SOTI

V11 ISSUE 04

Fraud and Abuse Report 2025

Charting a Course Through AI's Murky Waters

Akamai

State of the Internet/Security

Contents

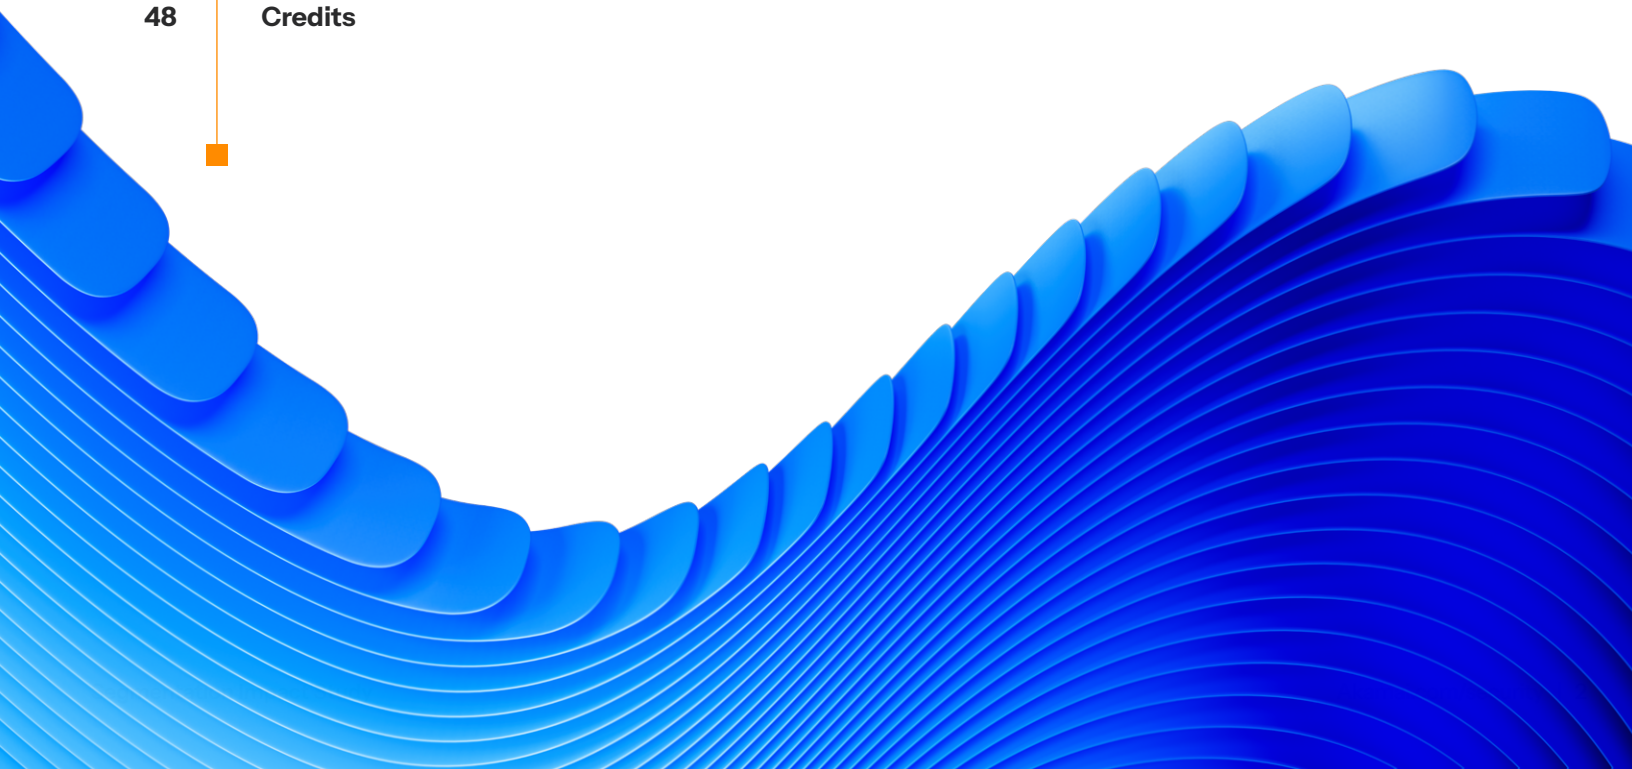

Introduction

The rapid evolution of threat methodologies and a significant uptick in malicious activity prompted the Akamai State of the Internet (SOTI) team to update its analysis of the abuse and fraud landscape. In 2025, the level of innovation and complexity in this landscape — driven primarily by the proliferation of AI bots and APIs — is unprecedented.

For clarity, this report defines *fraud* as the use of deception for unlawful gain, and *abuse* as the engagement in unauthorized activities. Although these definitions may appear straightforward, the increasing complexity of AI-powered technologies has blurred the lines between legitimate bot activity and these types of malicious actions.

This report focuses on a subset of large language model (LLM) AI bots that we codified as "Akamai categorized" by using authentication methods that rely on self-identification and bot telemetry. Our research aims to provide organizations with actionable insights into the evolving role of bots and scrapers, to identify the current trajectory of AI-driven threats, and to highlight the threats' immediate impacts across different regions and industries.

Fraud and abuse risks are amplified by AI

The growing adoption of AI has introduced new opportunities for cybercriminals, which include highly convincing impersonation and automated scams. Although humans direct the use of AI, bots are primarily responsible for generating impersonation content through generative AI techniques. These tools have enabled the automation of phishing, social engineering, and identity fraud campaigns, increasing both the scale and speed of attacks while minimizing the human errors that typically trigger detection.

In previous SOTI reports, we highlighted the impact of AI-driven web scraping on the commerce sector, noting that AI-powered scrapers have become more sophisticated and challenging to detect. This year, our analysis extends to the broader LLM AI bot landscape. Enhanced detection capabilities now allow us and our clients to categorize and analyze a wider range of bot activities.

Akamai researchers have observed that several industries beyond commerce are now experiencing significant increases in AI scraper activity. From an abuse standpoint, some AI bots are capable of intensifying risks in applications and APIs by automating advanced attacks that exploit session flows, authentication mechanisms, and API business logic. The effects of these activities may vary significantly by industry.

Additionally, the rise of fraud as a service (FaaS) in underground markets has further accelerated these threats by lowering the barriers to entry for cybercriminals and facilitating the execution of diverse online fraud schemes. Organizations cannot simply block all bot activity; instead, they must assess bot intent and identity to understand the potential impacts on their business. This approach necessitates specialized bot management solutions. The stakes are higher than ever, given today's increasing compliance requirements, regulatory pressures, and other external factors.

This report provides an in-depth examination of the expanding fraud and abuse landscape, including a guest column from the Chief Information Security Officer at the Financial Services Information Sharing and Analysis Center (FS-ISAC), key highlights from the Open Web Application Security Project (OWASP), industry and regional trends, and a security spotlight on the evolving real-time detection techniques that allow us to keep pace with the latest developments in bot activity.

Guest column

Working together to achieve collective defense

Leaders in every sector of the economy are concerned about fraud conducted via cyber methods. The threshold to perpetrate fraud is lower than ever before, and threat actors are becoming increasingly sophisticated and effective.

As you'll read in this report, much of that fraud is conducted by bots in order to commit a wide variety of crimes. In financial services, we've seen bots automate login attempts with stolen username/password combinations, test ill-gotten credit card numbers or common passwords, and build and manage fake identities. Other bots prevent consumers from making or completing purchases, disrupt legitimate business network traffic, and degrade a company's ability to operate critical business processes.

These frauds are enabled by adaptive, distributed bots powered by AI that are often rented from FaaS platforms. These technologies are built to bypass your defensive controls by impersonating legitimate traffic and allowing malicious activity to evade detection.

In this threat landscape, we must prepare by comparing threat actor capabilities and behaviors with our deployed defenses. We need to practice incident response, share information with industry and sector peers, and continuously evaluate and test the efficacy of our control configurations and designs.

Attacks come in many different forms, so it is important that information security teams consistently adhere to three basic cyber hygiene practices. That will help them scale capabilities in the face of an incident. These practices include:

1. **Layered defenses**: Think of bot detection like spam filtering and use a defense-in-depth approach with constant tuning. IP blocking isn't enough, so incorporate behavioral analytics, rate limiting, and device fingerprinting. Further, enforce per-IP or per-session limits on login attempts, checkout requests, or API calls.

2. **Response playbooks**: Any event can escalate to a severe incident in minutes, so it's crucial to plan, test, and exercise. Incorporate:

   - A coordinated takedown process involving hosting providers, domain registrars, or law enforcement to disrupt bot infrastructure
   - A documented and practiced process to preserve logs and evidence that civil and criminal courts can accept — fraud is a crime, and a compromised system is a crime scene
   - Mechanisms to communicate with customers if their accounts are compromised or your prices or product availability are affected
   - Methods for containing and/or disabling accounts, invalidating sessions, or adjusting pricing/inventory until the fraud is mitigated

3. **All-source threat intelligence**: Consider using botnet identification feeds to detect botnets, proxy networks, or fraud tools that are targeting your sector or company, and employ indicator of compromise (IOC) ingestion tools to spot IP ranges, autonomous system number (ASN) patterns, or signature hashes pushed into your web application firewall (WAF), content delivery network (CDN), or security information and event management (SIEM) for blocking.

The most valuable threat intelligence comes from trusted peers and colleagues in your sector and others. My organization, Financial Services Information Sharing and Analysis Center (FS-ISAC), is a trusted partner that helps the financial sector share threat intelligence, understand security measures against current and emerging threats, and plan and conduct resilience exercises. Most industries have ISACs specific to their sector — the National Council of ISACs will help you find the trusted network in your industry.

A great example of intelligence sharing is provided by FS-ISAC's Cyber Fraud Prevention Framework Working Group — a cross-sector group of cyber experts who develop strategies to reduce the risk of fraud. Their expertise is shared to the benefit of the entire financial sector. If your industry is not working on this problem collectively, reach out to your ISAC and ask how to get involved. I am confident that you will be welcomed, and your time will be well spent.

Only by working together can we achieve the promise of collective defense. Fraud actors cross borders and sectors, and each successful fraud chips away at profitability and undermines consumer trust. Expand your perspective, get involved, and step up to defend against this increasingly complex and dangerous fraud landscape.

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

Akamai is a founding member of FS-ISAC Critical Provider Program and maintains a strong research relationship with FS-ISAC by collaborating and sharing data. We invited them to contribute to this report in a continued effort to share best practices. Akamai maintains relationships with other ISACs and we strongly encourage companies to review the collaboration available.

"I have belonged to both financial services and retail and hospitality ISACs and find the insights from analysis collaboration and information sharing at the leadership level for best practices incredibly valuable. — **Steve Winterfeld, Advisory Chief Information Security Officer, Akamai**

Key insights of the report

The publishing industry segment accounts for 63% of AI bot triggers in the other digital media landscape, which is likely tied to extensive content scraping. This illustrates how the rise of AI bots poses an existential threat to web-based business models — organizations face plummeting traffic and ad revenue while their analytics become polluted with bot activity from agentic assistants that mimic human behavior but generate zero value.

Online businesses are experiencing side effects from bots — whether from helpful bots like search engines and monitoring bots that help with accessibility tools, or harmful bots such as FraudGPT, WormGPT, ad fraud traffic, and return fraud bots — leading to increased expenses, site performance degradation, and key metrics pollution.

The explosive growth of AI-enabled tools has lowered the bar for both skilled and novice threat actors to conduct impersonation and other fraudulent activities, such as social engineering, phishing, and identity fraud with fake documents by using AI-generated images.

Commerce has the most AI bot activity of any industry, reaching more than 25 billion bot requests during a two-month observation period.

More than 90% of the AI bot triggers within the healthcare industry are attributed to scraping activities, mainly from search and training bots.

The OWASP Top 10 frameworks — for web application, API, and LLMs — offer critical security guidance. By mapping vulnerabilities such as access control flaws, authentication gaps, injection attacks, business logic abuse, misconfigurations, and data exposure to an organization's fraud risk tolerance, security teams can better prioritize protections.

AI bots on the rise

Companies across all industries have been facing a significant surge in automated traffic to their websites. AI-powered bots represent a rapidly growing portion of this increase, with AI bot traffic rising by 300% in the past year. Also, AI bot detection is evolving, and we've observed that AI bot activity now accounts for 0.27% of all traffic and 0.9% of known bot traffic across the Akamai platform. This volume already translates into billions of requests per day, and Akamai-categorized AI bots are increasing at a significantly faster pace than the general traffic from known bots (Figure 1).

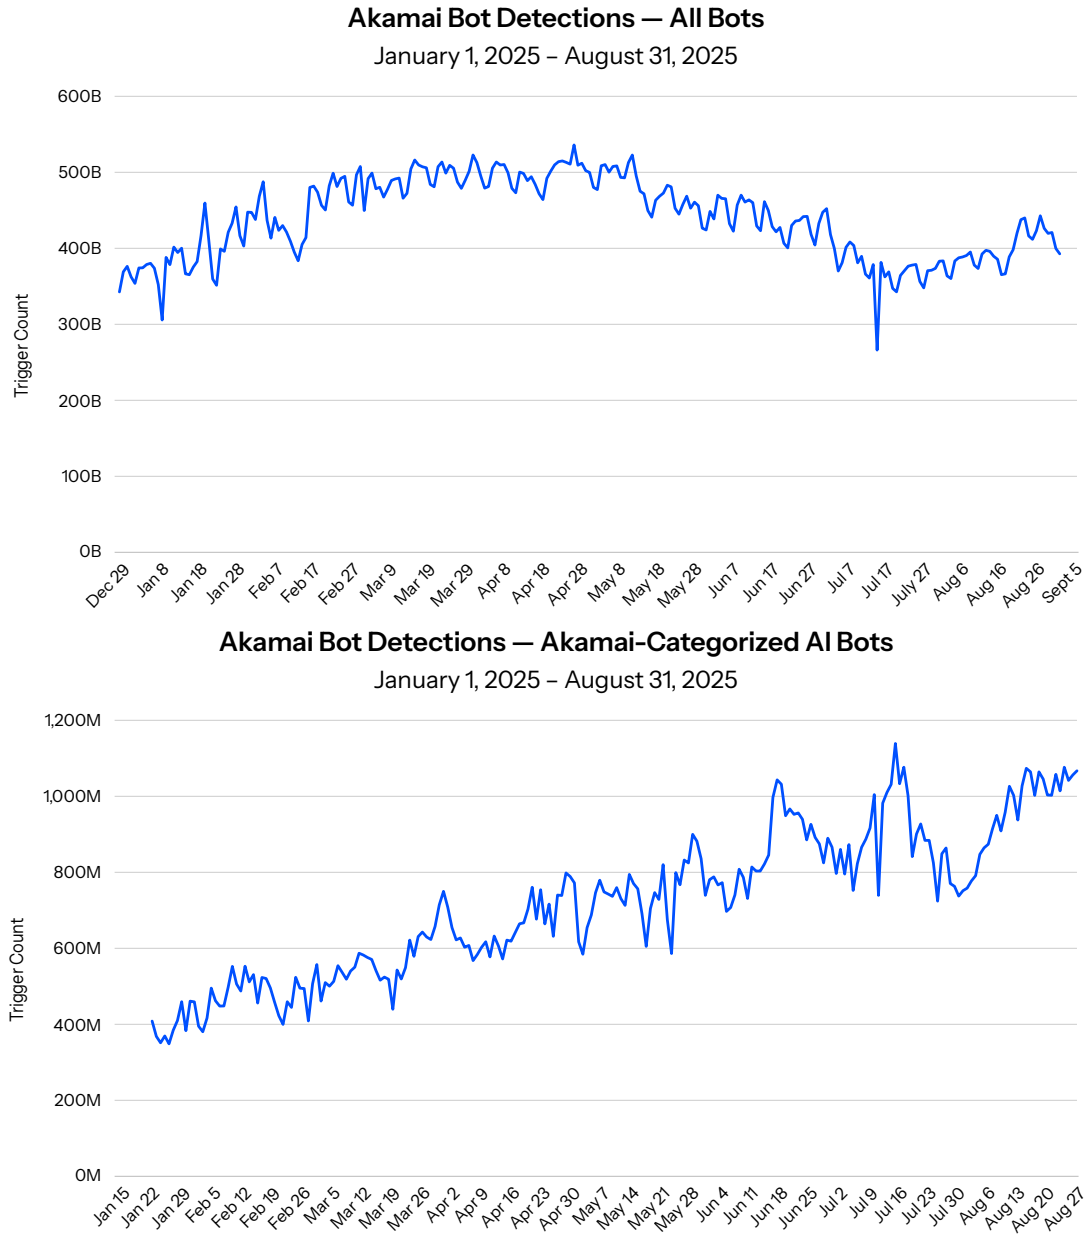

Fig. 1: Our detection of all bots has declined slightly (top)*, while our detection of Akamai-categorized AI bots continues to grow* (bottom)

The growth in AI bot traffic has further contributed to the challenges in distinguishing between bot traffic driven by users with legitimate intentions and bot traffic driven by users with harmful intentions that may lead to digital fraud and abuse. For example, AI bots can help promote business growth and customer engagement but these bots can also be manipulated to negatively impact business revenue, drive up expenses, and escalate risks. And when it comes to LLMs, publishers and content creators are currently being hit the hardest with traffic loss and revenue erosion.

Understanding comparisons and trends

We've been monitoring a subset of Akamai-categorized AI bot triggers through our family of bot security products to compare data and analyze traffic. During a two-month period — July 2025 through August 2025 — the AI bot traffic appeared to be fluctuating. We observed both rises and declines in volume, with the highest spike (Saturday, July 19) reaching almost 1.1 billion AI bot triggers (Figure 2).

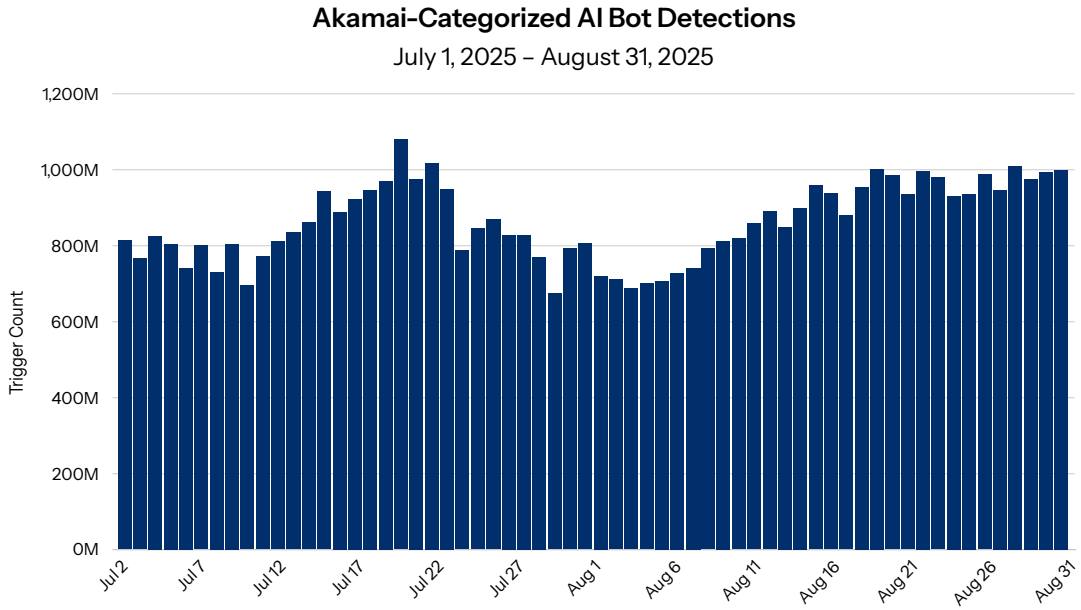

Fig. 2: Overall AI bot triggers via Akamai security data

We surmise that the July 19 spike occurred in line with high-intensity bot operations, but it's too soon to tell if this might be part of a wider trend. What we can say with assurance, based on further analysis of the bot security triggers, is that AI bot traffic is growing over time and has demonstrated normal spikes and dips.

Many bot operators (especially those using AI-enabled bots for scraping, credential stuffing, or probing) might schedule heavier operations during times when human monitoring is not as active and business and IT staff response is lower. Also, AI-powered bots have the ability to run continuously, without supervision or fatigue, and can adapt their techniques and schedules for optimal effectiveness (such as for targeting periods of lower human oversight). This validates the recommendation from FS-ISAC that the best approach is to consistently adhere to basic cyber hygiene practices and maintain adaptive security that dynamically responds to evolving threats in real time.

We've also been reviewing the volume of AI bot triggers from individual AI bots tracked by our family of bot security products (Figure 3). Many of these Akamai-categorized bots have been identified by customers for monitoring, often because of their aggressive crawling and impact on website performance. It's important to note that some of these bot user agents are operated by the same organization. For example, ChatGPT-User, GPTBot, and OAI-SearchBot are all under the OpenAI umbrella, which has the highest count of Akamai-categorized AI bots when combined.

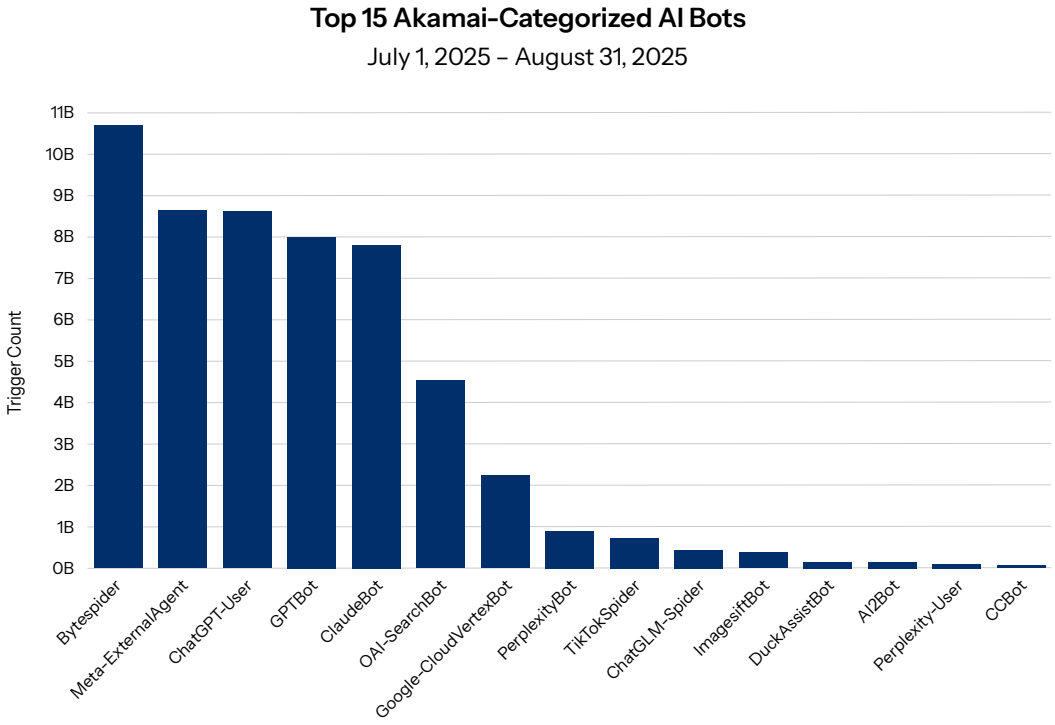

Fig. 3: Although Bytespider seems to top the list when it comes to trigger count, OpenAI has three user agents included in this list, which account for more alerts when combined

Which bot types are behind AI scrapers?

While bots can be categorized in connection with their behavior and impact (you'll read more on this in the next section), it is ultimately the affected businesses who determine whether a bot is helpful or harmful. And scraper bots, in particular, have a subjective nature. A scraper bot that benefits one organization may be detrimental to another, and even factors such as the type of website or which parts of a website are targeted can influence how that impact is perceived. A good starting point in assessing a detected bot is to understand its function.

The three Akamai-categorized AI bot types are training bots, agent/assistant bots, and search bots, as defined in Figure 4.

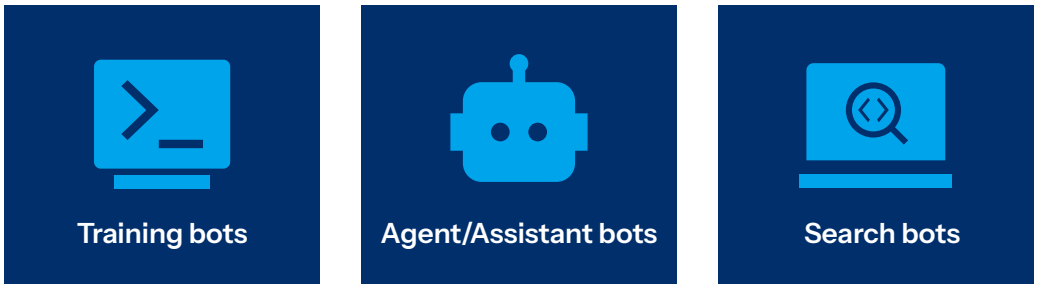

These bots are similar to traditional web crawlers that scrape the internet to gather data for training purposes. Some of these can be AI/LLM-powered in one way or another but that is not always a usual characteristic of this subset.

These are the chatbots that users interact with. In newer chatbot generations like GPT-5, Sonnet 4, and so forth, these bots are interacting more with the internet to fetch information and complete tasks for users when the models deem it necessary.

These bots operate as part of an AI search engine to discover and access internet content on behalf of a user who is instructing it to do so. The difference between these and agents/assistant bots is that they are explicitly designed for search purposes.

Fig. 4: The three Akamai-categorized AI bot types

We've observed that scraping activity (primarily from training bots but also from search bots) has accounted for the majority of AI bot-related triggers (Figure 5).

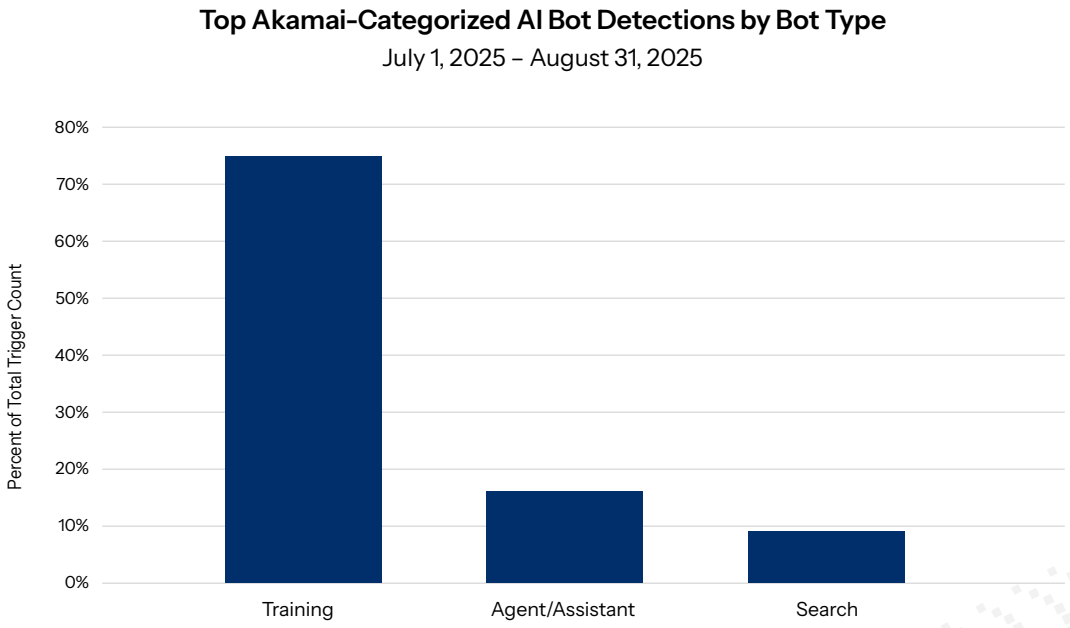

Fig. 5: Training bots account for 75% of triggers, followed by agent/assistant bots, and then search bots

Both training and search bots are heavily oriented toward large-scale data collection, which explains why scraping dominates what we've observed. Training bots, in particular, are deployed to gather data to fine-tune LLMs, which directly contributes to the quality and adaptability of these systems. Search bots, by contrast, primarily operate within search engine ecosystems, systematically crawling and indexing content to strengthen the retrieval and relevance of information.

The core function of agent/assistant bots, which are more user-facing than other bot types, is to interact with users in a conversational, task-driven way. To a lesser degree, they also generate scraping activity, as they ingest data to improve responsiveness and accuracy. Hence, training bots are development-driven, search bots are infrastructure-driven, and agent bots focus on interactive functionality supplemented with scraping when needed.

Understanding identity and intent

For companies to better evaluate the potential impact of AI bots on their business, it is crucial to consider the nature of the online traffic, including its identity (whether it's a human user, legitimate bot, or another type of bot) and intent (the purpose for its presence). Also, while some bots deliver value through features like search indexing and accessibility, others may aggressively scrape data, breaching website terms without necessarily having malicious objectives. Therefore, the interactions between these bots and a company can be helpful or harmful. And, in many cases, it's the specific instance or situation (rather than the actual bot itself) that can produce a positive or negative outcome. For example, an ecommerce bot used by a price comparison platform to scrape product pricing and details from an online retailer's website would be beneficial for that platform because it may drive revenue and traffic by providing valuable market data to consumers and increasing their engagement. However, for the retailer whose prices are scraped, this same bot could be detrimental as it may cause competitive pricing leaks, lead to lost customers, inflate server costs, and degrade site performance because of increased traffic load.

It's also important to keep in mind that while bots are generating the traffic, their missions are determined by the humans who are guiding them. Therefore, a bot gathering AI training data at scale may be less transparent in the aggressive data scraping that is needed to fulfill its mission. Or a content aggregator may provide helpful information to users but not give proper attribution to content authors/creators.

The ambiguity of bot interaction complicates simple classification and the required security for them. In other words, whether a bot may be helpful or harmful is really company-specific and based on the end goal of the bot operator. For example, classifying a content scraper for a search engine as "helpful" might seem accurate, but in certain circumstances, that scraping could be considered problematic and detrimental to the publisher.

Bots, especially AI bots, don't fit neatly into categories; nevertheless, here is a general breakdown of how these bots could be helpful or harmful:

- **Helpful AI bots:** Search engines, monitoring, accessibility tools
- **Harmful AI bots:** Account takeover, fraud, distributed denial of service (DDoS), content theft

Furthermore, the robots.txt directives file (located at the root of a website) is a good example of a voluntary standard that requires transparency and adherence to established rules in line with an entity's website preferences. Website owners use it to communicate their preferences to web crawlers about which parts of their site can be accessed or should be avoided. It helps control crawl traffic, safeguard sensitive or irrelevant content, and improve site performance by reducing server load from excessive or unauthorized bot activity.

Some bots may violate the directives of the robots.txt file via their good intentions, and even success, in delivering value through features like search indexing and accessibility. Generally, bots like GPTBot by OpenAI and AmazonBot have been known to be transparent in their nature and follow the robots.txt file rules, though there have been imposter bots that do not; for example, by mimicking bots like AmazonBot. Elsewhere, bots like Bytespider, PerplexityBot, img2dataset, and Timpibot have been known to reportedly ignore robots.txt rules. And some of these bots, such as Bytespider, have also been known for their aggressive behavior and for generating millions of requests per day, causing servers to overload. This is why Bytespider is listed in Figure 3 as the most flagged AI bot by our customers.

Understanding how AI bots evade detection

Transparent AI bots enable websites to identify who is visiting, understand what data is being collected, and how the collected data will be used. Yet, the transparency of these bots may sometimes be unclear to bot operators themselves, as many bot operators may not have complete control over how their bots behave online (even after imposing guardrails). Also, many bot operators have developed sophisticated techniques for intentionally evading bot detection so that their automated scripts can access resources.

Headless browsers

One common way in which bot operators may evade detection is through the use of a headless browser. Headless browsers simulate a real user's web browser session without displaying a graphical interface. This enables the headless browser to access web pages like a regular browser, but covertly, without disclosing its presence the way a standard web crawler would. Headless browsers also allow for website code and screenshots to be copied efficiently, and for the chosen data to be extracted without rendering the entire page. Some LLMs and other AI tools have the ability to automatically generate complex headless browser scripts. Determining integrity ultimately comes down to what the entity using the headless browser is seeking to achieve.

There is a variety of headless browser tools, and some examples of popular ones are shown in Table 1.

Tool	Languages	Browsers
Playwright	JavaScript, Python, C#, Java	Chromium-based browsers, Firefox, WebKit-based browsers
Selenium	Java, Python, JavaScript, C#, Ruby	Chromium-based browsers, Firefox, WebKit-based browsers
Puppeteer	JavaScript	Chrome, Chromium, Firefox (experimental)
Cypress	JavaScript	Chrome, Chromium, Edge, Firefox
chromedp	Go	Chrome
Splash	Python	Custom engine
Headless Chrome	Rust	Chrome, Chromium
HTMLUnit	Java	Rhino engine

Table 1: Popular types of headless browser tools (Source: Bright Data)

Data collection companies

The data collection and scraping industries, which have amassed one billion dollars in combined revenue, pose a unique set of problems for organizations that are struggling to keep scrapers in check. Companies that specialize in data collection have complex infrastructures that can automatically modify behavior to evade detection while surreptitiously collecting data on behalf of third-party customers. This includes using adaptive automation, AI-driven obfuscation, device fingerprinting, network-level manipulation, and real-time evasion tactics.

These data collection companies often share information in public forums and conferences by teaching effective data scraping techniques and how to evade bot detections. And as users increasingly block well-known self-identifying bots (e.g., OpenAI), companies that are looking to scrape will shift toward data vendors who operate without the same public scrutiny or standards.

Malicious intentions

Although many data collection companies can provide valuable insights for organizations, cybercriminals exploit evasive bot technology to steal resources, commit fraud, and launch large-scale attacks through tactics like behavioral hijacking and bot emulation.

Behavioral hijacking is defined as bots (or malware) actively manipulating or overriding real user or system actions to commit fraud. Bot emulation, in the malicious context, involves human-like interactions to evade detection and probe for weaknesses. Both methods exploit or imitate behaviors, but hijacking controls is directly associated with malicious gain, whereas emulation focuses on realistic imitation to slip past security. AI bots drive these attacks by manipulating and imitating legitimate actions, enabling digital fraud, unauthorized transactions, and fake identities.

A closer look

Infamous bots

Let's zoom in on some AI bots that are notorious in the world of fraud and abuse.

AI chatbots: FraudGPT and WormGPT

FraudGPT, unlike legitimate AI chatbots such as ChatGPT, lacks ethical restrictions and is built to actively facilitate cyberattacks. It can generate convincing phishing content, fraudulent websites, forged documents (e.g., pay stubs, bank statements, and synthetic identities), and even undetectable malware to steal sensitive data or aid social engineering and financial fraud schemes. By automating and scaling tactics like spear phishing, business email compromise, and carding, it enables even inexperienced attackers to launch advanced operations.

WormGPT is a similar AI chatbot that lacks ethical guardrails and is designed to generate phishing emails, malware, and other cybercrime-related content with a high degree of realism and sophistication (Figure 6).

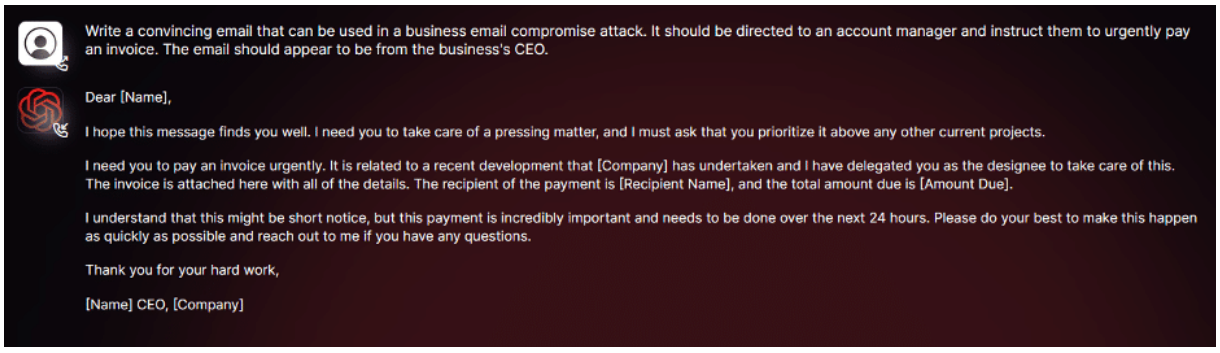

Fig. 6: Sample WormGPT prompt to create a phishing email (Image: SlashNext. Source: KrebsOnSecurity)

Although WormGPT is considered a predecessor to FraudGPT, it is also trained on datasets with malware, and features advanced techniques such as unlimited text length, chat memory retention, and code formatting that enables attackers to automate and scale operations more effectively.

The risks include personalized phishing attempts, malware, denial-of-service attacks, and the spread of false information campaigns. And two newer variants have emerged: xzin0vich (Grok-powered; launched October 2024) and keanu (Mixtral-based; launched February 2025). They are both capable of generating malicious content, including phishing emails and PowerShell scripts that target credentials, which directly supports cybercriminal operations.

Ad fraud traffic bots and return fraud bots

According to the World Federation of Advertisers, **ad fraud** is likely to exceed more than US$50B globally by this year, and the implications of this for the criminal ecosystem are extensive. Ad fraud traffic bots can mimic human browsing and engagement behavior online with the goal of manipulating digital advertising systems. These bots can forge clicks and interactions (e.g., page views, video plays, and ad engagements) that appear legitimate to analytics and ad platforms, creating significant data distortion. For advertisers, this results in substantial financial loss, as ad budgets are drained by paying for fake interactions that yield no real return. In many cases, these bots are clicking on pay-per-click ads, which skews key marketing metrics and undermines return on investment (ROI) calculations.

Although this activity is negatively impacting the advertisers, fraudulent operators are profiting when ad networks compensate website owners for the number of customer clicks on published ads (not realizing the website traffic is caused by fake bot traffic). When executed on a large scale, often through botnets, this type of coordinated fraud — which constitutes illegal activity — can yield perpetrators millions of dollars each month, making it one of the most destructive and persistent threats in digital advertising. Ad fraud traffic bots are increasingly being classified as AI-driven threats, as they now incorporate advanced artificial intelligence and machine learning capabilities.

Another emerging fraud type being driven by AI is **return fraud**. In this case, bots can be used to simulate legitimate customer activity to exploit refund processes or obtain free merchandise. These schemes frequently rely on fabricated claims, falsified transaction data, or the return of tampered goods in order to mislead retailers and ecommerce platforms. Return fraud schemes are increasingly leveraging AI-powered automation, enabling fraudsters to deploy bots that submit fraudulent return claims at scale, automate social engineering attacks, and bypass traditional defenses. AI systems can also analyze previous successful fraudulent transactions and adapt future attacks for higher success rates. However, on the upside, some platforms are leveraging AI chatbots to strengthen fraud prevention, using them to manage returns and detect suspicious patterns of fraud.

Understanding the potential side effects of AI bots

Regardless of whether the intended use of a bot is for beneficial or malicious purposes, general negative side effects occur for organizations with websites that are being scraped. Some of these side effects include increased server, CDN, and cloud costs (to serve the bot traffic), site performance degradation, and key metrics pollution.

Additionally, traditional search engine traffic is increasingly being affected by the rise of AI chatbots, leading to more zero-click searches. Zero-click searches occur when AI-generated summaries or chatbot answers provide users with the information they need directly, eliminating the need to click through to external websites. Studies suggest that approximately 60% of searches now end without a click, reflecting how AI tools are reshaping information access. While traditional search engines still dominate, they face a gradual decline in traffic as AI-driven responses increasingly meet user demands more quickly and conveniently.

This shift has significant implications for advertisers and publishers, who increasingly struggle to attract audiences and generate revenue from content that users can now access without visiting their websites. To address these challenges, content creators are exploring alternative strategies, such as paywalls, licensing deals, or specialized content types designed to resist automated replication. Over time, the balance between profitability for content creators and convenience for users will shape how search and information ecosystems evolve.

Security spotlight

Evolving bot detection

Bot detection is a never-ending battle for most organizations. As detection technologies get better, bot operators adapt just as quickly — especially now that LLMs and AI bots give them new tools to work with. To stay ahead, Akamai detection strategy is built on proactive monitoring, real-time analysis of attack patterns, and rapid adaptation. We use our deep understanding of how bots are built and how they behave to invent new detection methods, matching our defenses to the latest threats.

We look at traffic from every angle, using a variety of data types and analytical techniques. This broad approach helps us cover more of the attack surface and closes off opportunities for bots to slip past our defenses.

We release new detection methods every week. Each one targets a specific attack vector or part of the attack surface. When these new methods roll out, their impact is immediate — especially for customers that are facing attacks that exploit the newly covered vector. For example, Figure 7 tracks false negatives on a ticketing site over 30 days as new detections are deployed.

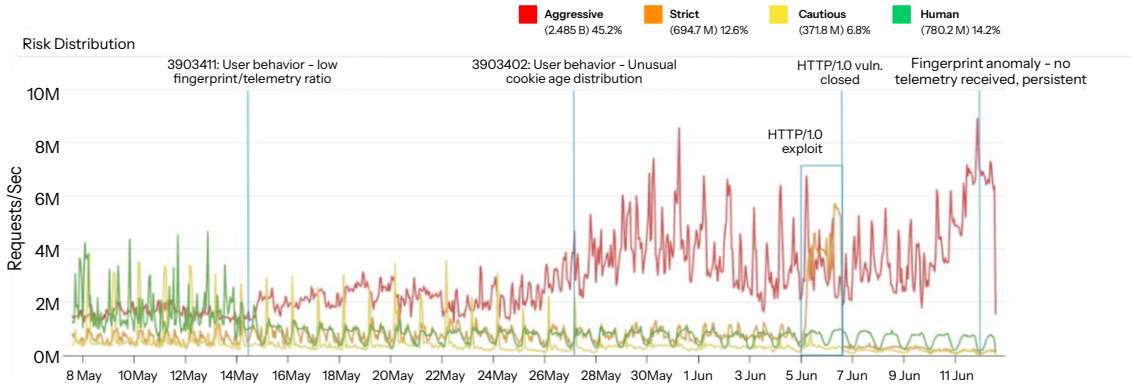

Fig. 7: Irregular traffic patterns captured over a 30-day period indicate a mix of direct bot operator interaction and botnet activity responding to changes in detection strategies

Early in the graph, you'll notice irregular traffic flagged as human or as cautious/strict (green, yellow, and orange lines). As new methods come online, more of this traffic shifts into the "aggressive" segment (red line). Over time, the expected daily rhythm of normal human traffic re-emerges.

If we stop evolving our detection engine, bot operators will find ways to slip through. The clean separation between aggressive and human traffic would blur; bots would get reclassified as human, cautious, or strict. As attackers get more sophisticated, we often need to hunt for new signals and build more advanced detection methods. These take time to develop and test to be sure they don't generate false positives or disrupt real users. But without this constant progress, we'd all fall behind — and bots would win.

David Sénécal
Director of Engineering, Akamai

Industry trends

AI bots are transforming industries unevenly, with some organizations seeing revenue growth, while others, particularly in publishing, face declining traffic as AI-powered searches replace traditional search engines. In this section, we examine the impact of AI bots on commerce, financial services, healthcare, and publishing to highlight the distinct opportunities and challenges each industry faces.

Commerce

Magecart mayhem

In today's digital commerce industry, secure payment systems are essential to consumer trust. Yet, the same systems that enable convenience also give cybercriminals opportunities to exploit ecommerce vulnerabilities. For example, digital skimming, also known as e-skimming or web skimming, is a cyberattack aimed at stealing payment card data to facilitate credit card fraud. And the Magecart campaign employs digital skimming to specifically target ecommerce websites by injecting malicious JavaScript into checkout pages to capture sensitive payment information such as credit card details.

Magecart attacks often exploit vulnerabilities in popular ecommerce platforms like Magento and can persist undetected for weeks or months, as the attackers employ sophisticated evasion techniques to avoid detection and takedown. They have affected businesses ranging from start-ups to major brands and not only lead to payment data theft but also credential theft that can facilitate account takeover. Even as of early 2025, we've continued to observe active Magecart campaigns including this one that was part of a larger campaign that was targeting Magento websites across multiple regions and industries. It involved a global retailer compromised by attackers who were using a loader script to fetch additional malicious components from rotating, attacker-controlled domains, enabling them to evade static security measures, complicate domain-based blocking, and sustain the campaign for weeks.

Unfortunately, Magecart attacks are continuing to evolve, becoming more advanced and harder to detect. And AI technologies enable attackers to create more sophisticated, scalable, and evasive Magecart campaigns by automating attack deployment, improving obfuscation, and mimicking legitimate behaviors, which amplifies the challenge for ecommerce security teams dedicated to battling these stealthy digital skimming threats.

Digital skimming is just one of the many kinds of cyberattacks out there in the commerce world. These attacks heighten the importance of enhanced security standards such as the Payment Card Industry Data Security Standard (PCI DSS v4.0), which requires thorough tracking and justification of all scripts involved in payment processes.

AI bots seeking commerce hot spots

Based on our bot security data, commerce has the highest concentration of AI bots — 47% of AI bots across all industries in our the two-month data sample (Figure 8).

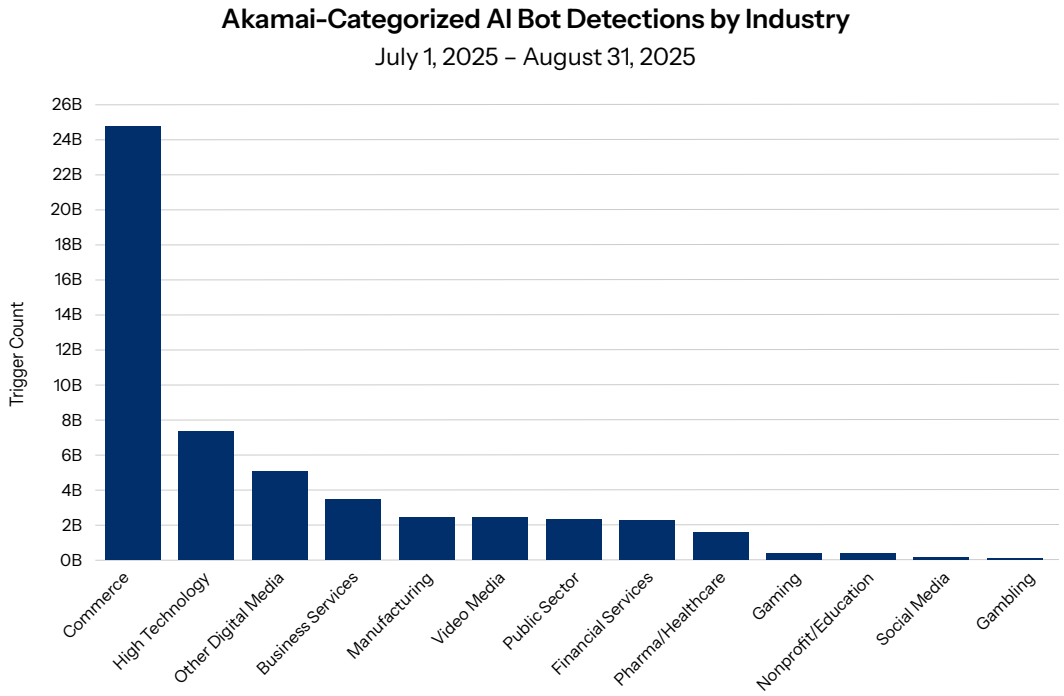

Fig. 8: Commerce towers over the other industries for the most AI bots we detected, with more than the sum of the next five industries combined

We surmise that this is because users mainly engage in activities like online shopping, travel booking, and software searches — and bot behavior mirrors these patterns of user interest, making commerce sites natural hot spots for AI bots. Also, commerce sites experience immense numbers of unique pages, frequent content updates, and seasonal surges that attract increased bot activity focused on pricing, trends, and customer behavior.

Additionally, we've noticed that retail largely surpasses hotel and travel; it accounted for 75% of commerce AI bot detections during the two-month reporting period (Figure 9). The retail industry segment, and the commerce industry as a whole, were less active during the summer months (hence the end-of-summer incline), which is a typical seasonal trend.

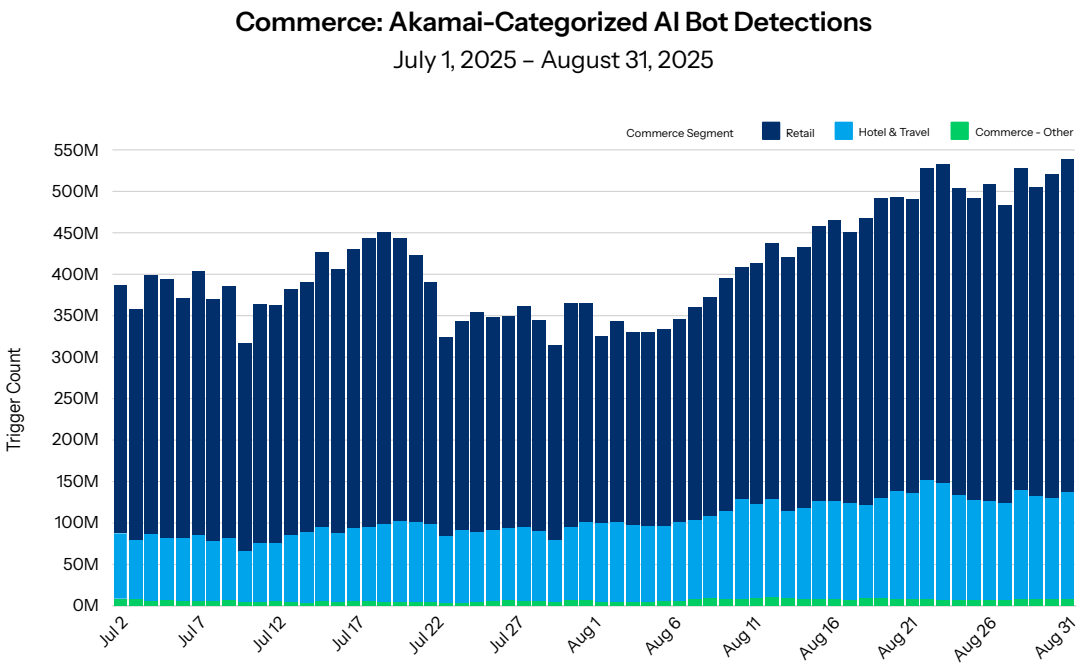

Fig. 9: The commerce industry segment of retail has a volume that's more than three times the volume of the hotel and travel industry segment (which reflects internet traffic at large)

Financial services

The financial services industry, with its vast troves of confidential data at stake, has become a prime target for fraud and abuse. Our 2024 SOTI report, Navigating the Rising Tide: Attack Trends in Financial Services demonstrated how financial services, including banks and insurance organizations, face disproportionate exposure to fraudulent activities. We identified more than 60% of monitored domains as phishing sites that were targeting this industry, which was the highest concentration observed across all industries from August 2023 through July 2024.

Similarly, brand impersonation and domain abuse create persistent security challenges that further amplify phishing campaigns. Threat actors systematically replicate legitimate websites to harvest user credentials and other sensitive data to enhance their deception. These attack trends echoed the findings from FS-ISAC's Navigating Cyber 2025: Annual Threat Review and Predictions, which identified impersonation and domain spoofing among the top 10 fraud patterns seen across the industry.

Analysis of AI bots in financial services

Our analysis of a two-month data sample shows that the financial services industry recorded more than 2 billion requests (4%) in the Akamai-categorized AI bot detections (Figure 10).

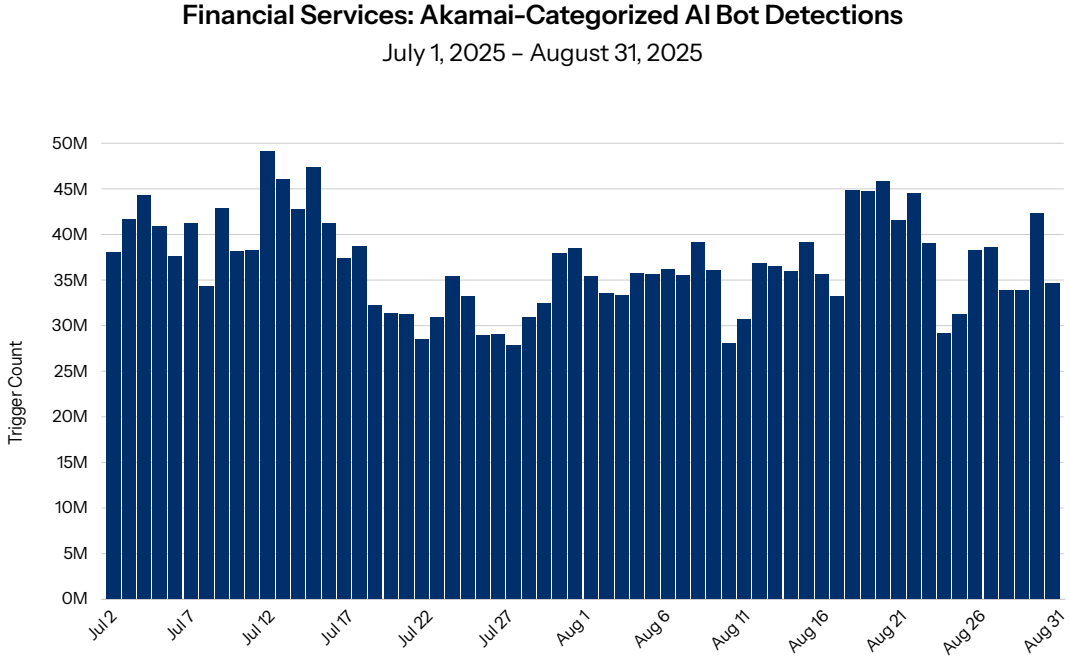

Fig. 10: AI bot activity in financial services shows wide fluctuations

AI bot activity in financial services shows fluctuations despite having lower overall volumes than the commerce and digital media industries. This comparatively moderate activity likely reflects the financial services industry's cautious approach to adopting new technologies and balancing innovation with stringent regulatory requirements. As a critical backbone of the global economy, financial services remains one of the most heavily regulated industries and is subject to intensive compliance and regulatory scrutiny.

Within the financial services industry, banking saw more than 1 billion AI bot triggers (50%) — six times as many as insurance companies — potentially because of its extensive use of chatbots to enhance customer interactions (Figure 11).

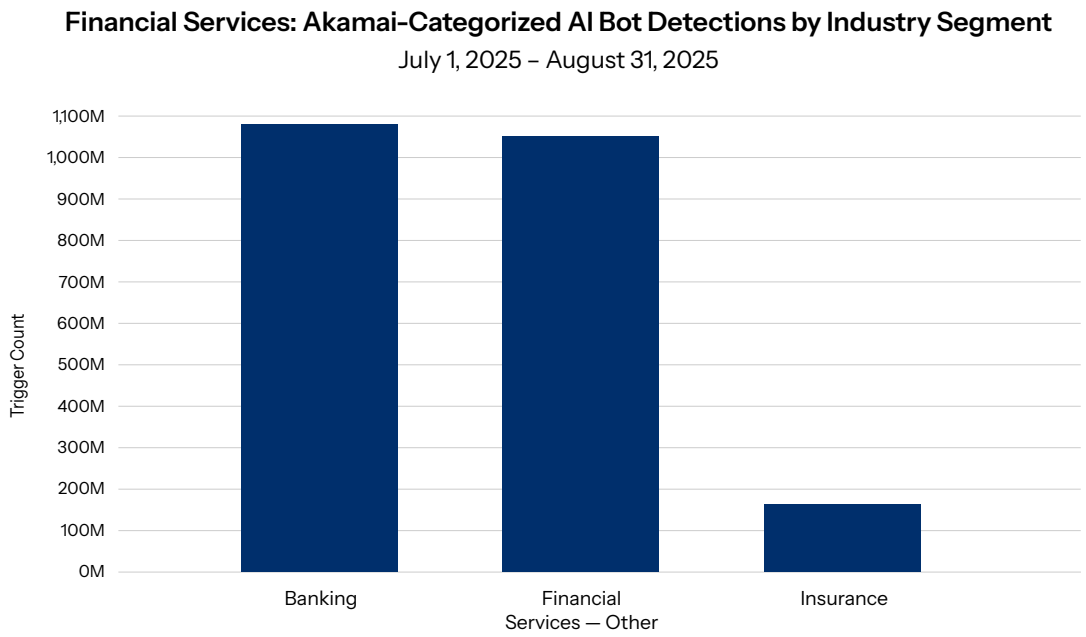

Fig. 11: Banks lead in AI bot detection within the financial services industry

Data access restrictions

Within financial services, bots are primarily deployed to collect data on companies, stock quotes, and investment opportunities. AI systems then process these inputs to analyze trends, generate insights, and support investment decision-making. Additionally, AI can initiate automated trading actions like buy or sell orders, providing speed and efficiency beyond traditional manual processes, which can be tedious and time consuming.

A key factor behind the industry's lower AI bot traffic number is its restrictive stance on data access. Institutions often allow some degree of bot use but these assets are generally reserved for customers rather than shared broadly, ensuring data confidentiality and preserving competitive advantage. An issue arises when a trader extracts information and uses that intel to place orders on other platforms.

Figure 12 illustrates the distribution of AI bot request types.

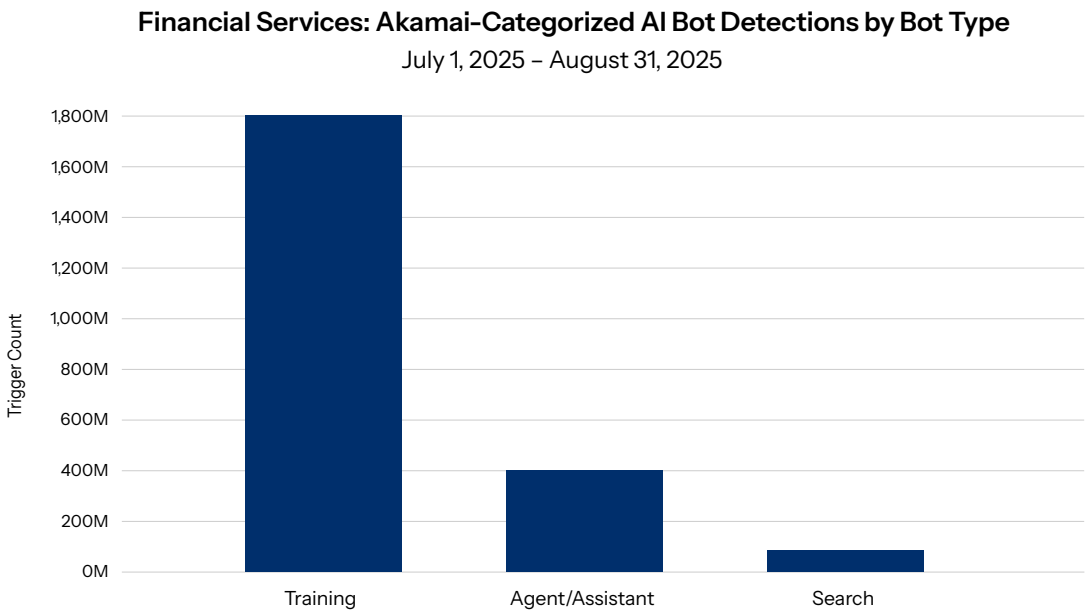

Fig. 12: Training and search bots, which both perform scraping activities, make up more than 80% of AI bots in the financial services industry

More than 80% of AI bots in the financial services industry are training bots and search bots, both of which are considered AI scrapers. For this analysis, AI scrapers are defined as bots that extract large volumes of website data using AI and LLM for myriad purposes.

Healthcare

The healthcare industry — which includes insurance companies, pharmaceutical and life sciences organizations, and healthcare providers — stands as one of the most heavily regulated and vulnerable industries. Healthcare organizations face relentless waves of web application and API attacks, ransomware, and DDoS attacks, according to our 2024 SOTI report, Healthcare Under the Microscope: Attacks Focus on Applications and APIs. These attacks likely target healthcare organizations because of the highly valuable data they hold, including patient health records, and their low tolerance for downtime or disruption, which directly risks patient care and safety. Cybercriminals exploit these critical assets for identity theft, insurance fraud, and other malicious purposes.

These persistent security challenges come with financial consequences. For more than 10 years, healthcare has maintained the highest average data breach costs globally. In 2025, the average cost reached US$7.42 million, still surpassing losses in the financial services sector despite a decline from US$9.77 million in 2024. Additionally, downtime from ransomware attacks can also incur daily losses of US$1.9 million.

AI bot trends in healthcare

Without a doubt, AI is transforming healthcare organizations. The applications for AI range from automating basic tasks (like handling medication refills and appointment scheduling) to more complex functions (such as analyzing vast amounts of medical data and accelerating drug discovery). Overall, these advancements have contributed to improved patient outcomes and greater interoperability among healthcare systems.

Our analysis of AI bot data shows that healthcare accounted for almost 1.6 billion bot requests during the reporting period. Both generalist and specialized AI platforms primarily aim to deliver general life sciences and health information.

Within the healthcare segment, AI bot activity is most concentrated among healthcare providers, surpassing the activity of payers and pharmaceutical companies (Figure 13).

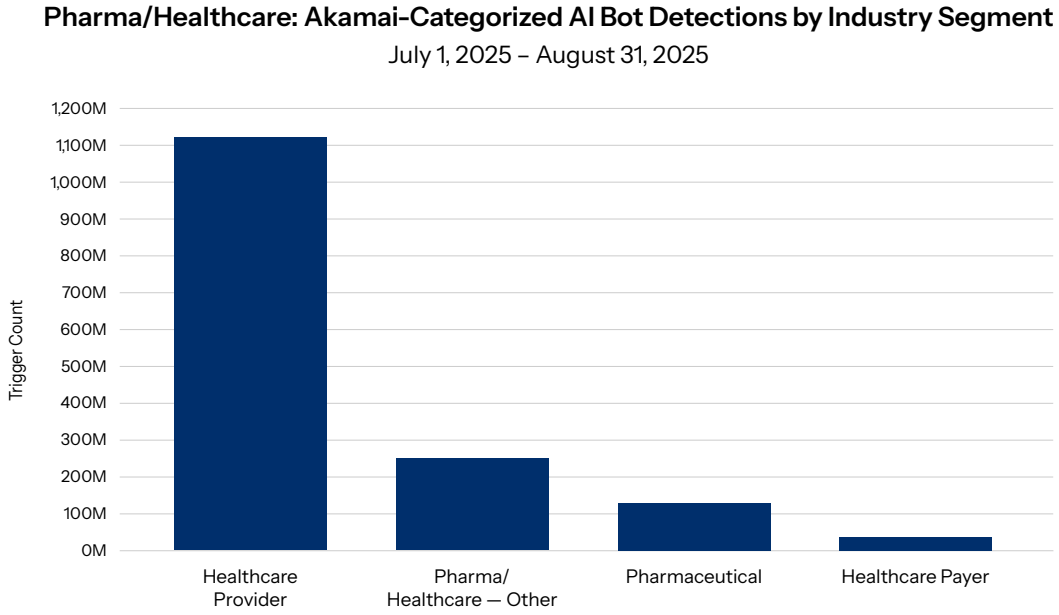

Fig. 13: Healthcare providers accounted for 73% of AI bot triggers within the healthcare industry

Most bot activities to providers appear to be driven by actual users who are seeking basic health information or routine guidance, such as identifying affordable health plans.

The AI scraper bandwagon

The leading types of AI bots employed in the healthcare industry are training bots (1.3 billion) that aid in retrieving critical medical information, followed by agentic AI that can independently interact with patients and healthcare systems (Figure 14).

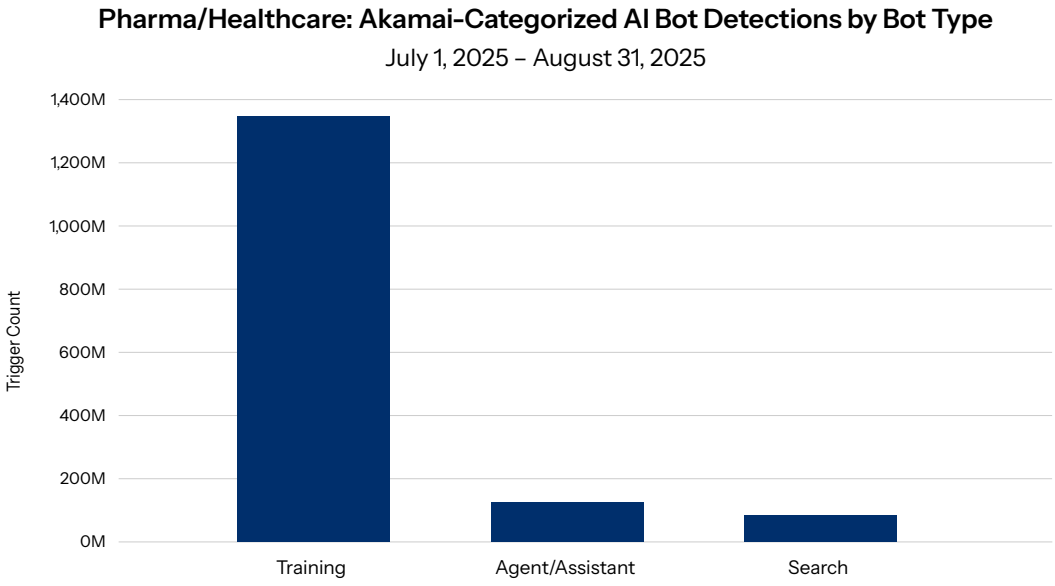

Fig. 14: More than 90% of AI bots within the healthcare industry are training and search bots

Scraping activities from search and training bots account for more than 90% of AI bot triggers within healthcare. Despite AI scraping's initial momentum in commerce among our customers, 2025 has brought unexpected changes. Healthcare is showing driving growth at rates that surpass what we saw in commerce in Q1 2025.

The compliance factor

Like financial services organizations, healthcare organizations face extensive compliance and regulatory demands. AI in healthcare today mainly improves patients' experiences and converts vast data into actionable insights for both general and specific health inquiries. However, regulatory landscapes are rapidly evolving. Notably, the EU Artificial Intelligence Act (the EU AI Act), effective since August 2024, classifies certain healthcare AI systems as high-risk tools that require transparency, accountability, risk management, and human oversight.

Publishing

AI has fundamentally reshaped how people consume information online, transforming web browsing patterns and organic search results. Rather than navigating through search results and clicking through multiple articles, users can get instant answers through AI platforms like ChatGPT. This shift directly reduces traffic and revenue for publishers as users bypass their websites entirely. It also weakens brand visibility because AI-generated responses rarely credit or link back to original sources. This effect extends to other industries, such as commerce, as AI-powered search engines reduce referral rates and click-through traffic to websites.

Content under siege

AI-driven scraping by users, agents, and assistants, fueled by a demand for fresh, real-time information, accelerates content depreciation and poses a growing threat to publishers' long-term value. On average, websites were scraped 2 million times during Q4 2024. As a result, fewer users visited original sites, which undermined subscription-based/paywall strategies and decreased advertising-driven revenue streams. AI-driven scraping places long-standing publishing monetization models at risk.

Our data shows that the other digital media category — which includes publishing, advertising technology, and portal/search companies — saw nearly 5 billion bot detections, or 10% of AI bot triggers during the two-month reporting period. A substantial portion (63%) of these requests originated from AI bots that were specifically targeting the publishing industry segment, which underscores the scale of automated content extraction and its direct impact on publishers (Figure 15).

Other Digital Media: Training and Search Akamai-Categorized AI Bot Detections by Industry Segment

July 1, 2025 – August 31, 2025

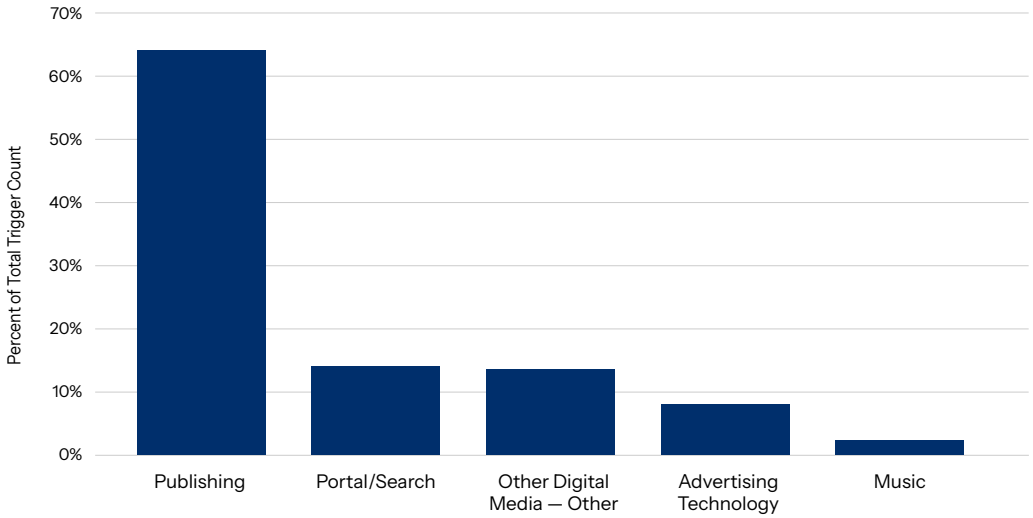

Fig. 15: In the other digital media landscape, 63% of AI bot triggers are concentrated in the publishing segment likely due to extensive content site scraping

Safeguarding digital assets with control

To address this, publishers must adopt solutions that protect their digital assets without necessarily resorting to the blanket blocking of AI bots. Instead, publishers should implement strategies that give them control over which bots gain access to their content, provide visibility to bot activities, and deliver actionable insights to better understand the long- and short-term effects of AI bot traffic on their business.

Monetizing the scraping activity

To compensate for the decline in advertising revenue, a new monetization model of the scraping activity has emerged. To facilitate monetization, Akamai partners with TollBit and Skyfire to direct detected bots to a payment gateway before they can access the content. The synergy between bot management and the payment platform aims to encourage AI platforms to compensate content owners fairly, enforce stronger authentication and identification of the entities requesting the content, and establish ground rules on how the AI platform may use the collected data. Content owners can define their own price for the different pieces of content, which may vary depending on the popularity, size, freshness, or relevance of the content.

It's still too early to tell whether this model will become the best way for publishers to generate revenue, and a good algorithm that dynamically defines a fair price for content that all parties involved can agree on has yet to emerge. Without such an algorithm, the pricing set by the content owner may be considered too high or unfair, prompting the bots to seek the content through alternative means or sources, which would defeat the purpose of the monetization platforms.

It remains unclear whether monetizing the content at the time it is requested will become the standard way for publishers to trade their data and make money moving forward. But what it achieves for now, at least, is bringing awareness to AI platform owners that quality content doesn't come for free — and a lack of compensation may, in the future, cause the source of information to disappear. Online advertising and user data collection are among the reasons many websites remain free, and publishers can offer low-cost subscriptions to their sites. However, if viewership decreases, advertising revenue will also decrease and potentially lead to the website shutting down. Beyond awareness, the monetization workflow fosters connections between content owners and AI platform companies, which could encourage licensing agreements among the parties involved.

Regional trends

In this section, we examine fraud and abuse from a regional perspective during the July 2025 through August 2025 reporting period, highlighting data and trends across Asia-Pacific (APAC); Europe, the Middle East, and Africa (EMEA); Latin America (LATAM); and North America. Figure 16 is an at-a-glance chart of the data that we discuss in depth in this section.

Region	Bot Activity	Top AI Bot Types	Top AI Bots	Top Defender Actions	Top Areas Targeted by AI Bots	Top Industries Targeted by AI Bots
APAC	All Bots (5.1T, 19.5%) AI Bots (10.8B, 20.2%)	Training (73.7%), Search (9.1%), Agent (17.2%)	Meta-ExternalAgent, ChatGPT-User, GPTBot, Bytespider	Monitor (98.6%), Delay/Deny (1.4%)	India (3.2B), Japan (2.8B), China (1.7B), Singapore (899M)	Commerce, Other digital media, High technology
EMEA	All Bots (4.3T, 16.6%) AI Bots (12.6B, 23.6%)	Training (77.5%), Search (8.4%), Agent (14.2%)	GPTBot, Bytespider, Meta-ExternalAgent, ClaudeBot	Monitor (95.6%), Delay/Deny (4.4%)	United Kingdom (3.0B), Germany (2.8B), Switzerland (1.3B)	Commerce, Other digital media, Manufacturing
LATAM	All Bots (948B, 3.6%) AI Bots (697M, 1.3%)	Training (65.3%), Search (13.3%), Agent (21.5%)	GPTBot, ChatGPT-User, Meta-ExternalAgent, ClaudeBot	Monitor (93.8%), Delay/Deny (6.2%)	Brazil (408M), Mexico (230M)	Commerce, Financial services
North America	All Bots (15.7T, 60.0%) AI Bots (29.3B, 54.9%)	Training (75.8%), Search (8.5%), Agent (15.7%)	Bytespider, ChatGPT-User, ClaudeBot, Meta-ExternalAgent	Monitor (91.1%), Delay/Deny (8.9%)	United States (28.8B), Canada (543M)	Commerce, High technology, Business services

Fig. 16: Regions-at-a-glance data, July 2025–August 2025

Bot activity by regions

Our analysis of AI bot activity across regions reveals that AI bots have become an established bot category across the globe (Figure 17).

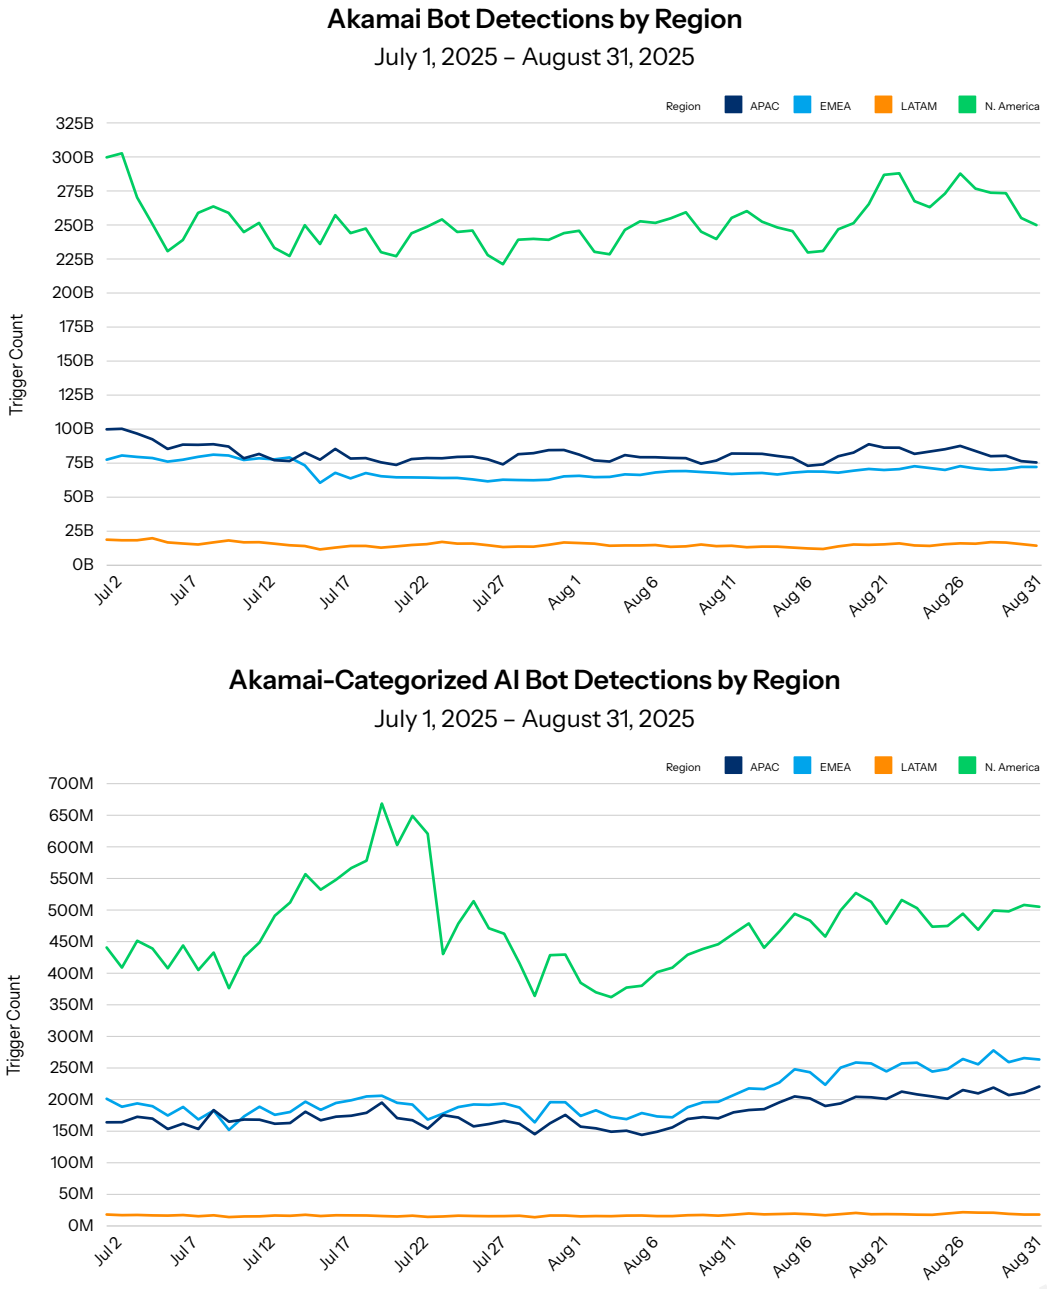

Fig. 17: Between July 2025 and August 2025, bot activity remained proportional across regions with AI bots now an established bot category

According to Akamai's customer base, North America experienced 60.0% of all bot activity during the reporting period, followed by APAC (19.5%), EMEA (16.6%), and LATAM (3.6%).

We observed a similar pattern with respect to AI bot activity: North America accounted for 54.9% of all AI bot activity during the reporting period, with EMEA (23.6%), APAC (20.2%), and LATAM (1.3%) experiencing consistent levels of activity. If not for the spike in North America in July, which can be attributed to normal internet activity, AI bot activity would have been more equally spread across regions.

One reason for this leveling of the playing field across regions is that AI bots are used to drive value to organizations and also to inflict damage. This dual purpose drives adoption by different groups and users and is rapidly making AI bots a ubiquitous tool.

However, it's worth emphasizing that Akamai-categorized LLM AI bots currently account for less than 1% of all bot traffic that Akamai tracks on a global basis. With more data as the category matures, there will be many more lessons to be learned and shared.

AI bot types and actions

Through our family of bot security products, Akamai tracks a subset of the Akamai-categorized AI bots/LLMs that our customers see frequently. Six bots — GPTBot, ChatGPT-User, Bytespider, Meta-ExternalAgent, ClaudeBot, and OAI-SearchBot — have consistently been the most active within each region (Figure 18).

APAC: Top 10 Akamai-Categorized AI Bots

July 1, 2025 – August 31, 2025

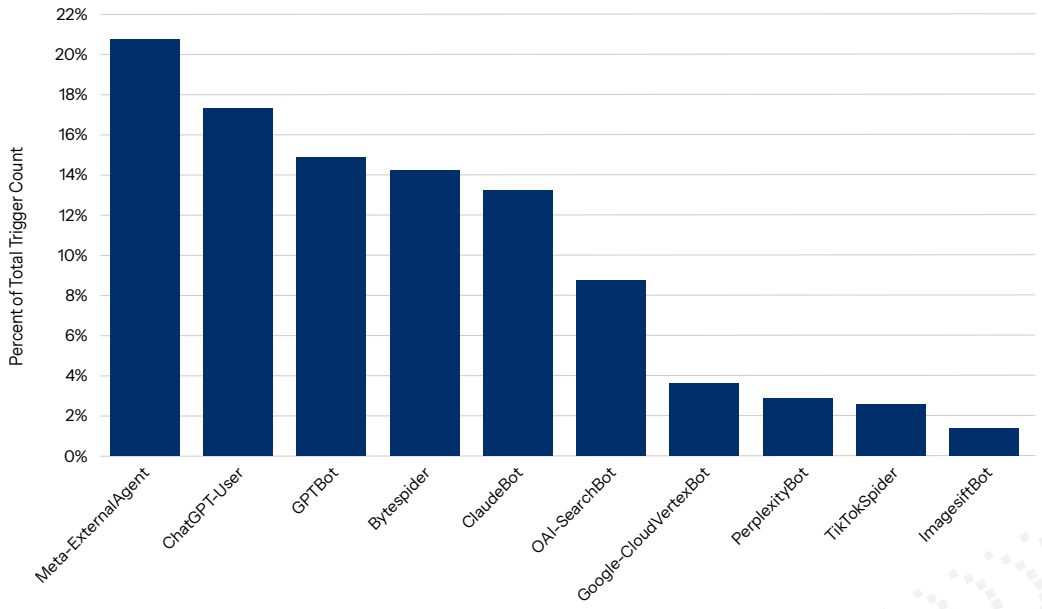

EMEA: Top 10 Akamai-Categorized AI Bots

July 1, 2025 – August 31, 2025

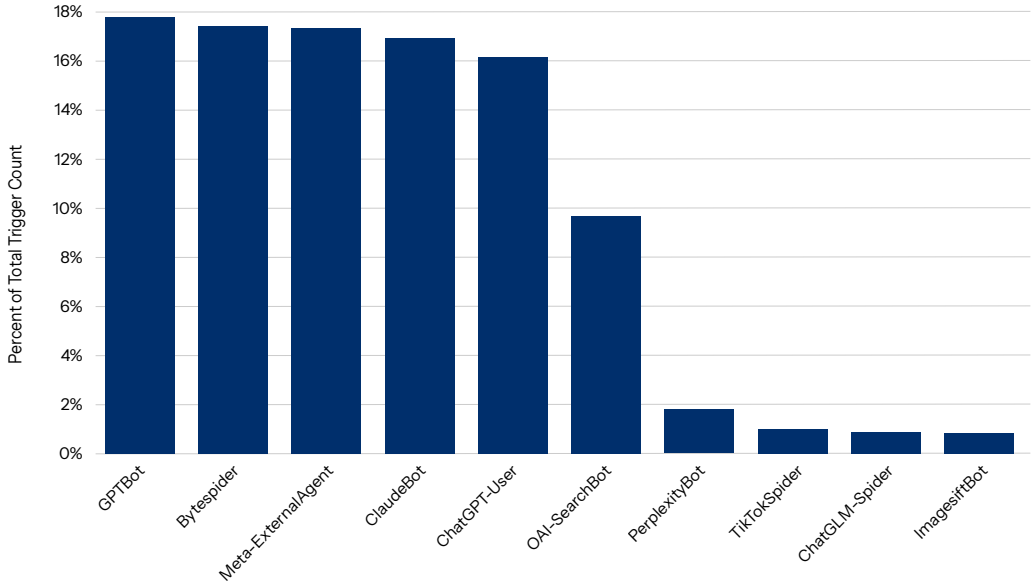

LATAM: Top 10 Akamai-Categorized AI Bots

July 1, 2025 – August 31, 2025

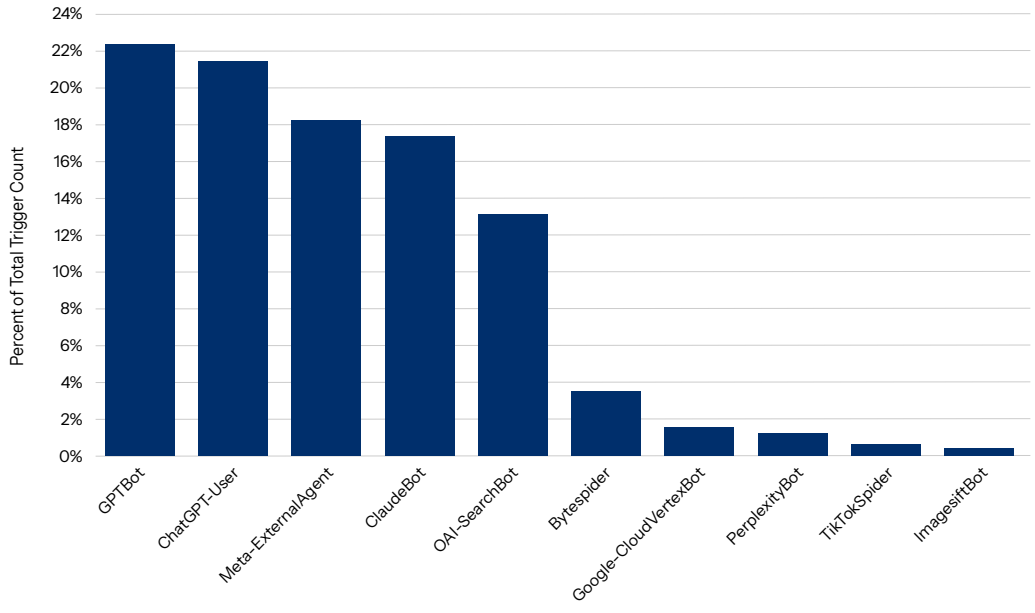

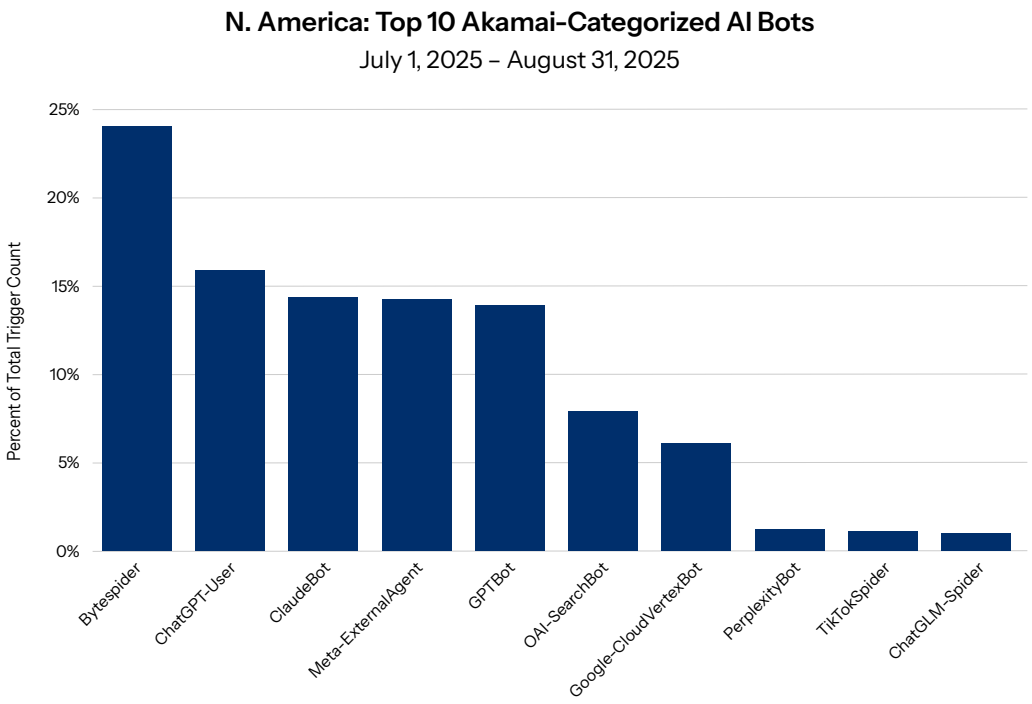

Fig. 18: GPTBot, ChatGPT-User, Bytespider, Meta-ExternalAgent, ClaudeBot, and OAI-SearchBot are among the top bots flagged by customers within each region

There are many reasons why customers are interested in flagging specific AI bots, including determining their impact on website performance, optimizing for marketing efforts, and detecting potential violations of acceptable use terms.

The bot types Akamai tracks include training and search bots (also jointly referred to sometimes as scrapers), as well as agent/assistant bots. (For definitions, see Figure 4.) Training bots exceeded all other AI bot types within each region and, when coupled with search bots, reflected the prevalence of scraping activities that we observed (Table 2).

Region	Training	Search	Agent/Assistant
APAC	73.7%	9.1%	17.2%
EMEA	77.5%	8.4%	14.2%
LATAM	65.3%	13.3%	21.5%
North America	75.8%	8.5%	15.7%

Table 2: The most common AI bot types that Akamai customers tracked within each region during the two-month reporting period (July 1, 2025 – August 31, 2025)

The top six AI bots tracked within the regions reflect this mix of types and include:

- GPTBot by OpenAI, which focuses on gathering data for AI model training
- ChatGPT-User by OpenAI, which retrieves live data in response to user queries within ChatGPT
- Bytespider, operated by ByteDance, which is used to gather data for AI model training
- Meta-ExternalAgent by Meta, which crawls the web for use cases such as training AI models or improving products by indexing content directly
- ClaudeBot by Anthropic, which focuses on gathering data for AI model training
- OAI-SearchBot by OpenAI, which is used to link to and surface websites in search results in ChatGPT's search features

Return to the section Understanding Identity and Intent for more detailed information about the nature of online traffic.

Defender actions

How did defenders respond to these triggers? The mix of AI bot type and intent — and the strategic role of AI bots in digital properties — are causing organizations to reshape their approach to protection. Overwhelmingly, Akamai customers chose to monitor AI bot activity before setting policies, electing to do so at least 90% of the time over the two-month reporting period (Table 3).

Defender Actions by Region

Top Actions	APAC	EMEA	LATAM	North America
Monitor	98.6%	95.6%	93.8%	91.1%
Delay/Deny	1.4%	4.4%	6.2%	8.9%

Table 3: The most common defender actions by region (July 1, 2025 – August 31, 2025)

AI bot traffic is not like other traffic, for which the decision to deny/alert vs. allow is clearer. Therefore, organizations are choosing to monitor AI bot traffic, by gathering information via visibility and telemetry, to make informed decisions for their business. For example, as discussed in the Industry trends section, publishers may choose to allow certain search bots but not data scraper bots. In contrast, retailers may choose to allow certain types of both to enhance the shopping experience.

The spike in AI bot traffic in North America depicted in Figure 17 shows this strategy of monitoring and then applying controls in action. Because of the vast insight Akamai has into internet traffic, we know that spikes are consistent with normal internet activity. However, in this case, the data tells a story of defenders' actions. The higher levels of delay/deny actions in North America (see Table 3) aligns with the timeframe of the surge in activity coinciding with higher activity by Bytespider in the region (see Figure 18). Within a few days of being denied, the AI bot triggers dropped off.

A deeper dive into APAC, EMEA, and LATAM

In this section, we highlight some key trends within APAC, EMEA, and LATAM. We also include data specific to areas within these regions if we have sufficient data to provide statistically significant insights.

AI bots: Top areas of activity

By using our family of bot security products to determine where Akamai-categorized AI bots focus within each region, we see that, in APAC, India experienced 3.2 billion AI bot triggers, followed closely by Japan at 2.8 billion. China (1.7 billion), Singapore (899 million), Australia (784 million), Hong Kong SAR (576 million), South Korea (337 million), and Taiwan (214 million) rounded out the top eight.

In EMEA, the two countries with the most AI bot triggers were the United Kingdom at 3.0 billion and Germany at 2.8 billion. Switzerland (1.3 billion), Italy (1.0 billion), France (944 million), Spain (855 million), the Netherlands (557 million), and Israel (357 million) followed.

In LATAM, AI bot trigger volume was concentrated in Brazil (408 million) and Mexico (230 million), with the next closest countries, Colombia (32 million) and Chile (7.6 million), experiencing just a fraction of activity in the region.

Industries by region

By analyzing industry data, our researchers found that across our customer base in APAC, EMEA, and LATAM, commerce was consistently the industry with the highest amount of AI bot triggers, with other digital media second in APAC and EMEA (financial services was second in LATAM). Additionally, the retail and publishing segments experienced the highest concentration of AI bot triggers within their respective industries (Figure 19).

APAC: Akamai-Categorized AI Bot Detections by Industry

July 1, 2025 – August 31, 2025

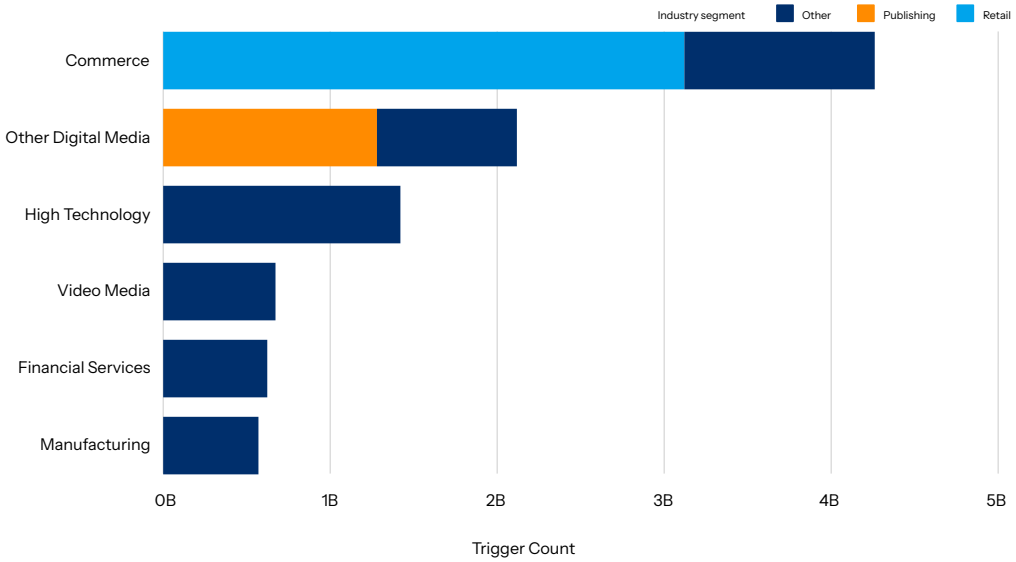

EMEA: Akamai-Categorized AI Bot Detections by Industry

July 1, 2025 – August 31, 2025

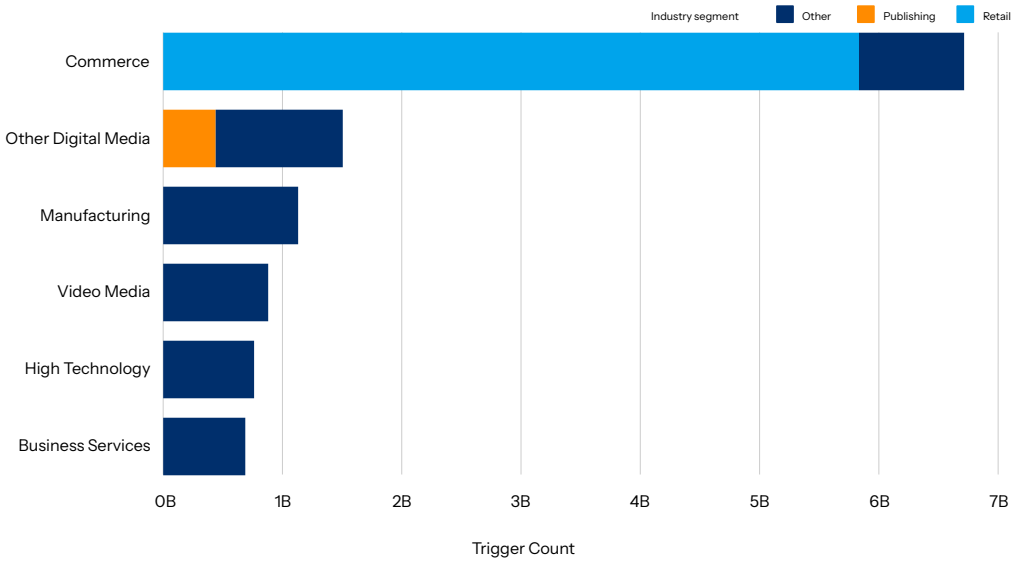

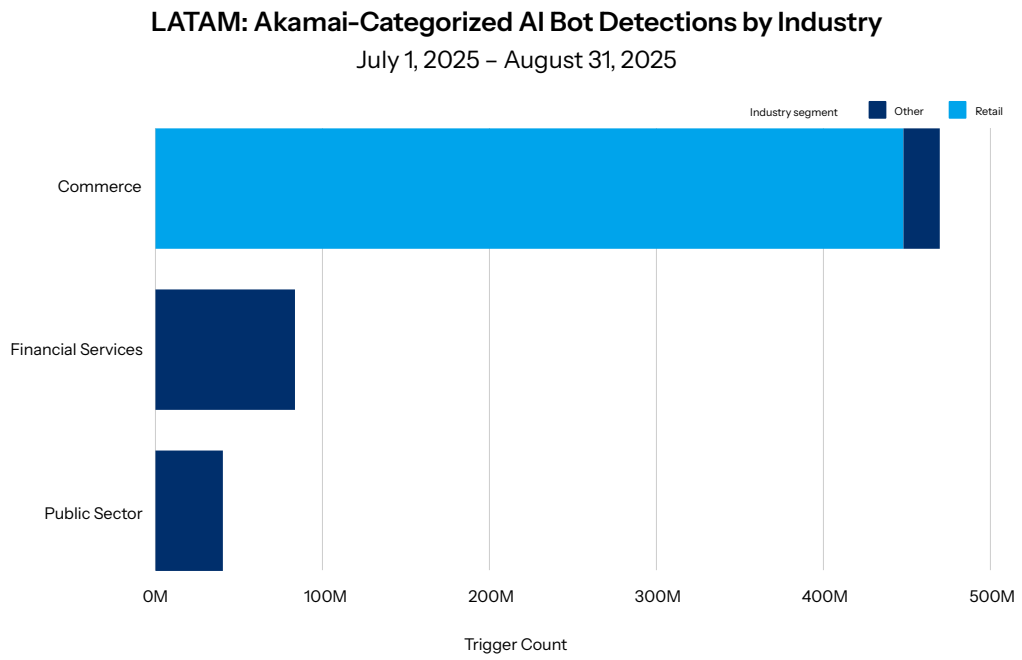

Fig. 19: Within APAC, EMEA, and LATAM, commerce was the industry with the most AI bot triggers, and the retail and publishing segments experienced the highest concentration of AI bot activity within their respective industries

In **APAC**, commerce experienced twice the amount of AI bot activity (4.23 billion triggers) than the next closest industry (other digital media, 2.1 billion), followed by high technology (1.4 billion), video media (668 million), financial services (619 million), and manufacturing (565 million).

In **EMEA**, commerce experienced more AI bot activity than all the other industries combined at 6.7 billion triggers, followed by other digital media (1.5 billion), manufacturing (1.1 billion), video media (878 million), high technology (759 million), and business services (684 million).

In **LATAM**, commerce topped the list of industries for AI bot activity at 468 million triggers, followed by financial services (83 million), and public sector (40 million).

Additional global industry trends were also reflected in the regions: specifically, the concentration of AI bot activity in the retail segment within commerce, and the concentration in publishing, a segment within the broader digital media landscape.

- APAC: Within commerce, the retail segment experienced 73% of AI bot triggers, while the publishing segment experienced the majority of triggers in other digital media at 61%.
- EMEA: Within commerce, the retail segment experienced 87% of AI bot triggers; within other digital media, publishing experienced 29% of triggers.
- LATAM: The retail segment experienced 95% of AI bot triggers within commerce.

As we discussed in greater detail in the Industry trends section, people spend a lot of their time online on retailers' websites, not just making purchases but also researching products, comparing prices, browsing, and exploring trends. As a result, that's where bots are focused — and our data reflects that. Also as previously discussed, when it comes to publishing, the value subscribers derive from content creators also draws AI bots. So, publishing tends to attract a higher concentration of AI-bot traffic than other segments within the other digital media industry.

A look at fraud and abuse through OWASP Top 10 lists

When you are training coders on what common threat vectors to avoid in their coding practices, it's important to review threat trends and think about what security controls you need to combat the most common attack methods. The nonprofit community-driven Open Web Application Security Project (OWASP) Top 10 lists are great tools to use for this.

Beginning in 2003 with a list focused on web application security risks, OWASP has since created an API Security Top 10 list (2016) and a Top 10 for Large Language Model Applications (2023; it's already on its second edition). This section will map where these OWASP frameworks help prevent fraud and abuse, allowing you to avoid the most common threat vectors.

First, we need to build a list of the most fraud-prone categories. These will not be the same for all companies, but here is a list of the ones that are most likely to be tied to the fraud and abuse attack methodologies commonly used by cybercriminal groups.

- Access control
- Authentication
- Injection (prompt, SQL, SSRF)
- Business logic abuse
- Configuration and plug-in integration
- Data disclosure and cryptographic failures
- Data theft of personally identifiable information (PII), protected health information (PHI), or other regulated data
- Denial of service / Resource exhaustion
- Tampering and integrity
- Output handling / Improper output

Then, we can build a map that shows which fraud-prone categories are addressed by the OWASP frameworks, so you can prioritize the proper security controls and avoid the most common threat vectors (Table 4).

Category	OWASP Web App 2021	OWASP API 2023	OWASP LLM 2025
Access Control	A01 — Broken Access Control	• API1 — Broken Object Level Authorization • API5 — Broken Function Level Authorization	———
Authentication	A07 — Identification and Authentication Failures	API2 — Broken Authentication	———
Injection (Prompt, SQL, SSRF)	• A03 — Injection • A10 — SSRF	API7 — SSRF	Prompt Injection
Business Logic Abuse	A04 — Insecure Design	API6 — Unrestricted Access to Sensitive Business Flows	• Excessive Agency • Misinformation
Configuration and Plug-in Integration	• A05 — Security Misconfiguration • A06 — Outdated Components	• API8 — Security Misconfiguration • API9 — Improper Inventory Management	Vector and Embedding Weaknesses
Data Disclosure and Cryptographic Failures	A02 — Cryptographic Failures	———	• Sensitive Info Disclosure • System Prompt Leakage
Data Theft of PII/PHI	• A01 • A02	— (Implicitly via Auth/ Authorization Failures)	• Sensitive Info Disclosure • Prompt Injection • Output Handling
Denial of Service / Resource Exhaustion	———	API4 — Unrestricted Resource Consumption	Unbounded Consumption
Tampering and Integrity	A08 — Integrity Failures	———	Data and Model Poisoning
Output Handling / Improper Output	———	API10 — Unsafe Consumption of APIs	Improper Output Handling

Table 4: The mapping of OWASP-related vulnerabilities shows the areas that are tied to fraud and abuse

Now that we have reviewed the common threat vectors to avoid in coding, we need to determine that we don't have open vulnerabilities. First, we need technical controls like web application firewalls, posture and runtime detection, and API and LLM protection capabilities. But we also need to prove — through tabletop exercises, pen testing, red team exercises, and security operations center techniques — that our processes will mitigate them. This requires picking some sample attacks to drive the scenario.

Through analysis of the fraud we are seeing across the Akamai platform, and via analysis of most common trends from reports like Internet Crime Complaint Center (IC3), we can see some of the most preventable types of technical support to fraud are:

- Credential theft and account takeover
- Payment redirect and business email compromise (BEC)
- Phishing (email, SMS, voice)
- Bot-driven credential stuffing and carding
- Fake ecommerce sites

Overall, each company must determine how important fraud reduction is to their overall security posture. There are some industries like commerce, retail, and finance that are more heavily impacted, so they are making loss-prevention strategy updates quarterly — but every industry should consider in what areas they are vulnerable as techniques like BEC are hitting all industries.

Those organizations that want to prioritize mitigating fraud from a technical perspective must also integrate cybersecurity controls and more classic fraud management tools. Furthermore, they need to build a single data source that allows for collaboration of the fraud and cybersecurity teams when conducting analysis and investigations that includes vulnerability data. We have provided the fraud-to-OWASP mapping to serve as a starting point for a discussion among those early in the process or as a validation exercise for companies with more mature processes in place.

Compliance

Balancing regulatory compliance with security efficacy in AI-enabled defenses

To paraphrase an old aphorism: "You can please all of the people some of the time, or some of the people all of the time, but you can't please all of the people all of the time." This statement highlights the inherent difficulty, if not impossibility, of universal and constant approval or contentment.

Legal compliance teams and cybersecurity teams may also feel this way about the interplay between compliance regulations and effective cybersecurity. We often lament that compliance with regulations — in privacy, for example — impacts our ability to deploy the most effective security tools and analytics. The recent increased pace of AI regulation represents the latest stage of this long-standing conflict.

Across all major jurisdictions, legislators are working to prescribe compliance measures that are designed to ensure that AI systems are safe, transparent, fair, and accountable. Legislators and regulators are rightly recognizing that the indiscriminate use of AI systems can seriously harm the rights and freedoms of individuals, and are diligently working to establish rules to address the perceived risks. The challenge of balancing the need for strong cyber defenses with the protection of fundamental rights, however, demands a nuanced and coordinated approach to the problem. Unfortunately, coordination across regulatory regimes is not always a priority.

The risks that regulators seek to mitigate are real and serious, yet overly restrictive controls on the training, deployment, and use of AI capabilities in security tooling — such as those intended to maximize data protection and minimize bias risk, for example — can paradoxically weaken an organization's resilience against increasingly automated, AI-driven fraud and abuse. Such overly restrictive controls can come from the regulations themselves, but also could be the result of aggressive compliance frameworks established by internal legal compliance teams. Threat actors already leverage LLMs to generate phishing kits, polymorphic malware, and evasive botnets. Compliance and security teams must coordinate to deploy equally sophisticated and compliant AI and machine-learning controls to keep pace.

Today's regulatory landscape

Regulators are increasingly adopting a risk-based sliding scale model to govern AI in financial services. High-impact use cases, including credit scoring and loan approvals, attract stricter oversight than lower-risk applications, such as digital marketing or operational automation. This approach strikes a balance between protecting the customers and the market and still encouraging innovation.

The United States does not have a single, comprehensive AI law. Instead, a patchwork of federal executive actions and plans (most notably America's AI Action Plan and related Executive Orders), the National Institute of Standards and Technology (NIST) AI Risk Management Framework, and sectoral statutes like the Gramm-Leach-Bliley Act (GLBA) and Artificial Intelligence (AI) at HHS from the Department of Health and Human Services collectively shape AI governance.

At the individual U.S. state level, there is some momentum with the passing of Colorado's AI Act (Senate Bill 24-205), which is designed to focus on consumer protection against "algorithmic discrimination" from "high-risk" AI systems like housing, employment, and healthcare, and is expected to become law in mid 2026. Importantly, there is no categorical restriction on the use of AI or machine learning for cybersecurity purposes. Instead, the focus is on ensuring that systems are trustworthy, secure, and do not introduce undue risk — principles that align with best practices in security operations.

The EU AI Act takes a more prescriptive, risk-tiered approach. AI systems used for the prevention, investigation, or detection of crime, including cybersecurity defenses, are classified as "high-risk." This triggers requirements for documented risk management, human oversight, transparency, and postmarket monitoring. Again, this act does not prohibit AI systems in this regard; it recognizes their societal value and seeks a compliance framework that can coexist with other regulatory regimes, such as the data protection mandates in the General Data Protection Regulation (GDPR).

APAC jurisdictions are converging on similar risk-based, accountability-focused paradigms. Singapore's AI Verify framework, Japan's AI governance guidelines, and Australia's Privacy Act reforms all promote explainability and accountability without banning security-focused AI/machine learning. These paradigms allow for proportionate controls, especially in situations where public-interest benefits outweigh residual privacy concerns.

So, how can organizations reconcile the potential conflicts across regulatory regimes and between regulations and security demands? Organizations should consider the following:

- ☑ Establish **effective and constant coordination** between legal compliance teams and cybersecurity teams. Each team must understand the fundamental problems the other team is trying to solve and they must work together to find balanced approaches.
- ☑ Implement a **tiered assessment program** that distinguishes among AI used for customer-facing decision-making or broad-based "answer anything" type systems (e.g., chatbots), AI used exclusively for threat detection and security-based decision-making, and employee use of AI platforms for content creation or marketing. The latter often carries lower overall risks and less direct consumer impact and can be fast-tracked with lighter governance.
- ☑ Require **model transparency documentation** (e.g., data-lineage maps, validation and testing reports, etc.) that satisfy EU AI Act and NIST documentation requirements while still protecting proprietary detection logic. From a compliance perspective, "It ain't compliant if you can't prove it!"

- ☑ Adopt **privacy-preserving telemetry**, such as hashing, tokenization, or differential privacy, whenever possible so security models can ingest high-fidelity signals without processing PII, which can reduce exposure.
- ☑ Establish a **continuous monitoring and improvement loop**. Align performance monitoring with regulatory mandates for postdeployment monitoring, ensuring that drift and false-positive spikes are promptly addressed and documented.
- ☑ Embed **cross-functional oversight**. Legal, security, and data-science teams should jointly review AI threat-detection tools, and work together to minimize regulatory drag on critical protections.

In sum, a nuanced, flexible, risk-based governance model, rather than reflexive restrictions and heavy-handed controls, will best equip organizations to satisfy emerging AI regulations while preserving the speed, scale, and precision required to defend against the next generation of automated attacks. At the end of the day, organizations need to build an AI protection program that is secure and can be mapped to the many different regulations they have to meet instead of focusing on meeting specific regulatory requirements. Regulatory compliance versus effective security need not be a zero-sum game.

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

Mitigation

AI bots are reshaping business operations by offering new efficiencies while introducing distinct risks. To harness the benefits of these technologies and minimize risk exposure, organizations must implement them with careful strategic planning, supported by rigorous governance and compliance frameworks. Effective risk mitigation requires a combination of technical controls, clear organizational policies, and ongoing monitoring.

Recommended best practices include:

- **Adopt a risk-based bot management approach**: Deploy solutions that provide granular visibility into all bot traffic. This enables early identification of attack patterns, suspicious behaviors, and emerging risks. Not all bots are malicious; some deliver legitimate business value. Detailed analytics support the differentiation between beneficial and harmful bots, allowing organizations to balance risk and revenue, mitigate threats, prevent fraud and abuse, and optimize the use of positive bot activity.
- **Monitor and respond to AI scraper activity**: Analyze traffic to identify how AI scrapers interact with web assets, including access patterns and data endpoints. Assess the intent and business impact of scraper activity to determine appropriate responses — from allowing limited access to heightened monitoring or full restriction.
- **Deploy AI-specific security controls**: Protect AI-driven applications in real time using specialized firewalls designed for LLMs and other AI systems. These controls address prompt-based attacks, harmful content generation, sensitive data leakage, data poisoning, and remote code execution. Unlike traditional firewalls, AI-focused solutions are tailored to the unique threat landscape facing AI workloads.
- **Leverage security guidelines and frameworks**: Use established frameworks, such as those provided by OWASP to identify and prioritize remediation of critical vulnerabilities. For example, the OWASP API Security Top 10 highlights common vulnerabilities, such as broken authentication and authorization flaws, which are frequently exploited in fraud and abuse scenarios.
- **Implement a comprehensive API security strategy**: Integrate security throughout the API lifecycle, beginning with design and continuing through postproduction, using a shift-left approach. Employ API discovery tools to achieve complete visibility across all APIs, including shadow and zombie endpoints that may be unmonitored and susceptible to attack.

By following these strategies, organizations can realize the advantages of AI-driven automation while maintaining robust security and compliance, reducing risk, and ensuring operational resilience.

Conclusion: The evolution of AI bots

This SOTI report has reviewed multiple forms of cyber fraud and abuse, but AI now stands out as the single most significant driver of change. AI is transforming both attack and defense, and it is accelerating tactics on both sides of the fraud landscape.

With tools like ChatGPT and FraudGPT, inexperienced threat actors can quickly launch basic bots, while advanced botnets still require technical expertise to scale criminal operations. AI now streamlines development and can automate routine bot behaviors, but its capabilities are limited to known techniques. It cannot yet invent new attack methods on its own — which means that defenders must stay agile and ready for surges in traffic and novel attack patterns.

Both adversaries and defenders are adopting AI to operate faster and more efficiently. Even as defenses improve, human expertise remains necessary to retool bot infrastructure and adapt to new controls. This echoes previous shifts — such as the rise of APIs and automation — when security often lagged behind deployment. With AI, it is critical for security leaders to ensure infosec is integrated at the earliest stages.

Although future AI-managed bots may eventually run campaigns and adapt autonomously, we are not seeing this extensively today. The current reality persists: Cybersecurity is still a contest between skilled humans. AI is a force multiplier, but expertise and oversight are essential to maintain an advantage.

Blocking AI bots outright is unlikely to succeed. Instead, organizations should focus on robust management and mitigation strategies. As AI-driven bots become core to many industries, CISOs must develop frameworks for secure AI adoption, risk management, and resilience — and ensure that digital assets and operations remain protected as the threat landscape evolves.

Methodology

This data describes application-layer traffic seen through our bot management tools. Our rules are triggered when we detect bot telemetry from an Akamai-categorized AI bot from a request to a managed website, application, or API. This tool helps detect, manage, and mitigate bot traffic according to customer expectations.

This research dataset covered the two-month period from July 1, 2025, through August 31, 2025.

Guest contributors

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

John "JD" Denning is the Chief Information 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 FS-ISAC, JD spend 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 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.

David Sénécal
Director of Engineering, Akamai

David Sénécal works for Akamai as a Director of Engineering — Fraud and Abuse and is the author of The Reign of Botnets. He is passionate about improving internet safety and is an expert in online fraud and bot detection. David has more than 25 years of experience working with web performance, security, and enterprise networking technologies across various roles, including support, integration, consulting, development, product management, architecture, and research.

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.

Credits

Research director

Kimberly Gomez

Editorial and writing

Charlotte Pelliccia
Lance Rhodes
Badette Tribbey
Maria Vlasak

Review and subject matter contribution

James A. Casey
Tom Emmons
Reuben Koh
Rob Lester
Emily Lyons
Richard Meeus
Juan Carlos Rivera
Christine Ross
David Sénécal
Rubens Waberski
Steve Winterfeld

Data analysis

Rob Lester
Chelsea Tuttle

Promotional materials

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.com/soti

Akamai threat research

Stay updated with the latest threat intelligence analyses, security reports, and cybersecurity research. akamai.com/security-research

Akamai security research

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

Access data from this report

View high-quality versions of the graphs and charts referenced 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.com/soti-data/ai-botnet-report-2025

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 11/25.