Can Defenses Catch Up Before Agentic AI Catches On?

Enterprises doubled their AI security budgets and confidence in just 12 months. But 53% expect agentic AI attacks to hit within six months. When threats evolve faster than implementation cycles, who wins the race?

The Arkose Labs 2025 AI Maturity in Cybersecurity Report

ARKOSE LABS

CONTENTS

EXECUTIVE SUMMARY

Security teams have six months, or less, to prepare for agentic AI attacks. Implementation cycles take 12-18 months. The math doesn't work.

That timeline gap drives everything revealed in Arkose Labs' second annual AI maturity survey: a cybersecurity landscape where speed determines survival. Enterprises nearly doubled their "very well prepared" ratings from 23% to 44% year-over-year and committed 32% of cybersecurity budgets to AI solutions, up from 27% in 2024. The buy-versus-build debate? Settled. Eighty-one percent now gain more value from purchasing solutions than building them, up 19 points from 62% as vendor partnerships evolved into operational extensions of security teams. Integration barriers dropped 11 points, removing what was previously the primary obstacle to effectiveness.

Can confidence substitute for capability? The data says no. It's an arms race where adversaries adapt and improvise as fast as defenders deploy. Fifty-five percent still lose $10 million-$500+ million annually. Over half of enterprises already use agentic AI for cybersecurity, but 71% see bad actors' use of agentic AI as the next major threat—one that 53% expect to reach critical status within six months.

The window to close the preparation-protection gap is narrowing. This report examines why only half of enterprises realize measurable benefits despite universal adoption, how the threat landscape evolved from bots to human fraud farms to agentic AI and what security teams must do before the next wave hits.

As with any research project, we have a plethora of data. If you would like us to customize a report for you, please reach out directly to us.

Kevin Gosschalk
Founder and CEO
Arkose Labs
kevin@arkoselabs.com

Frank Teruel
COO
Arkose Labs
f.teruel@arkoselabs.com

agentic AI | \ā-ˈjen-tik ā-ˈī\

noun

: autonomous AI systems that can act independently to achieve goals, make decisions and take actions without constant human oversight

KEY FINDINGS

BAD ACTORS USE OF AGENTIC AI

Fear the Threat

ANNUAL LOSSES

$10M-$500M

TIMING

3-6 Months

Until Agentic AI is Critical Risk

SKILLS GAP ACKNOWLEDGED

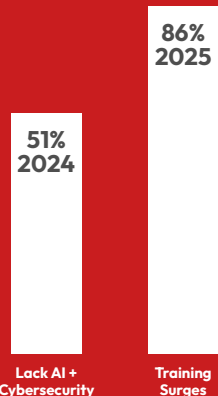

Lack AI + Cybersecurity expertise

Training Surges

OPINION ABOUT CONSUMER-CONTROLLED AGENTS

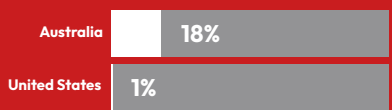

Net Negative for Our Business

AGENTIC AI USE W/O DEFENSES

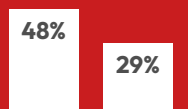

Australia United States

48% of Australian companies use agentic AI offensively without defensive posture versus 29% in the US, which is a 19-point gap

TIERS OF TOP ATTACK TYPE CONCERNS:

TIER 1: IDENTITY FRAUD

Fake Account Creation & Reverse Proxy Phishing

TIER 2: ACCESS & EXPLOITATION

Account Takeovers & API Abuse

TIER 3: VALUE EXTRACTION

Web Scraping, Loyalty Point Theft & Unauthorized Account Sharing

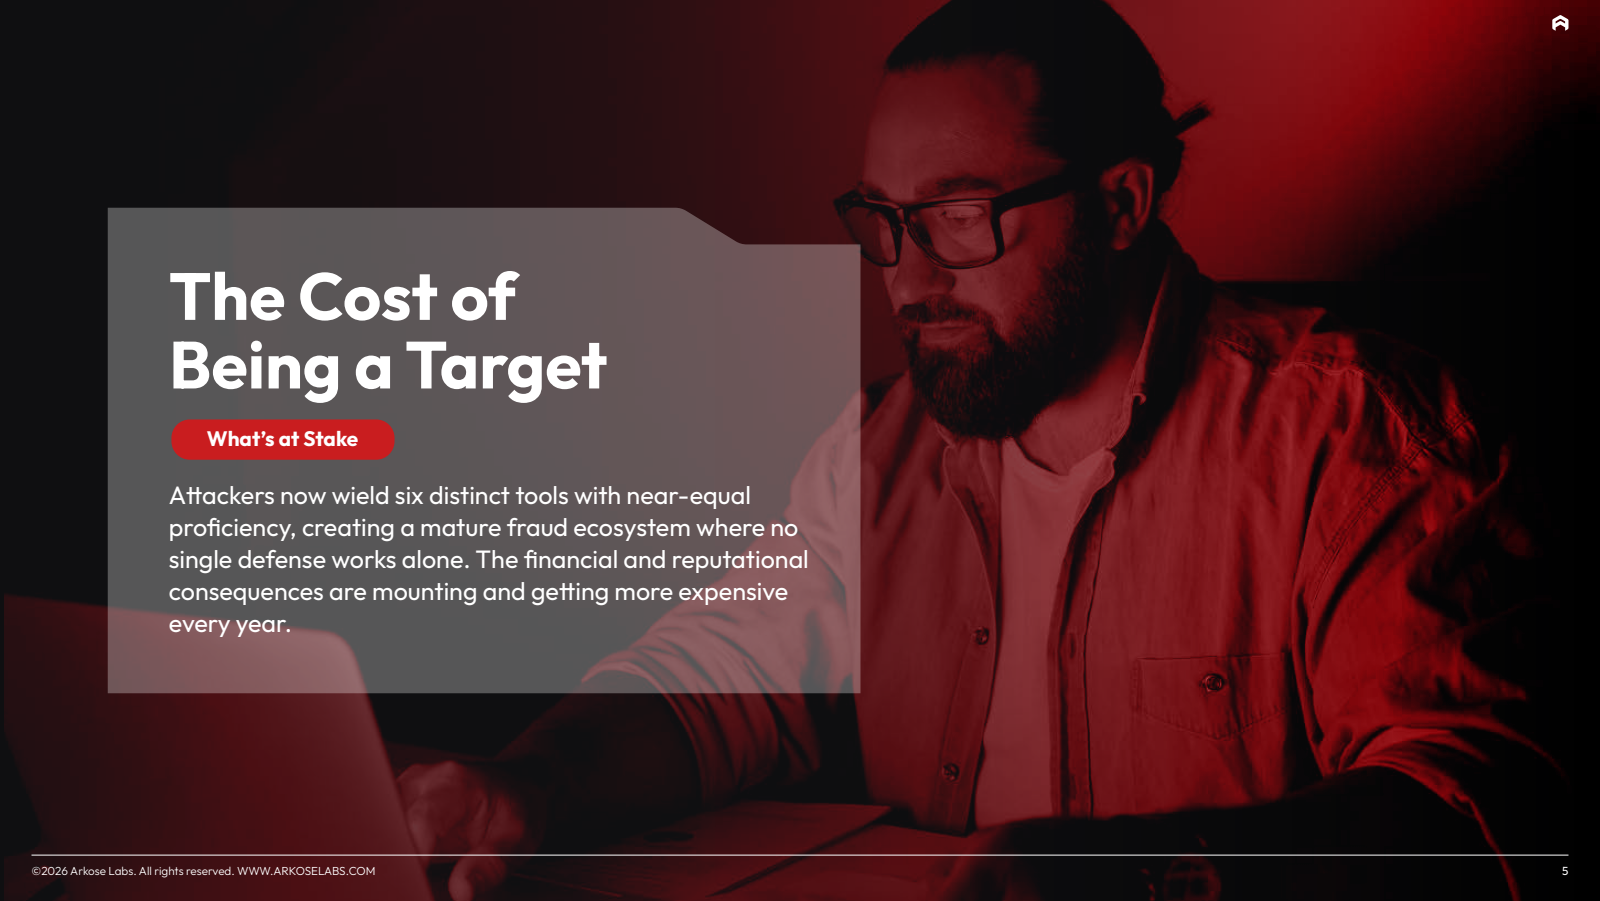

The Cost of Being a Target

What's at Stake

Attackers now wield six distinct tools with near-equal proficiency, creating a mature fraud ecosystem where no single defense works alone. The financial and reputational consequences are mounting and getting more expensive every year.

THE DISTRIBUTED ATTACK TOOL MATRIX

Distribution of the tools that bad actors use for cyberattacks is remarkably balanced. The tight 14%-20% distribution across six attack mechanisms reveals a nearly mature fraud ecosystem where no single technique holds a single advantage, yet.

Regional variations are minimal: US enterprises see slightly more bot traffic (21%) while Australian companies report marginally higher AI agent tools (16%).

Average Percentage of Cyberattacks from the Following Mechanisms

	Total (n=304)	Australia (A) (n=102)	US (B) (n=202)
Bots	20%	19%	21%
Human Fraud Farms	19%	17%	20%
AI-powered Bots	18%	18%	17%
Attack Automation Services	15%	16%	15%
AI Agents	15%	16%	14%
Low-and-Slow Manual Fraud	14%	14%	14%

Q3b: Consider the attack activity your enterprise has seen over the past 12 months, approximately what percentage of those attacks were from bots, AI-powered bots, AI agents, attack automation services, human fraud farms or low-and-slow manual attacks? Provide your best approximation.

"AI agents currently work primarily against weak targets. They're new and not yet at a point where they can solve or bypass advanced protection autonomously. But that is changing quickly. Bots are extremely effective but expensive for financially-motivated fraudsters to operate, and human fraud farms remain a go-to method because they're typically indistinguishable from legitimate human behavior, making detection considerably more difficult."
– Mitch Davies, Senior Data Scientist, Arkose Labs

For security teams, this even spread creates a persistent challenge. More enterprises are leveraging platform approaches that address the entire spectrum rather than optimizing for individual tools, because the data shows attackers already maintain some capability across all six methods.

AUTOMATED FRAUD AT SCALE: THREATS DRIVING 2025'S CONCERN INCREASES - MFA BYPASS BECOMING MORE PREVALENT

Concern levels jumped significantly from 2024 to 2025. Of note: Man-in-the-middle reverse proxy phishing now registers 73% (up from 57% in 2024) among global enterprises, and fake account creation edged into the top concern companies face.

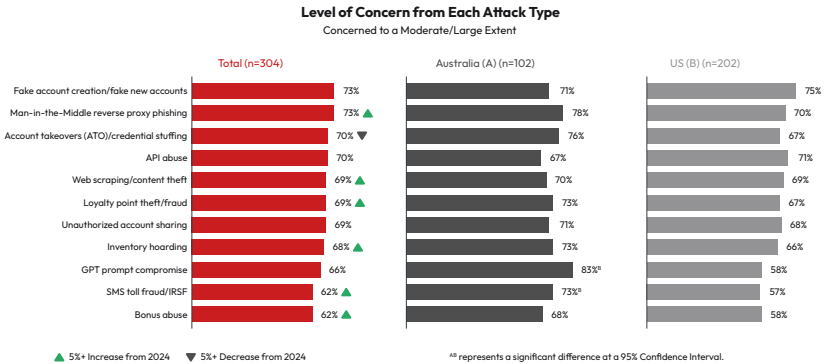

Q1: Think about the critical applications to your business – such as revenue-driving apps, websites/platforms, network systems, etc. How concerned are you with each of the following attack types to those critical business applications? (RATE EACH ROW)

Big increases reflect fraudster profitability. Man-in-the-middle reverse proxy phishing represents modern account takeover: sophisticated kits, like Veiled Marble, V3B, Labhost, that bypass MFA at scale without human social engineering required. This attack type generates convincing sites, intercepts MFA codes, steals session cookies and includes built-in generative AI to create templates, making MFA compromise super user-friendly for criminals. Single operations have compromised millions of accounts, causing hundreds of millions in losses.

Enterprises making progress to stop MITM reverse proxy attacks detect the automation patterns and behavioral anomalies these kits create, rather than relying solely on MFA controls that can now be circumvented.

It's no surprise fake account creation ranks as the top concern. Storm-1152 created 750 million Microsoft[1] fake accounts that the group then sold on the dark web for millions of dollars, showing just how lucrative this fraud has become. Most companies are now deploying AI-powered bot detection and mitigation because as agentic AI matures as a tool, the volume of fake account creations will skyrocket and be harder to detect and stop. Gartner predicts agentic AI[2] will halve account exploitation time by 2027 by automating the attack process with unprecedented speed and sophistication.

1. https://www.arkoselabs.com/actir/
2.https://www.gartner.com/en/newsroom/press-releases/2025-03-18-gartner-predicts-ai-agents-will-reduce-the-time-it-takes-to-exploit-account-exposures-by-50-percent-by-2027

WHEN CONSUMERS LEAVE FASTER THAN THEY ARRIVE

Enterprises worry uniformly (62%-73% concern) but suffer variably (47%-60% consequences). When attacks succeed, consequences cascade across three dimensions. Comparing 2025 over 2024, financial impacts held steady, operational impacts were split and reputational losses increased.

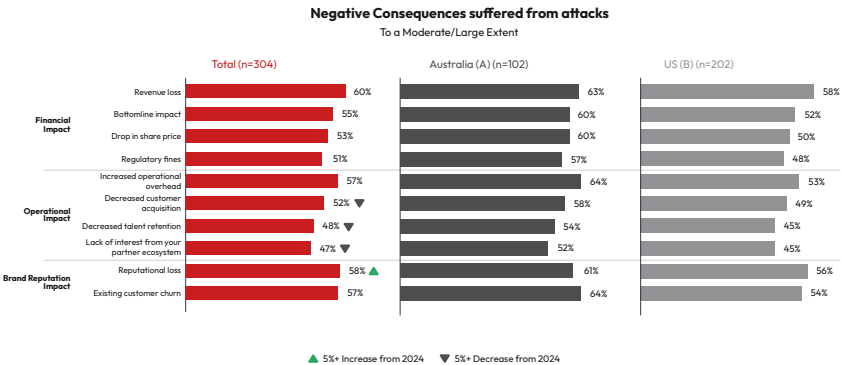

Q2. To what extent has your enterprise suffered negative consequences over the last 12 months in each of the following areas due to the attack types just considered? (RANDOMIZE) (RATE EACH ROW)

Companies are getting better at customer acquisition despite breach visibility, but each successful attack converts to customer churn that erodes the base. This creates compounding damage: shrinking customer bases absorb fraud costs across fewer accounts, increasing per-customer impact exponentially.

Reputational loss now cascades across investors, customers, regulators and partners simultaneously. For enterprises already losing revenue, customer churn accelerates decline. Each incident weakens organizational credibility needed to survive the next attack, while operational overhead (57%) strains the security teams managing fallout. The pattern shows why fraud prevention matters as much as response. The enterprises that stop attacks before they succeed avoid the entire consequence cascade, preserving customer relationships and brand trust.

THE REVENUE RATIO TRAP

Financial losses escalated from 2024 to 2025: enterprises reporting losses exceeding $10 million rose from 53% over the last two years to 55% within the last year. The $10 million-$99 million range held at 39%. Revenue distribution of survey respondents: 37% mid-market (under $1 billion), 31% large ($1 billion-$5 billion), 32% very large ($5 billion+).

Approximate Cost of Negative Consequences from Threats to Critical Business Applications

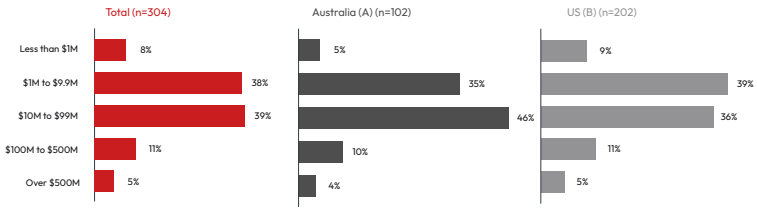

Q3a: What is the approximate DOLLAR quantification of these consequences over the past 12 months... that is, how much have these negative consequences COST your business to the best of your knowledge? (SELECT ONE)

The 39% plateau reveals different ROI profiles by revenue scale. Mid-market companies losing $10 million-$99 million face 1%-40% revenue impact. They will experience the highest ROI from fraud prevention as savings directly improve margins. Large enterprises face 0.2%-10% revenue impact with an optimal intervention window where losses justify fraud prevention investment before a crisis can occur. Very large enterprises face just 0.2%-2% impact, yet $50 million in losses still represents capital that could fund expansion instead of enriching attackers.

The same dollar loss means different things at different company sizes. Security teams presenting "$50 million lost to fraud" often struggle to get budget approval, but reframing it as "10% of our annual revenue" gets immediate board attention. When fraud exceeds 5% of revenue, companies need to redesign their security architecture, not just add more tools. For mid-market companies, the combination of high revenue-ratio losses and the broader pattern of customer churn creates compounding pressure that makes fraud prevention increasingly urgent.

The Investment Surge

What's at Stake

Enterprises are pouring resources into AI-powered defenses faster than anyone predicted, with budgets jumping ahead of schedule and confidence doubling in just 12 months. The question isn't whether companies are investing, it's whether they're building fast enough.

AI SECURITY SPENDING ACCELERATES BEYOND PREDICTIONS

Enterprise commitment to AI-powered security is accelerating faster than predicted, revealing urgency. In 2024, companies projected that 27% of their cybersecurity budget would be spent on AI solutions by 2026.

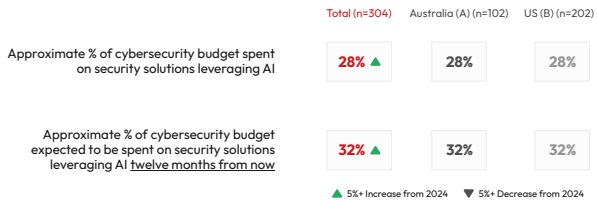

Q5c: Approximately what percentage of your overall cybersecurity budget is spent on security solutions leveraging AI (including generative AI and Agentic AI) today?

Q5d: And approximately what percentage of your overall cybersecurity budget do you estimate will be spent on security solutions leveraging AI (including generative AI and Agentic AI) 12 months from now?

AI threats are materializing fast, forcing budgets to catch up rather than follow planned roadmaps. The 32% figure roughly mirrors the 33% of attacks now using AI methods, suggesting enterprises pursue spending parity with adversaries—a dangerous assumption that equal investment produces equal capability.

The economics favor fraudsters: they need only one successful attack effort, while defenders must close every vulnerability. With one-third of security budgets concentrated on AI solutions and losses ranging from $10 million to $500+ million annually, enterprises will disrupt adversarial economics by driving up attacker costs. Ready-made AI bots require significant investment from fraudsters, including initial outlays plus recurring fees for proxies, training and hosting. Next-generation mitigation increases fraudster time and effort through adaptive challenges, eroding attack ROI until adversaries abandon the target for weaker alternatives.

PREPAREDNESS SURGE: FROM 23% TO 44% 'VERY WELL PREPARED'

Companies rating themselves "very well prepared" to defend against AI-powered attacks nearly doubled from 23% to 44% in 2025, while "moderately prepared" dropped from 52% in 2024 to 47%. Enterprises aren't gradually improving. They're jumping confidence tiers.

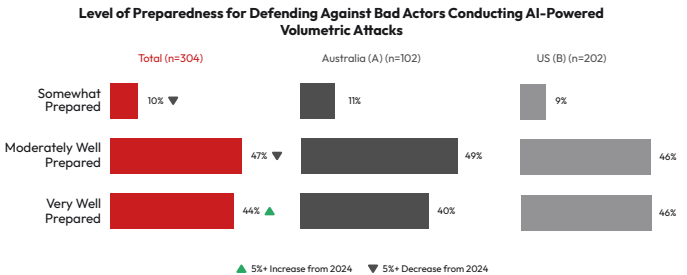

Q4a: How prepared would you rate your enterprise in terms of using AI to defend against bad actors using AI-powered bots to deploy volumetric attacks? (SELECT ONE)

Security teams have invested heavily—28% of budgets now fund AI solutions—and completed training programs that strengthen their capabilities. That deployment activity drives the confidence surge, and rightfully so. Building defenses requires investment first, results a very quick second.

The timing matters though. Confidence doubled while losses stayed at $10 million to $500+ million meaning companies measure success by what they've built rather than what they've stopped. As fraudsters scale AI-powered techniques to 33% of all attacks, security teams are deploying capabilities that drive up adversarial costs: adaptive bot mitigation that makes automated attacks expensive to execute and edge protection that stops threats before they consume resources or reach critical applications.

CAPABILITY BUILDING AFTER COMPROMISE

Enterprise AI security training investment reached 86% in 2025. GenAI threat acknowledgment rose from 56% in 2024 to 67% as enterprises encountered these attacks in production.

US enterprises invest in training at higher rates (92%) than Australian companies (75%), while also reporting more frequent GenAI attacks (72% versus 57%).

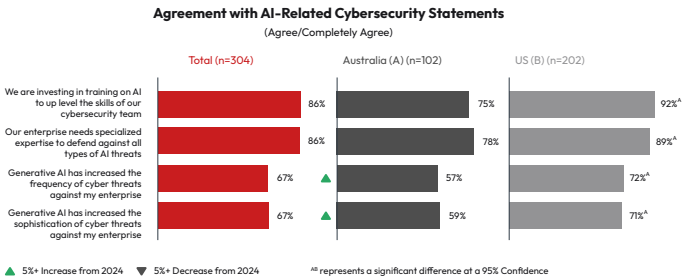

Q6. And how much do you agree with each of the following AI-related cybersecurity statements...

Typical training programs require six to 12 months before teams can apply new skills effectively. Current losses—55% of enterprises reporting $10 million to $500+ million—reflect attacks that succeeded before security teams completed training, not the capabilities they're building now.

This timing lag explains why confidence rose while losses remained flat. Teams invested in training and feel better prepared, but attackers targeted vulnerabilities during the learning period. US enterprises demonstrate this pattern most clearly: higher attack frequency drove higher training investment, yet losses haven't declined because trained teams are just now reaching operational readiness. The 2026 survey will show whether completed training programs translate to fewer successful breaches.

VENDOR CRITERIA OVERHAUL

Vendor priorities underwent a complete reordering in 12 months. Staying ahead of evolving threats surged 18 points from sixth to first position at 59%. Three capabilities now dominate: threat anticipation at 59%, industry-specific risk mitigation at 56% and AI-powered threat defense at 56%. This clustering reveals enterprises now evaluate vendors on adaptive capability rather than static feature sets.

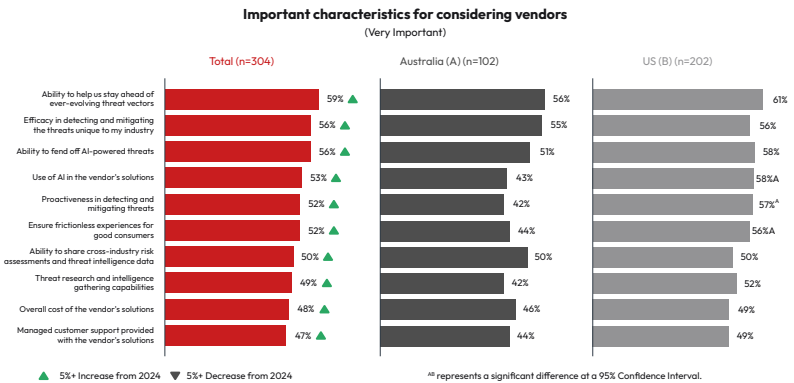

Q10. When considering a vendor to support your bot management device ID, API protection, etc. needs, how important are each of the following characteristics...

The 18-point shift signals enterprises fundamentally changed how they evaluate vendors. They switched from comparing features to assessing whether vendors can match adversary evolution speed. Every capability related to prediction and adaptation gained ground: threat research jumped 13 points, AI defense rose 11 points, proactive detection increased six points.

Cost rose 10 points to 48% but remains among the lowest priorities. When facing $10 million-plus annual losses, enterprises qualify vendors on effectiveness first, then negotiate price. The clustering at 56%-59% reveals companies now expect vendors to understand their specific attack landscape and adapt defenses continuously. Static solutions requiring manual updates or lengthy deployments no longer meet requirements, regardless of cost advantages.

> **When facing $10M+ in annual losses, effectiveness qualifies vendors first, then price gets negotiated.**

The Mobilization Year

What's at Stake

2025: the year AI security went mainstream. Enterprises deployed tools universally. Integration barriers fell. Yet only half see measurable results from what they built. The deployment-to-protection gap defines the turning point.

2025: THE YEAR ENTERPRISES GOT SERIOUS ABOUT AI SECURITY

AI security adoption surged uniformly across all capabilities in 2025. Faster incident response jumped 21 points to 92%, continuous monitoring climbed 26 points to 89%, real-time analysis rose 19 points to 88%, and predictive detection increased 18 points to 87%. Process automation and historical analysis gained significantly, clustering all six functions at 86-92%.

US enterprises lead adoption, particularly in faster incident response at 95% versus Australia's 85%.

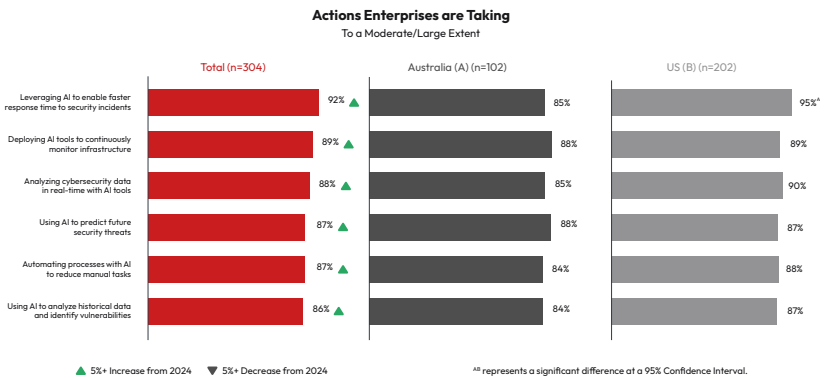

Q5a. To what extent is your enterprise taking the following actions...? (RATE EACH ROW)

The uniform surge signals 2025 as the year mainstream enterprises mobilized against AI-powered attacks. In the study last year, only a cohort we dubbed AI Enthusiasts—companies already under severe AI-powered assault—took action. Now the majority face these threats and are responding with comprehensive deployments treating all six capabilities as equally critical for defense.

The question is timing: did enterprises act fast enough? With agentic AI threats arriving within months, companies that started implementation in 2024 and deployed in 2025 are operationally ready. Those just beginning face a capability gap against adversaries already scaling autonomous attacks.

INTEGRATION BARRIERS FALL AS CONFIDENCE RISES

Two dramatic shifts reshaped AI security implementation in 2025. Buy-versus-build sentiment surged to 81%, up 19 points from 62% in 2024. Simultaneously, integration challenges plummeted to 38%, down 11 points from 49%. These movements connect. As enterprises solved integration friction, they gained confidence that purchased solutions deliver faster value than internal development.

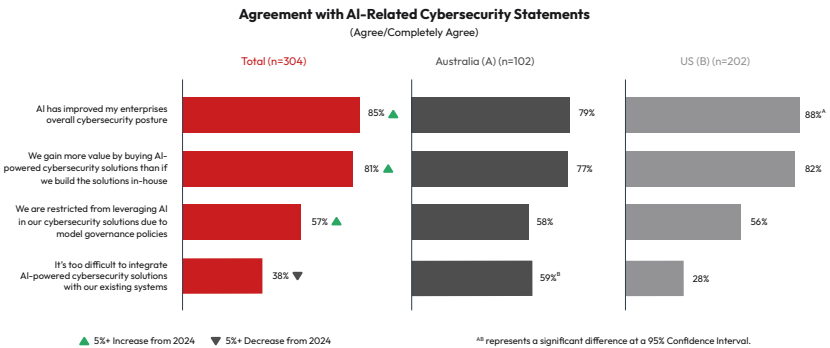

Q6. And how much do you agree with each of the following AI-related cybersecurity statements...(RATE EACH ROW)

Overall AI effectiveness confidence rose to 85%, up 17 points from 68% in 2024, validating the buy-side bet. The 19-point jump suggests security teams now view vendors as operational partners extending their capabilities rather than external products requiring extensive customization.

While most enterprises deploy smoothly, 59% of Australian firms still struggle versus 28% in the US, which may explain the 15-point benefit realization gaps in threat detection, web scraping and content theft and GenAI defense. Governance restrictions rose to 57%, up 13 points from 44% in 2024, as enterprises balance speed with control.

The integration breakthrough matters because vendor partnerships now determine defensive velocity. When integration difficulty drops 11 points and buy preference jumps 19 points in one year, the findings suggest that effective vendors accelerate capabilities faster than enterprises' internal teams can build them. That's critical when 53% expect new kinds of threats hitting them within six months fueled by agentic AI.

THE 50% REALITY: WHEN BENEFITS LAG ADOPTION

Benefit realization reached 50% in 2025: between 42%-51% of enterprises achieve measurable improvements from AI-powered solutions. Bot defense leads at 51%, up 6 points from 2024. Threat detection reached 47%, up 6 points. Phishing protection, API security and device identification cluster at 46%.

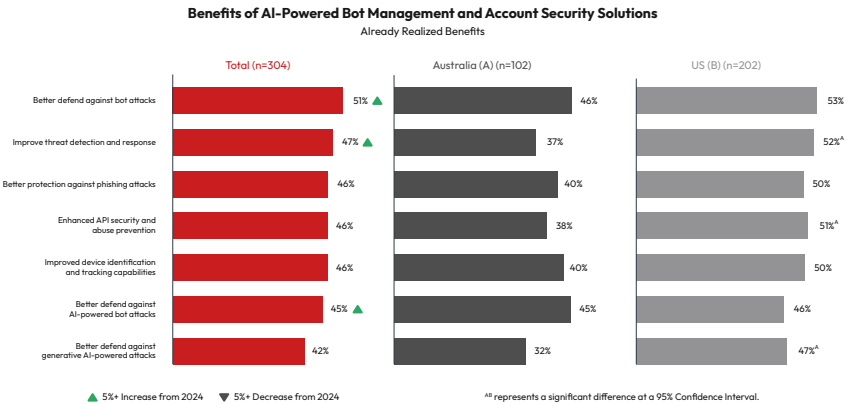

Q8. What benefits have you **realized** from AI-powered bot management and account security solutions?

Year-over-year progress shows the benefit range narrowed and rose: the floor lifted from 38% in 2024 to 42% in 2025, while the ceiling climbed from 46% to 51%. The 4 and 5 point improvements reveal AI security matured from experimental to reliable and is inching toward breakthrough transformation.

This creates a persistent divide. Half of enterprises with near-universal AI deployment (86%-92% page 16) now see concrete results. The other half deployed the same tools and are still translating adoption into outcomes. What separates them? Implementation rigor. The 38% facing integration challenges (page 17) fall predominantly into the struggling half.

The data quantifies the preparedness paradox. While 85% report improved posture (page 17) (up 17 points from 68% in 2024) only 45%-51% realize specific benefits. Security teams and vendors must work together on measurable outcomes: jointly tracking detection accuracy improvements, response time reductions and false positive decreases to close the deployment-to-protection gap.

EXPECTED BENEFITS CALIBRATE TO REALITY

Expected future benefits cluster at 48%-52%: human fraud farms (52%), attack automation services (51%), agentic AI attacks (51%) and generative AI attacks (51%). Enterprises view these emerging threats with equivalent concern.

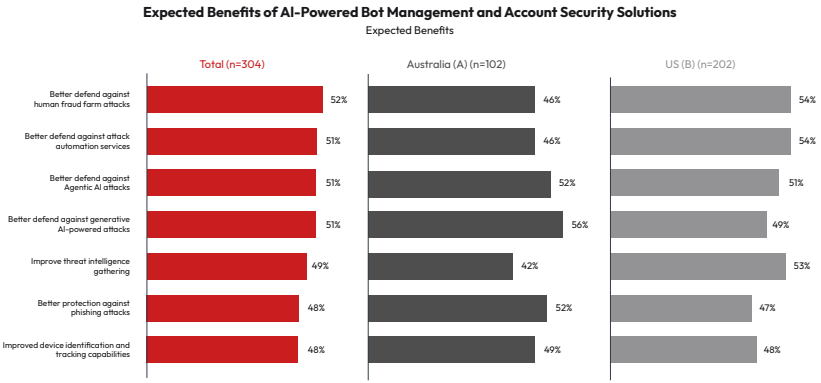

Q8a. What benefits do you **expect to realize** in the near future from AI-powered bot management and account security solutions?

Comparing 2024 expectations with 2025 reality reveals a critical insight: enterprises accurately predicted defensive capabilities but missed completely on cost reduction. They expected to defend against AI-powered bots at 45%—exactly matching 2025's realized 45%. They expected threat detection at 44% and achieved 47%. However, cost reduction led to expected benefits in 2024 at 46% yet vanished from 2025's top results. AI security adds capability was prioritized over spend.

The shift to human fraud farms as the top expected benefit (52%) reflects measurable progress: enterprises and their vendor partners got better at detecting and mitigating traditional bots, forcing attackers toward human fraud farms. This success creates the next challenge as agentic AI emerges to automate what fraud farms do manually.

Security teams are already prioritizing next-generation device intelligence and behavioral analysis that detect anomalous patterns regardless of attack mechanism. Deploying adaptive defenses rather than targeting specific threat types ensures protection as attackers evolve from bots to human fraud farms to agentic AI.

The Agentic AI Reckoning

What's at Stake

Most enterprises have less than six months before agentic AI attacks become critical business risks. Eighty-eight percent have already deployed the technology, but only half have built adequate defenses. The clock isn't just ticking, it's nearly run out.

WHEN AGENTIC AI BECOMES A CRITICAL RISK

Five percent of enterprises report agentic AI attacks are already critical, while 53% expect them within six months. Another 36% anticipate critical threats within 6-12 months, leaving just 6% seeing timelines beyond a year.

Estimated Duration Before Agentic AI-Powered Attacks Become Critical Risks

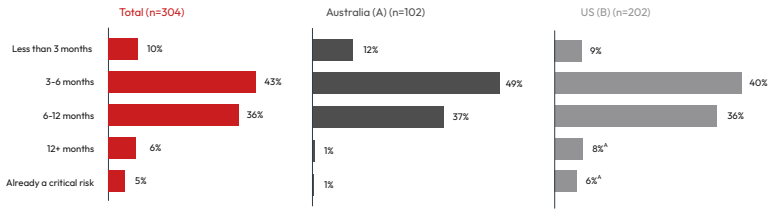

AB represents a significant difference at a 95% Confidence Interval.

Q13. How many months do you estimate your enterprise has before Agentic AI-powered attacks become a critical business risk?

Security teams are doing math that doesn't add up. Standard security transformations take 12-18 months from evaluation to deployment. The 53% expecting critical threats within six months face attacks arriving before defenses finish building.

The 43% clustering at 3-6 months aren't guessing. They're tracking adversary forums and attack tool releases, calculating when agentic AI becomes weaponized enough to cause damage. This timeline crunch explains why 88% of companies deployed agentic AI before determining what it's best at stopping. Security leaders made a calculated choice: deploy now with gaps or wait for comprehensive evaluation while adversaries gain months of advantage. The question wasn't between perfect and imperfect solutions, it was between imperfect protection now or excellent protection arriving after breaches succeed.

THE SPLIT-SCREEN REALITY: 86% TRUST PROVIDERS, 71% FEAR THE THREAT

Eighty-six percent of enterprises trust AI agent providers to prevent malicious use—29% completely, 57% mostly, 12% neutral. Yet 71% recognize agentic AI as the next major threat and 78% say attacks will be fundamentally different.

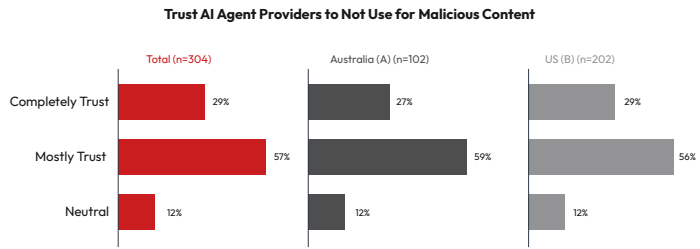

Q18b. How much do you trust AI agent provider(s) to ensure their agents are NOT used for malicious purposes?

The split reveals enterprises are answering two different questions as one. When security teams trust AI providers, they mean securing legitimate deployments: authentication, access controls, usage monitoring. That's reasonable for what providers can control.

What providers can't control is adversaries building their own agentic AI using open-source models, fine-tuning commercial APIs for attacks or developing autonomous fraud systems. The 86% trust in provider security doesn't protect against the 71% threat from adversarial AI.

Companies trust providers to secure the tools they deploy while building defenses against tools they don't control. The high trust enables adoption. The high threat recognition drives defensive investment. Both make sense when enterprises understand they're solving different problems—provider security for authorized agents, defensive capabilities for adversarial ones.

AUSTRALIA'S SPEED TRAP: FASTER ADOPTION, WEAKER DEFENSES

Eighty-eight percent of enterprises now use agentic AI for cybersecurity, but adoption outpaced protection. Fifty-two percent report comprehensive defenses against malicious agentic AI, while 36% admit they lack adequate defenses. Just 8% remain in evaluation.

Regional patterns diverge: 48% of Australian enterprises lack defenses versus 29% in the US. Yet 71% recognize bad actors using agentic AI as the next big threat, and 78% agree these attacks will be fundamentally different from traditional threats.

Current Enterprise's Agentic AI Reality

	Total (n=304)	Australia (A) (n=102)	US (B) (n=202)
We're actively using Agentic AI for cybersecurity/fraud prevention and have comprehensive defenses against malicious Agentic AI	52%	45%	55%
We're using Agentic AI for cybersecurity but lack adequate defenses against Agentic AI-powered threats	36%	**48%**[B]	29%
We're still evaluating the risks and benefits of Agentic AI in cybersecurity	8%	0%	**11%**[Atv]

Agreement with AI-Related Cybersecurity Statements

(Agree/Completely Agree)

	Total (n=304)	Australia (A) (n=102)	US (B) (n=202)
Bad actors use of Agentic AI to attack enterprises is the next big threat	71%	**73%**	70%
Agentic AI-powered attacks will pose a fundamentally different threat than traditional cyberattacks	78%	73%	81%

[AB] represents a significant difference at a 95% Confidence Interval. Only report result with 5% or above.

Q11. When it comes to Agentic AI, which statement best describes your enterprise's current reality?
Q6. And how much do you agree with each of the following AI-related cybersecurity statements...

The gap between adoption (88%) and adequate defense (52%) creates critical exposure. More than one-third of enterprises deployed agentic AI for offensive capabilities—faster response, predictive detection— without building corresponding defenses against adversaries using identical techniques.

This asymmetry explains the regional divide: Australian companies moved faster into adoption (0% still evaluating versus 11% US) but built fewer defenses. US enterprises adopted more cautiously, allowing time to develop protections. As 78% acknowledge agentic AI attacks will be fundamentally different, the 36% operating without adequate defenses face threats their current architecture wasn't designed to handle.

NO CLEAR WINNER: AGENTIC AI BENEFITS STILL UNDEFINED

When asked which potential benefit of agentic AI matters most, responses scatter widely. Faster threat response leads at just 23%, improved decentralized security reaches 20%, and task automation hits 18%.

Attack-type perception shows similar distribution: 39% cite fake account creation as where agentic AI excels, 30% point to GPT prompt compromise, and 29% identify account takeovers. Regional patterns hold consistent—US (41%) and Australia (36%) align on fake account prevention.

Top 3 Most Potential Benefits of Using Agentic AI

Top 3 Attack Types Agentic AI Is Best at Stopping

	Total (n=304)	Australia (A) (n=102)	US (B) (n=202)
Fake account creation/fake new accounts	39%	36%	41%
GPT prompt compromise	30%	28%	30%
Account takeovers (ATO)/credential stuffing	29%	25%	30%

Q14. Which potential benefit of using Agentic AI in your department is most important to your enterprise?
Q15. Which attack type do you think Agentic AI in your defense tech stack would be best at stopping?

The scattered responses reveal enterprises haven't identified what agentic AI uniquely solves. With no benefit exceeding 23% and no attack type exceeding 39%, security teams see potential across multiple use cases—response speed, decentralized protection, automation, fake account prevention—but lack conviction about where it delivers breakthrough value.

This distributed perception signals cautious experimentation rather than committed deployment. Companies are testing agentic AI in specific scenarios while evaluating whether results justify the architectural changes required. The even spread across benefits and attack types suggests the industry is still in discovery mode, building evidence for which applications warrant comprehensive adoption.

DOES AGENTIC AI MAKE SECURITY TEAMS BIGGER OR SMALLER?

Security executives are making hiring decisions right now, and they're splitting into two camps based on a single question: does agentic AI make security teams bigger or smaller?

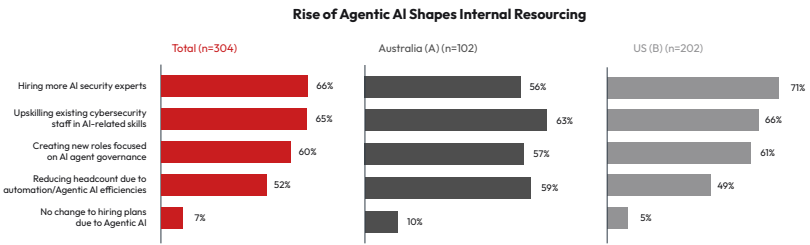

Q19: How is your enterprise adjusting hiring plans or internal resourcing in response to the rise of Agentic AI?

Geography reveals a philosophical divide. US enterprises pursue expansion with 71% hiring versus 49% reducing headcount. Australian CISOs take the opposite bet: 56% hiring but 59% also reducing, a balanced transformation. That's a 15-point hiring gap driven by different assumptions about whether AI agents multiply threat complexity or enable automation.

For leaders betting on threat sophistication outpacing automation, you're hiring specialists to address KYA capability gaps like the 25% of companies who can't verify agent authorization, the 23% who fail impersonation detection. For leaders betting on efficiency, you're automating operational tasks while selectively upskilling. But here's the problem: 93% are doing something, yet most haven't diagnosed which specific capability gaps require humans versus technology. Enterprises are staffing before strategizing.

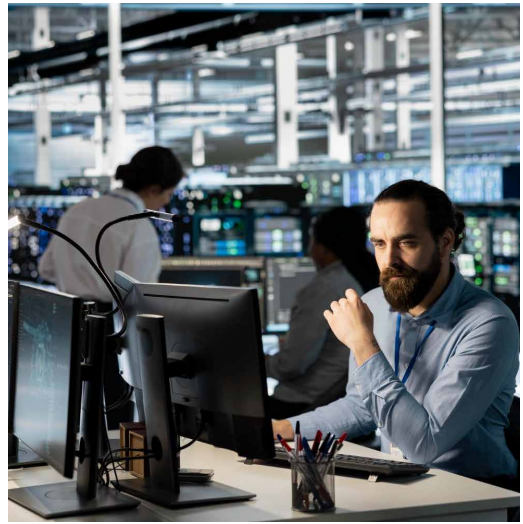

THE 2026 FORECAST: 93% EXPECT BREAKTHROUGH ATTACKS TO REWRITE DEFENSE STRATEGIES

Security professionals forecast competitive evolution through 2026. Ninety-three percent predict major agentic AI incidents will fundamentally change security approaches, 92% foresee open-source frameworks accelerating capabilities on both sides, and 88% anticipate defender improvements. Eighty-six percent expect sustained arms race dynamics, while 68% predict more dangerous attacks.

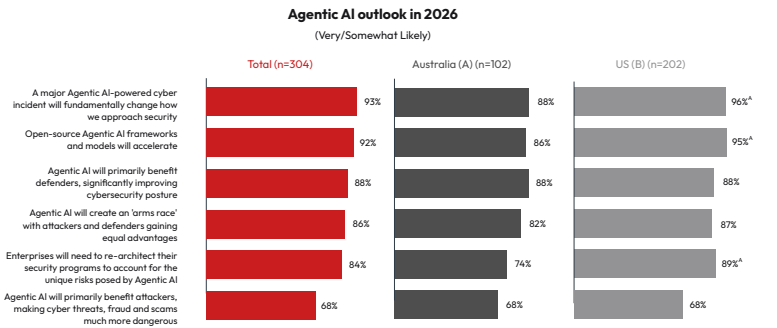

AB represents a significant difference at a 95% Confidence Interval.

Q12: Looking ahead to 2026, which scenario do you think is most likely regarding Agentic AI?

Regional predictions diverge: US respondents rate major incident likelihood at 96% versus 88% in Australia, with 89% expecting tech stack rearchitecture versus 74%—a 15-point gap reflecting higher perceived disruption in US markets.

The 93% forecasting fundamental change through major agentic AI incidents reflects how security innovation historically occurs—breakthrough attacks drive industry-wide learning and next-generation defenses. The balanced outlook (88% expecting defender improvements while 68% anticipate more dangerous attacks) shows neither side gains decisive advantage, creating sustained competition requiring continuous adaptation.

Budget allocation validates this forecast: AI's share of security spend currently sits at 28%, projected to reach 32% by 2026. This 4-point increase, identical across regions, demonstrates measured commitment, allocating one-third of cybersecurity resources toward the predicted threat landscape transformation.

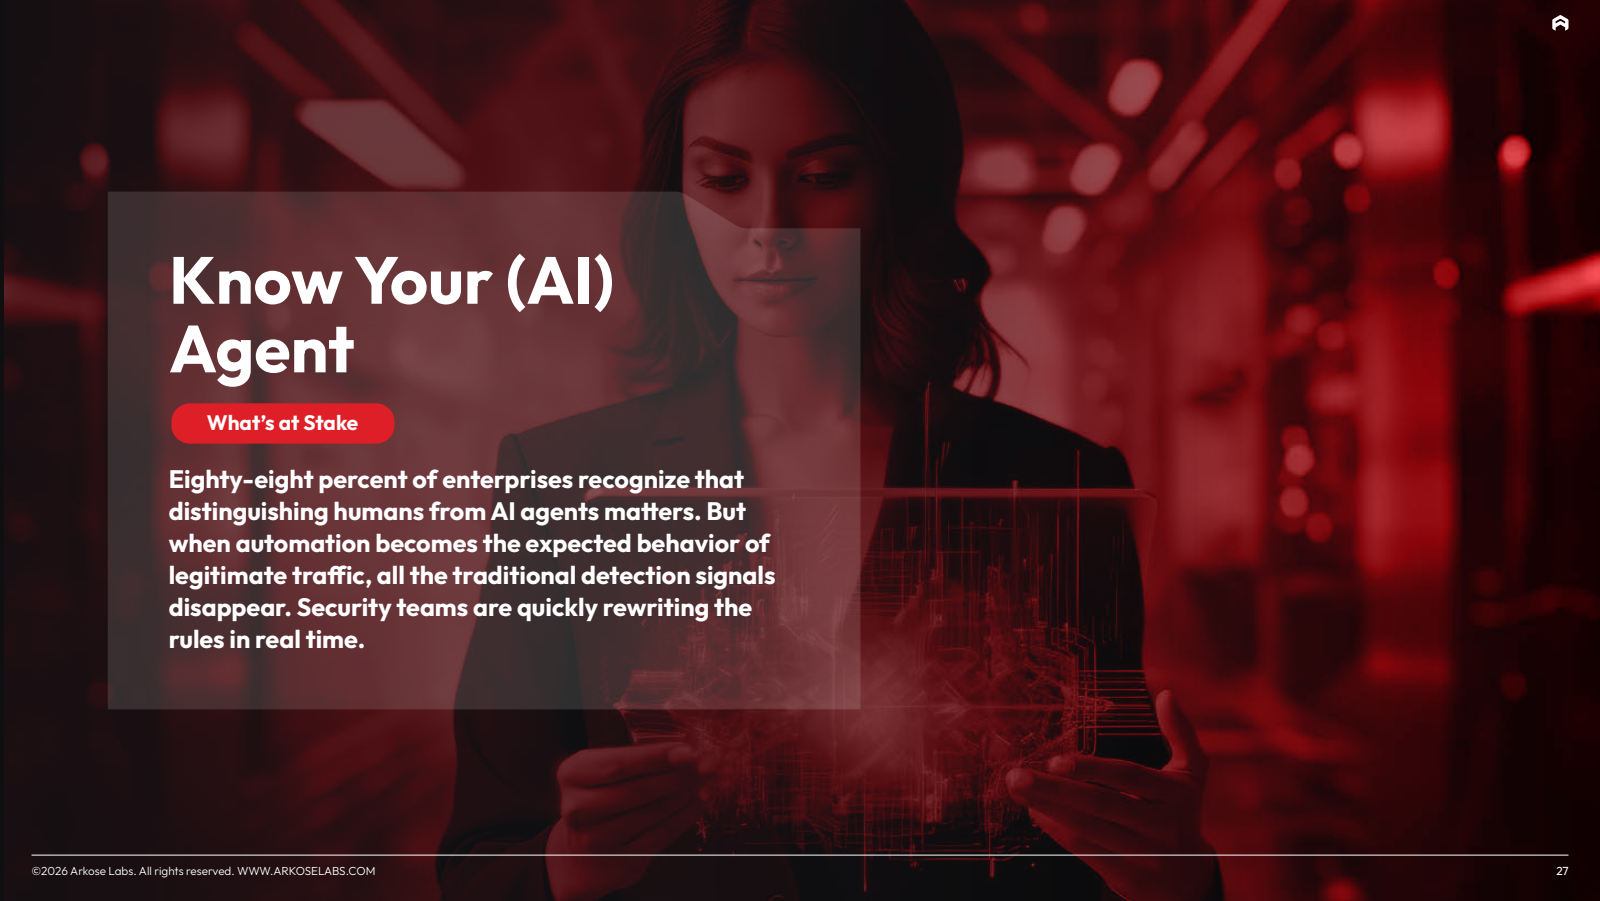

Know Your (AI) Agent

What's at Stake

Eighty-eight percent of enterprises recognize that distinguishing humans from AI agents matters. But when automation becomes the expected behavior of legitimate traffic, all the traditional detection signals disappear. Security teams are quickly rewriting the rules in real time.

AGENT IDENTIFICATION REQUIREMENTS

Eighty-eight percent of survey respondents recognize distinguishing human users from AI agents as important (34% critical, 54% important). This near-universal consensus establishes agent identification as an enterprise priority across industries and regions.

Authorization controls are split three ways: 39% require pre-authorization where humans approve agent capabilities upfront, 29% accept oversight with periodic review, 28% demand direct control with continuous monitoring.

Importance in Distinguishing Human Users and AI Agents

	Total (n=304)	Australia (A) (n=102)	US (B) (n=202)
Critical - we need to know definitively when an agent is acting vs. a human	34%	29%	36%
Important - we want visibility but may allow both types of interactions	54%	51%	56%

Level of Human Involvement when Customers Use AI Agents to Interact with Platforms

	Total (n=304)	Australia (A) (n=102)	US (B) (n=202)
Direct human control - a person must actively initiate and monitor each agent action	28%	21%	31%
Human authorization - a person must pre-approve what the agent can do, but doesn't need to monitor real-time	39%	**43%**	37%
Human oversight - agents can act autonomously within pre-set boundaries, with periodic human review	29%	31%	28%

Q23. How important is it for your enterprise to distinguish between human users and AI agents on your platforms?
Q24. When customers use AI agents to interact with your platforms, what level of human involvement do you require?

The nearly even split (39%/29%/28%) shows cybersecurity teams haven't converged on a single best approach yet, meaning if security teams are still deciding, they're not behind. Most companies (68%) accept that their consumers' AI agents need some autonomy to be useful, so the question becomes how to govern agent independence, not whether to allow it. Pre-authorization leads at 39% because it balances control and efficiency: approve what an agent can do upfront, then let it work without micromanaging. US companies feel more urgency (36% critical vs 29% Australia), suggesting American teams have tighter timelines to implement AI agent identification.

> **This near-universal consensus establishes agent identification as an enterprise priority across industries and regions.**

MALICIOUS AGENT DETECTION: WHY 71% ARE SEEKING NEW CAPABILITIES NOW

Ninety-eight percent of companies are confident in their current ability to distinguish legitimate from malicious agents (52% very confident), while 71% rate developing enhanced agent identity detection as extremely or very urgent. US enterprises show 99% confidence and 74% urgency, Australia shows 98% confidence and 66% urgency.

Both confidence and urgency run high across industries, showing security teams trust their current detection while prioritizing next-gen enhancements.

Confidence in Distinguishing Legitimate vs. Malicious AI Agents

Very somewhat / Confident

Total (n=304)	Australia (A) (n=102)	US (B) (n=202)
98%	98%	99%

Very Confident

Total (n=304)	Australia (A) (n=102)	US (B) (n=202)
52%	49%	53%

Urgency in Developing Capability that Distinguishing Legitimate vs. Malicious AI Agents

Extremely/Very Urgent

Total (n=304)	Australia (A) (n=102)	US (B) (n=202)
71%	66%	74%

Extremely Urgent

Total (n=304)	Australia (A) (n=102)	US (B) (n=202)
28%	23%	30%

Q16. How urgent is it for your enterprise to develop capabilities that can distinguish between legitimate consumer AI agents and malicious AI agents trying to appear as legitimate consumers?
Q17. How confident are you in your current security infrastructure's ability to distinguish between legitimate AI agents acting on behalf of consumers versus malicious AI agents trying to appear as legitimate consumers?

When 98% of security teams are confident in current capabilities and 71% call developing new capabilities urgent, this shows they are operating from a position of strength: confident enough in today's defenses to invest in tomorrow's without firefighting current failures. The 71% urgency isn't because detection is broken; it's because teams anticipate AI agent traffic will scale and malicious actors will become harder to distinguish from the legitimate AI agents of consumers, so companies want better capabilities in place before the challenge intensifies.

US teams prioritize this more strongly (74% urgent vs 66% Australia), likely seeing faster consumer AI agent adoption in American markets.

DETECTION DILEMMA: SPOTTING BAD AGENTS WHEN ALL AGENTS ACT NON-HUMAN

AI agent authentication is an active operational priority. Eighty percent of enterprises develop separate workflows for agent interactions, 78% monitor usage closely and 73% actively block suspected traffic. Challenges cluster tightly: 25% cite verifying legitimate agents are properly authorized, 23% struggle detecting malicious impersonation, and 20% face managing interaction scale.

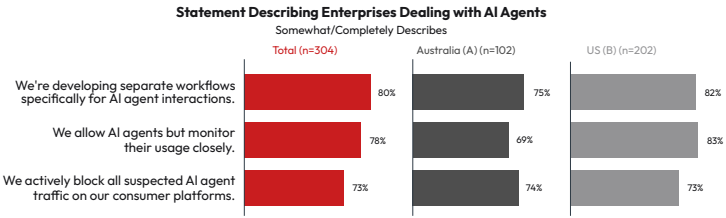

Know Your Agent (KYA) Greatest Challenge

	Total (n=304)	Australia (A) (n=102)	US (B) (n=202)
Verifying that legitimate AI agents are properly authorized by real users	25%	21%	27%
Detecting when malicious agents are impersonating legitimate customers	23%	30%[B]	20%
Managing the scale of AI agent interactions	20%	21%	20%
Ensuring compliance when AI agents act on behalf of users	17%	17%	17%
Distinguishing AI agent behavior from human behavior	14%	12%	16%

[AB] represents a significant difference at a 95% Confidence Interval.

Q18. Which aspect of AI agent authentication (also known as "Know Your Agent" [KYA]) poses the greatest challenge for your enterprise?
Q18c. To what degree does each statement describe your enterprise when dealing with AI agents on your platform?

Security teams face a problem that flips traditional bot defense upside down. Standard approaches flag automated behavior—speed, volume, linear patterns—as suspicious. AI agents are supposed to automate, creating a fundamental challenge: enterprises can't verify authenticity by spotting non-human patterns when agents are designed to be non-human, can't detect impersonation when legitimate and malicious agents act identically and can't manage scale through rate-limiting when customers expect machine-speed operations.

The 25%/23%/20% clustering shows these familiar challenges require different approaches when AI automation becomes an expected feature of legitimate traffic rather than a red flag for malicious activity.

Australian teams prioritize impersonation detection more (30% vs 20% US), likely reflecting the country's intense impersonation fraud environment where phishing losses tripled to AU$13.7 million[3] in early 2025, making agent impersonation an new layer of complexity on an already critical challenge.

The high engagement rates (80%/78%/73%) show enterprises aren't waiting to tackle AI agent authentication. Instead, they're proactively building workflows and monitoring now.

3. https://www.bankinfosecurity.com/australian-scam-losses-increase-28-in-2025-a-28580]

"Traditional detection looked for automation or suspicious fingerprints as the red flag. With AI agents, automation is expected and legitimate, so all the previous conventional detection signals are ineffective. We've had to fundamentally rethink what 'suspicious' means and look much closer at the intent of the agentic traffic."

– Mitch Davies, Senior Data Scientist, Arkose Labs.

AGENTIC AI INTERFACE DECISIONS: ONE PAIN POINT IS ENOUGH

Interface decisions trigger when specific pain points become acute. Performance optimization drives 35% of decisions, security concerns 28%, regulatory compliance 18%. No factor dominates, suggesting enterprises move forward when any single catalyst hits hard enough, not when all factors align.

Regional patterns show different catalysts. Australian enterprises moved primarily on performance pain, while US enterprises haven't yet encountered the forcing event that clarifies priorities.

Most Influential Factors Shaping Interface Approach for AI Agents

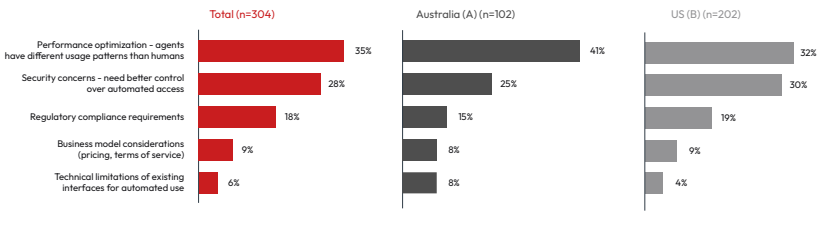

Q25. What would MOST influence your decision to create dedicated AI agent interfaces versus allowing agents to use existing user interfaces?

When consumers start to use AI agents on companies' websites or apps, security leadership has a tough decision: let agents use existing sites/apps (cheaper and faster but risky) or build dedicated agent sites/apps (expensive and slow but controlled). In the decision-making process, enterprises don't have to experience all three factors (performance, security, compliance) to justify building dedicated AI agent interfaces.

The 35%-28%-18% split reveals enterprises move when any single pain point hits hard enough. If enterprise teams are waiting for the perfect business case with all factors aligned, they may be overthinking it. For Australian teams, performance pain has clearly hit (41% vs 25% security). For US teams still evaluating (32% performance vs 30% security—virtually tied), they are likely waiting to see which breaks first: will AI agent traffic overwhelm existing interfaces (performance), will security incidents reveal gaps (security), or will regulations mandate separation (compliance)?

ENTERPRISES' OPINIONS ABOUT CONSUMER AGENTS

Opinions of consumer AI agents reveal a dramatic regional split: Australian enterprises show 18x higher negative rates than the US and lower positive rates (27% vs 34% US). Overall, 32% see net positive, 53% report mixed impact.

View of Customers Using AI Agents on Platforms

	Total (n=304)	Australia (A) (n=102)	US (B) (n=202)
Net positive for our business	32%	27%	34%
Net negative for our business	7%	18%[B]	1%
Mixed impact depending on context	53%	46%	56%
Too early to tell	9%	9%	9%

[AB] represents a significant difference at a 95% Confidence Interval. Only report result with 5% or above.

Consumer AI agents usher in the BYOA era—Bring Your Own Automation—where people instruct and authorize their own AI agents to do legitimate tasks on enterprise platforms rather than using only the tools enterprises provide. For example, a consumer's shopping agent browses an e-commerce site, compares products and completes purchases autonomously without the consumer clicking a single button.

How do companies view consumers bringing their own AI agents? The 53% seeing mixed impact are navigating where these agents create efficiency benefits in some contexts (account inquiries, routine transactions) but security risks in others (authorization abuse, impersonation).

The 32% seeing net positives have found ways to enable consumer automation while managing risks. Australia's dramatically higher negative opinion (18% vs 1%) aligns with the fact that Australian teams struggle more with agent impersonation detection (30% vs 20%). When enterprises cannot effectively distinguish legitimate consumer agents from malicious ones, the abuse and fraud risks of BYOA can outweigh the efficiency benefits, driving negative business perceptions.

The pendulum has shifted: Enterprises are scrambling to figure out how to accommodate automation tools consumers control, not just automation they control.

CONSUMER AI AGENTS: 72% ACCEPT WITH CONDITIONS, 24% BLOCK ENTIRELY

Major concerns split enterprises into BYOA (Bring Your Own Automation) accepters and rejectors. Twenty-four percent of enterprises want to block consumer agents entirely, while 72% accept them with varying control levels.

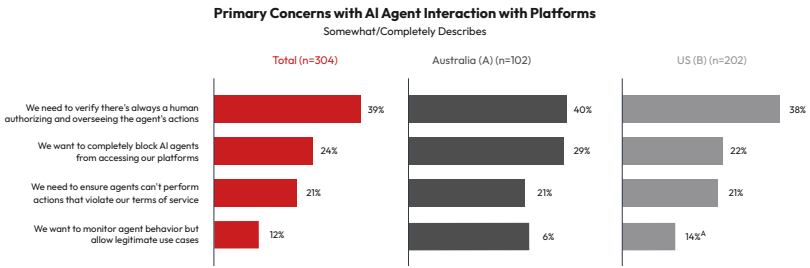

Q22. What is your primary concern when AI agents – acting on behalf of your customers – interact with your platforms? Please select your top concern.

The 72% accepting consumer AI agents are split into three camps based on risk tolerance.

-The 39% requiring verification for every action prioritize control over convenience because they're protecting high-value transactions where a single unauthorized action could mean significant financial or regulatory exposure, like an AI agent taking over a consumer's checking account.

-The 21% setting behavioral boundaries take a middle path: agents operate autonomously unless they trigger fraud signals or violate terms. For example, an OTA might let a consumer's AI agent search flights, book hotels and modify reservations freely, but flag the transaction if the agent tries to book last-minute international flights with mismatched passenger details or suddenly adds ten loyalty program accounts, allowing useful automation while preventing loyalty fraud or stolen credential use.

-If the company is a content platform or service with lower per-transaction risk, the 12% monitoring approach lets them accommodate innovation without burning resources on verification workflows, which drives business efficiencies.

The 24% who want to block agents entirely have decided the risk outweighs any consumer convenience benefit. A wealth management firm can risk a compromised agent executing unauthorized trades or accessing client portfolios which could trigger SEC violations and fiduciary liability that dwarf any efficiency gains.

Australia's higher blocking rate (29% vs 22% US) reflects genuine business impact concerns with 18% (compared to 1% US) expecting net negative outcomes from consumer agents. And that justifies rejection over accommodation.

WHY AGENT THREATS LEAD ALL AI SECURITY CONCERNS

Enterprises appear divided across six AI threat categories, but the data reveals a strategic contradiction. Seventy-two percent accept consumer AI agents on their platforms, yet when asked about TODAY's biggest threat, concerns split across generative AI (20%), AI-enhanced traditional attacks (19%), and three types of agent threats totaling 45% combined.

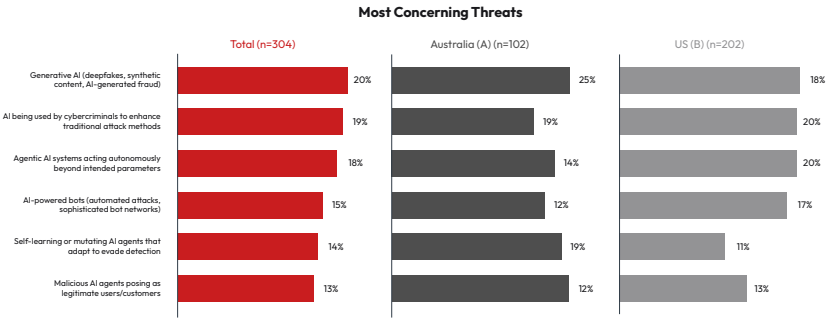

Q19a. Considering the various AI-based threats facing your enterprise, which one causes you the most concern due to its potential to significantly harm your enterprise today?

This isn't confusion. It's triage. Security teams defend against active damage today (deepfakes costing money now, AI making credential stuffing faster) while preparing for tomorrow's dominant threat. Australia's pattern offers a preview: 25% face generative AI fraud as their top concern today, while 19% are concerned about self-learning agents that adapt to evade detection, suggesting agent sophistication escalates quickly once autonomous systems establish themselves.

Enterprise security teams are building the airplane as they fly it. They have made a business decision to accept customer agents before finalizing their security approach to distinguish legitimate from malicious ones. The 98% confidence revealed earlier [Part 2] reflects defenses built for scripted bots; the 71% urgency acknowledges [Part 6] the need to upgrade to agent detection before the 45% who fear agentic threats are proven right.

Security teams should prioritize agent detection infrastructure now before the gap between deployment and verification becomes their largest vulnerability.

Airline Industry: High Confidence Meets Current Crisis

Airlines show the highest confidence yet face the harshest reality: 16% report agentic AI attacks already a critical risk, triple the average. "Hotels and airlines get hit harder by attackers because reselling airline tickets, loyalty points and room nights is easier than hacking bank accounts—and the federal consequences are lower," says Chris Staab, co-founder of Loyalty Security Alliance (LSA).

Integration difficulty sits at 28%, lowest of any sector. Airlines deploy tools fast despite the legacy systems they have in place. Budgets jump from 29% to 36%, steepest increase. The scorecards reveal why confidence meets

AIRLINE INDUSTRY AI MATURITY SCORECARDS

Airlines show the highest confidence (60% very well prepared versus 44% cross-industry) because integration friction stays low at 28%. "Most airlines still run legacy systems from years ago because they prioritize butts in seats and heads on beds. That's where the spend goes," notes Chris Staab of LSA. Yet somehow airlines deploy security tools faster than sectors with modern tech stacks, suggesting operational discipline compensates for aging infrastructure.

Evolving Threats

Level of Concern from Each Attack Type

Q1-moderate/large Extent	Total (n=304)	Airlines (n=50)
Man-in-the-Middle reverse proxy phishing	73%	62%
Account takeovers (ATO)/credential stuffing	70%	62%
Fake account creation/fake new accounts	73%	56%

Approximate Cost of Negative Consequences

Q3a	Total (n=304)	Airlines (n=24)
Less than $1M	8%	8%
$1M to $9.9M	38%	38%
$10M to $99M	39%	34%
$100M to $500M	11%	16%
Over $500M	5%	4%

Preparedness Rising

Approximate % of Cybersecurity Budget Spent on Security Solutions

Q5c,Q5d-average	Total (n=304)	Airlines (n=50)
Today	28%	29%
12 months from now	32%	36%

Level of Preparedness for Defending Against AI-Powered Volumetric Attacks

Q4a	Total (n=304)	Airlines (n=50)
Moderately well prepared	47%	36%
Very well prepared	44%	60%

Current Approaches to Learning About Agentic AI Opportunities or Risks

Q20 - Top 3	Total (n=304)	Airlines (n=50)
Conducting internal research and testing	64%	80%
Engaging with vendors or consultants to assess Agentic AI	68%	72%
Attending industry events or webinars focused on Agentic AI	59%	66%

Adoption & Benefits

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total (n=304)	Airlines (n=50)
AI has improved my enterprises overall cybersecurity posture	85%	84%
We gain more value by buying AI-powered cybersecurity solutions than if we build the solutions in-house	81%	78%
We are restricted from leveraging AI in our cybersecurity solutions due to model governance policies	57%	66%
It's too difficult to integrate AI-powered cybersecurity solutions with our existing systems	38%	28%

Benefits of AI-Powered Bot Management and Account Security Solution

Q8	Total (n=304)	Airlines (n=50)
Better defend against AI-powered bot attacks	45%	60%
Better defend against human fraud farm attacks	41%	54%
Better defend against Agentic AI attacks	41%	54%
Enhanced API security and abuse prevention	46%	50%

"Of all industries, airlines are most accurate about when agentic AI becomes critical, as 16% say it's already here because they're furthest along letting consumers transact through chatbots and have probably seen fraud in that channel," notes Chris Staab of LSA. This early exposure explains why 72% treat agent traffic the same as human traffic: airlines aren't building separate interfaces, they're letting consumer agents use existing booking systems, loyalty portals and customer service channels. This approach keeps integration friction low but creates scale risk. Of the airlines surveyed, 30% cite managing agent interaction volume as their greatest challenge.

The 30% reflects this early exposure: millions of booking transactions where consumer agents could overwhelm infrastructure.

Agentic AI Looms Large

Current Enterprise's Agentic AI Reality

Q11	Total (n=304)	Airlines (n=50)
We're actively using Agentic AI for cybersecurity/ fraud prevention and have comprehensive defenses against malicious Agentic AI	52%	50%
We're using Agentic AI for cybersecurity but lack adequate defenses against Agentic AI-powered threats	36%	32%
We're still evaluating the risks and benefits of Agentic AI in cybersecurity	8%	14%

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total (n=304)	Airlines (n=50)
Bad actors use of Agentic AI to attack enterprises is the next big threat	71%	72%
Agentic AI-powered attacks will pose a fundamentally different threat than traditional cyberattacks	78%	68%

Agentic AI Looms Large

Estimated Duration Before Agentic AI-Powered Attacks Become a Critical Risk

Q13	Total (n=304)	Airlines (n=50)
Less than 3 months	10%	16%
3-6 months	43%	24%
6-12 months	36%	32%
12+ months	6%	12%
Already a critical risk	5%	16%

Top 3 Most Important Potential Benefits of Using Agentic AI

Q14	Total (n=304)	Airlines (n=50)
Faster response to cyber threats	23%	26%
Automation of repetitive security tasks	18%	14%
Improved security for decentralized or remote environments	20%	14%

Top 3 Attack Types Agentic AI Is Best at Stopping

Q15	Total (n=304)	Airlines (n=50)
Fake account creation/fake new accounts	39%	34%
Account takeovers (ATO)/credential stuffing	29%	30%
Inventory hoarding	21%	30%

Know Your Agent (KYA)

Statement Describing Enterprises Dealing with AI Agents

Q18c-Somewhat/Completely Describe	Total (n=304)	Airlines (n=50)
We're developing separate workflows specifically for AI agent interactions.	80%	78%
We allow AI agents but monitor their usage closely.	78%	76%
We treat AI agent traffic the same as human user traffic.	60%	72%

Know Your Agent (KYA) Greatest Challenge

Q18	Total (n=304)	Airlines (n=50)
Managing the scale of AI agent interactions	20%	30%
Detecting when malicious agents are impersonating legitimate customers	23%	22%
Verifying that legitimate AI agents are properly authorized by real users	25%	20%

View of Customers Using AI Agents on Platform

Q18f	Total (n=304)	Airlines (n=50)
Net positive for our business	32%	38%
Net negative for our business	7%	4%
Mixed impact depending on context	53%	58%
Too early to tell	9%	0%

Banking on Defense: Why 85% Choose to Buy Rather Than Build

A majority of banks say AI improved cybersecurity. Only a few feel prepared for AI-powered attacks. That gap is the industry's defining vulnerability. The timeline makes it dire. A majority expect agentic AI attacks to become critical within 12 months. Security transformations can take 12-18 months. Banks are already behind schedule, so 85% choose to buy rather than build. When adversaries move at AI speed, internal development struggles to keep pace.

Budgets reflect the urgency: 28% of cybersecurity spending funds AI solutions today. Every percentage point matters. A single compromised account can move millions through laundering networks. One API breach exposes consumer data at enterprise scale. The scorecards ahead show where banking security leaders are placing their bets and reveal the gaps.

BANKING INDUSTRY AI MATURITY SCORECARDS

Pay close attention to the preparedness gap: while budgets are rising and AI adoption is strong, only 38% rate themselves as very well prepared for AI-powered volumetric attacks—6 points below the cross-industry average of 44%. This suggests banks may be investing in AI tools without fully operationalizing them for today's threat landscape. The 75% vendor engagement rate shows banks recognize they need external expertise, but the lower preparedness scores indicate deployment and integration remain works in progress.

Evolving Threats

Level of Concern from Each Attack Type

Q1-moderate/large Extent	Total (n=304)	Banking (n=80)
Fake account creation/fake new accounts	73%	79%
Man-in-the-Middle reverse proxy phishing	73%	79%
API abuse	70%	75%

Approximate Cost of Negative Consequences

Q3a	Total (n=304)	Banking (n=80)
Less than $1M	8%	8%
$1M to $9.9M	38%	34%
$10M to $99M	39%	43%
$100M to $500M	11%	10%
Over $500M	5%	6%

Preparedness Rising

Approximate % of Cybersecurity Budget Spent on Security Solutions Leveraging AI

Q5c,Q5d-average	Total (n=304)	Banking (n=80)
Today	28%	28%
12 months from now	32%	32%

Level of Preparedness for Defending Against AI-Powered Volumetric Attacks

Q4a	Total (n=304)	Banking (n=80)
Moderately well prepared	47%	50%
Very well prepared	44%	38%

Current Approaches to Learning About Agentic AI Opportunities or Risks

Q20 - Top 3	Total (n=304)	Banking (n=80)
Engaging with vendors or consultants to assess Agentic AI	68%	75%
Conducting internal research and testing	64%	59%
Following government/regulatory guidance	49%	56%

Adoption & Benefits

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total (n=304)	Banking (n=80)
AI has improved my enterprises overall cybersecurity posture	85%	90%
We gain more value by buying AI-powered cybersecurity solutions than if we build the solutions in-house	81%	85%
We are restricted from leveraging AI in our cybersecurity solutions due to model governance policies	57%	50%
It's too difficult to integrate AI-powered cybersecurity solutions with our existing systems	38%	40%

Benefits of AI-Powered Bot Management and Account Security Solutions

Q8	Total (n=304)	Banking (n=80)
Improve threat intelligence gathering	42%	50%
Improve threat detection and response	47%	46%
Better defend against bot attacks	51%	46%
Enhanced API security and abuse prevention	46%	43%

Banks are building the airplane while flying it—with 86% deploying consumer agent workflows as they rapidly solve how to verify legitimate agents or catch impersonators. And only 19% view consumer-controlled agents a net positive for business. Decades of security practice flipped: capabilities launched before defenses fully built out. Every transaction at machine speed runs without reliable fraud controls, and the timeline to fix this keeps shrinking.

> **"We've observed that enterprises are struggling to differentiate between the legitimate AI automation/agents and fraud farms running the same tech stack."**
> **– Sri Alapati, Threat Researcher, Arkose Labs**

Agentic AI Looms Large

Current Enterprise's Agentic AI Reality

Q11	Total (n=304)	Banking (n=80)
We're actively using Agentic AI for cybersecurity/ fraud prevention and have comprehensive defenses against malicious Agentic AI	52%	51%
We're using Agentic AI for cybersecurity but lack adequate defenses against Agentic AI-powered threats	36%	40%
We're still evaluating the risks and benefits of Agentic AI in cybersecurity	8%	8%

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total (n=304)	Banking (n=80)
Bad actors use of Agentic AI to attack enterprises is the next big threat	71%	76%
Agentic AI-powered attacks will pose a fundamentally different threat than traditional cyberattacks	78%	84%

Agentic AI Looms Large

Estimated Duration Before Agentic AI-Powered Attacks Become a Critical Risk

Q13	Total (n=304)	Banking (n=80)
Less than 3 months	10%	4%
3-6 months	43%	43%
6-12 months	36%	45%
12+ months	6%	4%
Already a critical risk	5%	5%

Top 3 Most Important Potential Benefits of Using Agentic AI

Q14	Total (n=304)	Banking (n=80)
Automation of repetitive security tasks	18%	20%
Improved security for decentralized or remote environments	20%	20%
Faster response to cyber threats	23%	19%

Top 3 Attack Types Agentic AI Is Best at Stopping

Q15	Total (n=304)	Banking (n=80)
GPT prompt compromise	30%	38%
Fake accounts creation / fake new accounts	39%	35%
Account takeovers (ATO)/credential stuffing	29%	33%

Know Your Agent (KYA)

Statement Describing Enterprises Dealing with AI Agents

Q18c-Somewhat/Completely Describe	Total (n=304)	Banking (n=80)
We're developing separate workflows specifically for AI agent interactions.	80%	86%
We allow AI agents but monitor their usage closely.	78%	85%
We actively blocked all suspected AI agent traffic on our consumer platforms.	73%	78%

Know Your Agent (KYA) Greatest Challenge

Q18	Total (n=304)	Banking (n=80)
Verifying that legitimate AI agents are properly authorized by real users	25%	26%
Detecting when malicious agents are impersonating legitimate customers	23%	24%
Managing the scale of AI agent interactions	20%	24%

View of Customers Using AI Agents on Platform

Q18f	Total (n=304)	Banking (n=80)
Net positive for our business	32%	19%
Net negative for our business	7%	8%
Mixed impact depending on context	53%	60%
Too early to tell	9%	14%

The Fintech Automation Tradeoff: Speed or Security First?

Fintech executives face a binary choice: deploy AI agents to match competitor automation and risk prompt injection attacks, or delay until defenses mature and surrender market share. Zero fintechs view consumer AI agents as net negative—fintech stands alone in unanimous acceptance—because payment platforms, digital lenders and robo-advisors that don't automate lose consumers within quarters.

Seventy-nine percent identify GPT prompt compromise as a major concern (highest of any industry). Deployment splits evenly between those that operate agentic AI with and without adequate defenses. Twenty-five percent struggle ensuring regulatory compliance when agents act on a consumer's behalf, where prompt manipulation could approve fraudulent loans, leak account data or bypass KYC controls.

FINTECH INDUSTRY AI MATURITY SCORECARDS

Fintech reports the highest GPT prompt compromise concern at 79% because AI systems directly control lending approvals, fraud detection and transaction processing—any of which could be manipulated through prompt injection to bypass controls and trigger regulatory violations. Despite 8% losing over $500M annually, only 38% feel very well prepared (below 44% benchmark), suggesting even catastrophic losses aren't driving readiness. At the same time, 25% expect agentic AI to become a critical risk factor within 3 months (vs 10% cross-industry), compressing response timelines.

Evolving Threats

Level of Concern from Each Attack Type

Q1-moderate/large Extent	Total (n=304)	Fintech (n=24)
GPT prompt compromise	66%	79%
Man-in-the-Middle reverse proxy phishing	73%	75%
Fake account creation/fake new accounts	73%	71%

Approximate Cost of Negative Consequences

Q3a	Total (n=304)	Fintech (n=24)
Less than $1M	8%	13%
$1M to $9.9M	38%	33%
$10M to $99M	39%	42%
$100M to $500M	11%	4%
Over $500M	5%	8%

Preparedness Rising

Approximate % of Cybersecurity Budget Spent on Security Solutions Leveraging AI

Q5c,Q5d-average	Total (n=304)	Fintech (n=24)
Today	28%	29%
12 months from now	32%	31%

Level of Preparedness for Defending Against AI-Powered Volumetric Attacks

Q4a	Total (n=304)	Fintech (n=24)
Moderately well prepared	47%	54%
Very well prepared	44%	38%

Current Approaches to Learning About Agentic AI Opportunities or Risks

Q20 - Top 3	Total (n=304)	Fintech (n=24)
Engaging with vendors or consultants to assess Agentic AI	68%	67%
Conducting internal research and testing	64%	67%
Attending industry events or webinars focused on Agentic AI	59%	50%

Adoption & Benefits

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total (n=304)	Fintech (n=24)
AI has improved my enterprises overall cybersecurity posture	85%	79%
We gain more value by buying AI-powered cybersecurity solutions than if we build the solutions in-house	81%	83%
We are restricted from leveraging AI in our cybersecurity solutions due to model governance policies	57%	50%
It's too difficult to integrate AI-powered cybersecurity solutions with our existing systems	38%	33%

Benefits of AI-Powered Bot Management and Account Security Solutions

Q8	Total (n=304)	Fintech (n=24)
Better defend against bot attacks	51%	67%
Better defend against AI-powered bot attacks	45%	63%
Improved device identification and tracking capabilities	46%	63%
Improve threat detection and response	47%	54%

Zero fintech firms view consumer-controlled AI agents as a net negative to their business. It's the only industry with unanimous acceptance, yet deployment splits evenly between those fintechs with adequate defenses (46%) and those without (46%). Twenty-nine percent struggle detecting malicious agent impersonation (vs 23%), while 25% face compliance challenges when agents act for consumers, which creates regulatory exposure where prompt injection could trigger both SEC violations and consumer trust erosion that ends fintech brands permanently.

Agentic AI Looms Large

Current Enterprise's Agentic AI Reality

Q11	Total (n=304)	Fintech (n=24)
We're actively using Agentic AI for cybersecurity/ fraud prevention and have comprehensive defenses against malicious Agentic AI	52%	46%
We're using Agentic AI for cybersecurity but lack adequate defenses against Agentic AI-powered threats	36%	46%
We're still evaluating the risks and benefits of Agentic AI in cybersecurity	8%	8%

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total (n=304)	Fintech (n=24)
Bad actors use of Agentic AI to attack enterprises is the next big threat	71%	67%
Agentic AI-powered attacks will pose a fundamentally different threat than traditional cyberattacks	78%	79%

Agentic AI Looms Large

Estimated Duration Before Agentic AI-Powered Attacks Become a Critical Risk

Q13	Total (n=304)	Fintech (n=24)
Less than 3 months	10%	25%
3-6 months	43%	42%
6-12 months	36%	25%
12+ months	6%	8%
Already a critical risk	5%	0%

Top 3 Most Important Potential Benefits of Using Agentic AI

Q14	Total (n=304)	Fintech (n=24)
Faster response to cyber threats	23%	29%
Automation of repetitive security tasks	18%	25%
Improved security for decentralized or remote environments	20%	25%

Top 3 Attack Types Agentic AI Is Best at Stopping

Q15	Total (n=304)	Fintech (n=24)
Fake account creation/fake new accounts	39%	46%
Account takeovers (ATO)/credential stuffing	29%	29%
GPT prompt compromise	30%	25%

Know Your Agent (KYA)

Statement Describing Enterprises Dealing with AI Agents

Q18c-Somewhat/Completely Describe	Total (n=304)	Fintech (n=24)
We allow AI agents but monitor their usage closely.	78%	79%
We're developing separate workflows specifically for AI agent interactions.	80%	79%
We actively block all suspected AI agent traffic on our consumer platforms.	73%	67%

Know Your Agent (KYA) Greatest Challenge

Q18	Total (n=304)	Fintech (n=24)
Detecting when malicious agents are impersonating legitimate customers	23%	29%
Ensuring compliance when AI agents act on behalf of users	17%	25%
Verifying that legitimate AI agents are properly authorized by real users	25%	17%

View of Customers Using AI Agents on Platform

Q18f	Total (n=304)	Fintech (n=24)
Net positive for our business	32%	33%
Net negative for our business	7%	0%
Mixed impact depending on context	53%	58%
Too early to tell	9%	4%

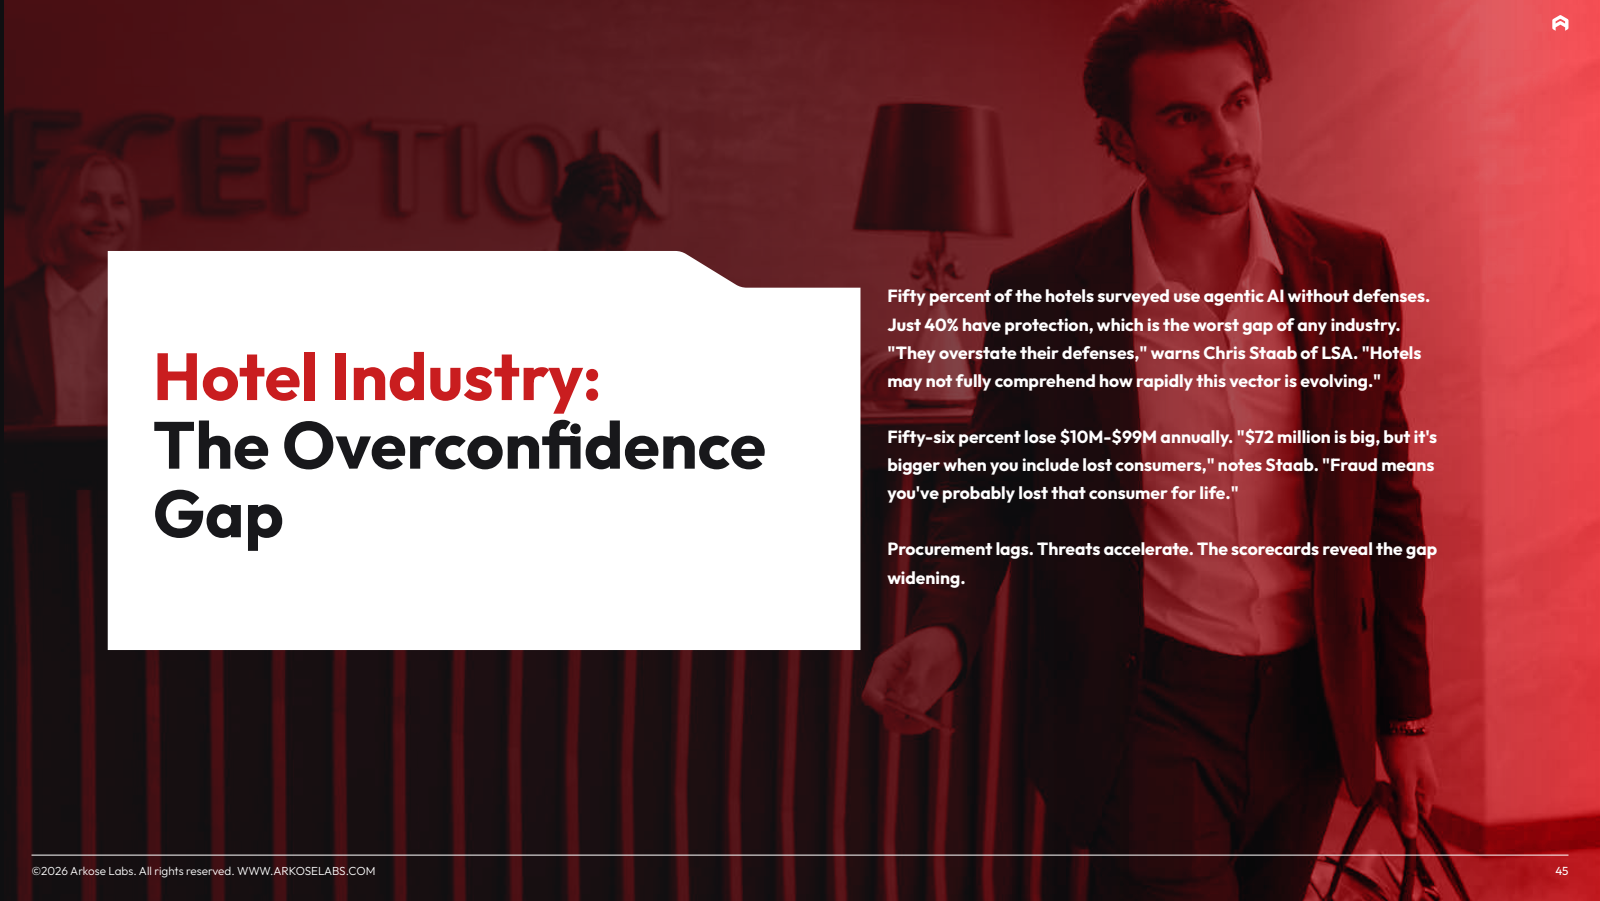

Hotel Industry: The Overconfidence Gap

Fifty percent of the hotels surveyed use agentic AI without defenses. Just 40% have protection, which is the worst gap of any industry. "They overstate their defenses," warns Chris Staab of LSA. "Hotels may not fully comprehend how rapidly this vector is evolving."

Fifty-six percent lose $10M-$99M annually. "$72 million is big, but it's bigger when you include lost consumers," notes Staab. "Fraud means you've probably lost that consumer for life."

Procurement lags. Threats accelerate. The scorecards reveal the gap widening.

HOTEL INDUSTRY AI MATURITY SCORECARDS

Hotels show the lowest confidence that AI improved security—just 68% versus 85% cross-industry—and the data validates this: only 40% have comprehensive agentic AI defenses, worst rate of any industry. "Hotels don't have the safety threat that drives airline cyber investment," explains Chris Staab of LSA. "Airlines must secure critical systems to operate aircraft. Hotels don't face that pressure." This motivation gap creates the preparedness gap.

Evolving Threats

Level of Concern from Each Attack Type

Q1-moderate/large Extent	Total (n=304)	Hotels (n=50)
Man-in-the-Middle reverse proxy phishing	73%	72%
Fake account creation/fake new accounts	73%	70%
Inventory hoarding	68%	70%

Approximate Cost of Negative Consequences

Q3a	Total (n=304)	Hotels (n=50)
Less than $1M	8%	2%
$1M to $9.9M	38%	24%
$10M to $99M	39%	56%
$100M to $500M	11%	16%
Over $500M	5%	2%

Preparedness Rising

Approximate % of Cybersecurity Budget Spent on Security Solutions Leveraging AI

Q5c,Q5d-average	Total (n=304)	Hotels (n=50)
Today	28%	28%
12 months from now	32%	33%

Level of Preparedness for Defending Against AI-Powered Volumetric Attacks

Q4a	Total (n=304)	Hotels (n=50)
Moderately well prepared	47%	52%
Very well prepared	44%	34%

Current Approaches to Learning About Agentic AI Opportunities or Risks

Q20 - Top 3	Total (n=304)	Hotels (n=50)
Engaging with vendors or consultants to assess Agentic AI	68%	62%
Attending industry events or webinars focused on Agentic AI	59%	58%
Conducting internal research and testing	64%	54%

Adoption & Benefits

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total (n=304)	Hotels (n=50)
AI has improved my enterprises overall cybersecurity posture	85%	68%
We gain more value by buying AI-powered cybersecurity solutions than if we build the solutions in-house	81%	74%
We are restricted from leveraging AI in our cybersecurity solutions due to model governance policies	57%	52%
It's too difficult to integrate AI-powered cybersecurity solutions with our existing systems	38%	42%

Benefits of AI-Powered Bot Management and Account Security Solutions

Q8	Total (n=304)	Hotels (n=50)
Better protection against phishing attacks	46%	54%
Improved device identification and tracking capabilities	46%	52%
Better defend against bot attacks	51%	48%
Better defend against attack automation services	39%	46%

Hotels deploy agentic AI with the worst defensive gap (50% lack adequate defenses) yet show the lowest threat recognition at just 64% believing attacks will be fundamentally different. "Hotels need to get ahead of this now because the procurement process is so slow. You've got to get the most modern tools in place now," urges Chris Staab of LSA. Slow procurement meets compressed timeline: 48% expect attacks critical within 3-6 months, but tools selected today may arrive after threats materialize.

Agentic AI Looms Large

Current Enterprise's Agentic AI Reality

Q11	Total (n=304)	Hotels (n=50)
We're actively using Agentic AI for cybersecurity/ fraud prevention and have comprehensive defenses against malicious Agentic AI	52%	40%
We're using Agentic AI for cybersecurity but lack adequate defenses against Agentic AI-powered threats	36%	50%
We're still evaluating the risks and benefits of Agentic AI in cybersecurity	8%	2%

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total	Hotels (n=50)
Bad actors use of Agentic AI to attack enterprises is the next big threat	71%	70%
Agentic AI-powered attacks will pose a fundamentally different threat than traditional cyberattacks	78%	64%

Agentic AI Looms Large

Estimated Duration Before Agentic AI-Powered Attacks Become a Critical Risk

Q13	Total (n=304)	Hotels (n=50)
Less than 3 months	10%	12%
3-6 months	43%	48%
6-12 months	36%	38%
12+ months	6%	2%
Already a critical risk	5%	0%

Top 3 Most Important Potential Benefits of Using Agentic AI

Q14	Total (n=304)	Hotels (n=50)
Automation of repetitive security tasks	18%	22%
Improved security for decentralized or remote environments	20%	20%
Freeing up analysts to focus on strategic threats	9%	16%

Top 3 Attack Types Agentic AI Is Best at Stopping

Q15	Total (n=304)	Hotels (n=50)
Fake account creation / fake new accounts	39%	32%
GPT prompt compromise	30%	30%
Inventory hoarding	21%	26%

Know Your Agent (KYA)

Statement Describing Enterprises Dealing with AI Agents

Q18c-Somewhat/Completely Describe	Total (n=304)	Hotels (n=50)
We actively blocked all suspected AI agent traffic on our consumer platforms.	73%	74%
We're developing separate workflows specifically for AI agent interactions.	80%	72%
We allow AI agents but monitor their usage closely.	78%	64%

Know Your Agent (KYA) Greatest Challenge

Q18	Total (n=304)	Hotels (n=50)
Verifying that legitimate AI agents are properly authorized by real users	25%	28%
Detecting when malicious agents are impersonating legitimate customers	23%	26%
Managing the scale of AI agent interactions	20%	22%

View of Customers Using AI Agents on Platform

Q18f	Total (n=304)	Hotels (n=50)
Net positive for our business	32%	30%
Net negative for our business	7%	10%
Mixed impact depending on context	53%	50%
Too early to tell	9%	10%

The Trust Tax: Why Sharing/Gig Platforms Pay More to Verify Less

Sharing/Gig platforms are very well prepared for AI attacks with the highest confidence in any industry surveyed. Fake accounts and unauthorized sharing both hit 90% concern. But the math doesn't work: supreme confidence meets extreme threat exposure.

Budgets reveal the squeeze. AI security spending sits at 24%, rising to 27%—below average despite higher attack volumes. One fake account cascades across multiple transactions, multiple victims, multiple fraud types. Thirty percent lack defenses against agentic threats they've already deployed. The difference: 65% see consumer agents as net positive, double the industry rate. Automation built these platforms. Consumer agents fit the model. The scorecards ahead show whether that optimism holds.

SHARING/GIG ECONOMY INDUSTRY AI MATURITY SCORECARDS

Gig economy platforms show a striking confidence-threat mismatch: 70% rate themselves very well prepared (26 points above cross-industry average), yet they face the highest concern levels for fake accounts and unauthorized sharing at 90% each—threats that directly undermine the trust-based business models these platforms depend on. The 80% who prefer to learn from and meet with vendors (vs 59% cross-industry) suggests these platforms recognize the value partners bring when building specialized fraud detection in-house while scaling operations.

Evolving Threats

Level of Concern from Each Attack Type

Q1-moderate/large Extent	Total (n=304)	Share/Gig Economy (n=20)
Fake account creation/fake new accounts	73%	90%
Unauthorized account sharing	69%	90%
API abuse	70%	85%

Approximate Cost of Negative Consequences

Q3a	Total (n=304)	Share/Gig Economy (n=20)
Less than $1M	8%	0%
$1M to $9.9M	38%	45%
$10M to $99M	39%	45%
$100M to $500M	11%	10%
Over $500M	5%	0%

Preparedness Rising

Approximate % of Cybersecurity Budget Spent on Security Solutions Leveraging AI

Q5c,Q5d-average	Total (n=304)	Share/Gig Economy (n=20)
Today	28%	24%
12 months from now	32%	27%

Level of Preparedness for Defending Against AI-Powered Volumetric Attacks

Q4a	Total (n=304)	Share/Gig Economy (n=20)
Moderately well prepared	47%	25%
Very well prepared	44%	70%

Current Approaches to Learning About Agentic AI Opportunities or Risks

Q20 - Top 3	Total (n=304)	Share/Gig Economy (n=20)
Attending industry events or webinars focused on Agentic AI	59%	80%
Conducting internal research and testing	64%	75%
Engaging with vendors or consultants to assess Agentic AI	68%	55%

Adoption & Benefits

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total (n=304)	Share/Gig Economy (n=20)
AI has improved my enterprises overall cybersecurity posture	85%	100%
We gain more value by buying AI-powered cybersecurity solutions than if we build the solutions in-house	81%	80%
We are restricted from leveraging AI in our cybersecurity solutions due to model governance policies	57%	75%
It's too difficult to integrate AI-powered cybersecurity solutions with our existing systems	38%	20%

Benefits of AI-Powered Bot Management and Account Security Solutions

Q8	Total (n=304)	Share/Gig Economy (n=20)
Improve threat detection and response	47%	75%
Better defend against bot attacks	51%	70%
Enhanced API security and abuse prevention	46%	65%
Reduce overall cost of securing my business	40%	55%

Gig platforms see consumer AI agents differently than other industries: 65% view them as net positive (versus 32% cross-industry) because automation built these businesses from the start. Yet this embrace creates exposure: platforms are deploying agent workflows while struggling with the same verification challenges everyone faces, and the 45% who expect attacks to become critical within 3-6 months means the window to solve authorization and impersonation detection is nearly closed.

Agentic AI Looms Large

Current Enterprise's Agentic AI Reality

Q11	Total (n=304)	Share/Gig Economy (n=20)
We're actively using Agentic AI for cybersecurity/ fraud prevention and have comprehensive defenses against malicious Agentic AI	52%	60%
We're using Agentic AI for cybersecurity but lack adequate defenses against Agentic AI-powered threats	36%	30%
We're neither using Agentic AI nor have specific defenses against Agentic AI-powered threats	1%	5%

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total (n=304)	Share/Gig Economy (n=20)
Bad actors use of Agentic AI to attack enterprises is the next big threat	71%	50%
Agentic AI-powered attacks will pose a fundamentally different threat than traditional cyberattacks	78%	80%

Agentic AI Looms Large

Estimated Duration Before Agentic AI-Powered Attacks Become a Critical Risk

Q13	Total (n=304)	Share/Gig Economy (n=20)
Less than 3 months	10%	0%
3-6 months	43%	45%
6-12 months	36%	45%
12+ months	6%	10%
Already a critical risk	5%	0%

Top 3 Most Important Potential Benefits of Using Agentic AI

Q14	Total (n=304)	Share/Gig Economy (n=20)
Automation of repetitive security tasks	18%	30%
Faster response to cyber threats	23%	20%
Improved security for decentralized or remote environments	20%	20%

Top 3 Attack Types Agentic AI Is Best at Stopping

Q15	Total (n=304)	Share/Gig Economy (n=20)
Account takeovers (ATO) / credential stuffing	29%	45%
Fake account creation / fake new accounts	39%	45%
API abuse	21%	20%

Know Your Agent (KYA)

Statement Describing Enterprises Dealing with AI Agents

Q18c-Somewhat/Completely Describe	Total (n=304)	Share/Gig Economy (n=20)
We're developing separate workflows specifically for AI agent interactions.	80%	65%
We allow AI agents but monitor their usage closely.	78%	60%
We actively blocked all suspected AI agent traffic on our consumer platforms.	73%	55%

Know Your Agent (KYA) Greatest Challenge

Q18	Total (n=304)	Share/Gig Economy (n=20)
Detecting when malicious agents are impersonating legitimate customers	23%	25%
Ensuring compliance when AI agents act on behalf of users	17%	25%
Verifying that legitimate AI agents are properly authorized by real users	25%	20%

View of Customers Using AI Agents on Platform

Q18f	Total (n=304)	Share/Gig Economy (n=20)
Net positive for our business	32%	65%
Net negative for our business	7%	5%
Mixed impact depending on context	53%	30%
Too early to tell	9%	0%

Social Media Platforms Lead in Agentic AI Defense

For decades social media platforms have battled volumes of abuse other industries never imagined. This constant warfare now makes them prepared for agentic AI threats. While most industries struggle to distinguish legitimate automation from malicious activity, social platforms already mastered this challenge—90% allow and monitor consumer AI agent behavior closely, well above the cross-industry rate. The scorecards ahead reveal an industry that leads by doing, not debating.

SOCIAL MEDIA/NETWORKING INDUSTRY AI MATURITY SCORECARDS

Social platforms deploy agentic AI defenses successfully—65% have comprehensive defenses versus the 52% cross-industry. Battle-tested infrastructure built to fight early types of automation extends naturally to defending against modern agentic AI. Social media platforms engage vendors selectively for specialized threats like web scraping while building in-house AI teams to extract maximum value from lower budgets through operational efficiency. The few vendor platforms they trust are deep experts at detecting and mitigating automation.

Evolving Threats

Level of Concern from Each Attack Type

Q1-moderate/large Extent	Total (n=304)	Social Platforms (n=20)
Web scraping/content theft	69%	90%
Account takeovers (ATO)/credential stuffing	70%	85%
API abuse	70%	85%

Approximate Cost of Negative Consequences

Q3a	Total (n=304)	Social Platforms (n=20)
Less than $1M	8%	15%
$1M to $9.9M	38%	60%
$10M to $99M	39%	20%
$100M to $500M	11%	5%
Over $500M	5%	0%

Preparedness Rising

Approximate % of Cybersecurity Budget Spent on Security Solutions Leveraging AI

Q5c,Q5d-average	Total (n=304)	Social Platforms (n=20)
Today	28%	22%
12 months from now	32%	24%

Level of Preparedness for Defending Against AI-Powered Volumetric Attacks

Q4a	Total (n=304)	Social Platforms (n=20)
Moderately well prepared	47%	45%
Very well prepared	44%	40%

Current Approaches to Learning About Agentic AI Opportunities or Risks

Q20 - Top 3	Total (n=304)	Social Platforms (n=20)
Engaging with vendors or consultants to assess Agentic AI	68%	85%
Attending industry events or webinars focused on Agentic AI	59%	65%
Following governments / regulatory guidance	49%	65%

Adoption & Benefits

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total (n=304)	Social Platforms (n=20)
AI has improved my enterprises overall cybersecurity posture	85%	80%
We gain more value by buying AI-powered cybersecurity solutions than if we build the solutions in-house	81%	75%
We are restricted from leveraging AI in our cybersecurity solutions due to model governance policies	57%	80%
It's too difficult to integrate AI-powered cybersecurity solutions with our existing systems	38%	60%

Benefits of AI-Powered Bot Management and Account Security Solutions

Q8	Total (n=304)	Social Platforms (n=20)
Reduce overall cost of securing my business	40%	55%
Better defend against attack automation services	39%	55%
Improve threat detection and response	47%	50%
Better protection against phishing attacks	46%	50%

While other industries debate whether to allow consumer AI agents, social platforms already solved it—90% enable them with active monitoring. That's 12 points higher than the cross-industry rate. This confidence comes from knowing exactly what malicious automation looks like, letting platforms greenlight consumer agents while catching attacks in real time. The question shifted from "should we allow agents" to "how do we enable them safely at scale."

Agentic AI Looms Large

Current Enterprise's Agentic AI Reality

Q11	Total (n=304)	Social Platforms (n=20)
We're actively using Agentic AI for cybersecurity/ fraud prevention and have comprehensive defenses against malicious Agentic AI	52%	65%
We're using Agentic AI for cybersecurity but lack adequate defenses against Agentic AI-powered threats	36%	20%
We're not using Agentic AI but have implemented specific defenses against Agentic AI-powered threats	4%	10%

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total (n=304)	Social Platforms (n=20)
Bad actors use of Agentic AI to attack enterprises is the next big threat	71%	75%
Agentic AI-powered attacks will pose a fundamentally different threat than traditional cyberattacks	78%	80%

Agentic AI Looms Large

Estimated Duration Before Agentic AI-Powered Attacks Become a Critical Risk

Q13	Total (n=304)	Social Platforms (n=20)
Less than 3 months	10%	10%
3-6 months	43%	55%
6-12 months	36%	30%
12+ months	6%	0%
Already a critical risk	5%	5%

Top 3 Most Important Potential Benefits of Using Agentic AI

Q14	Total (n=304)	Social Platforms (n=20)
Improved security for decentralized or remote environments	20%	30%
Faster response to cyber threats	23%	20%
Automation of repetitive security tasks	18%	15%

Top 3 Attack Types Agentic AI Is Best at Stopping

Q15	Total (n=304)	Social Platforms (n=20)
Fake account creation / fake new accounts	39%	60%
Account takeovers (ATO) / credential stuffing	29%	25%
GPT prompt compromise	30%	25%

Know Your Agent (KYA)

Statement Describing Enterprises Dealing with AI Agents

Q18c-Somewhat/Completely Describe	Total (n=304)	Social Platforms (n=20)
We allow AI agents but monitor their usage closely.	78%	90%
We actively blocked all suspected AI agent traffic on our consumer platforms.	73%	85%
We're developing separate workflows specifically for AI agent interactions.	80%	80%

Know Your Agent (KYA) Greatest Challenge

Q18	Total (n=304)	Social Platforms (n=20)
Detecting when malicious agents are impersonating legitimate customers	23%	25%
Ensuring compliance when AI agents act on behalf of users	17%	25%
Verifying that legitimate AI agents are properly authorized by real users	25%	20%

View of Customers Using AI Agents on Platform

Q18f	Total (n=304)	Social Platforms (n=20)
Net positive for our business	32%	30%
Net negative for our business	7%	20%
Mixed impact depending on context	53%	45%
Too early to tell	9%	5%

Streaming Services Embrace Consumer AI Agents

Streaming services see consumer AI agents as opportunity, not threat. Just 3% view them as net negative for business—the lowest skepticism of any industry. This enthusiasm reflects platforms' content recommendation DNA: they already use AI to personalize every user experience, so consumer agents that automate watchlist management and family sharing feel like natural extensions rather than invasive automation.

But openness creates exposure. Sixty percent expect agentic AI attacks to become critical within 3-6 months, a super compressed timeline. Streaming services racing to enable consumer convenience face adversaries who will exploit those same agent capabilities for credential stuffing, account sharing abuse and content access fraud. The scorecards reveal an industry betting that early experience with AI-powered personalization translates to defending against AI-powered attacks.

STREAMING SERVICES INDUSTRY AI MATURITY SCORECARDS

Streaming services operate with unusual confidence—53% rate themselves very well prepared versus 44% cross-industry. AI already powers their core business through recommendation engines and personalization algorithms, so security AI feels familiar rather than foreign. The 90% who prefer buying fraud detection recognize that stopping credential stuffers requires different expertise than suggesting your next binge-watch.

Evolving Threats

Level of Concern from Each Attack Type

Q1-moderate/large Extent	Total (n=304)	Streaming Services (n=30)
GPT prompt compromise	66%	80%
API abuse	70%	77%
Account takeovers (ATO) / credential stuffing	70%	77%

Approximate Cost of Negative Consequences

Q3a	Total (n=304)	Streaming Services (n=30)
Less than $1M	8%	13%
$1M to $9.9M	38%	50%
$10M to $99M	39%	30%
$100M to $500M	11%	7%
Over $500M	5%	0%

Preparedness Rising

Approximate % of Cybersecurity Budget Spent on Security Solutions Leveraging AI

Q5c,Q5d-average	Total (n=304)	Streaming Services (n=30)
Today	28%	26%
12 months from now	32%	31%

Level of Preparedness for Defending Against AI-Powered Volumetric Attacks

Q4a	Total (n=304)	Streaming Services (n=30)
Moderately well prepared	47%	43%
Very well prepared	44%	53%

Current Approaches to Learning About Agentic AI Opportunities or Risks

Q20 - Top 3	Total (n=304)	Streaming Services (n=30)
Engaging with vendors or consultants to assess Agentic AI	68%	67%
Attending industry events or webinars focused on Agentic AI	59%	63%
Conducting internal research and testing	64%	60%

Adoption & Benefits

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total (n=304)	Streaming Services (n=30)
AI has improved my enterprises overall cybersecurity posture	85%	93%
We gain more value by buying AI-powered cybersecurity solutions than if we build the solutions in-house	81%	90%
We are restricted from leveraging AI in our cybersecurity solutions due to model governance policies	57%	50%
It's too difficult to integrate AI-powered cybersecurity solutions with our existing systems	38%	33%

Benefits of AI-Powered Bot Management and Account Security Solutions

Q8	Total (n=304)	Streaming Services (n=30)
Better defend against bot attacks	51%	63%
Better defend against Agentic AI attacks	41%	57%
Reducing overall cost of securing my business	40%	50%
Better defend against generative AI-powered attacks	42%	50%

Ninety percent of streaming platforms believe agentic AI will fundamentally change cyberattacks—they're not just expecting faster bots, they're preparing for entirely new threat patterns. This recognition drives investment in speed: 27% cite a faster response to cyber threats as a benefit of using agentic AI for fraud prevention because when attacks scale across millions of distributed accounts, minutes matter. Streaming services managing users across devices, geographies and family sharing arrangements need agentic AI to solve scale problems manual security teams won't be able handle alone.

Agentic AI Looms Large

Current Enterprise's Agentic AI Reality

Q11	Total (n=304)	Streaming Services (n=30)
We're actively using Agentic AI for cybersecurity/ fraud prevention and have comprehensive defenses against malicious Agentic AI	52%	53%
We're using Agentic AI for cybersecurity but lack adequate defenses against Agentic AI-powered threats	36%	23%
We're still evaluating the risks and benefits of Agentic AI in cybersecurity	8%	17%

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total (n=304)	Streaming Services (n=30)
Bad actors use of Agentic AI to attack enterprises is the next big threat	71%	57%
Agentic AI-powered attacks will pose a fundamentally different threat than traditional cyberattacks	78%	90%

Agentic AI Looms Large

Estimated Duration Before Agentic AI-Powered Attacks Become a Critical Risk

Q13	Total (n=304)	Streaming Services (n=30)
Less than 3 months	10%	10%
3-6 months	43%	60%
6-12 months	36%	27%
12+ months	6%	3%
Already a critical risk	5%	0%

Top 3 Most Important Potential Benefits of Using Agentic AI

Q14	Total (n=304)	Streaming Services (n=30)
Faster response to cyber threats	23%	27%
Improved security for decentralized or remote environments	20%	20%
Automation of repetitive security tasks	11%	17%

Top 3 Attack Types Agentic AI Is Best at Stopping

Q15	Total (n=304)	Streaming Services (n=30)
Fake account creation / fake new accounts	39%	43%
GPT prompt compromise	30%	37%
Account takeovers (ATO) / credential stuffing	29%	30%

Know Your Agent (KYA)

Statement Describing Enterprises Dealing with AI Agents

Q18c-Somewhat/Completely Describe	Total (n=304)	Streaming Services (n=30)
We're developing separate workflows specifically for AI agent interactions.	80%	83%
We allow AI agents but monitor their usage closely.	78%	80%
We actively blocked all suspected AI agent traffic on our consumer platforms.	73%	73%

Know Your Agent (KYA) Greatest Challenge

Q18	Total (n=304)	Streaming Services (n=30)
Verifying that legitimate AI agents are properly authorized by real users	25%	27%
Distinguishing AI agent behavior from human behavior	14%	27%
Ensuring compliance when AI agents act on behalf of users	17%	27%

View of Customers Using AI Agents on Platform

Q18f	Total (n=304)	Streaming Services (n=30)
Net positive for our business	32%	30%
Net negative for our business	7%	3%
Mixed impact depending on context	53%	47%
Too early to tell	9%	20%

Tech Industry: Highest Losses, Lowest Confidence

Technology platforms spend the most of any other industry of their cybersecurity budgets on AI security. But they don't feel very well prepared to defend against AI attacks. Highest spending meets lowest confidence. Integration difficulty hits 50%, which is higher than the cross-industry rate. Purchased tools sit unused. Attacks succeed anyway. Thirteen percent lose over $500M annually—catastrophic losses no other industry approaches. The irony cuts: companies that are the most likely to build AI into their products struggle to deploy AI defenses. Threat recognition is crystal clear, as a majority say agentic attacks will be fundamentally different. Confidence lags behind capability, the data suggests.

TECHNOLOGY INDUSTRY AI MATURITY SCORECARDS

Tech platforms face the integration paradox: 30% AI security spending (highest of any industry) produces just 30% who feel very well prepared (lowest confidence), because 50% struggle to integrate purchased solutions with existing systems. This creates expensive shelf-ware where security tools sit unused while attacks succeed. The 93% who say AI improved their posture reflects aspiration more than operational reality—believing tools will help once deployment challenges get solved.

Evolving Threats

Level of Concern from Each Attack Type

Q1-moderate/large Extent	Total (n=304)	Technology Platforms (n=30)
API abuse	70%	83%
Fake account creation / fake new accounts	73%	83%
Web scraping / content theft	69%	83%

Approximate Cost of Negative Consequences

Q3a	Total (n=304)	Technology Platforms (n=30)
Less than $1M	8%	10%
$1M to $9.9M	38%	40%
$10M to $99M	39%	27%
$100M to $500M	11%	10%
Over $500M	5%	13%

Preparedness Rising

Approximate % of Cybersecurity Budget Spent on Security Solutions Leveraging AI

Q5c,Q5d-average	Total (n=304)	Technology Platforms (n=30)
Today	28%	30%
12 months from now	32%	32%

Level of Preparedness for Defending Against AI-Powered Volumetric Attacks

Q4a	Total (n=304)	Technology Platforms (n=30)
Moderately well prepared	47%	60%
Very well prepared	44%	30%

Current Approaches to Learning About Agentic AI Opportunities or Risks

Q20 - Top 3	Total (n=304)	Technology Platforms (n=30)
Attending industry events or webinars focused on Agentic AI	59%	73%
Conducting internal research and testing	64%	73%
Participating in industry groups or threat intelligence sharing networks	40%	63%

Adoption & Benefits

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total (n=304)	Technology Platforms (n=30)
AI has improved my enterprises overall cybersecurity posture	85%	93%
We gain more value by buying AI-powered cybersecurity solutions than if we build the solutions in-house	81%	77%
We are restricted from leveraging AI in our cybersecurity solutions due to model governance policies	57%	53%
It's too difficult to integrate AI-powered cybersecurity solutions with our existing systems	38%	50%

Benefits of AI-Powered Bot Management and Account Security Solutions

Q8	Total (n=304)	Technology Platforms (n=30)
Improve threat detection and response	47%	53%
Enahnced API security and abuse preventon	46%	53%
Better defend against bot attacks	51%	50%
Better defend against Agentic AI attacks	41%	50%

Tech platforms see the threat transformation more clearly than anyone—90% say agentic attacks will be fundamentally different versus 78% cross-industry—and they're responding with speed as the top priority: 43% identify faster threat response as agentic AI's most important benefit. But recognition without deployment creates vulnerability: platforms building AI products for customers struggle to integrate AI defenses for themselves, and the 33% who cite agent authorization verification as their greatest challenge reveals gaps in the Know Your Agent capabilities they're racing to enable.

Agentic AI Looms Large

Current Enterprise's Agentic AI Reality

Q11	Total (n=304)	Technology Platforms (n=30)
We're actively using Agentic AI for cybersecurity/ fraud prevention and have comprehensive defenses against malicious Agentic AI	52%	63%
We're using Agentic AI for cybersecurity but lack adequate defenses against Agentic AI-powered threats	36%	23%
We're still evaluating the risks and benefits of Agentic AI in cybersecurity	8%	7%

Agreement with AI-Related Cybersecurity Statements

Q6-Agree/Completely Agree	Total (n=304)	Technology Platforms (n=30)
Bad actors use of Agentic AI to attack enterprises is the next big threat	71%	83%
Agentic AI-powered attacks will pose a fundamentally different threat than traditional cyberattacks	78%	90%

Agentic AI Looms Large

Estimated Duration Before Agentic AI-Powered Attacks Become a Critical Risk

Q13	Total (n=304)	Technology Platforms (n=30)
Less than 3 months	10%	10%
3-6 months	43%	43%
6-12 months	36%	33%
12+ months	6%	10%
Already a critical risk	5%	3%

Top 3 Most Important Potential Benefits of Using Agentic AI

Q14	Total (n=304)	Technology Platforms (n=30)
Faster response to cyber threats	23%	43%
Proved security for decentralized or remote environments	20%	20%
Automation of repetitive security tasks	18%	10%

Top 3 Attack Types Agentic AI Is Best at Stopping

Q15	Total (n=304)	Technology Platforms (n=30)
Fake account creation / fake new accounts	39%	43%
GPT prompt compromise	30%	33%
Account takeovers (ATO) / credential stuffing	29%	27%

Know Your Agent (KYA)

Statement Describing Enterprises Dealing with AI Agents

Q18c-Somewhat/Completely Describe	Total (n=304)	Technology Platforms (n=30)
We allow AI agents but monitor their usage closely.	78%	87%
We're developing separate workflows specifically for AI agent interactions.	80%	83%
We actively blocked all suspected AI agent traffic on our consumer platforms.	73%	77%

Know Your Agent (KYA) Greatest Challenge

Q18	Total (n=304)	Technology Platforms (n=30)
Verifying that legitimate AI agents are properly authorized by real users	25%	33%
Detecting when malicious agents are impersonating legitimate customers	23%	23%
Distinguishing AI agent behavior from human behavior	14%	17%

View of Customers Using AI Agents on Platform

Q18f	Total (n=304)	Technology Platforms (n=30)
Net positive for our business	32%	37%
Net negative for our business	7%	3%
Mixed impact depending on context	53%	53%
Too early to tell	9%	7%

SIX STRATEGIC RECOMMENDATIONS FOR ENTERPRISE SECURITY LEADERS

1. Match Your Timeline to the Threat Window
Fifty-three percent of enterprises expect agentic AI attacks to become critical within six months, yet standard security transformations take 12-18 months. Companies that started in 2024 are ready. Those beginning now face capability gaps against adversaries already scaling autonomous attacks. Compress procurement cycles and deploy imperfect protection now rather than perfect protection arriving after breaches succeed.

2. Build Know Your Agent (KYA) Capabilities Before Consumer AI Agents Scale
Enterprises are enabling consumer AI agents (88% importance rating) before solving fundamental verification. Twenty-five percent cannot verify that legitimate agents are properly authorized, while 23% struggle to detect impersonation. Prioritize KYA infrastructure now: authorization verification, impersonation detection and behavioral analysis that distinguishes helpful consumer agents from credential stuffers and fake accounts. The gap between enablement and verification is your largest emerging vulnerability.

3. Don't Let Procurement Timelines Exceed Threat Timelines
"Hotels need to get ahead of this now because the procurement process is so slow. You've got to get the most modern tools in place now," notes Chris Staab of LSA. Establish pre-approved vendor frameworks and emergency procurement paths for critical security tools. When 36%-60% of industries expect attacks critical within 3-6 months, standard 6-12 month procurement cycles guarantee you're defending yesterday's threats with tomorrow's tools.

4. Decide Your Consumer Agent Strategy Now: Enable or Block, But Choose
Sixty-eight percent accept that consumers' AI agents need autonomy, but approaches are split: 39% require pre-authorization, 29% accept periodic oversight, 28% demand continuous control. Indecision creates inconsistent policies that confuse consumers and security teams. Decide whether consumer-controlled AI agents enhance your business model or threaten it, then build governance accordingly. The question isn't whether agents arrive, it's whether you're ready.

5. Partner With Vendors Who Provide Adaptive Defense Platforms and Cross-Industry Threat Intelligence
The 81% who prefer buying over building recognize specialized vendors deliver faster than internal teams. But not all vendors are equal. Prioritize partners who provide: comprehensive platforms that evolve with adversary techniques, offer cross-industry risk signals that identify emerging attack patterns before they hit your industry and provide real-time threat intelligence covering AI-powered bots, human fraud farms and agentic AI. Vendors with successful deployments at peer enterprises demonstrate proven capability at your scale.

6. Diagnose Whether You're Realizing Benefits or Creating Shelf-Ware
Between 42%-51% of enterprises achieve measurable improvements from AI-powered security solutions. The other half deployed the same tools but haven't translated adoption into outcomes. Implementation rigor separates them. The 38% facing integration challenges fall predominantly into the struggling half. If you've deployed AI security tools but haven't seen measurable improvements in detection accuracy, response times or false positive reductions, you have a deployment-to-protection gap. Work with trusted partners to track metrics that prove capabilities work, not just exist.

METHODOLOGY SECTION

This report presents findings from 304 cybersecurity decision-makers at global B2C enterprises with $250M+ in annual revenue—with 63% representing companies in the $1B to $20B+ range. Surveyed August 26–September 12, 2025, this is Arkose Labs' second annual AI security survey, expanding from U.S.-only coverage in 2024 to include Australia in 2025. This expansion enables year-over-year trend analysis and reveals how AI threats and defenses diverge across regulatory environments and market structures.

Sample: 202 U.S. and 102 Australian respondents across six industries: Banking & Fintech (n=104), Airline & Hotels (n=100), Streaming Services (n=30), Technology Platforms (n=30), Social Networking (n=20), and Share/Gig Economy (n=20). These sectors face disproportionate AI-powered fraud volumes.

Respondent credentials: 57% executive leadership (C-suite, CISO, CSO, CTO, SVP, VP, Head of Department) with enterprise-wide oversight; 43% operational leadership (Sr. Director, Director, Manager/Sr. Manager) with implementation responsibility. All were screened for active decision-making authority in cybersecurity strategy, budget allocation, vendor selection or technology implementation.

Research objectives: Measure enterprise AI threat awareness, assess defensive AI adoption maturity and benchmark preparedness levels for defending against AI-powered attack campaigns.

Arkose Labs partnered with nationally recognized research firm KS&R.

TAKE CONTROL WITH ARKOSE TITAN

Ready to make attacking you unprofitable?

Transform your security posture from reactive detection to proactive economic deterrence.
Schedule a call with an expert today.

ABOUT ARKOSE LABS

Arkose Labs is the leading global provider offering a proactive fraud deterrence platform purpose-built to neutralize modern attacks, including those powered by Agentic AI and large language models (LLMs). Its comprehensive solution combines proprietary device identification (device ID), behavioral analysis, phishing protection, email intelligence, scraping prevention, API defense and bot management. Trusted by the world's leading consumer brands—including two of the top three banks, Microsoft, Meta, Roblox, and many others—Arkose Labs stops account takeovers, fake account creation, LLM-driven scraping and SMS toll fraud. The platform actively undermines attacker ROI by introducing dynamic friction, making it economically unsustainable for adversaries to persist. Its Security Operations Center (SOC) provides actionable insights from an extensive cross-industry intelligence network, which monitors legitimate traffic and attack patterns across global enterprises. With unparalleled proactive support for internal security teams, Arkose Labs goes beyond conventional security by actively partnering with customers to disrupt organized fraud networks such as Storm-1152. Headquartered in San Mateo, California, the company maintains a global presence with offices throughout APAC, Central America, EMEA and South America.

TALK TO AN EXPERT

USA (San Mateo)

Australia (Brisbane)

United Kingdom (London)

Costa Rica (San José)

India (Pune)

Argentina (Buenos Aires)

APPENDIX A

Survey Questions:

Q1. Think about the critical applications to your business – such as revenue-driving apps, websites/platforms, network systems, etc. How concerned are you with each of the following attack types targeting those critical applications?

Q2. To what extent has your enterprise suffered negative consequences over the last 12 months in each of the following areas due to the attack types just considered?

Q3a. What is the approximate DOLLAR quantification of these consequences over the past 12 months… that is, how much have these negative consequences COST your enterprise to the best of your knowledge?

Q3b. Consider the attack activity your enterprise has seen over the past 12 months, approximately what percentage of those attacks were from bots, AI-powered bots, AI agents, attack automation services, human fraud farms or low-and-slow manual attacks? Provide your best approximation.

Q4a. How prepared would you rate your enterprise in terms of using AI to defend against bad actors using AI-powered bots to deploy volumetric attacks?

Q5a. To what extent is your enterprise taking the following actions…?
Q5c. Approximately what percentage of your overall cybersecurity budget is spent on security solutions leveraging AI (including generative AI and Agentic AI)?

Q5d. And approximately what percentage of your overall cybersecurity budget do you estimate will be spent on security solutions leveraging AI (including generative AI and Agentic AI) 12 months from now?

Q6. And how much do you agree with each of the following AI-related cybersecurity statements…

Q8. What benefits have you realized (or expect to realize in the near future) from AI-powered bot management and account security?

Q10. When considering a vendor to support your bot management device ID, API protection, etc. needs, how important are each of the following characteristics…

Q11. When it comes to Agentic AI, which statement best describes your enterprise's current reality?

Q12. Looking ahead to 2026, which scenario do you think is most likely regarding Agentic AI?

Q13. How many months do you estimate your enterprise has before Agentic AI-powered attacks become a critical business risk?

Q14. Which potential benefit of using Agentic AI in your department is most important to your enterprise?

These survey questions are copyrighted.

Q15. Which attack type do you think Agentic AI in your defense tech stack would be best at stopping?

Q16. How urgent is it for your enterprise to develop capabilities that can distinguish between legitimate consumer AI agents and malicious AI agents trying to appear as legitimate consumers?

Q17. How confident are you in your current security infrastructure's ability to distinguish between legitimate AI agents acting on behalf of consumers versus malicious AI agents trying to appear as legitimate consumers?

Q18. Which aspect of AI agent authentication (also known as “Know Your Agent” [KYA]) poses the greatest challenge for your enterprise?

Q18b. How much do you trust AI agent provider(s) to ensure their agents are NOT used for malicious purposes?
Q18c. To what degree does each statement describe your enterprise when dealing with AI agents on your platform?

Q18f. Which of the following best describes your view of customers using AI agents on your platform?

Q19. How is your enterprise adjusting hiring plans or internal resourcing in response to the rise of Agentic AI?

Q19a. Considering the various AI-based threats facing your enterprise, which one causes you the most concern due to its potential to significantly harm your enterprise today?

Q20. How is your enterprise currently learning about or evaluating Agentic AI as a cybersecurity risk or opportunity?

Q22. What is your primary concern when AI agents – acting on behalf of your customers - interact with your platforms? Please select your top concern.

Q23. How important is it for your enterprise to distinguish between human users and AI agents on your platforms?

Q24. When customers use AI agents to interact with your platforms, what level of human involvement do you require?

Q25. What would MOST influence your decision to create dedicated AI agent interfaces versus allowing agents to use existing user interfaces?