State of the Internet | V12 Issue 02 | 2026

Akamai Security

SOTI

AI-Empowered Botnets and API Visibility Gaps

Attack Trends in Financial Services

Contents

Akamai Security

Introduction

It's been two years since the State of the Internet/Security report analyzed the financial services industry. Since then, financial services has remained one of the most consistently targeted sectors for global threat actors. Digital transformations in the industry — including open banking and API-first architectures — have expanded the attack surface beyond traditional defenses. Financial institutions are currently bearing a disproportionate share of the global cyberattack burden and facing a wave of activity that is not only rising in volume but also increasing in complexity.

The year 2025 marked a critical inflection point in the industrialization of cybercrime. We observed a milestone for hyper-volumetric distributed denial-of-service (DDoS) attacks, with record-breaking packet rates and the proliferation of DDoS as a service (DDoSaaS). While major security victories — such as the takedown of the Aisuru and Kimwolf botnets — temporarily neutralized millions of compromised Internet of Things (IoT) devices, new adversaries are already rising to take their place.

Artificial intelligence (AI) is no longer an emerging capability; it's embedded in core business and security workflows. While financial firms leverage machine learning (ML) for fraud detection and personalization, attackers are manipulating AI to empower botnets that mimic legitimate behavior with near-perfect accuracy to defeat traditional defenses. We are witnessing a strategic move from simple automation to agentic AI autonomy. Today, malicious agents can navigate complex applications to perform deep contextual analysis and exploit sensitive data. Crucially, AI does not replace traditional security risks, it amplifies them — particularly through vulnerable API endpoints that serve as the connective tissue for these new models.

By analyzing trends over time, diving deeply into 2025 data, and looking ahead, this report helps financial services professionals understand and adapt to the evolving threat landscape. To stay resilient, financial organizations must move beyond basic compliance and embrace layered defenses powered by behavioral intelligence. In this report, we explain the theoretical risks and provide actionable insights, including the intelligence needed to maintain security and trust in an increasingly volatile environment.

Akamai Security

Guest column

Financial services in the age of AI: Interconnected risk, collective defense

The financial services sector's use cases for AI to improve cybersecurity, reduce fraud, and address compliance concerns are expanding constantly. From AI-enabled vulnerability detection to fraud detection to know your customer/anti–money laundering (KYC/AML) automation, AI is changing the pace and effectiveness of back-office work. Likewise, agentic AI is transforming customer-facing operations, enabling banking platforms to coordinate end-to-end workflows and support hyper-personalization.

As these tools increase the financial services sector's ability to operate at scale and speed, they also expand its attack surface and enable threat actors' capabilities and impact.

In response to the opportunities offered by AI-enabled cyber operations, as well as the risks posed by AI vulnerabilities and AI-fueled attacks, the Financial Services Information Sharing and Analysis Center (FS-ISAC) has made emerging technologies, including AI, a strategic priority, along with supply chain risk, geopolitical shifts, and fraud. As this Akamai SOTI Security report makes clear, these are not discrete issues — they are interrelated and reinforce one another.

For example, as geopolitical shifts increase, we've seen a growing number of nation-state actors, hacktivists, and ideologically driven coalitions employ agentic AI to compromise flaws with astonishing speed, such as generating functional exploits upon vulnerability disclosure or analyzing software patches to develop exploits before a CVE is formally disclosed.

In this evolving threat environment, risks aren't stacking — they're compounding. As the financial sector leverages AI, defenders need access to real-time threat intelligence across the financial ecosystem. We must exchange, learn from, and embed insights directly into detection, response, and resilience strategies.

The faster we share indicators, tactics, and lessons learned, the faster we raise the cost for attackers and reduce systemic risk across the sector.

In this hyperconnected threat landscape, effective defense is collective defense.

John "JD" Denning
Chief Security Officer, FS-ISAC

Akamai Security

Key insights of the report

- Financial services remains the most targeted industry by Layers 3 and 4 DDoS attack events, which are growing far more persistent and adaptive. The median attack duration is up by 738% since 2024 on a global basis and by 1033% in Europe, the Middle East, and Africa (EMEA), driven by AI-powered methods, legacy system flaws, and rapid digital banking expansion.
- Cyberattack methods against financial services vary significantly by region: EMEA is the primary target for Layers 3 and 4 DDoS (62%), Asia-Pacific (APAC) is the most targeted for Layer 7 DDoS (52%), and in North America, web attacks are the most prevalent (44%).
- Pro-Iran hacktivist activity, including both isolated and coordinated DDoS campaigns, has increasingly been targeting financial institutions through online banking, payment systems, and API disruptions.
- During 2025, banking was the primary target for web attacks: 60% of total web attacks and 83% of attacks against API endpoints targeted the banking vertical.
- Advanced bot activity surged by 147% in late 2025 — in one extreme case study, a staggering 96% of all site traffic was identified as malicious scraping bots.

DDoS remains finance's top threat

The financial services industry continues to face an escalating wave of DDoS activity that reflects both the industry's critical role in the digital economy and the growing sophistication of attackers. Financial institutions are bearing a disproportionate share of the burden, showing not just increasingly higher attack volumes but also more advanced, multilayered campaigns. The persistent nature of these attacks stresses the importance for financial institutions to maintain heightened awareness and adaptive defense strategies to counter increasingly sophisticated DDoS tactics.

Layers 3 and 4 attacks continue to rise

By comparing financial services with other industries, we see that it remains the most attacked industry by Layers 3 and 4 DDoS. Additionally, we've observed it to be the only industry in the top five that has experienced an increase of these types of attacks year over year (Figure 1). Three years ago, we wrote about how the number of Layers 3 and 4 DDoS attacks on financial services had just surpassed the number of those attacks on games, but now the difference in attack event count is even more significant.

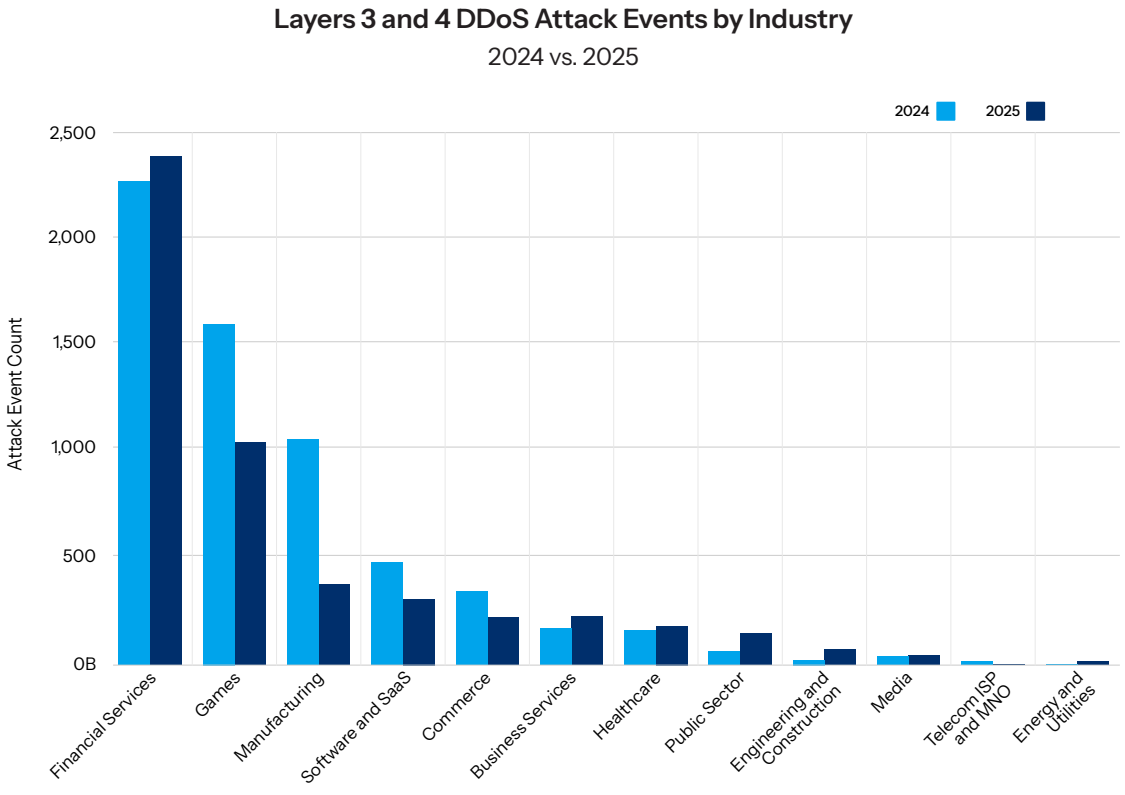

Fig. 1: Unlike other industries targeted by Layers 3 and 4 DDoS, the number of these attacks on financial services has increased by 5.2% year over year

Escalating threats: Bigger, longer, and more complex

In our recent Apps, APIs, and DDoS SOTI Security report, we analyzed the increase in Layers 3 and 4 DDoS attack events across industries over the last year, and we explored how DDoS attacks, in particular, have diversified in attack methodology. We see this same trend when zooming in on the financial services industry. In addition to the increasing Layers 3 and 4 event count, we've observed both the growth of massive attack event sizes (Figure 2) and attack event duration (Figure 3). This observation highlights the recent increase in intensity, complexity, and duration of DDoS attacks across the board.

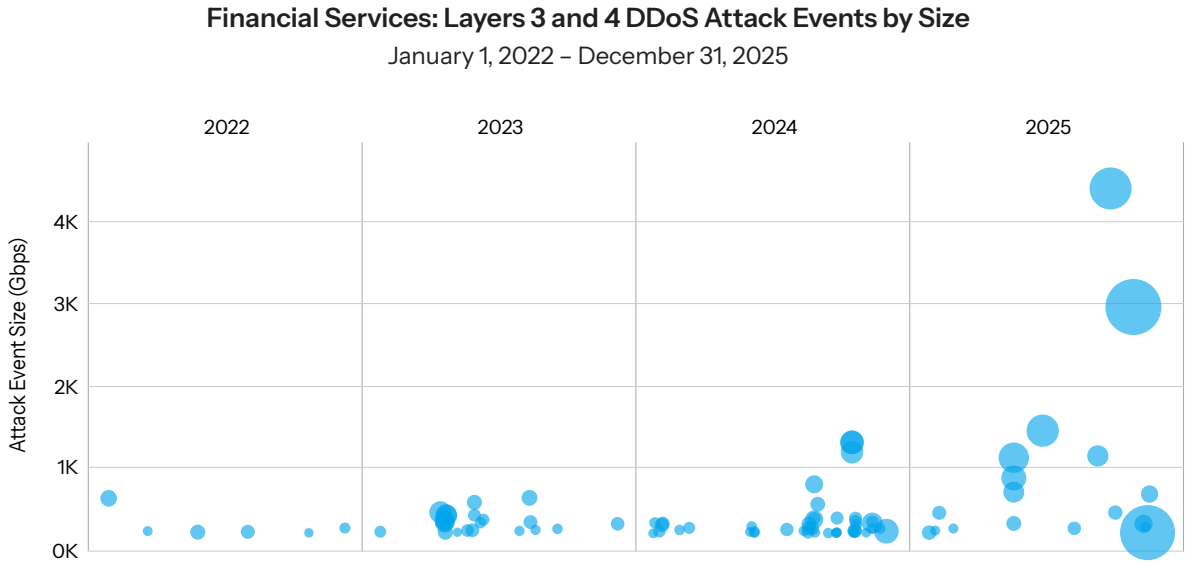

Fig. 2: The maximum DDoS attack event size against financial services increased by 236% from 2024 to 2025

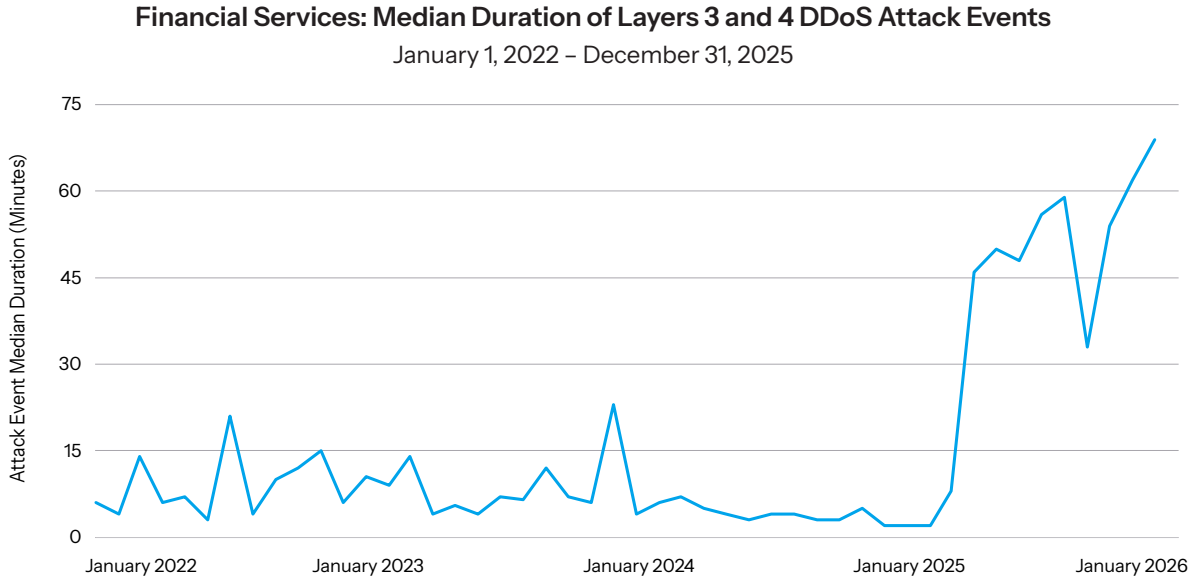

Fig. 3: The median duration of DDoS attack events against financial services increased by 738% between 2024 and 2025

Several structural factors likely contributed to this shift. Financial institutions remain high-value targets for attackers seeking disruption, financial gain, or reputational damage. Simultaneously, the rise of AI-enhanced attack frameworks allows adversaries to adapt to mitigation strategies in real time, dramatically increasing both the scale and complexity of threats. Legacy systems add more risk, as they often lack modern protocols and are difficult to patch or integrate securely with cloud infrastructures. Furthermore, the continued expansion of digital banking, the growing dependence on third-party services, and the global rise in hacktivism continue to broaden the attack surface.

These same conditions also help explain why financial institutions face a higher volume of DNS attacks annually than do most other industries. The complexity of these environments combined with the critical nature of their services also means that these attacks often take longer to mitigate.

DNS attacks: An increasing share of observed volumetric attack vectors

DNS floods continue to be a leading volumetric DDoS attack vector in the financial services industry, and the numbers are increasing (Figure 4). Attackers are drawn to this vector because it offers a high return on investment (ROI) and takes advantage of the ongoing evolution of online infrastructure and service architectures.

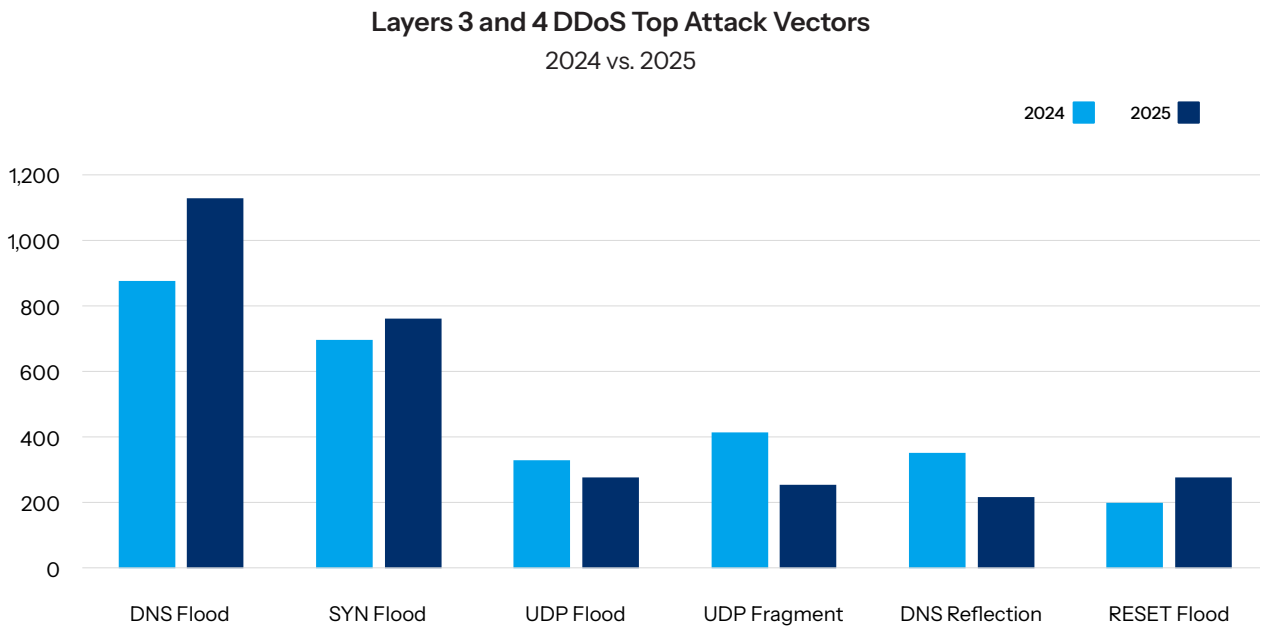

Fig. 4: The DNS flood attack vector has increased by 29% year over year and is the leading volumetric DDoS attack vector

By overwhelming DNS infrastructure with massive volumes of queries or exploiting amplification techniques, attackers can disrupt the foundational layer that supports all digital operations. DNS floods target a centralized point of failure which can take down an entire ecosystem of services, including websites, email systems, and internal tools, rather than just single applications. This significantly amplifies the operational impact.

As we've observed for the last several years, financial institutions are especially vulnerable because their main platforms (e.g., online banking, payment processing, and trading) depend on reliable and instant DNS resolution to preserve trust and transactional continuity. Moreover, financial service environments tend to be complex, interconnected, and tightly regulated. As a result, mitigation times can be longer than in other industries that more commonly experience multi-vector attacks.

Banking tops the charts in DDoS for financial services

Within the financial services industry, banking remains the leading vertical targeted by Layers 3 and 4 attacks, as well as by Layer 7 DDoS attacks, by a wide margin (Figure 5). This is due to its central role in the industry, including its combination of public visibility, high-value data, and deep interconnectivity with other financial and fintech services.

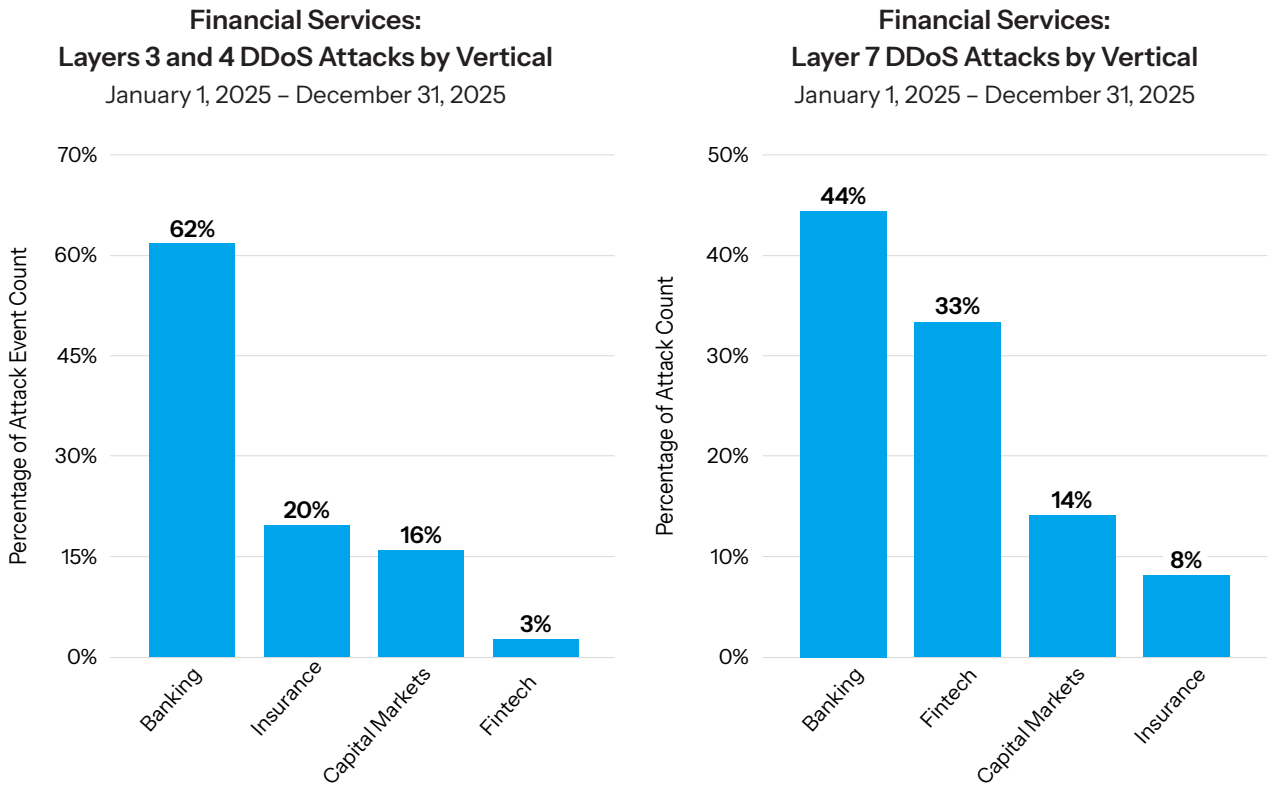

Fig. 5: Banking continues to remain the most targeted vertical in financial services, accounting for 62% of Layers 3 and 4 DDoS attack events (left) *and 44% of Layer 7 DDoS attacks* (right)

Banks serve as the primary gateway to the broader financial ecosystem. This means that attacking a bank can disrupt an entire network of dependent services that use bank credentials or APIs, such as third-party payment and investment apps. Cyberattacks in this space not only compromise direct operations but also can extend through third-party and supply chain credential breaches, amplifying the impact.

In addition to DDoS attacks, the critical, high-visibility infrastructure of the financial services industry as a whole makes it a prime target for ransomware attacks that generate extreme disruption and theft. The Akamai Segmentation Impact Study: Financial Services Industry report indicates that 79% of financial institutions have faced ransomware attacks in the past two years. Financial services institutions face some of the world's highest ransomware risks, yet less than half have adopted advanced technologies — a maturity gap that leaves them vulnerable.

Recent increases in geopolitical tensions have also fueled hacktivism, with attackers prioritizing banks to maximize visibility and brand exposure. For example, a bank website outage is far more public and alarming than disruptions within hidden, back-office financial systems.

Fintech's API growth heightens Layer 7 DDoS and third-party risk

While banking has a strong lead over other verticals in the financial services industry, the fintech vertical ranked particularly high in Layer 7 DDoS attacks for 2025, especially in comparison with insurance (Figure 6). The fintech vertical comprises technology-driven financial services that emphasize accessibility, speed, and user-friendly digital experiences.

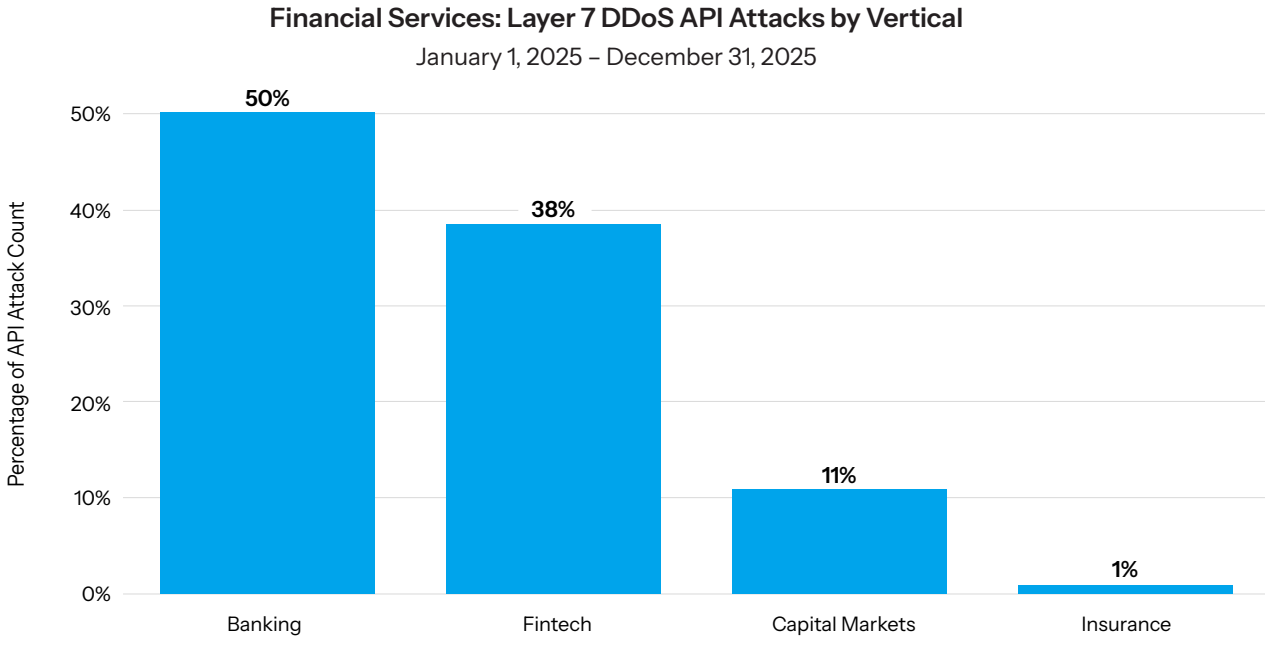

Fig. 6: The fintech vertical was hit with 38% of the Layer 7 DDoS attacks that targeted financial services API endpoints in 2025

Compared with other verticals, fintech is especially vulnerable to Layer 7 DDoS attacks and third-party service failures because of its API-first, latency-sensitive, and highly integrated business model. The proliferation of APIs has led to reduced development timelines, which often leads to shadow APIs — a major vulnerability and a primary magnet for attackers.

An attack on a single fintech API can potentially disrupt services across multiple partner institutions. Similarly, fintechs have an extensive blast radius given their nature (connecting banks and consumers). For example, a vendor breach cyberattack in December 2025 exposed data from more than 700 financial institutions, demonstrating how a single third-party exploit can maximize impact and allow attackers to extort multiple organizations at once. This attack illustrates how third-party risk in the fintech vertical can substantially compromise core banking data on a massive scale.

Regional trends: Layers 3 and 4 DDoS attacks

During 2025, EMEA was the most targeted region for Layers 3 and 4 DDoS attack events against the financial services industry, experiencing 62% of total attack events, compared with North America at 26% and APAC at 12% (Figure 7). EMEA's upward trajectory began in Q1 2024 when it first surpassed North America in Layers 3 and 4 DDoS attack event counts as reported in our SOTI report, Fighting the Heat: EMEA's Rising DDoS Threats. Since then, EMEA has maintained the lead.

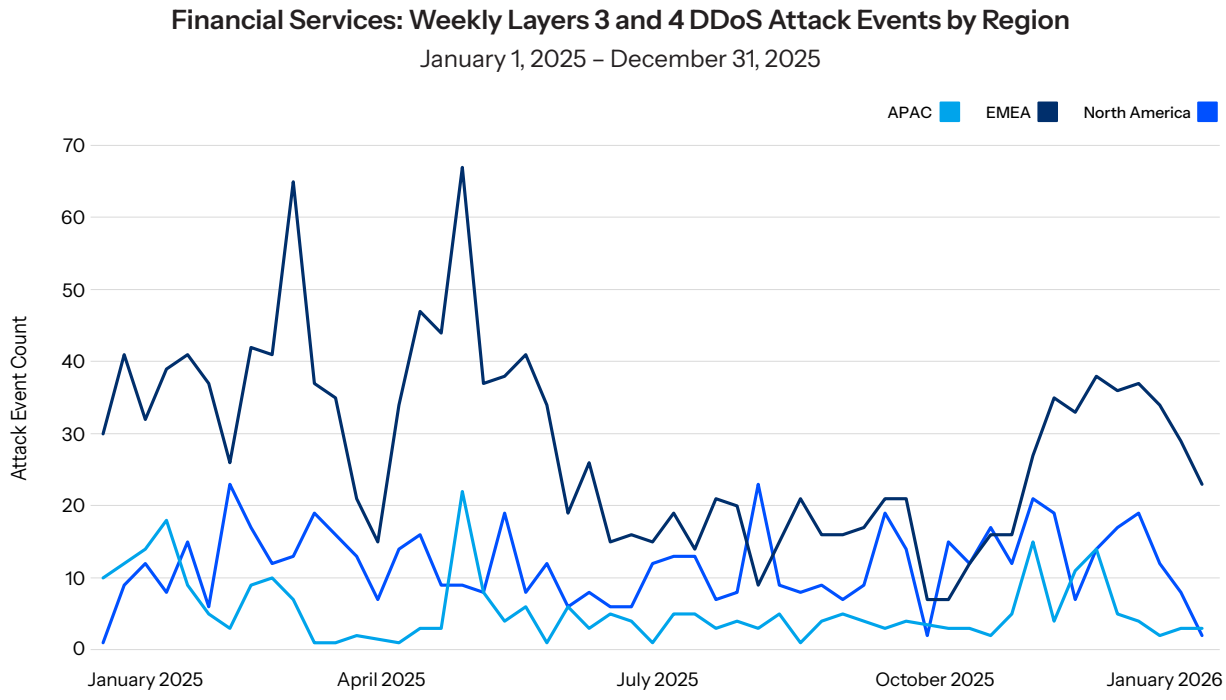

Fig. 7: During 2025, EMEA remained the primary target for Layers 3 and 4 DDoS attacks

When successful, Layers 3 and 4 DDoS attacks are loud and obvious which can be immediately felt by consumers and businesses who are unable to pay, borrow, or receive money because of high-friction access or no access at all.

The increasing size, frequency, and duration of attacks against the financial services industry contributed to the intensity of activity and aligned with threat actors' motives. For example, the median duration of attack events in EMEA from 2024 to 2025 increased by 1033%, from 3 minutes to 34 minutes (Figure 8).

EMEA Financial Services: Median Duration of Layers 3 and 4 DDoS Attack Events

January 1, 2022 – December 31, 2025

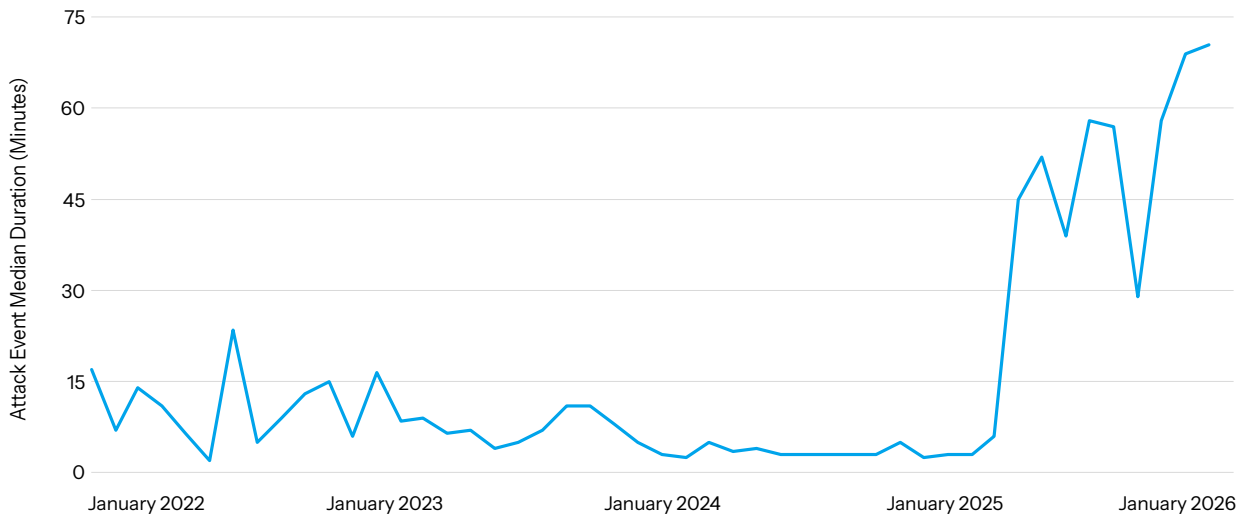

Fig. 8: The median duration of Layers 3 and 4 DDoS attacks in EMEA increased by 1033% in 2025 to 34 minutes

Attack activity related geopolitical events and hacktivist agendas

Surges in Layers 3 and 4 DDoS activity in EMEA often correlated to geopolitical events and hacktivist agendas, with malicious traffic generated from Iran and Russia aimed at the financial backbone of economies in the region. Opportunistic cybercriminals also used the intensifying unrest as a smoke screen for their own disruptive campaigns against financial services firms.

Within financial services, 71% of Layers 3 and 4 DDoS attack events in EMEA and 92% in APAC targeted the banking vertical. When directed at banking, the impact of attacks is amplified when the fallout spreads from economic to symbolic as hacktivist and nation-state groups alike aim to undermine public trust in the government's ability to maintain financial stability.

The geopolitical atmosphere in APAC is at its most unstable in decades, with relatively small-scale conflicts and scuffles leading to multiple splinter hacktivist groups coordinating attacks that target critical infrastructure like financial entities. During 2025, these included conflicts between India and Pakistan in May, Thailand and Cambodia in July, and persistent South China Sea territorial disputes between various countries.

As volatility continues across the global geopolitical theater, Akamai researchers see these Layer 3 and Layer 4 trends persisting.

Regional trends: Layer 7 DDoS attacks

During 2025, the financial services industry experienced several large-scale Layer 7 DDoS attack campaigns during the first half of the year in response to increased tensions in several geopolitical hot spots. Possible causes for a spike in numbers during February and March include an uptick of military drills in the Taiwan strait, naval standoffs in the South China Sea, and skirmishes along the borders of Kashmir (Figure 9).

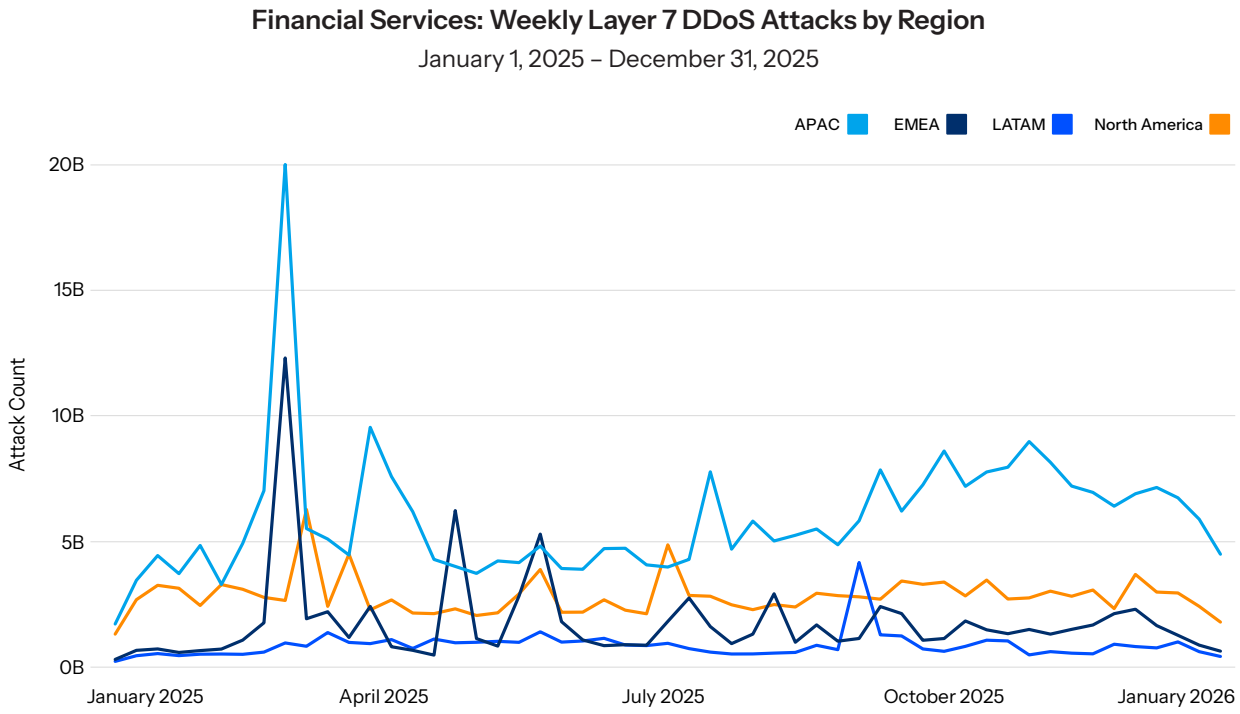

Fig. 9: Layer 7 DDoS attacks tend to ebb and flow with geopolitical events

Automation, AI, and DDoSaaS became force multipliers for threat actors, driving increases in attacks that don't seem dramatic when averaged over time, but with spikes that could have significant impact on networks.

On a global basis, Layer 7 DDoS attacks against financial services increased by nearly 25% year over year, and APAC experienced a notably higher level of Layer 7 DDoS attacks than other regions. In 2025, 52% of Layer 7 DDoS attacks against the industry were in APAC, 26% in North America, 14% in EMEA, and 7% in Latin America (LATAM). Even though attack activity in LATAM was comparatively low, the region experienced the second highest year-over-year increase in Layer 7 DDoS attacks at 40%. Additionally, the trend that we observed from 2023 through 2024 (and discussed in our 2024 financial services SOTI) persists in 2025: The number of Layer 7 DDoS attacks that specifically target APIs continued to increase (Figure 10).

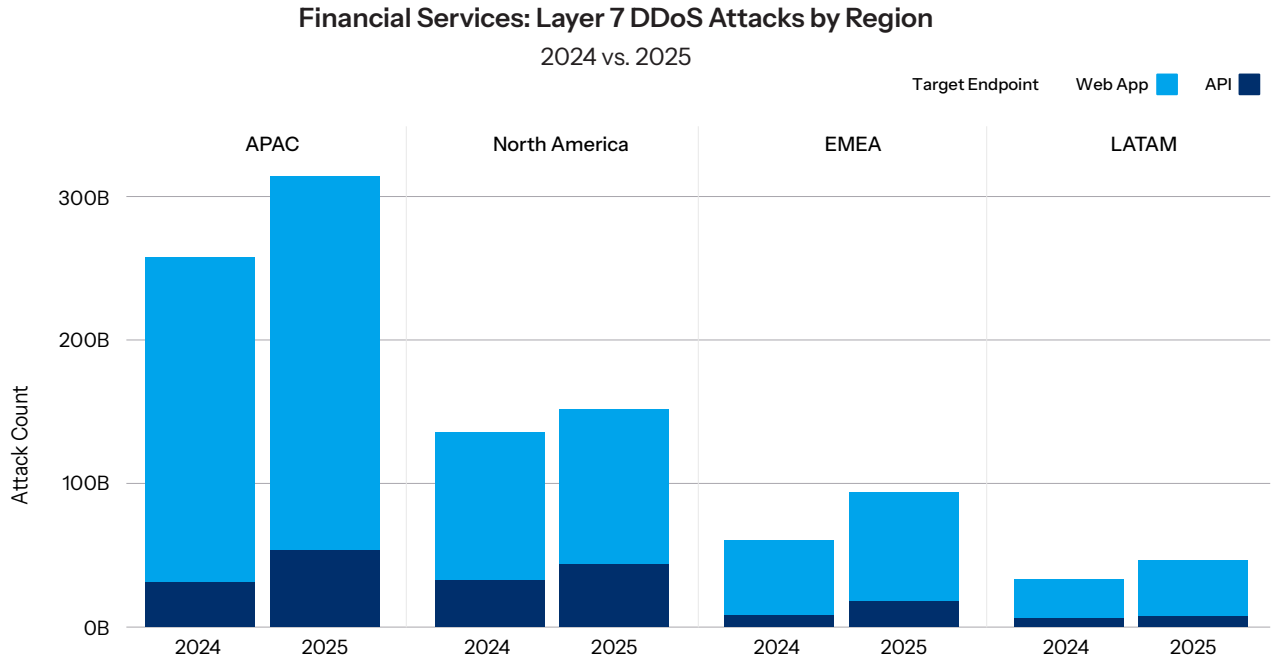

Fig. 10: Attacks targeting APIs continued to rise year over year, and APAC was the most targeted region for Layer 7 DDoS attacks in 2025, experiencing 314 billion attacks, more than double the next closest region (North America)

Banks in APAC are widely considered the fastest growing in terms of digital adoption, real-time payment volume, and neobanking (digital-only banks). APIs are powering everything from mobile wallets to open banking transactions as the region moves rapidly to transition to a cashless society.

In a very competitive financial market, where speed to release new services and products is a critical differentiator, development cycles shrink to deliver working code quicker. As a result, shadow APIs and the practice of AI-assisted coding (aka vibe coding) proliferate within the region. The fallout is troubling.

According to Akamai's 2026 API Security Impact Study, 77% of financial services IT and security leaders in APAC report having a full inventory of APIs, but only 27% also know which APIs return sensitive data. In mature financial markets like Singapore, banks have massive API portfolios that create visibility challenges; security teams often don't know exactly how many active endpoints they have. APAC's financial sector is pushing toward open banking and regional integration, creating unmapped API sprawl. As a direct result, the risks that concern them the most are APIs that leak sensitive information and attackers that are exploiting insecure API endpoints linked to AI technologies.

Additionally, insufficient DDoS protection of the digitized states of maturing organizations, alongside botnets for hire, other user-friendly tools, and a tense geopolitical climate, create favorable conditions for even less technologically savvy attackers to launch DDoS attacks.

How pro-Iran hacktivists are using multi-vector DDoS to disrupt digital banking

The persistence and increasing complexity of coordinated cyber campaigns suggests a permanent shift in the financial services threat landscape. In retaliation to Operation Epic Fury, pro-Iran hacktivists have been conducting multi-vector cyberattacks that include DDoS. While Iranian-aligned groups have a history of using DDoS tactics — most notably, Operation Ababil (2012–2013), which cost various financial institutions tens of millions of dollars — current anxieties are spiking due to the increasingly sophisticated threat landscape.

In March 2026, as the hacktivist groups began naming U.S. banks as specific targets, the U.S. Department of Justice and financial regulators like the Financial Industry Regulatory Authority issued high-priority alerts. Some of these hacktivist groups include Keymous+, DieNet, Handala, and the Cyber Islamic Resistance (CIR).

Keymous+ and DieNet: Targeting financial APIs

Both Keymous+ and DieNet launched coordinated volumetric and application-layer DDoS attacks targeting financial APIs and login portals. This prevented customers from accessing online banking and disabled mobile payment apps. More specifically, the campaign combined HTTP floods and DNS amplification to paralyze critical infrastructure throughout the United States and the Middle East.

Handala: Using DDoS as a tactical smoke screen

The group known as Handala has also been a strong driver of these disruptions, often using DDoS as a smoke screen for more destructive operations. Handala has often been recognized for its high-profile data-wiping attacks (such as on Stryker, a U.S. multinational medical tech company). But the group has also been linked to a surge of DDoS attacks mainly targeting U.S. and Israeli energy and telecom industries to cause high-visibility service outages.

Yet, Handala also continues to attack the financial services industry. For example, this group claimed to have hacked Verifone on March 11, 2026. They alleged the theft of transaction and financial data, causing widespread disruptions across various payment systems. Their operations also extended to a Hasidic Jewish community, where they reportedly stole documents related to financial cooperation.

CIR: Conducting retaliatory attacks on financial infrastructure

While Keymous+, DieNet, and Handala acted as high-volume spearhead attackers, CIR provided broader strategic coordination. Specifically, a coalition referred to as the Islamic Resilience Cyber Axis (via CIR) launched an "Electronic Operations Room." This "room" is not a physical location but a unified digital platform (primarily hosted on Telegram) designed to coordinate multiple disruptive operations with various proxy groups. Their DDoS attacks, alongside similar attacks by other pro-Iranian hacktivists, focus on financial infrastructure (as well as on the infrastructure of other industries) and disrupt banking portals and online payment systems.

When U.S. or Israeli forces launch military strikes, the group responds by sending massive volumetric DDoS waves that target major U.S. consumer banks and Israeli stock exchange infrastructure to trigger immediate economic disruption. As of now, the coalition continues to conduct these DDoS attacks in a retaliatory fashion.

Threats in depth

DNS is financial services' hidden attack surface

Across the banking, wealth management, insurance, and fintech industries, security leaders noted that DNS infrastructure often outlives the systems, products, and business units it originally supported. Cloud adoption, software as a service (SaaS) platforms, open banking integrations, and developer self-service have accelerated the creation and abandonment of DNS records. In an industry where customer trust and service availability are central, DNS represents a persistent, high-impact risk to the external trust layer.

The U.S. Cybersecurity and Infrastructure Security Agency (CISA), the Federal Bureau of Investigation (FBI), and FS-ISAC have consistently warned that DNS manipulation and domain infrastructure abuse play a central role in phishing, fraud, and infrastructure hijacking campaigns. Within the financial sector, FS-ISAC has also highlighted domain abuse and subdomain takeover as recurring enablers for brand impersonation and credential harvesting attacks that are targeting banks and payment platforms.

Data from Akamai DNS Posture Management reinforces these concerns. At enterprise scale, basic DNS misconfigurations — such as dangling CNAME records, leaked internal network information, and disabled registry protections — continue to leave large numbers of domains vulnerable to takeover and abuse. Within regulated sectors such as financial services, more than 85% of observed domains failed at least one foundational DNS control, including Start of Authority (SOA) integrity, Certificate Authority Authorization (CAA) enforcement, or DNSSEC, while nearly half lacked modern email authentication protections.

Security teams frequently described incidents using familiar operational shorthand: "It's not DNS. It can't be DNS. Oh, [expletive], it was DNS." What the data reveals is often not a new category of vulnerability, but how frequently unnoticed DNS drift directly contributes to incidents discovered elsewhere.

DNS and DNS-adjacent risk areas

Mergers and acquisitions

This pattern becomes particularly visible during periods of organizational change. Mergers and acquisitions — common across global banking and fintech — repeatedly surface previously unknown DNS exposures. During integration efforts, institutions often uncover large volumes of orphaned DNS records associated with decommissioned cloud workloads, legacy trading platforms, retired CDNs, marketing domains, and former technology vendors. DNS entries can remain resolvable long after the underlying services have disappeared, creating potential exposure to subdomain takeover, impersonation attacks, and unauthorized certificate issuance. In financial services environments where customers rely on domain trust for online banking, payments, and account communications, these issues can quickly translate into fraud risk and reputational damage.

Modernization efforts

Outside of mergers and acquisitions, modernization efforts create similar conditions. Cloud transformations frequently leave behind stale CNAME records and abandoned domains. Email migrations disrupt SPF, DKIM, or DMARC alignment long after projects are considered complete. This is often the result of marketing campaigns, fintech partnerships, and regional product launches that were never retired. Over time, security teams begin to treat DNS posture as a proxy for broader operational discipline. Additionally, when DNS ownership is unclear or configuration drift accumulates, it often reflects deeper breakdowns in governance, asset lifecycle management, and change control processes.

For financial services, the most common issues in 2025 were misconfigured SOA, missing CAA records, missing DNSKEY records, unnecessary wildcard DNS records, and disabled registry locks (Figure 11). These issues collectively represent critical vulnerabilities in DNS infrastructure security and domain governance programs. These vulnerabilities weaken the cryptographic trust chain and expose financial institutions to impersonation, hijacking, and unauthorized data interception.

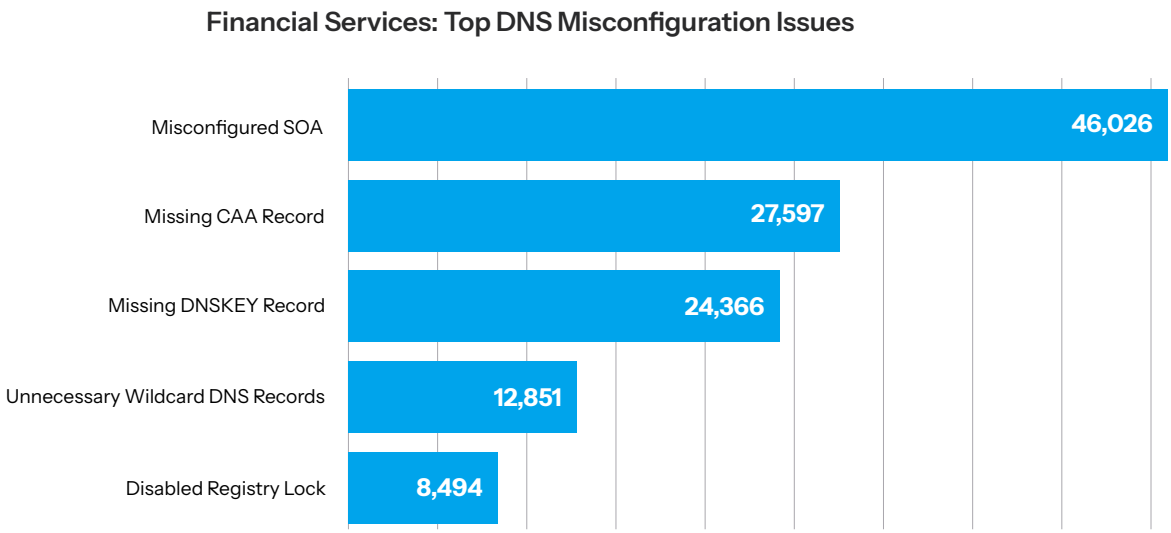

Fig. 11: The top five misconfiguration issues from DNS management analytics of financial sector hygiene practices

Certificate and cryptographic posture

Certificate and cryptographic posture has also emerged as a key DNS-adjacent risk. Nearly 40% of observed certificates expired or were nearing expiration (within 90 days) in 2025, while trust relationships were increasingly concentrated among a small number of certificate authorities. At the same time, fewer than 4% of observed enterprise domains showed indicators of post-quantum cryptographic readiness, highlighting the growing gap between current deployment practices and future resilience requirements. As certificate lifetimes continue to shrink, lifecycle management has shifted from a periodic task to a continuous availability and security requirement.

The impact of DNS on financial core services

What ultimately elevates DNS posture from a technical issue to a board-level operational risk is its impact on financial institutions' core services. DNS failures rarely appear as "DNS incidents." Instead, they manifest as online banking outages, failed payment transactions, disrupted API services, broken email communications, or brand impersonation campaigns targeting customers. By the time these symptoms are visible, the operational and reputational costs are, in many cases, being invisibly absorbed by the business.

For financial institutions, DNS posture is no longer a background hygiene issue. It is a control point for digital trust, fraud prevention, and service resilience.

Web attacks persist in financial services

Financial services' digital transformation relies heavily on customer-facing websites, back-end web applications, and APIs for open banking integrations, account payments, and more. Despite being one of the most fortified industries, attackers relentlessly target financial services as demonstrated by the 11% surge of web attacks from 2024 through 2025 on a global basis (Figure 12). Organizations in the financial services industry hold some of the most confidential and sensitive data, such as personally identifiable information (PII), account credentials, and payment card details, which makes them lucrative targets for threat actors.

Financial Services: Monthly Web Attacks

January 1, 2024 – December 31, 2025

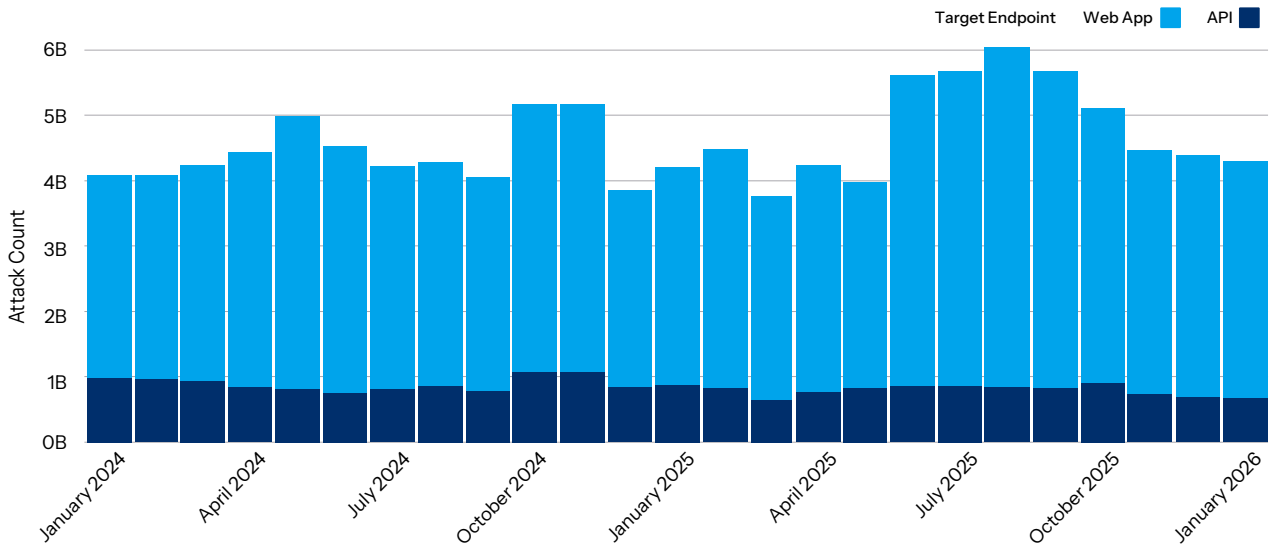

Fig. 12: Between 2024 and 2025, web attacks against the financial services industry grew 11% on a global basis

Additionally, most established financial services companies have been driven by the industry's shift to digital-first strategies to run hybrid environments that combine complex, legacy systems with modern applications, API endpoints, and AI systems. This creates security gaps that attackers can leverage to breach the company's security perimeter and establish a foothold.

From January 1, 2024, through December 31, 2025, the financial services industry remained the second most targeted industry, after commerce, for web attacks (Figure 13).

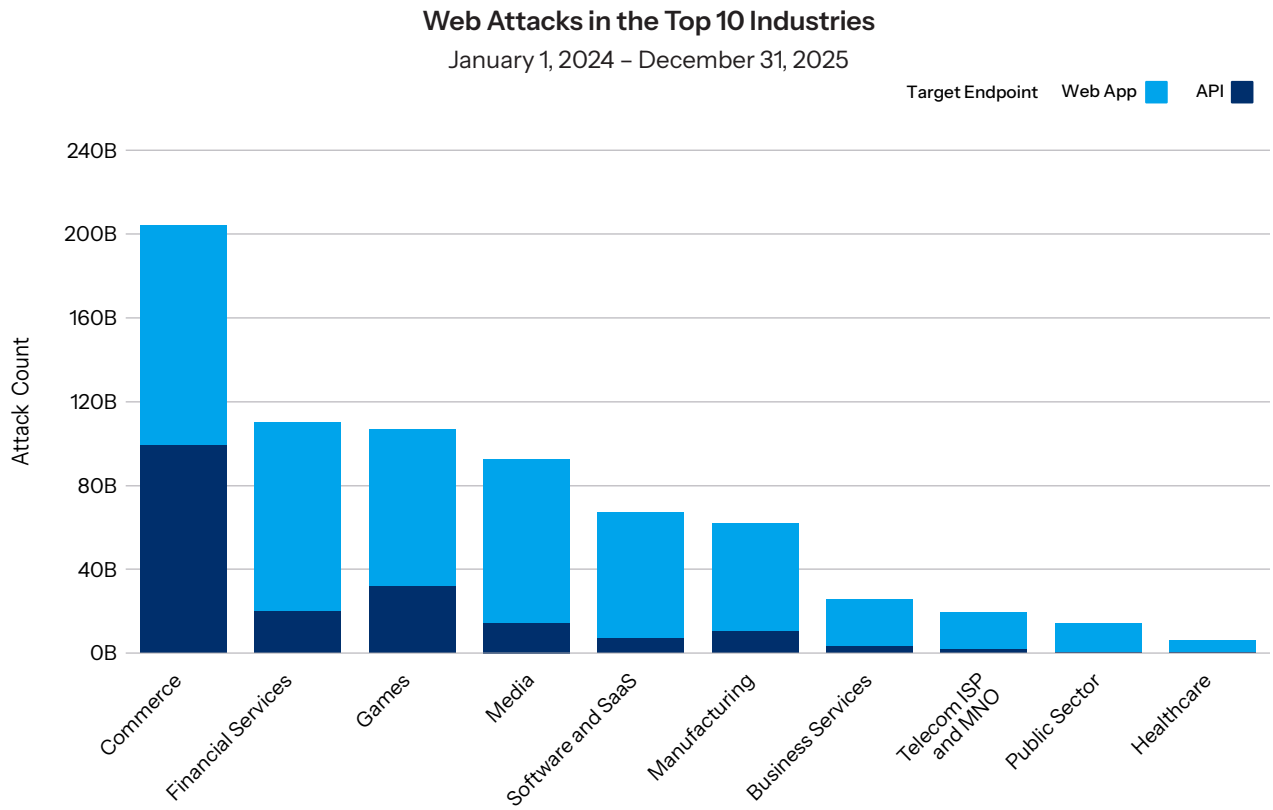

Fig. 13: The financial services industry ranks second, after commerce, with 110 billion web attacks over the two-year period

During that two-year period, financial services experienced 20 billion web attacks on API endpoints, fewer than those in the commerce and games industries. Ecommerce, particularly retailers, connects to thousands of merchants, shipping providers, social media platforms, and local payment systems, which creates a massive API sprawl that expands their attack surface. And in the less-regulated games industry, APIs can expand with fewer checks.

Financial services, by contrast, operates in a highly regulated environment and has a narrower yet deeper API footprint that is focused on performing complex business logic, from KYC/AML verification providers to interbank clearing houses, fintech service providers, and credit worthiness validators. These factors force attacks to be more targeted and also help to reduce the attack surface.

APIs make banking the primary target

As with DDoS attack targets, banking was by far the most targeted financial services vertical by web attacks (Figure 14). For attackers that are looking to make a profit and/or create disruption, attacks against banks provide threat actors with the opportunity for a higher ROI.

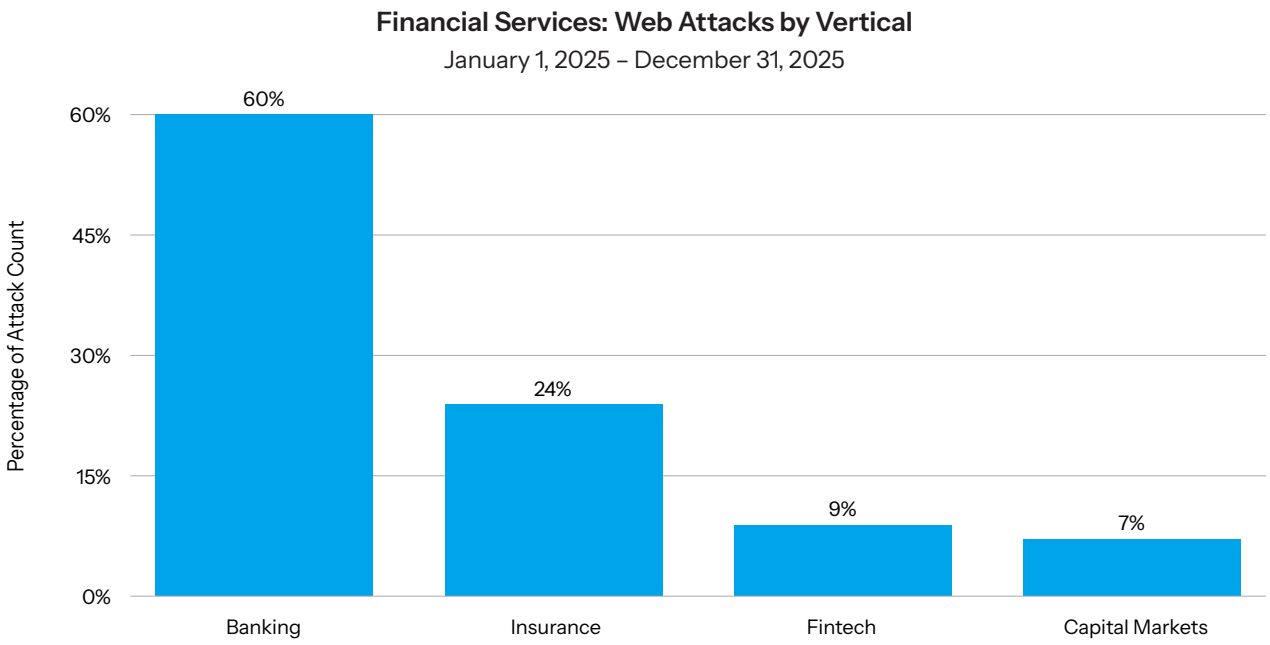

Fig. 14: During 2025, banking was the most targeted financial services vertical by web attacks

Additionally, across all verticals in all industries, banking ranked among the top five most targeted by web attacks during 2025 (Figure 15).

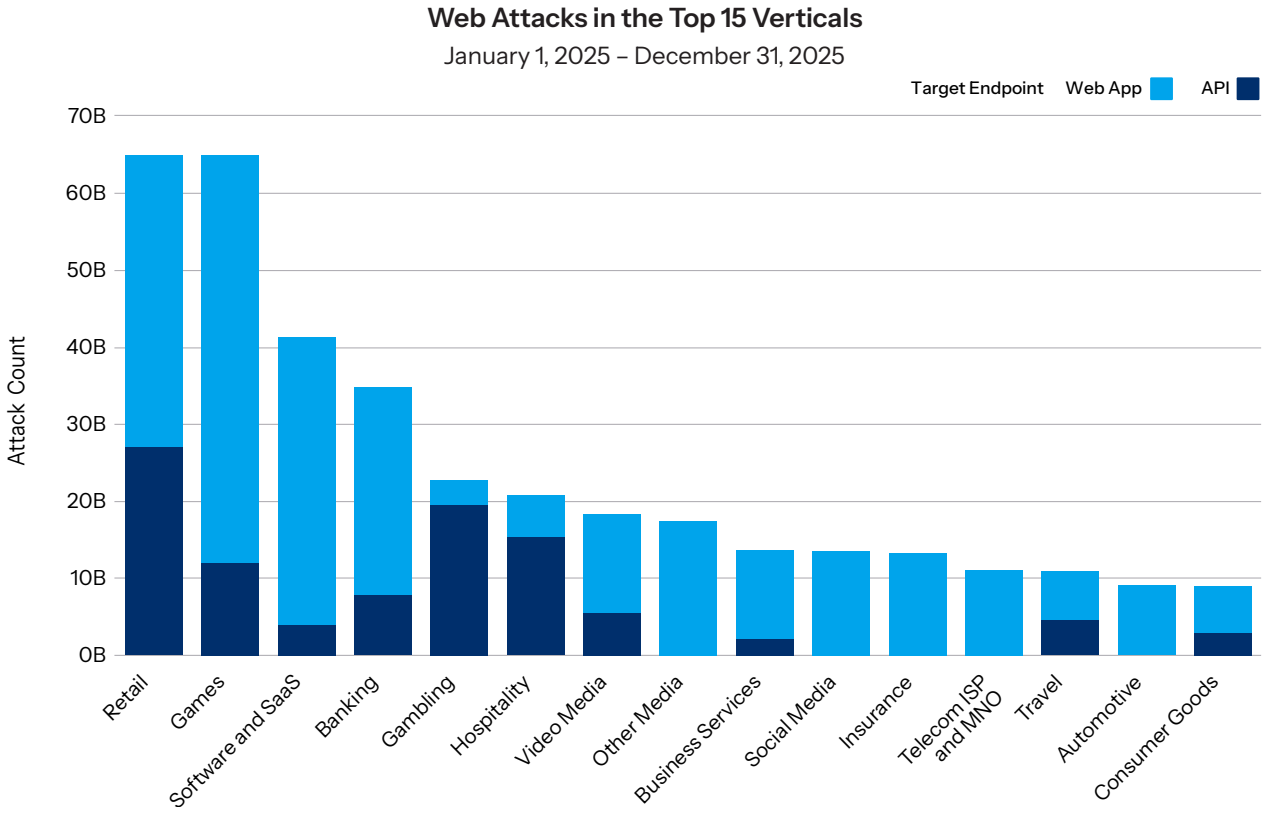

Fig. 15: Banking ranked fourth among all verticals in web attack targets

Within financial services, the banking vertical also led in web attack volume against APIs (Figure 16). APIs play a critical role in open banking by enabling banks to expose standardized APIs that allow third-party providers, such as fintechs, aggregators, and payment initiators, to access financial data and initiate payments with consumer consent. Additionally, digital payment systems and apps like Singapore's PayNow, India's UPI, and Japan's PayPay, further fuel the proliferation of APIs.

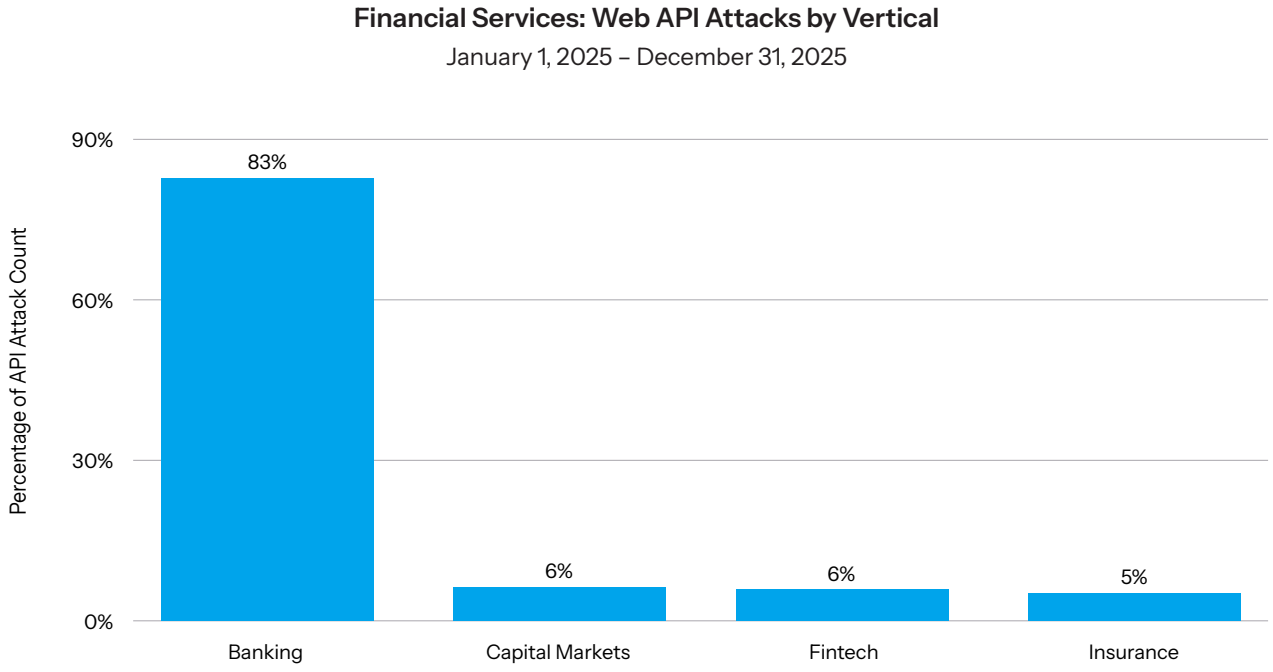

Fig. 16: Banking experienced 83% of web attacks targeting financial services API endpoints

Securing these APIs is essential, as vulnerabilities in authentication, authorization, input validation, or rate limiting enable data leaks, fraud, and account takeover. Regulatory frameworks like Payment Services Directive 2 (PSD2) and the Federal Financial Institutions Examination Council (FFIEC) mandate protection of all financial data, whether in transit or at rest. Attack trends show that, despite these guardrails, financial services remains a prime target.

Regional trends: Web attacks

During 2025, North America was the most targeted region for web attacks against the financial services industry, experiencing 44% of total attack events, followed by APAC (30%), EMEA (15%), and LATAM (11%). It's worth noting that since the publication of our 2023 financial services SOTI, the insurance vertical has jumped positions to move behind banking in North America and APAC and, in EMEA, the two verticals are nearly deadlocked for the top position (Figure 17).

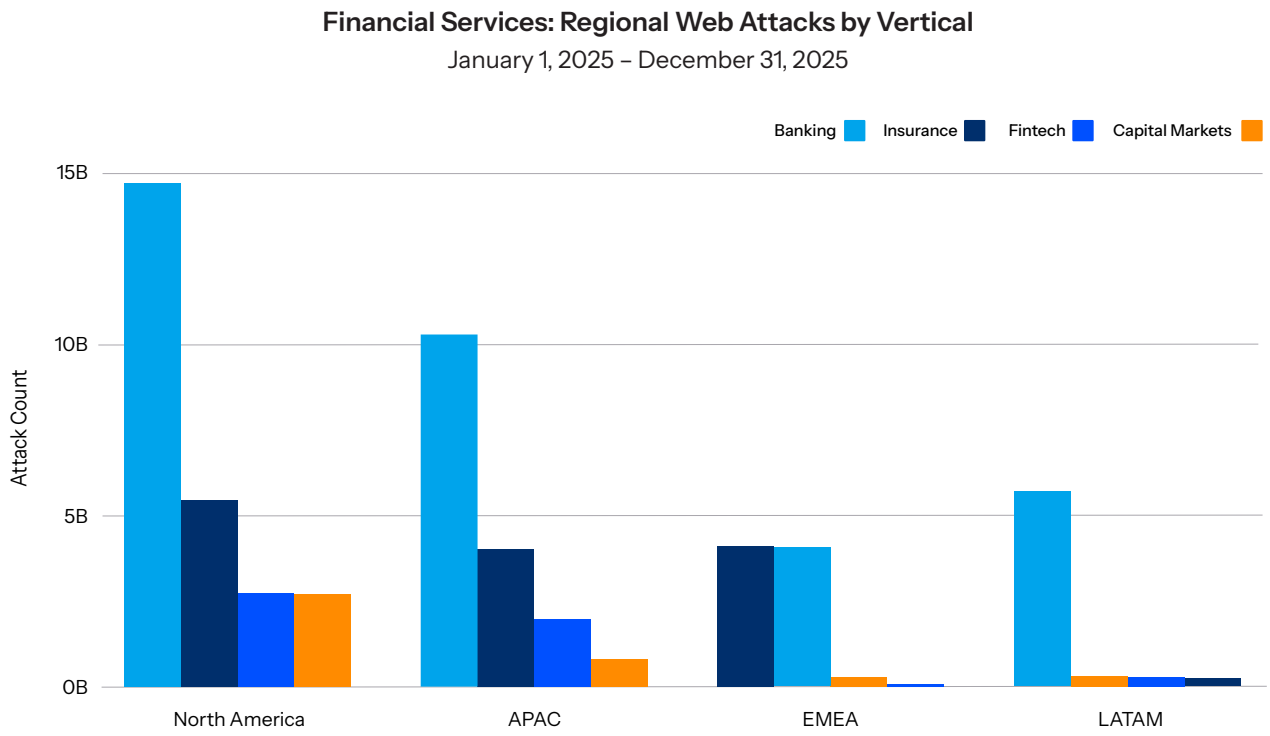

Fig. 17: Web attacks against the insurance vertical surpassed banking in EMEA and rose to second in North America and APAC

Insurers hold massive amounts of PII and, in some cases, protected health information (PHI), making them prime targets for data breaches and identity theft.

During 2025, threat actor groups such as Scattered Spider, Lapsus$, ShinyHunters, and Qilin focused on executing large-scale attacks against insurance companies by leveraging tactics that are part of a broader trend of exploiting supply chain vulnerabilities and third-party relationships. A complex ecosystem of external partners and small suppliers leads to security gaps and vulnerabilities that threat actors can exploit.

The shift to behavior-based threats

While tried-and-true methods, like local file inclusion, Structured Query Language (SQL) injection, and cross-site scripting, are still having an impact on financial institutions, our recently published Apps, APIs, and DDoS 2026 SOTI Security report highlighted a significant shift to behavior-based tactics that target vulnerabilities like broken object level authorization (BOLA) and broken object property level authorization (BOPLA) that remain difficult to detect with traditional, signature-based tools. The report also flagged the continued risk of sensitive data leaks coming from unmanaged APIs (e.g., zombie and shadow APIs) and those connected with AI integrations, as well as from common issues like misconfiguration and broken access control. With AI relying heavily on APIs for data exchange, the volume of data traveling through these connective tissues has increased significantly.

API security challenges in financial services

The 2026 API Security Impact Study highlights these trends, with 96% of global respondents in the financial services industry reporting at least one API security incident in the past 12 months — the most among all industries. According to the soon-to-be-published API Security Risks in Financial Services survey, data leaks and attacks against APIs related to AI technologies topped the list of the most common types of API security incidents in financial services at 39% each, followed by exploitation of insufficient access controls at 35%.

Another crucial problem for most organizations, including financial services institutions globally, is the lack of visibility into which APIs expose sensitive data. This raises concerns about how visibility gaps can hide the risks of misconfigured or overlooked API endpoints that can provide attackers with access to data. Moreover, this introduces compliance risks as AI systems pull troves of data via APIs.

Why dedicated API solutions matter

Respondents attributed their API security incidents primarily to misconfigurations (65%), lack of dedicated API security tools (48%), and vulnerabilities in AI-related APIs (39%). This result underscores how AI amplifies existing risks, and why securing APIs must be the first line of defense for AI protection. Furthermore, it highlights why organizations need web application firewalls (WAFs) for in-line protection in addition to visibility and detection of both conventional web attacks and behavioral risks.

Akamai Security

Threats in depth

Botnets and other threats to AI and APIs

The financial services sector is facing a strategic shift in the threat landscape, characterized by the integration of AI into attack cycles and the industrialization of botnet infrastructure. Akamai's research across its global telemetry security platform indicates that defenders must transition from static perimeters to adaptive, AI-aware security architectures to counter increased attack volume and sophistication.

DDoS: The industrialization of disruption

In 2025, Akamai recorded more than 6,500 significant Layer 7 DDoS attacks. Some of the attack campaigns in financial services have reached billions of malicious requests. Layers 3 and 4 attacks now constitute 37% of the total attacks targeting financial institutions, with median durations increasing by 900% year over year. We observed direct correlations between regional conflicts and attack spikes, notably following U.S./Israeli strikes in February 2026, which triggered increased DDoS activities that targeted Middle Eastern and Western infrastructure.

One underscoring observation in the threat landscape is that attackers are elevating DDoSaaS and massive IoT/residential botnets (e.g., Kimwolf). Beyond boosting volume, actors are exploiting protocol edge cases like HTTP/2 vulnerabilities (e.g., RapidReset) and persistent, sequential multi-vector, multi-surface, multi-day campaigns that pivot among web, DNS, and infrastructure floods.

Bot operations: AI-empowered evasion

Advanced bot activity surged by 147% in late 2025, and in one case, 96% of site traffic was identified as scraping bots. Threat actors continue to use classic low-and-slow tactics; for example, dropping request rates to less than one request per second to evade threshold-based detection. In addition to these classic tactics, attackers are now leveraging AI and advanced headless browsers, making defending against threats even more difficult.

Threat actors are more easily able to seamlessly blend human and attack traffic by using AI-driven botnets to mimic legitimate browser behaviors. When basic volumetric bots fail, these AI-empowered operators dynamically pivot to transparent browser impersonation.

Fighting evasive bots isn't always straightforward. Mitigation should move from basic signature blocking to behavioral heuristics and user-risk telemetry to identify fraudulent identities within transactional flows.

IoT botnet disruption

In March 2026, Akamai assisted authorities in halting DDoS attacks from the world's largest IoT botnets and shut down their related DDoS-for-hire services. The collaboration among law enforcement and industry partners dismantled infrastructure capable of record-breaking hyper-volumetric DDoS attacks. These botnets, which include Aisuru and Kimwolf, controlled more than 3 million compromised IoT devices.

The takedown of these botnets represent a major security victory for the financial services industry by effectively neutralizing a major DDoS threat aimed at crippling banking infrastructure. Furthermore, this operation reduces the risk of extortion and service disruption while strengthening the defenses against fraud and account takeover attempts that target financial institutions.

The rise of autonomous agency

AI-related traffic continues to grow dramatically, with an average of 2.5 billion daily AI bot requests across our network. Within this surge, AI-related traffic experienced hypergrowth, nearly doubling in the second half of 2025 alone. This signifies a fundamental transition from automation to autonomy; AI agents can now navigate complex applications to perform deep contextual analysis and, in some cases, make decisions.

Although the industry encourages responsible AI use via open door policies, malicious use is rising. According to private threat intelligence data (TLP Amber advisory September 2025), an abused search bot launched a volumetric flood that targeted 570,000 unique URLs in a single window. And another massive-scale AI training crawler targeted transactional API endpoints to harvest proprietary data, hitting close to 7,000 hostnames and 37,000 unique paths in less than 7 days. This aggressive scraping behavior demonstrates that AI bots can now extract valuable proprietary information at scale, potentially exposing sensitive business logic, customer data, and competitive intelligence.

We are seeing a trend of attacks that leverage AI-specific vulnerabilities on customers that deploy their own AI. These attacks include prompt injection, sensitive information disclosure, and other OWASP Top 10 for Large Language Model Applications.

Threats are moving from the front door to the model's brain and data. Therefore, a full AI security stack becomes a critical defense imperative to protect against AI as a threat vector and against threats to our customers' AI/LLM services. Specifically, organizations must deploy a firewall capability to inspect LLM inputs (to block prompt injections/jailbreaks) and to filter outputs (to prevent data leaks and toxic content) to ensure safe AI integration.

Threat actor profiles: High-scale adversaries

The modern attack surface was dominated by adversaries operating at unprecedented scales in 2025. Hyperscale botnets, political hacktivists, and social engineers dominated as the top threats in financial services. Understanding these categories allows security operations centers to build out playbooks that are tailored to the threat methodologies.

- **Hyperscale botnets:** Kimwolf controls millions of devices, using residential proxy forwarding to turn downstream home networks into anonymized DDoS proxies and neutralize standard IP reputation blocks.
- **Political hacktivists:** Groups like NoName057(16) and DDoSia use gamified platforms for continuous, politically motivated strikes.
- **Social engineers:** Groups like Scattered LAPSUS$ Hunters (a supergroup of cybercriminals formed in 2025 by merging ShinyHunters, Scattered Spider, and Lapsus$) target DevOps personnel to harvest OAuth tokens by bypassing multi-factor authentication through session hijacking.

Beyond immediate service unavailability, these high-scale adversaries trigger cascading operational failures. Volumetric spikes may serve as a smoke screen to distract security operations from low-and-slow account takeover attempts or PII exfiltration. For the financial services sector, the impact includes SLA breaches with downstream partners, heightened regulatory scrutiny, and lasting brand damage when downtime is viewed as a reputational failure, not just a tech failure.

To continue to level up protection against evolving challenges, institutions must adopt a resilient, layered, and deep defense posture. This includes comprehensive protocol coverage, adaptive rate limiting, and integrated visibility to mitigate automated, high-velocity, multi-vector botnet and AI-empowered attacks.

Akamai Security

Security spotlight

Using MITRE capabilities to improve security posture

In our State of Apps and API Security 2025 report, we covered how to use the Open Worldwide Application Security Project (OWASP) to enhance a cybersecurity program. Now, let's explore how to use the threat analysis capabilities of the ATT&CK/ATLAS frameworks from MITRE.

MITRE was established to advance national security in new ways and serve the public interest as an independent adviser. It is a Federally Funded Research and Development Center (FFRDC) that provides technical expertise, systems engineering, and research services to critical infrastructure industries like aviation, defense, healthcare, emerging technology, public sector policy, and cybersecurity.

The MITRE ATT&CK (Adversarial Tactics, Techniques, and Common Knowledge) framework plays a critical role in helping financial institutions navigate the threat landscape. By mapping real-world adversary tactics and techniques, it enables security teams to understand how attackers operate, test their defenses against constant attacks from various adversaries, and prioritize security investments where risk is highest to strengthen resilience against the threats most likely to impact their environment.

A key cyber capability that we covered in previous SOTI reports is the MITRE ATT&CK Enterprise Matrix (Adversarial Tactics, Techniques, and Common Knowledge). Additionally, there is a newer matrix named ATLAS™ (Adversarial Threat Landscape for Artificial-Intelligence Systems) that focuses on GenAI, specifically LLMs. The MITRE knowledge bases list the tactics or methodologies that cybercriminals can use to break into a network. Under each tactic is a list of techniques and subtechniques that may be used to execute different attacks (Figure 18).

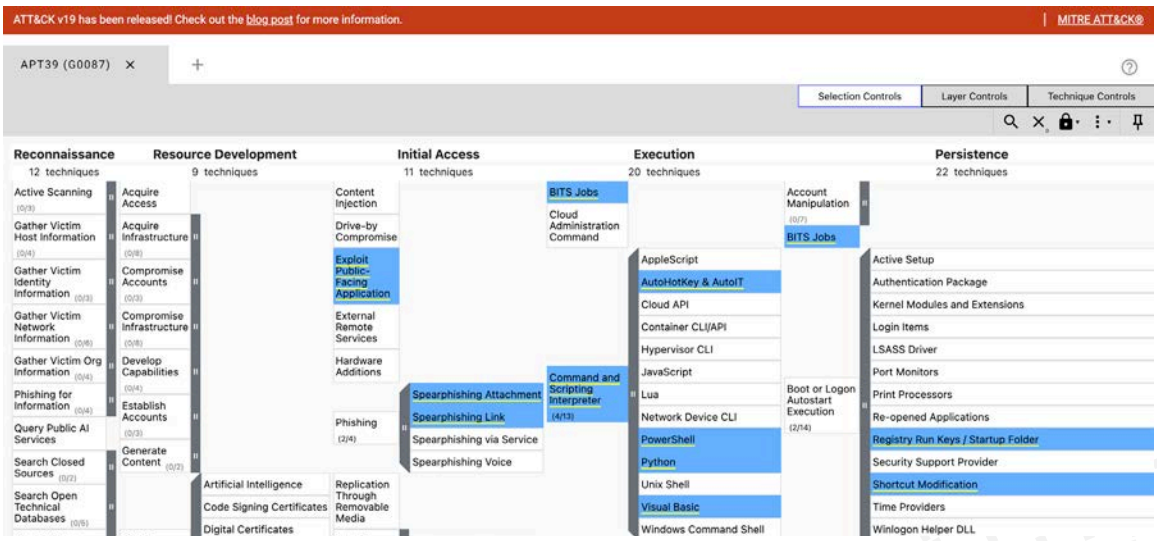

Fig. 18: Sample image of ATT&CK overlaid with criminal group methodology

Defend against real-world tactics

The reason that MITRE matrices or frameworks have become so useful is that they shift cybersecurity discussions from theoretical risks to observable attacker behavior. Instead of focusing only on vulnerabilities or compliance requirements, the ATT&CK/ATLAS frameworks focus on how adversaries actually operate once they begin targeting an organization. Organizations that understand the tactics and techniques used by threat actors can evaluate whether their defenses align with how attacks are conducted in the real world. This allows security teams, including those in the financial services industry, to prioritize having controls that disrupt known attack vectors rather than simply checking compliance boxes.

These frameworks are similar to the cyber kill chain in that they both present multiple opportunities to disrupt an attacker's methodology. The difference is that ATT&CK goes into much more detail to both identify the steps and define the specific actions attackers would take to accomplish the intrusion.

Using MITRE for defense, identification, and prioritization

There are a few ways to take advantage of these matrices. One way that many CISOs find helpful is to use them to evaluate the maturity of their cyber programs. By mapping security tools, detection rules, and response procedures to specific attacker techniques, financial services organizations can quickly identify visibility gaps in their environments. For example, many infosec teams discover that they have strong perimeter defenses but a limited ability to detect credential abuse, lateral movement, or persistence techniques once an attacker gains a foothold.

Using MITRE-based coverage analysis can help teams prioritize where to invest in new detection capabilities, automation, or threat hunting to fill those gaps. The MITRE matrices can also help determine where an organization may have too much coverage, which allows teams to reduce costs by cutting out tools with low ROI.

MITRE also tracks adversary activity clusters that they call Groups and maps their known attack methods to techniques in the matrix. Various naming conventions and migrating hackers can make it difficult to track the different organizations. MITRE provides activity clusters that are tracked by a common name. This allows analysts to track these Groups using various analytic methodologies.

Using MITRE for training and testing

To see the Groups' activities overlaid on the ATT&CK framework, one can use the ATT&CK Navigator tool, which shows the septic techniques used by groups like APT39 or The Lazarus Group. It can also be customized and annotated by threat intelligence teams for use in research and analytics.

Infosec teams can use the information to train security operations center analysts by having them study — and, in some cases, teach classes to other departments about — how poor coding is being used by cybercriminals to build out attack campaigns. Similarly, the pen test and red teams can mirror the attack groups shown in the Navigator tool as templates for their testing plans. This would prove that the security posture could stand up to real-world scenarios. Finally, this information can be used to run tabletop exercises that provide a more realistic experience.

Other MITRE tools that financial services companies can use include:

- **AADAPT ™ (Adversarial Actions in Digital Asset Payment Technologies)** — An outline of adversary tactics and techniques for digital asset management systems
- **CREF (Cyber Resiliency Engineering Framework)** — A relational tool that distills the complex concepts and relationships from NIST SP 800-160 Vol. 2 Rev. 1
- **FiGHT™ (5G Hierarchy of Threats)** — A knowledge base of adversary tactics and techniques for 5G systems
- **Adversary Emulation and Red Teaming** — A common language and framework for red teams
- **Caldera** — Automated security assessments

As cyberthreats increasingly use automation and AI-driven techniques, structured knowledge bases like MITRE's, become even more critical. Attackers are able to iterate and scale their operations faster than ever, which means defenders need equally structured frameworks to understand and anticipate adversary behavior. By using MITRE resources alongside other frameworks, like OWASP, financial services organizations can build a more comprehensive security strategy — one that addresses both application vulnerabilities and the real-world attack techniques used to exploit them.

Akamai Security

Cloud spotlight

Understanding and protecting the different AI architectures

As financial institutions accelerate their adoption of AI, many are discovering that AI is not a single architecture or risk category. In practice, organizations are deploying two distinct classes of capabilities:

1. Generative AI (GenAI) systems, such as LLMs
2. Classic ML-driven AI systems, which are typically nontransformer architectures that execute real-time decisions

Akamai also breaks out **inference AI** as a category. This is the category for which there is a need for a high-performance AI platform that brings the power of ML to the edge. Although inference cloud is not designed exclusively to support smaller ML-focused models, it is the more common deployment because of the need for real-time responsiveness.

Although no single authorized source tracks deployed AI models, some quick open source research shows that there are many more classic ML instances than GenAI or LLM instances today. These ML-based models power fraud detection, recommendation engines, personalization platforms, and autonomous decision systems across banking, payments, and financial services. That said, most reporting on investment is focused on GenAI or LLMs.

Understanding the difference between these architectures is critical because the infrastructure and security models that are required to protect them are fundamentally different. While there is some overlap at the most basic level, the difference between the attack methodologies is that GenAI attacks are logical layer attacks while ML attacks are based on classic malware attacks against the data or supporting code.

LLMs are a specific subset of GenAI designed to create or transform content. They rely on extremely large datasets and complex neural networks to generate human-like text, code, or media. Their security challenges tend to focus on model behavior and data integrity — including issues such as prompt injection, data leakage, and model manipulation.

Classic ML systems operate differently. In a classic ML deployment, a trained model is executed to produce real-time predictions, classifications, or decisions. These models are often embedded within applications and services and are frequently deployed on distributed infrastructure close to the user or transaction source to minimize latency.

For financial institutions, classic ML systems power critical capabilities such as:

- Fraud detection
- Transaction scoring
- Credit risk analysis
- Customer personalization
- Automated trading signals
- Data-driven recommender and cross-sell systems

In these environments, speed and availability are just as important as model accuracy. This architectural requirement often pushes classic ML workloads toward distributed or edge infrastructure, where models can execute closer to the application or user.

Because classic ML models are integrated directly into applications and APIs, their primary risk surface is often operational infrastructure. While LLMs are vulnerable to semantic manipulation (prompt injection), classic ML is vulnerable to feature manipulation (adversarial evasion), which is often executed via automated probing of the APIs.

Finally, all AI models must have the ability to be audited and undergo a cybersecurity investigation. Regulated industries, especially those operating in multiple regions, need to make sure that regulatory and compliance requirements are built into the design as well.

AI does not replace traditional security risks — it amplifies them

While the MITRE ATLAS framework catalogs more than 50 potential attack scenarios against AI systems, only a small subset represent confirmed real-world compromises. Many of these incidents combine elements of traditional ML attacks with more conventional infrastructure exploitation techniques.

In practice, many AI systems are compromised through the same mechanisms that affect modern digital applications:

- DDoS attacks
- Automated bot abuse
- Vulnerable APIs
- Exposed infrastructure services
- Compromised software dependencies

This trend highlights an important reality for security leaders: AI does not replace traditional security risks — it amplifies them. Classic ML systems are particularly exposed because they are frequently deployed as high-availability services connected to external APIs and real-time data feeds. Attackers can exploit these systems through automation, API abuse, or infrastructure exhaustion long before they attempt to manipulate the model itself.

As AI adoption grows, organizations are rapidly increasing the number of deployed models. While the financial services industry narrative often centers on a handful of massive LLM deployments, the operational reality is that enterprises may run hundreds or thousands of smaller classic ML models across distributed systems.

This creates a large and complex attack surface that includes:

- Data ingestion pipelines
- Model hosting environments
- Inference APIs
- Partner integrations
- Edge compute infrastructure

The security challenge is further compounded by modern development practices. AI capabilities are often assembled from multiple external services, cloud platforms, and data providers, making visibility and control more difficult. Security teams must therefore address not only the models themselves but also the surrounding ecosystem of APIs, services, and data pipelines. This requires AI firewalls that are purpose built to integrate with API security capabilities and traditional WAFs.

Best practices for protecting AI deployments

For cyber leaders, protecting AI deployments begins with understanding where risk actually resides in the architecture. Several foundational capabilities are essential, and situational awareness is the starting point. AI systems rely on data collected from numerous sources including APIs, sensors, user interactions, logs, and external feeds. Maintaining visibility across these data pipelines is essential for identifying abnormal behavior and detecting manipulation attempts.

Clarity around the role of AI models is also critical. Some deployments rely on massive labeled datasets and generalized models, while others use curated datasets designed for narrow operational missions such as fraud detection. These differences influence both risk tolerance and required controls.

Optimizing, meeting green initiatives, and managing costs can also affect a security posture. Model tuning and infrastructure decisions can significantly impact latency, resource consumption, and exposure to automated abuse.

Akamai Security

Finally, security and visibility must be integrated into the AI deployment pipeline. This includes:

- Protecting APIs used for model inference
- Defending against bot-driven automation and abuse
- Ensuring resilience against DDoS attacks
- Maintaining secure DNS and service discovery
- Identifying unused or zombie APIs
- Using microsegmentation to minimize impacts
- Monitoring model behavior and system performance in real time
- Having the ability to audit and investigate incidents

For example, there is a need for visibility or situational awareness with regard to model drift, data drift, and training-serving skew. These metrics, however, need to be observed and managed, which can have an adverse effect on the security of systems downstream. Additionally, companies need to integrate more traditional security capabilities like microsegmentation, which minimize the potential back-end impacts. Insurers and regulators are also moving to make those capabilities part of basic expectations for the sector.

At an industry level, the rapid growth of AI has elevated it to the top of many financial services organizations' risk registers. However, the most significant risks often stem not from exotic model manipulation attacks, but from the operational complexity of deploying AI systems at scale. For most enterprises, the challenge is not securing a single model but protecting an entire distributed ecosystem of inference services, APIs, data pipelines, and partner integrations.

When assessing potential solutions to protect an entire distributed ecosystem:

- Evaluate the type of infrastructure down to the CPUs used
- Consider capabilities like performance if there is a need for low-latency applications
- Determine whether the solution will provide
    - An orchestrator that brokers AI requests
    - Access at the edge with semantic caching and serverless services
    - The ability to leverage vendors' security platforms of tools and capabilities

Akamai Security

Cyber leaders who understand the architectural differences between LLM-based GenAI systems and classic ML-driven AI applications will be better positioned to design appropriate defenses and select the right infrastructure. For example, low latency should be built on an inference cloud. In practice, this means combining traditional application and infrastructure protections with emerging controls designed specifically for AI workloads.

IT and cybersecurity leaders need to determine which infrastructure and security capabilities are best for the type of AI models they are using. With 50% of enterprise apps expected to reside outside major public clouds by 2027, now is the time to start understanding both where to build and how to secure the next generation of customer-facing capabilities.

Ultimately, the financial services organizations that succeed in deploying AI safely will be those that treat security and observability as foundational capabilities of the AI platform itself, rather than as controls applied after the models are deployed.

Compliance

For financial institutions, cybersecurity is no longer solely an operational concern, it is a regulatory imperative. The global compliance landscape for the financial services industry's cybersecurity has continued to intensify in 2025 and 2026. Regulators worldwide have tightened requirements around data protection, incident reporting, third-party risk management, and now AI governance.

Security leaders must continue to navigate a complex and rapidly evolving patchwork of laws, regulations, and frameworks that directly affect how institutions protect their networks, deploy AI, and manage digital risk. This section provides a high-level overview of key regulatory developments that financial services organizations should have on their radar and suggests a few ways to navigate this regulatory landscape.

Global cybersecurity developments and an acceleration of enforcement efforts

Globally, the years 2025 and 2026 have seen significant enforcement across multiple regulatory domains that impact financial services. The New York Department of Financial Services (NYDFS) Cybersecurity Regulation saw its first annual certification deadline on April 15, 2026. The NYDFS has signaled an active enforcement posture, with fines of up to US$250,000 per day for ongoing noncompliance. Additionally, the SEC's 2025 Examination Priorities flagged cybersecurity governance as a top focus area.

Outside the United States, the Digital Operational Resilience Act (DORA) has been fully in effect since January 17, 2025, making 2026 the first full year of enforcement with significant penalties for noncompliance. Alongside DORA, the Network and Information Security Directive 2 (NIS2) expands the scope of cybersecurity obligations across the European Union, and Additionally, the General Data Protection Regulation (GDPR) enforcement continues to intensify, with cumulative fines exceeding 6.7 billion euros as of late 2025.

The global AI governance challenge: A new compliance frontier for financial services

Of course, financial institutions are attuned to the increasing security and privacy regulatory activity. Today's challenge is navigating the emergence of AI governance as a formal compliance domain. While financial institutions rapidly deploy AI systems in important areas (e.g., antifraud, credit scoring, personalization, and threat analysis), their use cases typically fall squarely within regulatory crosshairs.

The EU AI Act has set a standard for AI governance that will seriously impact financial services. The Act classifies AI systems into four risk tiers with escalating compliance obligations. The most significant deadline for financial institutions is August 2, 2026, when, absent regulatory delay, requirements for high-risk AI systems (such as credit scoring, behavioral profiling, and recruitment) become fully enforceable. Providers and deployers of such systems must implement risk management systems/mechanisms and conformity assessments before deployment. As is typical of modern EU regulations, penalties for noncompliance are severe and the law's reach is extraterritorial.

On a positive note, the European Commission has issued its Digital Omnibus regulation proposal, which aims to simplify and harmonize the EU's digital regulatory landscape by amending and consolidating compliance frameworks, including the EU AI Act, GDPR, NIS2, DORA, and the Data Act. The proposal introduces a single incident reporting point and aligns breach notification thresholds and timelines across frameworks, which could significantly simplify the navigation of multiple overlapping regimes.

Globally, AI governance approaches vary but are converging around common themes, with privacy, cybersecurity, and safety all playing central roles. In the United States, where there is no comprehensive federal AI law, the voluntary NIST AI Risk Management Framework has become the de facto standard for AI governance in both the private and public sectors. In addition, states such as California and Colorado have taken steps to impose obligations on developers and deployers of high-risk AI systems in areas including financial services, employment, and healthcare. Federal agencies, like the Securities and Exchange Commission (SEC), the Federal Trade Commission (FTC), and banking regulators have moved aggressively into AI oversight within their respective domains.

In the APAC region, Singapore continues to lead with its Model AI Governance Framework and the Monetary Authority of Singapore's (MAS's) technology risk management guidelines on the responsible use of AI in finance. Japan has established a dedicated AI Safety Institute and enacted the AI Promotion Act, maintaining a voluntary, industry-cooperative approach, while South Korea enacted its AI Basic Act, which established a risk-based classification system with lighter compliance requirements than in the EU. China has adopted the most operationally prescriptive approach, requiring algorithmic impact assessments, AI content labeling, and security assessments before deploying GenAI services.

Action plan: Integrating AI governance into financial services compliance programs

The regulatory trajectory is clear: Cybersecurity and AI governance requirements will continue to expand, enforcement will intensify, and penalties will grow more severe. For financial institutions, the convergence of cybersecurity regulations, data protection laws, and AI-specific governance frameworks, creates a complex compliance landscape.

Financial services organizations should move decisively to build integrated governance programs that address this convergence rather than treating each regulation in isolation. To manage risk efficiently and prepare for the continued expansion and intensification of enforcement across these domains, organizations should use the following checklist:

- ☐ **Establish a principles-based governance program that aligns privacy, security, and AI governance programs** to manage risk and reduce operational drag.
- ☐ **Inventory and classify all AI systems by risk tier**, mapping each system across applicable regulatory frameworks and standards.
- ☐ **Implement robust third-party risk management programs** that account for the expanding regulatory reach into the technology supply chain.
- ☐ **Maintain comprehensive technical documentation and audit trails** for all AI systems, ensuring conformity with regulatory frameworks, such as the EU AI Act, and the requirements for risk management systems, data governance, and conformity assessments ahead of coming deadlines.
- ☐ **Ensure meaningful human oversight of automated decision-making** by embedding governance controls that satisfy AI-specific and broader cybersecurity and privacy obligations.
- ☐ **Build security, privacy, and AI governance into the development lifecycle from the outset**. Adopt a "governance by design" approach.
- ☐ **Monitor and prepare for emerging regulatory developments** to ensure the organization remains ahead of new requirements as AI capabilities continue to evolve.

As AI capabilities evolve from predictive models to autonomous agents, the regulatory frameworks governing them will continue to mature. Financial services institutions that invest proactively in governance infrastructure now will not only meet current compliance obligations but also will be better prepared for whatever requirements emerge next.

Akamai Security

Mitigation strategies

True mitigation in the modern financial ecosystem requires moving beyond basic compliance toward a posture of active, continuous defense. The data reveals a stark reality: Despite widespread API adoption, only 27% of organizations with full API inventories have clear visibility into where their sensitive data lives. To survive today's threat landscape, leaders must bridge this gap while using proven strategies like microsegmentation, which can reduce response times by 33%. In a world where DNS, APIs, and AI models are under constant, automated surveillance, resilience is found in the ability to adapt as quickly as the adversary.

For most organizations, protection against DDoS is built around managing a service provider. The three keys to making sure this functionality is optimized are:

1. Validating that the capabilities are sufficient based on current peak attacks
2. Ensuring the processes for notifications and reporting are working by conducting exercises across Layers 3 and 4, Layer 7, and DNS infrastructure
3. Stress testing the systems to validate that all systems are covered and the capabilities are working

Fraud mitigation has become a major priority for banks as financial fraud grows faster and gains the ability to be more targeted. Cybercriminals are increasingly bypassing traditional security measures by manipulating users into sharing verification codes or approving fraudulent actions to gain access to accounts and financial systems.

APIs are an integrated part of AI so the most effective security controls address them together. API infrastructure should have the ability to conduct discovery, label sensitive data, detect logic/behavioral attacks, alert on abuse, and set responses based on a company's leadership's risk tolerances. Customer-facing AI deployments need to have security designed into them before deployment, have firewalls designed for logic/behavior–based attacks and use microsegmentation as a safeguard.

While adversaries use GenAI to map infrastructure and probe for weaknesses, defenders can leverage the same technology to automate the discovery of complex vulnerabilities. AI-driven testing can identify logic flaws in proprietary banking applications that traditional signature-based scanners often miss.

Finally, the hidden dangers around DNS need careful monitoring and management. Based on lessons learned from the Akamai DNS protection team, we developed the following checklist to provide security, infrastructure, fraud prevention, and compliance teams with a shared strategy to combat DNS risks.

Hygiene checks

- ☐ Use distributed, cloud-hosted DNS services designed to withstand major DDoS attacks
- ☐ Segregate internal and external DNS environments
- ☐ Routinely remove expired or unused records
- ☐ Keep DNS infrastructure independent from application and website hosting
- ☐ Use role-based access control, multi-factor authentication, audit logging, change management, and DNSSEC
- ☐ Manage time to live values deliberately to balance caching

Posture management and compliance checks

- ☐ Automate discovery of DNS assets across all environments
- ☐ Analyze configurations for drift, misconfigurations, and ownership gaps
- ☐ Manage certificates for expiration, vulnerabilities, and lifecycle compliance

Emerging threats checks

- ☐ Monitor for typosquatting, look-alike domains, and brand impersonation
- ☐ Prepare for post-quantum cryptographic transition

The findings in this SOTI Security report provide the intelligence necessary to review and reinforce modern cybersecurity programs. With the speed and scope of attacks setting new records annually, the transition to adaptive, layered protections is now a baseline requirement for financial stability.

By closing visibility gaps and prioritizing the strategic mitigations discussed in this report, leaders can safeguard critical services and ensure long-term resilience. Ultimately, staying proactive remains the most effective defense against the increasingly sophisticated tactics that are targeting the digital financial ecosystem.

Conclusion

The year 2025 marked a significant shift as cybercrime evolved from simple automation to more autonomous, AI-powered systems. Although the financial sector continues to face record-breaking DDoS volume and a surge in API exploits, this era also offers a unique opportunity to modernize the foundations of digital trust. Beyond the immediate technical challenges, maintaining high uptime and secure connectivity is what keeps customer loyalty strong and keeps regulatory conversations positive.

The takeaway is that while AI amplifies traditional risks like DNS flooding and DDoS, it also provides the tools to build more intelligent, adaptive defenses. To stay ahead, the financial services industry is already moving past baseline requirements toward security architectures that are as agile as the modern market. By prioritizing proactive protection and rapid restoration, leaders can do more than just mitigate risk — they can safeguard their institution's market value and ensure that the next wave of financial innovation is built on a secure, resilient, and trusted foundation.

Methodology

Web application and Layer 7 DDoS attacks

This data describes application-layer alerts on traffic seen through our web application firewall (WAF). The web application attack alerts are triggered when Akamai detects a malicious payload within a request to a protected website, application, or API.

The Layer 7 DDoS alerts are triggered when we detect volumetric anomalies in the number of requests to a protected website, application, or API. These alerts can be triggered by both malicious and benign requests. Typically, the requests themselves are benign, but the high volume of requests indicates malicious intent. The alerts do not indicate the successfulness of an attack. Although these products allow for a high level of customization, the data analyzed here does not consider custom configurations of the protected properties.

The data was drawn from an internal tool for analysis of security events detected on Akamai Cloud, a network of approximately 340,000 servers in more than 4,000 locations on nearly 1,300 networks in 130+ countries. Our security teams use this data, measured in petabytes per month, to research attacks, flag malicious behavior, and feed additional intelligence into Akamai's solutions.

Layers 3 and 4 DDoS attack events

Akamai Prolexic Routed protects organizations from DDoS attacks by blocking malicious traffic before it reaches applications, data centers, and cloud and hybrid internet-facing infrastructure (public or private), across all ports and protocols. Experts in the Akamai Security Operations Command Center (SOCC) deploy proactive mitigation controls to detect and stop attacks instantly, and conduct live analysis of the remaining traffic for further action. These mitigated attacks are organized into attack events, and recorded for analysis.

Guest contributors

James A. Casey
Vice President, Chief Privacy Officer

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 Indigenous peoples in the United States and around the world.

Chris Finch
Senior Technical Solutions Architect

Chris Finch is a Senior Technical Solutions Architect in North America, focused on AI/ML, high performance computing, low-latency workloads, and games industry use cases. Chris has deep experience spanning the telecom, financial services, marketing, life sciences, and game design industries. He is responsible for helping some of Akamai's largest and most creative customers build and grow on Akamai's global platform, while driving thought leadership and strategic thinking in Akamai's AI Center of Excellence as a leader of the product vision working group. Chris's current focus centers on helping customers understand how to use the Akamai Inference Cloud securely for their own workloads.

John "JD" Denning
Chief Security Officer, FS-ISAC

John "JD" Denning is the Chief Security Officer at the Financial Services Information Sharing and Analysis Center (FS-ISAC), owning the internal cybersecurity and risk management functions. In this role, he also works across the sector to curate and disseminate critical baseline cybersecurity practices, bringing learnings from the most mature cyber defense programs to the entire sector.

Prior to his work at FS-ISAC, JD spent 13 years at Bank of America; 11 years within Global Information Security. His most recent role within the bank was Global Compliance and Operational Risk Executive, where he was responsible for the second line of defense for Global Markets Technology, Global Markets Operational Technology, and Global Banking Technology, focused on risk identification and reduction. He also served as Senior Vice President of Cyber Crime Prevention, Identity, and Access Management and led the Cybersecurity Threat Intelligence team.

Prior to his time at Bank of America, JD was the Director of External Affairs for the US Department of Homeland Security's Office of Cybersecurity and Communication and spent 11 years as a congressional staff member focused on cybersecurity, telecommunications, and critical infrastructure protection.

Akamai Security

Ryan Gao
Strategic Engagement Manager

Ryan Gao serves as a Strategic Engagement Manager at Akamai, where he uses his technical expertise in a strategic capacity within Global Services. With more than 10 years of experience, Ryan has collaborated with financial institutions in the banking, payment card services, brokerage, wealth management, and fintech sectors. He has played a pivotal role in ensuring customer success in service and support, focusing on web performance, cloud computing, and security. Ryan is deeply invested in research and advocacy related to threat intelligence, financial industry metrics, and security operation trends. His leadership in these areas has positioned him as thought leader, driving forward the understanding and implementation of advanced security measures in the financial industry.

Jay Jenkins
Chief Technology Officer, Cloud Computing

As the Chief Technology Officer for Akamai's cloud computing services, Jay Jenkins helps organizations solve their biggest problems by embracing new technologies. His teams create or recreate applications using modern technologies to co-create the future for customers and partners. His teams have deep skills in building applications to take advantage of distributed cloud computing.

Jay has more than 20 years of experience in cloud transformation across a wide range of industries around the globe. Prior to working at Akamai, Jay was a tech strategist and evangelist at ByteDance and Google. Jay has also worked at global consulting firms to transform the finance and government industries.

Brent Maynard
Senior Director for Cybersecurity Strategy

Brent Maynard is the Senior Director for Cybersecurity Strategy at Akamai. With more than 17 years of experience in driving innovation in cybersecurity, Brent has led teams across the financial services sector and major cloud service providers, developing groundbreaking security solutions and advancing the industry's approach to threat detection and response.

Brent's contributions include holding a patent for automated security investigations and shaping transformative products that enhance the security operations center experience. As a trusted advisor to the intelligence community and federal law enforcement, Brent has guided high-profile cyber investigations and collaborated on solutions to complex security challenges.

Steve Winterfeld
Advisory Chief Information Security Officer

Steve Winterfeld is Akamai's Advisory CISO. He has a strong background in building operational security programs that are compliant with industry regulations. Before joining the team, Steve served as CISO for Nordstrom Bank, Managing Director of Incident Response and Threat Intelligence at Charles Schwab, and Senior Technical Director Cybersecurity & Group CTO at Northrop Grumman. Before working in the commercial sector, he was an Airborne Ranger in the United States Army and built out the first regional emergency response center (today called security operations centers) for Southern Command.

Steve focuses on collaborating with Akamai's customers to enable them to be successful in defending themselves and their customers. He also helps determine where Akamai should be focusing its security platform's capabilities. Steve has published a book on cyber warfare and holds CISSP, ITIL, and PMP certifications.

Credits

Research director

Kimberly Gomez

Writing and editing

Charlotte Pelliccia
Lance Rhodes
Badette Tribbey
Maria Vlasak

Review and subject matter contribution

James A. Casey
John "JD" Denning
Chris Finch
Ryan Gao
Jay Jenkins
Reuben Koh
Brent Maynard
Richard Meeus
Stas Neyman
Steve Winterfeld

Data analysis

Chelsea Tuttle

Promotional materials

Ashley Linares
Ellen O'Brien

Marketing and publishing

Georgina Morales Hampe
Kimberly Gomez

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 Security protects the applications that drive your business at every point of interaction, without compromising performance or customer experience. By leveraging the scale of our global platform and its visibility to threats, we partner with you to prevent, detect, and mitigate threats, so you can build brand trust and deliver on your vision. 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 05/26.