White Paper

Enterprise, Cloud, Telecommunications
Cybersecurity, Network Security

Next-Generation High-Density DDoS Mitigation System for Tbps-Scale Traffic

Reference architecture, data plane engineering, and production benchmarks on the Intel® 6+ Processors

CONNYON

Table of Contents

Modern DDoS Threat Landscape 1
Limits of Legacy Mitigation Architecture 2
Intel® Xeon® 6+ Processor Hardware Foundations.....2
The Siware-Kernel Data Plane.....3
Staircase Vectoring & Data Plane Optimization.....4
Multi-Stage Filtering Funnel Architecture 4
Out-of-Band AI Strategy Generation ..5
BGP Flowspec Cooperative Scaling .. 6
Empirical Benchmarks & Production Validation 6
Conclusion & Architectural Recommendations.....7
Key Design Takeaways 7

Executive Summary

As the volume and sophistication of Distributed Denial of Service (DDoS) attacks continue to scale, legacy security architectures are encountering architectural boundaries. Many network security solutions are deployed through systems based on specialized Application-Specific Integrated Circuit (ASIC). While performant, these systems are constrained by long update cycles and high capital costs for the solution providers, due to the nature of ASIC economics. Alternatively, utilizing conventional CPU-based general purpose servers may have lacked the compute density required for line-rate packet processing.

This whitepaper presents the architecture and performance characteristics of the Connyon Nova Shield NS-X4000, a software-defined, next-generation high-density DDoS mitigation appliance capable of scrubbing up to 4.0 Tbps of sustained traffic and processing 2,400 Mpps of small packets (64-byte) at sub-5 µs latencies on a standard general-purpose server platform. The solution is enabled by a dual-socket platform configured with Intel® Xeon® 6+ processors (formerly Clearwater Forest), manufactured on the Intel 18A process and delivering up to 576 Efficiency-cores (E-cores) in a single server system.

Central to the software architecture is the Siware-Kernel, a unified execution environment that replaces fragmented DPDK, VPP, and eBPF layers. It combines a 5-stage staircase vector scheduler, lockless localized data structures, adaptive funnel filtering, and out-of-band AI policy synthesis to deliver hardware-class throughput with Intel architecture flexibility.

1. The Modern DDoS Threat Landscape

Modern DDoS attacks are becoming shorter, denser, and more automated. The bottleneck is shifting from raw bandwidth to packet-rate processing and programmable filtering.



Figure 1. DDoS Threat Landscape

1. **Hyper-Volumetric Surges:** Peak volumetric attacks have broken historical records, reaching up to 31.4 Tbps. These short-burst, high-intensity attacks are designed to exhaust upstream pipes and saturate physical interfaces within seconds.
2. **Extreme Packet-per-Second (PPS) Density:** Attackers are increasingly leveraging stateless Layer 3 and Layer 4 packet surges, scaling peak packet rates up to 9.0 Bpps globally. High-frequency, small-packet floods (64-byte) target the processing capacity of network middleboxes rather than raw bandwidth, exposing CPU bottlenecks.
3. **Layer 7 and Algorithmic Complexity:** While Layer 3/4 volumetric floods represent 78% of the total attack footprint, Layer 7 application-layer attacks comprising the remaining 22% have grown in complexity by over 60% year-over-year. These attacks mimic legitimate user traffic, requiring deep packet inspection, stateful tracking, and intensive cryptographic verification.

Modern security gateways must absorb massive L3/L4 packet-rate pressure while retaining enough compute for L7 inspection, state tracking, and cryptographic validation.

2. Limits of Legacy Mitigation Architectures

Faced with this threat landscape, network operators have traditionally been trapped in a trilemma between **Performance**, **Programmability (Flexibility)**, and **Unit Cost**.

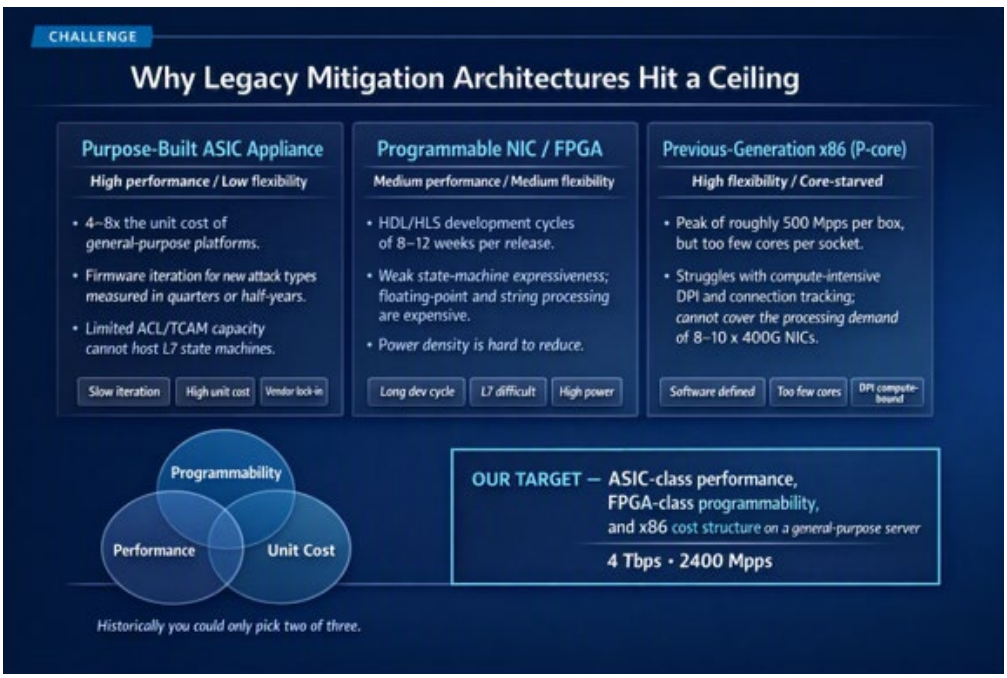


Figure 2. Legacy Mitigation Architecture Ceilings

Option A: Specialized ASIC-Based Hardware Firewalls

ASICs offer high packet processing performance but suffer from rigid hardware-defined pipelines.

- **Capital Expenses (CAPEX)**: Proprietary appliances cost 4x to 8x more than general-purpose x86 hardware.
- **Rigidity:** Updating an ASIC pipeline to mitigate a zero-day attack vector requires microcode or firmware rewrites, resulting in deployment cycles exceeding six months.
- **Stateful Limitations:** Hardened access control list (ACL) tables are bounded by static on-chip memory (TCAM), rendering them ineffective for large-scale Layer 7 state machines and dynamic IP blacklists.

Option B: FPGA-Based Programmable SmartNICs

FPGAs offer hardware-level acceleration and partial programmability, but introduce severe development bottlenecks.

- **Development Complexity:** Designing pipelines in HDL/HLS requires specialized engineering skillsets, with feature development loops averaging 8 to 12 weeks.
- **Resource Constraints:** FPGAs possess limited logic cells, restricting their ability to execute complex state machines, pattern matching, and deep payload analytics simultaneously.
- **Thermal and Power Densities:** High power consumption per card increases operational complexity in high-density data centers.

Option C: Previous-Generation and Intel® Xeon® Processors with P-Cores

Traditional standard servers rely on a limited number of high-frequency Performance-cores (P-cores).

- Core Starvation: Under small-packet flood conditions, a typical dual-socket P-core system hits processing ceilings at approximately 500 Mpps due to limited parallel execution contexts.

- Compute Jitter: Deep Packet Inspection (DPI) and cryptographic validation cause severe CPU cache thrashing, resulting in significant packet drops. To support 4 to 8 x 100G network interfaces, operators are forced to deploy multiple redundant servers, eroding cost efficiencies.

3. Intel® Xeon® 6+ Processor Hardware Foundations

Intel® Xeon® 6+ processor based on E-cores resolves this trilemma by offering an optimized architecture specifically for parallel, high-throughput network workloads.

The E-Core Paradigm for Network Workloads

Unlike conventional compute-intensive applications, network packet processing is **embarrassingly parallel**. Rather than high single-thread instruction-per-cycle (IPC) performance, network workloads require massive concurrency to scale. Each independent packet or flow can be processed in isolation:

- Independent packet or flow processing enables near-linear scaling across dedicated E-cores.
- By packing 288 physical E-cores onto a single die, Intel® Xeon® 6+ processor is architected to eliminate thread context-switching and scheduling jitter. In a dual-socket system, 576 dedicated execution engines are available, enabling a strict 1:1 mapping between hardware receive (RX) / transmit (TX) queues and dedicated CPU cores.

Cache and Memory Architecture

Intel® Xeon® 6+ processor provides 576 MB of Last Level Cache (LLC) per socket. Combined with **Intel® Direct Data I/O (Intel® DDIO)**, network interface cards DMA packet descriptors and payloads directly into the L3 cache rather than into the system memory. This keeps critical session tables, IP blacklists, and flow states cached hot in LLC, reducing main-memory fetch latency from approximately 100 ns to sub-10 ns.

With DDR5-8000 memory per socket, the system provides up to 1.5 TB/s of aggregate memory bandwidth, preventing memory bus saturation during heavy Layer 7 stateful tracking.

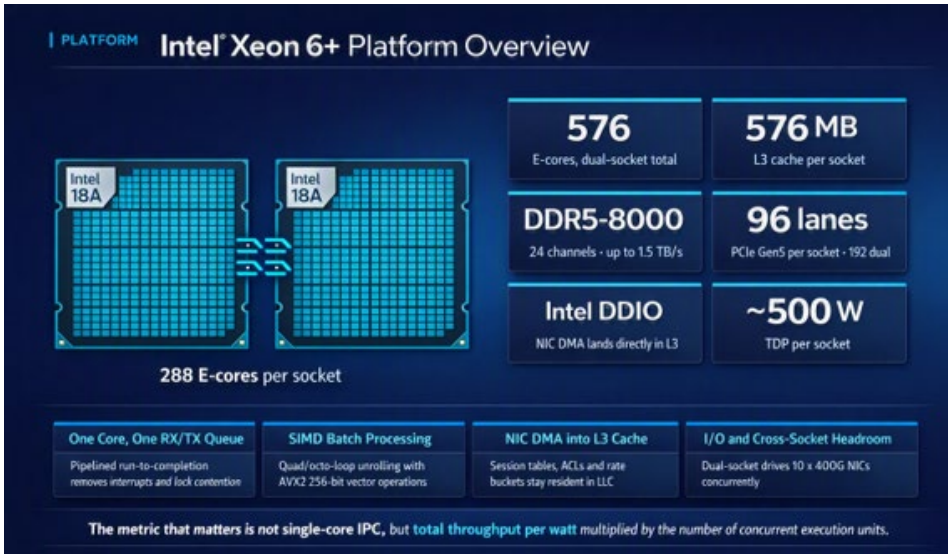


Figure 3. Intel® Xeon® 6+ Hardware Foundations

I/O and Socket Interconnect

Each socket features 96 lanes of high-speed PCIe Gen5 connectivity, allowing a dual-socket system to directly drive up to ten 400G network interfaces at line rate. High-speed UPI coherent interconnects operating at 24 GT/s over 16 lanes maintain cache coherency across sockets without bottlenecking cross-core telemetry and flow synchronization.

4. The Siware-Kernel Data Plane

Standard software-defined networking stacks typically chain DPDK (for kernel bypass), VPP (for vector packet processing), and eBPF (for dynamic filtering). This multi-tier architecture introduces significant context-switching overhead, packet-copy penalties, and Instruction-Cache (I-Cache) thrashing.

To overcome these performance penalties, this reference architecture introduces the **Siware-Kernel**, an integrated, high-performance general-purpose networking data plane solution.



Figure 4. Siware-Kernel Dataplane Architecture and Packet Processing Pipeline

Architectural Consolidation

The Siware-Kernel consolidated driver-level ingestion, vector dispatching, and stateful mitigation into a single unified execution context. It operates completely in user space, executing a **run-to-completion (RTC)** model on dedicated E-cores.

Key Data Plane Enablement

- **256-Packet Vector Processing:** Packets are processed in vectors of 256 packets. This amortizes packet-processing overhead across the batch, reducing instruction cache misses and branch misprediction rates.

- **One-Core RTC:** Packets are read, parsed, classified, mitigated, and forwarded on a single dedicated core, eliminating inter-thread message passing, lock contention, and NUMA boundary crossings.
- **Hot-Swappable Plugin Model:** Mitigation functions (e.g., Geo-Blocking, SYN flood protection, application-layer challenges) are compiled as modular runtime plugins. These can be hot-reloaded dynamically without restarting the packet processing engine.
- **Millisecond Config Reload:** Policy tables and ACL rules are updated via atomic pointer swaps with less than 200 ms reload time without interrupting active processing loops, ensuring zero packet loss during defense configuration updates.

5. Staircase Vectoring & Data Plane Optimization

To achieve a processing target of 2,400 Mpps, the Siware-Kernel data plane implements a **Five-Stage Staircase Vector Scheduling** pipeline, optimized for the E-core microarchitecture.

The Five-Stage Pipeline

This pipeline processes five batches of 256 packets in parallel at different stages of execution. While Batch n is in the forwarding stage, Batch n+1 is being mitigated, Batch n+2 is performing table lookups, and so forth. This staircase execution model ensures optimal instruction-level pipelining and hardware utilization.

- **Stage 1 - Parse Packet Headers (Intel® AVX2-enabled):** Extracts Layer 2/3/4 headers in parallel. Quad-loop architectures load 4 packet headers simultaneously into Intel® Advanced Vector Extensions 2 (Intel® AVX2) vector registers.
- **Stage 2 - Classify and Hashing:** Computes flow hashes using SIMD-accelerated 5-tuple Toeplitz hash instructions.
- **Stage 3 - Table Lookup (O(1) Array Indexing):** Maps hashes to flow states. Session matching is executed via pointer-free direct array lookups to avoid memory-indirect pointer chasing.
- **Stage 4 - Mitigate and Policy Enforcement:** Executes active filtering logic (e.g., matching against dynamic blacklists and active challenges).
- **Stage 5 - Forward Packet Handoff:** Places clean packets back into the TX ring of the respective network interface via DMA ring buffers.

Core Data Plane Optimization Principles

Lockless RCU (Read-Copy-Update): All shared state updates (such as dynamic blacklists and ACL rules) utilize lockless RCU patterns. Readers traverse lookup tables without locks, while the control plane performs atomic updates on copy-constructed tables before swapping pointers.

Per-Core Localized Memory: Shared global state is minimized. Session tables, packet metadata, and counter tracking are split locally per core. This architecture ensures zero cross-core communication during the fast-path packet cycle, eliminating UPI interconnect overhead.

Array Indexing over Pointer Chasing: Standard linked-list search techniques are replaced with pre-allocated, flat arrays. By indexing tables directly:

The engine resolves memory lookups in a deterministic O(1) cycle, eliminating cache misses associated with pointer-chasing structures.

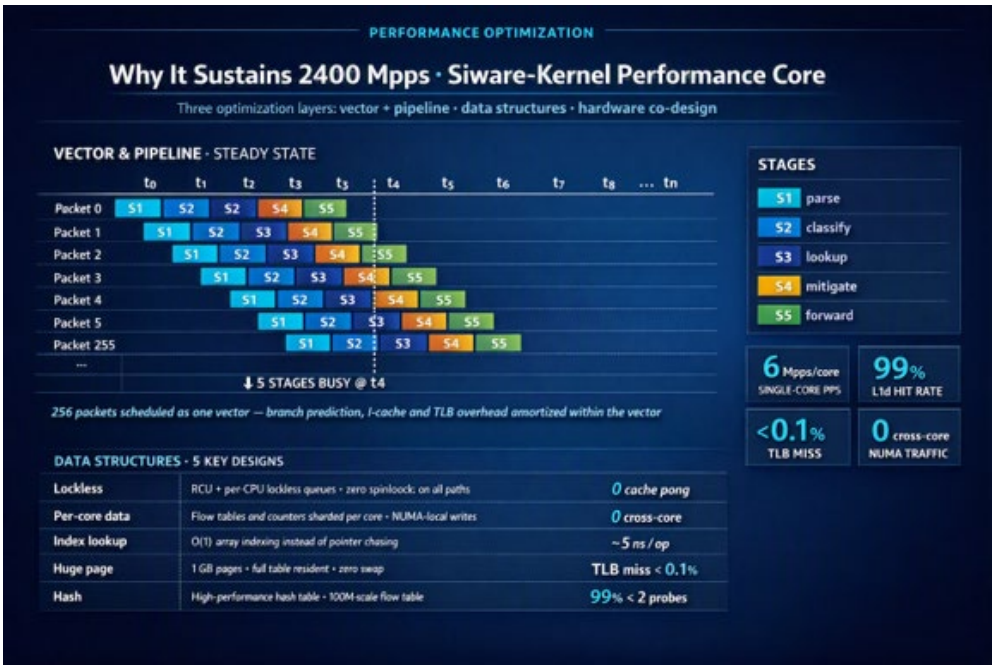


Figure 5. 5-Stage Staircase Vector Pipeline

1 GB Huge Pages: Standard 4 KB page tables are replaced by 1 GB Huge Pages, keeping the entire memory footprint resident in physical RAM. This reduces translation lookaside buffer (TLB) misses to near-zero during active volumetric attacks.

Perfect Hash Tables: The system utilizes perfect hash mapping optimized for dense networks. For 10M concurrent flows, lookups resolve in an average of 1.02 probes, preventing collision-induced latency spikes.

6. Multi-Stage Filtering Funnel Architecture

To process massive incoming traffic volumes cost-effectively, the reference architecture implements a multi-stage filtering funnel. The fundamental design principle is to drop malicious packets as early as possible in the processing path, minimizing compute overhead per filtered packet.

Level 01: Dynamic Match (100% to approximately 30%)

- **Mechanism:** Stateless matching against IP blacklists, blocklists, and geographical regions.
- **Implementation:** Executed on the SmartNIC utilizing hardware flow directors or in the early fast path.
- **Compute Cost:** Minimal O(1) lookup per packet.

Level 02: Feature Recognition (approximately 30% to approximately 8%)

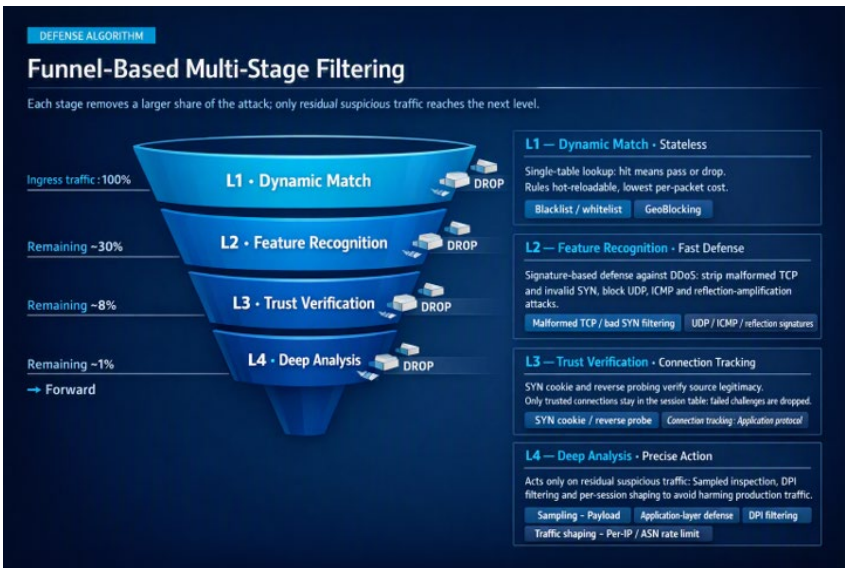


Figure 6. Multi-Stage Filtering Funnel (L1-L4)

- **Mechanism:** Malformed header detection, invalid TCP flags, invalid SYN-ACK sequences, UDP reflection fingerprints, and ICMP flood signatures.
- **Implementation:** Handled via Siware-Kernel vector execution loops using SIMD validation.
- **Compute Cost:** Extremely low.

Level 03: Trust Verification (approximately 8% to approximately 1%)

- **Mechanism:** Stateful active/passive challenges. For TCP connection requests, the engine executes cryptographic SYN cookies using: SYN cookie = H(src IP, dst IP, src port, dst port, timestamp, secret).

  It drop-forwards connection attempts unless the subsequent ACK contains a valid mathematical validation hash.
- **Implementation:** Siware-Kernel stateful handshake validation modules.
- **Compute Cost:** Moderate (requires cryptographic hashing).

Level 04: Deep Analysis (approximately 1% to clean egress)

- **Mechanism:** Full payload sampling , deep packet inspection (DPI) pattern matching, and adaptive rate-limiting per flow, IP, or autonomous system number (ASN).
- Implementation: Loaded on dedicated high-performance L7 engine threads inside the Siware-Kernel.
- **Compute Cost:** High (requires regex matching and stateful session tracking).

7. Out-of-Band AI Strategy Generation

While deep learning models excels at pattern matching, running real-time inline AI inference is incompatible with the nanosecond-level budget of a high-throughput data plane. Executing complex models within the packet loop introduces severe microsecond jitter and violates strict SLA latencies.

To address this, this architecture separates **inline forwarding** from **anomaly learning** using an out-of-band dual-pipeline model.

Dual-Pipeline AI Defense

- **Pipeline A - PCAP Payload Capture & Signature Generation:**

  During an active attack, the data plane samples raw packet payloads and streams them to the out-of-band engine. The AI subsystem runs unsupervised clustering models such as K-means clustering and Isolation Forest algorithms to extract attack fingerprints from malicious payloads. It then automatically synthesizes regular expressions and pushes them to the Level 4 DPI filter in the Siware-Kernel via atomic pointer swaps.

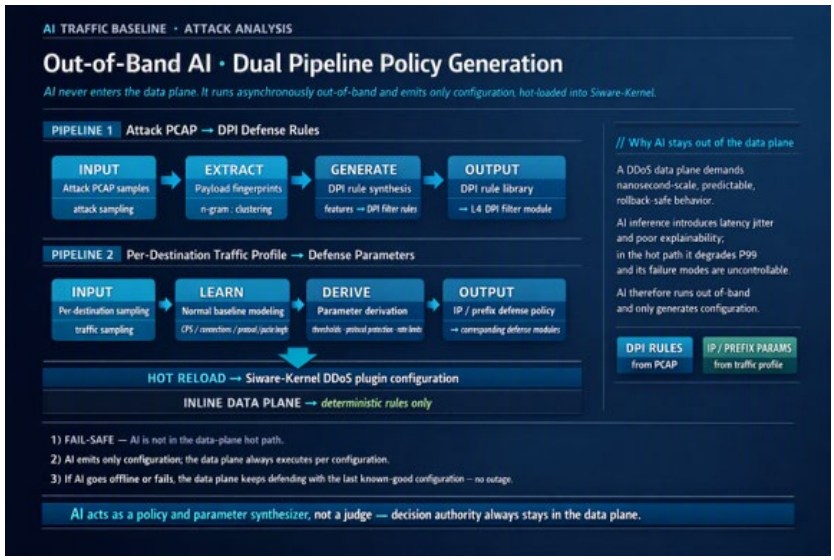


Figure 7. Out-of-Band AI Policy Generation (Dual Pipeline)

- **Pipeline B - Anomaly Detection & Adaptive Threshold Derivation:**
  The out-of-band engine continuously monitors normal traffic profiles for protected destinations. It trains baseline profiles (tracking metrics such as Connections-per-Second, Protocols, and payload length distribution). When deviations are identified, the AI engine derives dynamic parametric rate-limiting thresholds and pushes them down to the Level 4 rate limiter.

In this architecture, the inline data plane remains fast, standard, and completely deterministic, while the AI acts as a policy and parameter synthesizer, avoiding packet processing latency jitter.

8. BGP Flowspec Cooperative Scaling

Single-node mitigation capacity, even when running on highly dense Intel platforms, is ultimately bound by physical interface limits. For example, ten 400G links scale to 4.0 Tbps. To defend against extreme volumetric floods that exceed physical link capacities, the reference architecture integrates an automated upstream escalation mechanism via **BGP Flowspec (RFC 8955)**.

Cooperative Defense Orchestration

The local Intel® Xeon® 6+ processor scrubber continuously monitors its CPU utilization and packet buffer queues.

1. **Trigger Condition:** local load exceeds 80% for a sustained period.

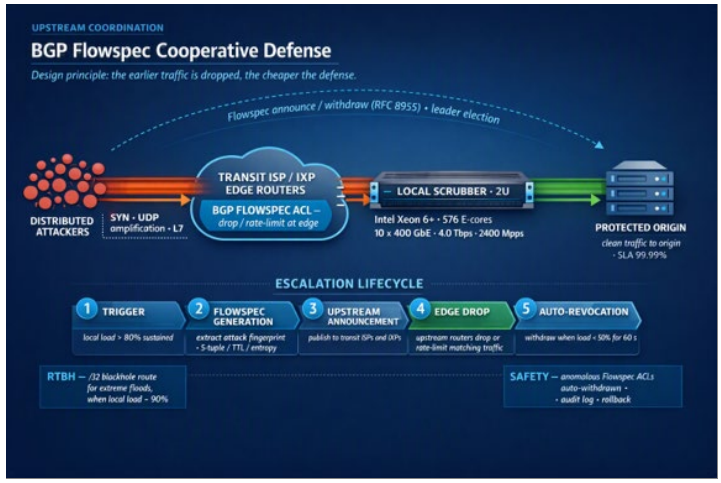


Figure 8. BGP Flowspec Cooperative Defense Topology

2. **Flowspec Generation:** extract the attack fingerprint.
3. **Upstream Announcement:** publish BGP Flowspec to transit ISPs and IXPs.
4. **Edge Drop:** upstream routers drop or rate-limit matching traffic.
5. **Auto-Revocation:** revoke when load drops below 50% for 60 seconds.

9. Empirical Benchmarks & Production Validation

The **Connyon Nova Shield NS-X4000** DDoS mitigation appliance was validated on a dual-socket server configured with Intel® Xeon® 6+ processors, 576 E-cores, 1 TB DDR5-8000 RAM, and 10 x 400G network interfaces.

Key Performance Metrics

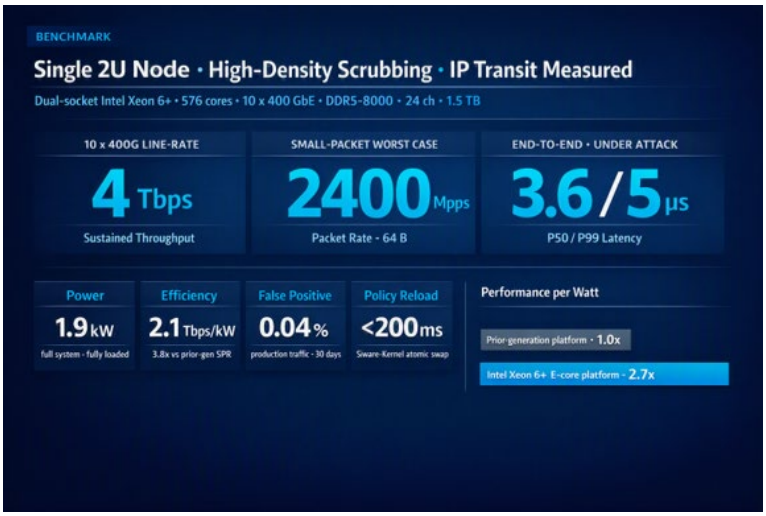


Figure 9. Benchmark KPI Dashboard Defense Topology

Latency Profiles Under Heavy Attack

By mapping packet streams directly to dedicated E-cores and bypassing traditional scheduling layers, the system maintains a stable, low-latency profile even during peak volumetric attacks (see Figure 10).

Power and Cost Efficiency Comparisons

Compared to prior-generation 4th Gen Intel® Xeon® Scalable processor (formerly Sapphire Rapids), Intel® Xeon® 6+ processor delivers substantial efficiency gains for packet scrubbing workloads:

Performance/Watt improvement = 2.7x compared to previous-generation systems.

This efficiency gain is driven by E-core microarchitecture optimizations and the high physical core density of the Intel® Xeon® 6+ processor, allowing operators to shrink their deployment footprint and lower operational costs.

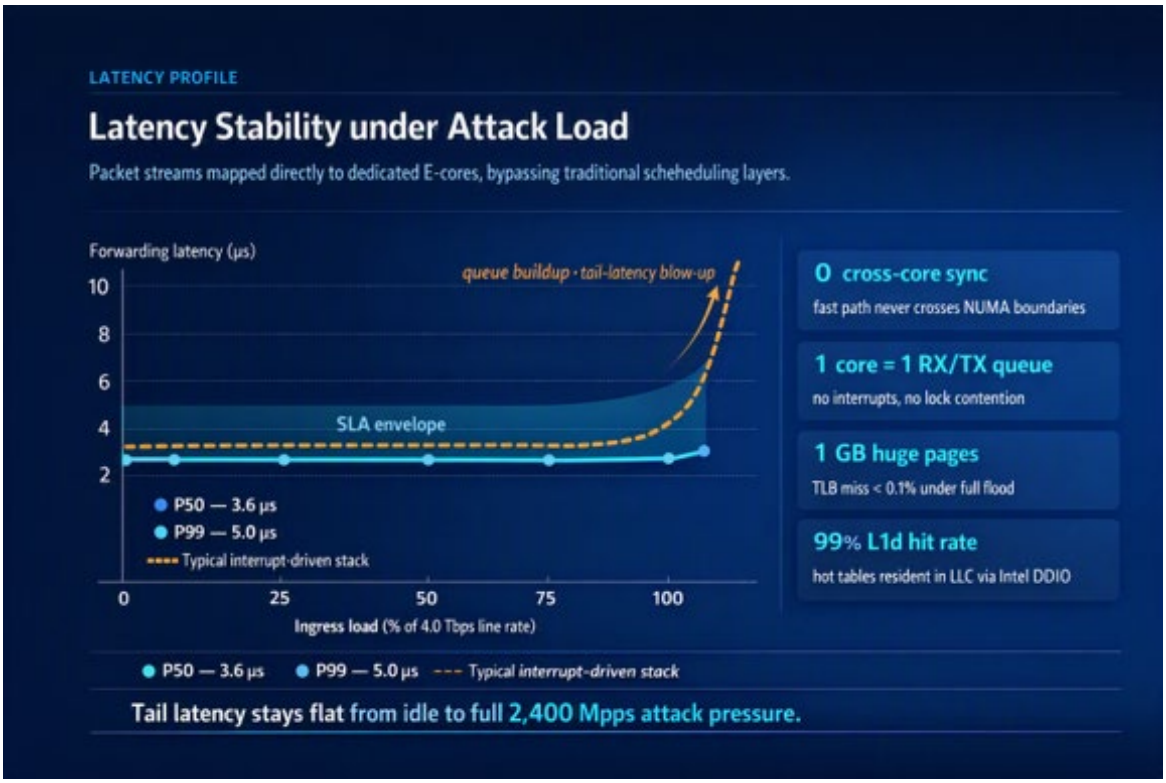


Figure 10. Latency Stability Under Attack Load

Conclusion & Architectural Recommendations

The high-density DDoS reference architecture on the Intel® Xeon® 6+ processors demonstrate that software-defined network appliances can match the performance of legacy hardware-hardened ASICs.

Key Design Takeaways:

1. Adopt high-density E-core architectures for parallel packet workloads.
2. **Consolidate software stacks** into one unified dataplane runtime.
3. Use lockless, pointer-free datastructures with Intel® AVX2 and Intel® DDIO.
4. Keep AI out of the hot path and use it for policy generation.
5. Scale beyond local links with automated BGP Flowspec orchestration.

Performance varies by use, configuration and other factors. Learn more on the Performance Index site. Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

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.

See our complete legal Notices and Disclaimers.

Intel is committed to respecting human rights and avoiding causing or contributing to adverse impacts on human rights. See Intel's Global Human Rights Principles. Intel's products and software are intended only to be used in applications that do not cause or contribute to adverse impacts on human rights.

© 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. Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

1126/PZ/MK/PDF 370430-001US