

2026 EDITION

THALES

CYBERSECURITY

BAD BOT REPORT

Bad Bots in
the Agentic Age

cpl.thalesgroup.com

#2026BadBotReport

Table of Contents

How to Read This Report	03
Executive Summary	04
Key Findings in Numbers	05
Bot Traffic Analysis	07
The Most Common Attack Types	09
Agentic AI & Malicious Bots – An Infrastructure Risk	12
Bad Bots & the API Attack Surface	19
Account Takeover Attacks	21
Bot Evasion Tactics	23
Bots Impersonating Browsers	24
Bad Bots by Industry	25
Top Targeted Industry by Bad Bots	25
Bot Traffic by Industry – Bad Bot vs Good Bot vs Human	26
Attack Sophistication by Industry	27
Top Targeted Industry by Account Takeover Attacks	28
Retail & Travel Lead in Targeted Business Logic Abuse	29
A Real-Life Bot Mitigation in the Financial Services Industry	31
A Real-Life Bot Mitigation in the Insurance Industry	32
Recommendations	33
Conclusion/Looking Ahead	38
About Thales	38
Definitions	39

How to Read This Report

The 13th annual Bad Bot Report looks back at full-year 2025 bot activity to explain how automation, now galvanized by artificial intelligence (AI), is reshaping digital infrastructure today, and what organizations should prepare for next. It combines data and insights from the Thales Threat Research and Security Analyst Services (SAS) teams, who investigate and mitigate bot attacks across industries every day.

Rather than focusing on isolated attack techniques, this report examines how bots interact with applications and APIs, at scale, with additional insights this year into the normalization of AI in internet traffic.

The findings are intended to help security leaders, risk teams, and business stakeholders understand not just what bots are doing, but how automation is impacting application availability, site performance, operational costs, business metrics, and trust, across modern digital services.



Executive Summary


AI-driven automation is no longer emerging; it is becoming foundational to internet traffic. Over the past 12 months, AI-driven bot activity has increased more than tenfold, with daily blocked requests rising from 2 million to 25 million¹ in 2025. While this growth is significant, the more fundamental shift is the normalization of AI-driven automation within internet infrastructure.

AI agents are emerging as a third category of automated traffic, alongside traditional good bots and bad bots. These agents interact directly with applications and APIs to retrieve data and perform tasks on behalf of users. As a result, automated activity that would previously have appeared anomalous is increasingly treated as expected behavior.

This shift is fundamentally changing how organizations must interpret traffic. The distinction between legitimate and malicious automation is becoming increasingly blurred, as both now operate through similar channels, workflows, and infrastructure.

At the same time, visibility into AI-driven activity remains limited. Current analysis is based on detectable traffic, meaning systems that identify themselves or trigger security controls. A much larger portion of AI-driven automation remains unverified, creating a growing gap between what organizations can observe and the true scale of AI-enabled risk.

Automation continues to dominate internet traffic. In 2025 bots accounted for over 53 percent of all traffic, reinforcing a structural shift in how digital services are accessed and consumed. Bot activity is no longer irregular or event-driven. It is a persistent and expected component of modern applications.



In 2025, 21 percent of all mitigated attacks aligned with OWASP automated threat categories, highlighting the scale of bot-driven abuse. APIs remain a primary target, with 27 percent of bot attacks directed at API endpoints, where attackers exploit business logic, access sensitive data, and interact directly with core application functionality.

The impact is most visible in high-value sectors. Financial Services was the most targeted industry, accounting for 24 percent of all bot attacks and 46 percent of account takeover incidents, underscoring the direct monetization potential of automated threats.

As AI-driven automation becomes embedded in digital infrastructure, the challenge is no longer simply identifying bots. Organizations must now distinguish between automation that enables business operations and automation that exploits them.

¹AI bot coverage expanded in 2025

Key Findings in Numbers



**17.2
TRILLION**

The number of bot requests blocked by Thales in 2025


21% Bot attacks targeting business logic



40% The percentage of internet traffic made up of bad bots

12.5x The year-over-year increase in AI-enabled bot attacks

20%



The percentage of AI bot attacks targeting Retail sites

53%

The percentage of internet traffic made up of bots

Key Findings in Numbers



42%


Simple bot attacks

41%

the percentage of bot attacks using Chrome to appear as legitimate traffic

24%

The percentage of bot attacks targeting Financial Services sites



Percentage of internet traffic attributed to human traffic



27%

Bot attacks targeting APIs



46%

The percentage of account takeover attacks targeting Financial Services


Bot Traffic Analysis

Bad Bot vs Good Bot vs Human Traffic in 2025

Automated traffic accounted for 53 percent of all observed web traffic in 2025. Of that, 40 percent was generated by bad bots, up from 37 percent the year before. The other 13 percent was benign automation, which includes harmless traffic such as search engines and AI crawlers. Human traffic, accounting for 47 percent of traffic, now represents a shrinking share of overall activity. These figures reflect more than a volume trend. They indicate a structural change in how digital services are consumed.

Bot traffic outpaces human traffic for the second year


BAD BOT VS GOOD BOT VS HUMAN TRAFFIC IN 2025



Bot Traffic Over The Years


Bot traffic patterns over the last ten years show a consistent increase in the percentage of bot traffic and bad bot traffic, which is up to 40 percent in 2025, a 3-percentage point increase from 2024 (37 percent). Human traffic to the internet is on a steady decline year on year.

GLOBAL INTERNET TRAFFIC PROFILE BETWEEN 2015 AND 2025



The Most Common Attack Types

This chart highlights the most common attack types targeting websites, applications, and APIs, showing how automation has become a major driver of modern cyber threats. The largest category is general automation which could include brute force automated attacks, or vulnerability scanning, where bots are used to probe systems, test credentials, scrape data, or repeatedly exploit weaknesses at scale, accounting for 29 percent of activity.



The data also shows that business logic abuse (21 percent) is a major threat. In these attacks, bots manipulate legitimate application workflows.

API violations (24 percent) are another significant category, including issues such as broken authorization or mass assignment that allow attackers to access data or functionality they shouldn't. API violations are often [OWASP API Top Ten](#) threats including: Broken Object Level Authorization (BOLA) or Unrestricted Access to Sensitive Business Flows.

Attack Sophistication

In 2025, advanced bot attacks (44 percent) and moderate bot attacks (14 percent) accounted for a combined total of 58 percent of all bot attacks, a 2-percentage point increase on the previous year. Bad bots became more adaptive and persistent, leveraging AI to mutate fingerprints, adjust interaction timing, and pivot rapidly when mitigations were applied. Security Analyst teams frequently observed campaigns that returned repeatedly, probing different endpoints and workflows until a viable path was found.


Simple bot attacks also prevailed in 2025, accounting for 42 percent of all attacks (dropping from 44 percent in 2024). However, the volume of simple bot attacks increased by more than 230% compared to 2024, driven by AI lowering the barrier to entry and enabling attackers with limited expertise to deploy automation at scale. While individually unsophisticated, simple bot attacks contributed to sustained infrastructure load and operational noise.

The interaction between these two types of attack produced a compounding effect. High-volume simple bots created constant background pressure, while advanced and moderate bots, targeted high-value workflows such as authentication, booking, checkout, and payments. Rather than short attack spikes, many organizations experienced continuous bot presence, effectively sharing their infrastructure with automated agents over extended periods.

Top Targeted Countries by Bad Bots


The United States was the most targeted country for bot attacks in 2025 targeted by 59 percent of all attacks, followed by Australia (12 percent), the United Kingdom (9 percent) and France (3 percent).

TOP TARGETED COUNTRIES



OWASP Attack Categories

In the past year, 21 percent of all attacks recorded and mitigated were automated threats, as defined by the OWASP. **The OWASP 21 Automated Threats** are a set of automated cyberattacks that leverage bots and scripts to exploit web application vulnerabilities at scale, bypass security controls, and disrupt businesses, and represent some of the most common and impactful forms of automated attack activity observed in modern application environments.



Agentic AI & Malicious Bots


An Infrastructure Risk

Artificial Intelligence is rapidly reshaping automated traffic on the internet. In 2025, the average number of AI-driven bot attacks mitigated increased more than tenfold (12.5x) compared to the previous year bringing the daily average of attacks blocked to 25 million. While this rise in AI-powered attacks is significant, the larger shift in 2025 was the normalization of AI and automation within internet infrastructure itself.

AI-driven attacks were observed across a wide range of industries and geographies, highlighting the global scale and reach of AI-enabled automation.


AI blurs the lines between legitimate and malicious automation

TOP 10 MOST TARGETED COUNTRIES BY AI BOTS



The United States was the most targeted country by AI bots.

TOP 10 INDUSTRIES TARGETED BY AI BOTS



Retail was the top targeted industry by AI bots.

A New Category of Automated Traffic

Traditionally, automated traffic has been grouped into two categories: good bots, such as search engine crawlers, and bad bots, including scrapers, credential stuffing tools, and scalpers.

AI introduces a third category: AI agents.


These agents access websites, retrieve data, and perform tasks on behalf of users. They are embedded in browsers, search platforms, and enterprise tools, interacting directly with applications and APIs at scale. As a result, automated activity that would previously have appeared undesired is increasingly treated as expected behavior.

Detectable AI Traffic is only the tip of the iceberg

The AI traffic analyzed in this report represents only detectable or declared AI clients. A much larger portion of AI-driven automation remains unverified.

Attackers can deploy self-hosted or modified large language models that do not identify themselves as AI agents and can be fine-tuned for malicious use. This creates a visibility gap between what organizations can detect and the true scale of AI-enabled activity.


What is observable today represents only a fraction of the total attack surface.



AI Bot Composition and Emerging Risk Patterns

Analysis of detectable AI traffic in 2025 shows two primary categories:

- AI crawlers (85 percent) gather data to train models
- AI fetchers (15 percent) perform actions in response to user prompts



More than 10% of AI fetcher sessions and nearly 9% of AI crawler sessions triggered bad bot detection rules, underscoring that AI-driven traffic is already crossing into behaviors typically associated with malicious automation.

RULE TYPE	AI CRAWLERS	AI FETCHERS
Bad Bots	8.8%	10.8%
Customer-defined rules	11.9%	7%
DDoS	4%	0.53%
SQLi	0.003%	0.003%
XSS, IRA	0.011%	0.005%
Backdoor & RFI	0.001%	0.1%

11.9% of crawlers and 7% of fetchers were blocked by customer-defined rules. This indicates that many administrators explicitly do not want AI bots crawling their sites, or they prefer to maintain strict, tailored control over their incoming traffic.


The fact that 4% of AI crawler sessions triggered DDoS rules further highlights the risk of AI-driven activity scaling into disruptive, attack-like traffic patterns.



AI Crawler Bot Traffic (sample view)

This chart represents the breakdown of AI crawler traffic by client over a seven-day period, showing which AI agents are most actively accessing applications. The distribution highlights a concentration of activity among a small number of dominant providers, reinforcing the growing influence of large AI platforms in automated web traffic.

AI CRAWLER BOT TRAFFIC BY SOURCE




AI Fetcher Bot Traffic (Sample View)

This chart represents the breakdown of AI fetcher traffic over a seven-day period, showing which AI agents are most actively retrieving content from applications. The distribution highlights a strong concentration of activity from a single dominant provider, with most traffic associated with ChatGPT user agents, reinforcing the growing role of AI-driven content retrieval in automated web traffic.



AI FETCHER BOT TRAFFIC BY SOURCE



Early Organizational Response to AI Agents

Organizations are beginning to define how AI agents interact with their systems. Almost 12% of AI crawler traffic is currently blocked using customer-defined rules, indicating early policy decisions around which AI tools are permitted access.

In practice, access is often determined by the sensitivity of the application. Public content may remain accessible, while authentication, transactional, and data-rich endpoints are more tightly controlled.

This reflects a broader shift: AI is both a business enabler and a source of risk that must be actively managed.

AI Is Reshaping the Threat Landscape

AI is not creating entirely new attack categories. Instead, it is accelerating and refining existing forms of automation.

Insights from Security Analyst Services teams show that AI is enhancing reconnaissance and data gathering, CAPTCHA solving and interaction timing, fingerprint evasion, and persistence with rapid adaptation to mitigation.

The result is more efficient, scalable, and resilient bot activity that is harder to detect using traditional methods.

The Blurring of Intent in Automated Traffic

As AI agents, crawlers, and automated assistants become standard components of web traffic, the distinction between legitimate and malicious automation is becoming less clear.

Modern bots increasingly use valid browser environments, follow expected workflows, and generate well-formed requests. As a result, intent can no longer be inferred from surface-level signals alone. The challenge is shifting from identifying bots to understanding how automation is interacting with business logic and infrastructure.

LLMs Now Warrant Their Own Security Framework

The rise of AI-driven systems has introduced new attack surfaces that are not fully addressed by traditional application security models.

The [OWASP Top 10 for Large Language Model applications](#) reflects this shift, highlighting risks such as prompt injection, retrieval poisoning, model manipulation, and supply chain attacks targeting training data. Attackers are also increasingly using LLMs to generate highly tailored attack payloads.

These threats extend beyond traditional bot and application security telemetry, reinforcing that AI is now a distinct and growing attack surface.

Monetization and the Formalization of AI Traffic

The normalization of AI agents is also driving new operational models, including the verification and potential monetization of AI-generated traffic.

AI-driven traffic is a potential new revenue channel

Verified AI bots use cryptographically signed headers, allowing organizations to authenticate and measure AI-driven access. This introduces a new layer of control, enabling organizations to distinguish between approved AI agents and unverified automation, and to define how and when these agents can interact with their applications and APIs.

As adoption matures, this is expected to evolve into enforceable access models, where AI traffic is not only identified but governed through policy, rate limiting, or commercial agreements, effectively turning AI-driven access into a managed and potentially billable channel.

Bad Bots & the API Attack Surface

Bad bots continue to target APIs, with 27 percent of bot attacks directed at API endpoints in 2025. As modern digital services increasingly rely on APIs to power core functionality, APIs have become a critical point of exposure for automated threats.


Based on bot mitigation activity in 2025, Security Analyst Services teams consistently observed malicious campaigns designed as API-first. Bots bypass user interfaces entirely, interacting directly with backend services using well-formed, authenticated requests. This allows attackers to operate at machine speed while avoiding many traditional web-layer controls.

If your APIs aren't protected your business logic is exposed

What makes API abuse particularly challenging is that it often does not appear malicious. Requests are valid, authentication frequently succeeds, and workflows are followed correctly. The impact comes from scale, persistence, and intent rather than malformed traffic.

Analysis of API attack traffic shows that the most common threats targeting APIs include data leakage (26 percent), business logic abuse (13 percent), and remote code execution or remote file inclusion attacks (13 percent).

MOST COMMON THREATS TO APIS



Data leakage occurs when APIs expose sensitive information such as customer data or internal records due to weak access controls. Business logic attacks exploit how an application is designed to function, allowing bots or attackers to manipulate workflows, bypass limits, abuse promotions, or hoard inventory. Technical attacks such as remote code execution or remote file inclusion exploit software vulnerabilities to run malicious code or load unauthorized files.

AI Agents Accelerate API Interactions at Scale

The rapid growth of AI tools and automated agents is further accelerating interactions with APIs because they (AI agents) rely heavily on APIs to retrieve data, a trend which is raising new concerns for developers and security teams.

According to the [2025 State of the API Report by Postman](#), 51 percent of developers worry about unauthorized or excessive API calls from AI agents, making it their number one API security concern.




Account Takeover Attacks

Account Takeover (ATO) remained one of the most persistent and damaging forms of automated abuse in 2025. Despite widespread adoption of controls, such as multi-factor authentication (MFA), bots continue to exploit credential reuse and API-driven identity workflows. The chart below shows a steady increase in ATO attacks since 2024, with a surge in attacks between May and July 2025, which could be attributed to increased usage of AI tools and agents. Comparing July 2024 to July 2025 account takeover attacks increased by 70 percent.

*Account Takeover is the fastest
path to automated fraud*

GROWTH OF ACCOUNT TAKEOVER ATTACKS 2024 TO 2025



Consequences of Account Takeover

A data breach resulting from an account takeover attack can lead to regulatory penalties, reputational damage and long-term financial losses. This list shows some of the regulations that could be triggered by a successful account takeover attack.

REGULATION	WHY IT MATTERS	PENALTIES
GDPR General Data Protection Regulation	Exposure of personal data is a reportable breach under GDPR.	Fines up to €20 million or 4 percent of global annual turnover for failure to protect personal data
CCPA California Consumer Privacy Act	Exposure of personal information can trigger fines and private lawsuits under CCPA.	Fines up to \$2,500 per violation or \$7,500 for intentional violations
HIPAA Health Insurance Portability & Accountability Act	Account takeover attacks involving protected health information are treated as serious security incidents.	Fines ranging from \$100 to \$50,000 per violation, with a maximum annual penalty of \$1.5 million
DORA (EU) Digital Operational Resilience Act	Failure to protect authentication systems and APIs may be viewed as a failure of digital operational resilience.	Significant administrative fines and remediation plans for failure to manage ICT risk appropriately.
PSD2 (EU) Payment Services Directive 2	Account takeover due to weak authentication increases SCA liability.	Financial sanctions, and restrictions for non-compliance with Strong Customer Authentication requirements.
NIS2 (EU) Network and Information Security Directive 2	Account takeover attacks that compromise systems or disrupt digital services can trigger NIS2 reporting and enforcement.	Fines up to €10 million or 2 percent of global annual turnover (whichever is higher)

Bot Evasion Tactics

Evasion in the Age of AI-Driven Automation

Based on insights from Thales Security Analyst Services and Threat Research, the following evasion tactics were most prevalent in 2025:

- **AI-Assisted Evasion and Learned Adaptation.** AI is used to analyze application workflows, refine attack logic, and rapidly adjust behavior when blocked. Analysts observed bots adapting within hours of mitigation deployment. AI-driven iteration enables attackers to test defensive responses and refine tactics at speed.
- **CAPTCHA Solving at Scale.** AI-assisted solving has weakened traditional CAPTCHA controls, but human CAPTCHA farms have not disappeared. Many services now blend AI with incentivized human solvers, enabling bots to bypass challenges at scale while increasing friction for legitimate users.
- **Privacy Tools.** Capabilities such as iCloud Private Relay obscure user identity, making it more difficult to distinguish between genuine user activity and automated bot traffic.
- **Adaptive Polymorphic Behavior.** Modern bots continuously modify timing, request sequencing, and infrastructure. Rather than relying on a fixed approach, they adapt dynamically in response to mitigation controls, making persistence a defining characteristic of automated attacks.
- **Advanced Fingerprinting and Identity Evasion.** Bots increasingly present valid, consistent browser identities that closely mirror real users. Analysts observed dynamic fingerprint manipulation across headers, device attributes, and execution patterns. In more advanced cases, bots leveraged token reuse, as well as cookie and session ID reuse from legitimate sessions to bypass controls.
- **API-First Attack Execution.** Bots increasingly bypass front-end interfaces and interact directly with authentication, search, booking, and payment APIs. These requests are often well-formed and authenticated, allowing attackers to operate at machine speed while avoiding traditional web-layer controls.
- **Infrastructure Masking Through Proxies.** The use of residential and mobile proxy networks remains highly effective. By routing traffic through legitimate user devices and ISP networks, bots blend into normal geographic and behavioral patterns, reducing the effectiveness of IP-based detection.
- **Headless and Automation Frameworks.** Tools such as Puppeteer, Selenium and Playwright remain widely used, now enhanced with AI-generated scripting and real-time behavioral tuning. These tools enable bots to execute JavaScript, maintain session state, and replicate multi-step user journeys.
- **Bots as a Service Platforms.** The commercialization of bot capabilities continues to expand. Service-based platforms offering scraping and scalping at scale, often enhanced with AI-driven refinement
- **Multi-Threaded and Persistent Attack Models.** Bots increasingly operating through parallel threads or fallback mechanisms. When one attack vector was mitigated, activity continues through lower noise channels making persistence a defining tactic.

Bots Impersonating Browsers



To avoid detection, bad bots frequently disguise themselves as legitimate web or mobile browsers commonly used by real users to bypass basic security controls.

Chrome continues to be the browser most impersonated by attackers. In 2025, 41 percent of bad bot traffic declared itself as Chrome, up slightly from 39 percent in 2024.

Android Browser remained the second most impersonated browser at 17 percent, reflecting the continued importance of mobile traffic as a disguise for automated attacks.

BOTS IMPERSONATING BROWSERS 2025 VS 2024

BROWSER	2024	2025
Chrome	39%	41%
Android Browser	17%	17%
Firefox	11%	10%
Internet Explorer	12%	9%
Mobile Safari	7%	9%
Safari	5%	6%
Mobile Chrome	4%	4%
Microsoft Edge	4%	3%
Opera	1%	1%
Samsung Browser	0.2%	0.14%
UCBrowser	0.18%	0.09%



When bots are indistinguishable from human traffic, traditional detection breaks down


Bad Bots by Industry

Bot activity varies significantly across industries to reveal distinct patterns in attack volume, composition and sophistication, providing insight into where automated threats are most prevalent and impactful.

Top Targeted Industry by Bad Bots

In 2025 the top targeted industry was Financial Services accounting for 24 percent of all bot attacks.

TOP TARGETED BY BAD BOTS




*Bots interact directly with the same APIs
that power core banking operations*

The sector remained one of the most consistently targeted for bot-driven abuse in 2025. While this is not new, what has changed is how deeply automated traffic is now embedded within financial infrastructure. Bots are no longer testing the edges of financial systems. They are interacting directly with the same APIs, identity services, and workflows that power core customer transactions and digital banking operations.

Bot Traffic by Industry – Bad Bot vs Good Bot vs Human

While Financial Services remained the most targeted industry by overall bot attack volume, analysis of traffic composition, comparing bad bots to good bots and human traffic, reveals a different picture. Telecoms, Society (Non-Profit), Computing & IT, Travel, and Business sectors rank among the top five industries with the highest proportion of bad bot traffic.

BOT TRAFFIC BY INDUSTRY




Attack Sophistication by Industry

Advanced bot activity increased across multiple industries in 2025. In Law & Government, 71 percent of attacks were classified as advanced, up from 68 percent in 2024. Gaming saw the most significant shift, with advanced attacks nearly doubling to 50 percent (up from 27 percent).

Education also experienced a sharp rise, with advanced attacks increasing from 3 percent to 12 percent. In Healthcare, 42 percent of attacks were classified as advanced or moderate, up from 16 percent, reflecting a substantial increase in sophisticated bot activity.

ATTACK SOPHISTICATION BY INDUSTRY




Top Targeted Industry by Account Takeover Attacks

Financial Services was the primary target for account takeover attacks, accounting for 46% of all incidents. This reflects the high value of financial accounts and the direct monetization opportunities they present to attackers. Authentication systems, login APIs, and identity workflows remain a key point of exposure, making the sector particularly vulnerable to credential stuffing and automated fraud.

Bots go where the money is – almost half of all Account Takeover attacks target financial services


TOP TARGETED INDUSTRIES BY ACCOUNT TAKEOVER ATTACKS



Retail & Travel Lead in Targeted Business Logic Abuse

Analysis of business logic attacks in isolation reveals a different pattern. Retail and Travel emerge as the most targeted industries, reflecting a higher concentration of targeted, fraud-driven activity such as loyalty abuse, transaction manipulation, and exploitation of application workflows.

TOP TARGETED BY BUSINESS LOGIC ABUSE



With AI services in Retail and Travel, attackers can masquerade as legitimate AI agents

Travel & Airlines

In 2025, the Travel and Airline sectors faced sustained bot activity targeting APIs that power search, pricing, availability, and booking. Analysts observed high-frequency queries to fare, seat, and route APIs, often far exceeding normal customer behavior.

Fetch-style automation continuously polled systems during peak periods, masquerading as legitimate agents to inflate look-to-book ratios and distort demand signals used for pricing and capacity planning. Seat spinning and denial-of-inventory attacks created artificial scarcity, reducing availability for legitimate customers and impacting both revenue and customer experience.

Retail

Retail organizations faced similar automation pressures, particularly those with dynamic pricing, limited inventory, or high-demand promotions. Bots frequently interacted with retail APIs to monitor availability, pricing, and promotions at high frequency, often using fetch-style automation to poll product endpoints at machine speed. Analysts also observed inventory denial tactics, where bots add items to carts without completing purchases, creating artificial scarcity during major launches or sales.



A Real-Life Bot Mitigation in the Financial Services Industry

The Challenge: Account Takeover (ATO)

A Financial Services organization experienced a sudden spike in automated login traffic targeting its customer authentication systems. As this occurred before they had a bot protection solution configured, the automated traffic was initially able to pass through without triggering mitigation controls.

Detection

The bank's security team detected the activity after noticing an abnormal surge in login traffic interacting with authentication APIs. Monitoring systems flagged the unusual volume of requests, prompting the team to investigate and quickly escalate the issue to the Security Analyst Services (SAS) team for support.

Business Impact

Account takeover attacks present significant risks for financial institutions. Successful attacks can lead to fraudulent transactions, financial losses, regulatory exposure, and damage to customer trust. Even short-lived attacks can undermine customer confidence and create operational disruption.

Mitigation

Working with SAS, the organization rapidly configured Account Takeover Protection (ATO), implemented custom rules within Advanced Bot Protection (ABP), and enabled out-of-the-box mitigation capabilities. These controls helped block malicious login attempts and strengthen defenses against future automated attacks.

While ABP focused on detecting and mitigating automated traffic at scale, ATO provided behavioral analysis of individual users and the broader customer base, identifying deviations from expected login patterns that indicate account takeover attempts.

Key Takeaways

- Implement both Account Takeover Protection and Advanced Bot Protection to secure authentication workflows.
- ABP detects and mitigates automated traffic, while ATO uses behavioral analysis to identify deviations in user activity and uncover account takeover attempts.
- Ensure baseline mitigation controls are enabled during deployment.
- Monitor login traffic closely and respond quickly to sudden spikes in activity.
- Engage security experts early to accelerate mitigation and reduce exposure.

A Real-Life Bot Mitigation in the Insurance Industry

Challenge: SMS Pumping

A global health insurance provider experienced a sustained automated attack targeting its SMS-based authentication system. The platform uses SMS One-Time Passwords (OTPs) to secure services such as account registration and login. Attackers exploited this process through an SMS pumping attack, where bots repeatedly triggered OTP requests via authentication APIs. The automated traffic generated large volumes of SMS messages sent to accounts created with temporary or invalid email addresses, driving up messaging costs.

Detection

The attack was identified after analysts noticed abnormal messaging patterns across the authentication infrastructure including unusual spikes in outbound SMS traffic and messages being sent to regions that did not align with the organization's typical customer base. Further investigation revealed high-frequency automated requests targeting OTP generation APIs, confirming that bots were systematically triggering authentication workflows.

Business Impact

Although no customer data was compromised, the attack generated significant financial impact. Because each OTP message incurs a telecom cost, the automated requests resulted in approximately \$300,000 in SMS charges over the 20-day period.

Mitigation

Following the incident, the organization implemented multiple controls to reduce exposure. These included blocking registrations from disposable email services, introducing strict rate limits on OTP requests, deploying advanced bot detection models, and applying additional security controls across authentication APIs.

Key Takeaways

- Block disposable or temporary email registrations during account creation.
- Apply rate limiting to OTP generation endpoints and messaging services.
- Monitor authentication APIs for abnormal traffic patterns.
- Integrate bot protection controls directly into authentication workflows.

Recommendations

Defending Digital Infrastructure Against AI-Driven Automation

Bot mitigation has entered a new phase. Automation is now embedded in everyday internet traffic and is often difficult to distinguish from legitimate activity. AI is accelerating this shift, enabling bots to adapt, persist, and blend into normal application behavior.

The following recommendations are based on 2025 data and real-world mitigation experience from Thales Product Management, Threat Research and Security Analyst Services teams.

1. Implement an AI Strategy for Protection and Optimization

Organizations must define a clear strategy for managing AI-driven traffic, including which AI agents are permitted to access their applications and which should be restricted. Allowlisting trusted AI agents while applying controls to unverified or high-risk automation is critical to maintaining visibility and protecting sensitive workflows.

Not all AI traffic should be treated equally. Organizations must determine where AI access is appropriate, particularly across APIs, authentication systems, and business-critical pathways. Without clear governance, AI-driven automation can introduce risk at scale while appearing legitimate.

This requires proactive action. Organizations should begin these discussions now, rather than waiting for an incident to expose gaps. Security and application teams must actively review available visibility tools and dashboards to understand how AI agents are interacting with their applications, enabling informed, risk-based decisions.

Equally important is cross-functional alignment. Armed with this visibility and control framework, security teams should engage business stakeholders to help define acceptable use, balance innovation with risk, and lead strategic decisions on how AI-driven automation is governed across the organization.

2. Design Defenses for Bots That Learn, Not Bots That React

One of the most important shifts observed in 2025 is that bots are no longer static tools responding predictably to controls. Analysis from Threat Research and Security Analyst Services teams shows that bots learned application workflows, analyzed mitigation responses, mutated fingerprints, and returned with refined behavior. This learning capability, driven by AI and agentic automation, fundamentally alters the defensive equation.

Organizations must therefore design defenses with the assumption that bots will study and adapt to them. Static rules, fixed thresholds,

and one-time deployments will increasingly fail as bots iterate faster than defenders can respond manually. Effective bot mitigation in 2026 requires adaptive controls that evolve over time, detect behavioral drift, and respond dynamically as attacker behavior changes. This recommendation represents the single biggest departure from last year's guidance and reflects how AI has changed the nature of automated threats.

Advanced bot protection solutions, such as Imperva Advanced Bot Protection, provide persistent bot identification capabilities that enable tracking of automated activity across sessions, even as attackers rotate IPs, devices, or fingerprints.

3. Shift from Defense to Governance

The rise of legitimate AI bots, crawlers, and fetch agents introduces a new category of risk. While not inherently malicious, this unmanaged automation can place sustained load on systems, distort analytics, and expose sensitive business logic at scale.

In this environment, organizations that focus solely on blocking attacks will struggle. Effective defense requires a shift toward governance, combining visibility across applications and APIs, adaptive and intent-aware controls, and human expertise to continuously evolve defenses.

Organizations must establish clear policies and technical controls that distinguish

acceptable automation from abuse. Allowlisting without ongoing monitoring is no longer viable. Even benign automation requires governance when interacting continuously with critical systems. Striking the right balance between openness and control will be essential as AI-driven automation becomes the norm.

Imperva Advanced Bot Protection offers granular visibility into AI tools, agents, and crawlers to provide a detailed, real-time view into which AI tools are accessing your websites, applications, and API endpoints.

4. Treat APIs as Critical Digital Infrastructure

APIs emerged as the dominant attack surface in 2025, reflecting their central role in modern digital services. Bots increasingly bypass user interfaces entirely and interact directly with backend APIs that handle authentication, pricing, booking, payments, and data access. Analysts overwhelmingly identified APIs as the most critical and consistently targeted surface.

Organizations should treat APIs with the same priority and rigor as public-facing applications. This includes continuous API discovery, strong authentication and authorization, behavior-based monitoring, and controls designed for sustained machine-speed interaction. As more organizations adopt API-first architectures, failure to protect APIs effectively becomes a systemic business risk rather than a technical gap.

5. Secure Identity Systems as Core Infrastructure

Account takeover continued to be one of the most damaging and widespread forms of automated abuse in 2025, particularly in Financial Services, Retail, and Travel. Insights from security analysts highlighted the growing sophistication of credential-based attacks, driven by AI-enabled automation that closely mimics legitimate login behavior.

Authentication and session management systems should be treated as critical digital infrastructure rather than isolated security controls. This requires layered protection, behavioral monitoring, and close coordination between security, fraud, and operations teams. Multi-factor authentication remains essential, but it must be complemented by controls designed to detect automated abuse of identity workflows rather than relying on MFA alone.

6. Protect Business Logic as a Primary Security Objective

In 2025, the most damaging bot attacks did not exploit software vulnerabilities. They exploited how applications are designed to function. Login workflows, search features, booking paths, checkout processes, and loyalty systems were repeatedly abused through valid requests executed at invalid scale or frequency.

Defenders must therefore shift focus from protecting individual pages or endpoints to

protecting business logic itself. This requires understanding how workflows are intended to operate, what normal usage looks like, and how automation can undermine outcomes such as availability, pricing integrity, analytics accuracy, and customer trust. Analysts consistently observed that without this business context, bot activity often appeared legitimate until operational or financial impact became unavoidable.

Imperva Advanced Bot Protection analyzes request payloads and transaction-level behavior to detect abuse of application workflows, including manipulation of parameters, scraping, and automation targeting APIs and critical business processes.

7. Assume Bots Are Indistinguishable from Humans

A recurring theme across analyst observations was how difficult bots have become to identify, describing bots using legitimate browsers, valid fingerprints, residential or mobile proxies, and realistic interaction timing. In many cases, bots were only detectable through persistence, scale, or downstream impact rather than obvious technical indicators.

Defensive strategies must therefore assume that bots will appear human at the surface level. Reliance on IP reputation, user-agent strings, or simple rate limiting alone is no longer sufficient. Effective detection increasingly depends on behavioral analysis, session consistency, and

historical patterns across workflows. Organizations that continue to rely on surface-level signals risk either missing attacks or blocking legitimate users.

Imperva Advanced Bot Protection detects residential proxy networks and CAPTCHA-solving services by analyzing propagation patterns, timing inconsistencies, and infrastructure signals. It also leverages advanced browser and device fingerprinting to identify inconsistencies across client attributes and detect automation designed to mimic legitimate users.

8. Plan for Persistent Automation, Not One-Time Attacks

Analysts consistently reported that bots do not disappear once mitigated. Instead, they return with modified tactics, new infrastructure, or alternate workflows. AI has accelerated this persistence by enabling bots to learn from failed attempts and adapt rapidly.

Organizations should design bot mitigation strategies with persistence in mind. Baseline protections should remain active even during quiet periods, and mitigation effectiveness should be reviewed continuously rather than only during incidents. Treating bot defense as a one-off exercise or an event-driven activity leaves organizations exposed to long-running, low-visibility attacks that cause cumulative damage over time.

9. Reduce Exposure Across the Entire Digital Surface

Bots increasingly exploit gaps between systems rather than weaknesses within a single application. Analysts described attacks that moved laterally across websites, APIs, mobile applications, and third-party integrations, taking advantage of inconsistent controls and visibility.

Organizations should adopt a holistic approach to bot mitigation that covers all digital touchpoints. This includes maintaining accurate inventories of applications and APIs, enforcing consistent authentication and authorization, and ensuring that monitoring and mitigation span the full request path. Fragmented defenses create opportunities for automation to bypass controls entirely.

10. Combine AI-Driven Detection with Human Expertise

While AI is essential for detecting and responding to bot activity at scale, analysis shows that bot mitigation cannot be fully automated. Attackers increasingly use creative, adaptive techniques that exploit edge cases and business context in ways that automated systems alone may miss.

Effective bot defense therefore requires a combination of AI-driven detection and experienced human oversight. Skilled analysts are needed to interpret intent, tune controls, manage false positives, and respond to novel

attack patterns. Organizations that rely solely on automated tools risk falling behind attackers who adapt faster than static models can be retrained.

Imperva Security Analyst Services (SAS) combine AI-driven detection with expert human oversight to stay ahead of adaptive, evasive bot attacks.

11. Avoid Revealing Your Full Defensive Strategy

The insight that closed last year's report remains valid and is reinforced by 2025 experience. Deploying all mitigation techniques simultaneously across an entire platform exposes defensive patterns and accelerates

attacker adaptation. Analysts described attackers explicitly testing controls to understand how and when defenses were applied.

Organizations should take a strategic, selective approach to mitigation. Stronger controls should be reserved for high-risk workflows or periods, and defenses should be adjusted dynamically based on conditions. Maintaining unpredictability makes it harder for bots to learn, adapt, and evade controls over time.

Imperva Advanced Bot Protection, can apply adaptive challenge mechanisms, including computational or proof-of-work techniques, to increase the cost of automation while minimizing friction for legitimate users.



Conclusion/Looking Ahead

The lessons from 2025 are clear. Automation now dominates how applications and APIs are accessed, and AI is accelerating its scale, speed, and sophistication.

This is not a temporary shift. As AI adoption expands, digital environments are becoming fundamentally more automated. The challenge for organizations is no longer distinguishing humans from bots, but distinguishing automation that aligns with business intent from automation that exploits it.

APIs and identity systems will remain primary targets, while AI continues to compress attacker learning cycles and blur the line between legitimate and malicious activity. Surface-level signals are no longer sufficient to determine risk.

In this environment, organizations that focus solely on blocking attacks will struggle. Effective defense requires a shift toward governance, combining visibility across applications and APIs, adaptive and intent-aware controls, and human expertise to continuously evolve defenses.

Ultimately, bot mitigation is no longer just a security problem. It is about controlling how automation interacts with digital infrastructure. Organizations that can enforce intent at scale and in real time will be best positioned to maintain trust, performance, and resilience as the web becomes increasingly machine-driven.




About Thales

Thales is a global leader in cybersecurity, helping the most trusted companies and organizations around the world protect critical applications, sensitive data, and identities anywhere at scale. Through our innovative services and integrated platforms, Thales helps customers achieve better visibility of risks, defend against cyber threats, close compliance gaps, and deliver trusted digital experiences for billions of consumers every day.

Definitions

WHAT IS A BOT?	In the context of the internet, a bot is a software application that runs automated tasks. Tasks can range from simple actions like filling out a form to more complex functions like scraping a website for data.
WHAT IS A BAD BOT?	Bad bots are malicious automated programs that exploit applications to extract data, commit fraud, and disrupt services. They enable activities such as credential stuffing, scalping, and DDoS attacks, with 21 attack types defined by OWASP.
WHAT IS AN AI BOT?	An AI bot is a software program that uses artificial intelligence techniques such as machine learning and natural language processing to perform tasks autonomously and adapt its behavior based on data and context. Unlike traditional bots, AI bots can make decisions, mimic human interactions, and execute complex, multi-step activities across applications and APIs.
WHAT IS AN AI CRAWLER BOT?	An AI crawler bot is an automated program that uses artificial intelligence to systematically browse websites and APIs to collect data, often for training AI models or powering search and discovery. Unlike traditional crawlers, AI crawlers can adapt their behavior, follow complex data paths, and extract structured and unstructured content at scale.
WHAT IS AN AI FETCH BOT?	An AI fetch bot is an automated program that retrieves specific pieces of content on demand often in response to a user query or an AI agent request to support real-time processing or generation. It focuses on targeted retrieval, pulling precise data from webpages, APIs, or databases as needed. An AI crawler bot, by contrast, systematically browses and collects data at scale, typically without a specific immediate query, to build datasets or indexes.



THALES

CYBERSECURITY

cpl.thalesgroup.com/bad-bot-report

For contact information, please visit
cpl.thalesgroup.com/contact-us

