HUMAN

The CISO's Guide to AI and Agentic Traffic

What security leaders need to know about the fastest-growing category of internet traffic and the threats converging around it.

8x

automated traffic growth vs. human traffic

7,851%

agentic AI traffic growth YoY

187%

Growth in AI Traffic YoY

0.5%

separates benign from malicious automation

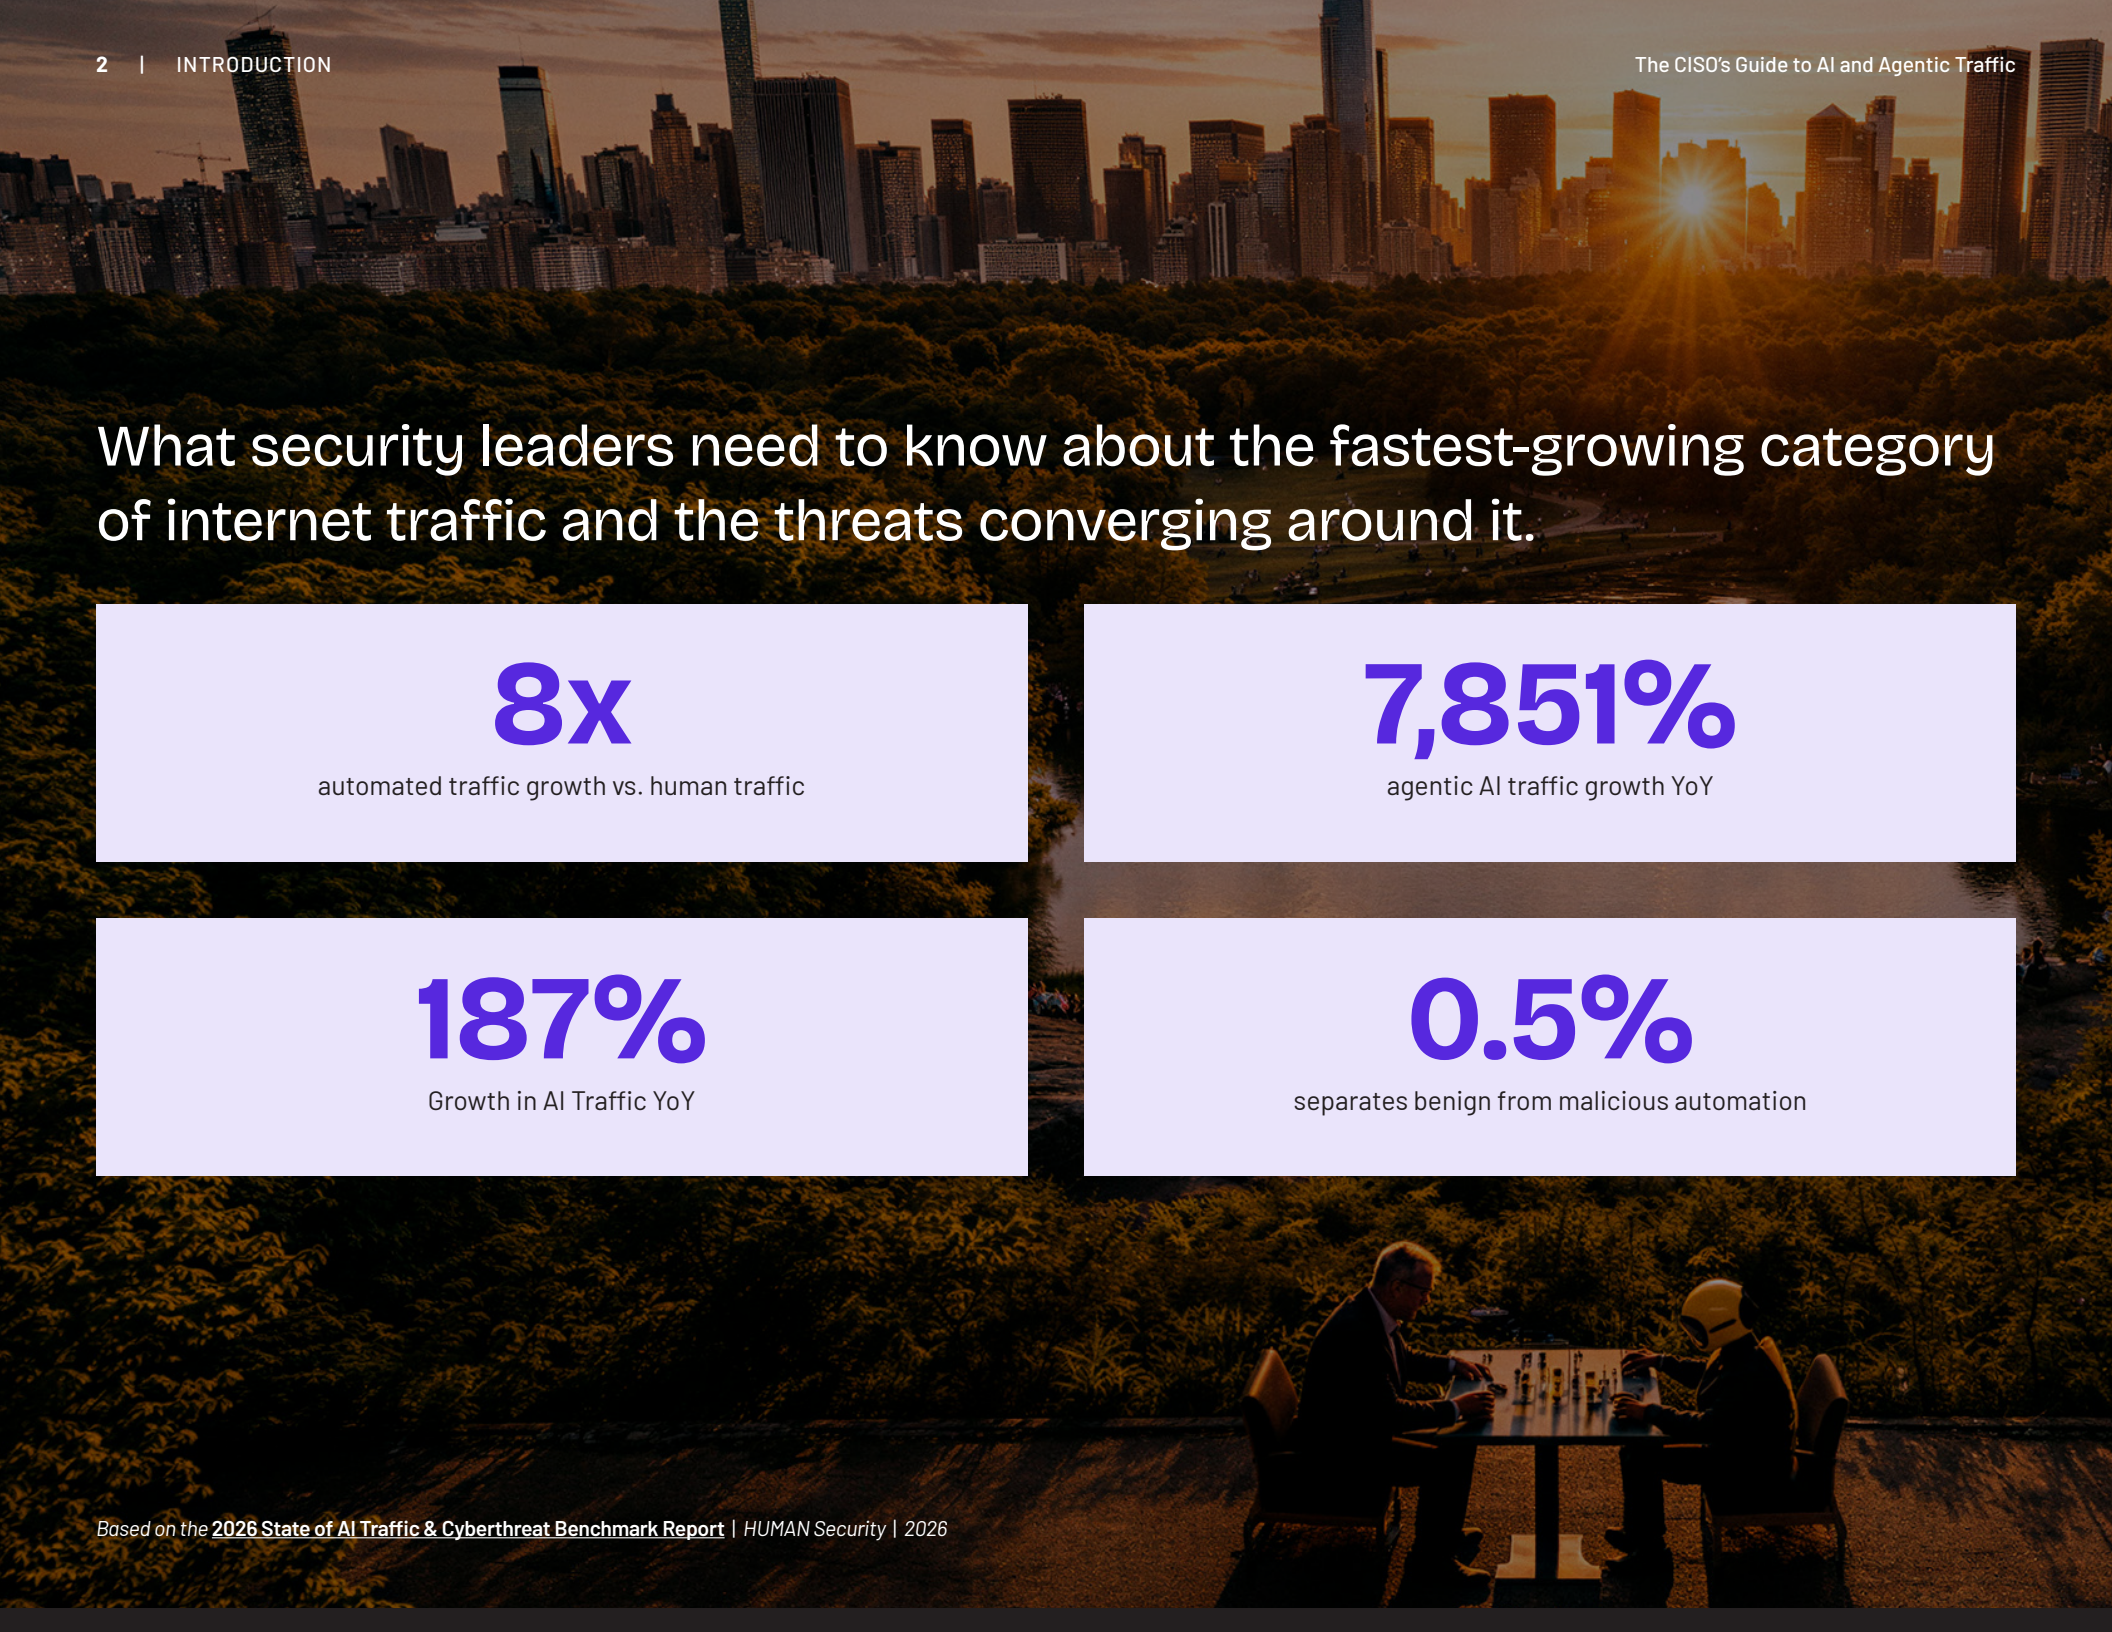

Based on the **2026 State of AI Traffic & Cyberthreat Benchmark Report** | *HUMAN Security | 2026*

The Landscape Has Changed

In 2025, the Human Defense Platform processed more than one quadrillion interactions. The data reveals a structural shift: automated traffic is growing eight times faster than human traffic, AI-driven traffic nearly tripled over the course of the year, and for the first time, AI systems are not just reading the web but transacting on it.

For CISOs, this changes the threat model. The behaviors that once reliably indicated an attack—rapid page navigation, programmatic form completion, automated checkout—may now be a legitimate consumer workflow powered by an AI shopping assistant. At the same time, threat actors are exploiting these same patterns to conduct fraud at machine speed.

AI-Driven Traffic By Month

Indexed to January, 2025. Includes training crawlers, scrapers, AI agents, and agentic browsers.

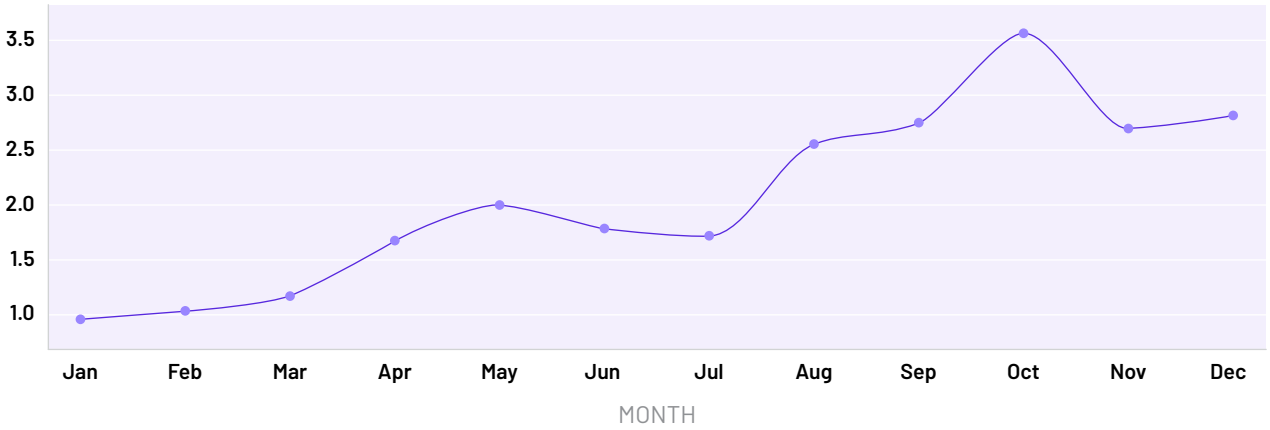

Figure 1: Monthly rate of AI-driven automation, January – December 2025.

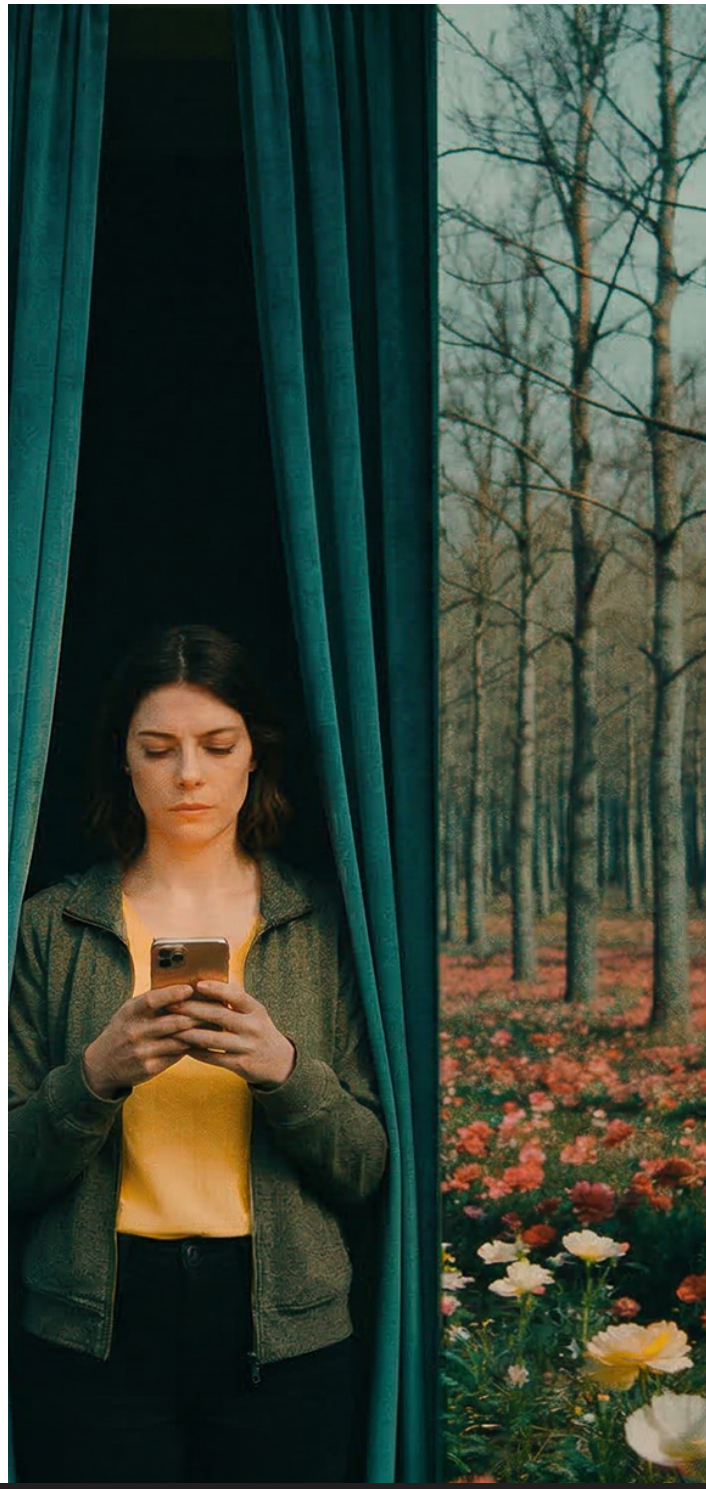

The numbers that matter

187%
growth in monthly AI-driven traffic, January to December 2025

69%
of AI-driven traffic generated by OpenAI's bots alone

95%+
of AI traffic concentrated in retail, media, and travel

The operator concentration is a critical planning input. Access policy decisions about three companies—OpenAI (69%), Meta (16%), and Anthropic (11%)—determine the vast majority of any organization's AI-driven traffic profile. This concentration creates both a governance lever and a single-point dependency.

AI-Driven Traffic By Operator

Top five operators by share of observed AI-driven traffic, January–December 2025.

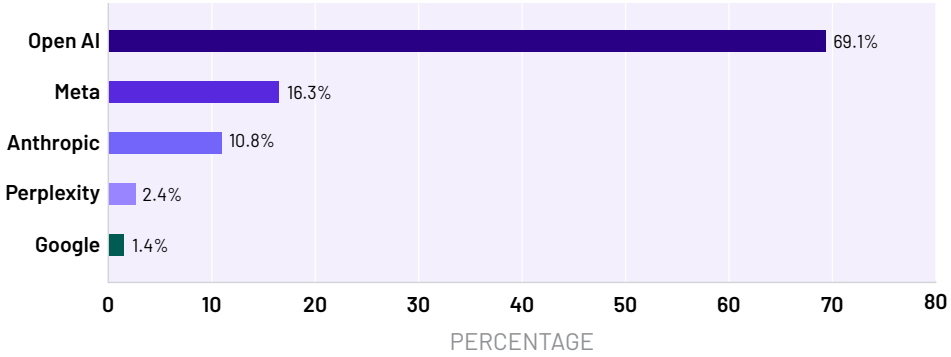

Figure 2: Percent of AI-driven automation in 2025 by operator.

AI-Driven Traffic By User Agent

Top user agents by share of observed AI-driven traffic, January–December 2025.

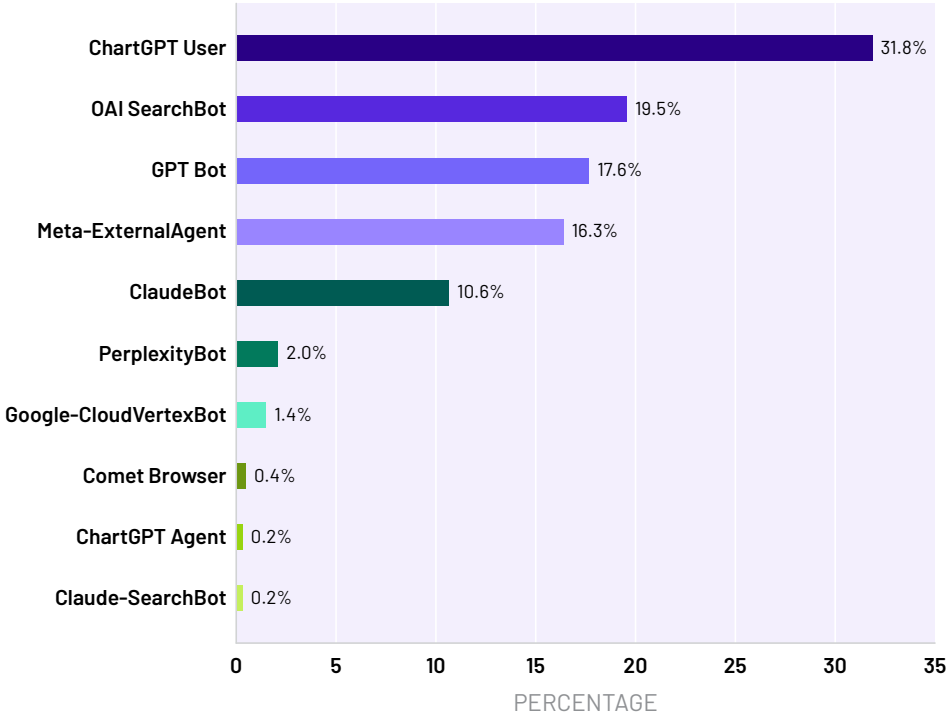

Figure 3: Percent of AI-driven automation in 2025 by user agent.

What's different about agentic AI

The most structurally significant development is the emergence of agentic AI. Unlike crawlers and scrapers, which read the web, AI agents interact with it. They navigate pages, fill forms, compare products, manage accounts, and execute transactions—autonomously.

In 2025, 77% of observed agentic AI activity occurred on product and search pages. But 8.8% was on account pages, 5% on authentication flows, and 2.3% on checkout pages. That last figure is small in share but significant in intent: agents are now completing transactions without direct human involvement. This was largely theoretical before 2025. The data confirms it is now operational.

Agentic AI Traffic By Page Category

Share of observed agentic traffic by destination page type.

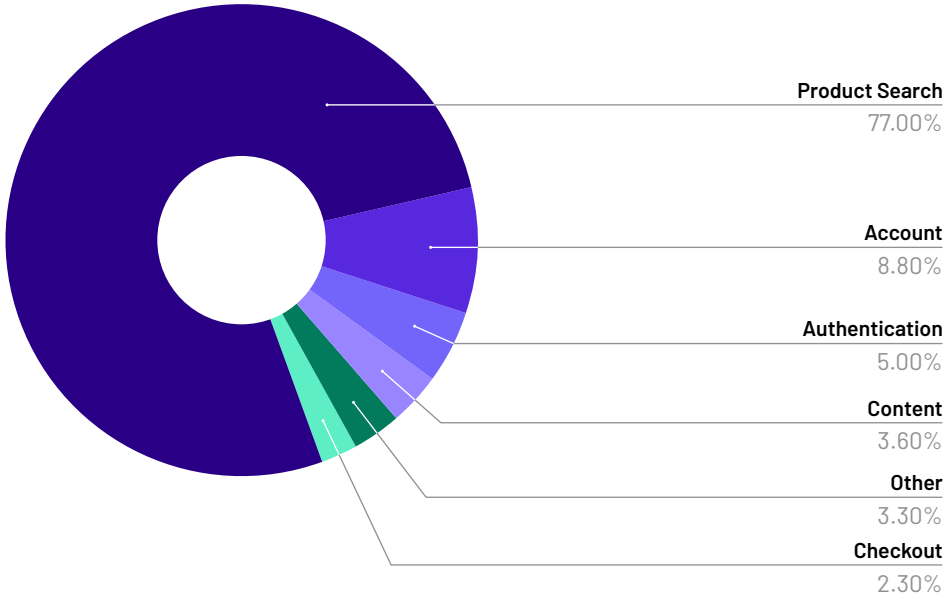

Figure 4: Agentic traffic by page category, 2025.

Share of Agentic AI Traffic By Industry

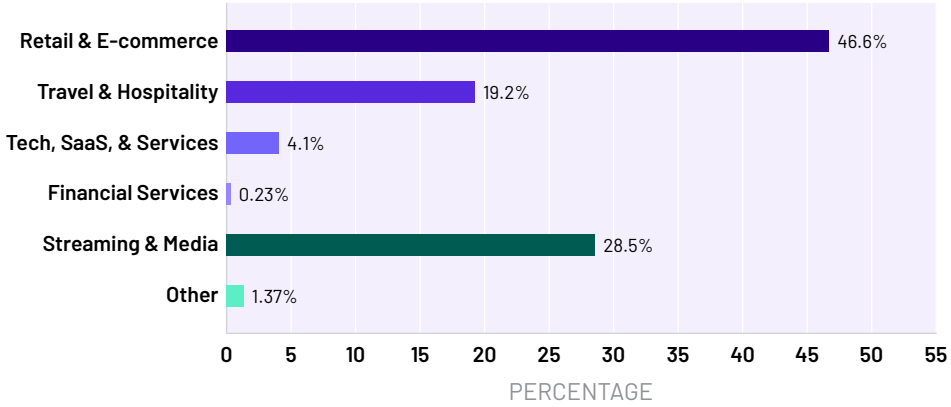

Figure 5: Agentic traffic by vertical, 2025.

Why this matters for your security posture

Agents that create accounts, manage sessions, and complete transactions are taking actions that carry financial and contractual weight. An AI agent rapidly browsing products and completing a checkout may be a consumer's shopping assistant or an automated fraud operation. The behavior is identical. The intent is not. Your security stack needs to distinguish between them.

Three Types of AI Traffic Security Leaders Need to Know

Not all AI traffic is the same. Security teams that treat "AI traffic" as a monolith will make poor access decisions. The table below distinguishes the three categories observed in HUMAN's 2025 telemetry, each with different risk profiles and governance requirements.

	 **Training Crawlers**	 **AI Scrapers**	 **Agentic AI**
What it does	Collects data in bulk to train or refine ML models	Extracts specific, timely data to power real-time AI features (RAG, search, comparison tools)	Navigates pages, fills forms, manages accounts, and executes transactions autonomously
Share of AI traffic	67.5%	31.9%	~1.7% (grew 7,851% YoY)
Top industries	Retail: 62.5% Media: 19.7% Travel: 16.6%	Retail: 36.7% Media: 40.9% Travel: 21.1%	Retail: 46.6% Media: 28.5% Travel: 19.2%
Security concern	IP extraction, competitive data loss, spoofing of declared identity	Content theft, real-time pricing exposure, paywall circumvention	Autonomous fraud: carding, ATO, fake account creation at machine speed

Source: Human Defense Platform telemetry, January–December 2025.

Declared identity is unreliable

HUMAN's Satori threat intelligence team analyzed the declared identities of AI training crawlers against behavioral and infrastructure signals and found that a significant portion of requests claiming to be ChatGPT, Mistral, and Perplexity bots did not originate from those operators' infrastructure. Attackers spoof user-agent strings to exploit the trust organizations extend to recognized AI crawlers, bypassing robots.txt allowlists and rate-limit exemptions.

Operational takeaway

Organizations that allowlist crawler traffic based solely on user-agent strings are granting access to an unknown number of unauthorized actors. Effective AI traffic management requires behavioral validation beyond declared identity—correlating network properties, browser authenticity signals, and interaction behavior to verify who is actually making the request.

Autonomous agents as attack infrastructure

HUMAN's Satori team analyzed traffic from publicly exposed OpenClaw gateways and found patterns spanning routine automation to clear abuse. Instances were observed generating synthetic referral traffic with fabricated social media UTM parameters, the tracking tags marketing and advertising teams use to tie site visits and conversions back to specific ad campaigns and digital marketing spend. Other activity included conducting high-velocity directory brute-forcing against web applications, and serving as targets for infostealer malware adapted to exfiltrate agent configuration secrets. Tools like OpenClaw lower the skill threshold for internet fraud, enabling users with no security expertise to conduct attacks that previously required hands-on technical knowledge.

The Threat Convergence

The digital surfaces that AI agents are reshaping—product discovery, account management, checkout—are the same surfaces that threat actors target most. The connection is structural: as more commercial activity moves through automated channels, the attack surface expands with it.

Account takeover: shifting to post-login compromise

While overall ATO volume fell more than 30%, post-login account compromise attempts more than quadrupled year over year, reaching an average of 402,000 per organization. Attackers abuse session tokens, manipulate account settings, or exploit weak step-up controls to maintain access after legitimate login. The decline in login-stage attacks suggests point-of-login defenses are working, forcing threat actors to get more technical. EMEA-sourced ATO traffic exceeded 13% of login attempts, versus less than 3.5% globally.

Carding: volume surge, agent-mediated attacks emerging

Checkout interactions blocked grew 20%+ from 2024 and 250% from 2022. US-based IPs account for a majority of carding attempts for the second consecutive year. Critically, Satori researchers documented carding behavior mediated by an AI agent: 11 card additions and 6 payment attempts across two sessions, then a pivot to loyalty-point redemption after cards failed. A verified agent can still be misused. The question is not only who the agent is, but what it is being allowed to do.

Web scraping: approaching 20% of all traffic

The median percentage of traffic attempting a scraping attack nearly doubled from 2022 to 2025, approaching 20%. For heavily targeted organizations, scraping exceeded 61% of traffic. Volume grew 47% YoY and 138% since 2022. The rise coincides with AI-driven automation growth, and only 0.5% separates the rate of benign from malicious automation across all interactions HUMAN analyzed.

Fake account creation: 89% growth

Fake accounts detected per organization grew 89% from 2024 to 2025, on top of 259% growth the prior year. These accounts serve as infrastructure for incentive abuse, fraudulent orders, and fake reviews.

402K

post-login compromise attempts per org (4x YoY)

250%

carding volume surge since 2022

~20%

of all traffic is a scraping attempt

The 0.5% gap

Across all interactions analyzed by the Human Defense Platform, only one half of one percent separates the rate of benign automation from the rate of malicious automation. This is the fundamental challenge: the margin between helpful and harmful automation is razor-thin, and closing it requires understanding intent, not just detecting presence.

Why "Bot or Not" No Longer Works

For years, digital defenses asked a binary question: is this traffic human or not? That framing no longer holds.

When a consumer's AI shopping assistant rapidly navigates product pages, fills a checkout form, and completes a purchase, it looks identical to a carding attack. When a legitimate AI scraper retrieves real-time pricing data to power a comparison tool, it looks identical to competitive intelligence theft. The same company can operate across all three AI traffic categories simultaneously, meaning operator-level access decisions don't map cleanly to behavioral distinctions.

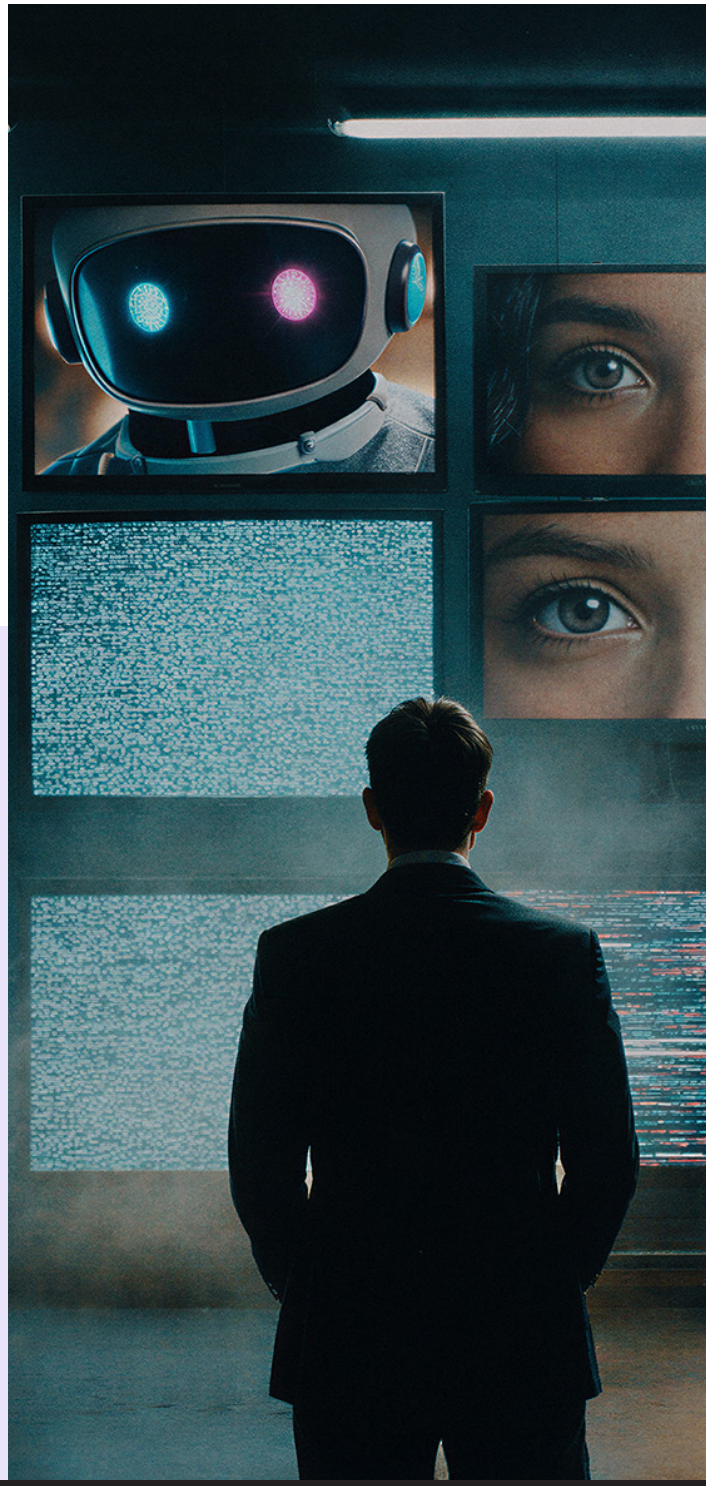

Block everything automated

Organizations that treat all automation as hostile will block revenue. AI agents represent a high-converting new channel—businesses accessible to agents capture demand that others miss.

Allow everything automated

Organizations that allow automation unchecked will absorb fraud. Autonomous transactions introduce unique security risks and fraud surfaces that differ qualitatively from simple information retrieval.

The path forward: intent-based trust

The question is no longer whether traffic is automated. It is whether a given interaction is trustworthy, regardless of whether it comes from a human, an AI agent, or an agentic browser. This requires three capabilities that most security stacks lack today:

1. Actor-level visibility

You need to identify and classify every actor on your surfaces, whether human, bot, or AI agent, and attribute them to known providers or platforms. User-agent strings alone are insufficient; behavioral validation, browser authenticity checks, and infrastructure correlation are required.

2. Behavioral context across the session

Evaluating agent actions in isolation offers very little security value. A single password change looks harmless unless preceded by dozens of login retries. A loyalty redemption is benign unless it follows a cascade of failed credit card checks. You need visibility into the full chain of actions.

3. Adaptive, dynamic trust enforcement

Trust is not a score assigned once. It is a continuous evaluation that evolves based on what the agent does, not just what it claims to be. Controls should tighten or relax based on real-time behavioral signals, allowing legitimate agents to transact while constraining suspicious sequences before they reach sensitive operations.

Four Priorities for Security Leaders

01.

Audit your AI traffic exposure now

Most organizations do not know what percentage of their traffic is AI-driven, which operators are hitting their properties, or what those systems are doing. Start with visibility. Understand the volume, composition, and behavior of AI traffic across your web properties. Access policy decisions about a small number of AI companies have outsized effects on your overall AI traffic profile. You cannot govern what you cannot see.

02.

Stop allowlisting on declared identity alone

HUMAN's research confirms that a significant portion of traffic claiming to be known AI crawlers does not originate from those operators' infrastructure. Attackers spoof user-agent strings to bypass robots.txt allowlists and rate-limit exemptions. Require behavioral validation—network properties, browser authenticity, interaction patterns—before granting trust. The industry is moving toward cryptographic verification (e.g., HTTP Message Signatures per RFC 9421) as the standard for reliable agent identity.

03.

Build an AI traffic governance policy, not just a blocking rule

Blocking all AI traffic is a business decision, not just a security one, and increasingly it is a bad one. Users who arrive via AI-generated answers convert 4.4x better than those from traditional search (Semrush). Among consumers who have tried agentic shopping, 85% said it improved the experience (Adobe). Gartner predicts that by 2028, AI agent machine customers will replace 20% of interactions at human-readable digital storefronts. Your commerce teams will push back hard on blanket blocking. Instead of a binary allow/block, define granular policies: which agents can access which surfaces, what actions are permitted at each stage of the customer journey, and what rate limits and behavioral boundaries apply. Treat AI traffic governance as you would any access control framework, with roles, permissions, and audit trails.

04.

Plan for the convergence of AI commerce and fraud

The same digital surfaces AI agents are reshaping are the primary targets of every attack category documented in this report. As agentic commerce scales, so does the attack surface. One immediate gap to assess: most fraud and chargeback tools were built for human commerce and rely on device fingerprints, IP reputation, and cardholder history. These signals are weak or missing when traffic comes through agents. Evaluate whether your existing stack can distinguish trusted agents from bots and adapt models to use agent identity, cryptographic signatures, and machine-specific behavioral baselines. Early 2026 data confirms the momentum has not slowed. Build cross-functional alignment between security, fraud, product, marketing and commerce teams before the volume forces reactive decisions.

How HUMAN Helps

HUMAN Security is the global leader in AgenticTrust. Powered by one of the world's largest behavioral signal networks, HUMAN analyzes over a quadrillion digital interactions each year to distinguish legitimate activity from fraud, abuse, and automated manipulation.

HUMAN Sightline Cyberfraud Defense

Unified, real-time visibility across bots, humans, and AI agents. Multi-method detection covers the entire customer journey, differentiating good from bad bots, risky from legitimate AI, and fraudulent from authentic human behavior. Includes advanced reporting and investigation tools to isolate and track distinct attacker profiles over time, plus custom fraud models that combine HUMAN's signals with your first-party data.

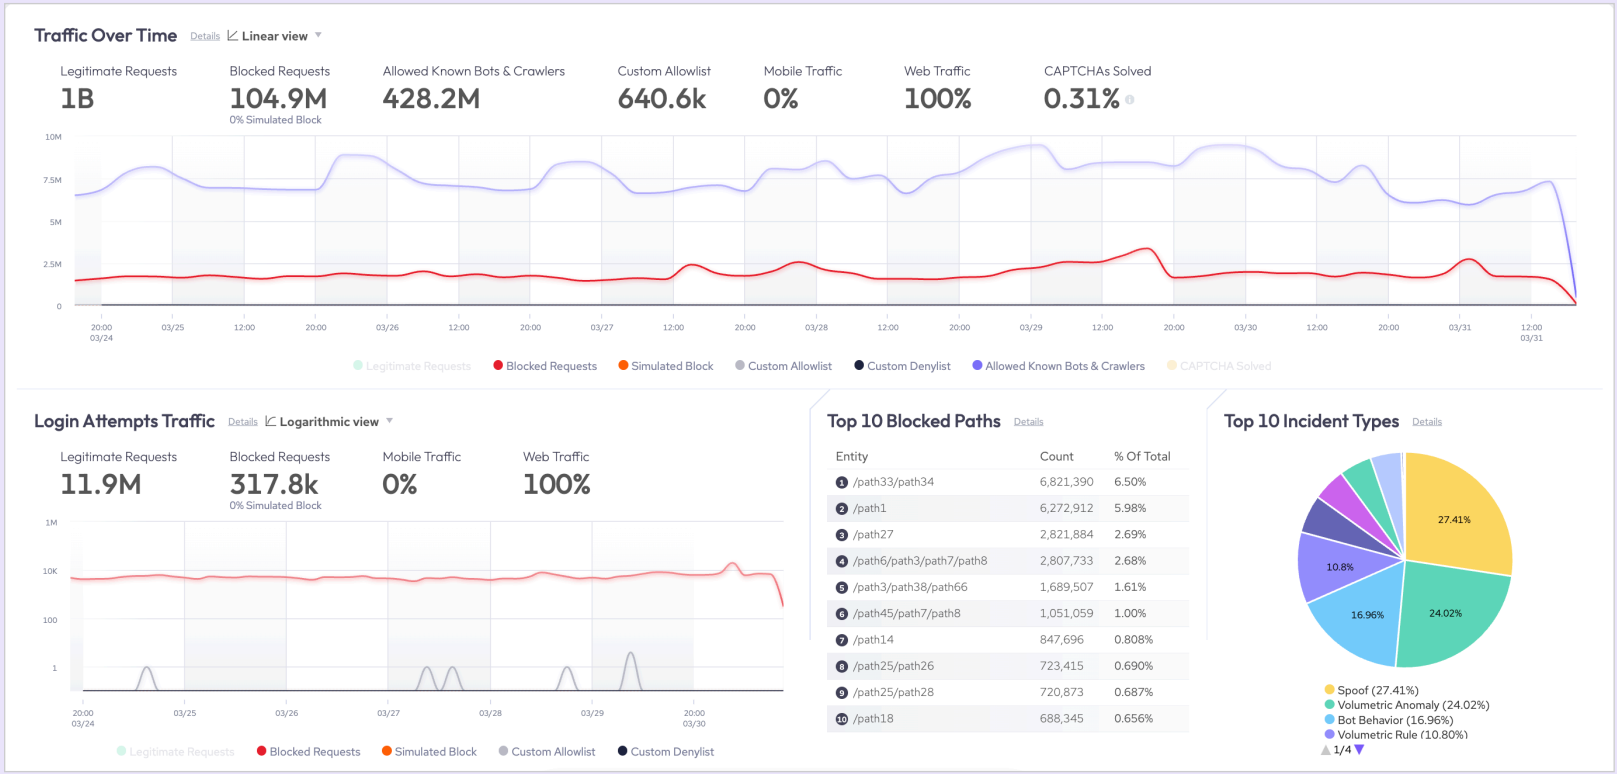

Figure 6: HUMAN Sightline Dashboard. Example data, not indicative or representative of any agent, customer, or organization.

AgenticTrust

The trust and control layer for agentic AI, built as a module within Sightline. AgenticTrust turns unknown AI traffic into visible, controllable, and trusted interactions.

Detect

Identify which AI agents are interacting with your applications, the paths they travel, and the actions they take—including agents that do not self-identify.

Verify

Authenticate agent identity using cryptographic digital signatures that cannot be spoofed. Automatically evaluate agent trustworthiness with AI-generated behavioral insights.

Govern

Set rules for critical flows. Define what agents can and cannot do at each stage of the customer journey. Enforce rate limits, permission boundaries, and adaptive controls that tighten or relax based on real-time behavior.

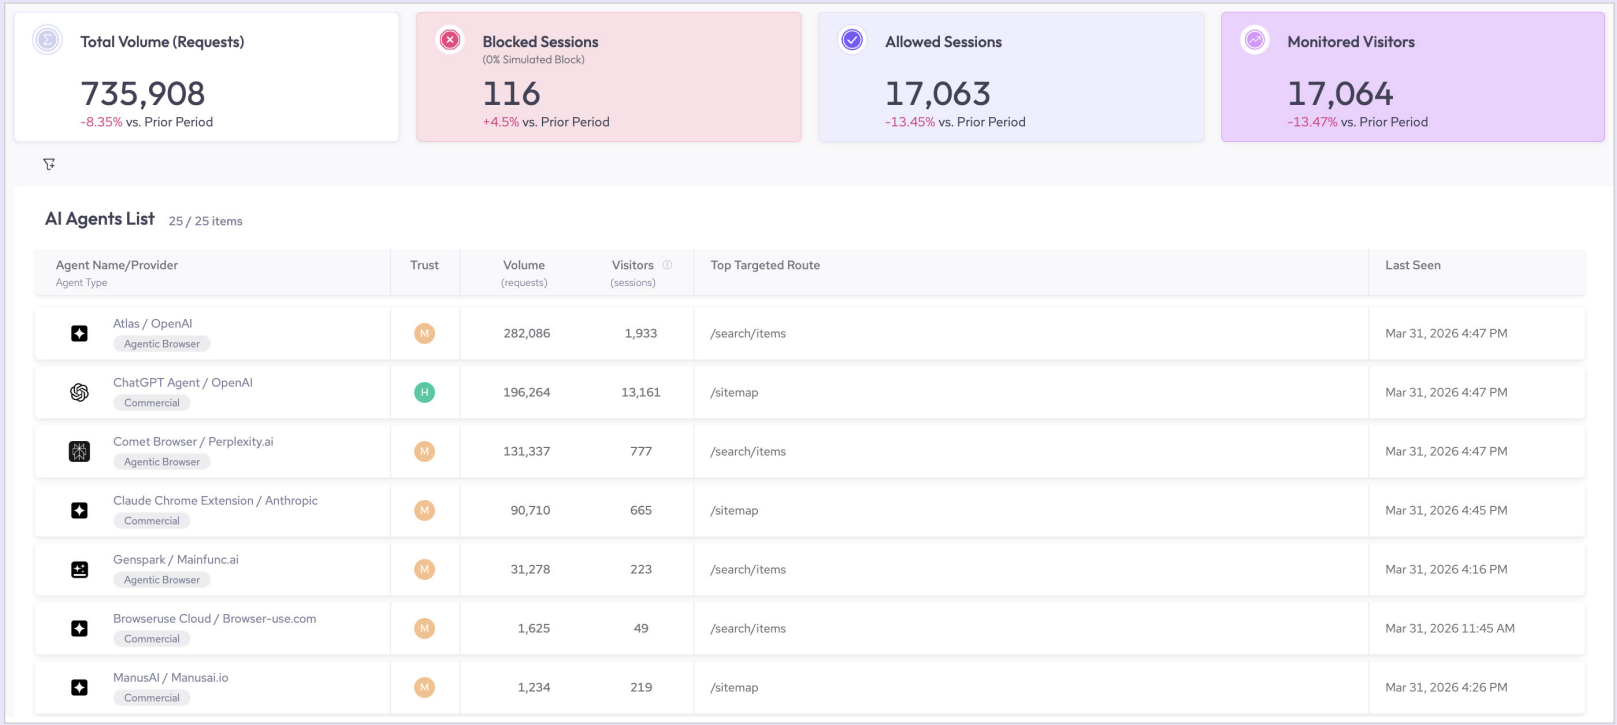

Figure 7: HUMAN AgenticTrust Dashboard. Example data, not indicative or representative of any agent, customer, or organization.

Sightline + AgenticTrust

Together, they provide holistic visibility, governance, and control across bots, humans, and AI agents through a single trust framework, enabling organizations to safely embrace agentic commerce while reducing fraud and abuse.

View the full 2026 State of AI Traffic & Cyberthreat Benchmark Report for complete analysis, industry-specific benchmarks, dark web pricing data, and detailed methodology.

VIEW NOW

To see which AI agents are impacting your traffic today, contact HUMAN at humansecurity.com.

It All Runs Better On Trust.

HUMAN

About HUMAN

HUMAN is a leading cybersecurity company committed to protecting the integrity of the digital world. We enable trusted interactions and transactions across the full spectrum of online actors: humans, bots and AI agents. HUMAN verifies 20 trillion digital interactions, providing unparalleled telemetry data to enable rapid, effective responses to the most sophisticated threats. Recognized by our customers as a G2 Leader, HUMAN continues to set the standard in cybersecurity. For more information please visit www.humansecurity.com.