



The Mind of the CISO

The New Insider: Rethinking Risk in the AI Era

How AI agents and autonomous systems are reshaping insider risk, creating new visibility and accountability gaps, and driving the need for machine-speed detection and defense.



Contents

Introduction..... 3

Key Findings..... 4

Quantitative respondents..... 5

Qualitative respondents..... 6

Section 1:
AI is redefining the insider 7

Section 2:
The visibility gap 9

Section 3:
The cost of trusted access..... 12

Section 4:
AI raises the stakes..... 15

Section 5:
Insider risk reshapes the CISO role 17

Conclusion 19

Recommendations..... 20

Additional resources..... 21

Boilerplates..... 23



Introduction

What happens when the insider is no longer always human? Employees, contractors, and other trusted users remain central to insider risk, but AI agents, autonomous systems, and machine identities are increasingly joining them. For CISOs, this expands an already difficult problem: more activity to understand, more ways for sensitive data to move, and more pressure to distinguish legitimate behavior from genuine risk. With AI agents acting on behalf of a user, it becomes difficult to distinguish between the two.

This edition of The Mind of the CISO explores how the shift toward AI as the new insider changes how organizations approach insider risk. It examines the visibility gaps emerging as AI becomes part of daily work, the consequences when trusted access is misused or goes wrong, and the growing role AI is playing on both sides of the security equation. It also examines what this means for CISOs, who are being asked to manage a broader set of risks, provide clearer answers to boards, and take on more accountability when incidents occur.

Inside this report:

- AI is redefining the insider
- The visibility and accountability gap
- The cost of trusted access
- AI raises the stakes
- Insider risk reshapes the CISO role

Across these themes, one issue is clear: traditional assumptions about trusted access are being tested. AI can help employees work faster and provide security teams new ways to detect and respond to threats, but it can also introduce activity existing controls were not designed to see or govern. At the same time, attackers are using the technology to move faster, exploit weaknesses, and take advantage of the trusted identities and systems organizations rely on.

For CISOs, navigating this change will require judgment as much as technology. The critical decisions, like where AI can be trusted, where oversight is needed, how risk should be explained to the board, and what good governance looks like as the technology evolves, cannot be deferred. This report examines how organizations are approaching those decisions today and where change is needed to embrace AI without losing control of the trusted access it depends on.


From emerging risk to real-world reality

Case Study: In June 2026, an agent model conducting research into public medicine spending gained unauthorized access to an Australian Government Medicare statistics portal, accessing both public and non-public files. No personal medical information is believed to have been accessed. Still, the incident raises a bigger question: As AI agents take on more responsibility, how do organizations maintain visibility and control over their actions? And are existing security strategies ready for trusted activity that no longer comes from people alone?


> Case Study Reference

Key Findings

AI is redefining the insider




78% of organizations include actions taken by AI agents, autonomous systems, or machine identities within their definition of insider risk, and 42% say it's among the most important security risk for their organization to address over the next 12–24 months.




52% do not yet have a fully implemented insider-risk strategy, despite 89% having increased their focus on insider risk over the past 12 months.

The visibility gap



84% agree they need to improve their ability to detect risks related to AI tools and non-human identities.



74% agree AI has introduced risks traditional security controls were not designed to address.

The cost of trusted access



Insider risk accounted for 39% of all security incidents experienced by organizations over the past 12 months, on average.




95% of organizations experienced an insider incident linked to the actions of an AI or autonomous system.


On average, organizations take **15 hours** to detect and respond to an insider-related incident.

Insider-related security incidents cost organizations **\$2.4 million** on average over the past 12 months.

AI is reshaping both sides of the battlefield




79% say they expect their organization to be targeted by an AI-assisted attack exploiting trusted users, identities, or systems over the next 12 months.




99% say AI and automation will be essential to managing insider risk over the next 12–24 months.

41% plan to prioritize investment in detection and response for AI-related or non-human identity risks.

Insider risk reshapes the CISO role



92% say their responsibility for managing insider risk has increased over the past 12 months, while 91% experience growing accountability for insider-related incidents.



69% are concerned about their own liability in the event of an insider-related security incident.


“The whole insider risk universe is expanding. As well as employees, we have third-party contractors and AI—but AI is newer and potentially more dangerous.”

— Director of Security Operations, Public Sector, Saudi Arabia (Survey Response)


Quantitative respondents

600 CISOs (or equivalent) were surveyed in August and September 2026, split in the following ways:



1. In four regions:
America, Europe,
Middle East & APAC



2. Across sectors, including finance,
healthcare, public sector, energy,
manufacturing, etc



3. Were from organizations with
more than 1,000 employees




All interviews were conducted using a rigorous multilevel screening process to ensure only suitable candidates participated.


Qualitative respondents

We conducted five qualitative interviews with CISOs (or equivalent) in September 2026, split in the following ways:

... by region



In the US, UAE and Singapore



x 3



x 1



x 1


... by organizational sector



Financial Services x 3



Utilities x 1



Manufacturing x 1

... by organizational size



1,000 - 2,999 employees x 3



5,000+ employees x 2

All interviews were conducted using a rigorous multilevel screening process to ensure only suitable candidates participated.

Section 1

AI is redefining the insider

The insider is no longer just human

Malicious intent and human error have long defined insider risk, and both remain significant threats. But AI is changing the equation. As AI agents proliferate across the enterprise, the very nature of the insider is changing.


The shift is evident in how organizations define insider risk today. Around eight in ten (78%) now include actions taken by AI, autonomous systems, or machine identities in their definition, alongside malicious, negligent, or accidental human actions.

How organizations define insider risk

Actions taken by AI agents, autonomous systems or machine identities operating with legitimate access



Malicious actions by employees, contractors or trusted third parties



Accidental or negligent actions by employees, contractors or trusted third parties



Figure 1. Which of the following does your organization include within its definition of “insider risk?” Base: 600

This broader definition comes as insider risk moves higher up the security agenda. Around two-thirds (64%) of CISOs believe it will be among the most important security risks their organization needs to address over the next 12–24 months—ahead of external threats (59%). Priorities span both human and non-human risk. More than half (55%) rank human insider risk among their most important security risks, while 42% point to non-human risks from AI tools, autonomous agents, and machine identities.


Organizations must now manage malicious intent and human error alongside AI agents that can act on behalf of trusted users with increasing autonomy and scale. This expansion is already causing concern, with three-quarters (75%) of CISOs at least moderately concerned about risks posed by the actions of AI-driven or autonomous systems.

“ As business processes become increasingly digitized, new risks such as human error, misuse of internal data, and AI systems are making internal risk management more challenging. ”


— CISO, Public Sector, Japan
(Survey Response)

Strategy is struggling to keep pace

As the boundaries of insider risk expand, the need to keep pace is becoming harder to ignore. Almost nine in ten (89%) CISOs say their organization has increased its focus on insider risk over the past 12 months, with a similar proportion (91%) expecting the focus to increase further over the coming year. Yet a heightened focus does not necessarily mean greater readiness.



89% of organizations have increased focus on insider risk over the past 12 months



91% of organizations will increase focus on insider risk over the next 12 months

Over half (52%) still do not have a fully implemented insider risk strategy. For some, formal strategies remain only partially implemented, while others continue to manage insider risk through broader security programs. This leaves a clear gap between the attention insider risk is receiving and organizations' readiness to manage it.

AI raises the bar further, and understanding why matters. An AI agent's value is that it can process large amounts of data and make good, informed decisions far faster than a human. Leveraging this efficiency adds risk. Policies to protect sensitive data were designed with humans in mind, and can't simply be extended to create proper AI guardrails. Organizations need to create AI agent-specific policies.

“ AI agents are not just tools; they are identities. They authenticate, access sensitive systems, and act on behalf of trusted users. Until organizations govern them with the same rigor they apply to human employees, they are operating with an unmonitored insider at scale. ”

— David Soto, CISO, Trellix

The visibility gap

Visibility is struggling to keep pace

Knowing where insider risk comes from is one thing; recognizing it as it unfolds is another. Almost every organization (99%) struggles to detect and respond to insider risk, with risks related to AI tools and non-human identities among the most common. As AI becomes part of everyday operations, the detection gap leaves organizations less certain about when its activity starts to present a risk.

Organizations need to understand activity across increasingly complex environments, while often operating with visibility limitations. Meanwhile, AI and previously unknown vulnerabilities are introducing risks that are harder to anticipate, placing further pressure on detection.

Taken together, these challenges point to a need for greater visibility and context, not simply more data. Insider activity often begins with legitimate access, so an individual action may reveal little on its own. Security teams need to connect activity across users, identities, devices, and environments to build a clearer picture of what is happening and identify when legitimate access begins to present a risk.

And organizations know there is room to improve. The vast majority (99%) say they need to improve visibility across users, devices, and environments, with 63% reporting their organization lacks sufficient visibility into how employees use AI tools.. Closing this gap becomes especially important as AI introduces more activity security teams need to understand and, increasingly, new ways sensitive data can be exposed.




Figure 2. What challenges, if any, does your organization face when detecting and responding to insider risks? Base: 600

Case Study: When AI agents operate beyond their boundaries

In July 2026, AI agents broke out of their intended environment and compromised Hugging Face’s production infrastructure. Although this happened during controlled cybersecurity testing with reduced safeguards, the outcome was significant: AI discovered previously unknown weaknesses and used them to compromise real-world infrastructure.

The incident illustrates why governance and guardrails are essential, and why visibility into AI agent behavior is not optional. When AI agents operate beyond their sanctioned boundaries, traditional monitoring, detection, and response cannot keep pace. It also underscores the case for deploying agentic AI within a controlled security boundary.

> Case Study Reference

The visibility gap

AI is creating new points of exposure

The way people work with sensitive data is changing as AI becomes part of everyday business. Employees are using AI to work faster, while AI-driven and autonomous systems increasingly play a role in how information is handled. This brings clear benefits, but it also creates more opportunities for sensitive data to move beyond the controls organizations have traditionally relied on.

Nearly all organizations (99.5%) see some risk of data leakage from using AI or autonomous tools. And that risk can emerge at different points—through the tools employees choose to use, the data AI-driven systems access and process, or the information AI produces. As AI reaches further into everyday workflows, organizations need much clearer visibility into where sensitive data goes and how it is used.

Greatest risk of data leakage when using AI or autonomous tools

Employees will naturally use tools to help them get things done faster. Some tool use will happen through approved channels, while some may not. The challenge is to support productivity without losing sight of how sensitive data is being used. The vast majority (84%) agree they need to somewhat or significantly improve their ability to detect risks related to AI tools and non-human identities. Organizations therefore need to make sure the safe route is also the easy route, rather than forcing employees to choose between security and getting the job done. But directing employees to approved tools is only part of the answer. Highly sensitive data can still be exposed through a sanctioned AI tool if it lacks the right protections. Approval and protection are not the same thing, and organizations need both.

However, AI-related data exposure is not just about what employees put into public tools. Automated systems can also access and process sensitive information, pass it through third-party AI services, or include it in AI-generated outputs. For organizations in regulated industries or public sector environments, the inability to control where that data travels, including whether it reaches a public large language model, is itself a compliance risk. This leaves organizations with more points of exposure to protect and more activity to track.

For many organizations, AI safeguards still haven't caught up. Around two-thirds (68%) say their security and governance policies do not address the risks associated with AI and autonomous tools "very well." At the same time, 74% agree AI has introduced risks traditional security controls were not designed to address.

The issue isn't simply having policies in place. It is also about knowing how AI is actually being used across the business. Without this insight, organizations risk losing control of sensitive information even when the intention behind its use is entirely legitimate.




Figure 3. Where do you see the greatest risk of data leakage within your organization as a result of using AI or autonomous tools? Base: 600

“ Our greatest challenge is not to prevent employees from using AI tools, but rather how to prevent data leakage without stifling productivity. ”

— Chief Security Officer, IT/Computer Services, US
(Survey Response)

The visibility gap

Monitoring trusted activity comes with limits

Insider risk sits in a sensitive space: Organizations need enough oversight to identify potentially harmful activity while respecting employee privacy and meeting legal and regulatory requirements. Finding this balance is already proving difficult, with privacy, legal, and compliance constraints on monitoring (46%) emerging as the most commonly reported barrier to detecting and responding to insider risk.

For CISOs, this is becoming increasingly difficult to navigate. Around eight in ten (79%) agree balancing insider-risk monitoring with privacy and regulatory requirements is becoming increasingly challenging. Better visibility matters, but monitoring must remain proportionate to the risk.

There is also trust to consider. Monitoring should focus on activity that genuinely presents risk, without placing unnecessary scrutiny on employees going about their day-to-day work. Get that balance wrong, and security efforts can create friction elsewhere in the organization.


Better visibility, then, is not about watching everything. It is about having enough context to know what deserves attention.

“ Ensuring compliance with increasingly strict global privacy regulations while implementing necessary monitoring technologies to detect insider threats forces us to balance security needs with employee trust constantly. ”

— CTO, Energy, Oil/Gas and Utilities, Singapore

“ I know some of what is happening, but a lot of risky behavior is getting through. We are chasing the wrong people and chasing false positives. Too many low-level alerts, so we need to prioritize. ”

— Chief Security Officer, Public Sector, Colombia
(Survey Response)



The cost of trusted access

AI is already contributing to insider incidents

Insider risk is no longer a problem organizations can afford to view in the abstract. It is already translating into security incidents with consequences reaching across the business. In fact, nearly all organizations (98%) experienced an insider-related security incident over the past 12 months, with insider risk accounting for 39% of all security incidents during that period, on average.

Human behavior remains a major part of the incident picture, whether through mistakes or malicious actions. AI-driven and autonomous systems are also firmly part of it. Almost all (95%) experienced an incident linked to the actions of an AI agent or autonomous system, with these systems accounting for 28% of all insider-related incidents, on average. With AI use continuing to grow, organizations need to implement the appropriate guardrails, policies, and governance now to strengthen their risk posture, and do so now, before AI-related incidents become the dominant share.

Average percentage of insider-related incidents by cause

Accidental or negligent actions by employees or trusted users



Malicious actions by employees or trusted users



Actions taken by AI-driven or autonomous systems



Figure 4. Of the insider-related incidents your organization experienced over the past 12 months, approximately what percentage were accidental or unintentional, malicious, or related to AI-driven or autonomous systems? The organization must have experienced a security incident as a result of an insider-related risk in the last 12 months. Base: 587



The cost of trusted access

Whether an incident starts with human error, malicious intent, or an AI system, the consequences can quickly spread beyond security. Almost every organization (97%) experiencing an insider-related incident reported an impact on the wider business, including data exposure, operational disruption, compliance, financial loss, and reputation.





Figure 5: What impact, if any, have insider-related incidents had on your organization over the past 12 months? The organization must have experienced a security incident as a result of an insider-related risk in the last 12 months. Base: 587

On average, insider-related security incidents cost organizations \$2.4 million over the past 12 months, but the disruption, regulatory consequences, and loss of customer trust can continue long after.

With so much at stake, how quickly an organization can identify and contain an incident matters.

Slow detection raises the stakes

The longer harmful activity goes unnoticed, the more time it has to affect data, systems, and the wider business. Detecting risk early and acting quickly is therefore critical. Yet, organizations take an average of seven hours to detect an insider-related security incident, followed by a further eight hours to respond.



Respondents take 7 hours to detect an insider-related incident, on average



Respondents take 8 hours to respond to an insider-related incident, on average

“ A seven-hour detection window and eight-hour response window is not a gap; it’s an invitation. When AI agents operate at machine speed, and attackers leverage AI to accelerate every stage of an attack, the organizations closing this gap will be the ones avoiding the \$2.4 million average cost. This requires AI-driven defense and, more specifically, an agentic SOC. ”

— David Soto, CISO, Trellix

The cost of trusted access

Acting faster isn't just about moving faster. Security teams first need to know where to focus. A third (33%) report too many alerts or false positives as a challenge when detecting and responding to insider risk. When genuine threats compete with false alarms for attention, teams can lose valuable time before an investigation even begins.

False positives are not the only problem slowing teams down. More than three-quarters (79%) agree AI and autonomous tools are increasing both the complexity and time required to investigate insider-related incidents. With limited visibility into how these tools are used day-to-day, security teams may have more to unravel when something goes wrong.

“ I know some of what is happening, but a lot of risky behavior is getting through. We are chasing the wrong people and chasing false positives. Too many low-level alerts, so we need to prioritize. ”

— Chief Security Officer,
Public Sector, Colombia

“ Incident investigation gets exponentially harder when complex autonomous workflows obscure lines between routine behavior and malicious network anomalies, straining our response teams. ”

— Chief Security Officer,
Financial Services, Singapore

Speed, then, is not just about generating alerts faster. It is about giving security teams enough context to identify what matters, investigate it quickly, and contain the threat before impact spreads



AI raises the stakes

AI is accelerating the external threat

The same technology reshaping insider risk is also strengthening external threats. AI enables attackers to move faster, find weaknesses, and exploit existing routes into the organization. For CISOs, that prospect feels increasingly real, with 79% saying they expect their organization to be targeted by an AI-assisted attack exploiting trusted users, identities, or systems over the next 12 months.

The concern is not necessarily that AI creates a new way into the organization, but that it gives attackers an advantage across existing routes. More than eight in ten (83%) CISOs agree AI is accelerating the discovery and exploitation of previously unknown vulnerabilities. This leaves security teams with less time to understand what they are seeing and respond, particularly when visibility is already incomplete and closing the detection-to-response gap becomes a matter of organizational survival.

The pressure becomes even greater when the vulnerability itself is unknown. With no known fix available when a zero-day is first exploited, organizations are already operating with limited warning.

AI becomes part of the defense

As AI becomes more deeply embedded in the attacker's toolkit, it's vital organizations use AI for defense. Almost all CISOs (99%) say AI and automation will be critically or very important to managing insider risk over the next 12–24 months, while 41% plan to prioritize investment in detection and response for AI-related or non-human identity risks.

This shift reflects the scale and complexity of activity security teams are now being asked to manage. AI and automation can connect signals, identify patterns, and focus attention on the highest-priority risks.

Adoption, however, comes with its own hurdles. Almost every organization (98%) faces, or expects to face, barriers when using AI or automation to manage insider risk



78%

Over three-quarters (78%) believe traditional security controls are becoming less effective against AI-assisted attacks exploiting zero-day vulnerabilities, making the ability to recognize and respond to unfamiliar activity increasingly important.

AI raises the stakes

Barriers faced when adopting AI or automation to manage insider risk




Figure 6: What barriers, if any, does / would your organization face when adopting AI or automation to manage insider risk? Base: 600

For all its potential, AI cannot just be added to the security stack and be expected to deliver. Organizations still need the right skills, confidence in AI decisions, and a way to address concerns around privacy, accuracy, and integration. In other words, this is far from plug and play. AI needs to work as part of the wider security operation, drawing on the systems, data, and context teams already rely on.

This makes a fragmented approach difficult. Adding standalone AI tools risks giving security teams another layer to manage, when what they need is a clearer, more connected view of risk across their entire environment. Increasing visibility across an organization’s security platform is imperative, where threat intelligence and AI play a pivotal role in enabling faster decision-making and risk mitigation.

For CISOs, these choices are becoming another part of an already expanding remit. As insider risk grows more complex, and AI changes both sides of the threat landscape, so too does the responsibility for managing it, and explaining its importance to the wider business.

“The organizations I talk to know they need AI in their defense. The ones that are stalled aren’t resistant to the technology, but to the trade-offs. For the agentic SOC, sending sensitive security telemetry to a public cloud, absorbing unpredictable token costs, operating without visibility into how AI decisions are made—those are real barriers. An agentic SOC has to work on the organization’s terms, not the vendor’s.”

— David Soto, CISO, Trellix

Insider risk reshapes the CISO role

Boards expect greater clarity

Managing a more complex insider-risk landscape is one challenge; explaining it to the board is another. As insider risk increasingly spans human activity, AI and autonomous systems, CISOs are being asked to turn a fast-moving security issue into something senior leaders can understand and act on.

More than eight in ten (82%) say senior leadership wants clearer reporting and metrics on insider risk than they can currently provide, while 74% say their board still lacks a clear understanding of the risk and its potential impact. That leaves CISOs in a difficult position: Expectations for clarity are rising as insider risk becomes harder to define, measure, and explain.

Much of that challenge comes from the fact that insider risk no longer fits neatly into familiar categories. CISOs need to distinguish between malicious and non-malicious activity, explain emerging risks such as AI and non-human identities, and translate the consequences into terms the wider business can understand. The challenge is not just reporting more, but turning a complex security problem into a clear business story, one that shows where the risk sits, what it could mean for the organization, and where attention is needed.




Figure 7: What are the biggest challenges you face when communicating insider risk to senior leadership or the board? Combination of responses ranked first, second, and third. Base: 600

Accountability continues to rise

The growing focus on insider risk is putting more responsibility directly on the CISO, with 92% reporting their accountability for insider-related incidents has increased over the past 12 months. Yet AI agents complicate who is responsible for what. Traditional security assumes a human actor: someone with an identity, permissions, and clear ownership of their actions. Agentic AI breaks that model. An agent may autonomously execute tasks, use delegated credentials, and make decisions no human individually approved, leaving accountability without a clear owner. And when ownership is unclear, that ambiguity falls somewhere.

That additional responsibility is landing on a role already under considerable pressure. In the previous edition of The Mind of the CISO, seven in ten (70%) CISOs said they were concerned about their future in the role due to expanding responsibilities. Now, as insider risk stretches across people and AI, CISOs are being asked to account for activity they may not directly control and risks that can have consequences well beyond security.

For many, that accountability is becoming personal. More than two-thirds (69%) are concerned about their own liability in the event of an insider-related security incident. The expanding remit of the CISO is therefore not only bringing more responsibility, but raising the stakes attached to the decisions they make.

Insider risk reshapes the CISO role

CISOs look for greater guidance

With more responsibility landing on CISOs, they need firmer ground to stand on. AI is changing how insider risk emerges and how organizations respond, but the guidance around managing that change has not always kept pace. It is little surprise that almost all (98%) CISOs see a role for governments or regulators in helping organizations manage insider risk.




Figure 8: What role, if any, should governments or regulators play in helping organizations manage insider risk? Base: 600

CISOs are not simply asking for more regulation; they want greater clarity on how to put it into practice. The focus on guidance, clearer requirements, and AI-specific frameworks reflects the decisions security leaders are already making as AI agents gain greater access to data and systems. They need a clearer sense of what good governance looks like, where the boundaries should sit, and how responsibility should be shared. Without it, CISOs are being asked to manage, and ultimately answer for, risks that are evolving faster than the rules around them.

Ultimately, the goal is to give organizations the confidence to use AI without losing control of the risks that come with it. As AI moves deeper into everyday operations, CISOs need enough certainty to make decisions around access, monitoring, and accountability without standing in the way of legitimate use. For a role carrying more responsibility when things go wrong, that certainty is becoming increasingly important.

Case Study: When guidance falls short

The frustration with a lack of clarity came through clearly in an interview with a US-based CISO at a financial operations provider. They described fragmented state requirements as “a nightmare.” They argued broad “embrace AI carefully” messaging offers little practical help as AI agents begin to act through employee credentials and access. Instead, they called for “intelligent guidance,” more consistent high-level regulation, and a “unified matrix of regulations” that would help organizations “safely use AI for legitimate business purposes.”

Conclusion

As this edition of The Mind of the CISO reveals, the boundaries of insider risk are shifting. Employees and trusted users remain a major source of risk, but AI agents are now firmly part of the picture. Security strategies built primarily around human behavior are no longer enough for an environment where AI is taking on a growing role across the business.

Maintaining visibility will be critical as AI creates more ways to access, process, and share sensitive information. As legitimate behavior becomes harder to distinguish from potential threats, organizations need a clearer view across users, identities, devices, and data, connecting signals and context to understand where to focus.

At the same time, attackers are using AI to find and exploit weaknesses faster, narrowing the window defenders have in which to respond. AI and automation will therefore have an important role in helping organizations act at the same pace, accelerating investigation and response before the impact spreads.

For CISOs, this comes as expectations and accountability rise. They are being asked to manage a broader risk landscape, explain it more clearly to the board, and make AI decisions that will shape how securely it is adopted across the business.

For organizations, the challenge now is to embrace what AI can offer without losing control of the data, access, and systems it touches. Doing so will require security solutions, like an agentic SOC, that can see more, understand more, and respond faster as the boundaries of insider risk continue to move.

Recommendations

Based on the findings of this Mind of the CISO report, several priorities stand out for CISOs and their organizations. Together, they reflect the need to rethink insider risk for an AI-driven world, strengthening visibility and control while enabling organizations to embrace AI securely.

Redefine insider risk for an AI-driven organization

Insider risk strategies must now extend beyond employees and trusted third parties to include AI agents, autonomous systems, and machine identities. Organizations should establish clear ownership, access controls, and oversight as they would for human users, ensuring AI operating through trusted access is governed with the same rigor, reviewed with the same frequency, and revoked with the same authority.

Build visibility around context, not more alerts

Insider risk rarely reveals itself through a single action. CISOs need connected visibility across users, identities, devices, data, and AI activity to distinguish genuine risk from legitimate behavior, reduce false positives, and respond faster. A more integrated security approach, like an agentic SOC, can provide the context teams need without simply adding more alerts or tools to manage.

Make secure AI adoption the path of least resistance

Employees will use AI where it helps them work faster, so security needs to enable safe use rather than stand in its way. Organizations should pair clear AI policies with appropriate data controls and monitoring, and consider whether sensitive workloads require a deployment model like on-premises, air-gapped, or private cloud that keeps data within the organization's security boundary and avoids reliance on public large language models.

Make AI part of the security operation

As AI accelerates the threat landscape, CISOs should use AI and automation to strengthen detection, investigation, and response. The greatest value will come from applying AI across existing security data, tools, and workflows, helping teams prioritize genuine risk and act faster rather than introducing another standalone layer into the security stack.

Share accountability and strengthen governance

CISOs cannot manage insider risk in isolation. Organizations should establish shared accountability across security, IT, HR, legal, business leadership, and the board, supported by clearer governance and reporting. Policymakers can reinforce this by providing practical AI governance frameworks that keep secure adoption front and center while enabling speed rather than hindering it, to help organizations protect sensitive environments while continuing to advance and innovate.

Additional resources

Mind of the CISO research series:

- **The Future of Cyber Resilience:** Over 500 CISOs worldwide share their views on building resilience through hybrid infrastructure, navigating OT and IT convergence, and responding to increasingly sophisticated AI-driven threats, while managing the growing demands placed on the CISO role.
- **Closing the Gap between Reaction and Readiness:** Over 500 CISOs worldwide share their views on the evolving threat landscape, the role of AI and automation in combating threats, and the value of peer communities in navigating complexities and driving clarity.
- **CISO Crossroads:** Over 500 security leaders worldwide share their views on cybersecurity regulation, the CISO role, and their interactions and challenges when reporting to their organization's board.
- **Decoding the GenAI Impact:** Trellix engaged with 500 security leaders across North America to understand how GenAI and AI are evolving the threat landscape and the CISO role to reshape the future of cybersecurity in the workplace.
- **Behind the Breach:** To shed light on the challenges CISOs face after a breach, Trellix surveyed over 500 security leaders worldwide who have managed a major cyber incident, revealing strategic insights, key stats, and lessons for the best path forward.
- **Understanding the CISO's Struggle:** Trellix engaged with over 500 security leaders to understand what's holding SOC teams back, how they work amid a tumultuous threat landscape, which business functions hold them back, what tools and support they need to succeed, and how best to move forward.

Additional resources

Trellix Advanced Research Center Digest

Subscribe to get the latest cybersecurity trends, best practices, security vulnerabilities, and more.

Trellix Advanced Research Center Threat Reports

Read the latest Trellix threat reports for insights into the evolving threat landscape and the tactics, techniques, and methodologies actors are using.

Reporting to the Board: CISO Best Practices

As cybercrime and regulatory pressure on cybersecurity grow, the CISO role is evolving from a technical expert to a business-focused leader. Many CISOs today are asked to report to their organization's board regularly. How can you make the most of your time in front of the board? Check out this guide from the Trellix CISO Council to learn best practices for presentations.



Trellix is a global cybersecurity leader helping organizations build cyber resilience for the AI era. Combining AI-native security, operational threat intelligence, and automated detection and response, Trellix provides defense-grade protection at machine speed across on-premises, air-gapped, operational technology, hybrid, and cloud environments. Trellix helps customers anticipate threats, scale effective security operations, maintain business continuity, and adapt their security posture for the modern threat landscape.

More at www.trellix.com
Follow Trellix on [LinkedIn](#) and [X](#)



Vanson Bourne is an independent specialist in market research for the technology sector. Their reputation for robust and credible research-based analysis is founded upon rigorous research principles and their ability to seek the opinions of senior decision makers across technical and business functions, in all business sectors and all major markets.

For more information, visit www.vansonbourne.com