radware

H1 2026 Global Threat Analysis Report

Navigate the Latest Network, Application and AI Attack Trends



H1 2026 Global Threat Analysis Report

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)

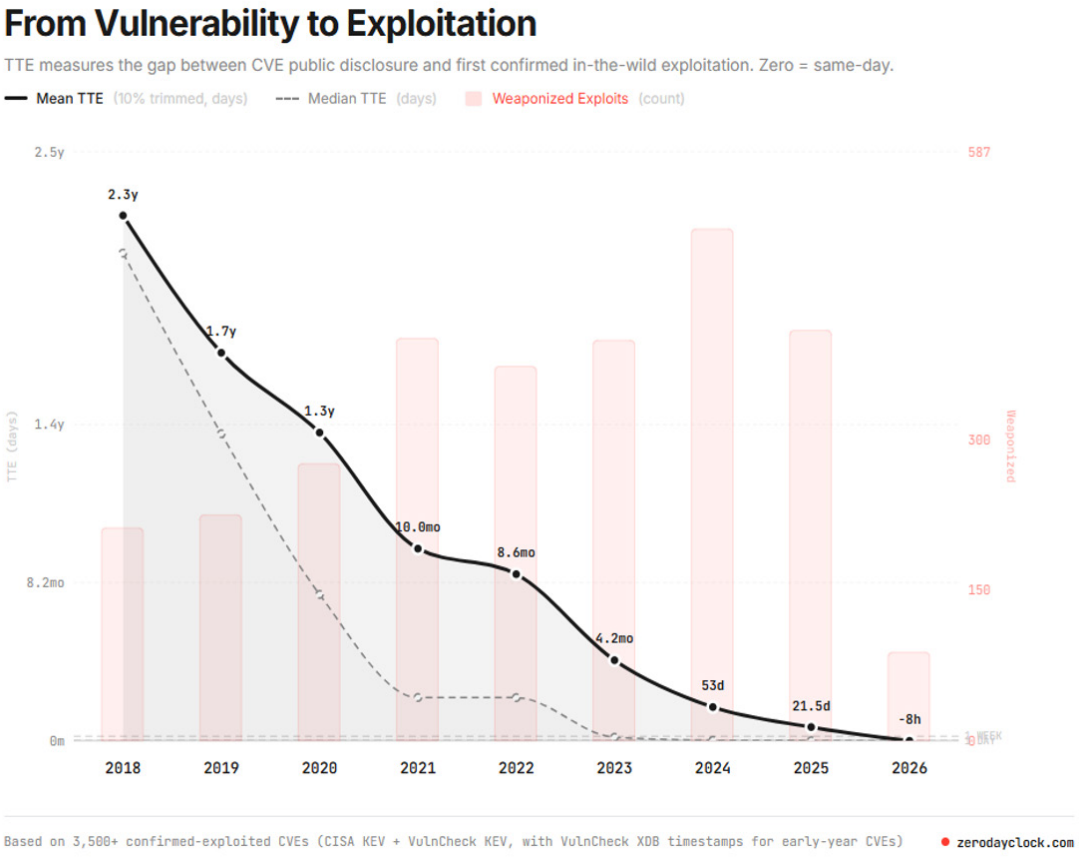
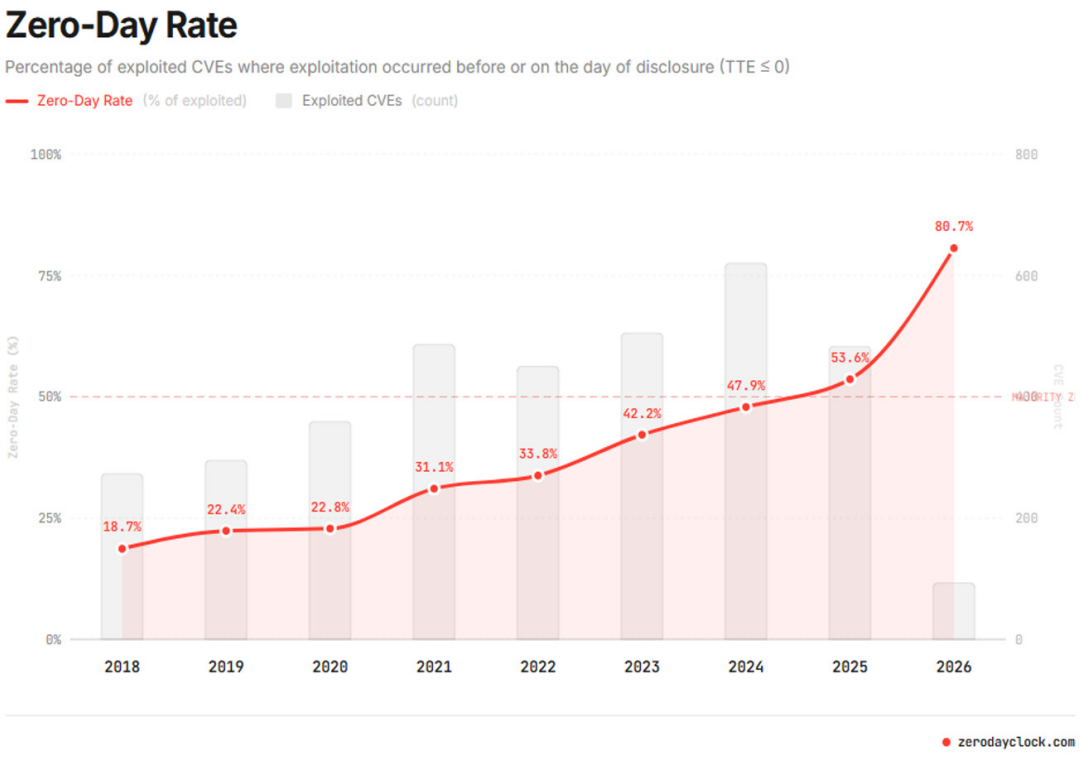


Figure 2:

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.

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.

THIS REPORT CONTAINS ONLY PUBLICLY AVAILABLE INFORMATION AND IS PROVIDED FOR GENERAL INFORMATIONAL PURPOSES ONLY. ALL INFORMATION IS PROVIDED "AS IS" WITHOUT ANY REPRESENTATION OR WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES THAT THIS REPORT IS ERROR-FREE OR ANY IMPLIED WARRANTIES REGARDING THE ACCURACY, VALIDITY, ADEQUACY, RELIABILITY, AVAILABILITY, COMPLETENESS, FITNESS FOR ANY PARTICULAR PURPOSE OR NON-INFRINGEMENT. USE OF THIS REPORT, IN WHOLE OR IN PART, IS AT USER'S SOLE RISK. RADWARE AND/OR ANYONE ON ITS BEHALF SPECIFICALLY DISCLAIMS ANY LIABILITY IN RELATION TO THIS REPORT, INCLUDING WITHOUT LIMITATION, FOR ANY DIRECT, SPECIAL, INDIREC, INCIDENTAL, CONSEQUENTIAL, OR EXAMPLARY DAMAGES, LOSSES AND EXPENSES ARISING FROM OR IN ANY WAY RELATED TO THIS REPORT, HOWEVER CAUSED, AND WHETHER BASED ON CONTRACT, TORT (INCLUDING NEGLIGENCE) OR OTHER THEORY OF LIABILITY, EVEN IF IT WAS ADVISED OF THE POSSIBILITY OF SUCH DAMAGES, LOSSES OR EXPENSES. CHARTS USED OR REPRODUCED SHOULD BE CREDITED TO RADWARE.

H1 2026
Global Threat Analysis Report

Download Full Report



This document is provided for information purposes only. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law. Radware specifically disclaims any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. The technologies, functionalities, services, or processes described herein are subject to change without notice.

© 2026 Radware Ltd. All rights reserved. The Radware products and solutions mentioned in this document are protected by trademarks, patents and pending patent applications of Radware in the U.S. and other countries. For more details, please see: https://www.radware.com/LegalNotice/. All other trademarks and names are property of their respective owners.

RWI6424 | 08.17.2026