© 2026 HyperNext Data Center Limited · hypernxt.com

HyperNext Data Centers

RESEARCH PAPER

HN-RP-010

Autonomous DDoS scrubbing at line rate.

Dual-algorithm filtering, sub-60-second autonomous mitigation, edge absorption, and packet-level enforcement for the AI-native data centre.

A technical description of the HyperNext Scrub architecture and the shape of the DDoS threat in the frontier-AI era.

This is what is Next.

SERIES	HyperNext Research
PAPER	HN-RP-010
ISSUED	03 July 2026
VERSION	1.0
CLASSIFICATION	Public release
CITATION	HyperNext Research, HN-RP-010

CONTENTS

What this paper covers

Methodology and figures are stated openly so other operators can reproduce the analysis on their own facilities. Citation as "HyperNext Research, HN-RP-010" is welcome.

§1 · WHAT CHANGED

Autonomous by necessity

ABSTRACT

The economics of distributed denial-of-service have inverted. Public threat reporting through 2025 records a near-continuous escalation of peak attack volume, from 6.5 Tbps early in the year to 31.4 Tbps in November 2025, the largest yet measured, sustained for roughly 35 seconds. At the same time, 99 per cent of attacks remain under 1 Gbps and most last under a minute. The defensive consequence is unambiguous: an attack now begins and ends faster than a human operator can read the first alert, so mitigation must be autonomous. This paper describes HyperNext Scrub, a sovereign scrubbing service built on four mechanisms. A **dual-algorithm filter** combines a deterministic volumetric stage with an adaptive machine-learning stage. An **autonomous mitigation pipeline** detects, classifies, diverts, scrubs, and returns clean traffic in under 60 seconds without a human in the loop. An **absorption model** carries the everyday band on owned edge capacity and reaches on-demand burst capacity for the rare multi-terabit tail. A **line-rate enforcement datapath** drops malicious packets at the scrubbing core and re-injects clean traffic to origin. Every packet is inspected inside India. The paper closes with the threat that motivated it: the arrival of frontier AI capable of discovering and weaponising novel attack vectors autonomously, and the defensive posture that answers it.

On present industry practice the defender operates two orders of magnitude slower than the attacker. That gap is the subject of this paper.

The response window has collapsed

For most of the history of DDoS mitigation, the operating assumption was that a person would be in the loop. An analyst watched a console, recognised an attack, and applied a countermeasure. That assumption held while attacks lasted minutes to hours and grew at a pace a team could track. It no longer holds. The median volumetric flood in 2025 lasted between 35 and 45 seconds. The largest single event on public record lasted about 35 seconds. An attack that peaks and subsides inside a minute is over before an alert is triaged. The only defence that operates on that timescale is one that requires no human decision at the moment of impact.

This is the same structural shift that HyperNext Research examined for the banking system in HN-RP-009: the interval between the appearance of a threat and its weaponisation has collapsed from months to hours, and soon to minutes. For DDoS the collapse is already complete. HyperNext Scrub is designed around it.

Four mechanisms, one design principle

- **Dual-algorithm filtering.** A fast deterministic stage handles the volumetric majority; an adaptive learning stage adjudicates the ambiguous, low-and-slow, and application-layer minority. Neither alone is sufficient.
- **Autonomous mitigation.** Detection to clean-return in under 60 seconds, with the operator able to take manual control at any point rather than being required to.
- **Attack absorption.** Owned edge capacity carries the everyday band; on-demand burst capacity handles the rare multi-terabit tail, and anycast splits a distributed attack across sites before any filtering begins.
- **Line-rate enforcement.** A packet-level datapath that drops matched malicious traffic at the scrubbing core and re-injects clean traffic to origin, measured in packets per second, not only bits.

The design principle underneath all four is that the defensive loop must run faster than the offensive loop, and must keep running when no operator is watching.

§2 · THE THREAT

DDoS by the numbers

This section states the magnitude of the threat with reference to public data. Figures are drawn from published 2025 to 2026 threat reporting and are cited in section 10.

31.4 Tbps
LARGEST ATTACK ON RECORD · NOV 2025

99%
OF ATTACKS UNDER 1 GBPS

35–45 s
MEDIAN ATTACK DURATION

The escalation of peak volume

Peak attack volume climbed through 2025 in a series of record-breaking events, each surpassing the last within weeks. The trajectory matters more than any single figure: it shows a capability curve, not a set of outliers.

PERIOD	PEAK VOLUME	NOTE
Early 2025	6.5 Tbps	Then a record; UDP flood.
Mid 2025	7.3 Tbps	Multi-vector.
Q3 2025	11.5 Tbps	Reflection-heavy.
Q3 2025	22.2 Tbps	Short-duration burst.
Nov 2025	**31.4 Tbps**	Largest on record; about 35 seconds; roughly a 700 per cent increase on the 2024 peak.

The distribution, and who absorbs it

The volume records draw attention, but the distribution is where mitigation economics are decided. The overwhelming majority of attacks are small and short. A vanishingly small fraction are the multi-terabit giants. A design that provisions owned capacity for the giant is uneconomic; a design that cannot reach the giant is unsafe. The absorption model in section 6 resolves this by owning the band that carries almost every attack and renting the tail.

ATTACK BAND	SHARE OF EVENTS	ABSORBED BY
Under 1 Gbps	~99%	Owned edge, near-zero marginal cost
1 Gbps to 1 Tbps	~0.95%	Owned edge, with cloud burst above threshold
Over 1 Tbps	~5 in 100,000	Cloud burst pool

Vector mix and target concentration

Volumetric UDP floods and reflection or amplification remain the largest share by volume. Two shifts define the current period. Carpet-bombing spreads a flood across an entire prefix rather than a single address, defeating per-host thresholds. Application-layer and API abuse, measured in requests per second rather than bits, targets the expensive parts of a service with traffic that looks legitimate one request at a time. The financial sector is among the most heavily targeted categories, and India is among the most exposed regions. Attacks now routinely combine several vectors in one event, so a single-technique defence fails by construction.

> **WHY THIS FORCES AUTONOMY**
>
> A 35-second attack, spread across a prefix, combining three vectors, is not a problem a human operator can solve inside the attack window. It is a problem a machine solves in the first second and reports afterwards. The rest of this paper describes that machine.

§3 · ARCHITECTURE

The scrubbing centre

HyperNext Scrub is a scrubbing centre operated inside a HyperNext data centre, advertising protected prefixes by anycast and inspecting every packet on Indian soil before returning clean traffic to origin.

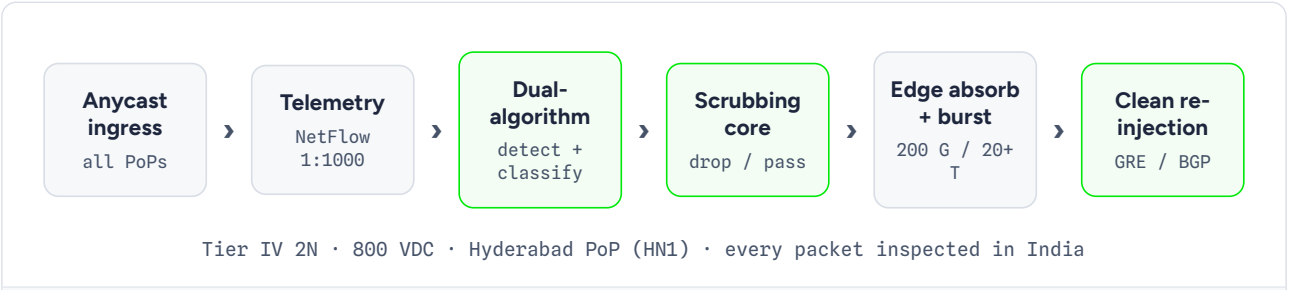

Figure 1. The scrubbing datapath. Anycast ingress feeds telemetry and the dual-algorithm detection and classification stage; the scrubbing core drops, challenges, or passes; the absorption layer carries owned-edge volume and bursts the tail; clean traffic is re-injected to origin.

The datapath, end to end

- **Anycast ingress.** A single scrubbing prefix is advertised from every point of presence. A distributed attack is split across sites by the routing system itself before any filtering begins, which is the first and cheapest layer of absorption.
- **Telemetry.** Sampled flow records (NetFlow and sFlow, typically 1:1000) and packet sampling feed the detection stage continuously, so the baseline is always current and deviation is visible within one sampling interval.
- **Dual-algorithm detection and classification.** The deterministic and adaptive stages described in section 4 produce a per-flow verdict.
- **Scrubbing core.** The enforcement datapath of section 7 drops, challenges, or passes each flow at line rate.
- **Absorption.** Owned edge capacity carries the everyday band; the cloud burst pool is engaged only above the divert threshold, as in section 6.
- **Clean re-injection.** Scrubbed traffic is returned to the tenant origin over GRE or a dedicated BGP path. The tenant sees clean traffic and a full record of what was removed.

The facility underneath

The service inherits the properties of the data centre it runs in: Tier IV 2N design, an 800 VDC power architecture, and the HyperNext operating posture. Sovereignty is a design property, not a marketing claim. Traffic is scrubbed inside India, telemetry is retained under Indian jurisdiction, and operational control rests with the operator. The console in Figure 2 is the operator's working view of a live event.

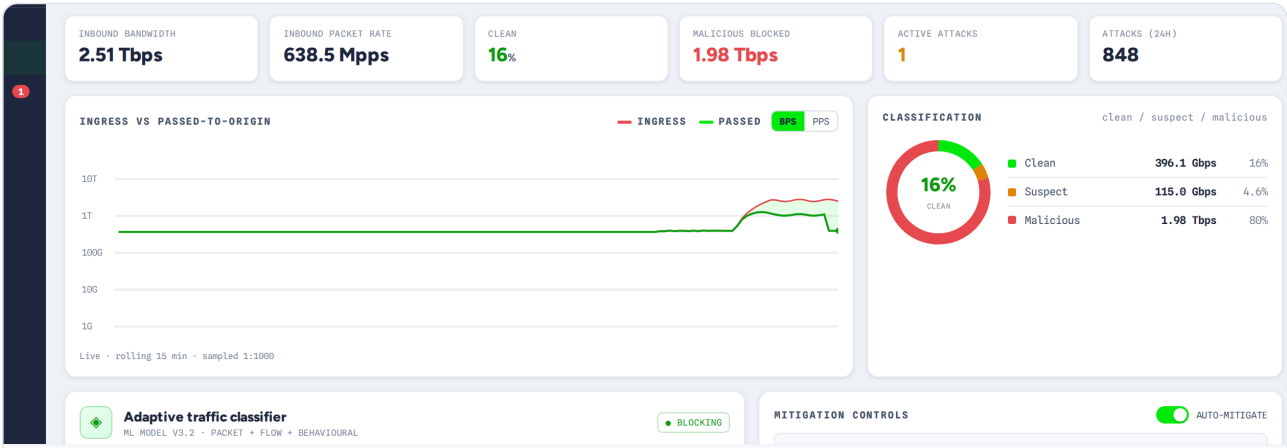

Figure 2. The HyperNext Scrub operator console during a 2.4 Tbps volumetric event on a protected object. Inbound bandwidth and packet rate rise, the classification mix shifts to malicious, and the passed-to-origin line stays flat as the attack is absorbed and dropped.

§4 · DETECTION

The dual-algorithm filter

No single algorithm covers the DDoS problem. Volumetric floods demand a fast, stateless, deterministic filter. Low-and-slow and application-layer attacks demand a model that understands behaviour. HyperNext Scrub runs both and fuses their verdicts.

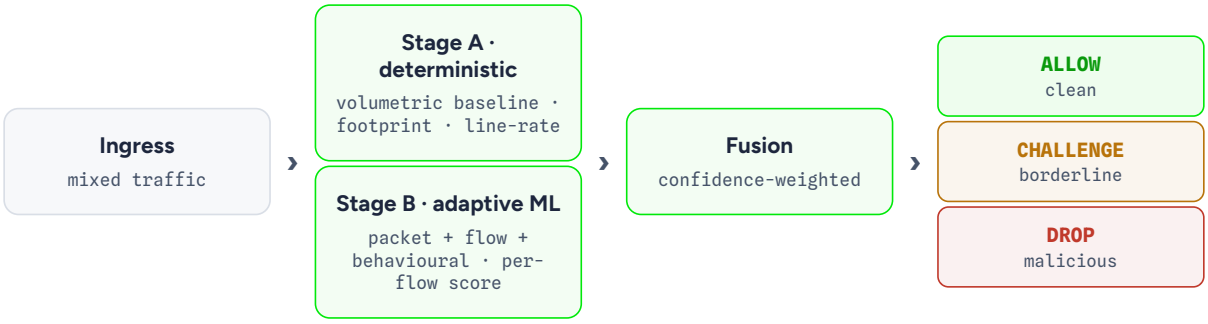

Figure 3. The dual-algorithm filter. Stage A is deterministic and line-rate; Stage B is an adaptive behavioural model. A confidence-weighted fusion step combines their outputs into a per-flow verdict of allow, challenge, or drop.

Stage A: the deterministic filter

Stage A maintains a continuously-learned baseline of normal traffic for each protected object across bandwidth, packet rate, protocol mix, and source distribution. When traffic deviates from the baseline, the stage generates a real-time signature, a footprint, that characterises the offending traffic precisely: protocol, source and destination ports, packet-size band, and time-to-live range. This is the same class of behaviour-based detection that mature mitigation appliances perform, generalised across the whole object. Stage A is stateless and runs at line rate, which is what allows it to hold the volumetric majority without a per-connection state table that an attacker could exhaust.

The footprint is the output that makes autonomous enforcement safe: rather than dropping by broad category, the core drops only traffic that matches the generated signature, which keeps collateral loss of legitimate traffic low. Figure 5 shows a footprint as recorded in the console.

Stage B: the adaptive classifier

Stage B is a machine-learning model that scores each flow on packet-level, flow-level, and behavioural features. It produces a continuous score between zero and one and a verdict of allow, challenge, or drop. Stage B is where the attacks that defeat thresholds are caught: application-layer floods that look like ordinary requests, low-and-slow connections that never breach a volumetric limit, and carpet-bombing spread thinly across a prefix. The model reports an anomaly score against its learned baseline, an inference latency in the sub-millisecond range, and a false-positive rate held near two hundredths of a per cent. Figure 4 shows the classifier adjudicating live flows.

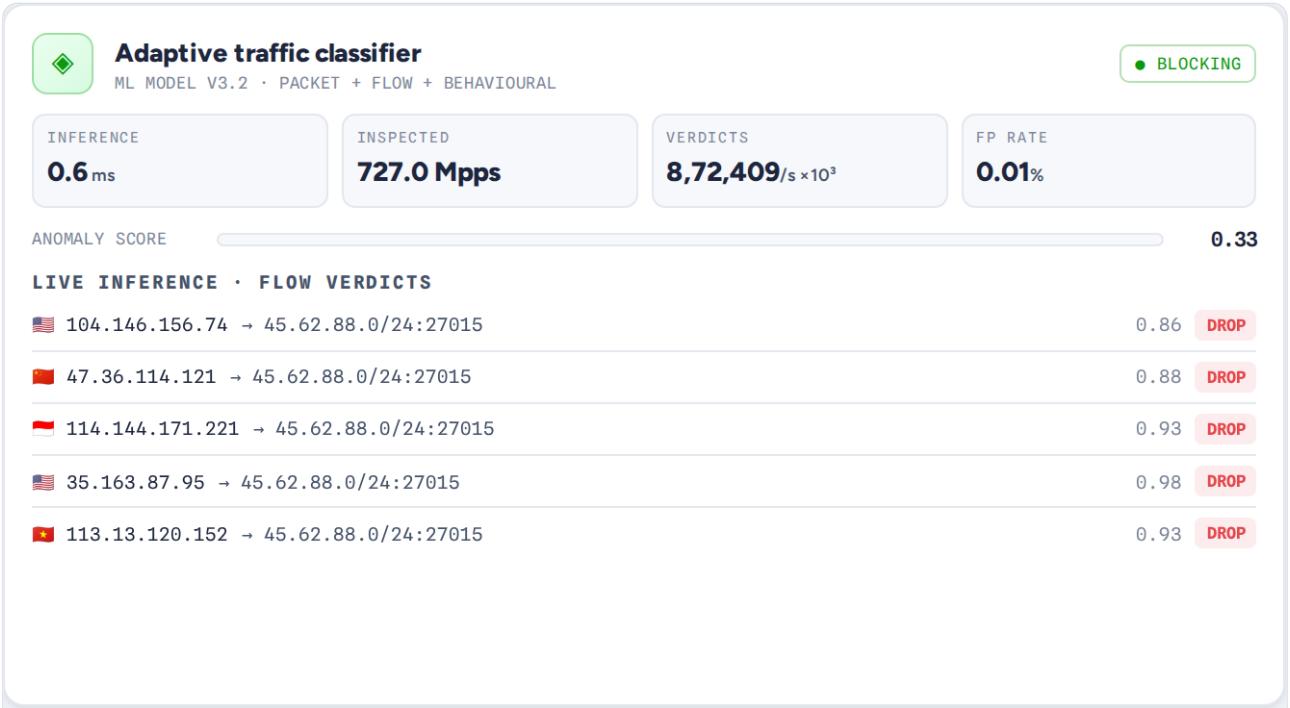

Figure 4. The adaptive classifier. Model status, inspected packet rate, and a live inference stream scoring each source network against a destination and returning allow, challenge, or drop. The anomaly score rises above baseline for the duration of an attack.

Fusion, and the challenge tier

The two stages are complementary, not redundant. Stage A gates volume: if traffic matches a high-confidence volumetric footprint it is dropped without waiting for the model. Stage B adjudicates ambiguity: where volume is normal but behaviour is not, the model decides. The fusion step weighs the two by confidence and applies a middle path for uncertain cases. Rather than a binary allow or drop, borderline traffic is issued a challenge, an out-of-band or protocol-level test that a legitimate client passes and an attack tool does not. The challenge tier is what keeps the false-positive rate low while the drop rate stays aggressive.

Figure 5. A real-time footprint (auto-generated signature) for a volumetric event, as recorded in the console: protocol, source ports, packet-size band, and time-to-live range. The core drops only traffic matching the signature.

§5 · MITIGATION

Autonomous mitigation

Mitigation runs as a pipeline with no human in the loop by default. The operator supervises and can intervene, but the system does not wait for a decision to act.

Detect, classify, divert, scrub, return

The pipeline has five stages. **Detect** registers a baseline deviation. **Classify** confirms the anomaly and generates the footprint. **Divert** announces the protected prefix from the scrubbing core so that traffic is drawn through the datapath. **Scrub** applies countermeasures and drops matched traffic. **Return** re-injects clean traffic to origin. On owned-edge events the whole cycle completes in under ten seconds; the service objective is under sixty seconds including the divert propagation for events that require it. Throughout, the passed-to-origin path is protected, so the tenant experiences continuity rather than an outage followed by a recovery.

Countermeasures, selected by attack class

- **Rate-limit** for floods that are high-volume but not spoofed.
- **SYN-cookie and first-packet drop** for state-exhaustion and spoofed-source floods.
- **Behavioural challenge** for application-layer and API abuse.
- **FlowSpec** announced to upstream transit for volumetric floods that are best stopped before they reach the edge.
- **Source-network blocks** for concentrated malicious origins.
- **Cloud burst** for the multi-terabit tail, as in section 6.

Autonomous by default, operator-controllable at will

Auto-mitigation is on by default. When it is on, the pipeline selects and applies the countermeasures appropriate to the classified attack and reports each action to an audit trail. When an operator turns it off, the attack is not mitigated until the operator acts: each countermeasure is applied by hand, coverage accrues against a threshold, and mitigation completes when coverage crosses it. This is a deliberate design choice. The safe default is autonomy, because the attack is faster than a person; the available override is manual control, because accountability for a mitigation decision must rest with a named operator when circumstances require it. Every action, automatic or manual, is written to an audit trail with an actor and a timestamp.

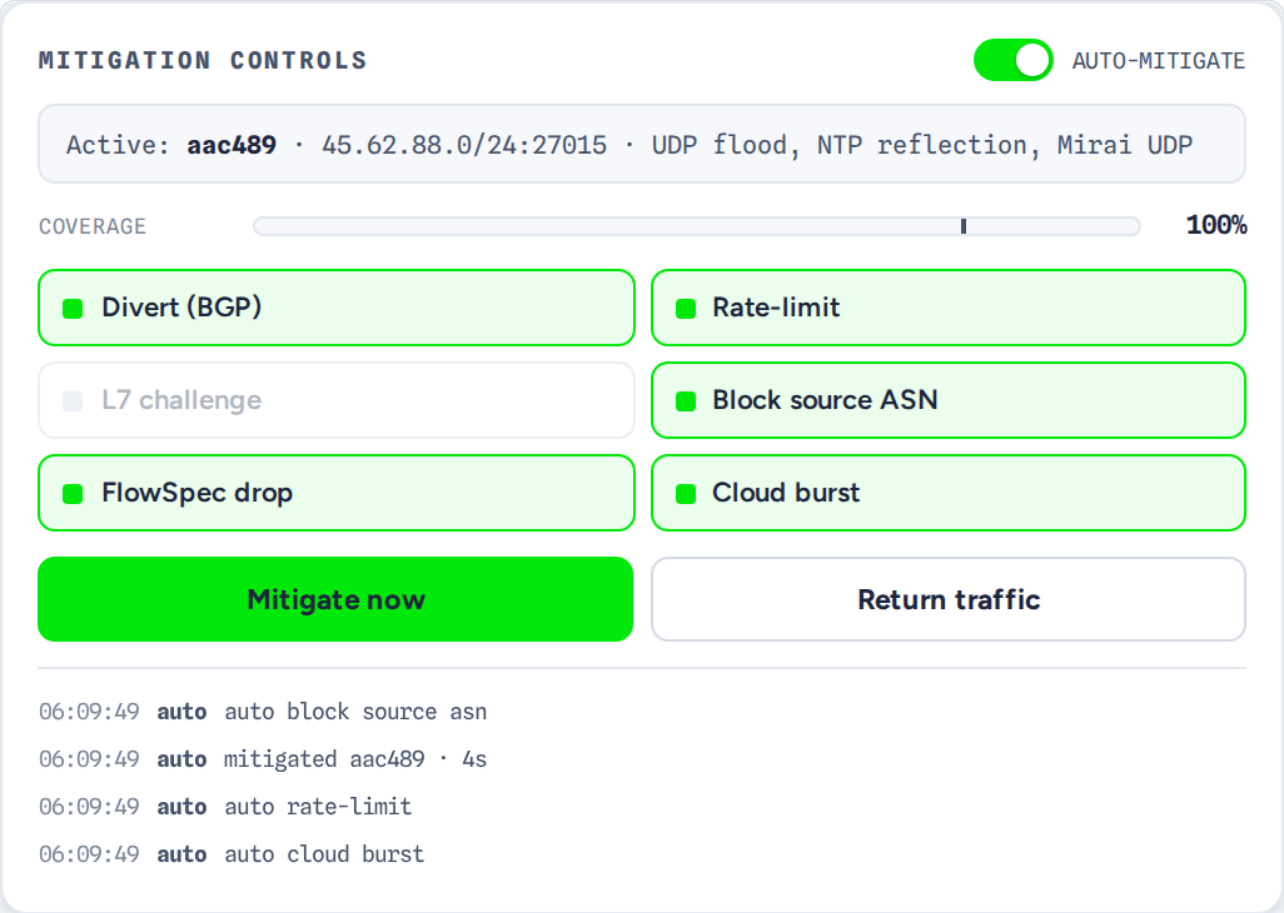

Figure 6. Mitigation controls. The auto-mitigate toggle, the applied countermeasures against a coverage threshold, and the audit trail recording each action as operator or automatic with a timestamp.

Figure 7. The mitigation timeline for a single event: baseline deviation, anomaly confirmation, footprint generation, BGP divert, countermeasures applied, mitigation, and clean return. The record is retained for forensics and for the customer.

§6 · ABSORPTION

Owned edge and cloud burst

The absorption model is the engineering core of the design. It provisions owned capacity for the band that carries almost every attack, and reaches on-demand capacity for the rare band that carries almost none.

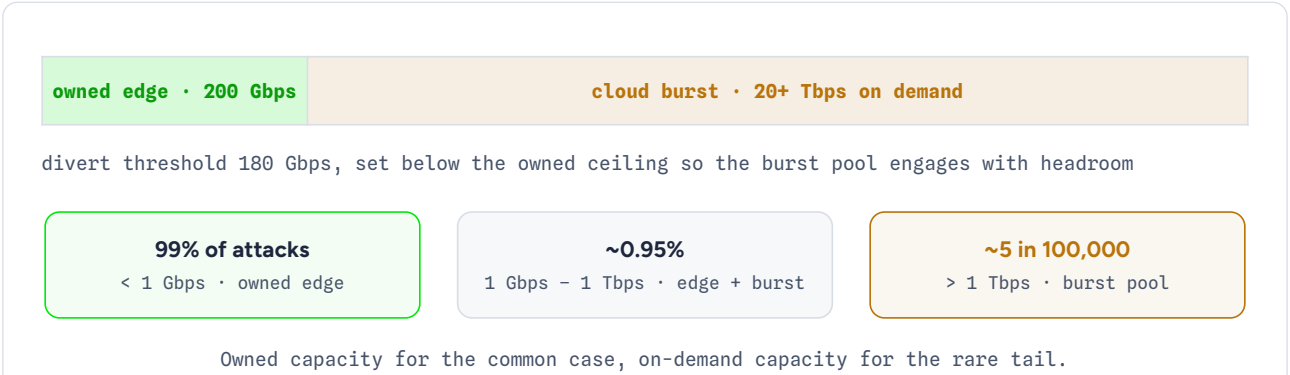

Figure 8. The absorption model. Owned edge capacity of 200 Gbps per point of presence carries the everyday band; the cloud burst pool of 20+ Tbps is engaged only above the 180 Gbps divert threshold. Almost every attack falls inside the owned band.

Owned edge for the everyday band

Each point of presence carries 200 Gbps of owned mitigation capacity. Because 99 per cent of attacks are under 1 Gbps and the great majority of the remainder are well under a terabit, owned capacity absorbs almost every attack without engaging the burst pool. Anycast makes this stronger: a distributed attack is split across every point of presence by routing before any capacity is consumed at a single site. The divert threshold is set at 180 Gbps, below the owned ceiling, so the system engages the burst pool with headroom rather than at the point of saturation.

Cloud burst for the rare tail

Above the threshold, the cloud burst pool of more than 20 Tbps is engaged on demand. The tenant origin never sees the tail; it is absorbed and scrubbed before re-injection. The burst pool is engaged only in the rare events that exceed the divert threshold. Figure 2 and the passed-to-origin line in the console show the effect: inbound volume rises into the terabits while the line delivered to origin stays flat.

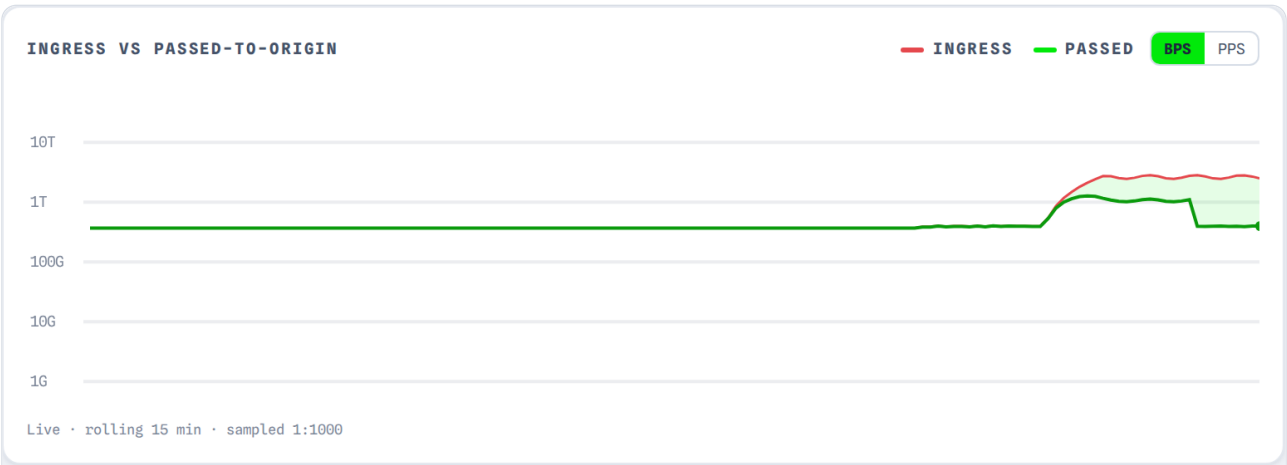

Figure 9. Ingress against passed-to-origin during an event. The upper line is total inbound volume; the lower line is clean traffic delivered to the tenant. The gap is what the absorption and enforcement layers remove.

§7 · ENFORCEMENT

The packet destroyer

The scrubbing core is the enforcement datapath. Its job is narrow and absolute: remove malicious packets at line rate and pass clean packets untouched. It is measured in packets per second, not only bits.

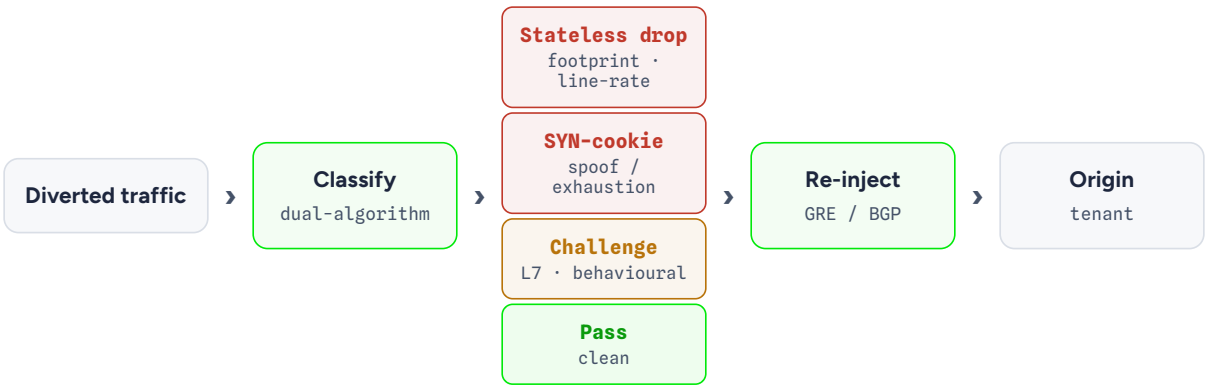

Figure 10. The enforcement datapath. Diverted traffic is classified, then each flow is dropped, challenged, or passed. Volumetric floods are dropped statelessly against the footprint; spoofed and state-exhaustion floods are met with SYN-cookies and first-packet drops; clean traffic is re-injected to origin.

Why packets, not just bits

Bandwidth is the headline number, but packet rate is often the harder problem. A flood of small packets can exhaust the forwarding capacity of a device long before it saturates the link, which is why the console reports packet rate alongside bandwidth and why the enforcement path is built to drop at packets-per-second scale. The largest events on record combine both: terabits per second of bandwidth and hundreds of millions of packets per second. The datapath is designed for the packet number.

How malicious packets are removed

- **Stateless footprint drop.** Traffic matching a high-confidence Stage A signature is dropped without per-connection state, at hardware line rate. This is the mechanism that holds the volumetric majority.
- **SYN-cookies and first-packet drop.** Spoofed-source and state-exhaustion floods are defeated by proving a client can complete a handshake before any state is committed.
- **Behavioural challenge.** Application-layer flows the model marks as borderline are issued a challenge; failure results in a drop, success in a pass.

- **FlowSpec upstream.** For the largest volumetric floods, a FlowSpec rule is announced to upstream transit so the flood is dropped in the carrier network before it reaches the edge, conserving owned capacity.
- **Hardware enforcement.** The highest-rate drops are performed in hardware at the edge, where an access-control decision costs a single forwarding cycle.

Clean re-injection

What survives the datapath is clean traffic, which is re-injected to the tenant origin over GRE or a dedicated BGP path. The tenant receives only legitimate traffic and, separately, a complete record of what was removed and why. Figure 11 shows that record at the level of individual source addresses.

TOP TALKERS — source IPs · NetFlow 1:1000 sampled

SOURCE IP	ASN	COUNTRY	BPS	PPS	PACKETS	SHARE
104.216.171.63	AS14061	United States	506.4 Gbps	141.4 Mpps	1.27B	20.8%
35.149.15.170	AS15169	United States	379.7 Gbps	106.0 Mpps	0.95B	15.6%
114.40.86.180	AS7713	Indonesia	347.4 Gbps	97.0 Mpps	0.87B	14.3%
113.222.164.5	AS45899	Vietnam	259.5 Gbps	72.5 Mpps	0.65B	10.7%
104.146.243.125	AS14061	United States	187.0 Gbps	52.2 Mpps	0.47B	7.7%
35.122.6.126	AS15169	United States	144.1 Gbps	40.2 Mpps	0.36B	5.9%
47.223.186.1	AS45102	China	136.1 Gbps	38.0 Mpps	0.34B	5.6%
113.223.98.81	AS45899	Vietnam	129.4 Gbps	36.1 Mpps	0.33B	5.3%
47.205.105.239	AS45102	China	107.5 Gbps	30.0 Mpps	0.27B	4.4%
114.198.25.90	AS7713	Indonesia	68.1 Gbps	19.0 Mpps	0.17B	2.8%
113.83.227.176	AS45899	Vietnam	63.8 Gbps	17.8 Mpps	0.16B	2.6%
104.207.180.170	AS14061	United States	51.9 Gbps	14.5 Mpps	0.13B	2.1%
114.197.195.229	AS7713	Indonesia	26.6 Gbps	7.4 Mpps	0.07B	1.1%
35.104.142.135	AS15169	United States	25.3 Gbps	7.1 Mpps	0.06B	1.0%

Figure 11. Top-talker forensics for a single event: the individual source addresses behind the flood, with per-source bandwidth, packet rate, packet count, and share. This is the evidence a tenant and a regulator receive.

§8 · OPERATIONS

Transparency and evidence

A mitigation service is only as trustworthy as the evidence it produces. HyperNext Scrub is built so that both the operator and the protected tenant can see exactly what happened, in full detail, during and after an event.

Two views, one record

The operator works in the console: a dashboard of live telemetry, a dense and filterable attack log, and a drill-down for each event. The tenant sees a scoped customer portal with the same underlying record for its own objects: where its traffic originates, how clean it is, what attacked it, where the attack came from, and how it was stopped. Nothing material is hidden from the customer. The same event produces the same facts on both sides.

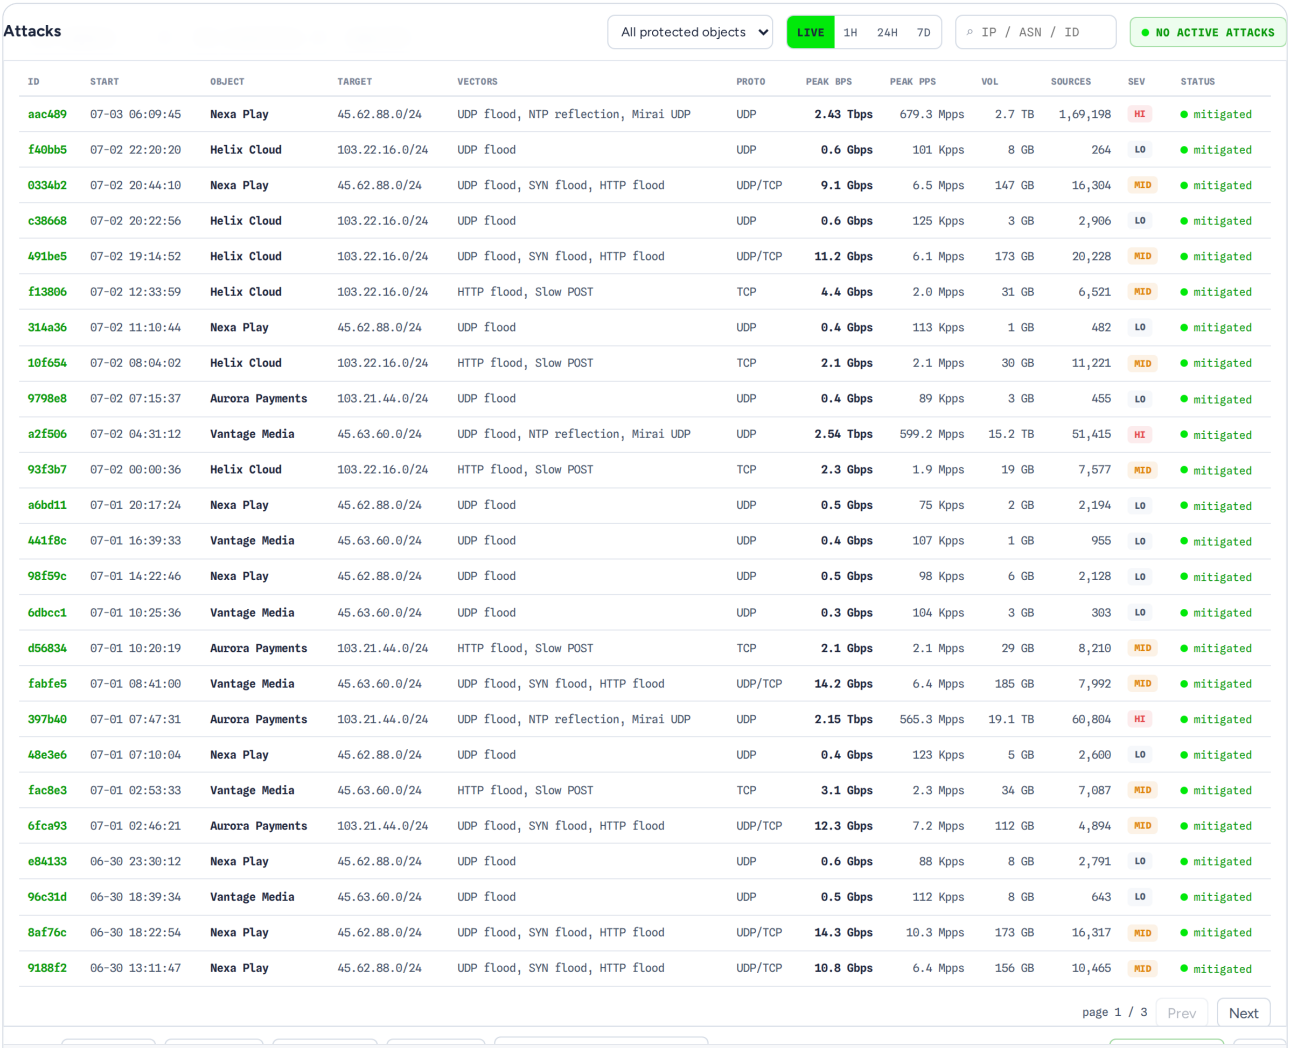

Figure 12. The attack log. Each event carries an identifier, target prefix, vectors, protocol, peak bandwidth and packet rate, unique-source count, duration, mitigation action, and status. Any row opens the forensic record of Figures 5, 7, and 11.

Forensics and export

Each event drills down to a complete forensic record: the auto-generated footprint, the mitigation timeline, the top-talker source addresses, per-attack statistics, and the geographic and network origin of the traffic. The record exports as machine-readable data, as a packet capture slice, and as a report. Reportable incidents are classified as such and filed within the six-hour window that Indian regulation requires, and the same record supports the software-bill-of-materials and audit posture that HN-RP-009 argued the financial sector now needs.

> **TRANSPARENCY AS A CONTROL**
>
> Full source visibility is not a courtesy to the customer. It is a control. A tenant that can see every source network and every dropped packet can reconcile the mitigation against its own logs, satisfy its own regulator, and hold the operator to account. The record is the product as much as the scrubbing is.

§9 · FUTURE SCOPE

AI and the next generation

The threat that motivates this architecture is not the DDoS of 2025. It is the DDoS that frontier AI will make possible, and the defence is an AI-native one.

What AI changes on the offensive side

HN-RP-009 documented the arrival of AI models able to discover previously-unknown software vulnerabilities autonomously and at scale. The same capability, turned toward denial-of-service, produces a class of attack the current generation of defences was not designed for.

- **Autonomous vector discovery.** A model that can read a protocol implementation can find new amplification and reflection vectors without human research, collapsing the lead time defenders currently get from public disclosure to nothing.
- **Adaptive botnets.** An agent coordinating a botnet can change vector, source distribution, and packet shape in response to the defence it observes, within the same attack, faster than a signature can be written by hand.
- **Adversarial evasion.** A model that understands a behavioural classifier can shape traffic to sit just inside the allow boundary, defeating a static model that is not itself learning.
- **Polymorphic application-layer abuse.** Generated request traffic that is individually indistinguishable from legitimate use, varied continuously so that no two requests share a signature.
- **Orchestrated multi-vector campaigns.** A planner that sequences volumetric, protocol, and application-layer phases to exhaust each layer of a defence in turn.

What AI must do on the defensive side

A static defence loses to an adaptive attacker by construction. The defensive roadmap is therefore to make every layer of the system learn, predict, and act autonomously, and to keep the defensive loop faster than the offensive one.

- **Online-learning classifiers.** The adaptive stage retrains continuously on live traffic rather than on a fixed model, so that an evasion that works in one minute is unlearned in the next.
- **Predictive pre-positioning.** Forecasting attack onset from early telemetry and pre-announcing divert paths, so the first packets of a flood arrive at a core already prepared for them.
- **Generative red-teaming.** Using the same generative capability defensively, to synthesise next-generation attacks in a simulator and train the classifier against them before they are seen in the wild.
-

Reinforcement-learning mitigation policy. Learning which countermeasure sequence minimises collateral loss and time-to-mitigate for each attack class, rather than applying a fixed playbook.

- **Graph and behavioural models for carpet-bombing.** Reasoning over a whole prefix and the relationships between sources, rather than per-host thresholds that a spread attack is designed to evade.
- **Federated, sovereign threat intelligence.** Sharing attack structure across points of presence and, under agreement, with peer national capabilities, so that an attack seen once is known everywhere, while the data stays under Indian jurisdiction.

A sovereign defensive posture

HN-RP-009 argued that India needs a sovereign defensive AI capability for the financial sector, near-term through partnership access to frontier models and medium-term through indigenous capability. The DDoS case is the same argument in a narrower domain. The defensive AI that answers AI-driven denial-of-service should run on Indian soil, under Indian control, on infrastructure that keeps traffic and telemetry within Indian jurisdiction. HyperNext Scrub is built to be exactly that substrate: an autonomous, learning, sovereign scrubbing capability that a bank, a payment operator, or a public service can stand behind.

> **THE ARMS RACE, STATED PLAINLY**
>
> When both offence and defence are driven by AI, the contest is decided by whose loop is faster and whose model learns first. A defence that requires a human at the moment of impact has already lost. The whole of this architecture is an argument that the defensive loop must be autonomous, must learn continuously, and must run where the defender, not the attacker, sets the rules.

§10 · CONCLUSION

Where this goes

DDoS mitigation has crossed a threshold. The attack is now faster than the operator, the largest events are measured in tens of terabits, and the next generation of both attack and defence will be driven by AI. A mitigation service that is not autonomous, not learning, and not sovereign is a service calibrated to the last decade.

HyperNext Scrub is the answer this paper describes: a dual-algorithm filter that pairs a deterministic volumetric stage with an adaptive behavioural model; an autonomous pipeline that detects, classifies, diverts, scrubs, and returns clean traffic in under a minute without a human in the loop; an absorption model that owns the everyday band and rents the rare tail; and a line-rate enforcement datapath that removes malicious packets and re-injects clean traffic to origin. Every packet is inspected inside India, and every event produces a complete forensic record shared with the protected tenant.

The absorption model in section 6 keeps owned capacity sized to the common case and reaches on-demand capacity for the rare tail. The case for building this now, rather than later, is set out in section 9: the AI that will drive the next generation of attacks is already here, and the only defence that keeps pace is one that is itself autonomous and learning.

References

Threat data and distribution figures are drawn from published 2025 to 2026 DDoS threat reporting, including Cloudflare DDoS Threat Reports (2025 quarterly and the November 2025 record disclosure), NETSCOUT Threat Intelligence Reports (2025), and Radware and vendor threat advisories (2025 to 2026). Mitigation techniques reference the established literature and vendor documentation: behaviour-based detection and real-time signature generation (Radware DefensePro, NETSCOUT Arbor Sightline and TMS), A10 Thunder TPS, BGP FlowSpec (RFC 8955, updating RFC 5575), and remotely-triggered black-hole and SYN-cookie techniques. Frontier-AI context references HyperNext Research, HN-RP-009, Frontier AI and the Indian Banking System (2026), and the public evaluations cited therein. Indian regulatory context references CERT-In incident-reporting directions and SBOM technical guidance (2022 to 2025).

About this paper

This paper is published by HyperNext Research. It describes an operating architecture and a working prototype console; the figures are captured from that console. Methodology is stated openly so other operators can reproduce the analysis on their own facilities. Correspondence on the methods, the figures, and the conclusions is welcome at hello@hypernxt.com.

HYPERNEXT RESEARCH · HN-RP-010

This is what is next.

Detect in the first second. Mitigate without a human. Keep the record.

We publish engineering papers because the Indian conversation about AI infrastructure needs more substance than marketing provides. Methodology is stated openly so other operators can reproduce the analysis on their own facilities.

REGISTERED OFFICE

HyperNext Data Center Limited

Mindspace, Jubilee Hills,
Hyderabad 500033, India

ONLINE

hypernxt.com/research

hello@hypernxt.com · +91 99784 23333

DOCUMENT

HN-RP-010

Frontier DDoS · Version 1.0 · 03 July 2026

© 2026 HyperNext Data Center Limited. All rights reserved.

hypernxt.com · Governed by policy, run with integrity.

HyperNext Data Centers

Designed by QUANTLOOP Technologies Pvt. Ltd.