



RONOG 11 · BUCHAREST · 2026 · PLATINUM SPONSOR

Stopping DDoS at the Door

AI/ML-driven distributed mitigation across a 50+ Tbps backbone

Filter where it enters. Zero attacks across the network.

Alexandru Tölgyi — CTO & Co-founder, ZET.NET

— THE QUESTION

How do you stop **multi-terabit DDoS** across dozens of PoPs on two continents — without your **own backbone** becoming the victim?

...and without buying €1.5–5M appliances for every site. This is the engineering version: the architecture, where AI/ML actually sits, what broke, and what it costs.

Global bandwidth, kept clean

40+

Points of presence · EU & US

3,780+

Direct peering partners

12

Tier-1 transit peers

50+ Tbps

Aggregate backbone capacity

50%+

Backbone headroom, always

AI routing

Best path on latency + stability

+

Nexonhost, acquired ~a year ago — years of battle-tested L4/L7 DDoS mitigation against real attack traffic.
Global bandwidth × mitigation intelligence, merged into one platform.

Hub-and-spoke scrubbing → **your backbone becomes the victim**

01

Backhaul saturation

You pay to transport junk instead of customer traffic.

02

Collateral damage

Other customers on that backbone see loss & jitter — for an attack that isn't theirs.

03

False congestion

Your routing AI reacts to “congestion” that's really junk that should never have entered.

04

Cost

€1.5–5M per 800 Gbps appliance — × a 50+ Tbps attack surface.


When 8–11 Tbps enters through dozens of peers on two continents, **all of it — junk included — must cross your backbone** to reach the scrubber.

— THE ONE DECISION EVERYTHING FOLLOWS FROM

Filter at the source — not the destination

✗ Destination · centralized


Pull everything to one scrubber
⚠ attack aggregate



40 streams **converge** → aggregate on shared links = **single bottleneck.**


✓ Source · distributed

Each PoP filters its own slice



Junk **dies at the PoP where it lands.** Backbone stays clean.

Telemetry → Profiling → ML detection → Distributed filtering




The brain decides. The data plane executes.

— STAGE 3 · WHAT'S GENUINELY DIFFERENT

A small brain that learns your traffic

- Legacy tools trigger on **3 signals** — pps, bps, protocol. That's why they false-positive.
- We seed a small model on known attack vectors + clean traffic (HTTP/S, VPN, gaming, VoIP), then feed it **live traffic** — it learns, **per IP**, what clean looks like.
- Like a child learning grammar: learn the rules, generalize, then **hear “wrong” instantly**.
- **50 Gbps** legit HTTPS? Fine. A **100 kpps** attack on that IP? Caught — **zero-day included**.
- Output = the exact **clean-vs-DDoS difference vector** — precisely what you need to filter.



No GPUs

Hundreds of patterns

Millisecond reaction

Self-generated zero-day signatures

Built in-house on **commodity hardware**

COST

~1/10 per Mpps

Commodity servers + 100G NICs — tens of thousands each, not million-euro boxes. Engineering counted too.

CONTROL

Fleet update in minutes

New vector? Push to the whole fleet in minutes — not hours, and not a firmware ticket that waits weeks.

DISTRIBUTION

No per-site license

Software on standard hardware. Replicate to 100 PoPs at hardware cost, not license cost.

Direct NIC access, kernel bypass, zero-copy, real batching → line-rate on commodity NICs, at hundreds of Mpps.

Dropped **exactly on its vector** — on the interface where it enters

COVERAGE

Every PoP

A filtering cluster installed at each site.

CAPACITY

20+ Tbps

Distributed filtering across the fleet.

SCOPE

All inputs

Including customer ports — blocked on ingress.

PRECISION

0 false-pos

Drop on the vector; clean traffic untouched.

The **zero-attacks** strategy: not “we absorb attacks well” — attack traffic never traverses the network at all. As far as we know, we're the first to operate this way.

Distribution · Visibility · Rollback

DISTRIBUTION

< 1 min convergence

Git-tracked. The control plane pushes to every cluster; drift protection flags any that falls behind.

VISIBILITY

One panel, per PoP

Per PoP, per client, aggregated, 30-day history. You see the attack at the entry point the moment it starts.

ROLLBACK

Blast radius = 1 PoP

Every deploy auto-verified. A misbehaving cluster rolls back itself — mistakes never go fleet-wide.

Correlation is **per-PoP, not per-scrubbing-center**: which peers, which ASNs, which vectors, where.

What actually breaks

HARDWARE

Line-rate is fragile

NUMA & CPU/NIC pinning, hugepages, isolcpus, BIOS/microcode. One mismatch silently costs an order of magnitude.

SOFTWARE

Every bug is an outage

At 100G, state must be consistent enough to work and distributed enough to scale. Schema changes become fleet migrations.

OPERATIONS

Live-fleet upgrades

No maintenance windows — rolling upgrades, canary PoPs, auto-rollback. Plus team discipline: write it down, review, respect prod.

None of it is unsolvable. All of it takes **engineering time** — that's the real cost of in-house, not the server invoice.

A real Nexonhost customer, under real fire

✗ Before · centralized

- Traffic pulled toward the Bucharest scrubber, across continents
- Backbone into Bucharest **saturated at every peak**
- Other customers on that backbone **degraded**
- Absorbed — at real operational + collateral cost

✓ Now · distributed + AI/ML

- ML hands each PoP its filter **vector**; dropped in-line at entry
- Backbone to the destination carries **only clean traffic**
- Backbone utilization during the attack: **like a normal day**
- Collateral damage **zero** · cost per Mpps ~**1/10**

Same event. **Completely different outcome.**

Where this goes next

01

Full fleet coverage

Every PoP on the platform soon.

02

ML ↔ routing coupling

React to fingerprints at millisecond timescale — combine routing shifts with filtering decisions.

03

Open-source the zero-day attack vectors

The observability layer first. No ISP should rebuild it from scratch. Core detection and mitigation IP stays proprietary for now.

Global bandwidth (ZET.NET) × mitigation intelligence (Nexonhost) — **one distributed platform.**

— THE TAKEAWAY

Put the decision **as close to the point of entry as possible** — so neither your backbone nor your customers ever feel the attack.

A small ML brain that outputs the exact filtering vector on commodity hardware at every PoP, makes terabit mitigation **architecturally, operationally, and economically sane** — today.

50+ Tbps where every attack packet is stopped where it enters — not carried into a bottleneck.



Thank you, RONOG.

Questions?

Alexandru Tölgyi — CTO & Co-founder · ZET.NET

AS6204 · peering@zet.net · zet.net · Platinum Sponsor, RONOG 11