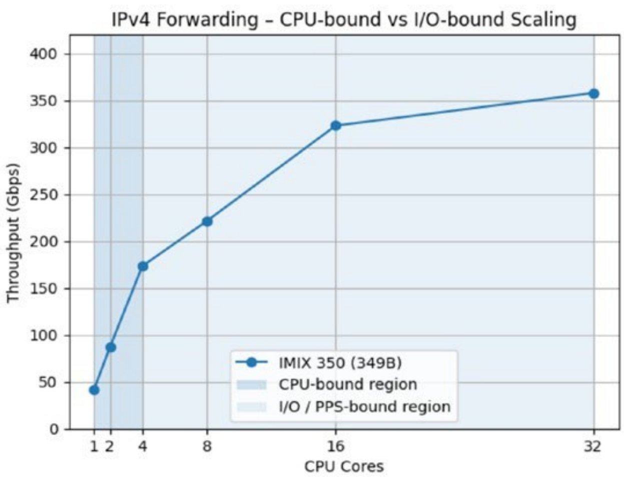

Figure 13: CPU-bount vs I/O-bound Scaling

IPsec Performance Analysis

The benchmark evaluates IPsec ESP tunnel performance using AES-GCM-256, with the objective of validating CPU cryptographic throughput, per-core efficiency, and scalability Unlike pure forwarding, IPsec introduces substantial per-packet cryptographic overhead, making it an effective stress test of CPU compute pipelines, vectorization, and memory bandwidth.

At low core counts, throughput is deliberately limited, reflecting the high computational cost of authenticated encryption. Even so, the results demonstrate meaningful throughput per core, particularly for larger packet profiles, indicating effective use of CPU crypto acceleration, instruction-level parallelism, and batching. As packet sizes increase, the amortization of crypto overhead becomes apparent, with significantly higher Gbps per core compared to small-packet profiles.

The following chart is built based on the IPsec performance results to illustrate the observed performance scalability:

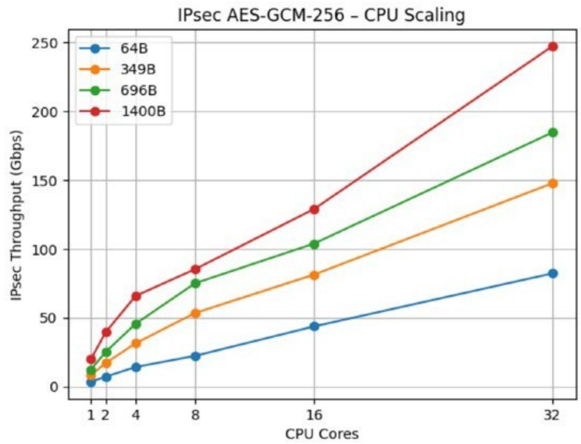

Figure 14: IPsec Performance - Near-linear scalability

Across all packet sizes, IPsec throughput increases steadily with the number of processor cores:

- From 1C to 32C, throughput scales predictably without abrupt flattening
- No early saturation or inflection points are observed
- Each additional core contributes incremental crypto capacity

This behavior confirms that the IPsec data path is CPU-bound across the tested core range, and that the multi-buffer crypto implementation scales efficiently across cores without contention or serialization bottlenecks

Power Efficiency Analysis

Taken together, the power-efficiency results from both IP forwarding and IPsec show that:

- IPv4 forwarding on AMD EPYC processors delivers excellent performance per watt, with a clear efficiency sweet spot before I/O saturation
- IPsec processing benefits from continued CPU scaling, improving both throughput and efficiency as more cores are applied
- The processor architecture efficiently handles both lightweight forwarding and heavy cryptographic workloads with predictable and explainable power behavior

These results reinforce the conclusion that the platform provides a balanced, energy-efficient architecture capable of scaling across diverse networking and security workloads.

The figure below illustrates the relationship between throughput and power-efficiency for IP forwarding and IPsec processing, highlighting the fundamentally different scaling characteristics of the two workloads.

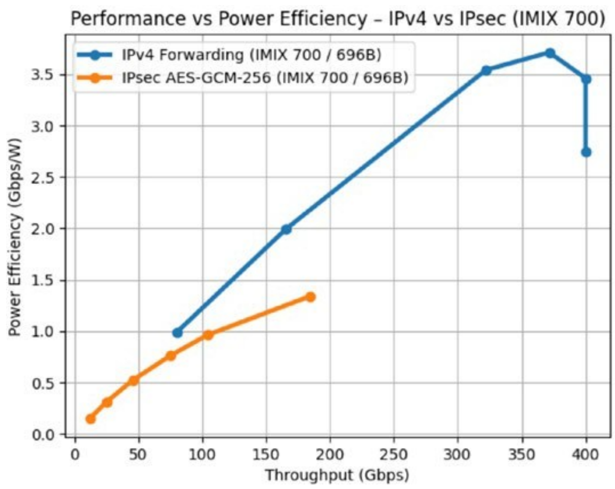

Figure 15: Performance vs Power Efficiency

For IP forwarding, power-efficiency increases rapidly as throughput rises, reaching a clear peak at mid-range throughput levels before declining as the system approaches the 400 Gbps NIC line rate.

This behavior reflects the transition from a CPU-bound regime, where additional performance is gained efficiently, to an I/O-bound regime, where throughput plateaus while power consumption continues to increase.

In contrast, IPsec shows a monotonic improvement in power-efficiency as throughput increases, indicating that the workload remains compute-bound across the measured range.

Each increase in encrypted throughput translates directly into more effective use of CPU cycles and power, with no efficiency collapse observed.

Together, the curves demonstrate that the platform efficiently scales both cleartext forwarding and cryptographic processing, while clearly exposing the different power-performance trade-offs driven by I/O versus compute limitations.

Implications for Deployment Scenarios

- ► Telco edge deployments benefit from high forwarding efficiency at low core counts, enabling compact and power-efficient designs.
- ► Cloud gateways can elastically scale throughput by allocating additional CPU resources.
- ► Secure WAN and SASE use cases benefit from linear IPsec scalability without efficiency degradation.

Conclusion

The benchmark results demonstrate that 6WIND VSR running on AMD EPYC processors delivers very high efficiency for software-based routing and security workloads. Forwarding performance scales predictably until reaching NIC line rate, while encrypted workloads remain compute-bound and continue to benefit from additional CPU cores, reflecting the combined behavior of the VSR data plane and the underlying processor architecture.

These characteristics allow operators to select configurations optimized for either maximum throughput or optimal energy efficiency, confirming the platform as a scalable, power-efficient foundation for modern cloud, edge, and service-provider networking deployments

References

1 6WIND Virtual Service Router (VSR) Retrieved from 6WIND Documentation: https://www.6wind.com/vrouter-vsr-solutions/

2 AMD EPYC Processors Retrieved from AMD website: https://www.amd.com/en/products/processors/server/epyc.html

3 AMD EPYC 8004 Series Processors Datasheet: https://www.amd.com/content/dam/amd/en/documents/products/epyc/amd-epyc-8004-series-processors-datasheet.pdf

4 AMD EPYC 8004 Purpose built CPU for Power efficiency: https://www.amd.com/en/partner/articles/epyc-8004-purpose-built-efficiency.html

5 AMD 8004 Tuning Guide: https://docs.amd.com/v/u/en-US/58310_amd-epyc-8004-tg-data-plane-dpdk

Notices & Disclaimers: Performance benchmarks vary by configuration and other factors Performance results are based on the testing environment described in the testbed description section.