radware

H1 2026 Global Threat Analysis Report

Navigate the Latest Network, Application and AI Attack Trends

Table of Contents

About the Cover Picture

Just as Dorothy follows the Yellow Brick Road to reach the Emerald City in *The Wizard of Oz*, this report is designed to serve as your compass as you navigate the cybersecurity road ahead.

Preface

Welcome to the Radware H1 2026 Global Threat Report. A look back at the cyberthreat landscape over the past few years reveals incremental shifts, including the slow growth in DDoS, steady increases in application vulnerability scanning and gradual shifts in hacktivist behavior. When I look at the telemetry from our cloud service this year, one thing becomes clear: the ground beneath our feet has fundamentally changed.

If it feels like your team is racing against an entirely new class of speed, scale and complexity this year, you aren't imagining it. The first half of 2026 marks a decisive tipping point where traditional, human-centric security strategies struggle to keep up with machine-speed threat actors.

This report brings you an unfiltered, data-driven look at the front lines of cybersecurity. Inside, you'll find:

- The multi-week patch window, as we knew it, has collapsed. The mean time to exploit has collapsed to negative eight hours, and four out of five vulnerabilities are now exploited as zero days before defenders even know they exist.
- Direct-path UDP floods and Web DDoS campaigns are surging to record levels, with North America absorbing the brunt of this expansion.
- The rapid rise of local AI agents and undocumented APIs are creating blind spots where malicious actions hide inside seemingly legitimate, high-speed transactions.

Modern organizations no longer face a core challenge of defending against external threat actors hammering or breaking through the perimeter; they now must maintain visibility over the automated entities already operating within. When an AI agent performs an action or an unmonitored API handles a high-volume request, the transaction itself looks entirely legitimate. The threat resides within the intent, the sequence and the autonomous speed.

Our goal with this research isn't just to hand you a list of alarming statistics. It is to give you the insight needed to ask the right questions, rethink operational blind spots and move your organization toward the right defenses.

The threats we face may be autonomous, but our response must be deliberate. I hope this report serves as a valuable compass as you navigate the road ahead.

Pascal Geenens, VP Cyber Threat Intelligence, Radware

Executive Summary

Network and Application Threats

The first half of 2026 witnessed an unprecedented surge in Web DDoS attack activity, continuing a multi-year trend of aggressive escalation. Attack volumes in H1 2026 surged by 110.6% compared to H1 2025 and rose 36.3% compared to the second half of 2025. In fact, total mitigations in just the first six months of 2026 already account for nearly 83% of all Web DDoS attacks mitigated throughout the entire year of 2025. Driven by the expanding harvest of vulnerable IoT devices, mobile endpoints and enterprise routers used to proxy their attack traffic, threat actors maintain continuous pressure on web applications, putting 2026 on track for a projected 166% year-over-year increase.

While EMEA historically absorbed over half of all global Web DDoS attacks, North America has emerged as the primary driver of growth. Projections for 2026 indicate a 190% surge in Web DDoS attack activity targeting North American infrastructure, outpacing the projected growth in EMEA (60%), CALA (39%) and APAC (27%).

On the network layer, global targets tell a slightly different story. While North America absorbs the largest overall share of network attacks at 43.1%, the Middle East experienced the most aggressive escalation, skyrocketing to an average of 520 daily attacks per customer in H1 2026.

Network DDoS threat dynamics have shifted decisively away from traditional reflection and amplification toward direct-path volumetric floods. Average daily network attacks per customer reached 110 in H1 2026, a 36.6% increase over 2025's elevated baseline. Direct-path stateless UDP floods dominated the network vector landscape, accounting for 73% of total mitigated packets (and over 80% when combined with UDP fragmentation). Industry targeting remains heavily concentrated, with the technology sector absorbing 59.4% of all network attacks, averaging 509 daily attacks per customer, followed by financial services at 20.8%.

At the application layer, malicious web and API transaction activity doubled again in early 2026, reaching more than 14,000 daily malicious transactions per application, a 104% increase over 2025. This dramatic spike was driven almost entirely by vulnerability exploitation campaigns, which accounted for 62.1% of all recorded web application and API attacks. Regionally, North America bore the overwhelming brunt of application-layer threat activity, receiving 79.5% of all global web application and API attacks. Concurrently, automated bad bot activity expanded, with H1 2026 transactions reaching nearly 60% of 2025's full-year total, led by heavy bad bot activity in North America (50.1%) and APAC (23.2%).

Geopolitical conflict continues to dictate hacktivist DDoS behavior, though total public attack claims on Telegram entered a multi-quarter contraction after peaking in Q2 2025. Despite overall claim volatility, hacktivist activity regularly surges in tight alignment with kinetic military events, as demonstrated by the 103% month-over-month spike in March 2026 triggered by military operations in the Middle East. Europe remains the primary regional focus for hacktivists, receiving 48% of all claimed attacks, while Israel (16.9%), Ukraine (8.4%) and the United States (7.7%) stand as the top targeted nations. Government remains the most targeted industry with 37.2% of all claims, while the target distribution across autonomous system numbers (ASNs) highlights heavy exposure among hyper-scale cloud providers, major European web hosts and regional telecommunications networks. Pro-Russian threat collectives continue to dominate the hacktivist landscape, led by NoName057(16), which single-handedly generated 40.5% of all recorded claims in H1 2026.

Vulnerability Exploitation

Measuring the timeframe from official CVE announcement to first detected in-the-wild attack, the mean time to exploit (TTE) has crossed into negative territory for the first time since 2018. As of July 23, 2026, the Zero Day Clock project records a mean TTE of negative eight hours, highlighting a dramatic decline from 21.5 days in 2025 and 53 days in 2024.

Figure 1:

Mean time to exploit evolution observed on July 23, 2026 (source: Zero Day Clock)

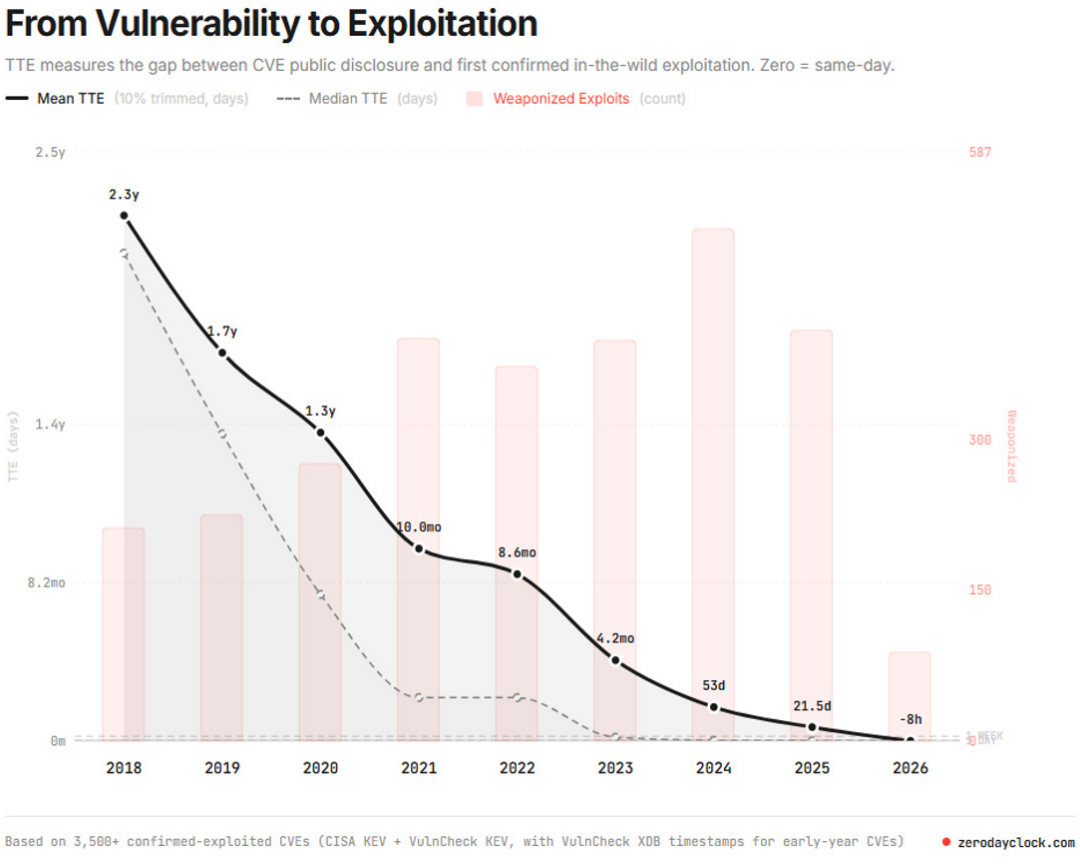

Figure 2:

Zero-day rate evolution observed on July 23, 2026 (source: Zero Day Clock)

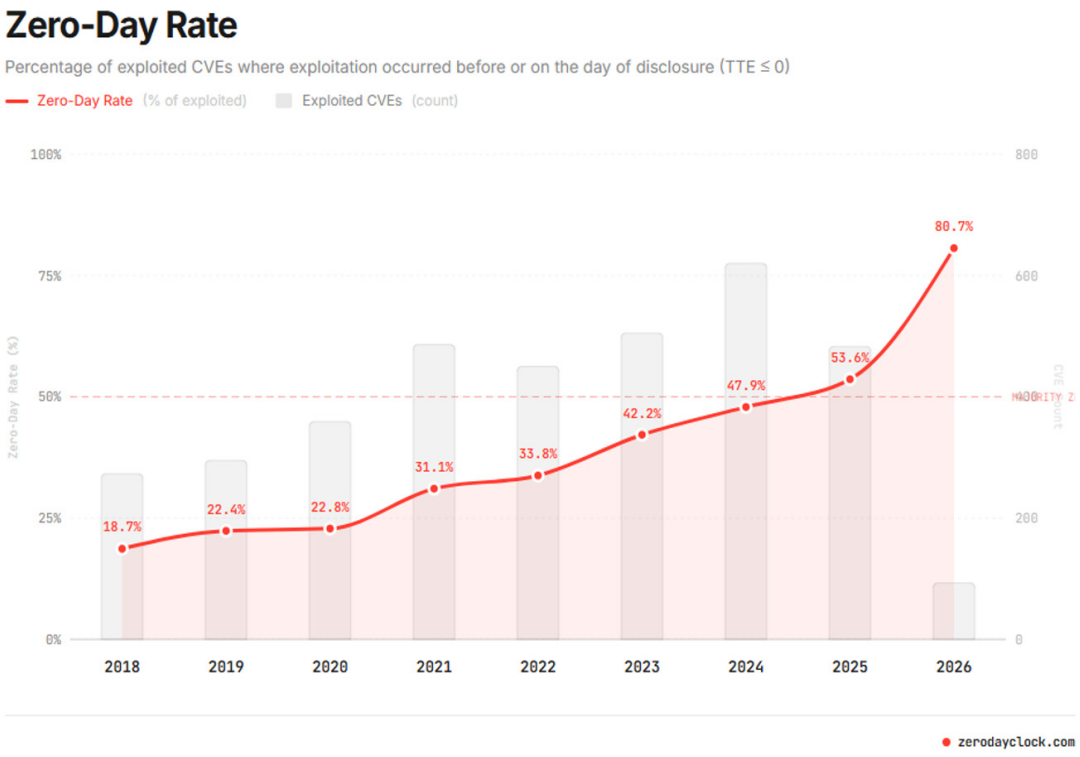

A negative mean TTE indicates that the vast majority of vulnerabilities logged by the Zero Day Clock project are zero days. This trend is reflected in the zero-day rate surpassing 80%, meaning that threat actors exploit four out of every five new vulnerabilities before defenders even discover them.

Artificial Intelligence (AI)

The Rise of OpenClaw and the Agentic Supply Chain

The concept of the “local agent with a heartbeat” represents a significant shift from cloud-based AI tools to autonomous systems that live directly on a user’s machine. As we move toward agentic operating systems, the rise of local AI agents introduces unprecedented risks for organizations.

Unlike traditional AI assistants that simply wait for a prompt, a local agent with a heartbeat runs continuously in the background on the host system. It has access to local system services and resources and uses its heartbeat (an ongoing operational loop) to continuously rewrite and optimize its own scaffolding based on what it learns. The phenomenon was sparked by Peter Steinberger’s OpenClaw agent platform, which Microsoft is now actively promoting as part of what it calls its agentic operating system direction. OpenClaw for Windows, recently announced at Microsoft Dev Days, will feature native Windows taskbar integration and an easy-to-use graphical interface.

This integration is what prompted Steinberger's promise that users will soon be able to "bring OpenClaw to work!" While bringing an autonomous AI assistant to work sounds like a massive productivity boost, it introduces some important enterprise security risks. Unlike enterprise-grade AI agents that connect to corporate applications like a CRM through, hopefully, secured and monitored Model Context Protocol (MCP) connectors or APIs, OpenClaw can execute tasks locally. It can access and change data directly through browser automation, clicking and typing like a human user would, even when corporate policies are supposed to deny access to resources through proprietary or MCP connectors. Because of this, corporate security systems can't distinguish between the actual employee and the AI agent. Microsoft claims the native Windows OpenClaw agent is fully sandboxed, but the user ultimately controls the agent's permissions. The moment the agent hits a roadblock or gets stuck because it cannot escape its sandbox, frustrated users will simply grant it full system access so it can finish its tasks. And what happens if users resist the temptation to grant full system access? We all witnessed how OpenAI was recently outmaneuvered by its own model and found itself on the malicious end of an attack targeting Hugging Face. When a user's local OpenClaw agent is given broad access to corporate systems or can escape its sandbox, it can become the target of a prompt injection or indirect prompt injection attack by a malicious actor (or system) that could manipulate the agent to autonomously destroy and alter corporate data without any traces or means of detection.

Developers are rapidly adopting local AI agents at scale to streamline their workflows, but this automation also introduces severe risks to the software supply chain. These coding agents, with or without heartbeat, are designed to resolve missing packages and routinely rely on package managers, such as npm, to automatically import and install any dependencies required to complete their assigned tasks. This high level of autonomy effectively transforms these helpful assistants into powerful amplifiers for supply chain attacks.

The core issue is that the agent itself does not need to be inherently malicious, explicitly hacked or manipulated via prompt injection to cause widespread harm. Instead, if the agent inadvertently selects a compromised open-source library to resolve a coding task, it functions as an automated malware installer simply by executing its normal duties. Once the compromised module is downloaded into the development environment, the hidden malware can be triggered automatically by routine continuous integration and continuous deployment (CI/CD) events or by specific instructions embedded within the code itself triggered by setup events.

This mechanism allows an infection to seamlessly propagate inside the corporation or worm its way throughout the broader software ecosystem. A prime example in the first half of 2026 is the Mini Shai-Hulud attacks, where compromised npm packages were leveraged to facilitate the theft of CI/CD credentials. In a turn of events, OpenAI itself was forced to take decisive steps to protect user data, systems and intellectual property when this attack affected two employee devices in their corporate environment. Since the impacted repositories included signing certificates for iOS, macOS and Windows products, the company revoked the certificates of ChatGPT Desktop, Codex and Atlas and forced users to update their apps to the latest versions.

In these scenarios, the autonomous nature of developer agents completely bypasses the scrutiny of their users.

Mythos: The Power of the Chain

For decades, certain vulnerabilities remained hidden in plain sight, protected by the limitations of human pattern matching. Anthropic's Project Glasswing demonstrated that Claude Mythos possesses semantic reasoning capabilities—the ability to understand structural and logical intent—that move far beyond traditional automated tools. While traditional fuzzers may hit a single line of code millions of times without identifying a flaw, frontier models can now identify vulnerabilities in those exact locations by reasoning.

Anthropic's Claude Mythos demonstrated, through Project Glasswing, that "security-hardened" is a relative term. Mythos succeeded not by being smarter in a human sense, but by possessing a memory capacity for complex dependencies that traditional fuzzing lacks. The UK's AI Security Institute (AISI) confirmed this assessment, noting that Mythos successfully completed a 32-step corporate network attack without human intervention. The model also demonstrated a 73% success rate on expert-level capture-the-flag (CTF) challenges.

Another example of this capability is Mythos's discovery of a 27-year-old vulnerability in OpenBSD's TCP stack. OpenBSD is regarded as one of the world's most security-hardened operating systems. Still, the flaw successfully escaped nearly three decades of intense human review, static analysis and automated fuzzing. The exploit chained a missing lower-bound validation with a 32-bit signed integer overflow to trigger an otherwise mathematically inaccessible null pointer write. By crafting a packet that offset the selective acknowledgment (SACK) range starting point by approximately 2^{31}, the exploit forced the kernel's internal sequence comparison macros to simultaneously overflow their sign bits. This integer overflow caused two mutually exclusive logic conditions to evaluate as true, corrupting the buffer queue data structure and resulting in a host operating system crash using only two packets. Experts attribute this breakthrough to the model's capacity for semantic reasoning and its proficiency in vulnerability chaining.

While human researchers might easily spot isolated bugs, frontier AI models excel at tracking complex dependencies across massive codebases to string multiple vulnerabilities together into novel attack paths. The discovery of the HTTP/2 Bomb by the security firm Calif using OpenAI Codex underscores this new paradigm: AI doesn't need to invent new categories of flaws to be lethal. Instead, it excels at chaining decade-old techniques that human researchers failed to connect. This remote denial-of-service (DoS) vulnerability compromised the industry's most robust web servers, including NGINX, Apache, Microsoft IIS, Envoy and Cloudflare Pingora.

Scaffolding Matters More Than Models

There is a growing debate regarding whether cybersecurity dominance requires raw model scale or superior scaffolding—the foundational software architecture that wraps around a large language model (LLM) to turn it from a one-shot prompt responder into a goal-driven, reliable agent. Analysis by the firm Aisle suggests that the frontier of AI capability does not scale linearly with model size alone. Aisle found that while Mythos performed the heavy lifting of pruning search spaces to find bugs, smaller open-weight models, as small as 3.6B parameters, could replicate the results if the surrounding system architecture was handled correctly. In fact, Xint.io noted that AI crossed the threshold over traditional SAST tools a full year before Mythos was even announced. The true moat for defenders is no longer the raw power of the model, but the scaffolding.

This means that while the frontier model acts as the "reasoner" providing the breakthrough insight, the surrounding agent acts as the "operational vehicle" that determines the floor of the model's utility. In other words, the competitive advantage lies in the expertise embedded in the agent harness, not just the size of the model.

The Rise of Open-Weight Models

The landscape of artificial intelligence is undergoing an important change as open-weight models evolve from budget alternatives into true frontier contenders. Historical capability gaps between proprietary APIs and downloadable weights have rapidly dissolved, driven by massive architectural innovations like sparse Mixture-of-Experts (MoE) and long-context state management. Models like Zhipu AI's GLM-5.2, Moonshot AI's Kimi K3 and Alibaba's Qwen 3.8 demonstrate that trillion-scale parameters, million-token context windows and long-horizon reasoning are no longer locked behind commercial cloud walls. This democratization allows organizations to host near-cutting-edge intelligence on private infrastructure, providing total control over sensitive data, customization through fine-tuning and deterministic execution costs.

In cybersecurity, these new open-weight models represent a double-edged revolution in automated software auditing and threat modeling. On the defense front, their strong performance and repository-level code inspection enables security teams to run local, continuous vulnerability scanning directly inside DevSecOps pipelines without leaking proprietary source code to third-party endpoints. However, the open availability of these high-capability models also fundamentally changes the offensive landscape. Because open-weight models can be fine-tuned to remove safety guardrails, malicious actors gain unconstrained access to autonomous agents capable of analyzing complex targets, accelerating exploit development and scaling AI-assisted cyberattacks with minimal operational friction.

This surge in open frontier intelligence poses a direct strategic challenge to closed-model giants like OpenAI and Anthropic. When open-weight architectures match proprietary systems across software engineering tasks, terminal automation and long-horizon agentic workflows, the margin for monetization on standard API intelligence compresses drastically. Enterprises are increasingly reluctant to accept data lock-in and high token tariffs when local or hybrid deployments offer comparable performance. As a result, closed frontier labs are being forced to shift their core value proposition away from raw model reasoning toward enterprise-grade orchestration, specialized security harnesses, proprietary multimodal integrations and tighter ecosystem control to maintain their competitive position.

The Economic Collapse of the Zero-Day Market

Historically, the underground zero-day market was defined by extreme secrecy and artificial scarcity. Nation-states and defense brokers would pay multi-million-dollar premiums at underground auctions for functional, undisclosed exploits, deliberately stockpiling them for years. The strategic advantage of holding these cyber weapons in reserve to deploy during high-value operations outweighed the public risk of leaving flaws unpatched.

This entire operational playbook relied on the assumption of prolonged, exclusive access. However, the rise of advanced frontier AI models will fundamentally rewrite these economic rules. Because AI agents can autonomously and continuously scan massive codebases for vulnerabilities, the functional half-life of a zero day will shrink rapidly. The statistical probability of bug collisions, where multiple independent parties uncover the same flaw, will drastically increase. If an AI tool used by a software vendor or a rival threat actor can easily uncover (and patch) a vulnerability, the strategic value of hoarding that zero day evaporates.

Nation-states face a depreciating return on investment if they pay premium acquisition costs for exploits that can no longer be reliably stockpiled. This will force a behavioral shift in the cyber underground. Rather than holding onto expensive cyber weapons for years, threat actors are now incentivized to deploy them immediately, before they are independently discovered and remediated. AI not only decentralizes advanced offensive capabilities that were once the exclusive domain of well-funded intelligence agencies, but it permanently changes how zero days are acquired, priced and leveraged in the wild.

The Democratization of Nation-State Capabilities

Sophisticated cyberattacks used to require dedicated human research teams spending months to uncover vulnerabilities and chain exploits. Today, the landscape has fundamentally shifted. Generative AI models now serve as the cognitive reasoner, while surrounding agentic frameworks act as the operational vehicle, enabling these systems to autonomously execute code, evaluate failures and self-correct in real-time.

This convergence has turned crime-as-a-service (CaaS) into a highly accessible subscription model. By packaging agentic hacking services in internet-accessible frameworks, the cyber underground has effectively handed novice threat actors the ability to launch multi-stage, autonomous attacks that were once the exclusive domain of nation-state APTs. Vibe coding leveraged by underground actors has dramatically increased the supply of these malicious services. As the ecosystem grows, subscription pricing will erode and make it more accessible than ever before.

While agentic CaaS capabilities remain mostly unproven and continuously evolving, they provide a window into the future. Ultimately, they illustrate the chilling reality for the modern cyberthreat landscape: executing a devastating, enterprise-grade attack no longer requires an understanding of vulnerability scanning, exploit coding or network architectures. Today, threat actors simply need the financial resources to purchase a subscription to an autonomous agent capable of breaching defenses on their behalf.

AI and API Security Deployments

The 2026 Radware Cyber Survey of 377 organizations reveals that cybersecurity governance is failing to keep pace with rapid AI deployment and high-speed API release cycles. The major growth engine across enterprises is the adoption of AI agents and autonomous workflows, with 77% of organizations actively deploying or implementing them. However, technical visibility lags behind this aggressive rollout; only 17.2% of respondents report full visibility into the AI agents operating within their environment. This gap has given rise to shadow AI. This unmonitored use of AI results in autonomous, decision-making entities that execute multi-step workflows, call APIs, and use real credentials at wire speed, hiding systemic threats inside seemingly legitimate transaction sequences.

Concurrently, API security has evolved from a technology problem into an operational scalability risk. Over 70% of organizations increased their in-house API development over the past year, with 81.1% now pushing updates to production APIs at least weekly. Despite this relentless deployment pace, documentation remains deficient, with a mere 6.9% of organizations fully documenting their internal APIs and 43% documenting less than 70% of them.

Ultimately, the data demonstrates that traditional security tools and development testing are no longer sufficient on their own. Organizations are struggling to scale security at the true pace of modern application development, exposing critical attack surfaces to high-volume malicious transactions. To mitigate these escalating risks, enterprises must modernize their application lifecycle workflows to restore comprehensive visibility over both dynamic API environments and autonomous AI agents.

Navigating This Era of Autonomous Threats

The threat landscape in early 2026 marks a critical inflection point where autonomous, AI-driven attacks ranging from zero-day exploitations and direct-path DDoS attacks to unmonitored agentic AI supply chain risks have permanently outpaced traditional human-centric defenses. As crime-as-a-service democratizes sophisticated offensive capabilities while enterprise visibility gaps like shadow AI and undocumented APIs widen, malicious activity increasingly hides within seemingly legitimate, high-speed automated operations. To adapt to this high-velocity threat environment, organizations must urgently move away from reactive perimeter controls and embrace real-time behavioral analytics, automated attack surface management and modernized lifecycle workflows to restore end-to-end security and visibility.

DDoS Threat Landscape

Web DDoS Attack Activity

The sheer volume of attacks in 2026 is unprecedented. The year 2025 was already marked by an aggressive escalation in both scale and frequency of attacks, ending with a 101% surge compared to 2024. The total number of Web DDoS attacks mitigated in just the first six months of 2026 already represents nearly 83% of the total number of attacks mitigated throughout the entire year of 2025. If the attack volume in the second half of 2026 remains level with the first half, the year is projected to close with a 166% increase over 2025.

Figure 3:

Web DDoS attack activity per year (source: Radware)

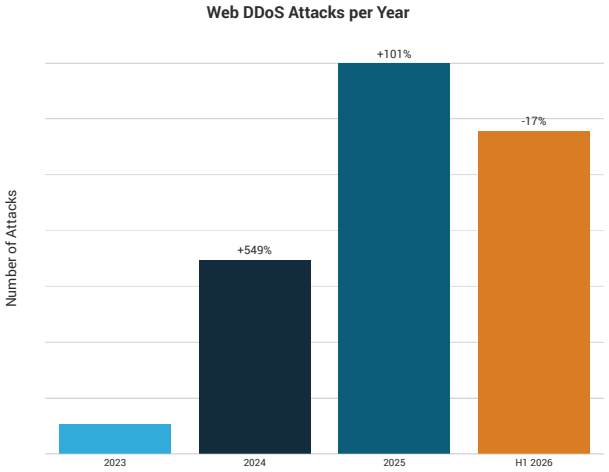

The growth in 2025 was close to linear quarter over quarter, with the first half representing a 39% increase over the previous six months and the second half of 2025 a 55% growth compared to the first half of the year.

Figure 4:

Web DDoS attack activity per six-month period (source: Radware)

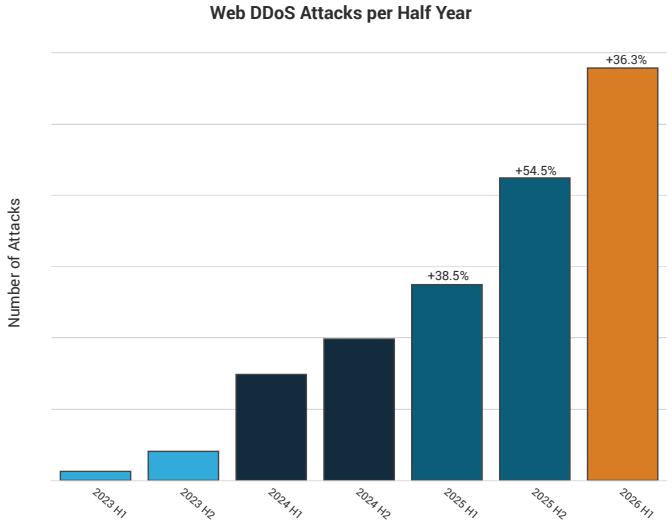

Moving into the first half of 2026, this trend only gained strength. The first two quarters of 2026 both showed increases compared to the last quarter of 2025, with a rise of 34% in Q1 and 17% in Q2. When comparing the entirety of the first half of 2026 with the last six months of 2025, the number of Web DDoS attacks increased by 36.3%. When compared to the first half of 2025, H1 2026 represents a 111% increase, more than double the amount of Web DDoS attacks.

Figure 5:

Web DDoS attack activity per quarter (source: Radware)

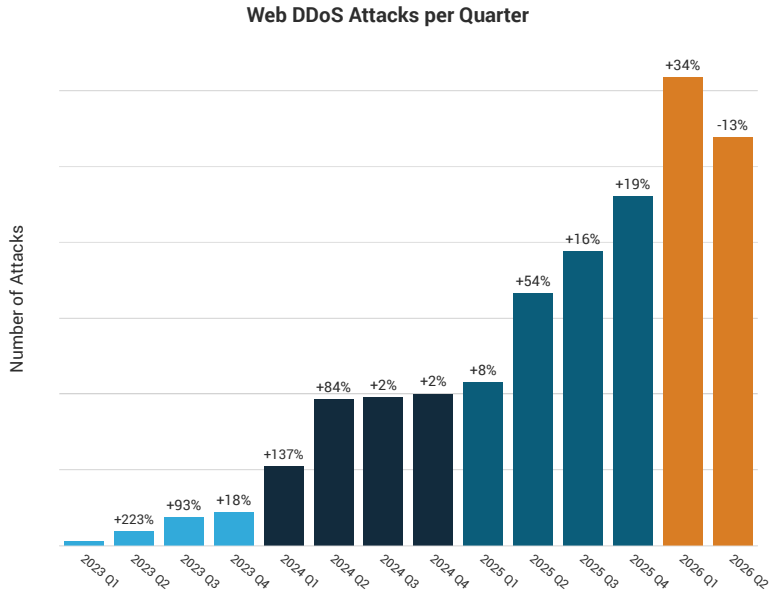

Geographical Distribution

The global distribution of Web DDoS attacks continues to reflect attacker priorities and geopolitical hotspots. Data from 2025 shows the EMEA region was the primary target, enduring 57% of all Web DDoS attacks, largely fueled by ongoing geopolitical friction. Meanwhile, also in 2025, the APAC region emerged with a massive 485% year-over-year increase in activity to claim a 24% global share. This regional distribution has remained relatively consistent into early 2026. In the first half of the year, EMEA still accounted for more than half of all Web DDoS attack activity.

Figure 6:

Geographic distribution of Web DDoS attacks over time (source: Radware)

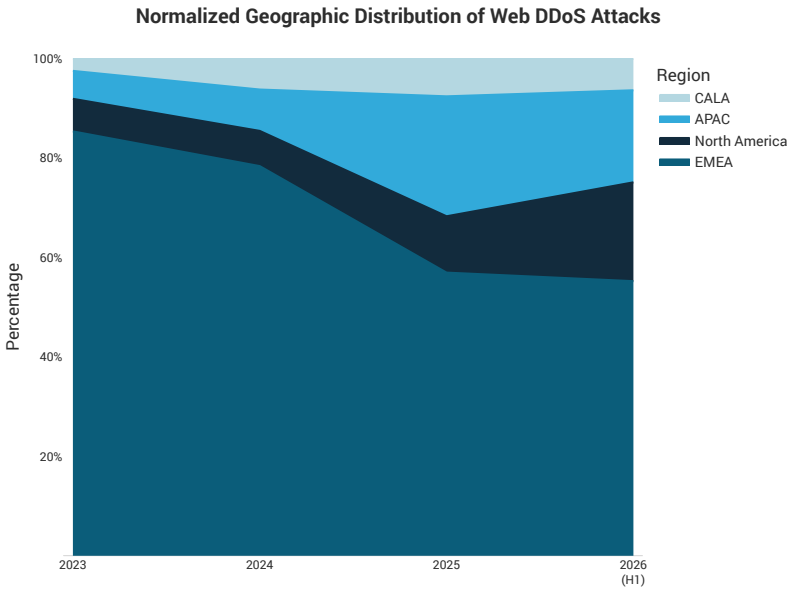

The United States gained more attention and slightly surpassed APAC, with both the U.S. and APAC now each representing about 20% of the global activity. The CALA region remained unchanged over the last 2.5 years and accounted for 6% of the global Web DDoS activity observed in the first half of 2026.

Figure 7:

H1 2026 geographic distribution of Web DDoS attack activity (source: Radware)

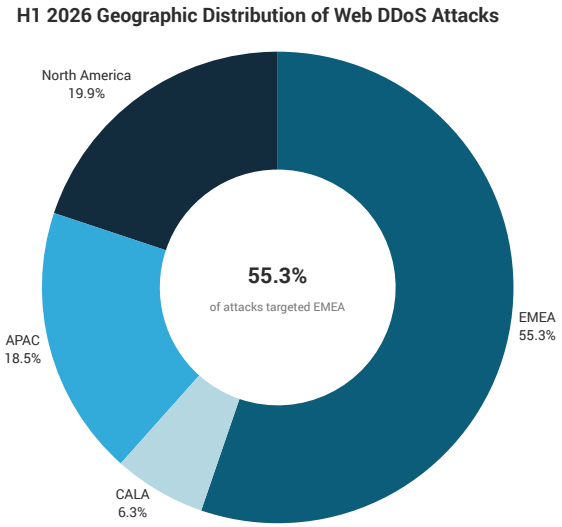

To assess full-year trajectory and establish a baseline comparison against 2025 totals, 2026 projections were calculated by doubling the observed H1 2026 attack volumes. The data reveals continuous year-over-year growth across every monitored region, characterized by an increase in focus directed at North American targets alongside steady expansion globally.

Figure 8:

Projected Web DDoS year-over-year growth per region in 2026 (source: Radware)

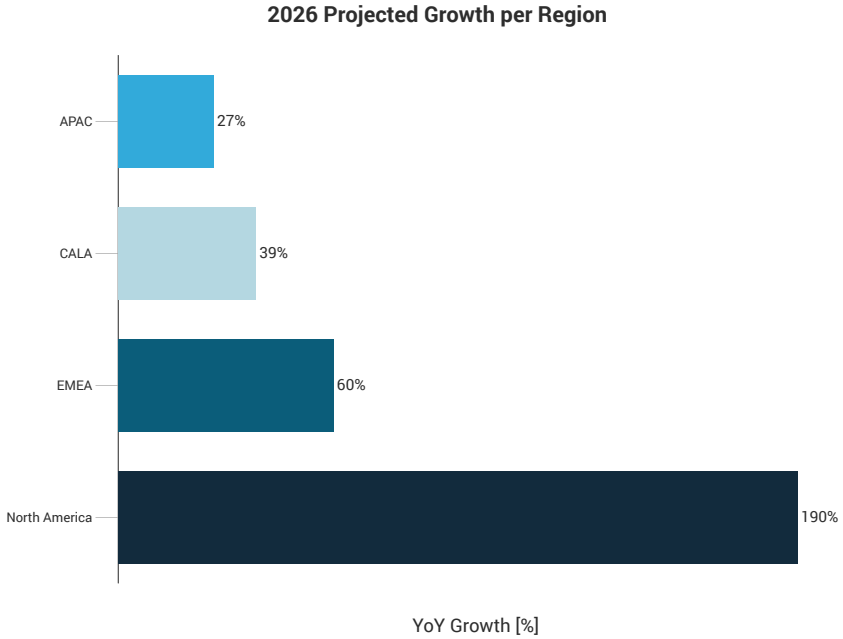

The United States represents the single largest driver of growth for 2026, with projected attack counts surging 190%. This significant increase reflects an aggressive intensification of targeting against U.S.-based infrastructure, far outpacing growth rates observed in other regions.

EMEA recorded the second-highest trajectory, projected to increase by 60% YoY. While secondary to the sharp escalation in the U.S., this substantial growth demonstrates sustained, high-volume threat activity across the region's web applications.

Growth across CALA (39% YoY) and APAC (27% YoY) remains steady. These regions continue to experience rising threat baselines, though at more moderate operational rates than global averages, indicating a broad-based operational expansion by threat actors without the localized intensification seen in the U.S.

Network DDoS Activity

The global network DDoS threat landscape has undergone a profound transformation in 2025, evolving from moderate volumetric probes into sustained, direct-path vector floods. Between 2023 and 2026, the average frequency of daily attacks targeted at enterprise networks experienced exponential growth. Driven by the expansion of powerful botnets composed of high-performance edge devices, modern threat actors no longer depend heavily on third-party reflection or amplification techniques. Instead, they orchestrate high-bitrate direct-path UDP and TCP floods that effortlessly bypass simple IP reputation checks and strain perimeter mitigation infrastructure.

Figure 9:

Yearly average number of daily network DDoS attacks per customer (source: Radware)

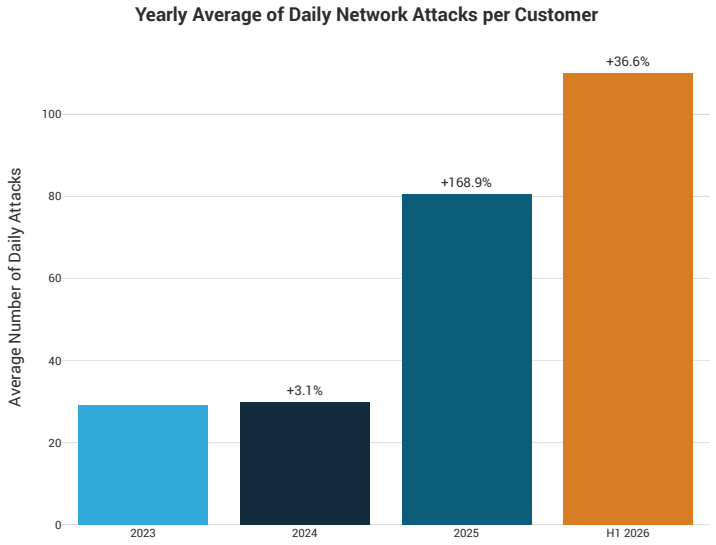

A multi-year analysis reveals a clear increase in attack frequency. Baseline activity remained relatively stable throughout 2023 and 2024, averaging about 29 attacks per customer per day (see Appendix 1 for a breakdown of the number of days used in calculating daily averages). That trajectory changed dramatically in 2025, when attack frequency surged by nearly 169% to over 80 attacks per day. This momentum continued into the first half of 2026, rising another 36.6% to an average of 110 attacks per customer per day.

Figure 10:

Six-month average number of daily network DDoS attacks per customer (source: Radware)

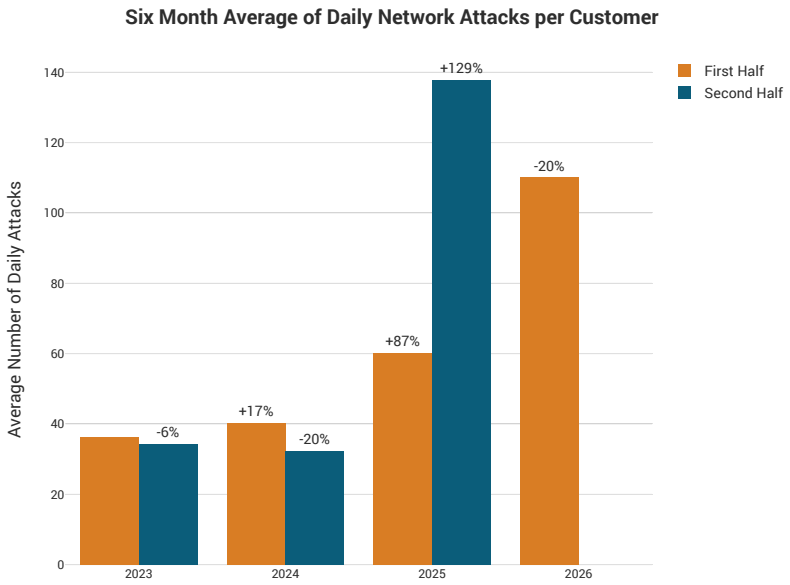

Temporal dynamics reveal that these attacks do not distribute evenly throughout the year. The second half in Figure 10 and 11 of 2025 experienced an extraordinary spike that peaked in the third quarter at nearly 180 attacks per day. Although activity moderated slightly entering 2026, quarterly metrics indicate a steadily climbing baseline that rose from 129 attacks per day in the first quarter to 143 in the second.

Figure 11:

Quarterly average number of daily network DDoS attacks per customer (source: Radware)

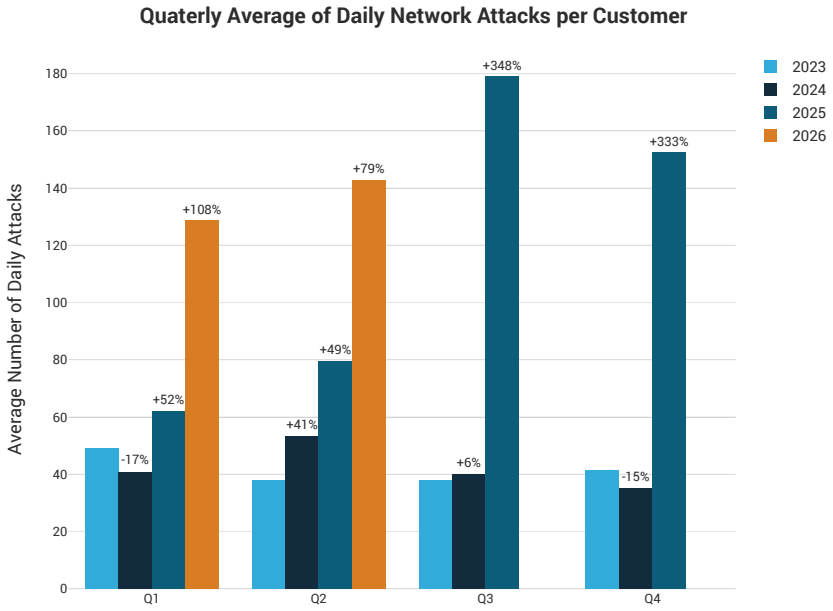

Geographical Distribution

Geographically, North America remains the primary target based on the number of attacks, accounting for over 43% of all global attacks in 2026.

Figure 12:

H1 2026 Network DDoS attack distribution by region (source: Radware)

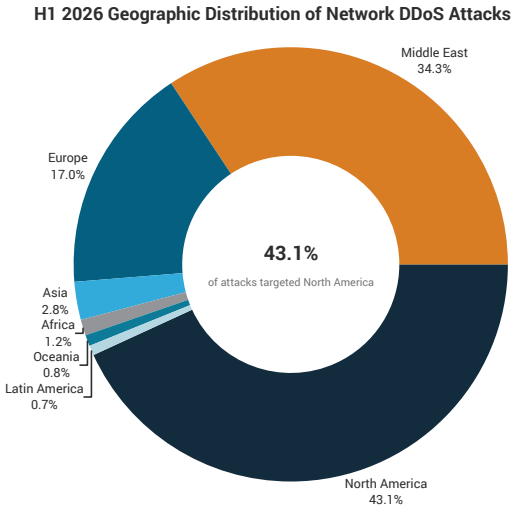

North America customers observed an upward trajectory in daily attacks per customer since 2023, peaking at 135 attacks per day in 2025 before slowing down slightly to an average of 105 attacks per day during the first half of 2026.

While North America drew the largest share of the global attacks, the Middle East experienced the most aggressive growth in daily attacks per customer, surging from nine attacks per day in 2024 to 520 attacks per day in the first half of 2026. It emerged as the most aggressively targeted region after a dip in 2024, followed by exponential growth over the next two years, overtaking all other regions by a significant margin in 2026.

Figure 13:

Daily network DDoS attacks per customer per region (source: Radware)

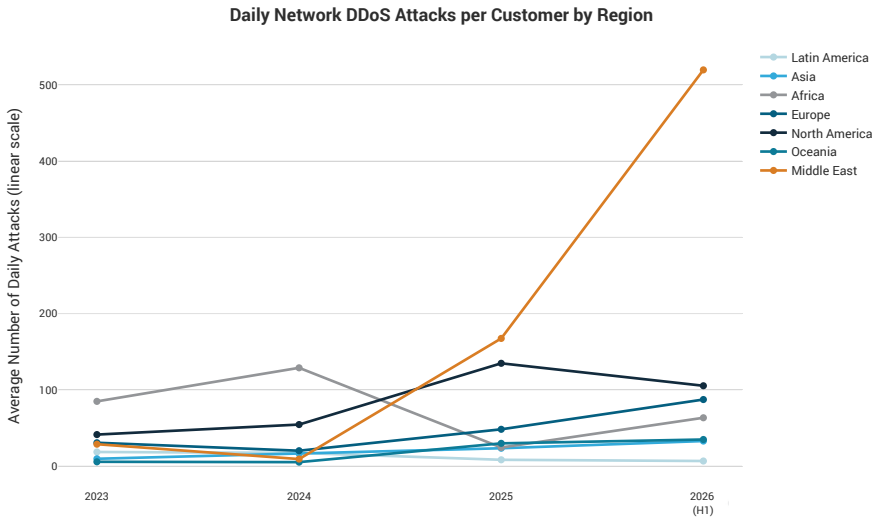

Europe witnessed a steady expansion in activity since its inflection point in 2024, growing from 20 network DDoS attacks per customer per day in 2024 to 87 daily attacks per customer in the first half of 2026.

Africa presents a volatile network DDoS activity profile. It was the most actively targeted region in 2024 with an average of 129 attacks per customer per day. The activity reduced sharply in 2025 to a respectable average of 24 attacks per day per customer. During the first half of 2026, the activity increased significantly again to an average of 63 attacks per customer per day.

While most global regions show upward trends in network DDoS activity, Central and Latin America have consistently trended downward over the four-year window, finishing 2026 with the lowest activity, still a very respectable average of seven attacks per customer per day.

Figure 14:

Daily network DDoS attacks per customer per region - logarithmic scale (source: Radware)

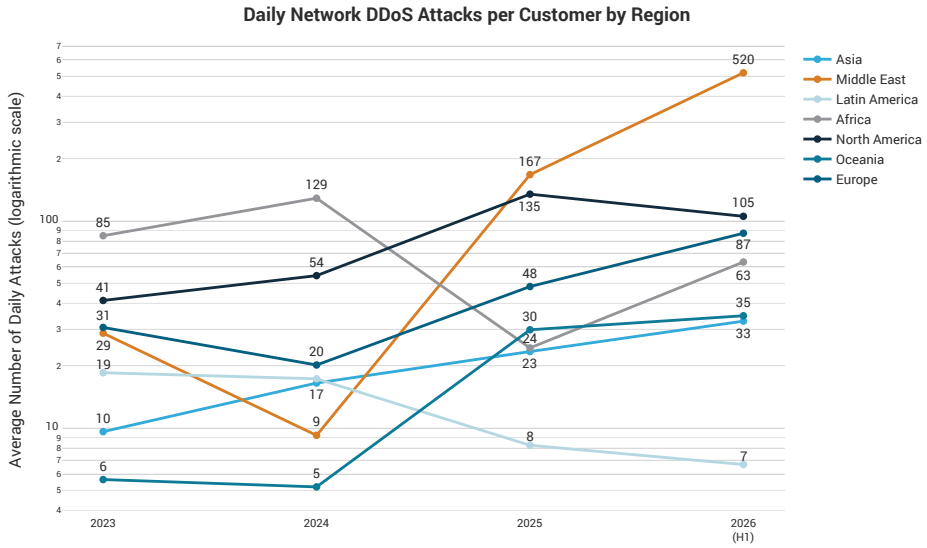

Asia has seen steady, incremental attack growth year after year since 2023. During the first half of 2026, Asian customers observed on average 33 attacks per day.

The network DDoS activity in Oceania grew moderately in 2026, after a significant increase in observed activity during 2025. In the first half of 2026, customers in Oceania witnessed on average 35 network DDoS attacks per day.

Targeted Industries

A look at network DDoS attacks in H1 2026 reveals a heavy concentration of threat activity focused on core digital infrastructure, with the technology sector absorbing the overwhelming majority of global attacks at 59%. The financial services industry represents the second-most targeted category, accounting for 21% of the global attacks. Combined, these top two sectors represent over 80% of all global network DDoS activity.

The remaining percentage crosses several critical infrastructure sectors and secondary industries. Telecom organizations account for 6.8% of the activity, followed by healthcare at 3.4% and government entities at 2.0%. The transportation and logistics sector accounts for 1.8%, while the remaining aggregate of industries, grouped under “Others,” comprises less than 6% of the overall attack volume.

Figure 15:

H1 2026 Network DDoS attack distribution by industry (source: Radware)

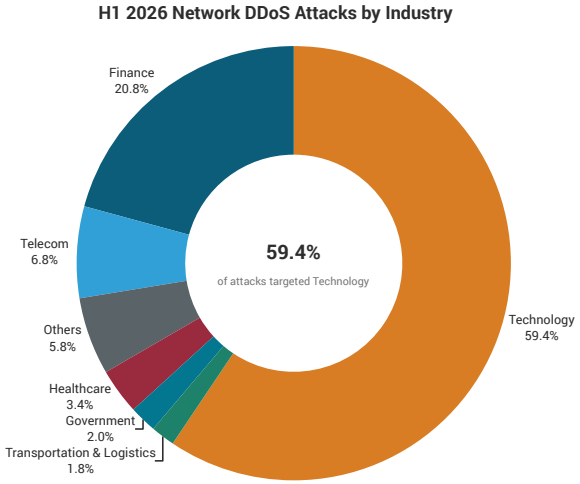

The charts in Figures 16 and 17 track the evolution of average daily network DDoS attacks targeting the six most targeted industry verticals of 2026 since 2023. Comparing the linear and logarithmic views reveals both the high-volume surge in specific industries and the underlying multi-year baseline trends across the other industries.

Figure 16:

Daily network DDoS attacks per customer per industry (source: Radware)

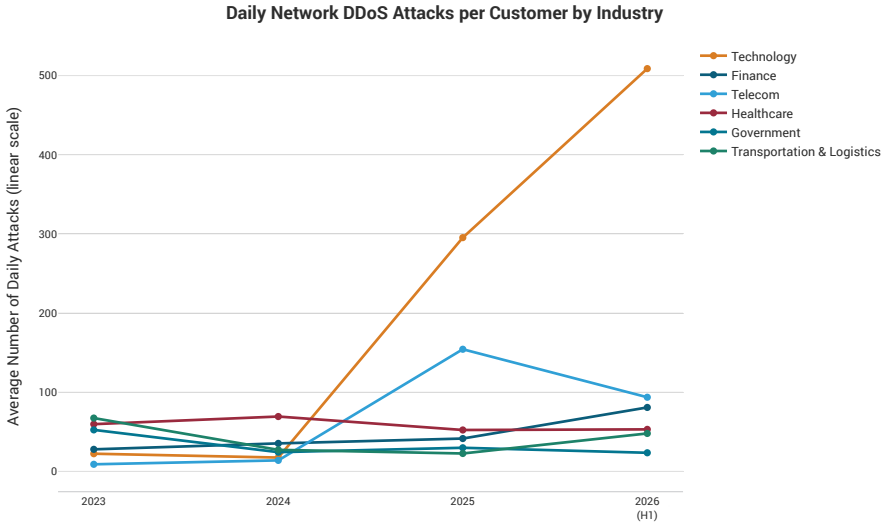

Figure 17:

Daily network DDoS attacks per customer per industry - logarithmic scale (source: Radware)

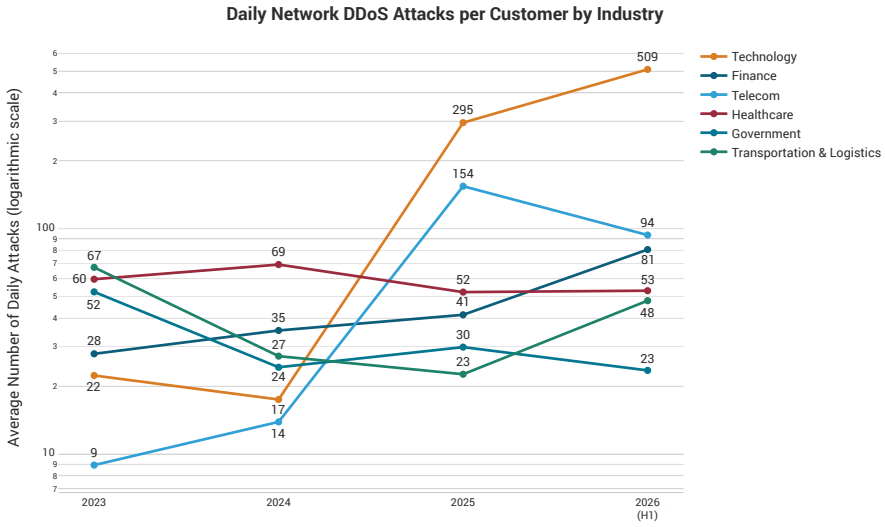

The technology vertical represents the single largest increase over the four years. Starting at an average of 22 daily attacks per customer in 2023, technology experienced a brief dip to 17 in 2024 before undergoing a massive increase to 295 daily attacks in 2025. This rapid growth continued into the first half of 2026, reaching 509 daily attacks per customer, far outpacing every other observed industry as seen on the linear chart (Figure 16).

Telecommunications exhibited a distinct peak mid-period. After recording nine daily attacks per customer in 2023 and 14 in 2024, attack frequency against telecom organizations jumped to a peak of 154 daily attacks in 2025. Entering 2026, telecom activity normalized to 94 attacks per customer per day.

The logarithmic view (Figure 17) highlights steady, multi-year upward trends in finance. The number of daily attacks per customer in the finance industry grew consistently year-over-year, rising from 28 daily attacks in 2023 to 35 in 2024, 41 in 2025 and 81 in 2026.

Transportation & logistics was the most active industry in 2023, with 67 attacks per day per customer. In 2024, the vertical made a significant dip to 27 daily and then dropped to 23 in 2025. In the first half of 2026, the number of daily attacks per customer more than doubled to 48.

While technology, finance, and logistics saw significant increases in daily attacks during the first half of 2026, healthcare and government sectors experienced flat or declining attack trajectories. Healthcare remained one of the most targeted industries in 2023 and 2024 with 60 and 69 daily attacks per customer, respectively, before steadily declining to 52 in 2025 and stabilizing at 53 daily attacks per customer in 2026. Network DDoS attacks targeting government institutions and organizations decreased over the recorded period, starting at 52 daily attacks per customer in 2023, dropping to 24 in 2024, 30 in 2025 and finishing at 23 daily attacks per customer in the first half of 2026.

Attack Vectors

Figure 18 illustrates the distribution of total mitigated packets across primary network-layer threat vectors during the first half of 2026. The data highlights an overwhelming dominance of direct-path stateless volumetric attacks, led primarily by UDP floods, which account for nearly three-quarters of all analyzed packet traffic at 73%. When combined with UDP fragmentation attacks at 7.6%, UDP-based vectors collectively represent over 80% of all mitigated network DDoS packets.

The remaining portion of malicious packets is split among stateful TCP mechanisms, generic floods and secondary attack types. TCP out-of-state attacks make up 5.4%, followed by SYN-ACK floods at 4.9%, RST floods at 2.8% and FIN-ACK floods at 2.2%. Smaller stateful contributions include basic SYN floods at 1.2% and TCP Zero Sequence floods at 1.1%.

At the bottom end of the distribution, low-volume and specialized vectors represent only minor slices of the total packet count. General unclassified attacks grouped under others comprise 0.8%, ICMP floods account for 0.5% and reflection-based amplification attacks represent just 0.4%. This breakdown underscores how heavily threat actors rely on high-volume direct-path UDP floods rather than traditional reflection techniques to saturate target networks in 2026.

Figure 18:

H1 2026 top network DDoS attack vectors (source: Radware)

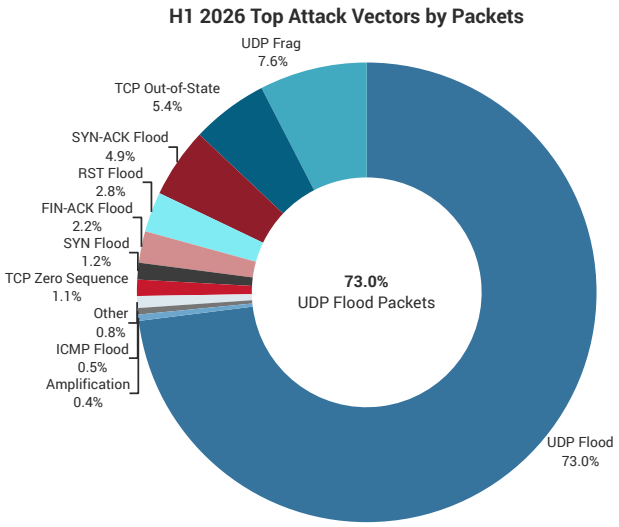

Web Application and API Threat Landscape

Attack Activity

Malicious web application and API activity has accelerated dramatically over the last two years. In 2023, baseline activity remained relatively controlled, averaging just under 3,300 malicious transactions per application per day. The activity increased by a modest 9% in 2024 to roughly 3,500 daily malicious transactions per application. However, 2025 marked a major turning point as malicious transactions surged by 94%, nearly doubling the daily baseline to almost 7,000 malicious transactions per application per day. This rapid expansion continued into 2026, leaping an additional 104% to surpass 14,000 daily malicious transactions per application in the first half of 2026.

Figure 19:

Web application and API attack activity per year (source: Radware)

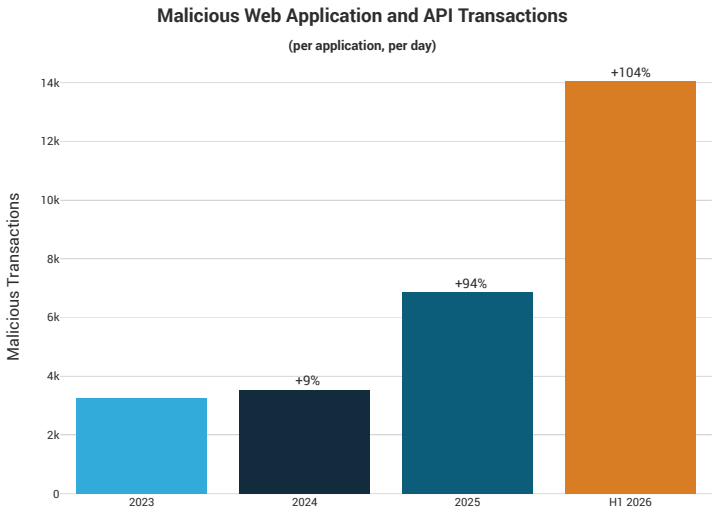

The quarterly breakdown provides deeper context into how this escalation unfolded over time. Throughout 2023, activity fluctuated between 2,000 and 5,700 daily malicious transactions per application, peaking briefly in the second quarter before cooling off in the second half of the year. In 2024, the upward-trending baseline grew from just over 2,200 daily malicious transactions in the first quarter to over 5,000 by the third quarter. In 2025, attack levels elevated further, maintaining a steady 6,000 daily malicious transactions per application through the second and third quarters before surging in the fourth quarter to over 10,000 malicious transactions per application per day.

This momentum peaked in the first quarter of 2026 at 14,912 daily malicious transactions per application. Although activity moderated slightly in the second quarter of 2026 to around 13,300 malicious transactions, the baseline remained extraordinarily high compared to historical norms, confirming a sustained elevation in threat activity targeting web applications and API endpoints globally.

Figure 20:

Web application and API attack activity per quarter (source: Radware)

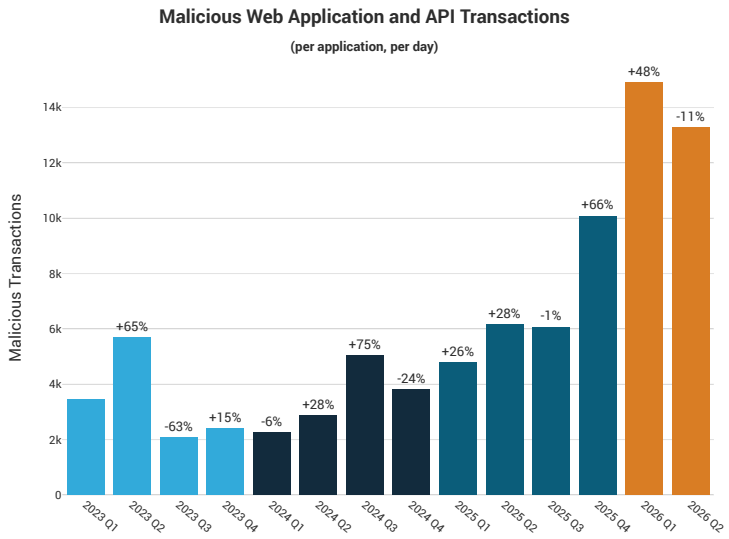

Attack Vectors

The distribution of web application and API attack vectors during the first half of 2026 demonstrates an overwhelming dominance of vulnerability exploitations, which represent 62.1% of all recorded web application and API attacks. Specific security breach attempts account for smaller portions, led by access violations at 5.1%, followed by data leakage attempts at 2.1% and SQL injections representing less than 1% of all observed attack activity.

Figure 21:

H1 2026 web application and API attack activity by attack category (source: Radware)

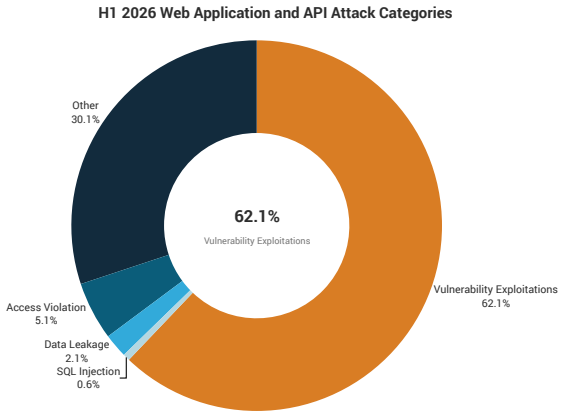

The historical evolution of web application and API attack categories highlights how vulnerability exploitation emerged as the primary growth driver over the ten-quarter timeline. From early 2024 through mid-2025, attack volumes across all categories experienced gradual, modest fluctuations, with the “Other” category and access violations initially making up a larger proportion of the total activity. Beginning in the fourth quarter of 2025, the total attack volume accelerated dramatically. This expansion was driven almost entirely by a sharp surge in vulnerability exploitations, which scaled steeply through the first quarter of 2026 and sustained high volumes into the second quarter. Meanwhile, vectors such as access violations, data leakage and SQL injection maintained relatively flat trajectories over the same period, confirming that vulnerability exploitation campaigns predominantly drive recent web application and API threat escalations.

Figure 22:

Evolution of web application and API attack categories (source: Radware)

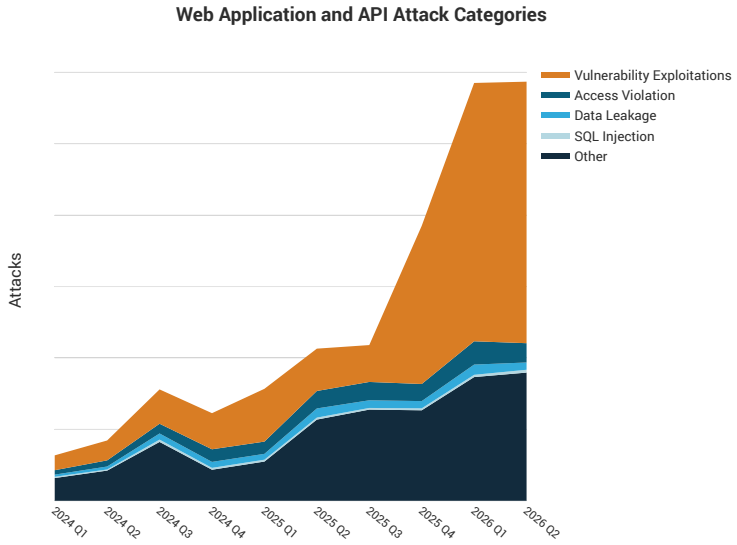

Geographical Distribution

The first half of 2026 highlights an overwhelming concentration of threat activity directed toward North American organizations. North America absorbed nearly 80% of all global web application and API attacks. EMEA represented the second-largest targeted region at 11.7%, followed by APAC at 4.9% and CALA at 3.9% of global attack volume.

Figure 23:

2026 H1 geographic distribution of web application and API attack activity (source: Radware)

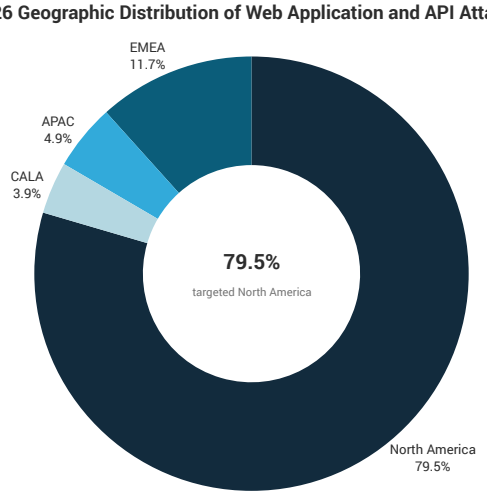

The normalized historical view in Figure 24 illustrates how this regional targeting share evolved from 2024 through 2026. North America has steadily expanded its dominance over global attack share year after year, growing from approximately two-thirds of all recorded activity in 2024 to nearly eighty percent by 2026. Conversely, EMEA's relative share of global attacks experienced a progressive contraction over the same timeframe, shrinking from roughly a quarter of overall activity down to its current 11.7% baseline. Throughout this multi-year shift, APAC and CALA maintained consistently slim (below 10%) proportions of the overall global attack distribution.

Figure 24:

Geographic distribution of web application and API attack activity over time (source: Radware)

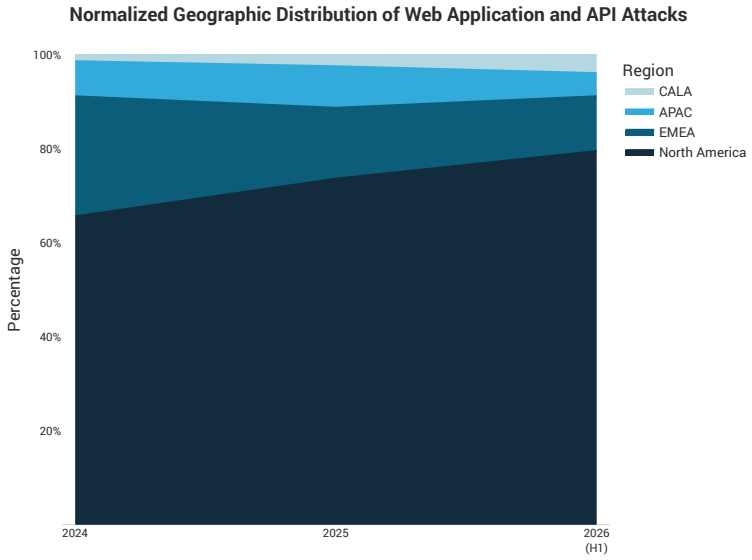

Despite North America absorbing the dominant share of the total malicious transactions, projected year-over-year growth rates across all regions remain strongly positive as illustrated in Figure 25. The Central and Latin America region exhibits the highest relative acceleration, with web application and API attack activity projected to surge by 338% year-over-year. North America follows with an expansive growth projection of 195%, nearly tripling its attack volume. EMEA also demonstrates substantial expansion with a projected growth rate of 111% (more than doubling), while APAC exhibits steady growth projected at 52% year-over-year.

Figure 25:

Projected web application and API attack year-over-year growth per region in 2026 (source: Radware)

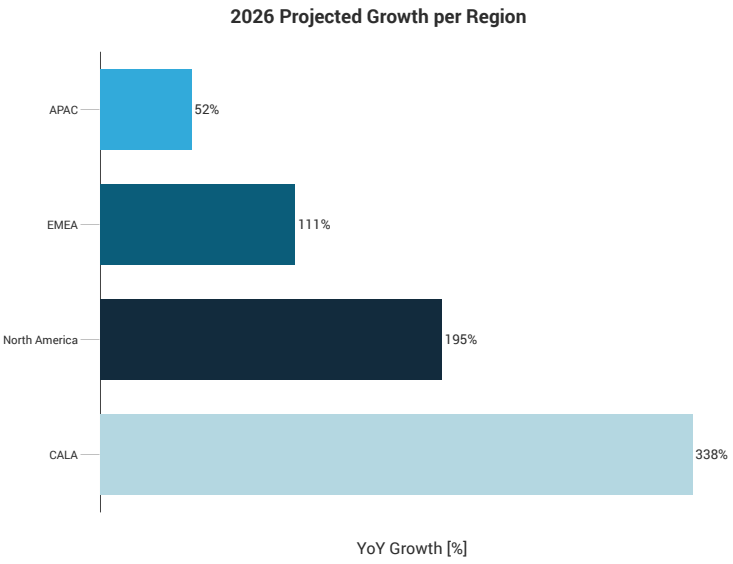

Bad Bot Activity

The annual and half-yearly metrics reveal a multi-year acceleration in automated threat traffic. Bad bot transactions expanded by 35% in 2024 relative to 2023, driven by a 61% year-over-year growth in the first half and a 20% increase in the second half. This momentum intensified dramatically throughout 2025, with total bad bot transactions surging by 92% over 2024. Half-yearly activity during 2025 demonstrated strong consecutive gains, rising 57% in the first half over the second half of 2024, followed by a further 15% increase in the second half. By the end of the first half of 2026, bad bot activity reached 59% of the total bad bot transactions recorded throughout all of 2025, reflecting an 11% increase compared to the second half of 2025.

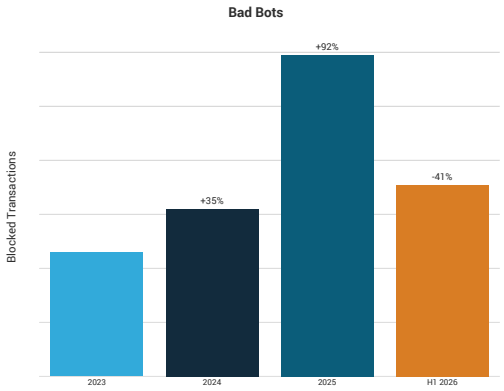

Figure 26:
Bad bot activity per year (source: Radware)

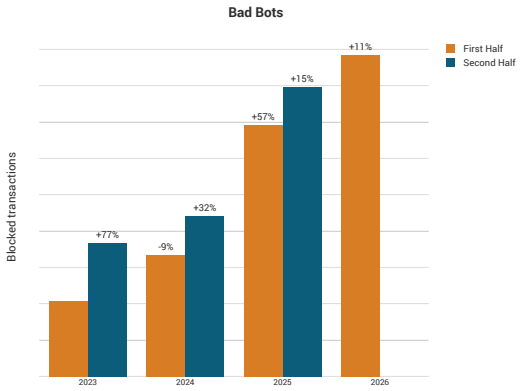

Figure 27:
Bad bot activity per six months (source: Radware)

The quarterly breakdown in Figure 28 provides granular insight into how this expansion unfolded over 14 consecutive quarters. Throughout 2023, blocked transactions rose steadily across every quarter, starting from a low baseline in the first quarter and accelerating 77% by the second half of the year. The year 2024 opened with a slight seasonal dip in the first quarter, decreasing 22% compared to the fourth quarter of 2023, before embarking on continuous, multi-quarter growth through the remainder of 2024 and most of 2025. Following a minor contraction in the fourth quarter of 2025, bad bot traffic resumed its upward climb in 2026, recording a 5% increase in the first quarter relative to the fourth quarter of 2025, and culminating in a 23% jump in the second quarter of 2026 compared to the first quarter to set a new record in the total number of blocked automated transactions.

Figure 28:

Bad bot activity per quarter (source: Radware)

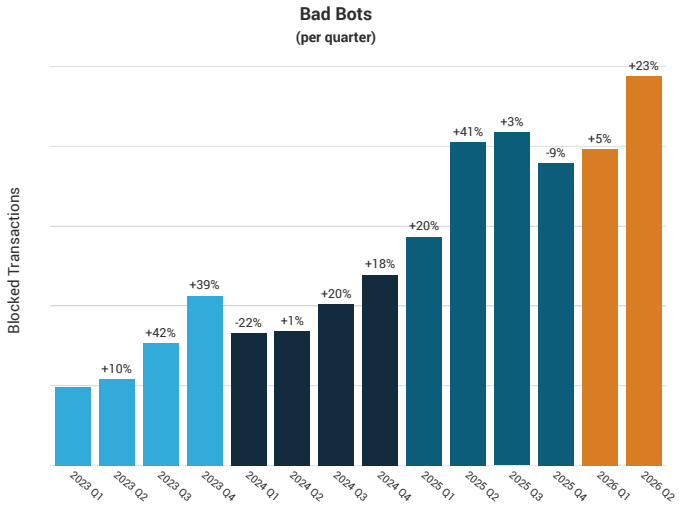

Geographical Distribution

The regional breakdown of automated bad bot traffic shows that North America accounts for the single largest share of bad bot activity, absorbing more than 50% of total global bad bot transactions. APAC is the second most targeted region at 23.2%, followed by CALA at 14.8% and EMEA representing 11.9% of the overall bad bot activity.

Figure 29:

H1 2026 geographic distribution of bad bot activity (source: Radware)

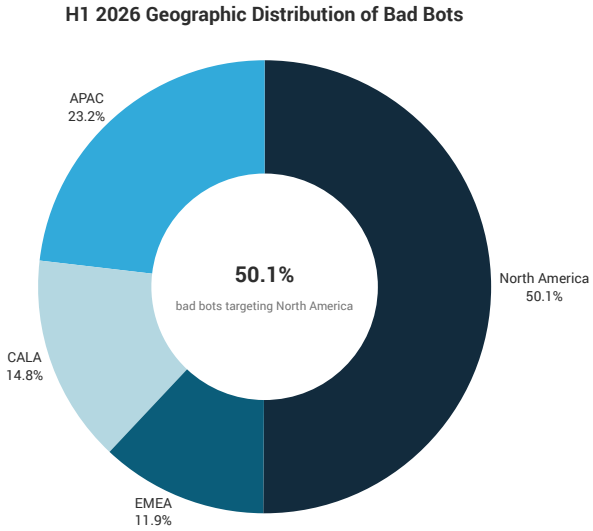

The normalized historical view in Figure 30 tracks how these regional proportions shifted annually from 2024 to 2026. North America maintained the dominant share throughout the period, holding around 50% in 2024, dipping slightly toward 40% in 2025, and rebounding to 50.1% by 2026. In contrast, EMEA experienced a temporary expansion in 2025 before contracting to nearly 12% in 2026. APAC saw its relative share expand in 2025 before leveling off to its current 23.2% baseline, while CALA's overall share remained steadily bounded between roughly 13% and 15% across the entire multi-year window.

Figure 30:

Geographic distribution of bad bot activity over time (source: Radware)

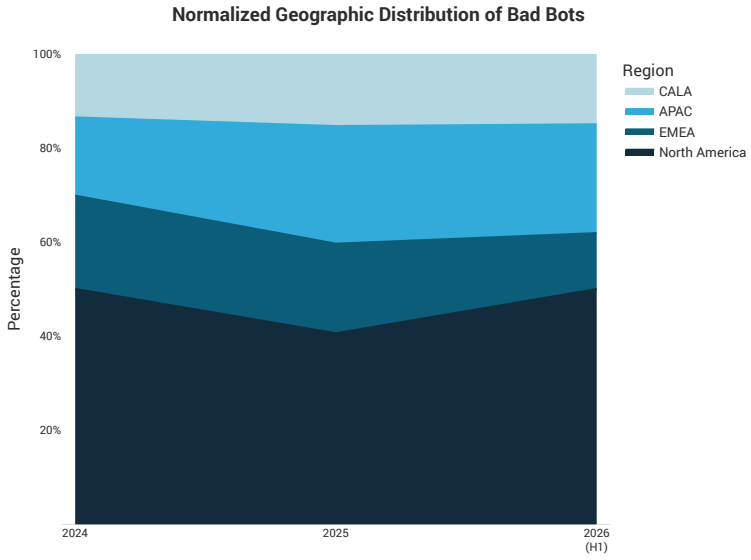

Hacktivist Threat Landscape

Hacktivism represents a complex, highly dynamic cyberthreat driven by political, religious and ideological convictions. Unlike financially motivated cybercriminals, hacktivist collectives leverage digital disruption to challenge nations, influence public opinion and promote specific geopolitical or ideological causes. Denial-of-service attacks, website defacements and data breaches remain their primary operational tactics.

To measure the true scope of hacktivist activity, our intelligence monitoring relies heavily on public messaging channels, predominantly Telegram. Threat actors regularly publish attack claims accompanied by check-host verification links, providing time-stamped proof of targeted availability checks. While check-host reports do not guarantee sustained target downtime, tracking these verified claims offers critical insight into actor intentions, campaign trajectories and changes in operational focus. Data from 2023 through the first half of 2026 reveals an evolving landscape characterized by major quarterly claim surges, heavy geographic concentration in European and Middle Eastern conflict zones, and dominant operational control exerted by pro-Russian threat groups.

Temporal Dynamics and Claim Volatility

Attack claims posted on Telegram frequently get forwarded and reposted on other channels. With the exception of Figure 31, which provides the total number and the unique claims, this section only counts the original claims posted by the responsible hacktivist group. This ensures that only unique attack claims and not the number of reposts or forwards are counted. Figure 31 also provides the number of unique DDoS attack claims per month as well as the total claims across more than 400 Telegram channels and close to 50,000 unique claims targeting more than 10,000 different organizations in 180 countries. The ratio of reposts to unique claims provides a measure of cooperation and affiliation among hacktivists who share each other's exploits.

Public hacktivist attack claims exhibit extreme quarter-over-quarter volatility, fluctuating in close alignment with major real-world geopolitical developments. Total claimed attacks peaked in the second quarter of 2025 at 16,214 claims, representing a 27.5% increase over the previous quarter. Unique claims during this same peak reached 5,067, reflecting a 21.8% rise over the previous quarter.

Beginning mid-2025, public hacktivist activity entered a sustained multi-quarter contraction. Total attack claims declined systematically from the Q2 2025 peak down to 6,190 claim posts in Q2 2026, marking a 38.3% decline from Q1 2026 alone. Unique claims followed a matching downward trend, falling to 1,982 in Q2 2026.

Figure 31:

DDoS attacks claimed per quarter on Telegram (source: Radware)

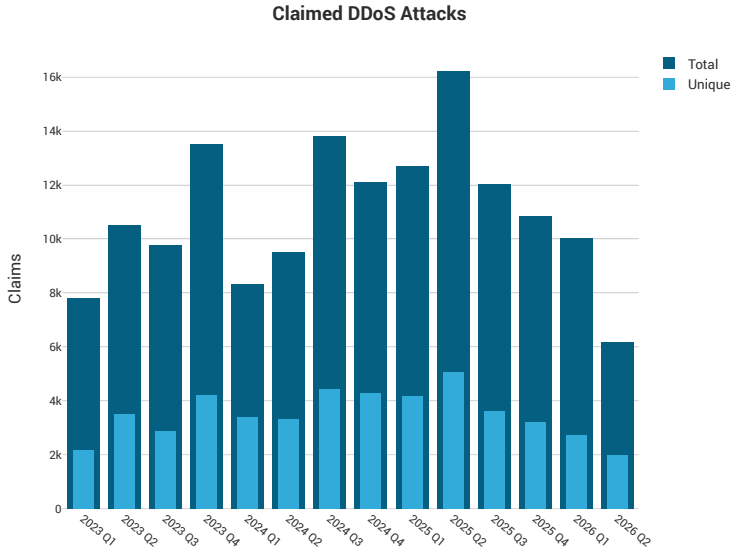

Granular monthly tracking of unique claims shows sharp operational bursts, often confined to single four-week windows:

(1) October 2023 Spike: Unique claims rose 58.8% month-over-month to 2,037, driven by the Hamas-led attacks on Israel on October 7, 2023, and the immediate outbreak of the Israel-Hamas war. This major kinetic conflict triggered an ad-hoc, global mobilization of pro-Palestinian, pro-Israeli and state-aligned hacktivist groups. The sudden explosion of claims was driven by coordinated Telegram campaigns targeting emergency alert systems, government portals, utility providers and news organizations in Israel, as well as counterattacks against Palestinian and Iranian web infrastructure.

(2) August and October 2024 Surges: The August 2024 spike of 1,807 unique claims followed regional escalations in the Middle East, specifically the high-profile targeted assassinations of Hamas political leader Ismail Haniyeh in Tehran and Hezbollah commander Fuad Shukr in Beirut. Simultaneously, Ukraine's surprise cross-border invasion of Russia's Kursk region in August 2024 mobilized pro-Russian collectives like NoName057(16) to retaliate against European and NATO infrastructure. The October 2024 peak of 1,800 claims was driven by Iran's October 1 missile barrage against Israel, the formal launch of Israel's ground campaign in southern Lebanon and symbolic hacktivist operational surges surrounding the first anniversary of the October 7 attacks.

(3) Mid-2025 Highs: The sustained high volume from March through June 2025, averaging more than 1,650 claims per month, reflects multi-front, sustained offensive campaigns across both the European and Middle Eastern theaters. During this window, pro-Russian hacktivist coalitions conducted systematic, continuous DDoS operations against European Union member states and NATO partners supporting Ukrainian defense initiatives. Concurrently, pro-Iranian and anti-Western collectives maintained persistent pressure on Israeli and U.S. government, financial and cloud infrastructure. The dramatic 34.9% drop in July 2025 marks a period of operational fatigue, botnet infrastructure disruption and tactical realignments following these prolonged spring campaigns.

(4) March 2026 Rebound: Following low winter baseline activity, March 2026 experienced a sudden 103.1% month-over-month surge to 1,308 unique claims that was directly ignited by the launch of Operation Epic Fury (also known as Operation Roaring Lion) on February 28, 2026. The joint U.S.-Israeli military and cyber campaign targeted Iranian command infrastructure and state leadership. In immediate response, around 60 hacktivist collectives, including Iranian state-aligned personas, anti-Western groups, and pro-Russian allies, activated widespread retaliatory DDoS campaigns. These campaigns targeted critical infrastructure in Israeli, U.S. and regional Gulf states, driving a dramatic resurgence in verified attack claims.

Figure 32:

Unique DDoS attack claims per month (source: Radware)

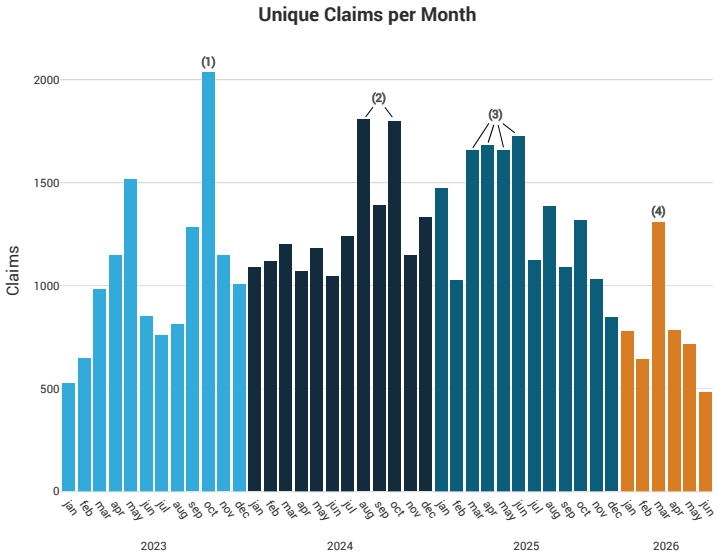

Geographic Target Distribution

The geographic targeting of hacktivist DDoS claims shows a concentrated focus on Europe and the Middle East, reflecting major regional geopolitical tensions. Europe consistently absorbs the majority of global hacktivist activity, accounting for almost half (48%) of all claims in the first half of 2026. Following an intense period of instability in the Middle East during the first half of the year, the region accounted for more than a quarter (26.1%) of all hacktivist activity. Asia ranks as the third-largest regional target at 8.6%, a notable decline from its 17.5% share in 2025. Last year's spike was driven by a convergence of escalating regional disputes, spillovers from Middle Eastern conflicts and direct threats to Western-aligned Asian allies. North America accounted for 8.2%, while Africa, Latin America and Oceania combined represent just below 10% of the overall targeting volume.

Figure 33:

H1 2026 most targeted regions (source: Radware)

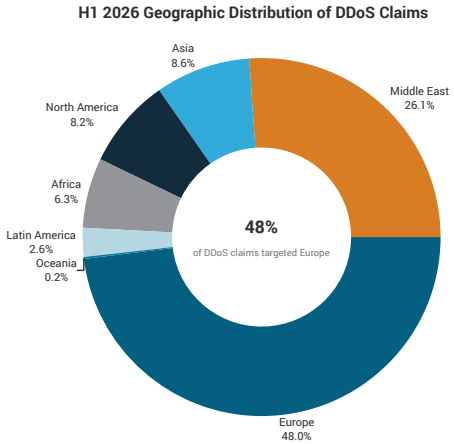

Historical normalized tracking from 2022 to 2025 shows Europe continuously maintaining more than half of all global hacktivist DDoS focus, while Asia and the Middle East have emerged since 2023.

Figure 34:

Evolution of geographic distribution of DDoS claim targets (source: Radware)

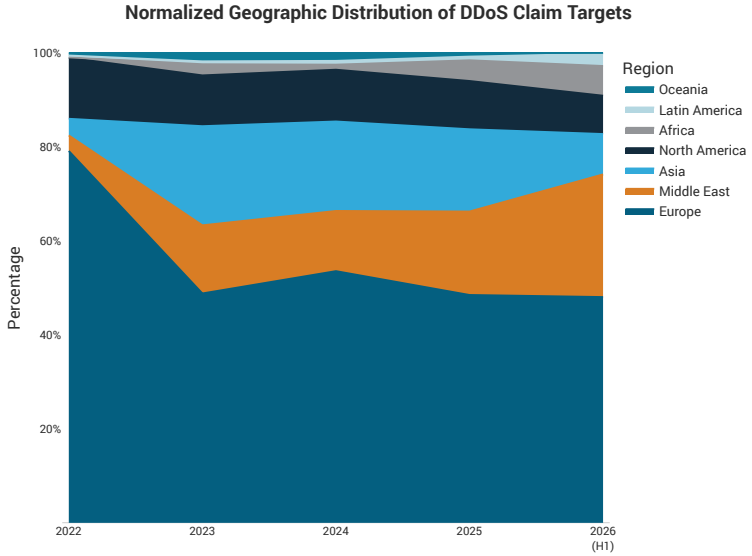

Country-level analysis reveals that state entities and private organizations in nations directly involved in active regional disputes bear the brunt of hacktivist campaigns.

Israel was the most heavily targeted nation during the first half of 2026, absorbing 784 recorded claims representing 16.9% of the global activity. Ukraine stood as the second most targeted nation with 390 claims (8.4%), followed by the United States with 359 claims (7.7%) and the United Kingdom with 285 claims (6.1%). European nations supporting Ukrainian defense efforts, including Denmark, Italy, Germany, Poland and France comprise a significant secondary tier of targeted states.

Figure 35:

H1 2026 most targeted countries (source: Radware)

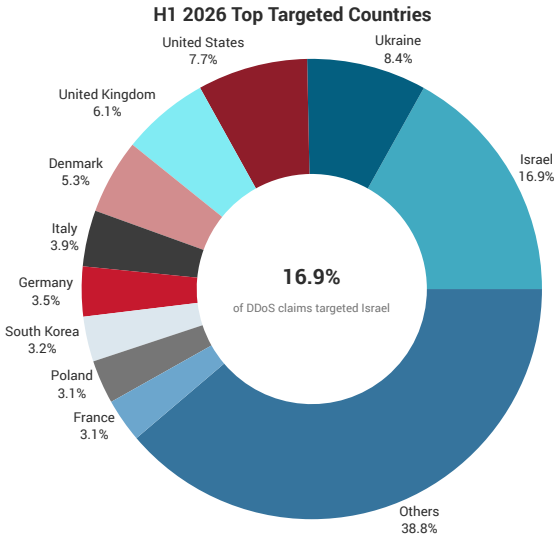

Figure 36:

H1 2026 most targeted countries (source: Radware)

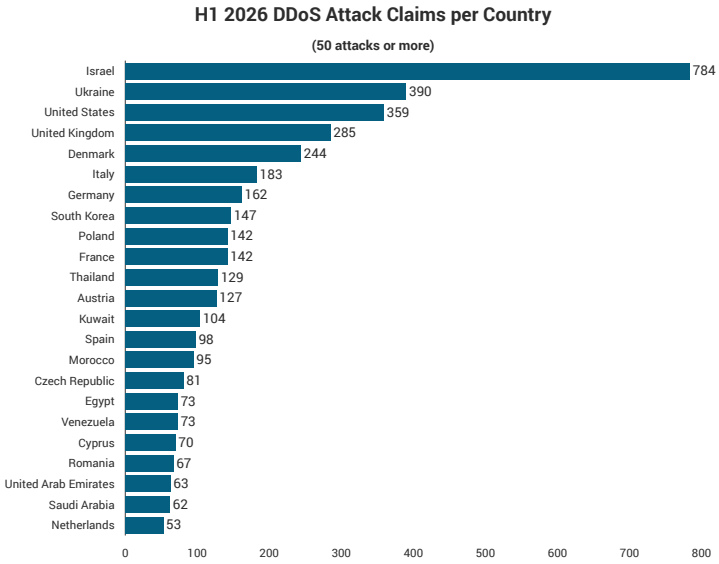

Targeted Infrastructure and Autonomous Systems

Hacktivists intentionally select industry targets to maximize public visibility, disrupt essential public services and pressure political leadership. Government institutions are the primary focus of hacktivist DDoS campaigns, receiving 1,426 claims and accounting for 37.2% of all targeted industry activity. Manufacturing represents the second-most targeted vertical at 7.8%, followed by business services at 6.5%, hospitality at 5.6%, transportation at 5.5% and finance at 5.2%.

Figure 37:

H1 2026 most targeted industries (source: Radware)

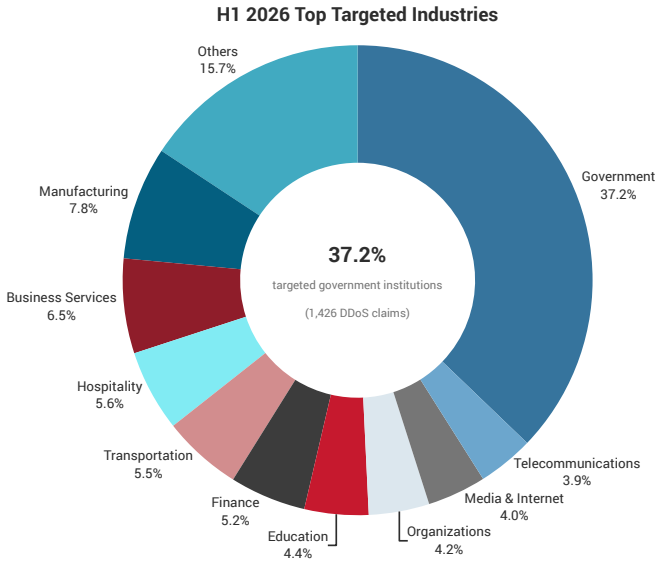

Target distribution across autonomous system numbers (ASNs) highlights heavy exposure among hyper-scale cloud providers, major European web hosts and regional telecommunications networks.

Figure 38:

H1 2026 most targeted ASO (source: Radware)

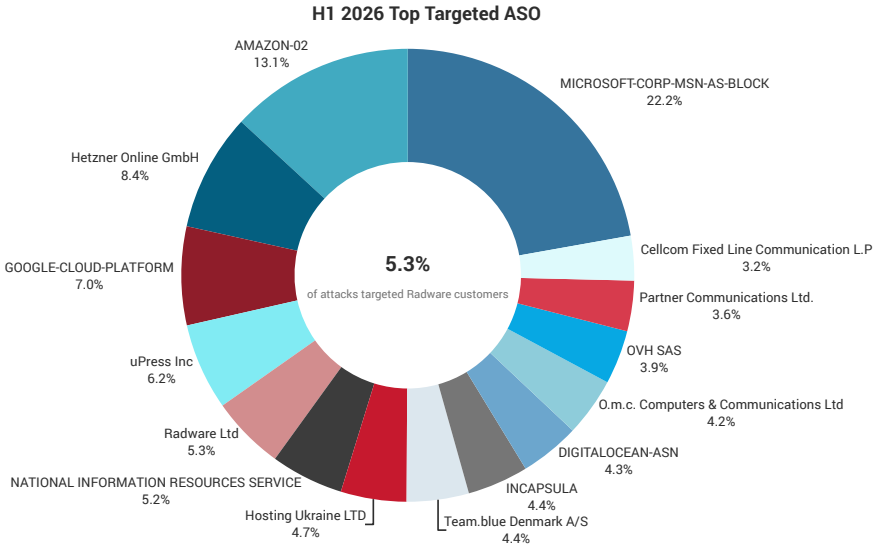

Microsoft infrastructure absorbed 22.2% of all targeted attacks, followed by Amazon Web Services at 13.1% and Google Cloud Platform at 7%. European hosting providers such as Hetzner Online GmbH (8.4%) and regional infrastructure operators supporting Ukrainian and Israeli networks, including uPress, Radware, Hosting Ukraine and Cellcom, accounted for substantial secondary target volumes. Radware-protected infrastructure and websites were the target of 5.3% of all global attack claims.

Threat Actor Groups and Campaign Operations

The hacktivist threat environment is heavily dominated by a small number of organized, pro-Russian and aligned threat actors that coordinate large-scale attack campaigns.

NoName057(16) remains the most active hacktivist group globally by a wide margin, responsible for 1,857 claims accounting for 40.5% of all recorded claims in the first half of 2026. Pro-Russian and anti-Western collectives, including DieNet (8.4%), Keymous+ (7.8%), 313 Team (5.9%), Conquerors Electronic Army (5.1%), Hider Nex (5.1%), RipperSec (4.7%), RuskiNet (3.3%), Serverkillers (2.6%) and BD Anonymous (1.9%), comprise the overwhelming majority (85.3%).

Figure 39:

H1 2026 top claiming hacktivist threat groups (source: Radware)

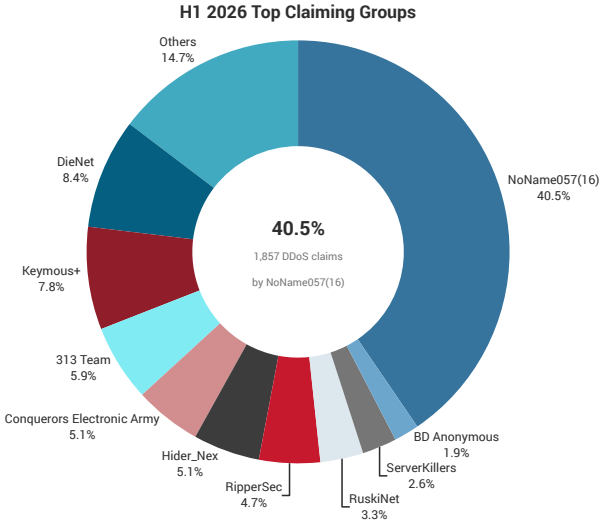

Hacktivist operations organize around dedicated operation hashtags (#OPs) on Telegram and social platforms to coordinate target selection and track campaign milestones.

The most frequent campaign hashtags center on active regional military conflicts, led by #OPUKRAINE with 557 mentions and #OPISRAEL with 469 mentions in the first six months of 2026. Secondary operations target European countries based on political declarations, aid packages or diplomatic summit meetings, including #OPITALY (361 mentions) and #OPDENMARK (308 mentions).

Figure 40:

H1 2026 top 20 used operation hashtags (source: Radware)

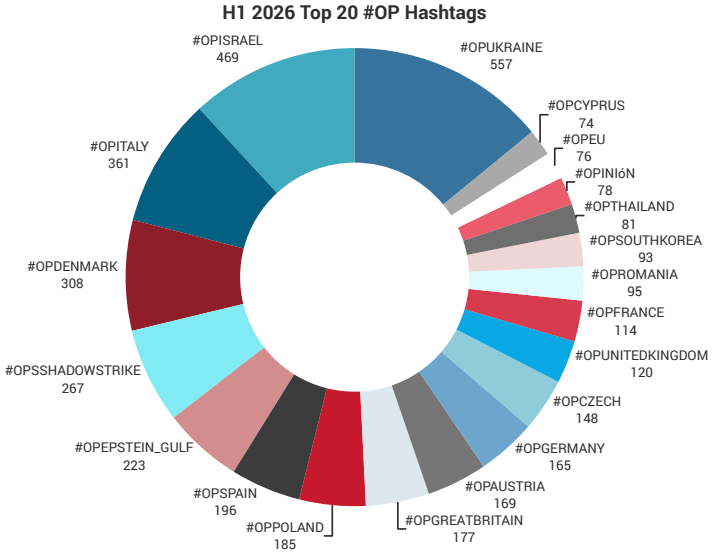

Tracking monthly campaign hashtag mentions throughout the first six months of 2026 illustrates the rapid escalation and dissolution of targeted operations. Mentions of #OPITALY surged sharply in February 2026 to nearly 180 mentions before quickly subsiding. Similarly, #OPISRAEL experienced a massive peak in March 2026, reaching over 250 mentions before declining sharply in April and May. In contrast, #OPUKRAINE maintained a steady baseline throughout the first half 2026 while #OPDENMARK reached peak mentions in February and March before trending down through Q2 2026.

Figure 41:

Evolution of top operation hashtags in H1 2026 (source: Radware)

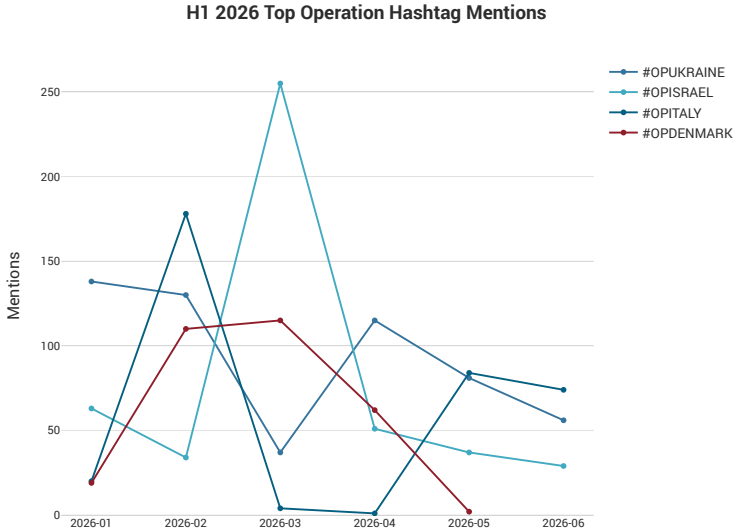

Inside the 2026 Radware Cyber Survey

New Trends in AI, Application and API Security

Radware explores the state of AI, application and API security through a survey of 377 randomly selected organizations conducted by Osterman Research. The resulting research paper, commissioned by Radware and released in July 2026, includes the following findings:

Organizations Are Creating a New Category of Shadow AI

AI agents and autonomous workflows stand out as major areas of growth, with the anticipated implementation timeframe running at about twice the rate of adoption of GenAI or LLM functionality. According to the survey, 77% of organizations are already using or implementing AI agents or autonomous workflows (see Figure 42).

Figure 42:

Plans for AI agents or autonomous workflows (source: Osterman Research 2026)

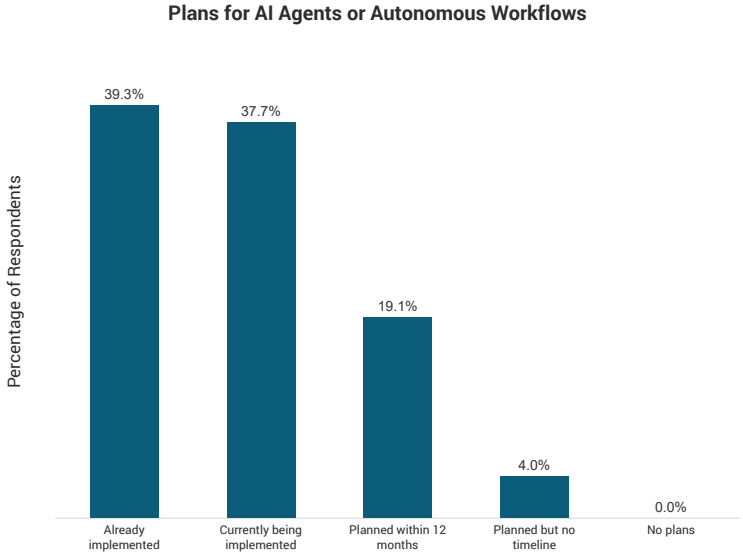

Visibility, however, significantly lags deployment cadence. Visibility into AI agents and AI-driven processes means having the technical ability to detect which agents and processes are running, which systems they are accessing (and with what credentials), and the actions they are taking autonomously. The most common answer in Figure 43 is “high visibility” (about 75%), cited by 58.4% of respondents. While this sounds promising, AI agents and AI-driven processes that hide in the remaining 25% of invisibility represent a significant, unquantified threat surface.

Figure 43:

Visibility into AI agents and AI-driven processes (source: Osterman Research 2026)

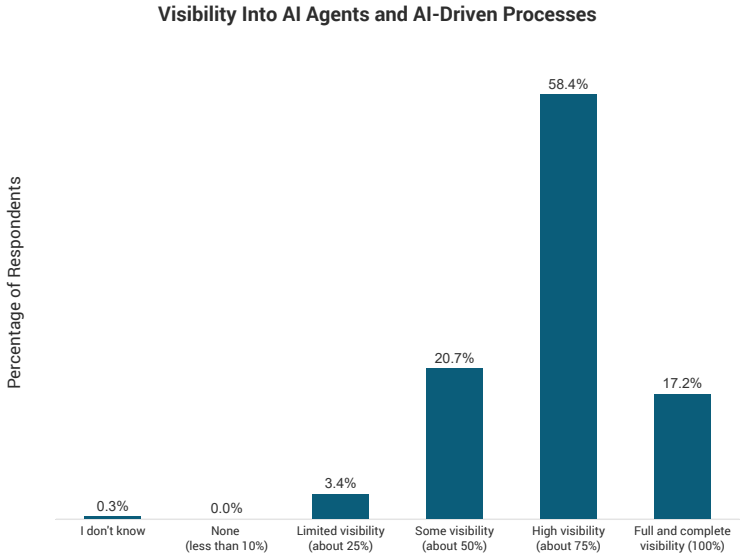

Only 17.2% of respondents say they have “full and complete visibility” (100%) of the AI agents or AI-driven processes operating in their organization, including internally developed and external agents. With 77% of organizations already using or implementing AI agents or autonomous workflows, but only 17.2% having complete visibility into these, organizations are already lagging significantly in security versus deployment cadence.

Shadow AI Is on the Rise

Shadow IT refers to employees using software or cloud services without the knowledge or approval of their central IT department. In 2026, the bigger challenge appears to be shadow AI.

Organizations face employees and application developers who deploy AI services and AI agents they cannot fully inventory. These shadow AI agents pose increased risk because they are decision-making entities that can access corporate and external services, use real user credentials, call APIs and execute workflows without tracking from corporate IT.

The visibility gap is particularly dangerous because AI agents are not static assets. Even a small percentage of shadow agents can generate thousands of actions across applications and APIs. Furthermore, every single transaction an agent performs is a legitimate transaction. The threat is hidden in the sequence of transactions and the intent, which are performed at wire speed.

API Security Has an Operational Scalability (Not Technology) Problem

APIs are a core part of modern applications, and the use of internally developed APIs has grown considerably over the past 12 months. More than 70% of respondents indicate that the use of home-grown APIs has increased significantly or moderately in 2026 (Figure 44). The share of organizations whose API usage stayed about the same fell threefold, from 24.1% in 2025 to 8.8% in 2026.

Figure 44:

How usage of APIs developed in-house changed over the past 12 months (source: Osterman Research 2026)

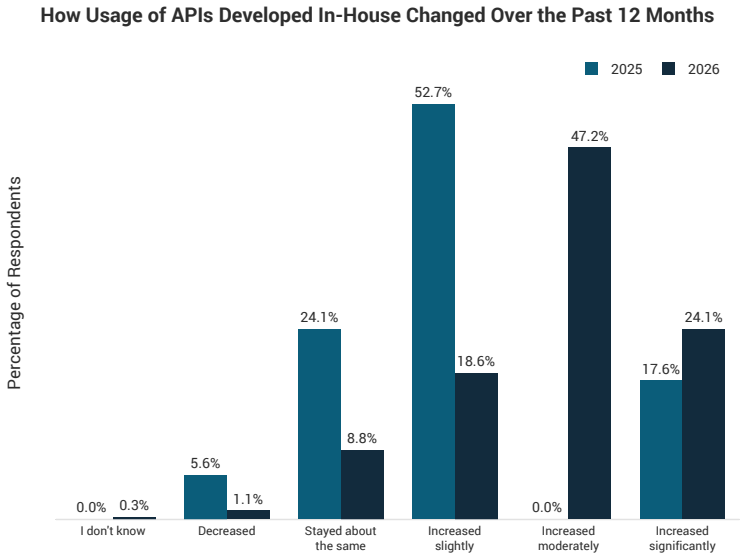

The next question asked organizations how often they update APIs for production use. The share that updates APIs at least weekly continues to increase, rising from 72.7% in 2025 to 81.1% in 2026 (Figure 45).

Figure 45:

Frequency of updating APIs for production usage (source: Osterman Research 2026)

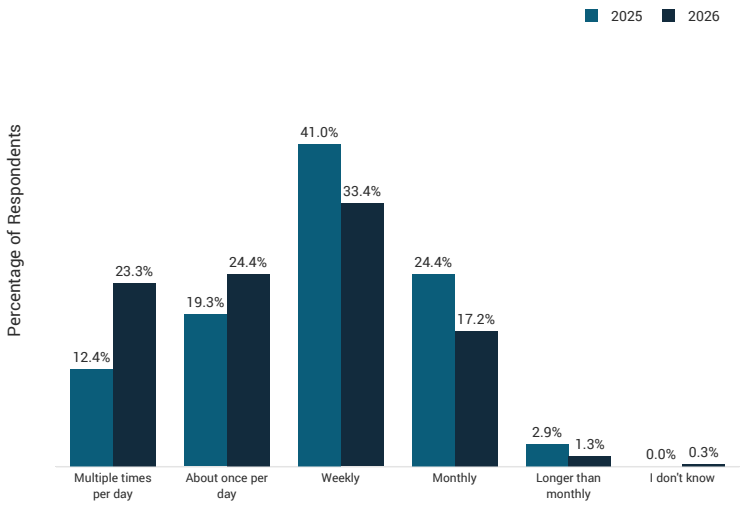

Even as the pace of API updates accelerates, API documentation lags. Only 6.9% of organizations claim to fully document all their APIs, a slight increase from 6.1% in 2025. 43% of organizations document fewer than 70% of their APIs, excluding third-party software APIs. See Figure 46.

Figure 46:

Percentage of documented APIs in the organization (source: Osterman Research 2026)

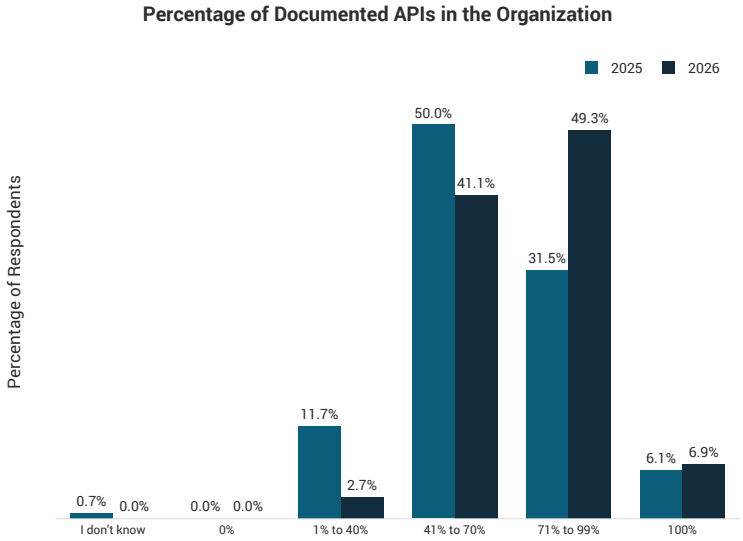

Cyber Survey Summary

With daily API updates becoming common, documentation remains extremely poor (only 6.9% are fully documented). We are back to the visibility challenge: it is impossible to protect what you cannot see. There is also little evidence that organizations lack API security tools. Most WAF solutions include API security capabilities, and over 70% of organizations integrate API testing into application development workflows.

The data suggests organizations cannot scale API security at the pace of application development. API security has become an operational and scalability problem that needs to be addressed by improving application lifecycle workflows. Shadow APIs and zombie APIs continue to grow at the pace of application features development and become a preferred attack surface for cybercriminals. This is backed by the sharp increase in blocked malicious web and API transactions.

Conclusion: The Convergence of Automated Velocity and Structural Blind Spots

The threat landscape in the first half of 2026 marks a fundamental tipping point: the speed, scale and complexity of attacks have permanently outpaced traditional, human-centric defense mechanisms. Driven by aggressive direct-path DDoS campaigns, automated vulnerability exploitation and improved AI agent harnesses, threat actors are operating at unprecedented velocity. At the same time, enterprise defenders struggle with compounding operational visibility gaps.

Key Strategic Dynamics Reshaping the Landscape

The Erasure of the Patch Window: The time between vulnerability disclosure and active exploitation has collapsed into negative territory, averaging negative eight hours, with over 80% of exploited CVEs now occurring as zero days. Advanced AI models, excelling at semantic reasoning and vulnerability chaining, have rendered traditional fuzzing and manual code audits obsolete.

Industrialized and Direct-Path Volumetric Attacks: Network and application threats have escalated to record highs, characterized by massive direct-path UDP floods and a continued surge in Web DDoS and malicious web application and API transactions. Threat actors have moved from indirect reflection toward high-volume, continuous pressure targeting critical infrastructure and tech sectors, with North America absorbing the majority of application-layer attacks.

The Dual Risk of Agentic AI and Supply Chain Amplification: The proliferation of local AI agents with continuous “heartbeats” and autonomous developer tools introduced severe supply chain risks. These agents execute tasks, call APIs and download unvetted dependencies at wire speed using legitimate user credentials. They act as unintentional, automated malware installers that bypass traditional sandbox restrictions and security perimeters.

Democratization of Nation-State Capabilities: Through crime-as-a-service (CaaS) platforms leveraging hijacked frontier models and jailbroken APIs, high-level offensive tactics such as dynamic social engineering, automated payload generation, and multi-step attack execution have become broadly accessible via online subscriptions. Technical skill is no longer a barrier to executing enterprise-grade cyberattacks, and with the rise of open-weight models combined with improved agent harnesses, this trend is only going to accelerate in the near future.

The Governance and Visibility Crisis: Enterprise security governance is failing to keep pace with deployment velocity. With the rise of open-weight models combined with improved agent harnesses, this trend is only going to accelerate in the near future. With 77% of organizations adopting AI agents (yet only 17.2% maintaining full visibility) and 81.1% releasing API updates weekly (while only 6.9% fully document them), organizations are accumulating massive shadow AI and zombie API attack surfaces.

Final Outlook

Security can no longer be solved by simply adding reactive tools or relying on legacy perimeter defense. Because every individual AI agent transaction or undocumented API call appears legitimate on the surface, the threat now hides in the intent, sequence and speed of automated operations. To survive this era of autonomous, AI-driven threats, organizations must urgently adopt intelligent, real-time, intent-based behavior analytics, modernized application lifecycle workflows, automated defenses and automated attack surface management. This allows for end-to-end visibility and protection across all threats.

Appendices

Appendix 1: Calculation of Daily Averages

Averages for annual, semiannual and quarterly periods in this report are calculated using an actual day-count convention, adjusting for leap years to faithfully represent time elapsed.

Breakdown of the number of days taken into account for each year:

Year	Days
2023	365
2024	366 (Leap Year)
2025	365
2026	365

Breakdown of the number of days taken into account for each half-year period:

Year	First half (Jan 1 – Jun 30)	Second half (Jul 1 – Dec 31)
2023	181 days	184 days
2024	182 days (Leap Year)	184 days
2025	181 days	184 days
2026	181 days	184 days

Breakdown of the exact number of days for each quarter:

Year	Q1 (Jan – Mar)	Q2 (Apr – Jun)	Q3 (JUL – SEP)	Q4 (Oct – Dec)	Total Days
2023	90 days	91 days	92 days	92 days	365 days
2024	91 days (Leap)	91 days	92 days	92 days	365 days
2025	90 days	91 days	92 days	92 days	365 days
2026	90 days	91 days	92 days	92 days	365 days

Appendix 2: Methodology and Sources

The data for DDoS events and volumes was collected from Radware devices deployed in Radware cloud scrubbing centers and on-premises managed devices in Radware hybrid and peak protection services, jointly denoted as **Radware's Cloud DDoS Protection Service**. Note that attack events and blocked events are considered the same for this report. All blocked volume is considered attack volume. An attack is a collection of several related attack vectors targeting the same customer and overlapping in time. Events correspond to attack vectors. Attack vectors consist of one or more packets. All packets of an attack vector generate a certain volume expressed in bytes. The volume generated by an attack vector is referred to as the blocked volume for that attack vector, which corresponds to the attack volume for that vector. The attack volume of all attack vectors from the same attack corresponds to that attack's attack volume.

The data for web application attacks and bot activity was collected from blocked application security events from the **Radware Cloud Application Protection Service**. Collected events were based solely on automatically detected and known vulnerability exploits and exclude any events that might be blocked or reported by custom rules added to a web application policy by managed services and/or customers.

Web DDoS attack details were collected from ERT SOC incidents relating to **Radware's Web DDoS Protection Service.**

Hacktivists openly publicize their actions on social media and public Telegram channels to gain media attention and raise awareness. They do not operate covertly or evade the media but instead reveal the names and resources of their targets and attempt to take credit for their attacks. Hacktivists utilize website monitoring tools to demonstrate the impact of their denial-of-service attacks on online resources and frequently share links to reports from online web monitoring tools in their messages. Through tracking and analyzing messages from several active hacktivist groups on Telegram, the **Radware Threat Intelligence** team is able to assess the global DDoS activity conducted by hacktivists.

Table of Figures

Author

Pascal Geenens | VP Cyber Threat Intelligence

Contributions

“Inside the 2026 Radware Cyber Survey” Section:

Ron Meyran, VP Strategic Alliances Marketing & Cyber Threat Intelligence

Executive Sponsors

Ron Meyran | VP Strategic Alliances Marketing & Cyber Threat Intelligence

Deborah Myers | Senior Director of Corporate Marketing

Production

Kimberly Burzynski | Senior Marketing Communication Manager

Jeffrey Komanetsky | Senior Content Development Manager

Mohan G | Senior Graphic Designer

About Radware

Radware® (NASDAQ: RDWR) is a global leader of cybersecurity and application delivery solutions for physical, cloud and software-defined data centers. Its award-winning solutions portfolio secures the digital experience by providing infrastructure, application and corporate IT protection and availability services to enterprises globally. Radware's solutions empower more than 12,500 enterprise and carrier customers worldwide to adapt quickly to market challenges, maintain business continuity and achieve maximum productivity while keeping costs down. For more information, please visit www.radware.com.

RWI6424 | 08.17.2026