corero
[NETWORK SECURITY]

WHITEPAPER

It's Not an Incident. It's Tuesday.

12.3 attacks a day — the sharpest annual rise in a decade. And most of them are too small for your defenses to notice. Here's what that means for how you build — and defend — your network.

12.3
Average DDoS attacks per day across monitored networks

82%
Of attacks under 1Gbps — below most detection thresholds

262%
Year-over-year increase in peak attack sizes

30–60s
How often sophisticated attackers cycle attack vectors

Corero Network Security
www.corero.com

March 2026

THE OPERATING ENVIRONMENT

You Are Under Attack Right Now.

Not probably. Not at risk of. Right now, at this moment, some fraction of the traffic hitting your network is hostile — probing for weaknesses, measuring your response time, testing whether your mitigation threshold kicks in before or after service degrades.

In our monitored networks, organizations faced an average of 12.3 DDoS attacks per day in 2025 — one attack every 117 minutes, around the clock, seven days a week, including holidays. The sharpest annual rise in a decade. The direction has not changed once in over 10 years.

12.3/day

Average DDoS attacks across our monitored networks — the **sharpest annual rise in a decade.** One attack every 117 minutes, around the clock.

This is not a threat landscape. It is the operating environment. The organizations that understand the difference are the ones building defense architectures that do not depend on someone noticing an attack is underway.

The real question is: how close to your infrastructure does your first line of defense need to be, and how fast does it need to respond, to stop what is actually hitting you right now? The answer changes everything about how you think about on-premises, cloud, and hybrid protection.

THE MYTH

DDoS attacks are large, obvious, and detectable. When you are under attack, you will know.

THE REALITY

More than 82% of observed attacks were under 1Gbps. They did not saturate links or trigger volumetric thresholds. They created latency, packet loss, and errors — the kind of noise most NOC teams write off as an unreliable ISP day.

"They were attacks — specifically designed to operate below the threshold at which most defenses engage, because the attackers know where that threshold is."

ATTACK PROFILES

The Attack You Won't See Coming

There is a persistent and comfortable myth in enterprise security: that DDoS attacks are large, obvious, and detectable. In our telemetry, more than 82% of observed attacks were under 1Gbps. They did not saturate links or trigger volumetric thresholds. They created latency, packet loss, and errors — the kind of noise most NOC teams write off as an unreliable ISP day.

82% Of observed attacks were under 1Gbps — **below the threshold most defenses are calibrated to catch.** Designed to probe and stay invisible while measuring your response.

ATTACK SIZE DISTRIBUTION — OBSERVED SHIFT OVER 5 YEARS

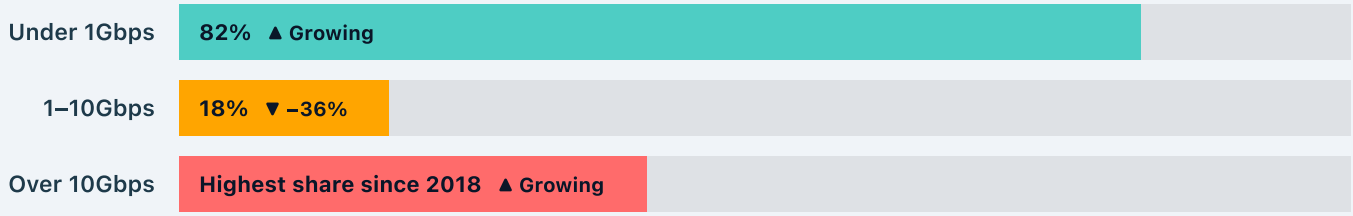

Attackers optimized the middle out of existence. The 1–10Gbps range was big enough to get flagged, but not big enough to guarantee disruption — poor return on investment. So they polarized: go small and stay invisible, or go large and leave nothing standing. If your detection posture is calibrated for mid-tier events, you are well-resourced for the attack profile that is disappearing — and underequipped for both ends of the spectrum that are growing.

CHAINED VECTOR ATTACKS

The Attack That Looks Like Noise

We have documented a growing pattern in sophisticated DDoS campaigns: coordinated sequences of different attack types — SYN floods, DNS amplification, UDP floods, HTTP mimics — cycling every 30 to 60 seconds. Not because any single vector is particularly powerful on its own. Because of what happens at the transitions.

30–60s How frequently sophisticated attackers cycle between attack vectors. Every switch forces your defense to pause, reclassify, and re-engage — creating a **permanent window of unmitigated traffic by design.**

This Is Not Noise. It's Engineering.

Every time an attacker switches tactics, your defense has to pause. It must analyze the new traffic profile, reclassify the threat, and initiate the appropriate mitigation response. Even advanced platforms can take 10 to 30 seconds to complete that cycle. With campaigns deploying 50 or more unique attack vector combinations, and the shortest observed attack waves lasting just 6 seconds, that detection-response lag is not an edge case — it is a permanent exploit window.

> ***The defenders best equipped against chained vector campaigns share one characteristic: their mitigation operates continuously and adaptively — not in detection-response cycles. Speed is not a performance metric. It is the entire defense strategy.***

DEPLOYMENT ARCHITECTURE

Where Your Defense Lives Determines What It Can See

Every DDoS protection deployment model makes a fundamental tradeoff: where does traffic get examined, and how long does that take? This has direct consequences for which attacks your defense catches — and which ones reach your infrastructure unchallenged.

CLOUD SCRUBBING

Off-Premises Inspection

Unmatched capacity for volumetric floods. But every packet makes a round trip — introducing latency that cannot be engineered away. For the 82% of sub-1Gbps attacks, it adds overhead without catching anything.

↳ BGP rerouting delay: 1–5 min

ON-PREMISES

First-Line Edge Defense

Traffic examined at the point of entry, at line rate, in-path. Sub-1Gbps probes caught by behavioral analysis. Vector switches handled in milliseconds. Limited only by hardware capacity against terabit-scale floods.

↳ Cannot absorb sustained terabit floods alone

HYBRID

The Point, Not the Compromise

Always-on on-premises mitigation handles daily attack pressure. Cloud capacity absorbs volumetric overflow. On-premises intelligence informs cloud response. Neither model can do this alone.

↳ The architecture the threat landscape demands

AI INFRASTRUCTURE & DDOS DEFENSE

Your AI Investment Changes Your Defense Requirement

There is a new and urgent reason that on-premises defense is moving from best practice to prerequisite — and it has nothing to do with DDoS specifically. Organizations are in the middle of a structural shift in AI infrastructure. After years of moving AI workloads to the cloud, a broad repatriation is underway.

- **Economics:** Cloud inference costs scale with usage in ways that were tolerable at proof-of-concept but unsustainable in production.
- **Sovereignty:** Data sovereignty requirements tightened. Model weights, training data, and proprietary fine-tuning created compliance pressure to keep AI infrastructure inside organizational control.
- **Latency:** Real-time AI inference — fraud detection, clinical decision support, autonomous operational systems — requires response times that cloud-based inference cannot consistently deliver.

<10ms

The response threshold required by real-time AI inference workloads. Cloud-based inference, even from regional endpoints, **cannot consistently meet it for latency-critical applications.** — Equinix, 2025

THE CONVERGENCE POINT

A successful DDoS attack against on-premises inference infrastructure does not cause generic service disruption — it collapses the latency guarantees the entire AI pipeline depends on. A 50ms latency spike from a sub-1Gbps probe attack causes real-time inference to miss response windows entirely. Not intermittently. Systematically, for as long as attack pressure continues.

Cloud scrubbing is architecturally incompatible with protecting sub-millisecond inference. You cannot route traffic to a remote mitigation center and back without introducing the very latency the infrastructure was built to eliminate. Organizations investing in on-premises AI infrastructure are accepting responsibility for on-premises defense — the question is whether that defense operates at the speed the workload demands.

APPLICATION LAYER THREATS

DDoS Is No Longer a Network Problem

It is a business problem, expressed through the network. Organizations that treat it as only a network problem are defending one layer of a three-layer attack surface. Application-layer attacks targeting HTTP endpoints, APIs, login portals, and resource-intensive web application logic are rising sharply. Unlike volumetric attacks that saturate bandwidth, Layer 7 attacks exhaust application resources — CPU cycles, session memory, database connections, API rate limits — with traffic volumes that would not register on a bandwidth graph.

LAYER	WHAT IT REQUIRES
L3 / L4 Network / Transport	Volumetric and protocol attack mitigation at line rate, inline, with zero-touch automation. Always-on first line handling the daily attack pressure that never stops.
L7 Application	Behavioral detection tuned to application-specific traffic patterns. Distinguishing attack traffic from legitimate users without false positives, at scale, in real time.
Intelligence Traffic Visibility	Deep visibility into what is hitting the network, where it is going, and how it compares to established behavioral baselines. Without this, L3/L4 and L7 defenses operate partially blind.

ZERO TRUST CONVERGENCE

Where Zero Trust and DDoS Defense Meet

Zero Trust operates on a principle most security teams know well: no traffic should be trusted by default. Every connection must be verified. Every access decision must be informed by identity and context. That principle requires exactly what advanced DDoS defense requires — comprehensive traffic visibility, accurate real-time classification, and enforcement that acts without introducing latency.

> THE RISK OF TREATING THEM SEPARATELY
>
> Organizations building Zero Trust architectures and those building modern DDoS defenses are solving adjacent problems with overlapping tools. Those who treat them as separate purchasing decisions end up with two sets of tools doing half a job each — at twice the operational overhead.

OPERATIONAL FINDINGS

The Right Architecture Fails When the Team Can't Act

Based on commissioned research by Merrill Research — security and network operations practitioners on real-world DDoS defense experience.

68%

struggle to demonstrate the ROI of DDoS protection to leadership

51%

cite cross-team coordination gaps as a key operational vulnerability

47%

report difficulty adapting existing tools to hybrid environments

>50%

are not confident mitigating advanced attacks without vendor guidance

These are not technology failures. They are operational failures — the kind that happen when the right capabilities exist but are not integrated into workflows that function at the speed and scale of modern attacks.

- **68% cannot justify DDoS investment to leadership** because low-visibility attacks produce low-visibility consequences — which produce under-resourced defenses, which produce better attack results. Attackers benefit from this cycle without ever having planned it.
- **51% cite cross-team coordination as a key vulnerability.** DDoS response requires network operations, security, platform teams, and leadership to act in concert — without pre-authorized response workflows, coordination overhead arrives at exactly the moment speed matters most.

THE OPERATIONAL IMPERATIVE

The organizations that handle attack pressure best have invested as much in operational alignment as in technical capability — clear ownership, pre-authorized response playbooks, and automation that does not require human approval for decisions that must happen in milliseconds. The technology is necessary. It is not sufficient.

EVALUATION FRAMEWORK

Selecting for the Threat That Actually Exists

Most DDoS protection evaluation frameworks were designed for the threat landscape of five years ago — emphasizing capacity, large attacks, and cloud scrubbing as the default. The threat landscape has moved. Evaluation criteria should move with it.

01 **Sub-Threshold Detection**	**ASK** At what attack size does the system begin responding — and what happens to the 82% of attacks that fall below typical volumetric thresholds? If the answer is "those don't cause outages," the evaluator has not yet understood the threat model.
02 **Response Speed**	**ASK** From the first anomalous packet to active mitigation, how long? In a chained vector environment, a 10-second gap is a permanent exploit window. Ask for documented response times against sub-1Gbps, multi-vector campaigns — not just maximum throughput figures.
03 **Multi-Layer Coverage**	**ASK** Can the system maintain effective mitigation across L3, L4, and L7 simultaneously — without separate tools or sequential response workflows? Multi-vector attacks do not pause to let you switch platforms.
04 **Behavioral Detection Depth**	**ASK** Does the system identify attacks based on traffic characteristics and continuously updated threat intelligence, or primarily on volume thresholds? Behavioral anomaly detection combined with access intelligence means the system recognizes what looks wrong and what is known hostile — without waiting to establish a baseline first.
05 **Operational Integration**	**ASK** Does this system reduce cognitive load on the teams operating it — or require dedicated management overhead to function at the level marketed? Research suggests this is where most implementations fail in practice.
06 **Hybrid Coherence**	**ASK** If the architecture spans on-premises and cloud enforcement points, do they operate as a coordinated system — sharing intelligence and handing off between layers — or as two separate products that happen to be sold together?

CONCLUSION

Business Continuity Is The Standard Now

DDoS defense has changed its job title. For most of its history it was a networking function — a way to keep bandwidth available when someone decided to flood a link. That framing is no longer adequate for what DDoS has become: a persistent, automated, multi-vector mechanism for creating business disruption at scale.

Attacks do not primarily target bandwidth anymore. They target availability, performance, user experience, transaction throughput, authentication systems, and the AI inference pipelines that organizations are building their next competitive advantage around. The organizations that have adapted their defense posture to this reality have stopped asking "which deployment model fits our budget?" and started asking "what does uninterrupted availability actually require?"

The answers are architectural:

First-Line Defense at the Edge
Always-on, inline mitigation that catches the daily attack pressure before it reaches infrastructure.

Behavioral Detection
Catches what volumetric thresholds miss — the 82% of attacks operating below the radar.

Application-Layer Awareness
Recognizes attack characteristics regardless of traffic volume or encryption state.

Operational Alignment
Pre-authorized playbooks and automation that let teams act at the speed attacks demand.

This is business continuity thinking. It is also, increasingly, the minimum viable standard for any organization that cannot afford to be offline. The signal is there. The question is whether your defense can see it.

See What Your Defense Is Missing

Send us your current architecture and we'll identify exactly where your gaps are — no sales pitch, just engineering analysis from our team to yours.

TALK TO AN EXPERT

The attacks aren't waiting. *Your defenses shouldn't either.*