

FORTUNE

CYBER60™

200 CISOS ON ENTERPRISE AI AGENTS,
THE CONTROLS THEY NEED, AND WHERE
THEY'LL INVEST NEXT

Three in four enterprises now run AI agents.

Security is losing ground in AI governance,
confidence in in-house skills has fallen, and
more CISOs are spending to close the gap.

TABLE OF CONTENTS

CISO SURVEY: SECURING THE AGENT ENTERPRISE	02
AI THREAT LANDSCAPE	03
AGENTS IN THE ENTERPRISE	05
ORGANIZATIONAL CONTEXT AND GOVERNANCE	07
AI IN SECURITY OPERATIONS	09
AUTONOMOUS RESPONSE	12
VENDOR STRATEGY AND PROCUREMENT	14
BUDGET AND INVESTMENT OUTLOOK	16
SKILLS AND TALENT	18
IN THEIR OWN WORDS	20
THE AI SECURITY OPPORTUNITY MAP	21
KEY CONCLUSIONS AND IMPLICATIONS	29
METHODOLOGY	31

WELCOME TO THE CYBER 60

FOURTH EDITION

At Lightspeed, we've been investing in enterprise founders for over 25 years. Cybersecurity investing is one of the pillars of our global enterprise practice, and a focus area of deep immersion and commitment. This report also shares the results of our annual CISO survey, conducted once again with Wakefield Research. Wakefield surveyed 200 chief information security officers (CISOs) at U.S. companies with \$500 million or more in revenue.

A year ago we asked CISOs about AI as a threat and a tool. This year we asked them about AI agents: software that can take actions inside the enterprise. Agents are already in production at **77%** of the companies we surveyed. Among companies using or exploring agents, only one in five say they have a full inventory of them, and a quarter use or would use human credentials to authenticate agents. Across the full sample, **88% already let AI take at least some response actions with no human in the loop.**

Security is formally involved in AI governance at 78% of these companies, down from 87% last year. The share where security leads the process fell from a third to a fifth. Confidence that in-house teams can handle AI-specific risk fell too. Budgets are still going up, and **85% of CISOs plan to hire or reskill for AI security within the year.**

THE SURVEY SHOWS
WHERE CISOs PLAN TO
SPEND AND HIRE, AND
WHERE THEY NEED
BETTER TOOLS.

CISO SURVEY: SECURING THE AGENT ENTERPRISE

The 2026–27 Lightspeed Cyber 60 CISO survey finds that companies are deploying AI agents while still building the controls to govern them. Our survey of 200 CISOs at companies with \$500M+ in annual revenue shows agents running in production at 77% of organizations. Among companies using or exploring agents, only a third use or would use dedicated agent identities with scoped permissions. Compromised or hijacked agents now tie AI-generated phishing and deepfakes as the top AI threat, and agent identity and access governance is

the tool category CISOs most plan to evaluate in the year ahead.

79% of CISOs report a confirmed or suspected AI-related incident in the past 12 months, up from 75%. Confidence that in-house teams can assess and mitigate AI risk dropped from 71% to 55%. Spending continues, with 85% expecting budgets to rise and 82% committing more than 5% of the security budget to AI.

77%

HAVE AI AGENTS IN PRODUCTION

79%

EXPERIENCED A CONFIRMED OR SUSPECTED AI-RELATED SECURITY INCIDENT IN THE PAST 12 MONTHS

88%

ALREADY ALLOW AI TO EXECUTE AT LEAST SOME RESPONSE ACTIONS WITHOUT HUMAN APPROVAL


55%

ARE FULLY OR MOSTLY CONFIDENT IN THEIR TEAM'S IN-HOUSE AI SECURITY SKILLS, DOWN FROM 71% LAST YEAR

AI THREAT LANDSCAPE

AI-RELATED SECURITY INCIDENTS (PAST 12 MONTHS)

AI-RELATED SECURITY INCIDENTS, 2026 VS. 2025




CRITICAL FINDING

79% of CISOs report a confirmed or suspected AI-related incident in the past year, up from 75%. Another 20% detected AI-related attack attempts without reporting an incident, leaving just 1% reporting neither. Within the incident total, confirmed incidents fell from 37% to 16%, while suspected but unconfirmed incidents rose from 38% to 63%. More CISOs suspect AI involvement, but fewer report confirming it. The survey can't say whether attackers are covering their tracks better or defenders are being more cautious about attribution. Only 1% report no incidents or detected attempts.

Companies at \$1B+ in revenue were more likely to report an incident (85%) than those between \$500M and \$1B (67%).

PERCEIVED THREAT LEVEL OF AI ATTACKS (NEXT 12 MONTHS)

THREAT POSED BY AI-GENERATED OR AI-AUGMENTED ATTACKS, 2026 VS. 2025




KEY INSIGHT

71% of CISOs rate AI-driven attacks a moderate or severe threat for the coming year, up from 67%. The share calling the threat severe rose by half, from 10% to 15%.

AI THREAT LANDSCAPE

AI-RELATED THREATS OF GREATEST CONCERN



KEY INSIGHT

Last year's list focused on phishing, synthetic identities, and LLM-powered malware. This year, compromised or hijacked agents tie AI-generated phishing and deepfakes at 51%. Three of the top five concerns involve agents: hijacking, ungoverned agent credentials (41%), and the tool and integration supply chain (40%). Prompt injection, which can be used to hijack agents, sits at 36%.

Concern about ungoverned non-human identities runs higher among CISOs who say a vendor's AI strategy is very or critically important to them (48%) than among those who say it matters less (27%).


"CISOS ARE WORRIED ABOUT AGENTS THEY DEPLOYED IN THEIR OWN ENVIRONMENTS BEING TURNED AGAINST THEM. HIJACKED AGENTS NOW RANK ALONGSIDE AI-GENERATED PHISHING AND DEEPFAKES AS THEIR TOP CONCERN. CONFIRMED INCIDENTS FELL WHILE SUSPECTED INCIDENTS ROSE SHARPLY. TO ME, THAT RAISES A QUESTION ABOUT HOW WELL SECURITY TEAMS CAN ESTABLISH WHEN AI WAS INVOLVED."

TAL MORGENSTERN
PARTNER, LIGHTSPEED

AGENTS IN THE ENTERPRISE

CURRENT USE OF AI AGENTS

WHICH STATEMENT BEST DESCRIBES YOUR COMPANY'S CURRENT USE OF AI AGENTS TODAY?




KEY INSIGHT

Agents are in production at 77% of surveyed companies, and 98% are using, piloting, or exploring them. Only 2% have no agent initiative at all. For most, deployment is still contained to a limited set of functions (56%); one in five runs agents across multiple business functions.

Companies whose security teams are mostly or fully confident in their AI skills are further along: 85% have agents in production, against 66% of companies with less confident teams.

AGENT INVENTORY

DOES YOUR ORGANIZATION MAINTAIN (OR PLAN TO MAINTAIN) AN INVENTORY OF THE AI AGENTS OPERATING IN YOUR ENVIRONMENT? (ASKED AMONG THOSE EXPLORING AGENTS OR WITH AGENTS IN PRODUCTION, N=196)



KEY INSIGHT

Nearly every company using or exploring agents keeps some inventory (94%), but only one in five calls it complete. Three in four are working from a partial list. Security teams need to know which agents are operating before they can govern their access.

AGENTS IN THE ENTERPRISE

HOW AGENTS AUTHENTICATE

WHICH BEST DESCRIBES HOW AI AGENTS AUTHENTICATE (OR WOULD AUTHENTICATE) TO SYSTEMS IN YOUR ENVIRONMENT?

(ASKED AMONG THOSE EXPLORING AGENTS OR WITH AGENTS IN PRODUCTION, N=196)



CRITICAL FINDING

Only a third of companies using or exploring agents select dedicated per-agent identities with scoped permissions as their current or intended approach. A quarter use or would use human credentials, which can give an agent the user's access and make its actions harder to distinguish in an audit log. Another 15% select shared service accounts, and 26% use a mix. Agent identity governance is the category CISOs most plan to evaluate this year (55%).


"A QUARTER OF THESE COMPANIES USE OR WOULD USE HUMAN CREDENTIALS FOR AGENTS. THAT CAN GIVE THE AGENT THE PERSON'S PERMISSIONS AND LEAVE SECURITY TEAMS STRUGGLING TO TELL WHO DID WHAT. AGENTS NEED THEIR OWN IDENTITIES AND LIMITS ON WHAT THEY CAN ACCESS. CISOs PUT IDENTITY GOVERNANCE FIRST ON THEIR LIST OF TOOLS TO EVALUATE."

GURU CHAHAL
PARTNER, LIGHTSPEED

ORGANIZATIONAL CONTEXT AND GOVERNANCE

SECURITY TEAM SIZE

FULL-TIME SECURITY PERSONNEL, 2026 VS. 2025




KEY INSIGHT

This year's respondents run smaller security teams than last year's: a median of 35 full-time security staff, against 50 in 2025, and only 3% with 100 or more. Nearly two-thirds (64%) run teams of under 50. Companies at \$1B+ in revenue report a median of 50; those between \$500M and \$1B, 25.

ROLE OF SECURITY IN AI GOVERNANCE

WHICH BEST DESCRIBES THE ROLE OF THE CYBERSECURITY FUNCTION IN YOUR COMPANY'S AI GOVERNANCE? 2026 VS. 2025



CRITICAL FINDING

Security is losing ground in AI governance. The share of companies where security drives the process fell from 33% to 21%, while the share where it is not formally involved rose from 13% to 21%. A year ago we wrote that security needed to move from participant to leader. This year, 57% describe security as one voice among many, alongside functions such as legal, compliance, and data.

Among CISOs who report no confirmed or suspected AI-related incident in the past year, 48% say security drives AI governance, against 13% of those who report one. Among CISOs whose teams are mostly or fully confident in their AI skills, 31% say security leads, against 9% of less confident teams. Security leadership is associated with fewer reported incidents and greater confidence; the survey can't establish whether it causes either.


"WHEN I LOOK AT HOW ENTERPRISES BUY, THE ORG CHART TELLS YOU MORE THAN THE BUDGET LINE. THIS YEAR, FEWER CISOS SAY SECURITY LEADS AI GOVERNANCE, AND AT A FIFTH OF THESE COMPANIES IT ISN'T FORMALLY INVOLVED. COMPANIES WHERE SECURITY LEADS ALSO REPORT FEWER INCIDENTS. FOR A FOUNDER SELLING INTO THIS MARKET, ONE OF THE FIRST QUESTIONS SHOULD BE WHO OWNS AI RISK AT THE ACCOUNT."

JUSTIN SHEN
PARTNER, LIGHTSPEED

AI IN SECURITY OPERATIONS

CURRENT AI TOOL DEPLOYMENT IN THE SECURITY STACK

WHICH AREAS OF YOUR SECURITY STACK ARE ALREADY LEVERAGING AI TOOLS INTERNALLY?
(SELECT ALL THAT APPLY)



KEY INSIGHT

AI is in use somewhere in the security stack at 99% of companies. AI-SPM leads at 43%, just ahead of SOC/SIEM and vulnerability management (42% each). Agentic SOC, where AI analysts triage and investigate on their own, is running at 30%. Non-human identity management has the lowest adoption at 26%, despite CISOs' concern about agent credentials.

CISOs with confident teams are ahead on AI-SPM (49%, versus 34% of less confident teams), and CISOs who say they have had no AI incident are more likely to run AI in vulnerability management (55%) and IAM (57%).

AI IN SECURITY OPERATIONS


PLANNED AI SECURITY TOOL EVALUATIONS (NEXT 12 MONTHS)

CATEGORIES OF AI SECURITY TOOLS PLANNED FOR EVALUATION IN THE NEXT 12 MONTHS
(SELECT ALL THAT APPLY)



FUTURE FOCUS

Every CISO surveyed plans to evaluate at least one category of AI security tool this year. Agent identity, authorization and access governance leads at 55%, followed by AI-powered risk scoring and behavior analytics at 49% and runtime guardrails for agents at 48%. Identity and runtime controls address the access and visibility gaps in the agent findings. Last year's leading categories were model access governance and secure inference.



CRITICAL FINDING
AI-generated code has added to the application security workload at 87% of companies, and for one in five the increase is significant. No CISO reported a significant decrease, and none said their company does not deploy AI-generated code. Among companies with agents in production, 92% report more AppSec work, against 70% of those still piloting. Only 36% of companies use AI in application security.

“THE SPEED AT WHICH AI AGENTS CAN INFILTRATE AN ENVIRONMENT DEMANDS A FUNDAMENTAL CHANGE IN HOW SECURITY TEAMS DETECT AND RESPOND TO ATTACKS. CISOS NEED TO RETHINK HOW QUICKLY THEIR TEAMS CAN IDENTIFY A THREAT, INVESTIGATE IT AND TAKE ACTION. THERE’S A SIGNIFICANT OPPORTUNITY FOR TOOLS THAT SHORTEN THE TIME FROM DETECTION TO CONTAINMENT.”


ANOUSHKA VASWANI,
PARTNER, LIGHTSPEED

AUTONOMOUS RESPONSE


AI EXECUTING RESPONSE ACTIONS WITHOUT HUMAN APPROVAL

ARE ANY SECURITY RESPONSE ACTIONS (E.G. BLOCKING TRAFFIC, ISOLATING ENDPOINTS, REVOKING CREDENTIALS) EXECUTED AUTONOMOUSLY BY AI IN YOUR ENVIRONMENT TODAY, WITHOUT PRIOR HUMAN APPROVAL?


Yes, broadly across our response workflows:




Yes, for a narrow set of low-risk actions:



No, a human approves every response action:



No, and we have no plans to allow it:



CRITICAL FINDING

AI is already taking response actions without human sign-off at 88% of companies. Most (71%) limit that autonomy to a narrow set of low-risk actions, while 17% allow it broadly across their response workflows. Only 12% require a human to approve every action, and 1% have ruled autonomy out.

Companies with confident teams appear to allow more autonomy (92% versus 82%), and the share requiring human approval for every action is more than twice as high among teams that emerge less confident (18% versus 7%).

WHAT WOULD IT TAKE

We asked CISOs, in their own words, what would have to be true for them to let an AI agent take containment or response actions without human approval. In a rough tally of the 200 answers, six in ten named proven accuracy, a track record or testing. Roughly three in ten named scoped permissions or predefined limits on what the agent can touch, about a quarter named audit logs and monitoring, and about one in five named rollback or reversibility.


AUTONOMOUS RESPONSE

Selected responses.	
"ONLY WITH BOUNDED LOW-IMPACT BLAST RADIUSSES, DETERMINISTIC ROLLBACKS, STRICT SCOPE CONSTRAINTS, REAL-TIME LOGGING AND EXCEPTIONALLY LOW FALSE POSITIVE ERROR RATES."	"AGENT CAN ONLY ACT WITHIN A NARROW SCOPE LIKE ISOLATING A COMPROMISED ENDPOINT OR BLOCKING MALICIOUS TRAFFIC. ANYTHING BIGGER STILL NEEDS HUMAN APPROVAL."
"THE AGENT WOULD NEED TO PROVE ITS DECISIONS ARE SAFER THAN WAITING FOR A HUMAN ANALYST."	"THE AI'S HANDLING LOGIC IS COMPLETELY TRANSPARENT AND EXPLAINABLE; THE SCOPE OF RISK IMPACT IS STRICTLY LIMITED; AND ALL AI OPERATIONS CAN BE TERMINATED WITH A SINGLE CLICK IN THE EVENT OF AN ANOMALY."
"IT WOULD HAVE TO ACHIEVE A PROVEN ZERO FALSE POSITIVE TRACK RECORD DURING A LONG TRIAL PERIOD IN OUR SANDBOX."	"REVERSIBILITY IS VERY IMPORTANT. IF THE AI COULD BE UNDONE, OR DIDN'T HAVE A LARGE EFFECT ON SOMETHING, THEN I WOULD BE MORE WILLING TO LET IT AUTOMATE SOME TASKS."
"VENDOR ACCEPTS LIABILITY FOR ANY AGENT-CAUSED DAMAGE."	"CONTAINMENT MUST BE STRICTLY ISOLATED FROM PHYSICAL MACHINERY CONTROL, PREVENTING ANY AUTOMATED RESPONSE FROM HALTING FACTORY LINES."
"ACTIONS MUST BE CONFINED ENTIRELY TO SANDBOXED BUILDS, REVOKING MALICIOUS ARTIFACTS WITHOUT BLOCKING CORE CI/CD RELEASE PIPELINES."	"SAFE AUTONOMOUS RESPONSE CAPABILITIES SHOULD BE VERIFIED BY INDEPENDENT VALIDATION."
"THERE SHOULD BE CLEAR RULES FOR HIGH-RISK SITUATIONS. I'D BE COMFORTABLE WITH AUTOMATIC CONTAINMENT FOR OBVIOUS THREATS, BUT ANYTHING INVOLVING SENSITIVE DATA OR MAJOR BUSINESS IMPACT SHOULD STILL REQUIRE A HUMAN."	

VENDOR STRATEGY AND PROCUREMENT


IMPORTANCE OF A VENDOR'S AI STRATEGY

HOW IMPORTANT WILL A VENDOR'S AI STRATEGY BE IN YOUR EVALUATION PROCESS IN 2026-27? 2026 VS. 2025



CONSOLIDATION VS. POINT SOLUTIONS

ARE YOU MORE LIKELY TO CONSOLIDATE VENDORS OR ADOPT NEW POINT SOLUTIONS IN RESPONSE TO AI-DRIVEN NEEDS? 2026 VS. 2025



KEY INSIGHT

A vendor's AI strategy still matters to nearly every CISO (98% call it at least somewhat important), but it matters less than it did. The share rating it very or critically important fell from 82% to 65%, and the share calling it critically important dropped from 23% to 13%. CISOs may now expect AI capability from every vendor, or they may be more skeptical of vendors' claims. The survey doesn't establish why the ratings fell. CISOs at larger companies assign less importance to vendor AI strategy: 57% of those at companies with 2,500 or more employees rate it very or critically important, against 72% at smaller companies. CISOs with confident teams are more likely to give those ratings (78%) than those with less confident teams (48%).

SPLIT STRATEGY

CISOs remain almost evenly divided between AI-native point solutions (53%) and consolidation into fewer platforms (48%). Preferences differ by incident history: 60% of CISOs who report a confirmed or suspected AI-related incident favor point solutions, while 73% of those who report neither favor consolidation. Among consolidators, 25% call the AI threat severe, compared with 6% of point-solution buyers.


"THE CONSOLIDATION-VERSUS-POINT-SOLUTION SPLIT HAS BARELY MOVED IN A YEAR. ABOUT HALF OF CISOS STILL FAVOR SPECIALIST PRODUCTS, AND THAT RISES TO 60% AMONG THOSE WHO REPORT AN AI INCIDENT. CISOS WHO REPORT NO INCIDENT ARE MORE LIKELY TO FAVOR CONSOLIDATION. THERE'S ROOM TO BUILD A LARGE COMPANY ON EITHER SIDE, AS LONG AS YOU KNOW WHICH BUYER YOU'RE SELLING TO."

JAMES EPHRATI
PARTNER, LIGHTSPEED

BUDGET AND INVESTMENT OUTLOOK

CYBERSECURITY BUDGET CHANGES

EXPECTED CHANGE IN CYBERSECURITY BUDGET OVER THE NEXT 12 MONTHS, 2026 VS. 2025




BUDGET MOMENTUM

85% of CISOs expect their budgets to grow in the next 12 months, down slightly from 88%. The shift is from significant to moderate growth: the share expecting a significant increase fell from one in five to 12%. Only 3% expect any cut. Smaller companies are more bullish, with 91% of those under 2,500 employees expecting an increase against 77% of larger ones, and 92% of companies between \$500M and \$1B in revenue against 79% of those at \$1B+.

AI BUDGET ALLOCATION

PERCENT OF SECURITY BUDGET ALLOCATED TO AI-RELATED TOOLS OR INITIATIVES, 2026 VS. 2025



KEY INSIGHT

Every CISO surveyed allocates some security budget to AI, and 82% allocate more than 5%. The 6%–10% band shrank from 59% to 52%, while the share allocating more than 10% rose from 27% to 31% and the share at 5% or under rose from 15% to 18%. Confident teams are more likely to commit over 5% (89%) than less confident ones (73%).

“FROM MY CONVERSATIONS WITH CISOS, I EXPECT SECURITY BUDGETS TO INCREASE SUBSTANTIALLY AS BOARDS GROW MORE CONCERNED ABOUT AI-DRIVEN ATTACKS. THE SURVEY CAPTURES CURRENT EXPECTATIONS, BUT THE URGENCY I’M HEARING SUGGESTS SPENDING COULD RISE MUCH FASTER AS COMPANIES RESPOND TO THESE THREATS.”

ANOUSHKA VASWANI
PARTNER, LIGHTSPEED

SKILLS AND TALENT

CONFIDENCE IN IN-HOUSE AI SECURITY SKILLS

CONFIDENCE THAT THE TEAM HAS THE IN-HOUSE SKILLS TO ASSESS AND MITIGATE AI-SPECIFIC SECURITY RISKS, 2026 VS. 2025

2025



2026



CONFIDENCE GAP

55% of CISOs are fully or mostly confident their team can assess and mitigate AI-specific risk, down from 71%. The share fully confident dropped from 16% to 9%, and the share only somewhat confident rose from 29% to 43%. Confidence is lower at larger companies (47% of those with 2,500+ employees are fully or mostly confident, versus 62% of smaller ones) and much lower among companies still piloting agents (35%) than those with agents in production (61%).

HIRING AND RESKILLING PLANS FOR AI SECURITY

ARE YOU PLANNING TO HIRE OR RE-SKILL STAFF SPECIFICALLY FOR AI-RELATED SECURITY WORK?

We already have:



Yes, within 6 months:



Yes, within 12 months



Yes, longer than a year from now



No plans to



SKILLS AND TALENT

TALENT URGENCY

99% of CISOs have hired or reskilled for AI security or plan to, and 85% plan to do so within the next 12 months. Only 12% say they already have. Half have already acted or plan to within six months. CISOs who have already hired or reskilled are concentrated among confident teams (17%, versus 6% of less confident teams) and companies with agents in production (16%, versus none of those still piloting).

"I THINK SOME OF THE DROP IN CONFIDENCE MAY COME FROM CISOS GETTING A CLEARER VIEW OF WHAT THEIR TEAMS CAN'T DO YET. NEARLY ALL HAVE HIRED OR RESKILLED FOR AI SECURITY, OR PLAN TO. HAVING BUILT AN ENTERPRISE INFRASTRUCTURE COMPANY MYSELF, I'D LOOK FOR PRODUCTS THAT GIVE A 35-PERSON SECURITY TEAM ACCESS TO EXPERTISE IT CAN'T EASILY HIRE."

GURU CHAHAL
PARTNER, LIGHTSPEED

IN THEIR OWN WORDS

WHAT CISOS WISH VENDORS AND STARTUPS WERE SOLVING


"What's one AI-related security challenge you wish vendors or startups were solving today?" In a rough tally of the 200 answers, about three in ten were about agents: their access, their permissions, their identities, and what they do when nobody is watching. About a quarter were about sensitive data leaking through AI tools. Prompt injection, AI-generated phishing and deepfakes, and visibility into what AI is doing each drew roughly one in ten.

Selected responses:	
"Easier management of agent identities and permissions so AI agents only access the systems and data they actually need."	"Automated software bill of materials and supply chain integrity attestation for open-weight model weights, dependencies and fine-tuning datasets."
"Better controls are needed to prevent AI agents from gradually expanding their access beyond their original permissions."	"Robust guardrails and telemetry for third-party plugins and connectors operating within OT/legacy SCADA interfaces."
"As multiple AI agents start working together, I think there needs to be more focus on securing those interactions. One compromised agent shouldn't be able to influence or compromise the others."	"Automatically identifying vulnerabilities created by AI-generated code before deployment."
"Securing dynamic agentic workflows by providing real-time runtime authorization and access control for multi-agent API calls and tools."	"I wish they would build tools that can accurately spot deepfake audio and video clones before they fool our team."
"I wish vendors had better tools for monitoring AI agents in real time without creating more noise for security teams."	"Knowing when to trust an AI alert and when it's just noise is still a real challenge."
"Vendors using real behavior to calculate agent trust scores."	"Vendors should be more open about the hazards associated with their own AI agents."
"Better safeguards against malicious instructions hidden inside documents, emails, and webpages."	"Secure data governance for AI tools used across hotel operations. Many teams adopt generative AI to draft guest communications and operational reports, creating risks of exposing PII and guest private data. We want vendors to build lightweight controls that prevent sensitive hotel guest data leakage into large language models."
"AI models and datasets should live in a trusted library like an app store where each one has a clear seal of authenticity so we know it's safe and untampered."	


THE AI SECURITY OPPORTUNITY MAP

WHERE CISO PAIN POINTS SIGNAL MARKET NEEDS

Last year's map identified five areas where CISOs needed better tools. Three of them, governance automation, workforce augmentation and shadow AI, remain concerns. This year, agents are in production at three in four companies, and CISOs rank hijacked agents alongside AI-generated phishing and deepfakes as their top threat. Among companies using or exploring agents, only one in five has a complete inventory and only a third select dedicated identities with scoped permissions. We see six areas for builders to help address these needs.



AGENT IDENTITY AND ACCESS GOVERNANCE



THE GAP

Among companies using or exploring agents, 34% use or would use dedicated identities with scoped permissions. 26% select human credentials, 15% shared service accounts, and 26% a mix. 41% of CISOs name ungoverned non-human identities and agent credentials among their top threat vectors. Only 26% use AI in non-human identity management, the lowest of any category in the security stack.


MARKET OPPORTUNITY

- Identity providers and directories built for agents, with per-agent credentials issued and revoked automatically
- Scoped, just-in-time authorization for agent actions, enforced at the tool and API layer
- Identity graphs that map which humans, services and agents can reach which data
- Detection of permission creep, where an agent's access expands past its original grant
- Attribution that separates agent actions from the human whose credentials it may be using

VALIDATION

Agent identity, authorization and access governance is the top category CISOs plan to evaluate in the next 12 months, at 55%. In the rough tally of open-ended responses, agent access and permissions was the most common request.

AGENT RUNTIME SECURITY



THE GAP

51% of CISOs rank compromised or hijacked agents acting inside their environment among their top threats, tied for first. 40% cite the agent tool and integration supply chain (MCP servers, plugins, connectors), and 36% cite prompt injection. Only 20% of companies using or exploring agents have a complete inventory.


MARKET OPPORTUNITY

- Discovery and inventory of agents, including those deployed outside the security function's view
- Runtime guardrails that constrain what an agent can do regardless of what it is told
- Prompt injection detection for indirect attacks arriving through documents, email and web content
- Security for the agent supply chain: vetting, sandboxing and monitoring of MCP servers, plugins and connectors
- Behavioral monitoring and trust scoring for agents, including agent-to-agent interactions

VALIDATION

48% plan to evaluate runtime guardrails for agents in production, 46% model firewalls and API-layer security, and 44% LLM red teaming tools. CISOs asked, in their own words, for "guardrails and telemetry for third-party plugins and connectors" and for ways to ensure "one compromised agent shouldn't be able to influence or compromise the others."

TRUSTED AUTONOMOUS RESPONSE



THE GAP

88% of companies already let AI take some response actions without human approval, but 71% limit it to a narrow set of low-risk actions. Asked what they would require for autonomous response, CISOs named proven accuracy and a track record (cited by roughly six in ten), scoped and predefined actions, audit trails, and reversibility. 30% already run an agentic SOC.


MARKET OPPORTUNITY

- Autonomous response platforms with built-in blast-radius limits, rollback and one-click termination
- Confidence scoring and evidence trails that explain each action before and after it is taken
- Simulation and sandbox environments where an agent earns autonomy on recorded performance
- Independent validation and benchmarking of autonomous response accuracy
- Policy engines that define which actions an agent may take alone and which escalate to a human

VALIDATION

Only 12% of CISOs require a human to approve every response action, and 1% have ruled out autonomy. Vendors can address CISOs' stated conditions for autonomous response by demonstrating accuracy, logging and reversibility.

APPLICATION SECURITY FOR AI-GENERATED CODE



THE GAP

87% of CISOs say AI-generated code has increased their application security workload, 20% significantly. None reported that their company does not deploy AI-generated code. Only 36% use AI in application security.


MARKET OPPORTUNITY

- Code review and vulnerability detection tuned to the failure modes of AI-generated code
- Security checks embedded in AI coding tools, so flaws are caught where the code is written
- Provenance tracking that records which code was generated, by what, from what prompt
- Policy enforcement for what AI-generated code may touch before it ships
- AI-driven AppSec that scales review capacity with the volume of code arriving

VALIDATION

More CISOs report increased AppSec work at companies with agents in production (92%) than at those still piloting (70%). One CISO described the need as "Automatically identifying vulnerabilities created by AI-generated code before deployment."

AI DATA LEAKAGE AND SHADOW AI CONTROL



THE GAP

42% of CISOs name shadow AI and unsanctioned employee usage among their top threat vectors. In the open-ended responses, sensitive data leaking through AI tools was the second most common challenge CISOs wished vendors would solve, after agent access. 34% have AI applied to DSPM and data security.


MARKET OPPORTUNITY

- Discovery and control of employee AI tool usage across browsers, endpoints and SaaS
- Data loss prevention built for prompts, uploads and agent tool calls
- Policy enforcement for what data may reach which models, including third-party and self-hosted
- Monitoring of what agents read, retain and send to outside systems
- Lightweight controls for regulated data in industries where frontline teams have adopted generative AI

VALIDATION

CISOs described the problem in specific operational terms, from "stopping sensitive business data from leaking through AI tools without employees realizing it" to preventing hotel guest data from reaching large language models.

DEFENSE AGAINST AI- GENERATED DECEPTION



THE GAP

AI-generated phishing and deepfakes tie with hijacked agents as the top threat category at 51%. Synthetic identity fraud is at 34%. Concern about phishing and deepfakes is highest among CISOs who report no confirmed or suspected AI incident (68%).

Market Opportunity

MARKET OPPORTUNITY

- Detection of AI-generated email, voice and video at the point of contact
- Identity verification that holds up against cloned voices and synthetic documents
- Behavioral analysis that distinguishes human from synthetic traffic and interaction
- Human risk platforms that train and test employees against AI-crafted social engineering
- Protection for high-value workflows (payments, access requests, executive communications) where impersonation pays

VALIDATION

49% plan to evaluate AI-powered risk scoring and behavior analytics in the next 12 months. CISOs asked for tools that "accurately spot deepfake audio and video clones before they fool our team" and that detect "highly convincing AI-generated phishing emails that easily bypass traditional spam filters."

CROSS-CUTTING THEMES

AGENT CONTROLS NEED TO WORK TOGETHER.

Identity, runtime security and autonomous response all require teams to know which agents are operating, what they can access and how to limit their actions. AI-generated code and data leakage add to the workload.

CONFIDENT TEAMS REPORT MORE ADOPTION AND INVESTMENT.

CISOs with confident teams are further along on agents (85% in production), more likely to have security leading governance (31%), more likely to run AI-SPM (49%) and more likely to have already hired or reskilled for AI security (17%). The survey doesn't establish whether those steps built confidence or confidence made them easier to take.

BUYING PREFERENCES DIFFER BY INCIDENT HISTORY.

The near-even split between point solutions and platforms has barely changed in a year. CISOs who report an AI incident favor specialists; those who report none favor consolidation.

MOST AUTONOMOUS RESPONSE REMAINS LIMITED.

88% already allow some autonomous response, 71% narrowly. CISOs most often named accuracy, scope, logging and reversibility as conditions for allowing AI to act without approval.

IMPLICATIONS FOR INNOVATION

- Address the agent risks CISOs identify. Hijacked agents (51%), ungoverned agent credentials (41%) and the agent supply chain (40%) concern two in five or more CISOs.
- Understand the buyer's incident history. CISOs who report an incident are more likely to favor a point solution (60%), to have hired or reskilled for AI security or plan to within six months (54%), and to already run an agentic SOC (35%).
- Build in controls for autonomous actions. Include the audit, rollback and scope controls CISOs named as preconditions.
- Help smaller teams handle the work. Median security headcount in this sample is 35, and 99% have hired or reskilled for AI or plan to.
- Show what the product does. The share of CISOs who rate a vendor's AI strategy critically important fell from 23% to 13%.

"THE SURVEY GIVES FOUNDERS SPECIFIC PROBLEMS TO WORK ON: IDENTIFYING AGENTS, CONTROLLING WHAT THEY CAN ACCESS, DETECTING WHEN THEY'RE COMPROMISED AND DECIDING WHICH ACTIONS THEY CAN TAKE ALONE. CISOs HAVE NAMED THE TOOLS THEY PLAN TO EVALUATE. I'D START WITH THOSE REQUIREMENTS.

TAL MORGENSTERN,
PARTNER, LIGHTSPEED

KEY CONCLUSIONS AND IMPLICATIONS

AGENT DEPLOYMENT IS WIDESPREAD; CONTROLS ARE INCOMPLETE.

77% of companies run AI agents in production. Among companies using or exploring agents, 20% have a complete inventory; 34% use or would use dedicated per-agent identities, and 26% select human credentials. CISOs rank hijacked agents alongside AI-generated phishing and deepfakes as their top threat and make agent identity governance their top evaluation priority.

SECURITY LEADS AI GOVERNANCE AT FEWER COMPANIES.

The share where security drives the process fell from 33% to 21%, and the share where it is not formally involved rose from 13% to 21%. Security leadership is associated with fewer reported incidents and greater confidence.

MORE INCIDENTS ARE SUSPECTED; FEWER ARE CONFIRMED.

79% report a confirmed or suspected AI-related incident, up from 75%, and 99% have at least detected attempts. Confirmed incidents fell from 37% to 16% while suspected ones rose from 38% to 63%. 71% rate the coming year's AI threat moderate or severe, and the share calling it severe rose from 10% to 15%.

MOST COMPANIES ALLOW LIMITED AUTONOMOUS RESPONSE.

88% already let AI take response actions without human approval; 71% only for low-risk actions. CISOs named accuracy, scoped permissions, audit trails, and reversibility as conditions for autonomous response.

CONFIDENCE IN AI SECURITY SKILLS FELL.

55% of CISOs are fully or mostly confident in their team's skills, down from 71%. 99% have hired or reskilled for AI security or plan to, and 85% plan to within the year. Teams are smaller than in last year's sample, with a median of 35 security staff.

BUDGETS CONTINUE TO GROW, MOSTLY AT A MODERATE PACE.

85% expect budget growth (down from 88%), with fewer expecting significant increases. 82% commit more than 5% of the security budget to AI, and the share committing more than 10% rose to 31%. The share rating a vendor's AI strategy very or critically important fell from 82% to 65%.

STRATEGIC IMPERATIVES FOR 2027

INVENTORY EVERY AGENT AND GIVE EACH ITS OWN IDENTITY.

Only one in five companies using or exploring agents has a complete inventory. A third select dedicated identities with scoped permissions as their current or intended authentication approach. Establish which agents are running and what each can access.

PUT SECURITY IN THE LEAD ON AI GOVERNANCE.

The share of companies where security leads fell from 33% to 21%, and at 21% it is not formally involved. Give security responsibility for agent permissions, monitoring and response.

TEST CONTROLS BEFORE EXPANDING AUTONOMY.

Require proven accuracy, scoped actions, full logging and rollback before allowing agents to take on more response work.

PLAN FOR AI-GENERATED CODE AS THE DEFAULT.

87% of security teams are carrying more AppSec work because of it. Review capacity has to scale with code volume.

HIRE OR RESKILL FOR THE WORK AHEAD.

85% of CISOs plan to hire or reskill for AI security within the year. Identify the expertise the team needs to govern agents and review AI-generated code.

METHODOLOGY

The Lightspeed CISO Survey was conducted by Wakefield Research among 200 U.S. CISOs at companies with a minimum of \$500M in annual revenue, between September 4th and September 14th, 2026, using an email invitation and an online survey. Where year-over-year comparisons appear, 2025 figures are from the Lightspeed CISO Survey conducted by Wakefield Research between July 29th and August 6th, 2025, among 200 U.S. CISOs at companies with a minimum of \$500M in annual revenue.

Results of any sample are subject to sampling variation. The magnitude of the variation is measurable and is affected by the number of interviews and the level of the percentages expressing the results. For the interviews conducted in this study, the chances are 95 in 100 that a survey result does not vary, plus or minus, by more than 6.9 percentage points from the result that would be obtained if interviews had been conducted with all persons in the universe represented by the sample. Subgroup findings with base sizes under 100 are directional.

Open-ended responses quoted in this report are lightly edited for spelling and punctuation.

PRESENTED BY  | 

FORTUNE

CYBER60™

LSVP.COM/CYBER60

CYBER60@LSVP.COM