

Elusive threats, elastic defense

Securing AI at scale

Contents

Foreword	3
Introduction	5
Part one	
The AI battlefield	8
Part two	
Dynamic defense through AI exposure management	13
Part three	
Making AI a platform for growth.....	18
Part four	
Aligning the ecosystem to secure AI at scale	22
Action guide	28

How IBM and Palo Alto can help

IBM Consulting and Palo Alto Networks have joined forces to deliver AI-powered, fully integrated, open, end-to-end security solutions to enterprises. From consultation through execution, we can help you modernize your cybersecurity program, saving time, money, and resources as well as enhancing your organization's resilience against today's complex threats. For more information, visit ibm.com/consulting/palo-alto.

Established security models can't keep pace with AI threats

Enterprise security is at a new breaking point as AI accelerates attack innovation and reshapes the threat landscape. Two-thirds of executives say their organization has been targeted by an AI-enabled attack in the past year. Welcome to the era of agentic AI threats, where attacks operate with machine speed and persistence, exploiting vulnerabilities and adapting in real time to outmaneuver traditional security operations.

Fighting the new breed of attackers requires a new enterprise security approach. AI is a force multiplier for threat actors. This means even highly automated defenses can no longer meet some mission-critical security requirements. The only way to best AI-powered attacks is with AI-powered defenses.

In short, your AI must be better than their AI. That demands more than incremental improvements to existing tools. It calls for rethinking security operations as an AI-powered, agent-automated, platform-based architecture that integrates data, intelligence automation, and governance across the enterprise. Securing increasingly AI-integrated operations involves protecting decisions and outcomes, not just assets. The attack surface is no longer defined by technology alone: It is increasingly contextual and semantic in nature. And in this new AI-powered environment, even best-of-breed point solutions will struggle to keep pace as adversaries learn and adapt, creating dangerous vulnerabilities.

When built intentionally, a unified AI-first platform approach can integrate everything from IT to security to business operations. This allows security teams to move from reactive defense to continuous, adaptive prevention, remediation,

and resilience. For example, agentic AI defenses can scan for and manage exposures, prioritize threats, recommend actions and—for mature use cases approved by security domain experts—even respond autonomously. Just as importantly, a platform approach provides a foundation for responsible AI governance, ensuring that innovation does not come at the expense of trust, transparency, or control.

Of course, technology alone is not enough. Real progress requires evolving operating models and partnerships. Security leaders must work closely with IT and business stakeholders as they embed AI into the fabric of the enterprise. Runtime security at this scale has no precedent. What is required is nothing short of a new shared responsibility model for securing AI that can handle the speed and uncertainties involved with AI-first ways of working. This is how AI can create a secure and trustworthy operations environment for driving growth opportunities across the enterprise.


Most organizations are not there yet, however. This report analyzes the fast-moving landscape and maps the route forward. We invite you to dive in and explore how enterprise security must evolve—now—to stay ahead in an AI-powered world.

Mohamad Ali

Senior Vice President
IBM Consulting

BJ Jenkins

President
Palo Alto Networks



Key takeaways

The new reality for security and operations leaders: conventional models are failing as AI reshapes the enterprise.

- Organizations are rapidly integrating AI across business operations, creating new attack surfaces and security vulnerabilities.

61% of surveyed leaders say their organization's AI models, assets, or data have been compromised.

- AI-powered threats are multiplying and evolving faster than defenses. Legacy security models cannot be retrofitted to address new AI-specific risks.

67% of surveyed executives say their organizations have been targeted by an AI-enabled attack in the past year. A patchwork of point security solutions poses a barrier to securing an increasingly diffuse enterprise boundary and a proliferation of non-human identities, among other new-breed risks.

- An integrated AI operations platform supports security at scale.

To secure the AI estate while scaling it, organizations are integrating AI operations and embracing AI exposure management. Four out of five leaders say their organization needs better integration across hybrid cloud, AI, and security platforms.

- Scaled and secure AI leaders are emerging.

About one in four (24%) organizations have begun to break out from the pack in terms of integrating AI into operations and security infrastructure, and adopting mature AI exposure management and shared AI governance practices, our analysis finds. Associated benefits include greater security resilience, ROI, and operational performance.

A look at our study

20
industriesAverage annual revenue
of surveyed organizations

~\$17B

17 countries across Asia Pacific,
Europe, Middle East, Latin
America, and North America600 CISOs, CTOs, CIOs,
CDOs, CAIOs400 CEOs, COOs, CFOs,
Chief Risk Officers

Security for AI, and AI for security

AI is changing the nature of both business operations and enterprise cybersecurity. As organizations integrate AI agents into day-to-day workflows, supporting human experts with new forms of digital labor, adversaries are creating new agentic threats.¹ They are fast, relentless, and precise. Anthropic revealed in November 2025 that hackers exploited vulnerabilities in its Claude AI model to attack 30 global organizations.² These and other threat actors are now using widely available AI models and tools, including agentic capabilities, to automatically execute novel attack methods.


To understand how AI is reshaping business operations, technology, and security challenges in the emerging agentic era, the IBM Institute for Business Value (IBM IBV) in partnership with Palo Alto Networks conducted a global survey of 1,000 C-level executives. Its dual focus: how large enterprises are securing AI-powered operations infrastructure today, and how they are using AI to secure the enterprise. This report's insights and recommendations are grounded in survey results as well as insights gleaned from interviews with security and technology executives.

One conclusion is clear: Conventional security models are failing as AI introduces new dynamic attack surfaces.³ Threat intelligence teams see clear evidence of cybercriminal tactics targeting AI infrastructure and operations.⁴ 61% of leaders say their organization's AI models, assets, or data have been compromised, our survey found.

The potential for disruption is immense. An AI agent might infiltrate a procurement system and auto generate fraudulent vendor invoices at scale. A rogue model might spread inaccurate inventory data across supply chain systems or hide nefarious activity in common workloads.⁵

As enterprises rush to scale AI across business units, however, they often fail to see where they are exposed and underestimate the AI-specific risks confounding current governance models.⁶ Many AI agents are enabled using one-time authentication and inherited permissions that cross organizational boundaries. Based on implicit trust, this kind of deployment carries significant risk.

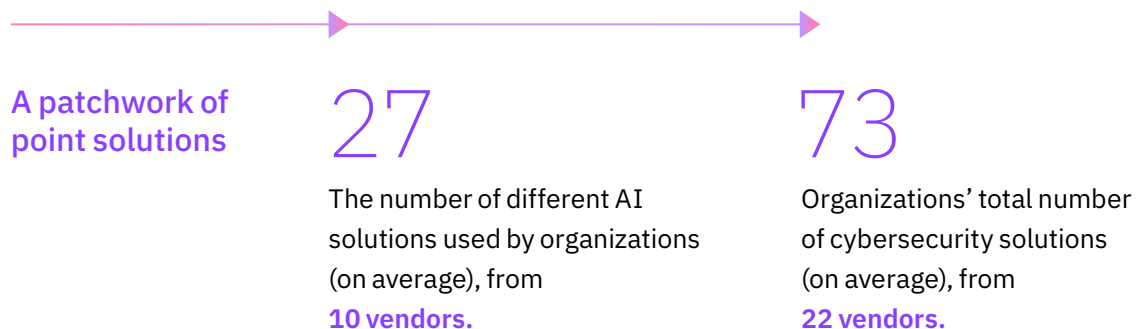
In this environment, enterprises are scrambling to figure out how best to update their security strategy. But many are unsure how they should prioritize AI investments to harden defenses while fueling growth. What is clearer: Efforts to secure AI will likely grow costlier in the coming years. Our research finds that annual cybersecurity spend resulting from efforts to secure AI will rise 55% on average by 2028.



This report aims to help leaders navigate the amorphous AI landscape. In Part One, we examine how AI is reshaping the battlefield for attackers and defenders—and why existing security models can’t simply be retrofitted for this new reality. Part Two outlines how organizations can mitigate the most consequential business and operational risks by focusing on AI exposure management. In Part Three we explore how an integrated approach to AI operations provides a foundation for next-generation security platforms. Part Four looks at how shared responsibility, interoperability, and a resilient ecosystem enable AI to be safely scaled. Each section ends with a “readiness assessment” offering leaders questions to ask. Finally, Part Five presents a step-by-step action guide.

“Security professionals are paid not to trust—so zero standing privilege and strong controls are essential to scaling autonomy responsibly. Success will be incremental: start with quick wins at Tier 1, prove safety, earn trust, then ratchet up autonomy.”

Matthew Bissell, Senior Director, Cloud Security & Response, Sabre



Part one

The AI battlefield

Today's cyberthreats are evolving faster than the defenses built to keep them at bay, according to 70% of surveyed executives. Bots that once probed host domains for vulnerabilities are transforming into a fleet of AI malware agents.⁷ These automated adversaries are weaponizing zero day vulnerabilities, moving faster than even the most highly skilled human defenders.⁸

Broadening AI usage across the enterprise—everything from employee “shadow AI” solutions to customer-facing AI applications to back-office integrations—is multiplying security risks. Half of executives (48%) report their organizations lack adequate protection for model inference and live AI usage, and 56% say the same about AI model development and training. When it comes to AI data management, 38% concede they have insufficient safeguards. There are similar shortcomings in access controls and identity protections (45%), as well as telemetry and observability (42%). Meanwhile, 76% of executives say they are seeing a rise in unsanctioned AI use by employees.

Attackers are now targeting these gaps, from overly broad AI permissions and data access to machine identities lacking granular controls and context-specific least privilege.⁹ The result is an asymmetrical battlefield. If organizations stick with the status quo, they'll spend resources and budget fire-fighting a barrage of increasingly capable autonomous threats. And they will find it difficult to keep up.

One critical limitation: cybersecurity operations beholden to old playbooks. Today's attackers evolve so rapidly that human-in-the-loop approaches are less viable. The alternative? A “human-on-the-loop” model where human experts supervise cyber defense agents as an integral part of autonomous threat operations.

“Agentic AI changes the threat model because it can operate continuously, quietly, and without triggering traditional controls. Threat actors don't need to be loud anymore. They just need to be patient.”

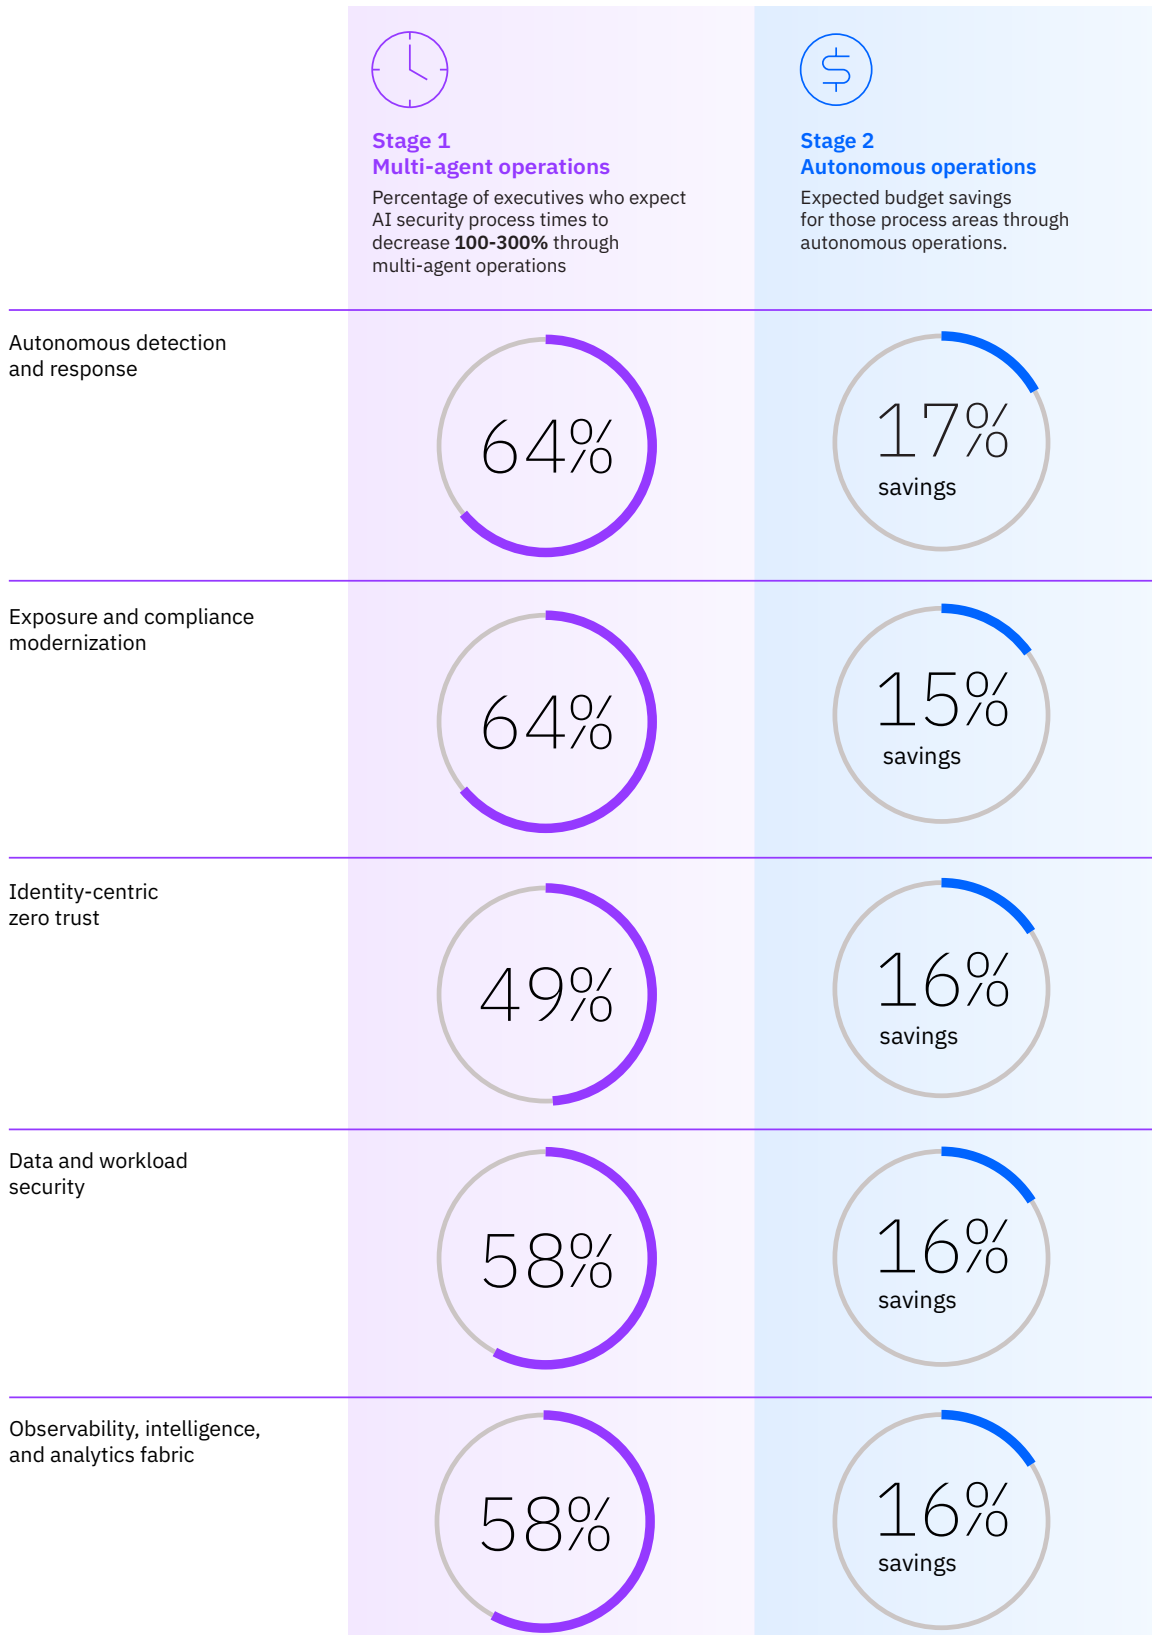
Steve Jablonski, CISO, TELUS Digital

FIGURE 1

Advanced AI at scale delivers significant time and cost savings

From 4 Hours

To 1 Hour

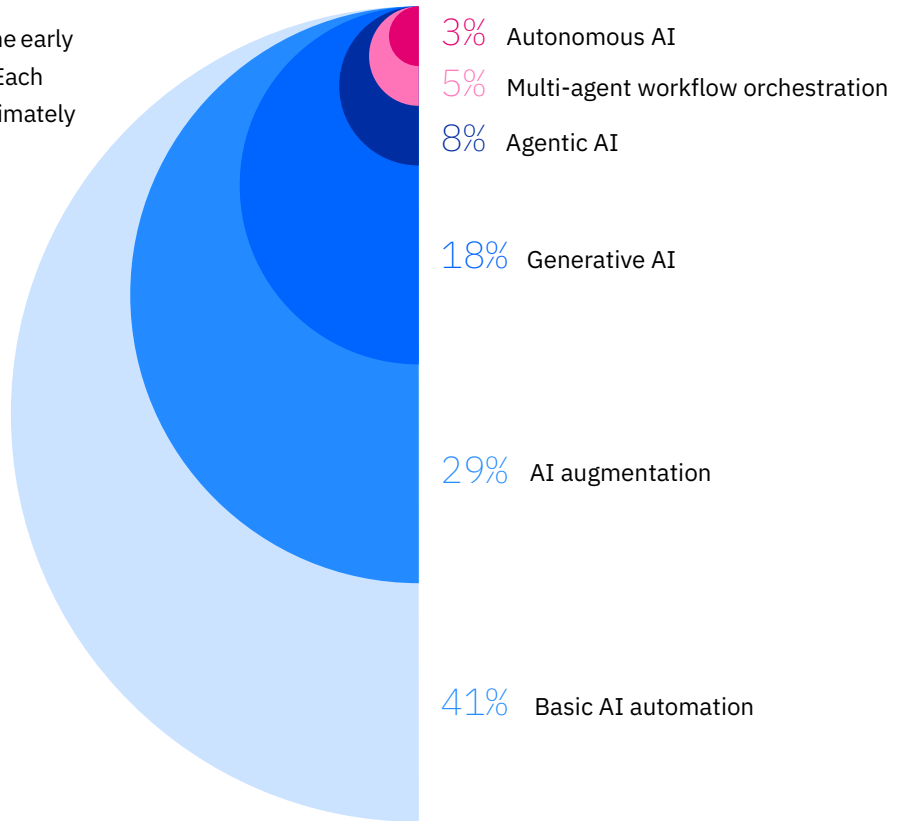


Qs: Which 6 functional areas would deliver the greatest value in the shortest time to your organization? What would be the estimated monetary benefit if these capabilities could be made fully autonomous (i.e. requiring minimal human supervision) at your organization? (n= 41-463)


FIGURE 2

The state of enterprise AI

Organizations are currently in the early stages of adopting AI at-scale. Each step up in maturity sees approximately one-third to one-half the rate of adoption (top). AI integration is still a work in progress, limiting the potential gains from AI investments (bottom).



AI integration



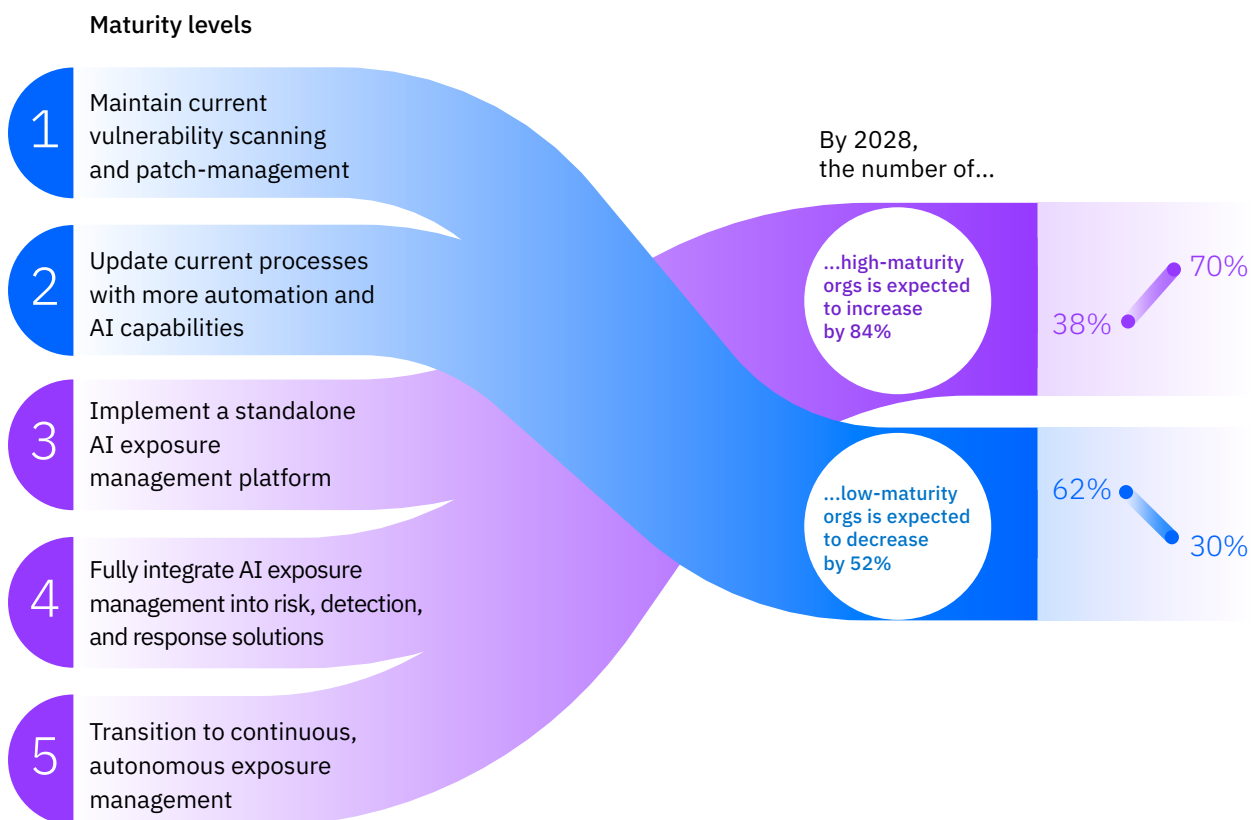
Q (upper graphic): Considering current operations across your organization, what percentage of your workloads are enabled with the following? Q (lower graphic): To what extent is AI integrated into the following areas of your organization? (n=1,000)

In this approach, humans proactively fortify defenses using the same highly automated, adaptive technologies wielded by attackers. When used strategically, agentic AI can continuously inventory AI attack surfaces across the enterprise and then generate policies that enforce granular, identity-based controls.¹⁰ This can block unwanted agents from inheriting permissions across sessions and enforce constraints on AI inputs, outputs, and access to sensitive data.¹¹ It can bind every action to a specific user, a specific intent, and a specific authorization path.

Yet few organizations have achieved such AI security maturity. Only one in 10 executives in our research says AI is fully integrated into the IT operating model and DevSecOps. Fewer than one-third describe their organization as “more mature than peers” or “best in class” with respect to AI security. In fact, 51% say their current infrastructure cannot securely support multi-agent or autonomous operations.

FIGURE 3

Pendulum shift? Leaders expect rapid AI exposure management evolution in the coming years



Q: How do you approach your exposure management capabilities in 2025, and how do you expect that to evolve over the next 3 years (by 2028)? n=1,000

“Risk isn’t merely rising, it’s compounding. Every new AI-driven integration multiplies the attack surface and increases complexity. At scale, this creates an exponential risk curve that escalates faster than leaders expect.”

Koos Lodewijkx, CISO, IBM



Readiness assessment

Building your battle plan

Key questions about your enterprise’s AI maturity and vulnerability.

- What percentage of current threat operation workflows are automated? Where is the organization reliant on human inputs, outputs, judgement and context? Are these areas subject to higher risks due to potential delays, resource or capacity constraints?
- Is your organization investing in AI for security as incremental tooling or as a foundational redesign of trust, identity and risk management across the enterprise?
- Do all AI agents have an identity with specified access rights and restrictions? Do those rights grant the appropriate levels of authority and autonomy to execute only those actions required for workflow objectives? Are AI agents governed with the same rigor as humans? Is each agent treated as distinct, with controls that span inputs, operations, and outputs?
- If mission-critical business operations depend on autonomous agents, how confident are you that suppliers and partners can meet the same AI security and governance standards? Do you have responsibility mechanisms in place to hold providers and suppliers accountable for safe agent behaviors?

Part two

Dynamic defense through AI exposure management

If widespread AI adoption and security governance are in tension, then addressing the friction becomes an urgent business imperative. So how can organizations best increase the tempo of their security operations? Most executives (61%) agree that the only way to secure AI is with AI.

When it comes to AI security, leaders need an approach that meets three criteria: discover, assessment, and protection. Organizations should have complete visibility of all AI deployed. They should be able to assess the risk of all deployments. And they should have the ability to protect AI at the point of production (aka, runtime), acting on threats in minutes given finite knowledge and limited context.

That puts AI exposure management on the frontline of cybersecurity defense. In our research, 76% of executives say AI exposure management will be their organization's primary method for unifying prevention, detection, and response activities by 2028. Put another way: With AI exposure management, autonomy can become a strategic advantage. Without it, autonomy will likely amplify risk.

Executives recognize that AI exposure management is a critical building block to secure the AI driven enterprise. Not just another security function, it offers full visibility of the organization's AI footprint. Continuous AI threat exposure management capabilities allow enterprises to adopt AI at scale without surrendering visibility, accountability, or resilience. Unlike traditional security focused on fixed infrastructure, AI exposure management is concerned with dynamic risk surfaces arising from non-human identities and non-deterministic AI behaviors. That requires moving away from reactive security operations and static risk inventories. The new goals, given how opaque AI operations are: continuously identifying and mitigating risks through clear AI policies and controls, behavioral testing of AI systems, and consistent AI governance across the ecosystem.

“We want our AI standards to shape expectations around speed and time to value. Moving fast only matters if the outcomes are sustainable. Otherwise, velocity becomes technical debt.”

Matt Lyteson, CIO, IBM

Aspiration versus reality

76%

of executives say that by 2028, exposure management will be their primary method for unifying prevention, detection, and response activities.

But today,

64%

of organizations are still dependent on human-based remediation of their AI security exposures.

Most organizations aren't there yet. A majority are not linking exposure insights to real-time business metrics (52%), not extending zero trust or posture management to non human identities (54%), and not automating AI threat remediation (64%). Half of organizations say their current AI security capabilities are best described as "fragmented" or "initial" policies. Few are developing platforms that automatically inform detection and response priorities (12%) or integrating exposure data into dashboards (21%).

The good news is that executives are increasingly prioritizing AI exposure management. They are looking to move beyond compliance reporting and incident response to proactive, continuous risk-based threat assessment. The goal: the ability to prioritize addressing vulnerabilities that pose the highest risk to the business, rather than trying to fix everything at once.

While only 38% of surveyed organizations say they had integrated AI exposure management into risk, detection and response solutions as of 2025, 70% expect to by 2028. To achieve this, major obstacles will need to be addressed, such as the complexity of integrating with existing tools and dependence on legacy infrastructure.


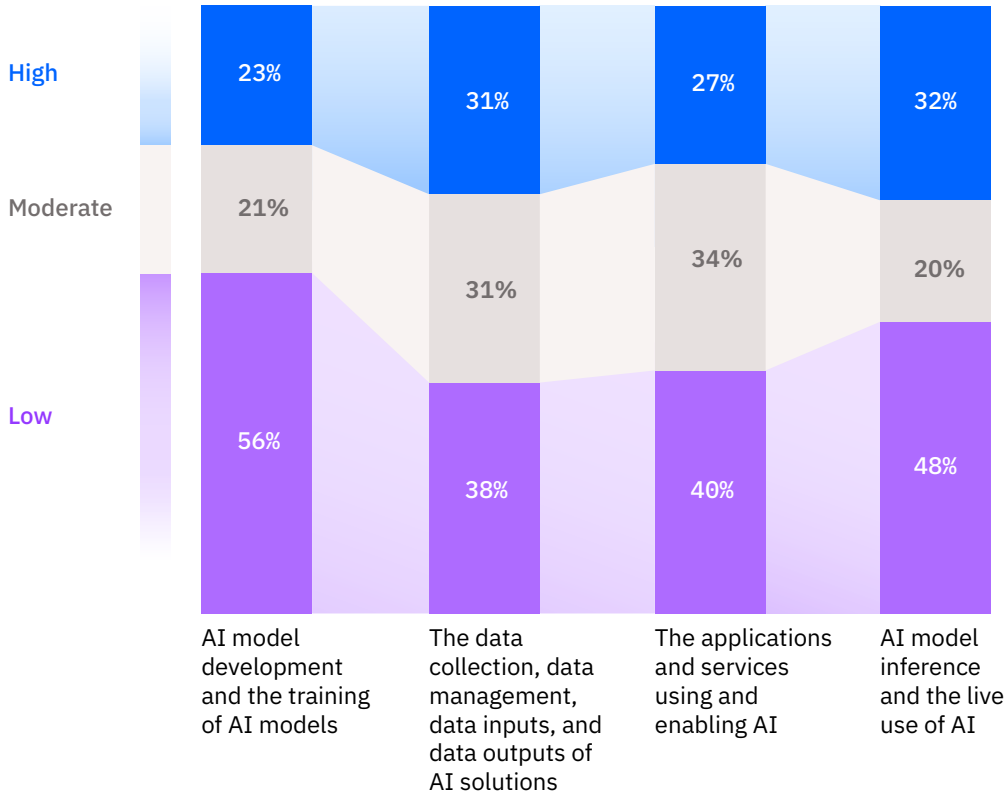
AI exposure management

AI exposure management is the discipline of discovering, contextualizing, and reducing security risk across an organization's AI and agentic applications footprint, including models, agents, data pipelines, integrations, identities/permissions, and runtime behaviors. Risks are reduced by correlating cybersecurity exposures to business impact and then prioritizing (increasingly automated) mitigation actions.

FIGURE 4

Haves and have nots: The state of enterprise AI security

Some organizations (blue) are starting to realize AI security benefits at scale. Others (purple) are at risk of being left behind.



Q: Estimate the current robustness of your security in the following [enterprise security] areas... Q: Estimate your organization's current maturities in the following [enterprise security] areas: n=1,000



Readiness assessment

Bolstering defense

Key questions about your enterprise's approach to AI exposure management.

- As AI adoption accelerates, how are you calculating (and calibrating) the tradeoff between AI speed to market and AI security readiness? What metrics or maturity frameworks help you determine when to move forward with measured risk versus pausing to implement additional safeguards?
- If you had real-time insight into your organization's AI attack surface, what decisions would that enable you to make differently? What tooling or procedural obstacles stand in the way of implementing real-time AI telemetry and real-time AI risk analysis?
- Given that AI systems introduce different security challenges around real-time authorization, auditability, integrity, and data provenance, how do you evaluate whether incremental improvements to your operating model are sufficient? What criteria would signal the need for transformative changes to your security architecture?

Perspective

On the cusp: multi-agent systems

One conclusion from our analysis: most organizations' governance and infrastructure cannot (yet) support secure multi-agent or autonomous operations. Yet the benefits of multi-agent systems (MAS)—which consist of multiple autonomous agents that interact with each other and their environment to achieve complex or collective goals—are undeniable. Their value derives from a combination of decentralization, collaboration and adaptability:

Scalability

Agents can be added to a MAS as needed to address larger, more complex problems

Robustness and fault tolerance

Failure in one component doesn't disrupt the business process.

Parallelism and efficiency

Complex problems can be decomposed into smaller, more manageable subproblems

Flexibility and adaptability

Agents can react to changes in the environment independently.

Modularity and upgradability

It is easy to update, replace or improve individual agents without disrupting the whole system.

Improved decision-making

Diverse agents may have different perspectives, models or heuristics, generating better outcomes via their collective intelligence.

Digital labor integration

Agents can be designed to support high-skill domain experts with specialized knowledge.

Some organizations have already deployed multi-agent systems, including in support of mature business processes such as incident detection and response, and threat remediation. When integrated into a unified AI platform spanning the enterprise, MAS can support end-to-end operations and security. In this context, AI becomes ambient: embedded, adaptive, and self-coordinating. MAS can scan for vulnerabilities, prioritize them, and then patch the most urgent issues—all automatically.

Part three

Making AI a platform for growth

While AI exposure management is the first step toward integrated, end-to-end AI security, next-generation defense relies on a unified platform for managing shared data, identity controls, and governance. A secure AI foundation is what business stakeholders need to power AI-first productivity and revenue growth.

Cybersecurity systems today are often far removed from that vision: a patchwork of limited visibility and operational disconnects that slow threat detection and response. In fact, organizations in our survey report having 27 different AI solutions on average, from 10 different vendors. Their average total number of cybersecurity solutions: 73, from 22 vendors. This level of complexity impedes efficient security operations.

For defenders to keep up with AI-powered attacks, they need to be able to rely on their own cyber-defense agents to take preventative and remedial actions in real-time. This is a big leap from current practice. For agents to become autonomous, they need to first be demonstrably trustworthy and reliable. A step toward this requirement: adopting an integrated AI operations platform that standardizes capabilities and integrations, reducing complexity and strengthening trust.

A unified AI platform is ambient in nature, supporting end-to-end operations by integrating IT, security, applications, cloud services, and other functional areas. This facilitates consistent enterprise-wide AI governance, such as enforcing sovereign data privacy and compliance requirements. Executives recognize secure and trustworthy AI requires collaboration:¹² Our research shows leaders in the C-suite are jointly responsible (to varying degrees) for security, technology, and AI budgets (see Figure 5).

“Uncontrolled AI tool sprawl isn’t just inefficient—it’s ungovernable. You can’t secure what you don’t understand, and you can’t understand fifty different tools.”

Steve Jablonski, CISO, TELUS Digital

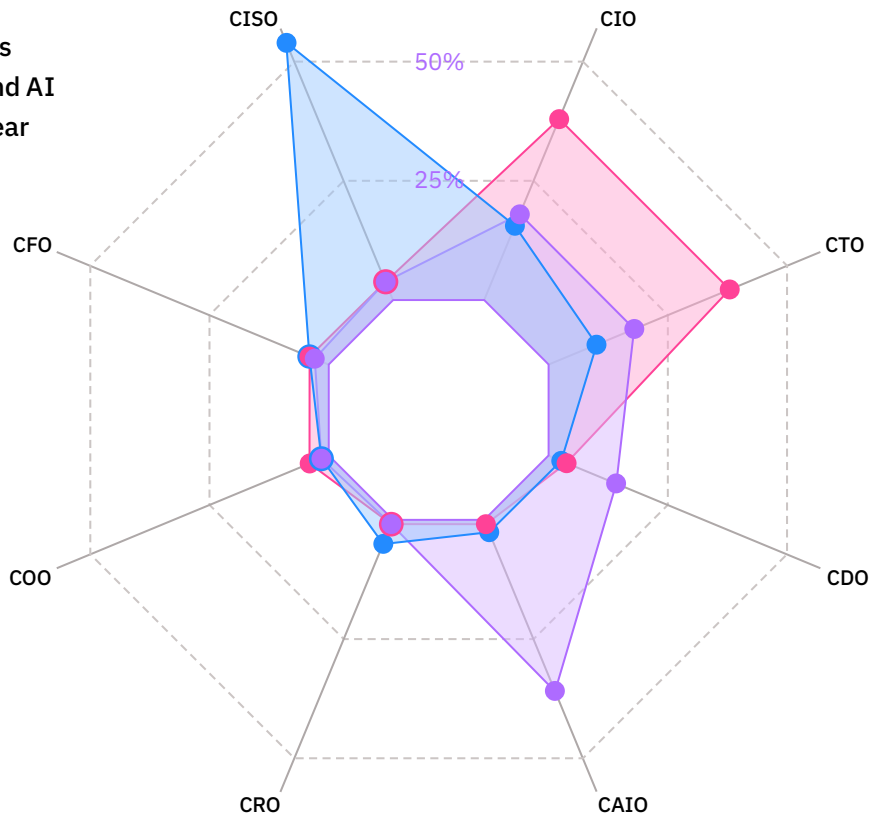
FIGURE 5

Budget responsibility by C-suite role

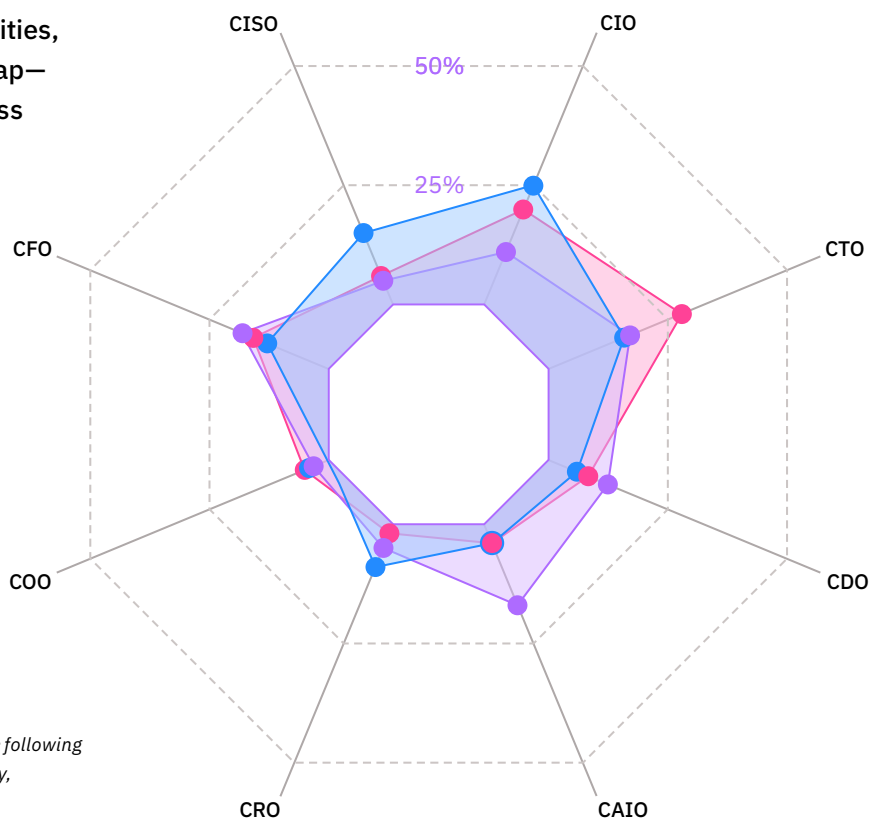
Collaboration enables integrated AI operations

● Cybersecurity ● Technology ● AI

Primary budget responsibilities across security, technology, and AI show some overlap, though clear distinctions by role.



Secondary budget responsibilities, however, show far more overlap—suggesting collaboration across roles is critical to success.



Q: Which executive(s) are responsible for the following budgets at your organization? (Cybersecurity, technology, AI) n=1,000

Leaders we spoke with see the merits in taking a platform approach to integrated AI operations. Most executives acknowledge the need for greater AI interoperability. Three out of four (77%) in our research say their organizations need better integration across hybrid cloud, AI, and security platforms.

Integration helps set the stage for AI agents earning the trust and reliability required to become truly autonomous. But safe autonomy must involve every agent operating as a managed, authenticated identity with specified permissions and associated controls. Supplemental governance practices such as contextual asset awareness and AI runtime security—which prevents AI model prompt manipulation—can further help secure agent identities.¹³

Leaders recognize that AI agents are emerging as a potential attack surface. About two-thirds (68%) of executives say their organization needs to treat AI agents as distinct identities. Failure

to do so increases an organization's exposure to cybersecurity attacks. Our analysis finds that organizations with elevated AI identity-related risks experience a 52% higher cybersecurity incident rate, even after adjusting for size. In other words, poor identity management has real security consequences.

Finally, an integrated AI operations platform facilitates common governance. This streamlines decision-making across functional boundaries. By integrating AI capabilities at an enterprise-level—across technology, security, and AI domains—leaders are connecting technical capabilities with operational outcomes.

Executives see a strong connection between governance and security. Nearly nine in 10 leaders agree that today's decisions about governance, policy, and infrastructure will directly determine how safely they can deploy multi-agent and autonomous AI over the next two years.

Three steps to AI resilience

- 1. Discover:** Establish total visibility and control of the AI system
 - Treat every AI component as an asset and identity
 - Continuously map how AI systems connect and act
 - Operationalize governance across domains and ecosystems
- 2. Assess:** Make AI exposure measurable, contextual, and actionable
 - Unify AI exposure into a single, contextual risk model
 - Continuously test AI systems against real-world adversarial behavior
 - Leverage AI to prioritize and remediate at machine speed
- 3. Protect:** Enforce security in real time as AI systems act autonomously
 - Secure agents as autonomous actors with enforceable guardrails
 - Apply runtime security to stop threats as they emerge
 - Build an integrated architecture for trusted autonomy at scale

At the same time, the pace of change in AI makes long-term planning difficult. New frameworks, tools, and attack techniques are emerging faster than most organizations can track. New agentic capabilities can appear quickly, often outside traditional security visibility. Because the threat landscape is evolving so rapidly, organizations cannot rely on static policies or point solutions.

The most effective approach is to secure AI by design. This means giving organizations the ability to discover AI systems and risks across the enterprise, assess vulnerabilities before they are exploited, and apply continuous, proactive protection that can detect and contain new and unexpected threats as they emerge.



Readiness assessment

Replacing complexity with capability

Key questions to assess if your organization is platform ready.

- Consider your current security and AI operations footprint. What's preventing your organization from moving from dozens of point security and AI solutions to an integrated AI operations platform? What would need to change in terms of technical architecture or operational support processes?
- Do you know how many AI agents are deployed across your organization? As AI agents proliferate, do you have standards for how they are designed, deployed, and managed? Who is responsible for managing agents as distinct identities—assigning permissions, monitoring behavior, and revoking access?
- What are the costs of complexity and fragmentation when it comes to your AI security solutions? Assess the impacts in terms of licensing and support, user experience, operational complexity from multiple data and service integrations, and indirect costs from managing shadow AI deployments across the enterprise.

Part four

Aligning the ecosystem to secure AI at scale

For AI-first operations, an integrated platform sets the stage for security at scale. But it's not enough. With AI and security talent in short supply, organizations rely on partners to fill gaps.

Given the reality that AI value derives from shared models, data, agents, and workflows, leaders need to assess AI operations in terms of shared responsibility. Interconnected AI infrastructure and operations demand that security be managed at an ecosystem level, bringing together partners that share responsibility for managing risk and ensuring accountability.

A large majority of surveyed executives (82%) agree that AI security is a shared responsibility across suppliers, providers, and partners. But there is a lack of consensus on who is primarily responsible for making sure AI agents work together and can scale across functions. 29% of leaders say primary responsibility lies with the organization itself, while 17% say it is shared among all parties. Nearly one in six (16%) aren't sure—a signal of how quickly the landscape is changing.

“Partners should run the tooling. We don't want to spend headcount maintaining platforms. Our people should focus on using insights to secure the enterprise. That shift amplifies our impact instead of pulling us into infrastructure overhead.”

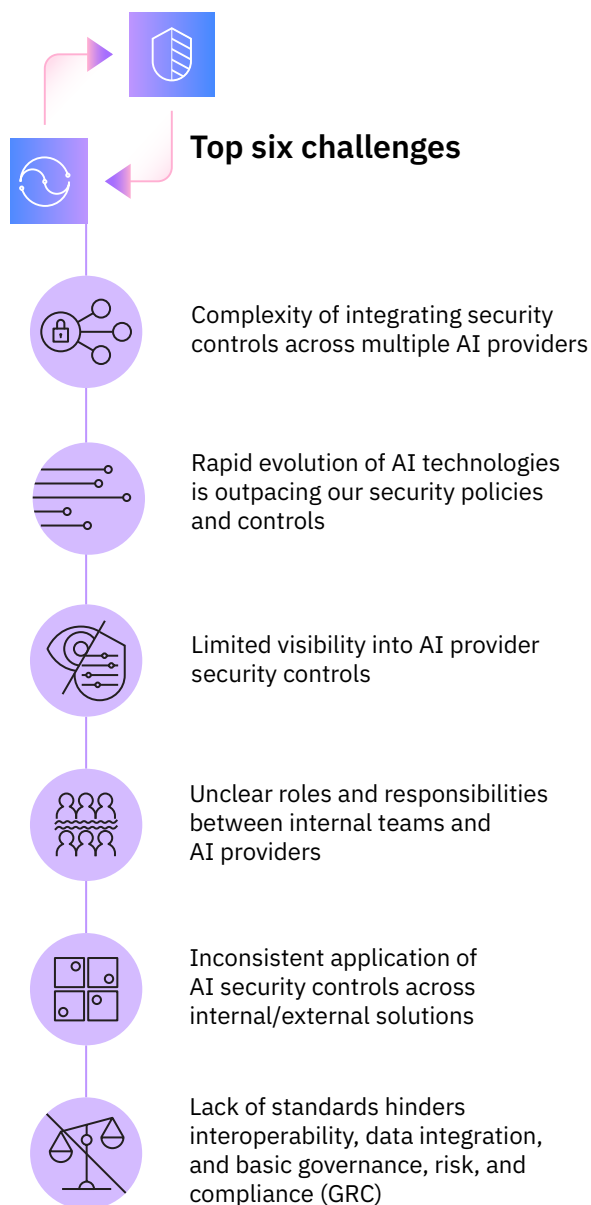
Scott Moser, CISO, Sabre

Four in 10 say their current approach to shared responsibility is either “minimally effective” or “not effective.” The three biggest obstacles are complexity across multiple AI providers, rapid evolution of AI technologies, and limited visibility into AI provider security controls (see Figure 6).

These obstacles underscore the importance of choosing the right partners and communicating clearly about AI standards and agent protocols. Much like the transition to cloud infrastructure, the shift to enterprise AI and trusted autonomy will take time. There is no shortcut to reach the finish line next quarter, or even next year. It’s a quarter-by-quarter multi-year journey. But with agentic threats at the doorstep, the urgency is clear. Trusted partners can help scale a secure enterprise AI approach, turning targeted investments into a growth engine for the business.

FIGURE 6

What’s standing in the way of AI shared responsibility?



Q: What are the primary challenges your organization faces in making shared cybersecurity responsibility with AI providers more effective? n=1,000



Readiness assessment

Identifying your allies

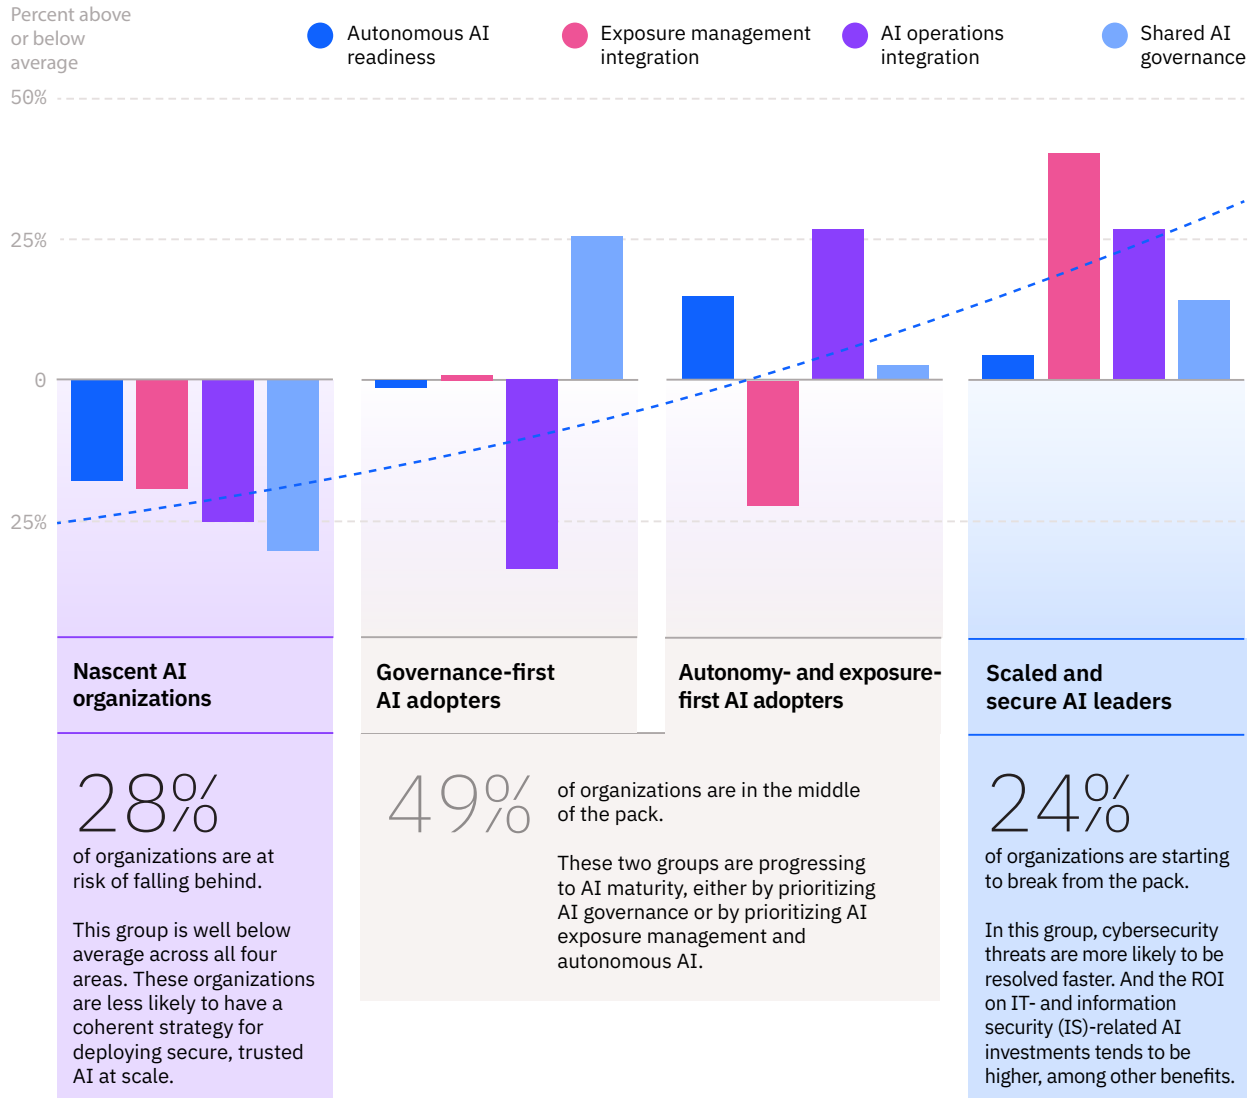
Key questions about your organization's shared-responsibility posture.

- Where do decision rights for AI security sit today—and where must they sit for agentic AI to scale safely across the enterprise and ecosystem?
- If an AI agent causes material harm today, can you clearly articulate who is accountable—internally and externally—and how that accountability is enforced?
- If you evaluated your AI partners tomorrow, how many could clearly demonstrate alignment to your secure AI strategy—across governance, visibility and accountability?
- As AI agents become more autonomous, what governance mechanisms must shift from policybased oversight to continuous, system level enforcement?

FIGURE 7

Secure AI leaders are emerging

What should organizations do to secure increasingly AI-infused operations? Our analysis points to four important areas of action—and identifies one group pushing ahead.



Good things happen when organizations integrate AI into their technology, security, and operations infrastructure, and when they adopt mature AI exposure management and shared governance practices (see Figure 7). Our analysis finds that cybersecurity threats are resolved faster. The ROI on IT- and information security (IS)-related AI investments rises. And advantages derived from technology and security operations grow.

We found that about one in four (24%) organizations—the scaled and secure AI leaders—have begun to break out from the pack across four key areas. In contrast, more than one in four (28%)—nascent AI organizations—are struggling to get started, and risk falling behind.

These two groups are a study in contrast between AI opportunity and AI disruption.



“We’re moving away from brittle, static playbooks toward context-based decisions powered by real-time signals. We see a fleet of specialized bots working in tandem as more powerful than rigid runbooks. The shift is learning how to use context to make good security decisions.”

Matthew Bissell, Senior Director, Cloud Security and Response, Sabre

Perspective

TELUS Digital: Securing AI-driven operations at global scale

As a global provider of customer experience and digital consulting, TELUS Digital operates at the intersection of sensitive data, advanced automation, and large-scale AI-enabled workflows. With more than 78,000 employees across 31 countries, the ability of the company—part of the Canadian telecom giant TELUS—to securely deploy and operate AI drives both customer trust and continued innovation.

TELUS Digital’s leadership recognized that traditional, tool-heavy security models were no longer sufficient for protecting AI-driven business processes. As AI adoption accelerated—across content moderation, cloud platforms, and proprietary AI solutions—the organization needed a consolidated security plus operations approach that could keep pace with machine-speed threats while reducing operator friction and improving security intelligence.

“We came to the conclusion that, given our size and spend, a platform approach would help us achieve what we wanted very quickly,” says Steve Jablonski, CISO of TELUS Digital. By embedding AI-driven analytics and automation directly into its security operations, machine learning–based threat detection and behavioral analytics now correlate signals across network, endpoint, cloud, and identity domains.

“Built-in analytics reduced the noise dramatically,” says Julio Vivas, Director of Cybersecurity at TELUS. “That shift alone changed how our analysts spend their time—less triage, more proactive defense.”

TELUS Digital is now expanding scope to evaluate secure AI runtimes and APIs. “What excites us is the visibility into how AI APIs are being used and where malicious activity may be occurring,” Jablonski explains. “That’s critical as AI becomes core to our operations.”

“Our approach has enabled us to accelerate the overall deployment of security to the organization, giving us significantly more visibility and control across a global landscape,” Jablonski says.¹⁴

65% faster time to respond to security incidents.

TELUS Digital achieved this through automated investigation and response workflows.

38 million threats blocked in just 30 days.

Using AI-driven inspection across upgraded firewalls and with a new cloud security management tool in place, the company gained centralized visibility across its entire network.

Action guide

Getting to secure AI at scale: what to do now

AI is rapidly expanding across the enterprise. Applications, agents, models, and data are increasingly connected and capable of taking action on behalf of users. This creates new opportunities for innovation and growth. But it also introduces new attack surfaces that traditional security approaches were not designed to defend.

Achieving secure AI at scale requires new thinking and action on multiple fronts. Old security assumptions and practices need to be cast aside. New infrastructure needs to be built and integrated to power the shift to secure agentic and autonomous AI. And the relationship between humans and technology needs to be reimaged as machine-speed threats and digital labor models evolve.

What follows is a guide to achieving greater security performance and resilience through AI investments. Its four components are not sequential. One area is not more important than others. Robust AI asset governance, an integrated AI-first operations and security environment, a mature AI exposure management posture, and a strategically aligned security ecosystem all work together to enable new value.




Tightly scope and govern every AI asset.

- **Enforce zero trust authentication** for every AI asset, service, and event. Define and segment AI workflows at a granular level.
- **Treat AI agents as first class identities:** authenticated, authorized, monitored, and supervised in real-time.
- **Build cross functional governance that spans domains, clouds, and business units,** ensuring autonomous actions remain visible, explainable, reversible, and compliant.
- **Use runtime security to secure the live use of AI,** preventing model manipulation and prompt injections.



Make AI exposure management your organization's security baseline and connective tissue.

- **Unify visibility across hybrid environments** so risk can be assessed on demand.
- **Partner with key providers, suppliers and industry peers to adopt common standards** for AI risk scoring and remediation.
- **Use contextual asset management insights to assess risks relative to AI models and agents.** Reconcile assets, configurations and vulnerabilities into a single, holistic picture. Build exposure scores grounded in business context and criticality.
- **Empower agents to act.** Let agents automate prioritization and remediation, and handle time sensitive insights at machine speed. Keep humans focused on governance, escalation, and training. Assess agent performance and update agents to reflect lessons learned.



Adopt an integrated architecture to lay the groundwork for trusted autonomy.

- **Work with your security, technology, and AI counterparts to define the integrated operations environment** where AI models, agents, data, and assets can scale securely, predictably, and reliably across domains.
- **Focus on IT/IS interoperability and common governance—not just tool integration.** Build the systemic framework that allows AI agents to collaborate across the security lifecycle. Plan for scale by adopting standards for AI agent communication and interaction (e.g. MCP, ACP, A2A).
- **Build resilience by weaving AI into core infrastructure and operations.** Eliminate AI, security, and technology silos by prioritizing these practices:
 - Define policies to manage AI exposure risk across the entire operations lifecycle
 - Create a reference architecture to support the secure administration of AI across cloud environments
 - Adopt a maturity framework to improve data governance for AI systems
 - Implement AI-specific governance, risk, and compliance controls specific to your sector/industry and key jurisdictions.



Build a secure AI ecosystem where responsibility is shared and accountability is operationalized.

- **Decide what AI infrastructure you want to own and what AI outcomes you want to outsource.** Then define a security posture consistent with that strategy.
- **Assess current AI procurement practices to rationalize your AI solution and vendor footprint.** Articulate your internal, second, third, and n-party governance practices. Choose partners aligned to your organization's AI strategy and success factors.
- **Pick strategic partners that fill critical operational and security gaps** and standardize security practices in collaboration with them. Engage with key suppliers to define key performance metrics and objectives, then create incentives to learn from doing.


IBM CISO/CIO Q&A

Bring IT and IS together for secure AI operations

IBM's senior technology and security leaders on their AI partnership

“While AI security remains deeply immature, AI for cyber defense is a huge opportunity,” says IBM CISO Koos Lodewijkx. His colleague, CIO Matt Lyteson, agrees. Yet neither of them are pointing fingers at the other. Here, the two leaders discuss their collaborative strategy for enterprise AI security—one platform, one source of truth—and how AI integrations can open the door to velocity with visibility.

How should leaders be thinking of AI at scale? It seems to be both an opportunity and a risk.



Koos Lodewijkx (CISO): It's both—and leaders need to strategize for that. Organizations are pushing AI from the top down, with boards and C-suites looking for immediate productivity gains. Our teams are tasked with delivering value almost overnight. That pace often outstrips the maturity of AI security standards and practices. And so risk isn't merely rising, it's compounding. Every new AI tool integration multiplies the attack surface. At scale across the organization, this looks like an exponential risk curve. We've already seen smaller AI integration vendors breached—attackers go where the credentials live.

Matt Lyteson (CIO): The acceleration of AI leads teams to skip essentials like least privilege and role-based access control. But AI doesn't exempt you from basic IT hygiene. If anything, it makes those basics even more important. AI can deliver massive value, but it's not a license to cut corners.

IBM CISO/CIO Q&A

So how can organizations balance two needs that are in tension: speed and safety?

Lodewijkx (CISO): It's a definite challenge, especially with shadow AI already here, creating basic perimeter and governance challenges. Employees can spin up consumer agents for a few dollars a month and unintentionally create significant risk. These tools now connect to calendars, inboxes, and collaboration apps, widening the blast radius. Proactive detection is a must.

Lytson (CIO): If someone wants a custom AI agent solution, we can support that but we put guardrails around it. We isolate it so it can't touch enterprise systems or sensitive data. We back that with procurement rules and an intake process for legitimate use cases, so innovation stays safe. And we know that many backend APIs were never designed for the identity and permission granularity AI requires. We're remediating those gaps as quickly as the standards evolve. Maintaining a consistent chain of identity from user to agent to backend is non-negotiable.

To move quickly and safely, IT [Information Technology] and IS [Information Security] must partner from the very beginning of an AI use case. If security comes in late, the architectural compromises are already baked in. Embedding IS early preserves both speed and integrity.

Lodewijkx (CISO): Right, and if IT and IS have diverging priorities, with respect to AI or otherwise, that's not a technical disconnect. It's a leadership disconnect. Both functions must serve the business with aligned goals and integrated strategies.

Is there anything that gets overlooked in the enterprise AI conversation?

Lodewijkx (CISO): Where you start matters. The biggest short-term productivity gains come when AI is applied at the workflow level, not at the consumer or individual gadget level. That's where value becomes measurable and sustainable. The real ROI doesn't come from personal productivity tools. It comes from redesigning enterprise workflows and embedding AI into them.

Lytson (CIO): I'll add one point to that. Given the speed and scale of change many are now looking for, you really need to maintain discipline and rebuild enterprise architecture muscle. IT leaders need to anchor on IT basics, platform adoption, business outcomes and architectural rigor to make AI sustainable. And teams need to transform in flight, actively acquiring new skills while they build. That's the only way to keep pace.

IBM CISO/CIO Q&A

From the CIO perspective, what's the winning strategy for minimizing AI-related security risks?

Lyteson (CIO): In my view, an enterprise platform approach is the only realistic way to maintain consistency at scale. Letting everyone build wherever they want is unsustainable. At IBM, our 'license to drive' program ensures non-IT teams know their responsibilities before developing agents.

Creating platform-wide guardrails has to involve robust tool and agent scoping and identify management. Identity is the key to controlling AI. You have to enforce narrow scopes to keep agentic systems auditable. We design agents so each can only do one task, and that tight scoping preserves accountability. You can't let an LLM roam freely. Guardrails must enforce trusted sources, trusted tools and trusted outcomes, all bound to a consistent identity chain.

What makes you optimistic about the outlook for cyber defenders?

Lodewijkx (CISO): For me the optimism comes out of AI itself. AI capabilities are advancing incredibly quickly. We are seeing support for new use cases every three to six months. AI can be a uniquely powerful tool in addressing factors like complexity and acceleration, but also skill and capacity constraints. We need to work closely with our counterparts in technology and the business to make that a reality.

We talk about the speed and scale of AI-enabled threats. But those same capabilities can also work to our advantage. We can work faster and with more reach than we've ever had. There's potential for us to gain an order of magnitude increase in capacity that will help us handle more complex, context-rich alerts.

If I can use AI to improve my signal-to-noise ratio to make sure I catch the small percentage of high-impact threats, we're no longer chasing threats that don't exist. I'm now looking to redesign all of our security business processes from the ground up—not for incremental 30% productivity gains, but for 5 to 10 times the throughput and quality. AI gives us the chance to achieve a new level of resilience.

About the authors

[Mark Hughes](#)

Global Managing Partner,
IBM Security Services
IBM Consulting
[linkedin.com/in/markhughesibm/](https://www.linkedin.com/in/markhughesibm/)

Mark leads IBM's team of experts helping organizations transform security into a business enabler. His role spans the sales and services delivery of threat detection and response, data security, cloud security, IAM, infrastructure, risk management and ecosystem partnerships. Mark's career spans over two decades, including roles as President of Security at DXC Technology and Chief Executive at BT Security, a leading global telecommunications provider.

[Ian Swanson](#)

VP, Product, AI Security
Palo Alto Networks
[linkedin.com/in/ianswanson/](https://www.linkedin.com/in/ianswanson/)

Ian Swanson is Vice President of AI Security Products at Palo Alto Networks, where he leads strategy and product innovation to secure the next generation of artificial intelligence applications. He joined Palo Alto Networks following its acquisition of Protect AI, the company he founded and led as CEO, which became the industry's leading platform for managing risk and securing AI and machine learning environments end-to-end.

[Srinivas Tummalapenta](#)

CTO, IBM Security Services
IBM Distinguished Engineer
[linkedin.com/in/srinivastummalapenta/](https://www.linkedin.com/in/srinivastummalapenta/)

As CTO of IBM Security Services, Sriniv partners with product partners, solution architects, and strategy leaders to define and deliver security solutions across the NIST Cybersecurity Framework.

[Gerald Parham](#)

Global Research Leader, Security & CIO
IBM Institute for Business Value
[linkedin.com/in/gerryparham](https://www.linkedin.com/in/gerryparham)

Gerry's research insights have appeared in publications such as *The Wall Street Journal*, *Forbes*, *CIO*, *Cyber*, and *Infosecurity Magazine*. His papers have been recognized as among the leading examples of thought leadership in the world.

Contributors

We would like to thank the following individuals for their significant contributions to these materials:

Jad Abdulsalam, Cybersecurity domain expert, industrial sector

Matthew Bissell, Sr. Director, Cloud Security & Response, Sabre

Steve Jablonski, Vice President, Information Security & Chief Security Officer, TELUS Digital

Scott Moser, SVP and Chief Information Security Officer, Sabre

Koos Lodewijkx, VP & Chief Information Security Officer (CISO), IBM

Matt Lyteson, Chief Information Officer (CIO), VP Technology Platforms Transformation, IBM

Robert Daniels, Global Director, GSI Business Development, Palo Alto Networks

Kenne Miller, Director of Global Brand Campaign, Palo Alto Networks

Joanna Cullinan, Senior Marketing Manager, Palo Alto Networks

Hitendar Sethi, Principal Product Marketing Manager, AI, Palo Alto Networks

Webbo Chen, Public Relations Manager, AI Security, Palo Alto Networks

Tim Van den Heede, Vice President, IBM Security Services (Palo Alto Networks partnership)

Shlomi Kramer, Executive Partner, IBM Security Services (Palo Alto Networks partnership)

Aastha Kaul, Marketing Manager, IBM-Palo Alto Networks Partnership, IBM Security Services

Ameena Azmath, Product Marketing Manager, IBM Security Services

Priya Kurien, Research Director IBM Institute for Business Value

Andrew Womack, Creative Director, IBM Institute for Business Value

Jeremy Gantz, Editorial Lead, IBM Institute for Business Value

Heba Nashaat, Business analytics consultant, Research Hub, IBM Institute for Business Value

Nathan Boudreaux, Design Lead, IBM Institute for Business Value

Research methodology

To examine how organizations are securing their AI estate and adapting cybersecurity operations in response to advanced AI threats, the IBM Institute for Business Value, in collaboration with Palo Alto Networks, conducted a global, cross-sectional executive survey during Q4 2025 and Q1 2026. The study looks at organizational practices, adoption levels, and performance outcomes related to AI, cybersecurity resilience, exposure management, and identity risk.

The survey collected responses from 1,000 senior executives, including 600 CISOs, CTOs, CIOs, CDOs, and CAIOs, and 400 CEO, COOs, CFOs, and Chief Risk Officers. Participants represent a broad range of industries from 17 countries across North America, Europe, Asia Pacific, the Middle East, and Latin America.

The research was designed to ensure coverage of diverse operating environments and AI maturity levels.

Survey measures include adoption percentages, agreement-based ratings, and maturity classifications. Where needed, responses were combined into summary measures to allow consistent comparison of organizations across derivative analyses.

1,000

C-level executives

17

Countries

20

Industries



Analytical frameworks

To assess cybersecurity operational resilience, we constructed a composite resilience index based on respondents reported mean times to identify, contain, and recover from cybersecurity incidents. These measures were converted to a common time basis and combined to align with a 0 to 100 scale, where higher scores indicate faster and more effective operational response. We then evaluated the relative influence of key organizational factors on resilience outcomes, using composite measures of AI workload penetration, AI governance maturity, AI integration into IT operations, policy and ecosystem alignment, security operations automation, and technology fragmentation.

To isolate the most influential underlying practices, we conducted a more detailed analysis within the strongest contributing areas. This allowed us to identify specific contributors to resilience, including generative AI workload penetration, AI integration into cloud infrastructure, and maturity in AI data governance and lifecycle management.

To assess the potential impact of advanced AI adoption on value realization, we measured the share of workloads currently enabled with agentic AI, multi-agent orchestration, and autonomous AI. We then modelled forward-looking adoption scenarios by scaling current usage levels within realistic bounds. Changes in AI adoption were translated into changes in realized ROI by estimating how incremental gains in AI adoption affect value creation.

We estimated this using a purpose-built model that accounts for the level of AI adoption, the sensitivity of ROI to relative changes in adoption, and a penalty margin for projects that were cancelled or delayed. The same modelling assumptions were applied across scenarios to ensure consistent and conservative comparisons.

To quantify organizational exposure to AI-related identity risk, we created an AI identity risk score on a 0 to 100 scale, where higher scores indicate greater AI risk exposure. The score reflects the strength of key protections and governance practices, including zero trust coverage for non-human identities, AI governance maturity, clarity of AI security policies, effectiveness of shared responsibility with AI providers, and the urgency of closing identity security gaps.

We then assessed whether organizations with higher AI identity risk experience more cybersecurity incidents, while accounting for differences in organizational size. Results are presented as expected incident levels per 10,000 employees to illustrate the operational impact of elevated identity risk.

Finally, organizations were segmented using standardized indicators capturing readiness for autonomous AI operations, exposure management integration, AI operations integration, and shared governance practices. Descriptive comparisons were used to contrast more mature and less mature groups across outcomes related to cybersecurity resilience, expected ROI, and competitive advantage.

About Research Insights

Research Insights are fact-based strategic insights for business executives on critical public- and private-sector issues. They are based on findings from analysis of our own primary research studies. For more information, contact the IBM Institute for Business Value at iibv@us.ibm.com.

IBM Institute for Business Value

For two decades, the IBM Institute for Business Value has served as the thought leadership think tank for IBM. What inspires us is producing research-backed, technology-informed strategic insights that help leaders make smarter business decisions.

From our unique position at the intersection of business, technology, and society, we survey, interview, and engage with thousands of executives, consumers, and experts each year, synthesizing their perspectives into credible, inspiring, and actionable insights.

To stay connected and informed, sign up to receive IBV's email newsletter at ibm.com/ibv. You can also find us on LinkedIn at ibm.co/ibv-linkedin.

The right partner for a changing world

At IBM, we collaborate with our clients, bringing together business insight, advanced research, and technology to give them a distinct advantage in today's rapidly changing environment.

About Palo Alto Networks

As the global cybersecurity leader, Palo Alto Networks (NASDAQ: PANW) is dedicated to protecting our digital way of life via continuous innovation. Trusted by more than 70,000 organizations worldwide, we provide comprehensive AI-powered security solutions across network, cloud, security operations and AI, enhanced by the expertise and threat intelligence of Unit 42®. Our focus on platformization allows enterprises to streamline security at scale, ensuring protection fuels innovation. Explore more at www.paloaltonetworks.com.

Related reports

Cybersecurity 2028: Your workforce, built for the AI frontier


ibm.com/thought-leadership/institute-business-value/en-us/report/cybersecurity-ai-workforce

Secure by design, smarter with AI: Redefining cyber resilience for the age of intelligent threats

ibm.com/thought-leadership/institute-business-value/en-us/report/ai-secure-design-cyber-resilience

Capturing the cybersecurity dividend: How security platforms generate business value

ibm.com/thought-leadership/institute-business-value/en-us/report/unified-cybersecurity-platform



Subscribe to our IdeaWatch newsletter

Just the insights. At your fingertips.
Delivered monthly.

Brought to you by the IBM Institute for Business Value, ranked #1 in thought leadership quality by Source Global Research for the second consecutive year.

Research-based thought leadership insights, data, and analysis to help you make smarter business decisions and more informed technology investments.

Subscribe now: ibm.co/ideawatch



Notes and sources

1. Tidy, Joe. "AI firm claims Chinese spies used its tech to automate cyber attacks." *BBC News*. November 14, 2025. <https://www.bbc.com/news/articles/cx2lzmygr84o>; Sabin, Sam. "Chinese Hackers Used Anthropic's Claude AI Agent to Automate Spying." *Axios*. November 13, 2025. <https://www.axios.com/2025/11/13/anthropic-china-claude-code-cyberattack>; Martin, Andrew and Millan, Carolina. "Hacker Used Anthropic's Claude to Steal Mexican Data Trove." *Bloomberg*. February 25, 2026. <https://www.bloomberg.com/news/articles/2026-02-25/hacker-used-anthropic-s-claude-to-steal-sensitive-mexican-data>
2. Anthropic. "Disrupting the First Reported AI-Orchestrated Cyber Espionage Campaign." November 13, 2025. <https://www.anthropic.com/news/disrupting-AI-espionage>; Derek Thompson. "Anthropic Was Hacked—and It Changed How We Should Think About AI and Cybersecurity." *The Atlantic*, November 2025. <https://www.theatlantic.com/technology/2025/11/anthropic-hack-ai-cybersecurity/685061/>; OpenAI. "Strengthening Cyber Resilience as AI Capabilities Advance." OpenAI, December 10, 2025. <https://openai.com/index/strengthening-cyber-resilience/>; Jubu Babu. "OpenAI Warns New Models Pose 'High' Cybersecurity Risk." *Reuters*, December 10, 2025. <https://www.reuters.com/business/openai-warns-new-models-pose-high-cybersecurity-risk-2025-12-10/>
3. Chuck Brooks. "AI Polymorphic Threats Are Forcing a Rethink of Cybersecurity." *Forbes*, February 21, 2026. <https://www.forbes.com/sites/chuckbrooks/2026/02/21/ai-polymorphic-threats-are-forcing-a-rethink-of-cybersecurity/>; Cyber Desserts. "AI Security Threats: Complete Guide to Attack Vectors." *Cyber Desserts Blog*, December 12, 2025. <https://blog.cyberdesserts.com/ai-security-threats/>
4. Udinmwun, Efosa. "'Weaponized AI' Could Be the Biggest Security Threat Facing Your Business This Year—Here's What Experts Say You Should Be on the Lookout For." *Tech Radar*. January 24, 2026. <https://www.techradar.com/pro/weaponized-ai-could-be-the-biggest-security-threat-facing-your-business-this-year-heres-what-experts-say-you-should-be-on-the-lookout-for>; Thompson, Derek. "Anthropic's Hack Shows How AI Is Changing Cybersecurity." *The Atlantic*. November 2025. <https://www.theatlantic.com/technology/2025/11/anthropic-hack-ai-cybersecurity/685061/>; McMillan, Robert. "AI Hackers Are Coming Dangerously Close to Beating Humans." *Wall Street Journal*. December 11, 2025. <https://www.wsj.com/tech/ai/ai-hackers-are-coming-dangerously-close-to-beating-humans-4afc3ad6>
5. Calvin Slangen. "The Invisible Threat: Why 79% of Cyberattacks No Longer Use Malware." *LinkedIn Pulse*, October 7, 2025. <https://www.linkedin.com/pulse/invisible-threat-why-79-cyberattacks-longer-izkge/>.

6. IBM X-Force. *X-Force Threat Intelligence Index 2026*. Armonk, NY: IBM Corporation, 2026. <https://www.ibm.com/reports/threat-intelligence>; Sam Sabin. "AI Is About to Supercharge Cyberattacks." *Axios*, October 25, 2025. <https://www.axios.com/2025/10/25/ai-is-about-to-supercharge-cyberattacks>; CybersecurityHQ Editorial. "When Trust Fails: The Collapse of Security Verification Across Hardware, Software, and Identity Layers." *CybersecurityHQ*, November 1, 2025. <https://dispatch.cybersecurityhq.com/p/when-trust-fails-the-collapse-of-security-verification-across-hardware-software-and-identity-layers>.
7. ThreatDown. *2026 State of Malware: The Dawn of Machine-Scale Cybercrime*. Santa Clara, CA: ThreatDown (Malwarebytes), 2026. <https://www.threatdown.com/dl-state-of-malware-2026/>
8. Charlie Osborne. "Google Spots Malware in the Wild That Morphs Mid Attack, Thanks to AI." *ZDNET*, November 6, 2025. <https://www.zdnet.com/article/google-spots-malware-in-the-wild-that-morphs-mid-attack-thanks-to-ai/>.
9. Palo Alto Networks Unit 42. *2026 Global Incident Response Report*. Santa Clara, CA: Palo Alto Networks, 2026. <https://www.paloaltonetworks.com/resources/research/unit-42-incident-response-report>.
10. Sudhakar Tiwari. "Identity Is the New Perimeter: Why IAM Is the Frontline of Cybersecurity." *Cyber Defense Magazine*, December 18, 2025. <https://www.cyberdefensemagazine.com/identity-is-the-new-perimeter-why-iam-is-the-frontline-of-cybersecurity/>.
11. CybersecurityHQ Editorial. "The Authorization Gap: When AI Acts Without Sanction." *CybersecurityHQ Newsletter*, December 13, 2025. <https://newsletter.cybersecurityhq.com/p/the-authorization-gap-when-ai-acts-without-sanction>.
12. Nate Rattner and Jason Dean. "Big Tech's \$670 Billion AI Push Dwarfs Spending on Moon Landing, U.S. Railroads." *Wall Street Journal*, February 7, 2026. <https://www.wsj.com/tech/ai/ai-spending-tech-companies-compared-02b90046>; Khadija Said. "Microsoft AI CEO Mustafa Suleyman Says Building Superintelligence Could Cost 'Hundreds of Billions.'" *Business Insider*, December 2025. <https://www.businessinsider.com/microsoft-ai-ceo-mustafa-suleyman-cost-hundred-billions-superintelligence-2025-12>.
13. Open Worldwide Application Security Project (OWASP). *OWASP Top 10 for Agentic Applications for 2026*. OWASP GenAI Security Project, December 2025. <https://genai.owasp.org/resource/owasp-top-10-for-agentic-applications-for-2026/>.
14. Palo Alto Networks. "TELUS Digital Secures Customer Trust and Innovation with Palo Alto Networks." <https://www.paloaltonetworks.com/customers/telus-digital-secures-customer-trust-and-innovation>



© Copyright IBM Corporation 2026

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America | March 2026

IBM, the IBM logo, [ibm.com](https://www.ibm.com), and IBM Consulting are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at: [ibm.com/legal/copytrade.shtml](https://www.ibm.com/legal/copytrade.shtml).

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

This report is intended for general guidance only. It is not intended to be a substitute for detailed research or the exercise of professional judgment. IBM shall not be responsible for any loss whatsoever sustained by any organization or person who relies on this publication.

Examples presented are illustrative only. Actual results will vary based on client configurations and conditions and, therefore, generally expected results cannot be provided.