Consultiply

• AI

MIP Resource Center

Put AI agents, approval lines and fix evidence into your client reviews.

a16z surveyed 100 CISOs at companies with more than 1,000 employees and found that the 2027 security agenda is built around AI. This guide covers what that signals for MSP owners and four changes to consider for your practice.

Source: a16z, The New Security Agenda, Oct 28, 2026. Independent reading; not affiliated with or endorsed by a16z.

What a16z found

On Oct 28, 2026, Andreessen Horowitz published The New Security Agenda, by Stephenie Zhang and Owen Colby. In a16z's words: "we surveyed 100 CISOs at companies with more than 1,000 employees about their priorities, budgets, and buying plans for 2027."

AI now drives the list. In the a16z survey, 87% expect cybersecurity spending to increase. a16z says the main drivers are AI-enabled threats and the larger attack surface from respondents' own AI deployments.

When asked to recall their 2026 priorities, respondents put exposure and vulnerability management and SIEM at the top, a16z reports. For 2027, a16z says all seven of the most-selected initiatives involve AI. The biggest gains are internal AI agent security, SOC automation and offensive security, with code security, remediation and AI governance close behind.

Headcount is mostly holding. In the a16z survey, nearly half expect security headcount to increase in 2027, another half expect it to stay flat, and only 4% expect to cut.

The AI is already in the building

AI use is already well ahead of AI security tooling. In the a16z survey, 92% report at least limited production use of employee generative AI tools, 83% report production use of coding assistants or agents, and 55% have MCP servers or other agent-to-application integrations in production.

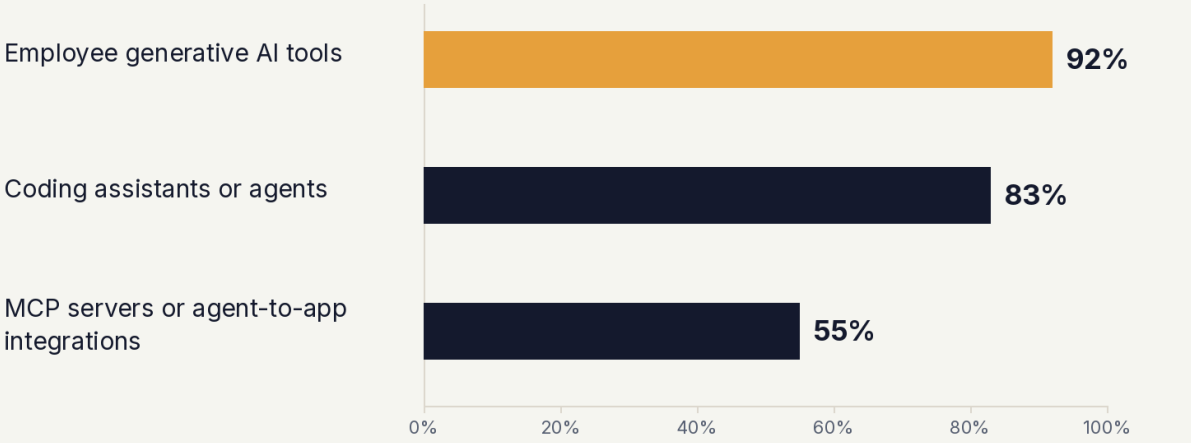
AI-based security tools trail far behind. a16z reports that only 18% use AI SOC tools in production, and 13% use either AI offensive security or AI exposure and vulnerability remediation in production.

a16z quotes a security leader at a restaurant software company who asked the question many teams have: “How do you really know what the agents are really doing?”

a16z's answer is that security teams need to know four things: which tools and agents are running, who owns them, what they can access, and what they actually do.

AI is already in production

Share of CISOs reporting at least limited production use



Source: a16z, The New Security Agenda (Oct 28, 2026), survey of 100 CISOs at companies with 1,000+ employees. Consultiply chart of a16z figures.

Consultiply chart of a16z's published figures: production use of employee generative AI tools 92%, coding assistants or agents 83%, MCP servers or agent-to-application integrations 55%. Source: a16z survey of 100 CISOs, Oct 28, 2026

transform@consultiply.ai

a16z's numbers, and what they are

All figures below are from the a16z survey and are a16z's own.

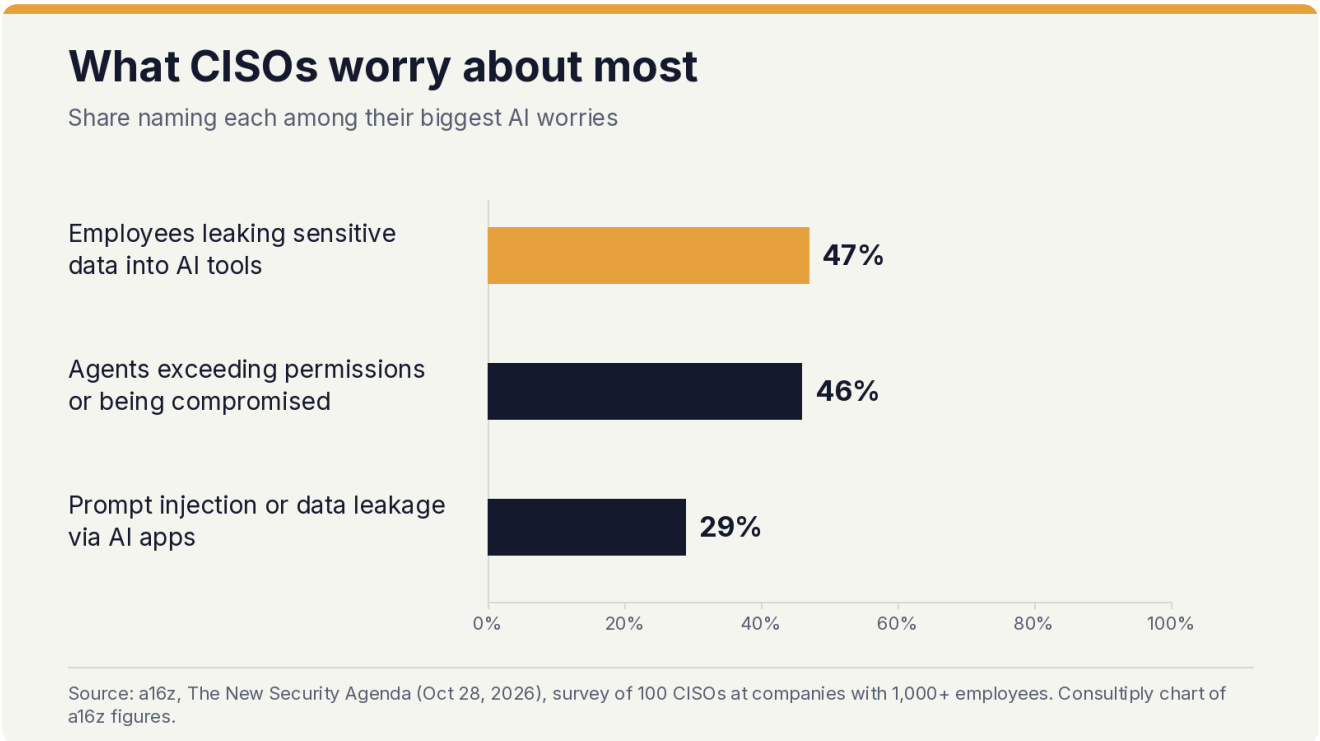
Claim	Figure	Status
Expect cybersecurity spending to increase	87%	a16z survey
At least limited production use of employee generative AI tools	92%	a16z survey
Production use of coding assistants or agents	83%	a16z survey
MCP servers or other agent-to-application integrations in production	55%	a16z survey
Agents exceeding permissions or being compromised among biggest worries	46%	a16z survey
Responses describing security slowing at least some AI adoption	63%	a16z survey
Want to retest and verify vulnerabilities were fixed	70%	a16z survey
Plan new purchase or upgrade for internal AI agent security within twelve months	45%	a16z survey

Why this matters to an MSP

a16z's sample is 100 CISOs at companies with more than 1,000 employees. Most MSP clients are smaller. Treat these numbers as a signal of where security buying and expectations are heading. They are not data about your SMB clients.

The signal is still close to your clients. They use the same kinds of tools: generative AI in the browser, coding assistants, and connectors that let an agent reach email, files and business apps. The worries a16z recorded attach to those same kinds of tools. In the a16z survey, 47% include employees leaking sensitive data into AI tools among their biggest worries, 46% cite agents exceeding their permissions or being compromised, and 29% name prompt injection or data leakage through AI applications.

The survey also shows friction. In the a16z survey, 63% of responses describe security slowing at least some AI adoption, through reviews, approval requirements and limits on particular uses. Our reading: clients who feel that friction need rules they can follow, not blanket bans. An MSP that already knows what runs in each environment is better placed to help write those rules.



Consultiply chart of a16z's published figures: employees leaking sensitive data into AI tools 47%, agents exceeding permissions or being compromised 46%, prompt injection or data leakage through AI applications 29%. Source: a16z survey of 100 CISOs, Oct 28, 2026

transform@consultiply.ai

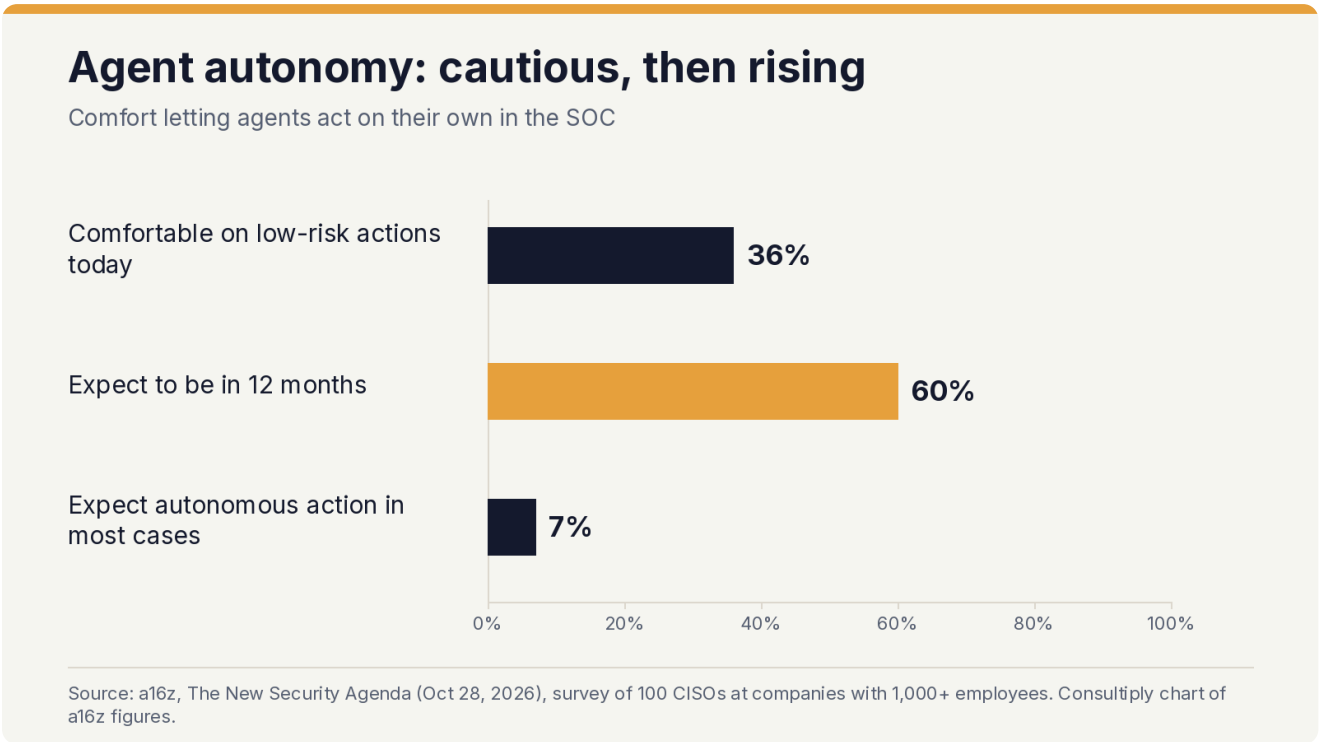
Decide what agents may do alone

a16z asked how much authority to give agents in the SOC. In the a16z survey, 36% are comfortable today letting agents act on their own for low-risk actions, 60% expect to be comfortable with that in 12 months, and only 7% expect to allow autonomous action in most cases.

a16z puts the line this way: "Teams can delegate much of an investigation while retaining approval for a response that could disrupt the business."

Identity is the other half of the problem. Lightspeed's survey of 200 CISOs at companies with $500M+ in annual revenue found that only a third of companies using or exploring agents use, or would use, dedicated agent identities with scoped permissions. The Lightspeed survey also reports that compromised or hijacked agents tie with AI-generated phishing and deepfakes as the top AI threat, at 51% each.

The practical point for an MSP is that an agent running on a person's credentials is hard to scope and hard to audit. You can't decide what an agent may do alone until you know what it can reach and whose access it is using.



Consultiply chart of a16z's published figures: comfortable with autonomous agent action on low-risk actions today 36%, expected in 12 months 60%, autonomous action in most cases 7%. Source: a16z survey of 100 CISOs, Oct 28, 2026

transform@consultiply.ai

Close the loop from find to fix

Finding vulnerabilities is getting faster, and fixing them is not keeping up. a16z quotes a fintech CISO: “With a lot of these AI tools, you’re able to find vulnerabilities faster, but you’re not able to patch faster.”

The a16z survey shows what buyers want: 70% want to retest and verify that vulnerabilities were fixed, 67% want to prioritize remediation using validated exploitability, and 57% want fixes or remediation tickets generated automatically.

Testing is also speeding up. In the a16z survey, 83% rate more frequent security posture testing as important. 18% say most of their offensive testing runs continuously against production assets today, and 52% expect it to within a year.

The top request is proof that a fix worked. Closing a ticket doesn't show that; a retest does. For a related read on review as the bottleneck, see our guide on the agent review bottleneck.

Four moves for your practice

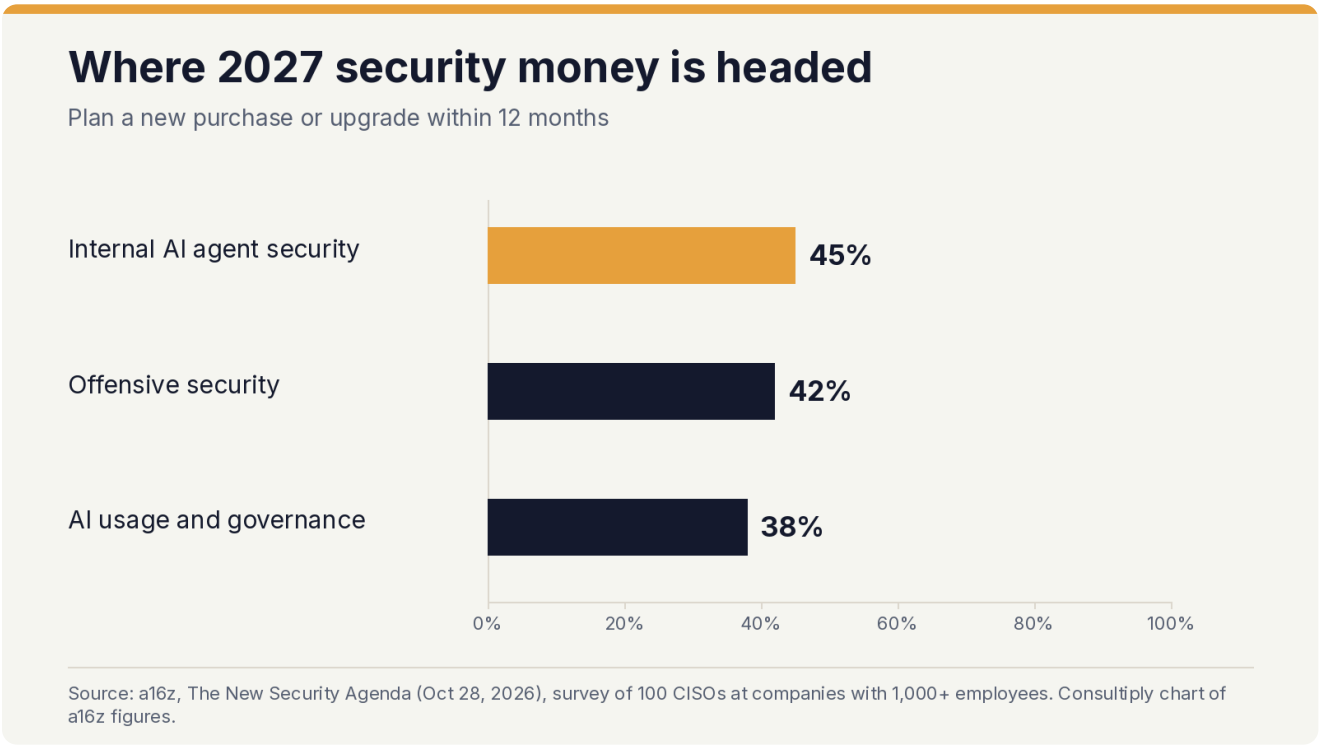
These moves are Consultiply's reading of the a16z data. They are not a16z's recommendations.

- **Inventory agents and what they reach**: For each client, list the AI tools, agents, extensions and MCP servers in use, who owns each one, and which data and systems it can touch. a16z notes

that software approved as an ordinary application can gain agentic capabilities in its next update, so keep the list current.

- **Write an approval line for agent actions**: Agree with each client which agent actions can run alone, such as gathering logs, and which need a named person to approve, such as changing a production system. Write it down. The client owns the systems, so the client owns the line.
- **Retest and keep evidence of each fix**: Close a remediation ticket only after a retest, and store the retest result with the ticket. This matches the top remediation request in the a16z survey, at 70%.
- **Add AI usage and governance to the review cadence**: Put AI usage, agent permissions and approval rules on the agenda of every regular client review. Revisit the inventory and the approval line each time.

Buying is close at hand at the companies a16z surveyed. Within twelve months, 45% plan a new purchase or upgrade for internal AI agent security, 42% for offensive security and 38% for AI usage and governance, according to the a16z survey. a16z also found that 44% look first to features from an existing security platform. So check what your current stack already covers before you add a tool.



Consultiply chart of a16z's published figures: planned purchase or upgrade within twelve months for internal AI agent security 45%, offensive security 42%, AI usage and governance 38%. Source: a16z survey of 100 CISOs, Oct 28, 2026

transform@consultiply.ai

Read with care

- a16z is a venture investor that backs security startups, and the article promotes companies it has invested in. This guide does not name them.
- The sample is 100 CISOs at companies with more than 1,000 employees, per a16z. It is not a survey of small businesses or of MSPs.
- The 2026 priorities are what respondents recalled in this survey. They come from memory, not from a prior survey.

- The figures are a16z's own and have not been independently verified. The Lightspeed figures are Lightspeed's own and come from a different sample.

> This guide is Consultiply's independent reading of a16z's published survey. It is not affiliated with or endorsed by a16z, and it is not a recommendation to buy any product.

Sources

The New Security Agenda (Andreessen Horowitz, Stephenie Zhang and Owen Colby, Oct 28, 2026)

2026-27 Lightspeed Cyber 60 CISO survey (Lightspeed Venture Partners, author not listed, accessed Oct 29, 2026)

---

transform@consultiply.ai