by **Evan Schuman**
Contributor

What do CISOs need to rest easy about future AI risks?

Feature
Sep 7, 2026 • 7 mins

Security execs who feel prepared to face what's coming hold several organizational factors in common, according to a recent report. But confidence and preparedness aren't necessarily the true measures of security.

Credit: insta_photos / Shutterstock

Security leaders' confidence in their ability to navigate the security risks AI will pose over the next two years rests on several clear factors, according to IANS analysis of its AI Security Survey, fielded earlier this year.

Of the 113 CISOs IANS surveyed in April and May, 41% expressed optimism about their organization's ability to manage AI security risks over the next 24 months, versus 38% who were pessimistic.

Delving deeper, IANS identified six strong organizational signals that differentiate CISOs who feel confident about their ability to address rapidly evolving AI risks and those CISOs who dread what's to come. And they have less to do with current controls and capacities and more to do with the CISO's standing and sway within the organization.
"The findings reveal a distinction between present-day risk perceptions and optimism about the manageability of AI risk in the future. Perceptions of present-day risk are associated most strongly with AI security maturity — an organization's operational capacity and security controls to manage its AI environment," the report said.
"Optimism about the future is associated more strongly with organizational readiness factors."

Those factors include:

- Leadership's understanding of AI risk
- Clearly defined AI governance ownership
- The security team's effectiveness with AI tools
- CISO ownership over the AI-security budget
- Sustainable security team workloads
- Sufficient security staffing

Above any present-day measure of AI security maturity or budgetary increases, CISO optimism about the hard road ahead appears to rest largely on whether organizational leadership listens to, empowers, and adequately resources them — at least according to IANS' poll.

sensationalized agentic attacks

By Lee Rossey • 6 mins

AI incidents need a new playbook. Here's how to build one

By Julie Brunias • 7 mins

[https://www.csoonline.com/?post_type=post&p=4196303&utm_source=miso&utm_medium=related&utm_campaign=thumbnail_list]

Why AI governance is failing — and what actually works

By Stephanie Overby • 7 mins

[https://www.cio.com/?post_type=post&p=4201343&utm_source=miso&utm_medium=related&utm_campaign=thumbnail_list]

A question of confidence

Cybersecurity consultants and analysts largely agree those elements will play a critical role in a company's AI cybersecurity effectiveness going forward, but with nuance and caveats.

Pearl Almeida [https://www.infotech.com/profiles/pearl-almeida], a research director at Info-Tech Research Group, underscores misalignment and leadership misunderstanding of AI risks as key factors that undercut a CISOs' ability to mitigate those risks.

"The risk-averse CISOs who resisted agent adoption have mostly been overruled by the business, and they [business executives] are moving forward whether or not they feel ready," Almeida says. "Only a handful of organizations I work with have strong controls in place. The rest are starting with fundamentals because deferring is no longer an option."
Moreover, many CISOs are contending with what Almeida described as "organizational blind spots."

CSO Smart Answers Learn more

Explore related questions

- How can a "data-first" design improve your enterprise AI governance framework?

- Who should hold the "pause authority" for your company's AI systems?
- What are the most common use cases for generative AI in security?
- Why is an AI risk register insufficient for your incident response plan?
- What are the seven essential competencies for security engineers using AI?

Ask a question

Ask

“In my conversations with CISOs, I have found that most leaders, security ones included, cannot explain how an agent works. They know that they need to use it, but don’t understand the mechanics: where the agent gets its information, how it calls on tools, how it makes decisions, how human-in-the-loop actually works,” Almeida says. “When you can’t describe the mechanism, you can’t right-size the exposure.”

As for AI governance, Almeida cautions that clear ownership isn’t enough on its own; the real issue is security culture.

“Assigning a named business owner only works if security is already embedded in how the business operates rather than bolted on at approval,” she explains. “Owners who won’t take accountability should understand the trade plainly: They are handing security the authority to shut their workflow down during an incident without consulting them.”

Rock Lambros [https://zenity.io/authors/rock-lambros], director of AI standards and governance at agentic security vendor Zenity, cautions that optimism can be misleading.

“Every number in the report is one person’s self-assessment measured against that same person’s self-assessment,” Lambros notes. “A CISO who just told you their program is ‘efficient’ on a five-point scale isn’t going to turn around two questions later and rate their risk a nine. People stay consistent with themselves.”
As for the factors identified, Lambros noted two clear groupings.

“Leadership understanding, clear governance ownership, and CISO budget control all rise and fall together because they’re all downstream of a single decision somebody made in the C-suite or the boardroom,” Lambros says, adding that workload and staffing are closely aligned, “measuring whether the team has room to breathe.”

The practicalities of preparedness

Sanchit Vir Gogia [https://greyhoundresearch.com/svg/], chief analyst at Greyhound Research, notes that, while the factors IANS identified may be indicators of organizational readiness, "readiness is not the same measure as security."

"They show whether the CISO has leadership understanding, budget control, and capacity to act, not whether the AI estate is under control," Gogia says.

Aman Mahapatra [https://www.linkedin.com/in/akm76/], chief strategy officer for Tribeca Softtech, a New York City-based technology consulting firm, echoes Gogia's concerns.

"A CISO with an informed board, clean governance ownership, budget control, and a fully staffed team will feel optimistic whether or not a single agent in production has bounded authorization," Mahapatra says. "That is a useful map of how CISOs currently think. It is not a security roadmap, and the risk is that it gets read as one." How a security team uses AI internally, however, provides practical indication of preparedness, Mahapatra suggests, because it "builds durable capability rather than favorable conditions."

And it's something CISOs should emphasize, as it "has to be earned over quarters, which is exactly why it compounds," he says. "Teams running AI in their own triage and detection tuning learn experientially where prompt injection bites and how agents fail under adversarial input, and that knowledge transfers directly to securing AI everywhere else."

Cybersecurity consultant Brian Levine [https://formergov.com/directory/brianlevine], executive director of FormerGov, advises CISOs to ensure that AI governance ownership includes a focus on third-party risk.
"Governance is useless if you don't actually know which models are being pulled into your environment and the risks associated with those models," Levine says. "For a CISO, the risk isn't just the AI you build. It's the AI you inherit via vendors and unauthorized use."

Justin Greis [https://acceligence.com/talent/profiles/justin-greis/], CEO of consulting firm Acceligence, says future AI risk preparedness hinges on gaining as much intelligence about AI system data and app interactions as possible.

"The CISO needs to know what identity that agent is operating under, what systems it can reach, what information it can retrieve, what transactions it can initiate, what other agents or tools it can invoke, whether every material action is observable and whether the organization can immediately terminate its authority," Greis says. "That starts looking much closer to managing a digital workforce than securing another software application."
Greis also was skeptical — along with various analysts and consultants — about whether a 41% optimism rate truly reflects what most enterprise CISOs are experiencing.

“Almost half of the CISOs are optimistic about the AI security future? Who the heck are these people?” Greis asks. “The enterprise CISOs I talk with are white-knuckling their way through the fastest-moving security challenge of their careers.”

And that begs the question of timing: Asked in April and May about their level of confidence, would those CISOs answer the same way given what they’ve seen since?

It’s possible that the further advancement of Mythos-class models and high-profile rogue acts by agents may be shaking that optimism. But it’s also true that AI amplifies foundational security shortcomings, with organizational dysfunction chief among them.

Artificial Intelligence[https://www.csoonline.com/artificial-intelligence/] Risk Management[https://www.csoonline.com/risk-management/] Security[https://www.csoonline.com/security/] CSO and CISO[https://www.csoonline.com/cso-and-ciso/] IT Governance[https://www.csoonline.com/governance/] IT Leadership[https://www.csoonline.com/it-leadership/]