State of the Internet | V12 Issue 06

Akamai Security

SOTI

AI-Driven Digital Deception

The Fraud and Abuse Report 2026



Contents

03 **Introduction**

04 **Expert insight: Integrating cyber and fraud defense in the agentic AI age**
(Patrick Sullivan, Chief Technology Officer, Security Strategy, Akamai)

06 **Key insights from the report**

07 **Protect your account integrity**

11 Specific fraud and abuse schemes by industry

12 High- or critical-risk login attempts by industry

14 Login activity and attack schemes by region

16 Regional and industry-specific account abuse schemes in action

18 **Protect your brand**

20 Regional trends reveal notable fraud and abuse patterns

25 **Threats in depth: High-profile sporting event spoofed in coordinated ticket impersonation attack**
(Moshe Cohen, Principal Security Analyst, Akamai, and Shiran Guez, Director of Engineering, Akamai)

27 **Beyond traditional phishing: Online scams and fraud**

31 Regional campaigns: Attackers build global, act local

32 **Security spotlight: Ransomware is about extortion, *not* encryption**
(Ofer Wolf, Senior Vice President and General Manager of Enterprise Security, Akamai)

36 **Expert insight: Agentic workflows: A new frontier for compliance**
(James A. Casey, Vice President and Chief Privacy Officer, Akamai)

39 **Mitigation strategies**

41 **Conclusion**

42 **Methodologies**

43 **Guest contributors**

45 **Credits**

Akamai Security

Introduction

When the State of the Internet team examined online fraud and abuse a year ago, our focus was squarely on the initial impacts of AI bots and the implications of the dramatic rise in automated web traffic. We still define *fraud* as the use of deception for unlawful gain, and *abuse* as the engagement in unauthorized activities. However, the financial stakes have sharply escalated over the last year and, in response, we have expanded our scope in this report to address the broader, multichannel ecosystem of digital deception, in which threat actors abuse brand trust, exploit the entire account lifecycle, and pivot toward these high-loss financial vectors.

An organization's attack surface is no longer bounded by its network perimeter. Today, cybercriminals bypass traditional access controls to target a company's most critical asset: brand trust. By leveraging recognizable logos, spoofing workflow tools, and deploying fake applications across social media and search, adversaries turn customer confidence into an entry point for brand abuse and systematic account takeover. Once credentials and identities are compromised through these deceptive channels, attackers hijack legitimate accounts to drain funds, steal sensitive data, and further exploit customer relationships from the inside out.

In this edition of the Akamai State of the Internet/Security report, we combine our global data insights with crossplatform threat intelligence from Bolster AI to deliver actionable insights for key operational teams that spend their days focused on the fraud and abuse problem. These insights include:

- How phishing as a service (PhaaS) and rapid infrastructure rotation are shrinking response windows and forcing back-end signal detection
- Why online scams are eclipsing traditional phishing by mapping risk across the account lifecycle from SMS pumping to point draining
- How look-alike domains, recruitment scams, and spoofed event portals erode customer trust and damage revenue
- Why low-volume, high-impact threats like ransomware leverage initial access vectors to target high-value assets and drive staggering financial losses

Whether you are defending enterprise networks, securing user accounts, or safeguarding brand reputation, this report provides data-backed analyses and strategies to help you stay ahead of automated, sophisticated fraud and abuse campaigns.

Akamai Security

Expert insight

Integrating cyber and fraud defense in the agentic AI age

The fight against fraud was once a siloed function, managed within finance or disparate business units. However, as online fraud evolved into a major revenue risk, fusion centers that integrate cybersecurity and fraud emerged. The concept was borrowed from government agencies, which first created fusion centers to break down walls and encourage organizations to share intelligence, resources, and expertise to address national threats.

Commercial fraud fusion centers bring diverse teams together to consolidate cyber and fraud operations. These centers unify data, processes, and authority into a cohesive capability to detect, investigate, and mitigate attacks across enterprise systems and customer transactions. Whether you use the term "fusion center" or not, security is a team sport — and in the era of the agentic web, a crossfunctional operating model is essential to fight fraud.

The evolving CISO role

When CISOs became part of the fraud solution they took on the responsibility of preventing fraud at the edge, while the fraud team stayed focused on supporting impacted customers. The CISOs' daily focus also shifted from purely technical risk to business outcomes — specifically safeguarding online revenue — and their security priorities changed as online fraud grew.

Within this new generation of fraud defense, CISOs initially tended to focus on stopping malware and compromise of systems to protect personally identifiable information (PII) and the payment card industry. Then, many prioritized stopping known fraud techniques to protect customer credentials and accounts.

CISOs' priorities continue to evolve as business models continue to change. Today, CISOs face new AI-driven agentic autonomous hacking attacks (large language model [LLM], agentic, and Model Context Protocol [MCP] endpoints) that are forcing them to rethink their role again. Autonomous AI agents are redefining how users and attackers interact with enterprise infrastructure. Navigating this new era requires CISOs to defend an increasingly complex, agentic attack surface.

The agentic web and expanding complexity

The agentic web has created new issues like the zero-click web, where your customers interface with your brand based on how you show up in LLM responses. When AI acts as the intermediary between people and company websites, it means fewer direct human visits and more bot/agent activity. Cybercriminals are starting to take advantage of this shift to conduct fraud and abuse schemes like zero-click data exfiltration via MCP. Customers are buying shoes, booking travel, or applying for a loan directly through an AI platform.

Cybercriminals are following the same path by using synthetic identities to impersonate legitimate customers. Interactions between these synthetic customers and brands introduce risk because it's increasingly difficult to distinguish legitimate customer traffic from malicious activity. Threats like business logic abuse, agent hijacking, and synthetic identity fraud have become critical concerns.

The agentic web also marks a shift to next-generation fusion centers, where additional coalitions between CISOs, marketing, sales operations, customer experience, vendor management, IT, fraud, and compliance teams are forged to combat new AI-driven attacks. For most companies, the working agreement should include the following core components:

- **Executive sponsorship/steering committee.** Appoint the CISO, head of fraud, and CMO to champion the collaboration, set the direction, and remove roadblocks.
- **Balanced security and digital experience.** Protect data while minimizing friction at critical touchpoints like logins, payments, and account creation.
- **Cross-team visibility.** Understand how attack schemes impact different functions like finance, marketing, human resources, and brand erosion.
- **Uniform metrics for ROI.** Establish fraud tracking baselines that align with strategic business objectives.
- **Clearly defined responsibilities.** Use a RACI (responsible, accountable, consulted, and informed) matrix to define responsibilities for processes and investigations, as well as to enable shared data and common playbooks.
- **Gap analysis and a prioritized plan.** Identify gaps in fraud defense and prioritize investments based on risk (i.e., protect against API abuse, then MCP abuse as an emerging vector with agentic AI).

Strategic recommendations for security leaders

The cybersecurity leader's role must evolve again to defend against AI-driven attacks: Cybersecurity and risk leaders should assume that autonomous AI interactions will become a normal part of digital business and plan security capabilities and governance accordingly. The security roadmap should focus on:

- Embracing agentic readiness by broadening your bot management strategy to include visibility and management of synthetic identities
- Managing and protecting APIs against account takeover, payment fraud, and other abuse attack schemes, and gaining visibility into MCP endpoints
- Expanding data integration within the fusion center to include digital, marketing, and agentic transactions
- Building governance and KPIs across security, fraud, digital, marketing, and customer experience to eliminate visibility gaps and cover the scope of a next-gen fusion center

Overlapping tools, unintegrated workflows, and conflicting metrics result in missed opportunities to prevent fraud. The security team sees the attack. The fraud team sees losses. The marketing team sees trust erosion. No single team has full visibility to connect the dots. A next-gen fraud fusion center provides the visibility and protection you need to secure the new digital frontier.



Patrick Sullivan
Chief Technology Officer,
Security Strategy, Akamai

Akamai Security

Key insights from the report

- The login phase is the most active stage of the account lifecycle journey and the entry point for malicious bot activity, so there is significant value in monitoring login activity. North America (N. Amer.) experienced the highest percentage of login attempts at 45.8%, while Asia-Pacific (APAC) and Latin America (LATAM) experienced the highest proportion of high- or critical-risk level login attempts at 36.4% and 37.8%, respectively.
- Account verification requests made up 40% of high/critical login traffic, demonstrating that adversaries may be validating whether email addresses map to legitimate accounts before launching credential stuffing attacks.
- Among all industries, commerce had the highest concentration of login attempts flagged as high or critical.
- The phishing attack lifespan declined by more than 50%: from 48 days (Q4 2025) to 22 days (Q2 2026). This leaves security teams with less time to detect and disrupt campaigns before attackers rotate infrastructure. This undermines traditional defenses.
- Online scams approximately doubled in volume, surpassing traditional phishing for the first time in Q1 2026, surging from ~3.2 million to ~ 6.4 million, marking a significant shift in threat tactics.
- Adversaries rely heavily on trusted brand identity to gain access and harvest credentials. The majority of the malicious domains that targeted enterprise networks during the reporting period (January 2026 through July 2026) impersonated high tech (41.6%), financial services (18.9%), and commerce (8.4%) brands.

Protect your account integrity

Fraud can occur at every stage of the customer journey, from the moment of account creation and authentication through password changes, account updates, and payment activity (Figure 1). When left unaddressed, fraud not only puts customers at risk but also drives financial losses, increases support costs, and introduces reputational damage and potential legal or compliance consequences for organizations.

Risks Across the Entire Account Lifecycle

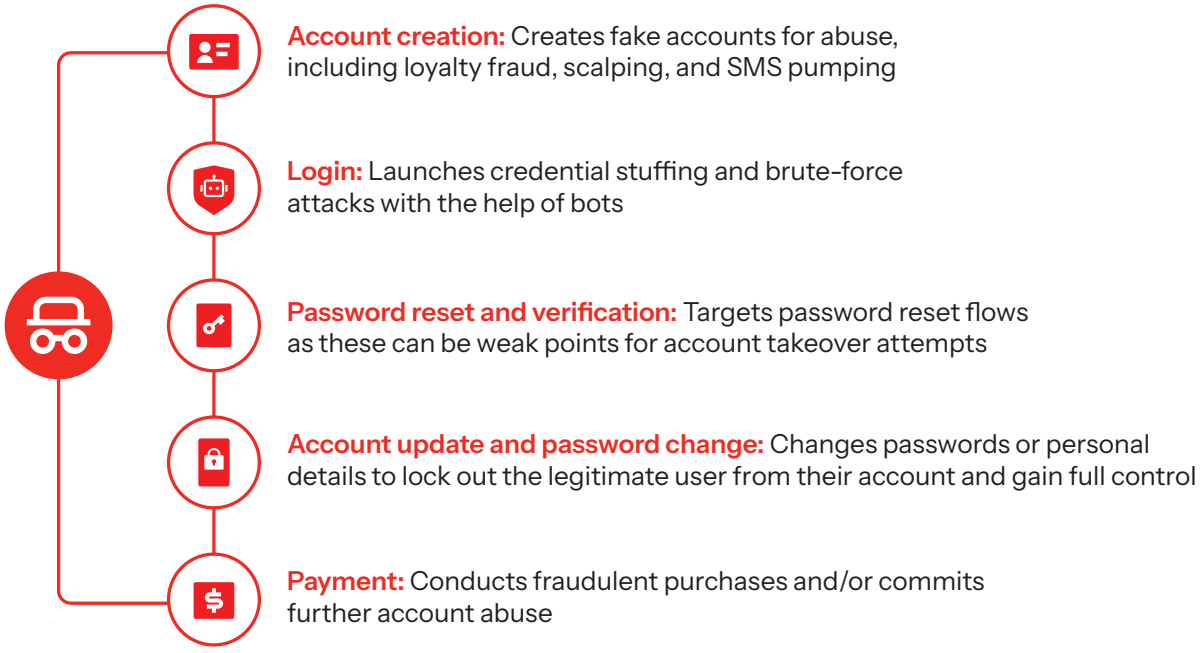


Fig. 1: Account lifecycle stages and associated risks to consumers

The consumer lifecycle often begins with account creation; attackers generate large volumes of fraudulent accounts to enable abuse at scale. These bogus accounts serve as the foundation for verifying stolen payment credentials through low-value "test" transactions before larger purchases are attempted. This fraudulent activity can yield chargebacks, processing fees, and penalties for affected organizations. With 8.3% of account creation attempts linked to suspected fraud in 2025, this phase has become one of the most vulnerable stages in the account lifecycle.

In the commerce industry, for example, the motivation behind fake accounts extends beyond payments. Fraudsters use them to reserve or hoard high-demand and/or limited-edition inventory, effectively reducing availability for legitimate customers and enabling resale at inflated prices. This directly affects revenue while simultaneously degrading customer experience and eroding trust in the brand.

As accounts accumulate, adversaries systematically exploit promotions, discounts, and loyalty programs designed to drive engagement and to maximize value. Harvested credentials fuel credential stuffing and brute-force attacks, allowing attackers to scale account takeover attacks efficiently. In 2025, the FBI's Internet Crime Report indicated a whopping US$359.7 million in losses due to account takeover attacks.

After gaining access, attackers establish persistence by modifying account details and recovery options, making remediation more difficult. Compromised accounts are then monetized through numerous fraudulent transactions (such as draining rewards and loyalty points) and continued exploitation of payment methods and promotional incentives. Beyond immediate financial losses, organizations may face chargebacks, reimbursement claims, and reputational damage, compounding the overall fraud impact.

Analyzing login activities

Organizations should closely monitor login activity not only to distinguish legitimate behavior from suspicious activity, but also because login activity can provide early warnings of an active threat. According to a 30-day sample data from Akamai Account Protector, high-risk and critical-risk login activities together accounted for 26.6% of all recorded attempts, with critical-risk attempts alone reaching 18.0% (Figure 2).

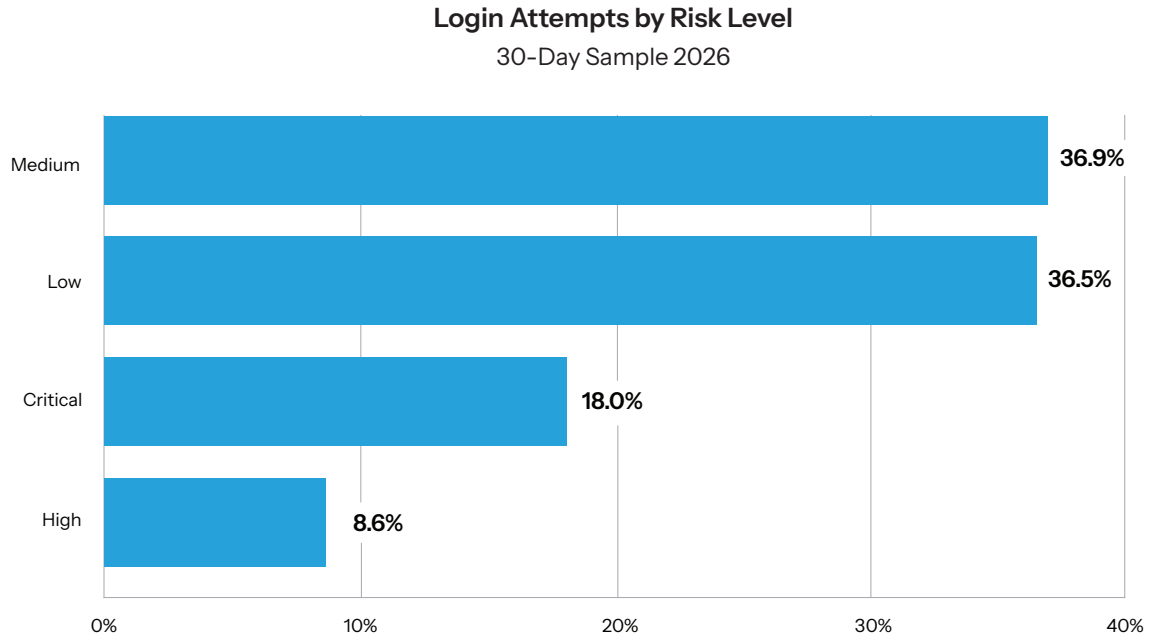


Fig. 2: More than 25% of login attempts are at the high- or critical-risk level

A high or critical risk level suggests a likelihood of brute-force attacks and credential stuffing attack attempts and requires defenders to investigate further to prevent successful account takeover. By contrast, the low- and medium-risk login attempts represent standard day-to-day user behavior from a known user device or IP.

Assessing login objectives

The login phase is the most active stage and generates the majority of requests across the entire account lifecycle. As such, we analyzed these login attempts to assess intent and determine whether patterns differed across all risk tiers or only in the high- or critical-risk categories. Across all risk levels (from low to critical), standard login attempts have the largest share of activity at 63.6%, followed by account verification requests at 16.1% (Figure 3).

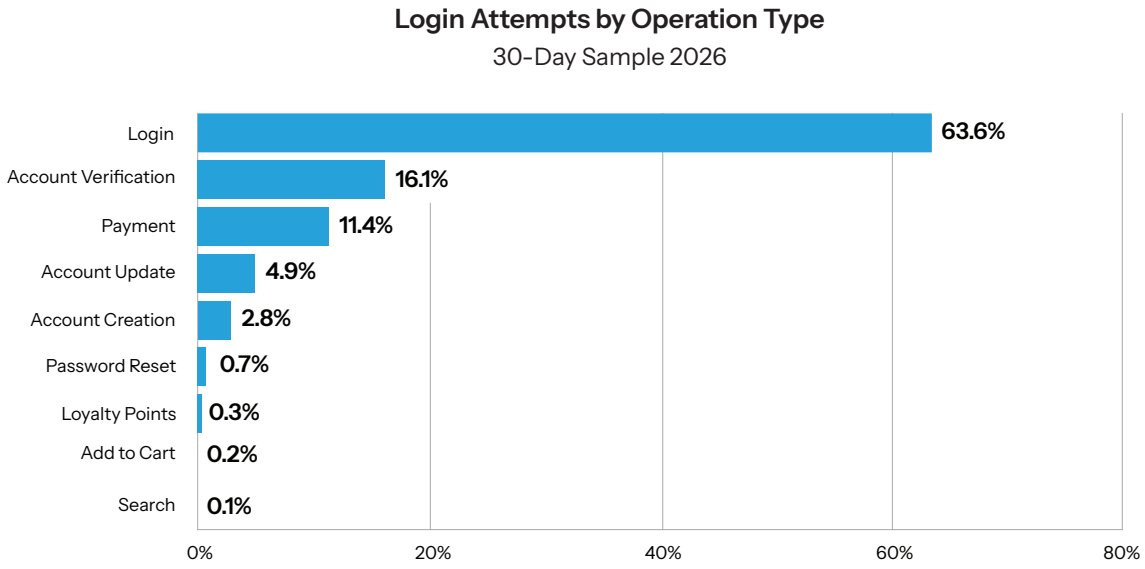


Fig. 3: Logins dominated the total activity regardless of risk level at 63.6%

When isolating login attempts categorized as high or critical risk, account verification requests ranked first at 40.0%, closely followed by logins at 36.2% (Figure 4).

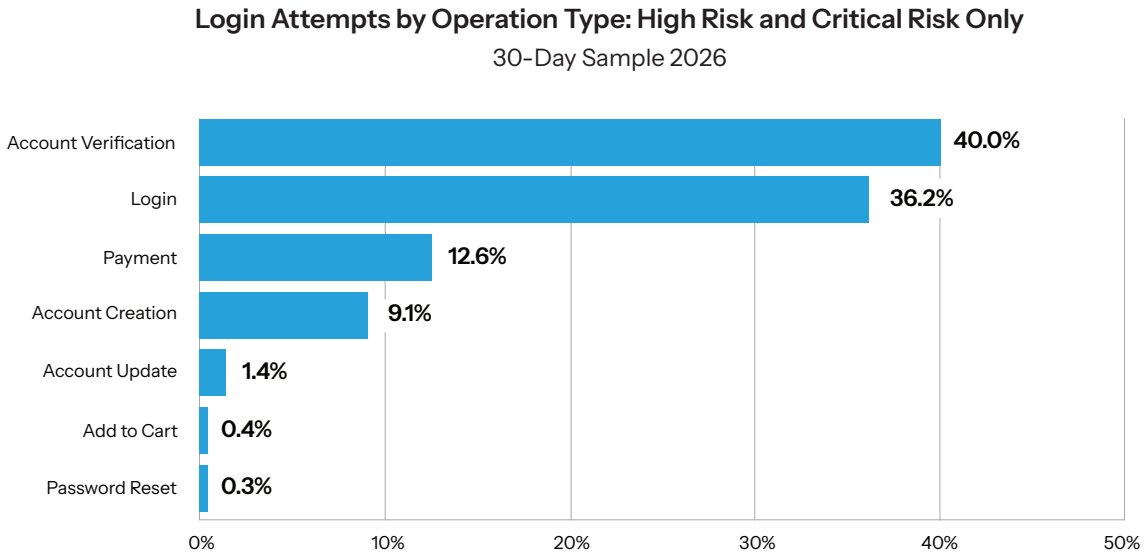


Fig. 4: Account verification led the high-risk/critical-risk login activity at 40%

This suggests that threat actors may use account verification requests to determine whether an email address is associated with a valid account before launching credential stuffing campaigns. In multistep authentication flows, these requests could also indicate automated multi-factor authentication (MFA) prompts or SMS-pumping activity, even when premium-rate numbers are not involved.

Payment activity ranked third in both views, but its share increased slightly from 11.4% to 12.6% among high- and critical-risk activities. This is concerning because elevated risk payment activity may indicate that attackers have gained account access and are attempting unauthorized purchases or other fraudulent transactions without the account holder's knowledge.

Assessing client type

According to the same 30-day sample data, web browsers were the most prevalent client type, representing 53.7% of login attempts, followed by Android apps at 30.5% (Figure 5).

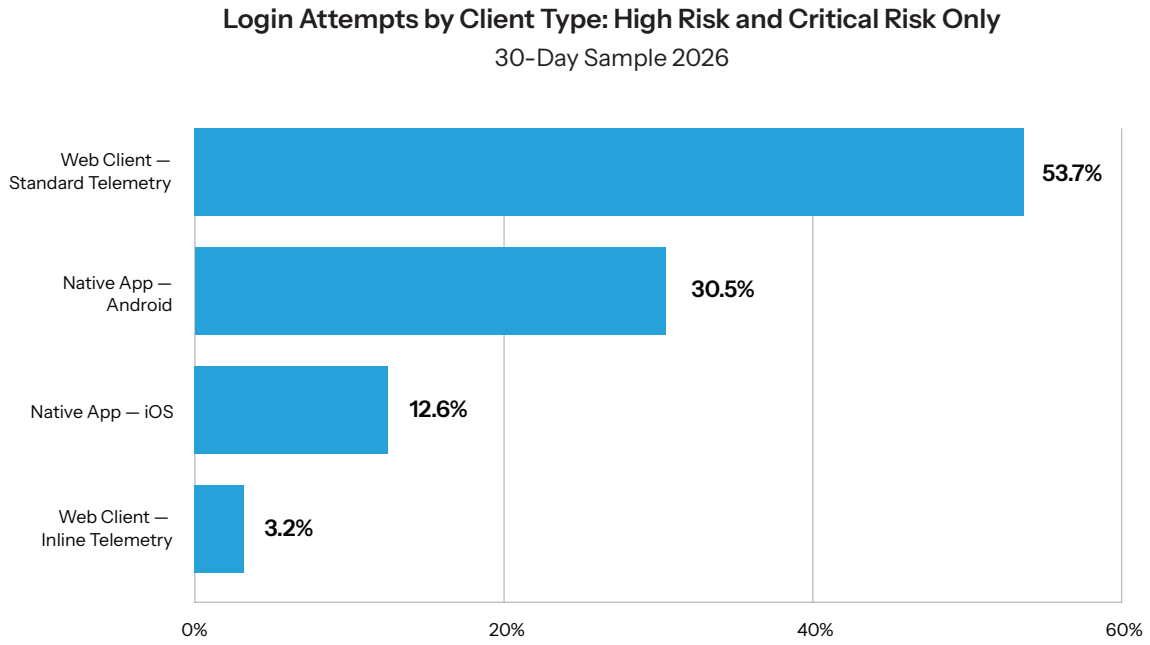


Fig. 5: Most login attempts tagged as high or critical originated from browsers, followed by Android apps

This distribution likely reflects the broader availability of web and mobile browser access since not all businesses have their own dedicated apps. However, this traffic volume does not imply lower exposure to fraud and online scams for iOS and Android app users.

Account opening abuse schemes

Account opening is a high-risk point in the consumer journey, prone to a plethora of abuse schemes that can cause substantial financial losses when not effectively mitigated. Two prevalent tactics — SMS pumping and welcome bonus fraud — were observed across customer environments during the reporting period and they each pose specific dangers to organizations.

SMS pumping

In SMS pumping schemes, adversaries exploit one-time password (OTP) triggers in account creation and password recovery workflows. By automating requests that send registration or password-reset OTPs, bots can create fake accounts at scale and direct messages to premium-rate international numbers. In some instances, the scheme involves rogue telecom carriers that share in the resulting revenue. This leaves victimized organizations with inflated SMS termination charges while fraudsters and participating rogue carriers split the payout.

Welcome bonus fraud

In welcome bonus fraud, often referred to as promo abuse, adversaries typically automate account registration to repeatedly collect incentives intended for new or first-time customers, including cash promos or discounts. The success of this scheme relies on the creation of a large pool of fabricated fake identities, enabling fraudsters to operate at a scale that manual registration won't be able to support.

In food delivery and ride-sharing services, fraudsters generate bogus accounts to harvest "First ride free" or "First meal free" offers. They may then resell the accounts, already loaded with promotional value, through channels such as Telegram. Meanwhile, in media and streaming, bots similarly register large numbers of free-trial accounts, which are bundled and sold on the dark web. The result is direct promotional loss for the organization, skewed customer metrics, and potential exposure to account abuse.

Specific fraud and abuse schemes by industry

Threat actors tailor fraud and abuse schemes to the industries and verticals they target. While methods may overlap, attacker objectives and the resulting impact to organizations may differ by sector. The table outlines schemes against commerce, financial services, and media and games.

Industry	Fraud and abuse schemes
Commerce	• **Promo and return fraud:** Adversaries hijack an account, place an order using a stored credit card (via an account takeover attack), intercept the package, then file a "did not arrive" claim for a refund, leaving the merchant with a total loss. • **Stored value extraction:** Attackers cash out saved gift card balances and loyalty points, which are untraceable on the dark web. • **Loyalty points as digital currency:** Frequent-flier miles and hotel points are highly valuable. Attackers compromise travel accounts either to book flights for resale or redeem points.
Financial services	• **Fast money movement:** Once attackers take over an account, they change contact details to block alerts, add payees, and use digital payments networks or wire transfers to send funds to mule accounts. • **Payout redirection:** Attackers use account takeover to access customer portals, change bank details, and divert payouts or refunds to their own accounts instead.
Media and games	• **Concurrent session abuse and resale:** In media and games, attackers hijack premium accounts or subscriptions, rotate access among viewing rings, or obtain rare in-game assets to peddle for real-world cash.

Table: Each industry encounters distinct fraud schemes

High- or critical-risk login attempts by industry

High- or critical-risk login attempts are often concentrated in industries that maintain large repositories of account data and have direct monetization opportunities. Based on our telemetry, commerce accounted for 41.1% of observed login attempts, which is more than double the volume of the second highest industry (Figure 6).

Login Attempts by Top 10 Industries: High Risk and Critical Risk Only

30-Day Sample 2026

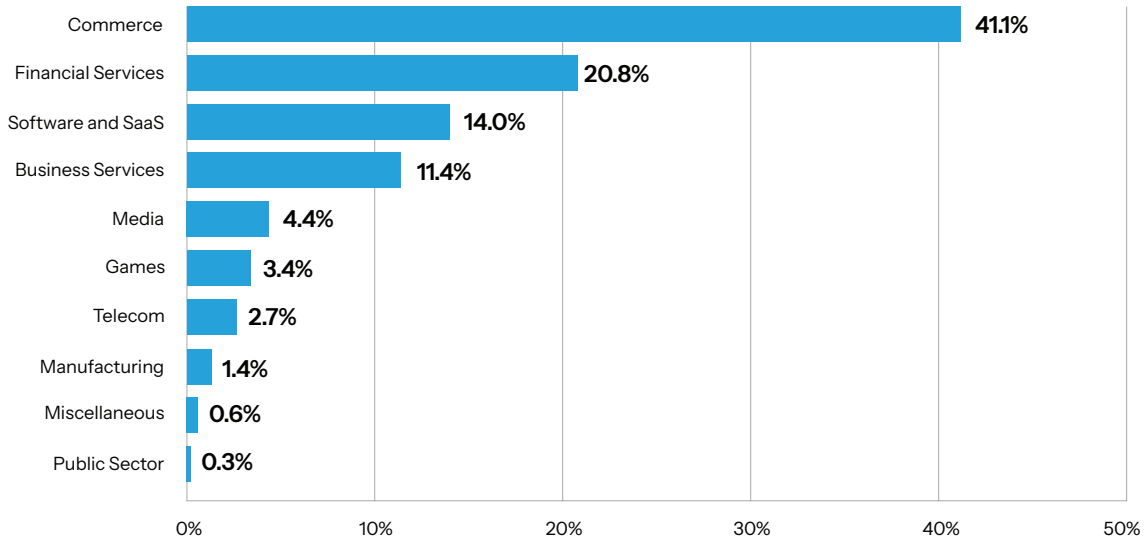


Fig. 6: Commerce had the highest percentage of high-risk and critical-risk login attempts

Retailers, in particular, manage large volumes of consumer accounts, making retail a viable testing ground for attackers who are validating credentials obtained through previous data breaches. A successful account takeover can enable unauthorized purchases, misuse of loyalty and rewards points, and other forms of fraud, including those outlined in the table.

Retail accounts also contain valuable data beyond credentials, including PII, order and purchase histories, and stored payment methods. Attackers can resell this information through dark web forums and other underground marketplaces.

Financial services accounted for 20.8% of observed high-and critical-risk login attempts. Like commerce, this industry attracts adversaries because accounts contain sensitive financial data, including account balances and transaction information. Once attackers compromise an account, they may transfer funds, redirect payments, or conduct other unauthorized activities.

Akamai also recorded 4.4% of high-and critical-risk login attempts targeting the media industry, which includes video, social media, and publishing companies. Meanwhile, games accounted for 3.4% of login attempts classified as high or critical risk. Compromised games accounts can support unauthorized in-game purchases, the theft and resale of high-value virtual assets, and even the resale of the accounts themselves.

Login activity and attack schemes by region

As we discussed in Analyzing login activities, the login phase is the most active stage of the account lifecycle journey. While most login activity is generated by legitimate users, login is also the entry point for malicious bot activity. So, there is significant value in monitoring this first interaction to get a head start on mitigating high- or critical-risk events. AI and machine learning can identify suspicious activity quickly and accelerate investigation and mitigation of all types of digital fraud, including AI-driven fraud.

A regional view of login activity during a 30-day sample of data from Akamai Account Protector reveals that N. Amer. experienced the highest percentage of login attempts at 45.8%, followed by APAC (25.1%), EMEA (23.1%), and LATAM (1%; Figure 7).

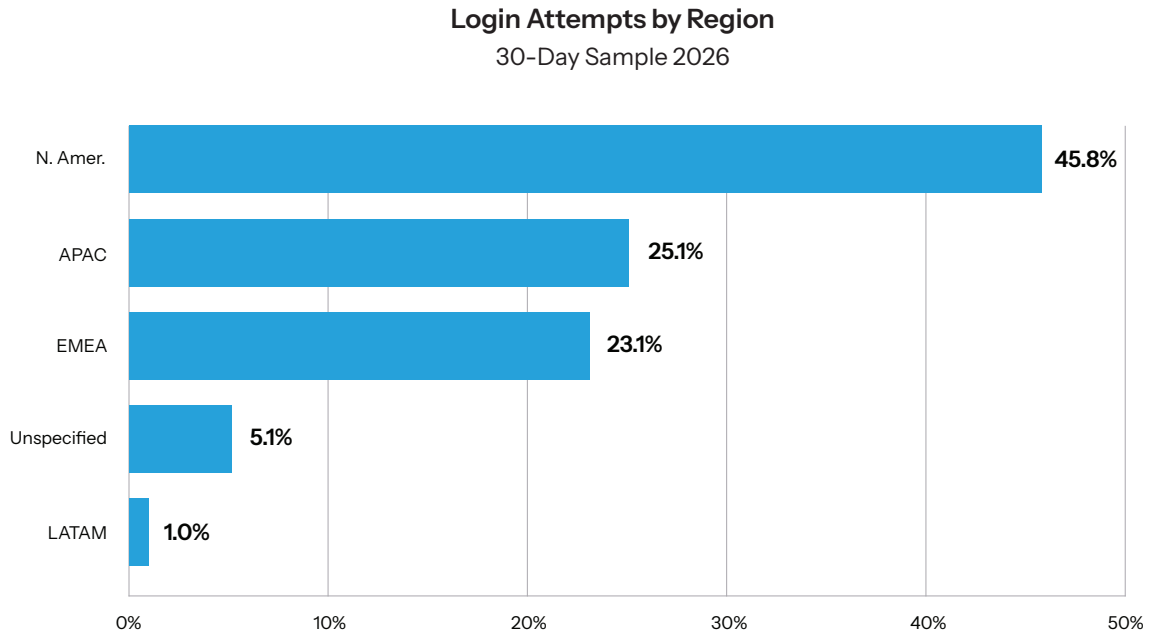


Fig. 7: N. Amer. experienced the highest proportion of login attempts followed by APAC, EMEA, and LATAM. Note: "Unspecified" are the attempts for which a customer isn't listed, so they cannot be classified by region.

A closer look at login attempts based on user behavior profiles and risk signals provides insights into potential adversarial bot activity against the network. The majority of login attempts posed low to medium risk across all regions. However, when looking only at high- to critical-risk levels, we see that APAC and LATAM experienced the highest proportion of these elevated risk level login attempts at 36.4% and 37.8%, respectively, versus N. Amer. at 23.6% and EMEA at 23.7% (Figure 8).

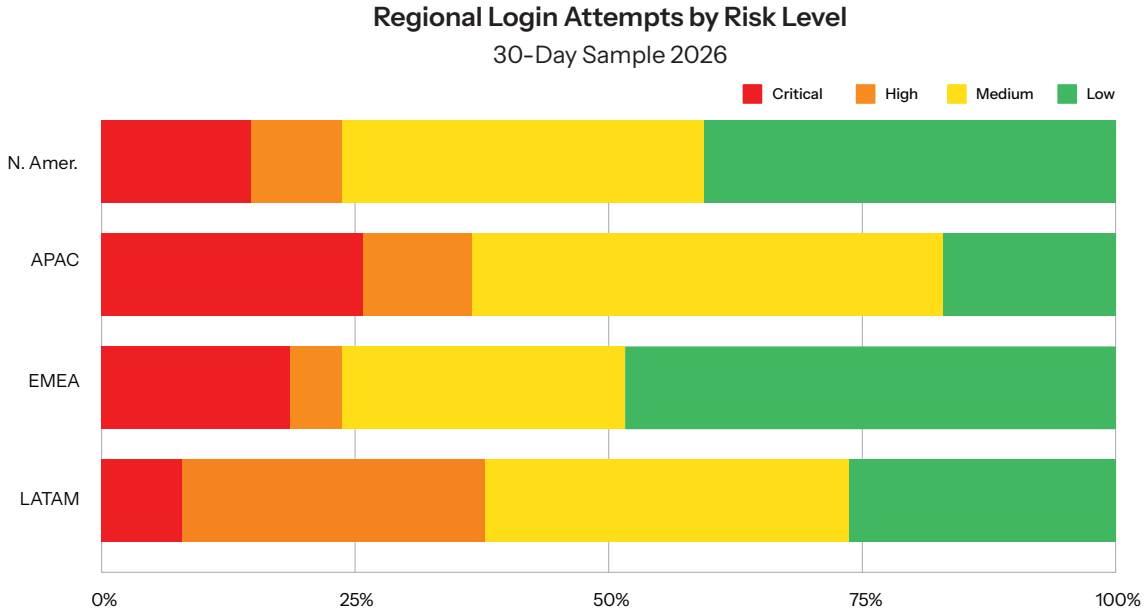


Fig. 8: The majority of login attempts within each region posed low/medium risk, with a higher concentration of high/critical risk login attempts in APAC and LATAM

Our analysis of daily data during the same 30-day period reveals a fairly consistent level of activity with dips during the weekends, most notably in N. Amer. and EMEA (Figure 9).

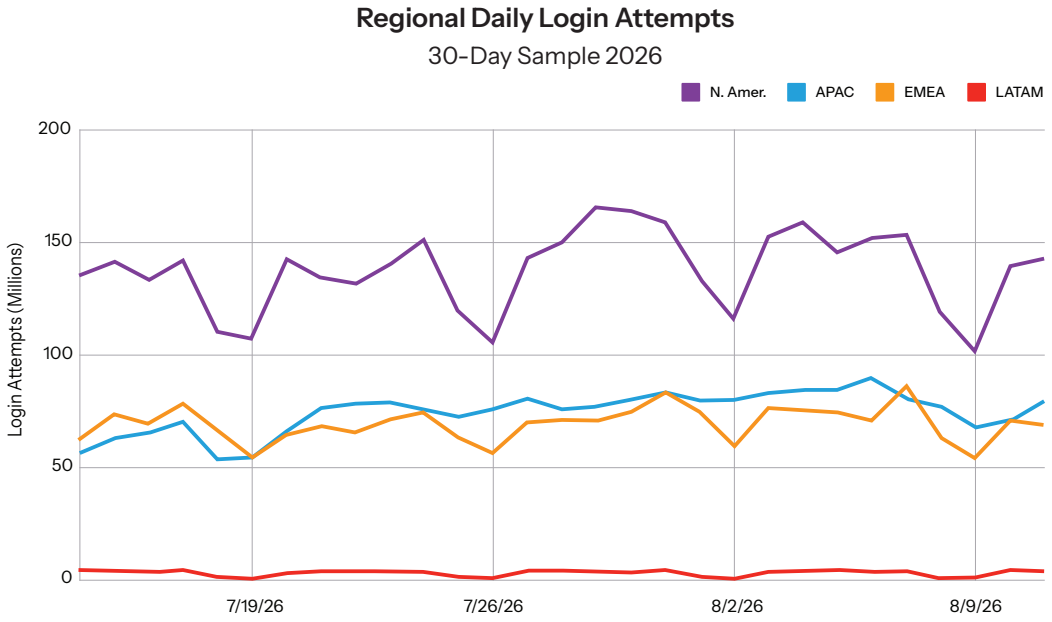


Fig. 9: This 30-day sample of data shows consistent activity with declines mostly during weekends

As discussed in our Commerce 2026 SOTI Security report, this trend confirms that, although automated, bots remain tethered to human business cycles and taper off during the weekend when human oversight is less active.

Regional and industry-specific account abuse schemes in action

Certain account abuse schemes are more aligned with certain industries. For example, money muling is an issue for financial institutions, while loyalty points are highly liquid in ecommerce. Akamai researchers observed many industry-specific examples of account abuse during 2026 in APAC, EMEA, and N. Amer.

APAC: Financial services account opening abuse

During the week starting July 29, 2026, Akamai researchers observed two attacks occurring simultaneously against a large financial institution in APAC. The attacks involved creating net-new accounts for distinct purposes and they unfolded at different speeds.

The first was a volumetric attack in which the cybercriminals leveraged botnets to send approximately 42,000 requests during a one-week period to the targeted financial institution (Figure 10).

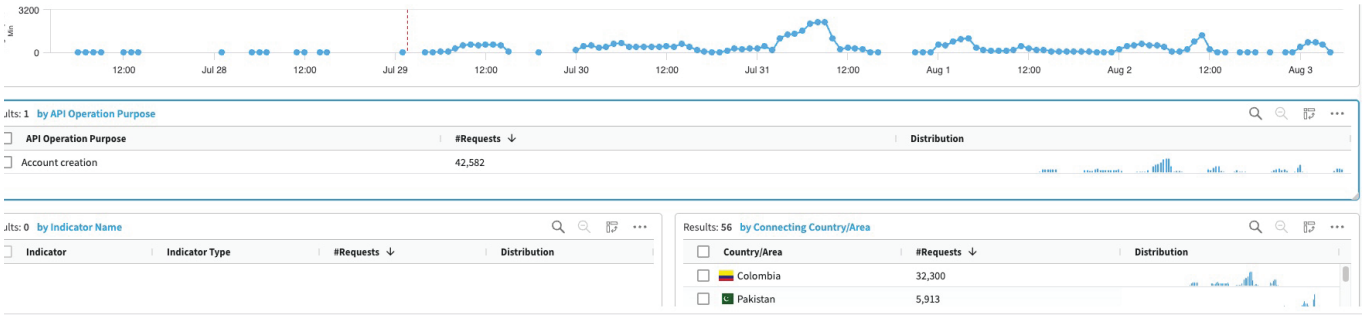


Fig. 10: Volumetric botnet attack activity from July 29, 2026 — August 3, 2026 to validate PII

The intent was likely to validate PII before launching a coordinated large-scale attack or selling the credentials. However, because the activity originated from two areas with which the financial institution does not normally interact, and the attack reused known fraudulent IPs, it triggered a high/critical risk score for investigation and was blocked.

Concurrently, fraudsters targeted the financial institution with a low-volume attack consisting of 164 requests, intended to operate below detection thresholds and likely for the purposes of creating accounts to serve as money mules or for synthetic fraud (Figure 11).

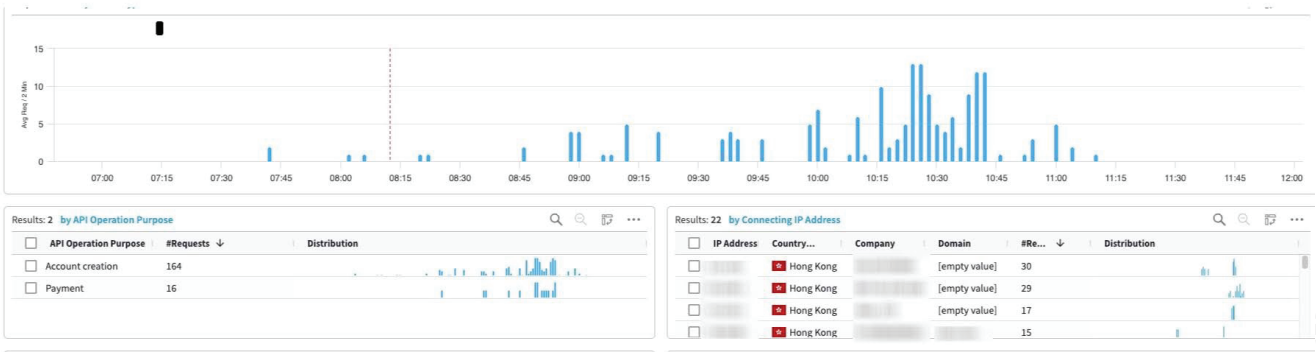


Fig. 11: Low-volume attack activity during a 24-hour period to set up payment fraud

Money mules shuffle stolen funds across institutions and payment channels. Meanwhile, synthetic accounts that may combine real data, such as a valid Social Security number, with a fake name and fake credit history can be used to abuse lines of credit.

These types of coordinated efforts provide a glimpse into the elaborate fraud networks in play and the broad and deep intelligence required to detect and block attacks while enabling legitimate activity.

EMEA: Event ticket hub account takeover

During Q2 2026, a ticket merchant based in EMEA was targeted by an account takeover attack. The attack was a low-volume, highly targeted attack. The threat actors likely would have followed a typical account takeover flow by using high-quality stolen credentials to gain unauthorized access to user accounts, immediately reset credentials, and use the privileges to engage in fraudulent activities. Had the attack not been blocked, the threat actors could have made fraudulent purchases on the site or elsewhere by using saved credit cards, draining loyalty points, and engaging in ticket hoarding and scalping.

N. Amer.: Online retail payment fraud

During July 2026, threat actors targeted a homeware retailer in N. Amer. with a payment fraud scheme (Figure 12).

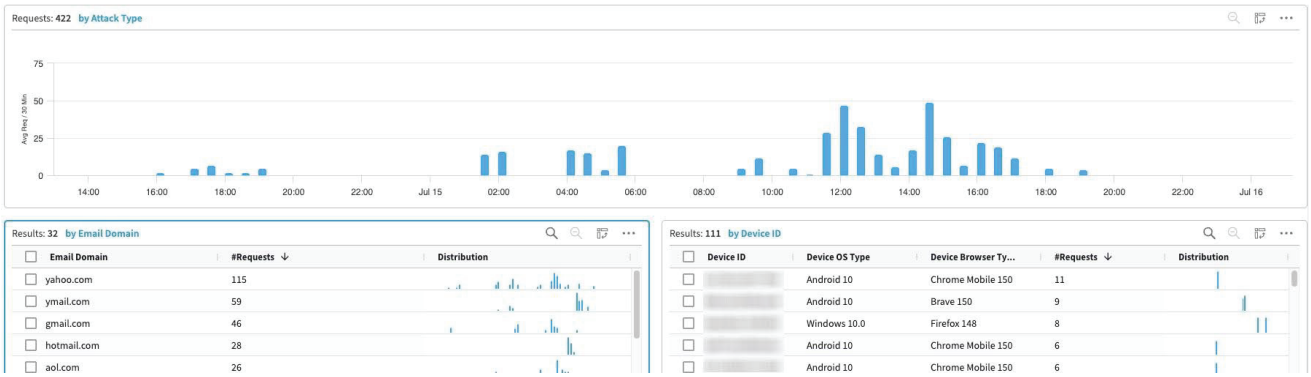


Fig. 12: Low-volume activity to validate credit cards

Akamai researchers believe it was likely that the cybercriminals obtained a list of stolen credit card numbers with expiration dates and the three- or four-digit security codes (CVV). Attempting to evade detection, the attack scheme used the guest checkout function, legitimate email addresses, and bots to conduct a low volume of transactions at low purchase amounts for card validation.

Had the attack been successful and the credit cards validated, cybercriminals could have either used them in large-scale fraudulent operations or sold them at a premium (e.g., 10 to 20 times the basic stolen card rate depending on the credit limit, expiration date, and geographic region).

Protect your brand

An organization's attack surface was once largely confined to the network perimeter: Adversaries had to breach that perimeter and persist within it to siphon confidential data. In recent years, however, attacks have expanded beyond internal systems and networks to encompass damage to the brand reputation.

Cybercriminals have increasingly capitalized on corporate trust, turning hard-earned brand loyalty into a vector for exploiting its customers and partners. When phishing schemes and logo abuse succeed, victimized consumers lose money, but the brand ultimately pays the price through shattered customer relationships, degraded lifetime value, and direct revenue loss. Moreover, phishing and multiple channel scams across social media, paid ads, and online marketplaces are often entry points for larger fraud schemes, including credential stuffing and account takeover.

Traces of fraud in the DNS traffic

The web has long served as the primary battleground for fraud and abuse, from the heyday of spam campaigns to modern sophisticated phishing and malware operations. Today, we can harness Akamai's unique visibility into billions of daily DNS requests to analyze the malicious traffic and highlight the queries or "early warning signs" that often precede data theft or system compromise.

During the first half of 2026, the number of phishing attacks surged, almost doubling between Q1 2026 and Q2 2026. As a percent of the total threat traffic each quarter, phishing increased from 8% in Q1 to 15% in Q2, jumping from seventh place to the number two spot overall (Figure 13).

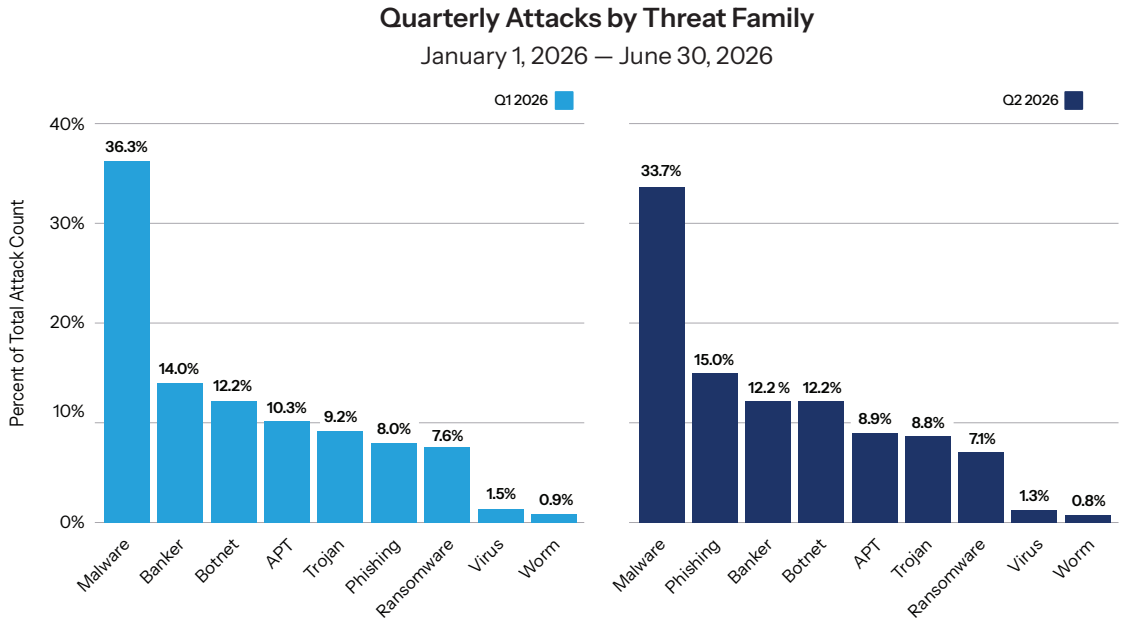


Fig. 13: Phishing attacks rose from Q1 to Q2 2026, surpassing both botnets and advanced persistent threats

Phishing activities impersonating high technology, commerce, and financial services, along with npm registry-abusing spear-phishing campaigns, drove this growth. However, malware remained the top threat type with more than double the volume of the second place category in both quarters.

Advanced persistent threats (APTs) represented 10.3% and 8.9% of observed attacks on organizations in Q1 and Q2 2026, respectively, posing critical risks to highly valuable crown jewels, like intellectual property and other sensitive data, through undetected lateral movement and data exfiltration. Botnets also continued to drive malicious traffic by automating and scaling credential stuffing and account takeover attempts.

Top phishing threats

By zooming in on phishing, our DNS telemetry data reveals a clear trend: Adversaries rely heavily on trusted brand identity to gain initial access and harvest credentials. During the reporting period (January 1, 2026 — July 31, 2026), the majority of the malicious domains that targeted enterprise networks during the reporting period impersonated high tech (41.6%), financial services (18.9%), and commerce (8.4%) brands (Figure 14).

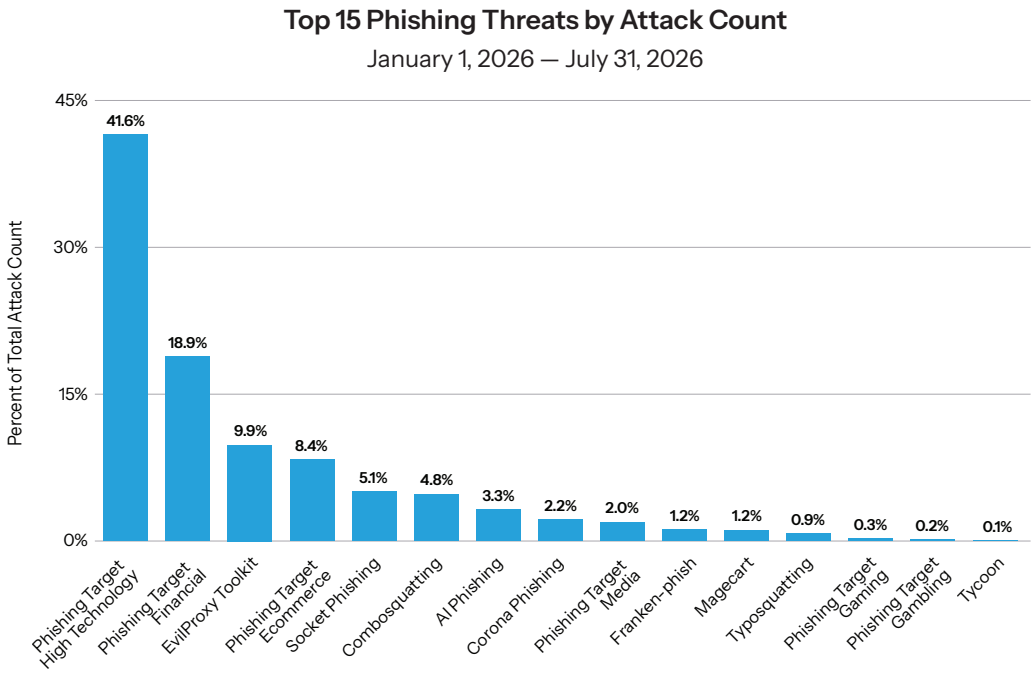


Fig. 14: Nearly 70% of phishing threats mimic high technology, financial services, and commerce brands

The phishing threat in third place (with almost 10% of attacks) is EvilProxy toolkit, a prominent PhaaS platform that allows threat actors with little to no technical expertise to launch their own sophisticated, large-scale campaigns that are capable of bypassing MFA and stealing active session cookies in real time.

Highly targeted spear-phishing campaigns known as “socket phishing” took fifth place (5.1%), with attackers spoofing daily workflow notifications, such as secure sharing and Microsoft login prompts, to obtain credentials. Reports indicate that these campaigns were aimed at sales and commercial staff within the critical infrastructure sectors in the United States and allied nations.

Beyond traditional phishing attacks, adversaries are capitalizing on the boom of generative AI in the workplace. This has given rise to “AI phishing” (seventh in our dataset), which is defined as attacks that include fake sites impersonating major platforms like OpenAI’s ChatGPT to pilfer employee credentials. Although AI phishing currently represents only 3.3% of observed activity, organizations should monitor it closely as a critical initial access vector. According to our Enterprise AI Usage Risk Report 2026, nearly half of enterprise employees have tried workplace AI tools, though only 18% use them weekly; those figures will likely rise steadily over time. As employee adoption inevitably scales, these low-volume lures will likely transform into high-yield entry points for enterprise compromise.

Regional trends reveal notable fraud and abuse patterns

Analyzing Akamai DNS telemetry across threat families and verticals reveals that malware, APTs, trojans, ransomware, and phishing were the most prevalent regional threats during the reporting period (Figure 15).

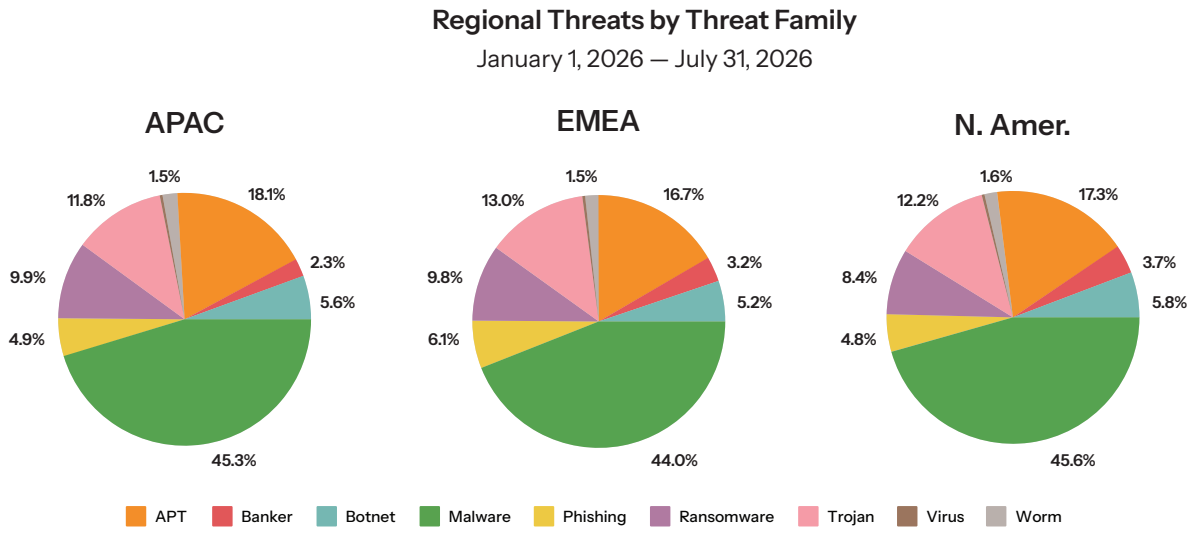


Fig. 15: Malware, APTs, ransomware, and phishing were the top threat types across regions

It is important for organizations to have visibility into both attack volume and threat types to ensure a comprehensive security architecture. For example, phishing and malware attacks operate based on volume, sending more queries or attack attempts to their targets (see Figure 13). In contrast, threat types like ransomware and APTs focus on select targets, move quietly within a network, and strike when they’ve established a foothold. These threat types currently rank lower in the number of attack attempts, but they are prevalent: Their ranking is increasing and their impact is rising. According to the FBI’s 2025 Internet Crime Report, ransomware complaints increased by 14.4% from 2024 and 2025 and losses jumped by 159%.

A view of regional data through a vertical lens reinforces that attackers don't discriminate. During the reporting period, several verticals were impacted by fraud and abuse, with little overlap among the top verticals targeted within each region (Figure 16).

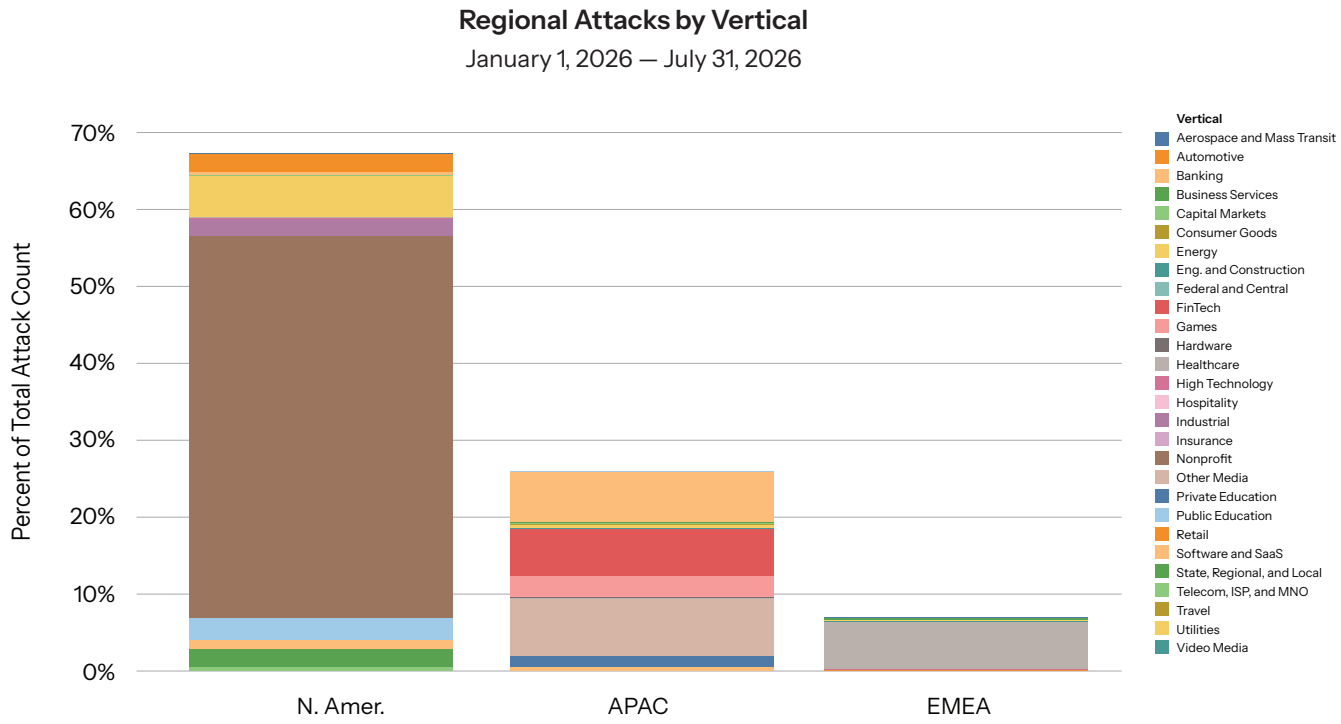


Fig. 16: Attacks primarily targeted the nonprofit, other media, banking, fintech, and healthcare verticals; different attack types were aimed at different regions

According to Akamai DNS telemetry, 49.6% of attacks targeted the nonprofit vertical in N. Amer., and consisted primarily of malware, phishing, and botnet queries. Nonprofit organizations maintain valuable data yet they typically don't have the same IT resources as organizations in other industries, such as financial services, which makes them valuable targets.

After nonprofit, the second most targeted vertical was found in APAC: other media at 7.5%. This vertical includes publishing, which is a prime target for bot traffic. APAC also had the third most targeted vertical: banking (6.4% of attack activity). The combination of the banking and fintech verticals made financial services the most targeted industry in APAC. Malware, APTs, ransomware, and botnets were among the primary attacks observed in the region.

EMEA had the fourth most targeted vertical globally: healthcare (6.2% of attacks), with threat actors relying primarily on a combination of malware, phishing, and ransomware.

Modern breaches often begin with phishing and they leverage an expanding ecosystem of tools to amplify attacks. A closer look at phishing attacks by region reveals that N. Amer. was the most targeted region by phishing attacks at 66.8%, followed by APAC (17.5%), and EMEA (15.7%; Figure 17).

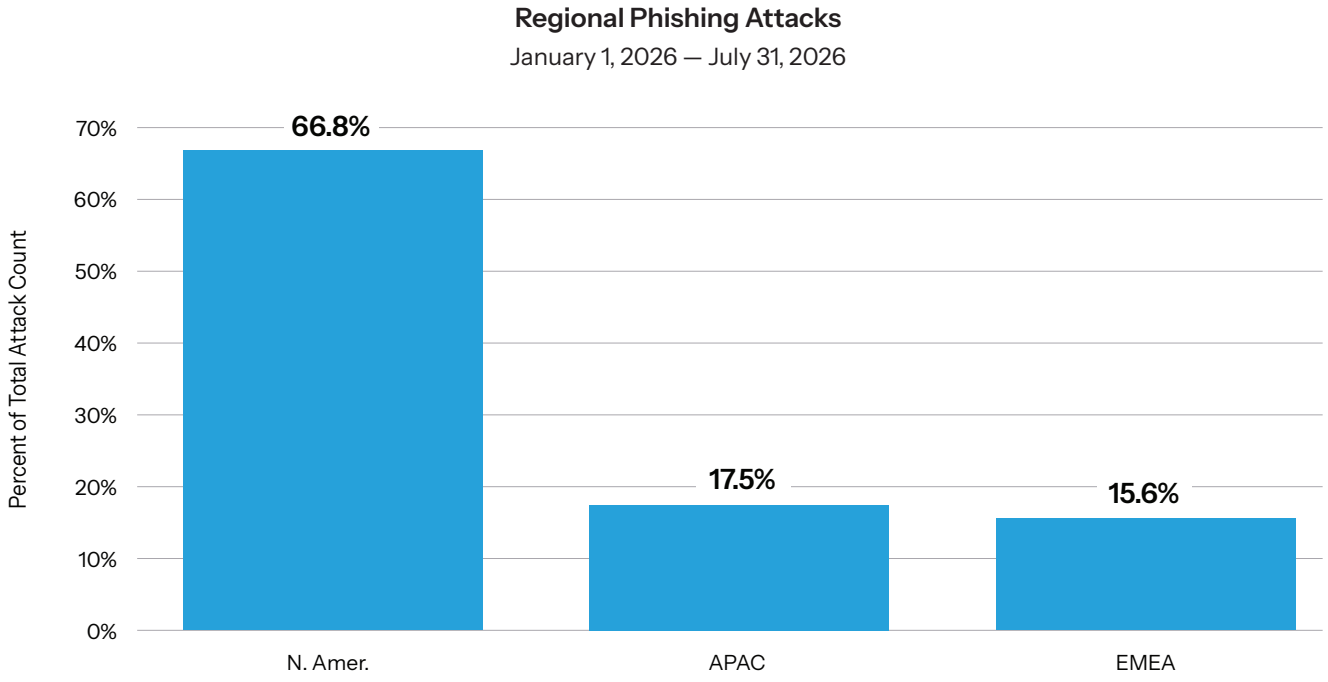


Fig. 17: More than two-thirds of phishing attacks targeted N. Amer.

A World Economic Forum study found that of the 72% of CEOs who reported they or someone in their network had been impacted by cyber-enabled fraud in 2025, 62% experienced a phishing attack, the most of any scheme. On the front end, attackers are iterating daily, using phishing kits to accelerate execution and LLM-polished lures to create hyperpersonalized emails and more convincing campaigns.

Additionally, Akamai DNS telemetry confirmed that attacks were further boosted by tactics that enable scale, speed, and localization to launch campaigns against a broader set of targets. These include bulk domain registration, shared infrastructure, brand impersonation, spoofed login and payment portals, and fake workflows.

Notable phishing campaigns

Over the past year, threat actors increasingly piggybacked on trusted brands to launch sophisticated phishing campaigns, with dire consequences for impacted organizations. We spotted three notable phishing attacks in the wild: A spoofed Microsoft campaign, an Amazon-themed campaign, and a financial services impersonation campaign. The Threats in depth section details a fourth notable attack that we observed this year: a campaign that impersonated tickets to a popular sporting event in 2026.

Spoofed Microsoft campaign

As one of the world's most widely trusted enterprise platforms, Microsoft remains a prime target for phishing and brand impersonation campaigns. By creating bogus Microsoft login portals, threat actors aim to harvest enterprise credentials that can be used to gain initial access, establish persistence, and ultimately pilfer high-value corporate data.

Between October 20, 2025, and March 19, 2026, Akamai observed an active credential phishing campaign centered on spoofed Microsoft login interfaces. Our analysis identified 78 deduplicated DNS connection attempts to attacker-controlled infrastructure, specifically within the meask[.]xyz and nsdjj[.]xyz namespaces, spanning multiple customer environments. The phishing operation consisted of more than 400 newly registered domains, indicating sustained activity, with threat actors reportedly conducting bulk domain registration.

Based on our findings, the campaign disproportionately targeted telecommunications and hosting/IT service providers. Organizations in the telecommunications vertical across the APAC region and Africa experienced sustained exposure throughout late 2025, while affected hosting and IT companies drove concentrated spikes in early 2026. Notably, a single organization in N. Amer. accounted for nearly half of all observed activity. The campaign's primary objective is consistent with Microsoft 365 credential harvesting, enabling account takeover and further compromise.

Amazon-themed campaign

Popular commerce platforms with integrated services to support the entire customer journey are attractive targets for phishing and brand impersonation. By creating Amazon-themed pages that are designed to mimic legitimate Amazon branding and content, threat actors stand up domains to launch phishing campaigns engineered for credential theft and payment card fraud.

During a widely publicized phishing spike between March 5, 2026, and April 1, 2026, Akamai DNS telemetry revealed a large wave of 54,283 DNS connection attempts across 181 customers. Within that activity, our analysis discovered a persistent phishing campaign spoofing Amazon, specifically from brasil108-6015.com[.]br and brasil106-5030.com[.]br. The two Amazon-themed phishing domains generated 640 DNS connection attempts. A third domain, brasil110-8020.com[.]br, expanded the cluster to 983 attempts with combined daily volume peaking at 63 attempts (Figure 18).

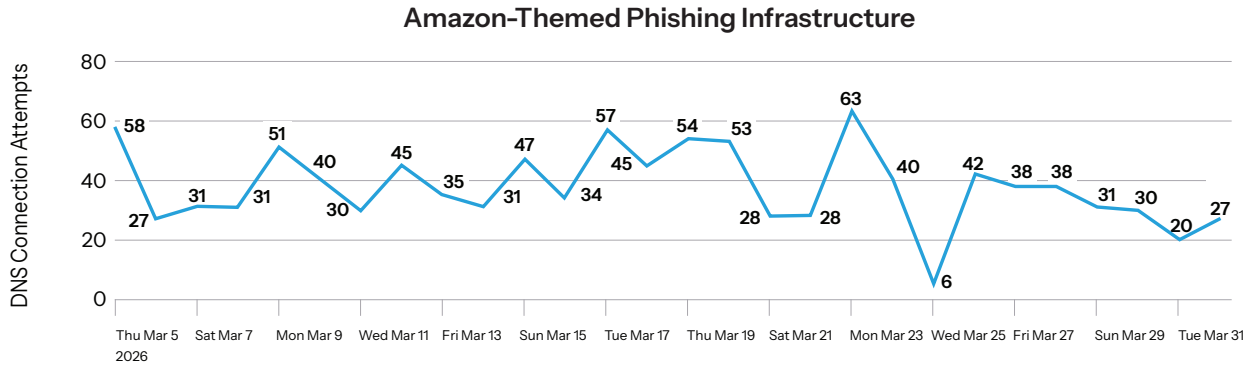


Fig. 18: Combined daily DNS connection attempts across two Amazon-themed phishing domains (Note: March 27–28 are estimated on the basis of the whole history)

The third domain, brasil110-8020[.]com.br, shares the same naming convention as the two domains associated with Amazon impersonation, although its involvement in the same phishing campaign remains unconfirmed.

Based on our findings, the target landscape was highly concentrated. In N. Amer., a single hosting/IT services provider experienced more than half of the DNS connection attempts across the three domains. Meanwhile, Brazil represented the most targeted region for attack attempts. After hosting/IT services, the most targeted industries included government/public sector, telecommunications, financial services, and retail/consumer goods.

Financial services impersonation campaign

The financial services industry continues to be a prime target for fraud and abuse. Using phishing infrastructure that impersonates banks, payment providers, and other financial services firms, threat actors harvest online banking credentials and engage in account takeover for financial gain.

Between May 20, 2026, and May 29, 2026, Akamai observed a financial phishing cluster with 862 DNS connection attempts, with activity peaking on May 23 (Figure 19).

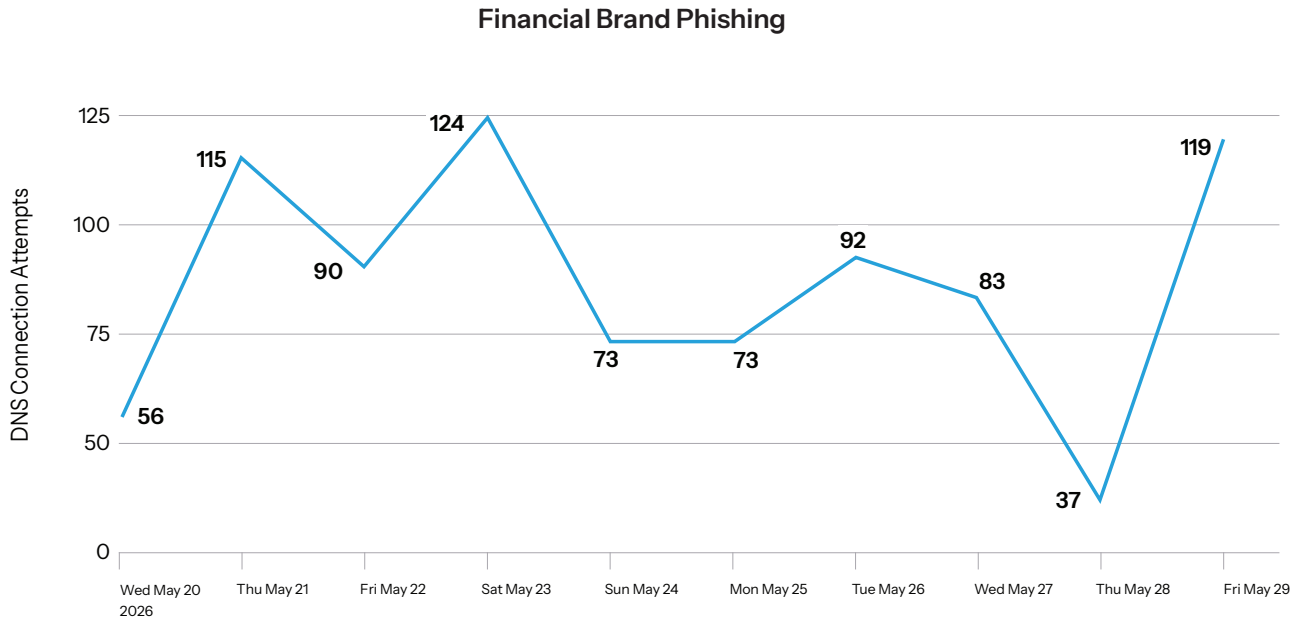


Fig. 19: Daily DNS phishing attempts during a multibrand crossregion financial services impersonation campaign

The phishing campaign impersonated major financial services brands based in different regions and used 45 destination domains and 76 destination IPs to target customers in APAC, EMEA, LATAM, and N. Amer. Our researchers observed a disproportionate focus on one construction/materials organization in LATAM, alongside broader recurring attempts against the telecommunications and public sector industries in APAC, LATAM, and N. Amer. In EMEA, the automotive and retail verticals were primarily targeted.

Akamai Security

Threats in depth

High-profile sporting event spoofed in coordinated ticket impersonation attack

Major sporting events create ideal phishing conditions: high demand, limited supply, emotional urgency, and fragmented user journeys across official and unofficial search results. Between May 17, 2026, and June 13, 2026, Akamai researchers identified a coordinated campaign impersonating official ticketing and hospitality portals for a major sporting event. Our analysis revealed a highly unified toolkit family using infrastructure reuse, such as cloned HTML layouts and shared tracking telemetry, rather than isolated copycat sites, to scale fake ticket sales globally.

Across 12 unique hostnames and 20 scrape captures, the sites presented a near-identical copy of the impersonated sites of the global sporting event, with prominent "BUY NOW" prompts that pushed visitors into fake purchase flows. The "BUY NOW" model served as the main conversion point. Once users engaged with the page, they were funneled into a fake purchase flow designed to collect sensitive information.

Keep your employees from becoming victims

Phishing emails tend to use a set of techniques to get employees to click on links. The techniques are similar to sales or marketing strategies. A great example of these strategies are Robert Cialdini's classic principles of influence that guide psychological shortcuts and human decisions.

1. **Reciprocity** — Create obligation by giving something first
2. **Commitment and consistency** — Exploit prior commitments to influence future actions
3. **Social proof** — Make malicious behavior appear normal or widely accepted
4. **Authority** — Impersonate trusted figures to compel action
5. **Liking** — Build rapport to lower skepticism and resistance
6. **Scarcity** — Create urgency by making opportunities seem limited
7. **Unity** — Exploit shared identity to establish trust quickly

Training employees to recognize these techniques can help them avoid being the victim of a phishing campaign.

What victims risk exposing

The fake checkout flows were built to collect high-value personal and financial data. Victims who followed the “BUY NOW” path may have exposed their:

- Full payment card data, including primary account numbers, CVV, and expiration date
- PII, including name, email, and phone number
- Passport or identity-document details in regional flows where such information is requested for ticket delivery or event access

Because the checkout flow is connected to actor-controlled infrastructure, the submitted data does not reach the official purchase site. Instead, it is sent to infrastructure controlled by the phishing operator.

That creates two layers of risk:

1. Immediate financial fraud via stolen payment card information
2. Longer-term identity abuse through harvested personal details, email addresses, phone numbers, and identity documents

The second risk can persist well beyond the event itself. Stolen identity data can be reused for account takeover, synthetic identity creation, immigration-themed scams, travel fraud, and other downstream abuse.

Shared checkout infrastructure

One of the strongest technical links in the campaign was the shared checkout back end. Multiple front-end spoof sites routed transactions to the exact same Google Cloud Run serverless back end to centralize payment theft.

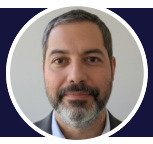
This indicates that multiple phishing sites were connected to one back-end payment-skimming API. In practical terms, the domains acted as interchangeable landing pages, while the check-out and data-exfiltration paths remained centralized. This architecture is common in scalable phishing operations. Operators can rotate, replace, or expand front-end domains while keeping the back-end collection infrastructure stable. If one domain is blocked or taken down, another can continue to send victim data to the same back end. For defenders, this means the back-end infrastructure may be more durable and more valuable to monitor than any single domain.

Additionally, operators behind this campaign used highly disposable infrastructure, including short-lived 90-day certificates and companies that provide poorly supervised reverse proxies, to systematically bypass traditional security reputation blocklists.

For defenders, the lesson is clear: Event-themed phishing should not be tracked domain by domain alone. The more durable signals often sit underneath the page in tracking IDs, back-end endpoints, certificate timing, content hashes, and reused third-party components. Those signals can reveal the operator infrastructure that is connecting what may otherwise look like separate phishing sites.



Moshe Cohen
Principal Security Analyst, Akamai



Shiran Guez
Director of Engineering, Akamai

Beyond traditional phishing: Online scams and fraud

We collaborated with Bolster AI to provide a more comprehensive view of evolving phishing and online scam trends, tactics, and their impact across organizations from July 2025 through June 2026.

In early 2026, the number of online scams surpassed that of traditional phishing for the first time. While *phishing* refers to activities designed to steal credentials, authentication information, or other sensitive data through brand impersonation, *online scams* focus on fraud activity with the goal of financial or consumer deception rather than credential harvesting (e.g., fake stores, investment scams, tech support scams, payment fraud, or similar schemes).

Data from Bolster AI demonstrated that online scam detections approximately doubled, surging from ~ 3.2 million to ~ 6.4 million, marking a significant shift in threat tactics. Since MFA has become a baseline security measure, attackers are pivoting toward higher-return scams for faster payouts. Traditional phishing, however, still grew by 61% by Q2 2026, according to Bolster.

Together, these numbers show that while credential theft remains a core driver, threat actors are aggressively expanding into social engineering tactics that can manipulate consumers into direct financial losses (Figure 20).

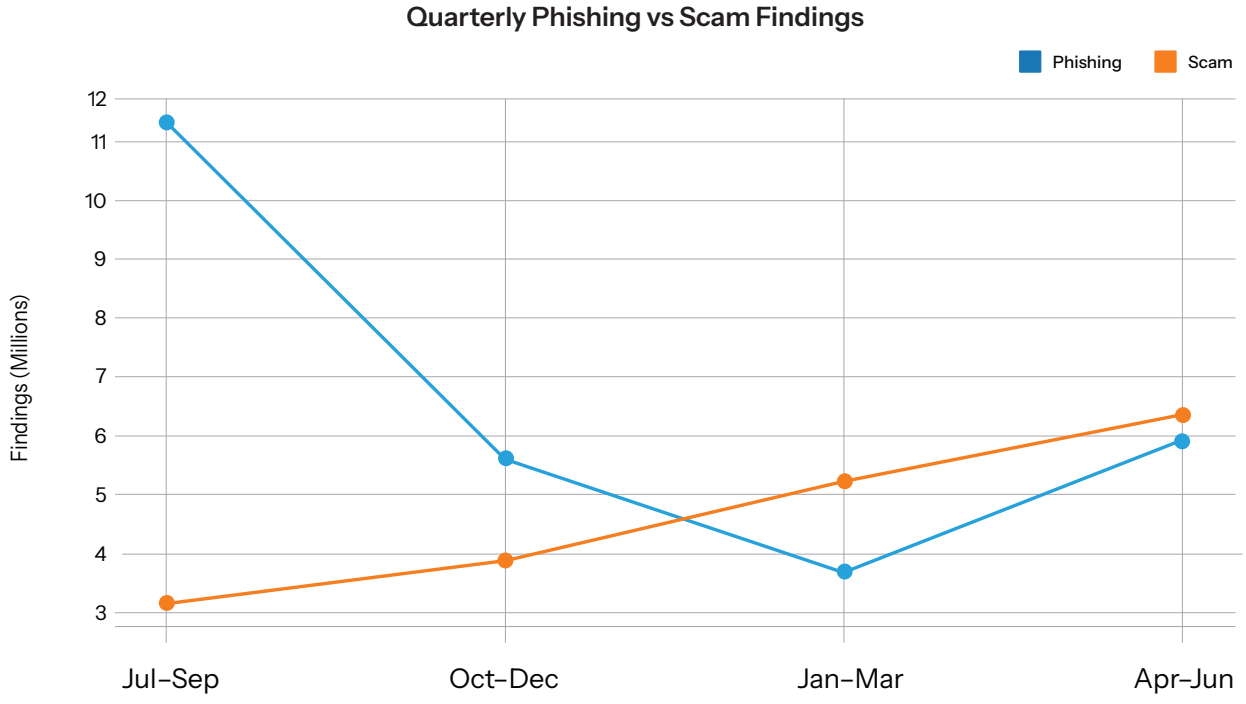


Fig. 20: Online scams are gaining traction and overtook phishing volume for the first time in Q1 2026

Industrialized infrastructure and crossplatform threats

Large-scale brand impersonation drove phishing across customer-facing sectors, led by gambling and betting, media and entertainment, and travel and transportation (Figure 21). Campaigns are increasingly industrialized through template sharing and infrastructure clustering. According to Bolster AI research, threat actors used more than 1,200 look-alike .shop domains in one retail campaign, sharing a single IP address and reusing the same HTML, JavaScript, and payment-page designs to harvest card numbers, expiration dates, and CVVs.

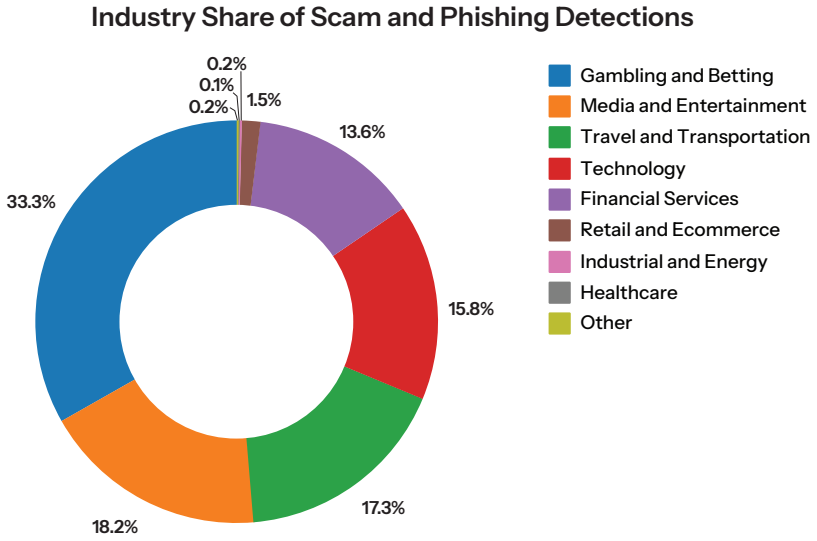


Fig. 21: The gambling/betting and media/entertainment industries experienced the highest frequency of phishing and online scams

Simultaneously, Bolster AI observed scams evolving into multichannel campaigns that spanned social media and apps to create a cohesive marketing funnel. Figure 22 shows one observed attack chain flow.

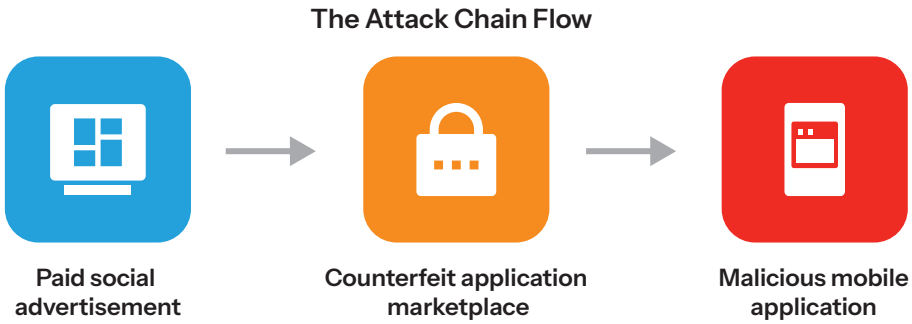


Fig. 22: The attack chain flow observed by Bolster AI in one campaign

The paid ad redirected consumers to a counterfeit app store where they downloaded a malicious app that collected sensitive data and abused device permissions. In a separate incident involving the manufacturing industry, Bolster AI security researchers traced two suspicious social media profiles back to a broader infrastructure network spanning four platforms: social media, encrypted messaging, anonymous email, and a link-aggregation service.

Recruitment scams were prevalent across four major social platforms, with fake hiring content to lure applicants toward phishing forms and collect sensitive data (Figure 23).

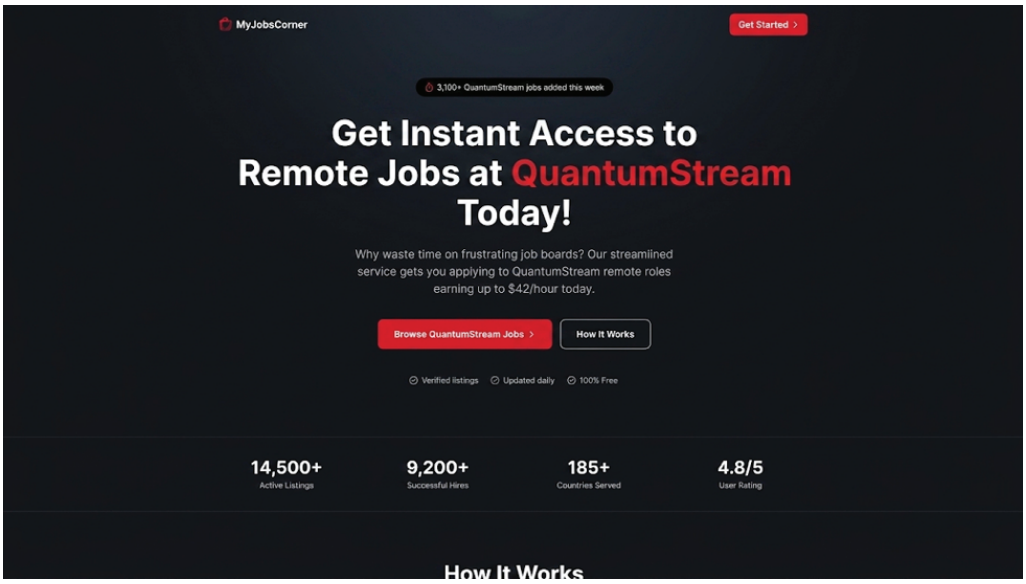


Fig. 23: An example of a recruitment scam

Adversary infrastructure and evasion tactics

Cybercriminals rely on proven infrastructure tactics to build convincing lures, evade security controls, and quickly replace burned assets across fake storefronts, impersonation scams, and credential harvesting operations via:

- **Strategic "domain aging":** Adversaries register look-alike domains in bulk and leave them dormant for weeks or months to build a clean reputation score before launching an attack.
- **Infrastructure reuse and rotation:** Attackers reuse phishing kits, hosting providers, registrars, and payment gateways across multiple campaigns to reduce logistical friction. Once a site is taken offline, replacements go live within hours, and threat actors rotate registrar sets every three to four quarters to avoid alerts and fly below the fraud radar.
- **Built-in page evasion:** Phishing sites employ cloaking, geolocation, user-agent filtering, and conditional redirection. If an automated crawler visits, the site serves benign content, revealing the malicious page only to verified targets.
- **Compromised legitimate domains:** Attackers compromise legitimate websites to host phishing pages, directly inheriting the host domain's established reputation and trust.
- **Dynamic URL patterns:** Attackers continuously alter URL structures to bypass signature-based detections (e.g., shifting between hxxps://sub-domain.directory[.]cfd/get/215962865 and hxxps://sub-domain.domain[.]info/r/Mn6OTy7ejdWeHgg0LZKMSQ).

Shrinking attack windows

Phishing campaigns are becoming increasingly agile as adversaries cycle through infrastructure at record speed. Bolster AI data revealed that the average phishing attack lifespan decreased by more than 50%, dropping from 48 days (Q4 2025) to 22 days (Q2 2026; Figure 24).

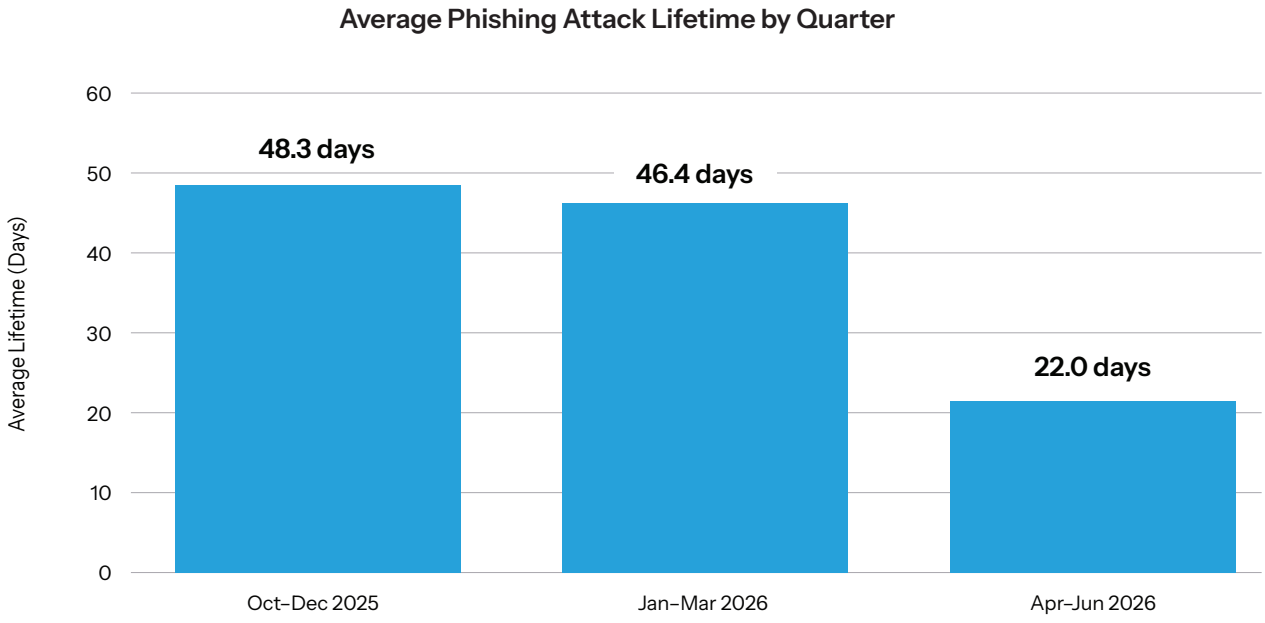


Fig. 24: The lifespan of a phishing attack dropped by more than 50%, from 48 days in Q4 2025 to 22 days in Q2 2026

This compressed window forces security teams into a race against time to detect and disrupt attacks before infrastructure shifts — rendering traditional, reactive defenses far less effective than in the past.

The AI (dis)advantage for both attackers and defenders

AI is transforming phishing into a battle of operational speed for both attackers and defenders. Bolster AI research connects widespread AI access to faster-adapting, multi-vector campaigns that reach larger target pools.

Adversaries leverage AI to generate dynamic templates, produce multilingual lures at scale, and deploy condition-aware evasion — drastically lowering costs while bypassing legacy defenses. Conversely, security teams use AI to compress investigation and reduce workflows from days to less than an hour through automated indicator analysis, infrastructure clustering, and evidence collection, disrupting attacks before victims are impacted.

Regional campaigns: Attackers build global, act local

From a regional perspective, Bolster AI researchers observed that although the supporting infrastructure that threat actors build for phishing and online scams is globally distributed, campaign content is frequently localized by region, industry, language, and payment method. At the same time, telemetry indicates that these are not narrowly targeted attacks, but internet-scale campaigns. Evidence of thousands of phishing URLs, widespread reuse of phishing kits, and repeated impersonation of globally recognized brands suggest that attacks were designed to reach very large audiences.

Analysis from June 2025 through June 2026 highlights how scams have evolved from simple brand spoofing into trusted digital interactions backed by advanced evasion.

APAC

By exploiting regional online gambling ecosystems, threat actors hijacked travel-related search queries to redirect traffic to Southeast Asian gambling platforms. Analysis of 2,157 URLs revealed campaigns targeting cryptocurrency investors, financial institutions, consumer brands, and commerce platforms to drive credential theft, payment fraud, and account takeover.

EMEA

A retail campaign impersonating seven major global brands with billions in annual revenue each, used geolocation and user-agent filtering to restrict access strictly to regional targets. External security analysts received benign pages or error codes, while malicious payloads executed only under specific user-agent conditions — maximizing campaign ROI while evading analysis.

LATAM

Activity in LATAM during the reporting period was notable for two reasons. First, Bolster AI researchers observed that rather than launching single, coordinated regional campaigns, threat actors deployed diverse phishing assets hosted across infrastructure in more than half a dozen LATAM countries. Second, researchers saw that these assets spoofed international shipping providers, enterprise email services, document workflow platforms, and government agencies to harvest credentials.

N. Amer.

Scammers targeted job seekers in engineering and construction with fake recruitment posts on third-party job boards and social media. These campaigns used brand impersonation and cloned login portals for widespread credential harvest and financial fraud.

Akamai Security

Security spotlight

Ransomware is about extortion, *not* encryption

A decade ago, ransomware was malware-centric with a goal of encrypting critical data. Today, however, some actors are skipping encryption if data theft alone provides sufficient leverage. Ransomware has become an operating model, not a type of malware, and the focus is on how to apply enough pressure to get victims to pay. The CISO's job continues to shift from stopping malware to disrupting a complex criminal operation.

Why does ransomware get so much attention? Because it is one of the few cyberattacks that can instantly attract public and regulatory attention. Along with business email compromise (BEC), ransomware is among the cyberattacks most likely to cause financial consequences serious enough to force companies out of business. This makes ransomware a material and recognizable cyber risk, and AI-based techniques are amplifying concerns.

Recent data from the Cloud Security Alliance (CSA) finds that generative AI (GenAI) has drastically compressed the timeline between vulnerability discovery and active exploit deployment from more than nine months to just nine hours. But there's more to the story.

AI is anticipated to be the most significant driver of change in cybersecurity in the year ahead. According to 87% of C-suite respondents to the World Economic Forum Global Cybersecurity Outlook 2026 survey, AI-related vulnerabilities were the fastest-growing cyber risk over the course of 2025, with cyber-enabled fraud and phishing (77%), supply chain disruption (65%), exploitation of software vulnerabilities (58%), and ransomware attacks (54%) also in the top five.

AI is showing up everywhere in our customers' networks, and attackers are using that same trend to blend in. Our Akamai Hunt team uncovered malware trying to hide its traffic behind LLM API endpoints, which is a subtle shift that organizations must watch for as exploitation could lead to full attacker control and data exfiltration. This shift is a reminder that as the ecosystem evolves, so do the visibility gaps. Staying ahead requires paying attention to the quiet changes that reshape risk, understanding the democratization of elite cyberattacks, and ensuring that when a penetration inevitably occurs, the adversary is immediately contained.

Akamai Security

Unpacking threat methodology

The most recent data from the Federal Bureau of Investigation Internet Crime Report 2025 shows that the FBI identified 63 new ransomware variants during 2025 — an average of 5.25 new variants per month. The 10 most frequently reported ransomware variants were:

1. Akira
2. Qilin
3. INC Ransom/Lynx/Sinobi
4. BianLian
5. Play
6. RansomHub
7. LockBit
8. DragonForce
9. SafePay
10. Medusa

As law enforcement agencies achieve more success in taking down criminal groups, big gangs get disrupted and re-form elsewhere, which has caused a fractured landscape as members switch groups and blend techniques. Beyond Akamai ransomware research, which shows threat actor resilience, we see similar findings by MITRE, a nonprofit organization that tracks groups like Akira, Play, Medusa, INC Ransom, and Wizard Spider, among others, and provides templates of their methodologies against their MITRE Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK) framework. The most common techniques include:

- Valid accounts
- External remote services
- PowerShell
- Remote access tools
- Archived or exfiltrated
- Disabled security tools
- Data encryption for impact

5 key questions for executives to ask about ransomware

1. If credentials are stolen tonight, what can the attackers reach? (Access)
2. How far can one compromised workload or endpoint move? (Movement)
3. Can attackers reach and exfiltrate your highest-value information? (Data)
4. Can you detect and contain the attackers before they establish leverage? (Time)
5. Can you continue to operate even if they succeed? (Recovery)

Companies need to think of the hallmarks of ransomware (leverage and extortion) as an outcome of a longer intrusion involving data encryption and/or data theft, access, identity, and movement (Figure 25).

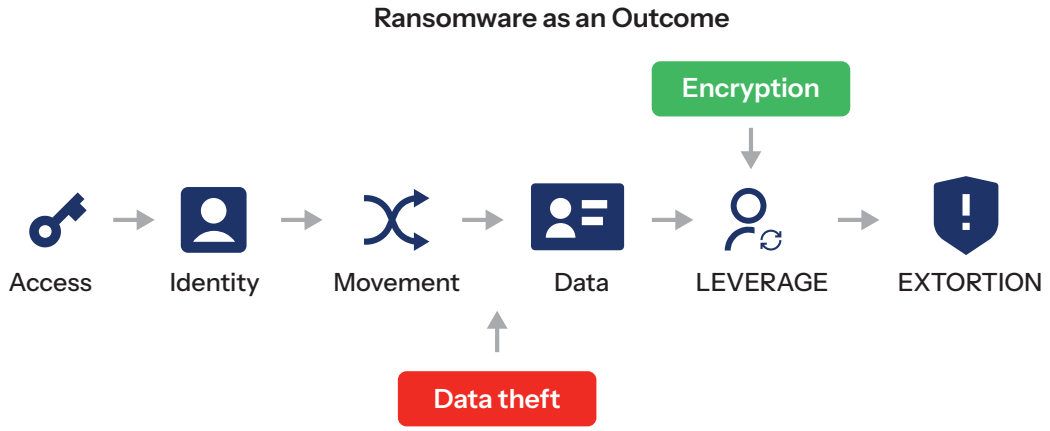


Fig. 25: Ransomware is an outcome of a longer intrusion involving multiple steps

Therefore, one of the best indicators of compromise is finding illegal command and control (C2) infrastructure through tools like recursive DNS, secure web gateways, segmentation, and DNS firewalls. Evidence of beaconing behavior, persistent communication, and live communication between infected hosts and attacker infrastructure confirms compromise.

Top C2 families seen on the Akamai security platform are Cobalt Strike (dominant) followed by APT40, PlugX, APT31/APT29, and BlackMatter Ransomware. Cobalt Strike infrastructure is typically reused across victims; APT-associated infrastructure also shows reuse. Infrastructure reuse is one of the technical manifestations of the ransomware as a business model. Ultimately, CISOs should treat persistent postcompromise C2 as an escalation event, even when there is no evidence of ransomware deployment.

Disrupting a complex criminal operation

The Akamai 2025 Segmentation Impact Study revealed that 79% of surveyed companies experienced at least one ransomware attack over a 24-month period. The companies using microsegmentation had 21.4% faster ransomware containment than average (32.6% faster for organizations worth US$1B or more).

Companies must treat ransomware as the outcome, not the attack, because ransomware isn't how they get in — it's how they get paid. Companies need to have annual exercises that address how to respond to both encryption and data compromise. Additionally, they should work with their legal team to determine if they will pay the ransom. Finally, a great reference for policy development is NIST IR 8374 Revision 1, a ransomware risk management profile aligned to the Cybersecurity Framework 2.0.

There is no single control that solves ransomware. Effective disruption requires a layered approach that prioritizes four core security capabilities: control access, locate and stop C2, defend at the application level, and enforce real-time containment and exposure analysis and response (Figure 26).

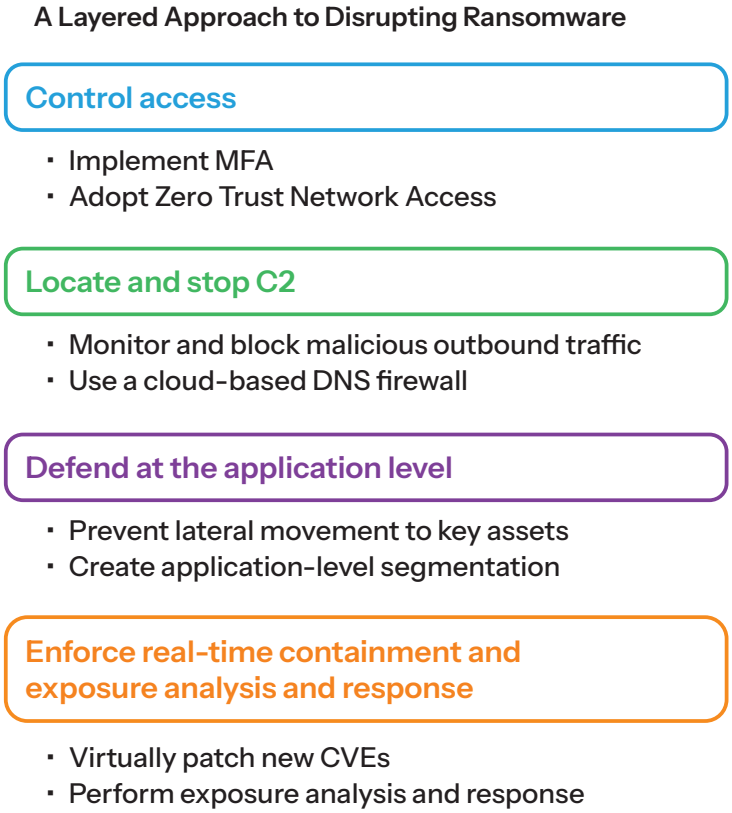


Fig. 26: Disrupting ransomware requires a layered approach that prioritizes four core security capabilities

To execute this, the first line of defense is minimizing the data that threat actors can access and the infrastructure they can infect with encryption malware. This is done with security controls focused on access (such as MFA and information movement visibility) and management (like segmentation). Following Zero Trust principles is worth the effort. Protecting web applications and APIs further reduces the attack surface, while continuous monitoring helps isolate malicious traffic. Finally, regularly testing of backups and targeted components of your business continuity plan validates every layer of this defense.

Ultimately, minimizing ransomware risk is about disrupting a complex criminal operation commoditized by AI, and that requires a focus on building resilience.



Ofer Wolf
Senior Vice President and General Manager of Enterprise Security, Akamai

Akamai Security

Expert insight

Agentic workflows: A new frontier for compliance

Autonomous software agents represent both an operational advance and a compliance frontier. As consumers trust agents to execute transactions on their behalf, and organizations deploy agents for automation, threat detection, and customer service, they introduce novel vectors for data misuse, privacy violations, and accountability gaps. Can existing regulatory frameworks address agents acting on behalf of humans in the best cases, and as malicious actors in the worst instances?

When organizations deploy agentic workflows, they can rely on known risks and familiar compliance approaches (e.g., assess the systems, the data they use, and the impact). When an agentic system processes personal data, it triggers the same obligations that apply to any data controller or processor. The EU AI Act requires determining whether a use case is prohibited. Other laws, such as Colorado's AI Act, establish disclosure obligations. And the NIST Artificial Intelligence Risk Management Framework continues to inform best practices. Yet, agentic workflows amplify familiar risks in unfamiliar ways (Figure 27).

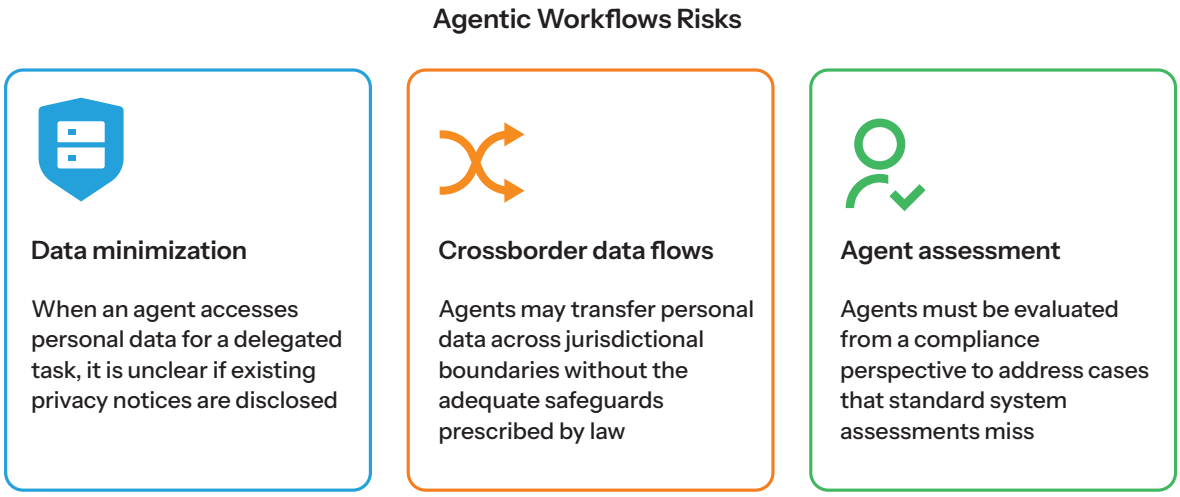


Fig. 27: Agentic workflows amplify familiar risks

Data minimization, for example, is difficult to enforce when an agent's utility depends on broad contextual access. Consent and notice present additional challenges. When an agent accesses personal data to complete a delegated task, it may be unclear whether existing privacy notices adequately disclose this processing, or whether additional notice is required each time an agent is deployed in each context. Crossborder data flows compound these concerns. Agents may transfer personal data across jurisdictional boundaries without the adequate safeguards prescribed by law.

Moreover, the system-focused compliance assessments with which we are comfortable fail if we do not consider that approved systems can easily be deployed by individuals in harmful ways. Agents are not colleagues, but we may need to think about them as though they are colleagues from a compliance point of view to address scenarios that the system-focused assessments don't reach.

Perhaps the most novel compliance challenge with agentic workflows is the emergence of "synthetic identities" that interact with an organization. These are nonhuman actors operating within systems designed for human users. Traditional identity and access management (IAM) frameworks distinguish between human users and service accounts, but agentic AI occupies an uncomfortable middle ground where they act with the autonomy of a human yet lack legal personhood or any innate accountability.

From a compliance perspective, the law increasingly will require corporations to distinguish between (a) unauthorized bots engaged in scraping, credential stuffing, or other wrongful access; and (b) AI agents acting on behalf of consumers with explicit authorization. The first category has familiar bases for defensive actions under cybersecurity and privacy laws. The second, a software intermediary through which a consumer exercises rights, occupies fundamentally different legal ground.

Does this mean an authorized agent's acts should be treated as if they are under the consumer's direct control? Not entirely. Consumer rights (access, deletion, opt-out, and automated decision-making protections) travel with the delegation, and providers cannot use the agent intermediary as pretext to collect more data. But full equivalence is inappropriate. Data protection obligations must account for agents' expanded risk profile. Consent portability remains unresolved, as no mechanism exists for carrying consent across interfaces. And a provider's automated decision-making obligations apply with full force. Consumers do not lose important rights, such as opt-out or human review rights, merely because the session was agent-mediated.

Providers will need to authenticate agents and verify delegation authority, creating a new processing activity that itself must have a lawful basis and be appropriately limited. This may require collecting agent identity attributes, delegation tokens, scope and expiration data, and audit logs. NIST's AI Agent Standards Initiative is working toward standards for auditable verification of agent identity and actions. Verification must now link a nonhuman identity to the consumer, the delegated scope, and the duration of authority. Privileged access management frameworks face similar strain, as the accountability chain from agent to delegating user to authorizing organization remains legally untested.

Finally, when an agent causes harm, existing frameworks offer no clear answer on whether liability rests with the deploying organization, the developer, or the delegating user. Senator Mark R. Warner's AI AGENT Act would impose duties of loyalty and care on AI agent providers, prohibit secondary commercial use of consumer data, and require auditable records of agent actions. Thus, the legislative trajectory may favor holding deployers and providers accountable.

The compliance landscape for agentic workflows is evolving rapidly, but organizations can take practical steps now to prepare, including:

- **Conduct agentic-specific impact assessments.** Treat agentic workflows as a distinct category requiring dedicated assessment, not merely an extension of existing AI or privacy reviews. Evaluate agents both as systems and as tools used by individuals.
- **Update privacy and AI notices.** Revise consumer-facing notices to disclose how personal data may be processed or shared when consumers use authorized agents, and when your organization deploys agents or use AI tools that process consumer data.
- **Implement agent-aware identity governance.** Extend IAM frameworks to treat AI agents as a distinct identity class. Establish protocols for verifying delegation authority, scoping access to delegated tasks, logging actions for audit, and revoking access when delegation expires.
- **Monitor the regulatory trajectory.** The compliance framework for agentic AI is being written now. Actively monitor developments, including the AI AGENT Act and NIST's AI Agent Standards Initiative.



James A. Casey
Vice President and Chief Privacy Officer, Akamai

Authorized agents are not a legal blank slate

In the United States — Senator Warner's AI AGENT Act discussion draft (June 2026) would establish a right of delegation, letting consumers designate custodial user agents to act on services under the same terms as the user.

In the United Kingdom — The Competition & Markets Authority confirmed in March 2026 that consumer law applies equally to human and AI-driven decisions.

Mitigation strategies

Mitigating modern, AI-driven digital deception requires CISOs to deploy proactive controls that address the full spectrum of threats detailed in this report — from rapid-rotation PhaaS and multichannel scams to ransomware extortion, account lifecycle abuse, and agentic AI compliance risks. The dramatic rise in phishing volume, combined with scams eclipsing traditional attack models and the shift of AI crawlers from read-only scraping to transactional execution, presents a complex risk profile for modern enterprises.

To maintain business continuity, safeguard enterprise reputation, and preserve customer trust, CISOs must evolve defensive architectures beyond static perimeter controls and reactive blocklists.

Areas of focus for security leaders

- **Intent-based bot management and account lifecycle abuse:** Evolve bot defenses beyond simple identity verification to intent-based behavioral inspection, stopping low-and-slow credential stuffing, automated account takeover, and credit card validation testing across every stage of the account lifecycle.
- **PhaaS, brand impersonation, and automated takedowns:** Defeat MFA-bypassing PhaaS toolkits and fast-rotating look-alike domains by tracking durable back-end infrastructure signals and automating investigation workflows to disrupt campaigns within shrinking 22-day attack windows.
- **Modern API and agentic AI endpoint abuse:** Safeguard GraphQL, REST APIs, and emerging Model Context Protocol (MCP) endpoints against autonomous AI agents by replacing passive "verified crawler" trust with deep inspection of transactional POST requests.
- **Ransomware extortion and lateral movement:** Counter modern data-extortion operations by deploying microsegmentation to block internal lateral movement, monitoring recursive DNS for C2 beaconing, and conducting routine business continuity and disaster recovery exercises.
- **SMS pumping and telecom exploitation:** Prevent inflated international toll charges driven by rogue telecom collusion by rate limiting automated OTP triggers and adding risk telemetry across account creation and password reset workflows.
- **Promo, welcome bonus, and stored-value abuse:** Protect digital margins against automated promo farming and loyalty point theft by implementing bot management controls that block synthetic account creation and monitor high-risk point redemption endpoints.

- **Search-function abuse and index poisoning:** Enforce deep payload inspection on verified AI crawler traffic to prevent attackers from using trusted bots to seed malicious query strings and poison search indexes.
- **Multichannel scams and counterfeit apps:** Neutralize social engineering funnels originating from paid social ads, fake recruitment boards, and rogue app stores by expanding crossplatform threat intelligence beyond the enterprise perimeter.
- **Shared back-end infrastructure tracking:** Defeat fast-rotating phishing campaigns by tracking durable back-end signals — such as short-lived 90-day SSL certificates, content hashes, and centralized serverless API endpoints — rather than relying solely on reactive domain takedowns.
- **AI agent risks and regulatory noncompliance:** Prepare for evolving regulations like the AI AGENT Act and EU AI Act by extending identity governance frameworks to verify consumer agent delegation tokens, scope access lifespans, and enforce strict data minimization.

By following these best practices, security leaders can protect enterprise assets, maintain operational resilience, and neutralize sophisticated fraud vectors without hampering business velocity or user experience.

As the threat landscape shifts toward autonomous AI tools and industrialized phishing operations, infosec teams must move from static reputation controls to dynamic, intent-focused capabilities. Implementing an integrated security framework ensures that organizations reduce risk exposure, uphold regulatory compliance, and maintain a robust posture against evolving threat vectors. This must be reinforced by processes and validation exercises.

By aligning crossfunctional teams with these comprehensive controls, security leaders can effectively contain lateral movement, neutralize high-yield fraud schemes, and establish a resilient defense capable of safeguarding enterprise revenue and customer trust in an increasingly automated web ecosystem.

Conclusion

Our final State of the Internet/Security report of 2026 has made it clear that digital fraud and abuse have undergone a profound evolution over the past year. What was once dominated by isolated phishing campaigns has matured into an industrialized ecosystem in which threat actors directly abuse brand trust, exploit automated workflows, and target the full lifecycle of modern digital interactions.

To navigate this shifting threat landscape into next year and beyond, organizations should be considering how to address three critical realities:

1. **Escalating loss and structural shifts:** Fraud costs businesses hundreds of billions of dollars globally, driven by synthetic identity fraud, account takeover, and multichannel scams that now surpass traditional phishing in volume.
2. **Infrastructure agility and evasion:** Attackers are leveraging PhaaS, short-lived certificates, and rapid infrastructure rotation to shrink response windows, making back-end signal detection and DNS telemetry vital defense components.
3. **The emerging agentic frontier:** As AI intermediaries and autonomous agents redefine digital interactions, security leaders face an expanded attack surface that encompasses zero-click interactions, MCP endpoints, and synthetic customer personas.

Ultimately, defending against internet-scale fraud in this new era will require security leaders and professionals to evolve beyond traditional technical perimeter defense and move toward crossfunctional collaboration and AI-ready governance. Protecting future revenue and reputation requires treating digital trust not as a static asset to safeguard, but as a dynamic capability that must be continually verified, defended, and earned across every digital interaction.

Methodologies

Akamai Account Protector data

The data in this section is generated by our Account Protector security solution, which protects against account abuse across the entire lifecycle of an account by using machine learning and a significant dataset of risk and trust indicators to determine the legitimacy of a user request. Account Protector has unique insight into real-time behavioral analysis to determine the risk level associated with each detected login attempt. When suspicious or anomalous behavior is detected, immediate mitigation options are available to prevent account compromise.

This sample of data covered a 30-day period between July 2026 through August 2026.

Akamai Secure Internet Access data

The data in this section was generated by our Secure Internet Access product, which is a cloud-based secure web gateway designed to help users easily connect their devices to the internet in a secure manner. Secure Internet Access has a unique visibility into billions of daily DNS queries from enterprise organizations with large amounts of users. This data was measured by both the number of queries and the number of distinct threats detected. Our security teams use this data in-house to research attacks, flag malicious behavior to notify customers, and feed additional intelligence into Akamai's security solutions.

This data covered the seven-month period from January 1, 2026, through July 31, 2026.

Bolster AI data

The data in this section was provided to Akamai through a collaboration with Bolster AI. Bolster provides AI defense against phishing and brand impersonation attacks using large structured datasets and AI models. Their security platform is trained on more than 10 billion data points and leverages LLMs and computer vision to analyze more than 4 million websites each day to detect and respond against threats in real time.

This data covered the 12-month period from July 2025 through June 2026.

Akamai Security

Guest contributors



James A. Casey
Vice President, Chief Privacy Officer, Akamai

James A. Casey is Vice President and Chief Privacy Officer at Akamai and heads the Akamai Global Data Protection team. Jim has served as in-house counsel for technology companies for the past 20+ years and has significant experience in supporting new technology and product initiatives in the internet, cybersecurity, information services and analytics, and telecommunications industries. Jim provides legal counsel in a variety of areas, including technology law and regulation, public policy, privacy and artificial intelligence governance, import/export and trade compliance, and cybersecurity. Prior to moving in-house, Jim's law firm experience focused on supporting technology regulation and initiatives in the data, telecommunications, and internet industries, both domestically and internationally, as well as supporting technology and telecommunications projects with native peoples in the United States and around the world.



Moshe Cohen
Principal Security Analyst, Akamai

Moshe Cohen is a Principal Security Analyst at Akamai. Since joining the company in 2018, he has specialized in maintaining products security posture, developing tools to discover emerging threats, and directly assisting customers with critical threat assessments. He is committed to continuous learning while maintaining a healthy balance between his passion for security and his personal life.



Shiran Guez
Director of Engineering, Akamai

With a career spanning more than two decades, Shiran Guez has established a solid networking and telecom industry foundation. A passion for technology and his entrepreneurial mindset have driven his success across various domains. In 2014, Shiran joined Akamai, managing information security and security operations within the Enterprise Security division. Shiran is a firm believer in continuous learning and relishes new challenges that enable him to expand his expertise. Outside the professional realm, Shiran is a devoted husband and father of three. He recognizes the significance of maintaining a balance between work and personal life, and his commitment to both aspects demonstrates his determination and adaptability in an ever-evolving industry.



Patrick Sullivan
Chief Technology Officer,
Security Strategy, Akamai

Patrick Sullivan is Vice President and Chief Technology Officer of Security Strategy at Akamai, where he has led security innovation for more than 20 years. Patrick is an award-winning leader with 30+ years of experience in cybersecurity who works with global enterprises to design security architectures and defend against evolving threats. A frequent speaker at RSAC, Black Hat, and other major conferences, Patrick brings unique insights from his frontline work responding to attacks targeting the world's top organizations. Patrick is a contributor to Infosecurity Magazine, holds multiple certifications (including CISSP), and previously held leadership roles at DISA, AT&T, Savvis, and Cable & Wireless.



Ofer Wolf
Senior Vice President and General Manager of Enterprise Security, Akamai

Ofer Wolf is Senior Vice President and General Manager of Akamai's Enterprise Security business within Akamai's larger Security Technology Group. He is responsible for Akamai's Zero Trust–focused security solutions, including microsegmentation, Zero Trust Network Access, and multi-factor authentication. He is based in Tel Aviv, Israel.

Wolf joined Akamai in October 2021 as part of Akamai's acquisition of Guardicore, where he was COO, in charge of International Sales and Global Operations. Before joining Guardicore, Wolf was CEO at Sentrix, a web application security start-up. He also held leadership positions in U.S. publicly traded enterprises, including Vice President of Worldwide Technical Services; was part of the senior management team at Check Point; and was Vice President and General Manager at Comverse, a leading value-added services and billing supplier.

Credits

Research director

Kimberly Gomez

Writing and editing

Charlotte Pelliccia
Badette Tribbey
Maria Vlasak

Review and subject matter contribution

Lauren Barraco
Tanya Belousov
James A. Casey
Moshe Cohen
Guy Gayer
Shiran Guez
Ariel Kibudi
Reuben Koh
Richard Meeus
Patrick Sullivan
Doron Tzur
Steve Winterfeld
Ofer Wolf

Data analysis

Vikas Jha
Afzal Sheik
Chelsea Tuttle

Promotional materials

Carolyn Lange

Marketing and publishing

Carolyn Lange
Georgina Morales Hampe

State of the Internet/Security

Read back issues and watch for upcoming releases of Akamai's acclaimed State of the Internet/Security reports.

Akamai threat research

Stay updated with the latest threat intelligence analyses, security reports, and cybersecurity research.

Akamai security research

Read the Akamai security research blog for a rapid response perspective on today's most important research.

Access data from this report

View high-quality versions of the graphs and charts shown in this report. These images are free to use and reference, provided that Akamai is duly credited as a source and the Akamai logo is retained.



Akamai is the cloud company that powers and protects an AI-driven world. Our cloud platform extends high-performance cloud computing from the core to the edge, enabling organizations to build and scale next-generation AI applications and providing comprehensive, multilayered security to safeguard enterprises against evolving cyberthreats. Learn more about Akamai's cloud computing, security, and content delivery solutions at akamai.com and akamai.com/blog, or follow Akamai Technologies on X, and LinkedIn. Published 10/26.