cloud security alliance®

Google Cloud Security

Survey Report

The State of AI Security and Governance

Acknowledgments

Lead Authors

Hillary Baron

Contributors

Stephen Lawton
Daniele Catteddu
Rich Mogull
John Yeoh
Anton Chuvakin
Douglas Ko

Graphic Design

Stephen Lumpe

About the Sponsor

Make Google part of your security team with unmatched threat visibility, a unified security platform, and Mandiant frontline experts — supercharged by AI.

Organizations can reduce digital risk and secure their AI transformation with the same cybersecurity specialists, capabilities, and secure enterprise platforms Google uses to keep more people and organizations safe online than anyone else in the world, powered by our industry-leading threat intelligence. AI enhances all of these components, enabling security teams to detect more threats, minimize toil, and take productivity to new levels.

Table of Contents

Executive Summary

The survey reveals a clear divide: **organizations with established AI governance are accelerating adoption with confidence, while the rest are moving quickly but without the structures needed to manage emerging risk.** As generative and agentic AI scale from pilots to production, the gap between governance "haves" and "have-nots" is becoming the strongest predictor of readiness. This year's CSA–Google Cloud survey shows security leaders stepping into a defining moment—working to secure AI systems even as they begin using AI to strengthen security itself. The market is evolving at remarkable speed, and governance is emerging as the foundation that determines whether adoption advances responsibly or outpaces an organization's ability to manage it.

> "As organizations move from experimentation to operational deployment, strong security and mature governance are the key differentiators for AI adoption."
>
> *– Dr. Anton Chuvakin, Security Advisor at Office of the CISO, Google Cloud*

Across every sector and region surveyed, organizations are now embedding AI into core operations and security workflows. However, the governance structures and talent pipelines needed to secure this adoption are still catching up.

By the Numbers – AI Security Snapshot

54%
of organizations **use public frontier LLMs** like GPT-4, Claude, or Gemini

52%
cite **sensitive data exposure** as their top security risk

27%
of respondents are **confident they can secure AI** used in core business operations

~60%
are using or plan to use **agentic AI** within 12 months

26%
have comprehensive AI governance policies; **44%** among large enterprises

Key Insights

1. Governance Is the Maturity Multiplier

Organizations with formal AI governance are significantly more advanced:

- **2x more likely** to adopt agentic AI
- **3x more likely** to train staff on AI security tools
- **2x more confident** in protecting AI systems

This reinforces governance as the foundation for responsible innovation—and a practical countermeasure to "shadow AI."

2. Security Becomes an Early AI Adopter

In a marked shift from past technology cycles, security teams are among the earliest adopters of AI. Over **90% of respondents are testing or planning to use AI for threat detection**, red teaming, and access control—highlighting the urgency and opportunity to embed AI into security from the outset.

3. LLM Consolidation within Multi-Model Strategies

Organizations are pursuing multi-model strategies—using an average of 2.6 models—but deployments are increasingly concentrated among the "Big Four": Gemini, Claude, GPT, and LLaMA. While this signals growing operational maturity, it also introduces new resilience, interoperability, and vendor lock-in concerns.

4. Executive AI Enthusiasm, Questions About Ability to Secure

Executive enthusiasm for AI remains high, yet most respondents (72%) were either not confident or neutral in their organization's ability to secure it. While **70%** report moderate to full leadership awareness of AI security implications, this gap underscores the need for deeper governance, education, and cross-functional collaboration.

5. AI Ownership Is Diffuse—Security Is Stepping Up

Responsibility for AI deployment is distributed across functions, but **security teams now lead AI protection in 53% of organizations**. Security is no longer an afterthought—it's emerging as both a stakeholder and a pioneer in responsible AI implementation.

6. Data Risk Takes Center Stage—But Model Risk & Safety Still Lags Behind

Organizations are prioritizing well-understood risks: 52% cite sensitive data exposure as their top concern, followed by regulatory compliance (50%). These traditional issues far outweigh newer AI-specific threats like model drift, prompt injection, and model theft—which remain acknowledged, but rarely ranked as top-tier. That's notable given that a data breach today carries an average global cost of **US $4.88**

million—making the stakes of treating AI security solely as an extension of existing privacy and compliance frameworks far too high. However, a deeper issue lies beneath the surface: just 21% of respondents call out model-level risks—including data poisoning, prompt injection, or other forms of model manipulation—as key concerns. Part of this reflects a maturity gap, but it also highlights a practical reality: model-focused risks are newer, and many organizations are still developing the skills and familiarity needed to use the emerging tooling that has only recently come to market.

What's Next?

This report calls on organizations to:

- **Accelerate AI governance** using frameworks like CSA's AICM or Google's SAIF, then supplemented where appropriate with independent assessments or advisory services.
- **Invest in AI+Cybersecurity skill building** through training, upskilling, and inter-team collaboration
- **Embed secure-by-design** principles into AI development workflows
- **Measure what matters**—from model integrity to policy adherence

In summary, the survey reveals a landscape where AI is moving faster—and security is catching up. The difference-maker is governance maturity: organizations that operationalize their policies today will be tomorrow's leaders in trustworthy AI adoption.

Key Findings

AI adoption is still early but accelerating rapidly, and this year's results show organizations moving from experimentation to meaningful operational use. What stands out is not just the pace of deployment, but the heightened awareness accompanying it: leaders are engaged, governance efforts are emerging, and teams are working to balance innovation with accountability. Across both dimensions of the survey—*securing AI systems* and *using AI to strengthen security*—organizations are making early progress, even as they continue to grapple with foundational challenges in risk understanding, data protection, staffing, and policy.

Key Finding 1:

Strong AI Governance Associated with Confidence, Risk Awareness, and Responsible Innovation

While organizations continue to build their AI security capabilities, the presence of formal governance policies stands out as one of the clearest predictors of maturity and readiness. **Only 26% of organizations report having comprehensive AI security governance policies in place,** but an additional **64% say they have some guidelines or are in the process of developing them**. These numbers show that while comprehensive governance remains the exception, most organizations recognize its importance and are taking steps to formalize it.

Has your organization established clear security governance policies and guidelines specifically for the development, deployment, and use of AI?

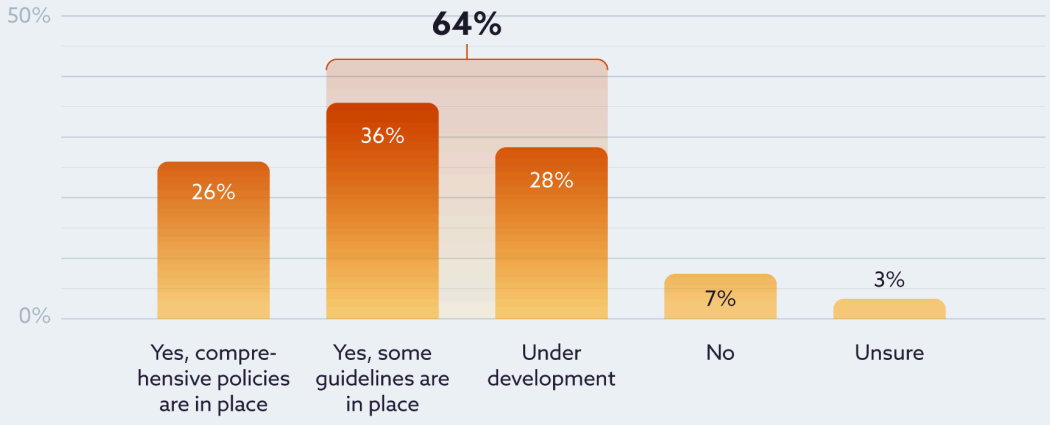

The data reveal a consistent pattern: mature governance is strongly associated with better outcomes across multiple dimensions of AI adoption and security.

2x

Organizations with comprehensive policies are nearly **twice as likely to report early adoption of agentic AI** (46%) compared to those with only partial guidelines (25%) or policies still in development (12%).

Organizations are **far more likely to have tested AI capabilities for security**, with 70% reporting experimentation compared to 43% of those with partial governance and just 39% among those still developing their policies.

Forty percent of organizations with mature governance are **already using agentic AI tools for cybersecurity**, versus only 11% with partial guidelines and 10% in early development.

Governance maturity is also tied to leadership awareness and organizational confidence. Among **organizations whose boards fully understand AI's security implications, 55% have comprehensive governance policies**. Those with established governance also report higher confidence in protecting AI systems — 48% describe themselves as confident, compared to 23% with partial guidelines and 16% still developing governance. These results show that formal governance helps align leadership visibility, risk understanding, and operational assurance.

55%

of organizations have comprehensive governance policies in place

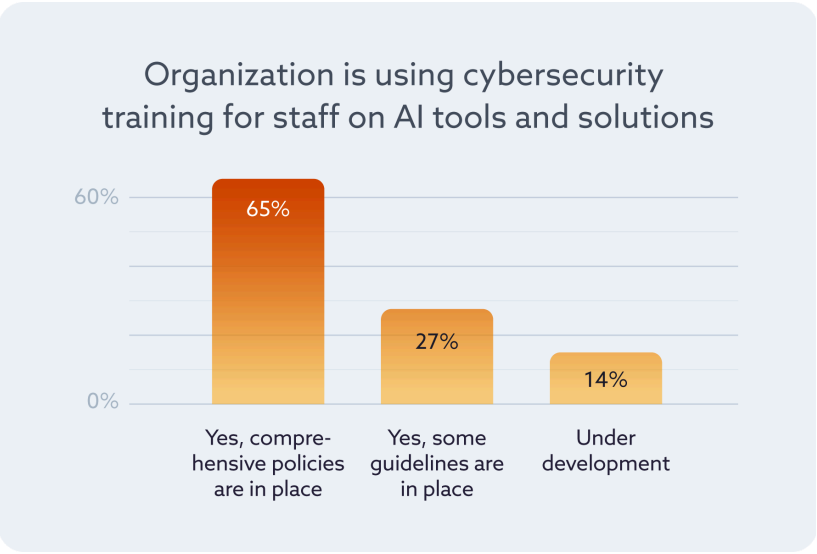

The connection extends to workforce readiness as well. **Sixty-five percent of organizations with comprehensive governance policies are already training staff on AI tools**, while just 27% with partial policies and 14% with developing policies are doing the same). Training is a key enabler of responsible AI adoption, and these numbers indicate that governance may help move organizations from awareness to action — ensuring staff know how to use AI tools securely and effectively.

Finally, robust governance may help organizations avoid the rise of "shadow AI" (unsanctioned or unmanaged AI use that introduces compliance and data privacy risks). As organizations formalize their governance, AI adoption becomes encouraged and structured rather than restricted, reducing the incentive for employees to use unapproved tools. This approach contrasts with early cloud and SaaS adoption cycles, where a lack of governance often led to uncontrolled use and security blind spots.

These findings highlight the central role of governance in advancing AI security maturity. Organizations that invest early in comprehensive governance frameworks are better positioned to innovate responsibly, maintain leadership alignment, and build staff confidence. Governance provides the foundation for sustainable AI adoption, bridging the gap between enthusiasm and execution, and ensuring that innovation moves forward securely.

Key Finding 2:

Security Becomes an Early Adopter of AI - Shift From Lagging to Leading

This year's results signal a turning point: security teams are becoming early adopters of AI, not followers. Historically, security functions have focused on securing implementations of emerging technologies. However, the 'AI for security' use case is so compelling, compared to past technologies, that apparently has increased the appetite for experimentation with the new technology and this in turn will help accelerate the maturity of 'Securing AI'. In fact, 13% of organizations report that security is responsible for adoption of AI. In this new paradigm, security has an opportunity to be embedded in AI adoption rather than an afterthought.

Nearly half of organizations (48%) report that they have already tested AI capabilities in security, and another 44% plan to do so within the next year. This means that over 90% are at least exploring how AI can improve detection, investigation, or response processes. The numbers are even more encouraging for agentic AI – autonomous or semi-autonomous systems used for incident response, red teaming, or adaptive access control. Nineteen percent are already using these tools, and another 47% plan to adopt them within the next year. With only 10% reporting no plans to invest, this represents a major inflection point: AI is not just a future concept for cybersecurity, it is becoming a near-term operational reality.

The growth in AI use for security stands in sharp contrast to 2024, when resource limitations and staffing shortages were the most frequently cited barriers to AI implementations in security. A year ago, a third of organizations said skill gaps and lack of knowledge. The current results suggest that organizations have made tangible movement toward implementation: **90% are actively providing or planning a combination of general security awareness and cybersecurity-focused training for AI tools**.

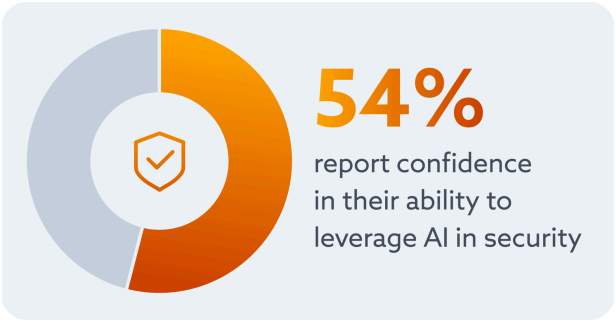

Confidence in using AI in security is also rising, particularly among organizations with comprehensive governance frameworks – **54% of those with formal governance policies report confidence in their ability to leverage AI in security**, compared to just 25% among those with partial guidelines.

The implications of this change are significant. Security's early embrace of AI could help close long-standing gaps between security and operations, creating a shared understanding of the technologies driving business innovation. As AI continues to transform digital environments, security professionals who use AI themselves will be better positioned to understand its risks, capabilities, and operational dependencies – making them more effective partners across the organization.

These findings suggest that **AI in security has reached an inflection point**. After years of being cautious followers, security teams are now among the earliest adopters of AI, demonstrating both curiosity and confidence. This proactive posture not only improves defensive capabilities but also reshapes the role of security – from a function that reacts to new technologies, to one that helps lead and shape how they are safely deployed.

Key Finding 3:

Enterprise LLM Adoption Accelerates Toward a Multi-Model Future Dominated by a Few Providers

Large language model (LLM) adoption has moved from experimentation to enterprise-scale deployment, marking a major inflection point in the evolution of AI strategy. **More than half of organizations (54%) report using public frontier models such as GPT-4, Claude, or Gemini**, while nearly half (47%) are using vendor-hosted private models through services like Google Vertex AI, Azure OpenAI, or Amazon Bedrock. Another 44% are leveraging self-hosted or open-source models in their own cloud or on-prem environments, and just 12% report no plans to use LLMs. In 2024, only 22% of organizations were actively using generative AI and 55% were still planning for adoption. One year later, that intent has clearly translated into action. GenAI has shifted from a forward-looking investment to an operational capability.

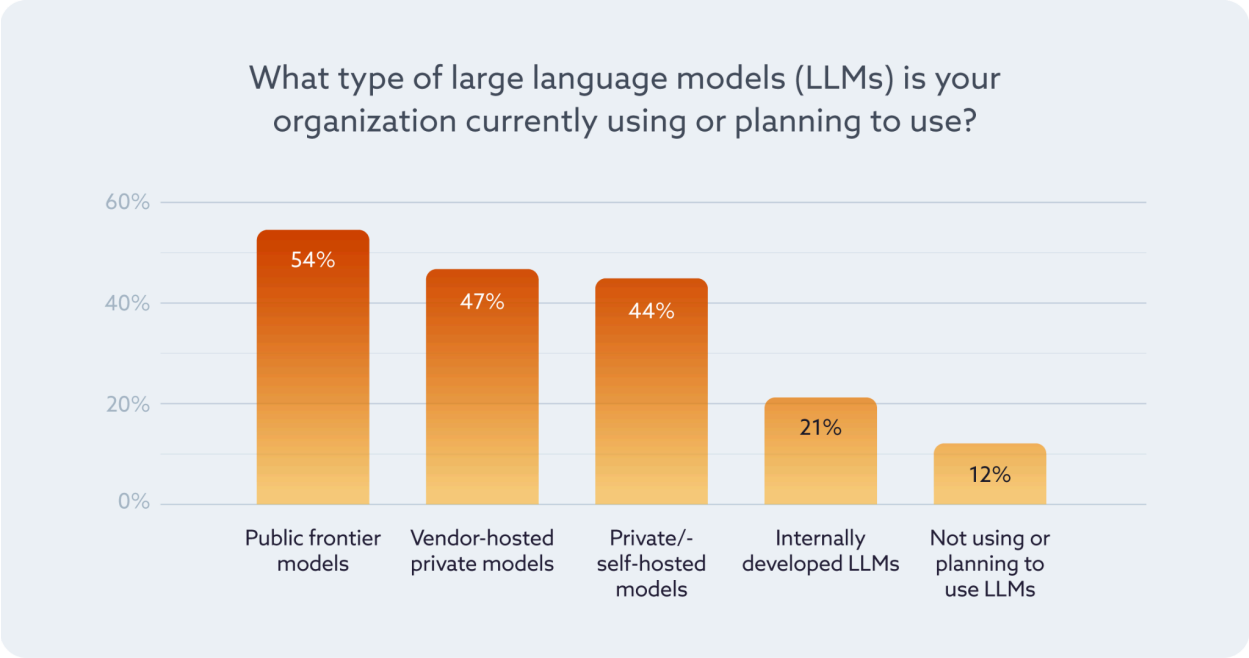

This expansion, however, is not evenly distributed. Adoption is rapidly consolidating around a small number of major providers. GPT (OpenAI) leads with 70% of organizations reporting use or evaluation, followed by Gemini (Google) at 48%, Claude (Anthropic) at 29%, and LLaMA (Meta) at 20%. Together, these "Big Four" account for the vast majority of enterprise deployments, signaling an ecosystem increasingly defined by a handful of dominant players. The concentration of adoption echoes earlier patterns seen in cloud computing, where early innovation gave way to consolidation around large hyperscalers.

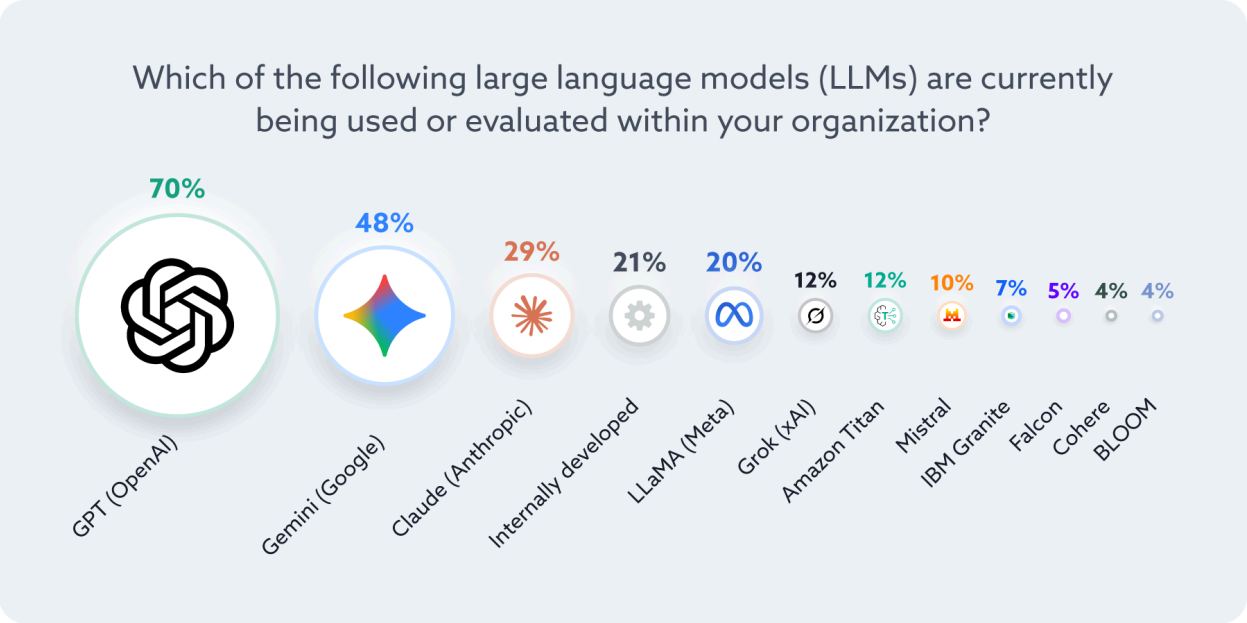

Organizations also report using an average of 2.6 different models, suggesting that **many are pursuing a multi-model**. Rather than standardizing on a single platform, they are combining different models for specific business use cases. This mirrors broader cloud strategies that blend public, private, multi, and hybrid environments—allowing organizations to balance innovation, data governance, and risk.

These trends illustrate a pivotal moment in enterprise AI maturity. LLMs are no longer an emerging technology; they are becoming foundational digital infrastructure. Yet with that growth comes new concentration risks, as dependence on a small number of providers introduces potential challenges in resilience, interoperability, and governance. As organizations continue to integrate LLMs into core operations, diversification strategies, risk frameworks, and open standards will be essential to ensuring that the next phase of AI adoption is not only scalable but also sustainable and secure.

As LLMs become foundational infrastructure, organizations now face the challenge of securing increasingly complex, multi-model environments.

Key Finding 4:

Leadership Enthusiasm for AI Outpaces Understanding of Potential Security Risks

While AI adoption has accelerated across the enterprise, many leaders remain more enthusiastic about its potential than aware of its risks. The majority of respondents **(72%) are neutral or lack confidence in their organization's ability to execute a security strategy for AI**, while 51% are neutral and 21% say they are not confident. In 2024, just 4% said they were not confident and a majority rated themselves confident (25%) or very confident (26%). These moderate confidence levels suggest that as AI systems move from pilot to production, organizations are recognizing the depth of the security challenge—and realizing they may not yet have the skills or resources to meet it.

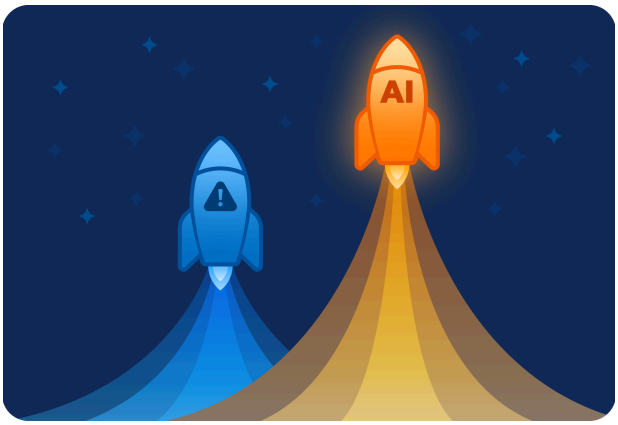

At the same time, leadership remains heavily invested in advancing AI adoption. In 2024, 82% of organizations said their executive leadership was actively pushing for AI initiatives. While 70% of organizations report moderate to full leadership awareness of AI's security implications, this awareness understandably remains a work in progress given the speed of technological change. **Executive enthusiasm for AI continues to outpace confidence in managing its risks**, underscoring the importance of strengthening governance capabilities over time.

These findings reveal a critical inflection point in organizational readiness. AI adoption has become a board-level priority, but understanding its security implications has not matured at the same pace. To close this gap, leadership must evolve from being champions of AI innovation to stewards of AI risk—fostering deeper collaboration with security teams, investing in specialized expertise, and integrating AI governance into enterprise risk management. Part of this uncertainty may also stem from still-evolving ownership structures around AI deployment and protection.

Key Finding 5:

Responsibility for AI Deployment Is Distributed Across Teams, but Security Ownership Is Clearly Emerging

Ownership of AI within organizations remains distributed, reflecting both the complexity of implementation and the early stage of operational maturity. When asked which **team is primarily responsible for AI deployment, 20% identified a dedicated AI or ML team**, followed closely by the IT department (19%) and cross-functional groups (16%). The remainder is divided among the security team (13%), senior leadership (9%), and data science or analytics teams (8%).

While AI governance responsibilities remain distributed across functions, early signs of consolidation may be emerging. In 2024, 74% of organizations reported plans to establish teams dedicated to governing the secure use of AI, and many now appear to be following through with the formation of AI and ML teams. This trend suggests that today's dispersed structures may mature into more formalized governance models over time, but they have yet to fully materialize as the technology is still in the early stage of adoption.

Security responsibilities, however, appear more clearly defined. **Over half of respondents (53%) say the security team is primarily responsible for securing AI systems**, with another 18% pointing to cross-functional teams and 11% to IT. Compared to deployment ownership, this represents stronger alignment with traditional cybersecurity structures.

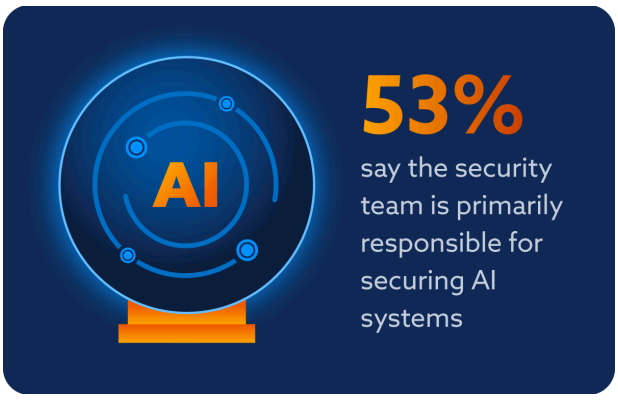

In many organizations, AI security is being integrated into existing governance frameworks rather than handled separately—mirroring earlier technology transitions such as cloud and SaaS adoption, where security teams gradually assumed responsibility once technologies matured.

Funding patterns provide additional insight into how accountability is forming. Nearly half (49%) report that the Chief Information Security Officer (CISO) oversees funding for AI security tools, followed by the CTO (36%) and CIO (33%). Business unit owners (25%) and emerging AI leadership roles, such as Chief AI Officers (11%), also play a role—indicating that financial responsibility for AI security is shared between operational and strategic leaders. This mix reflects an evolving governance model where AI security is treated as both a technical and business investment.

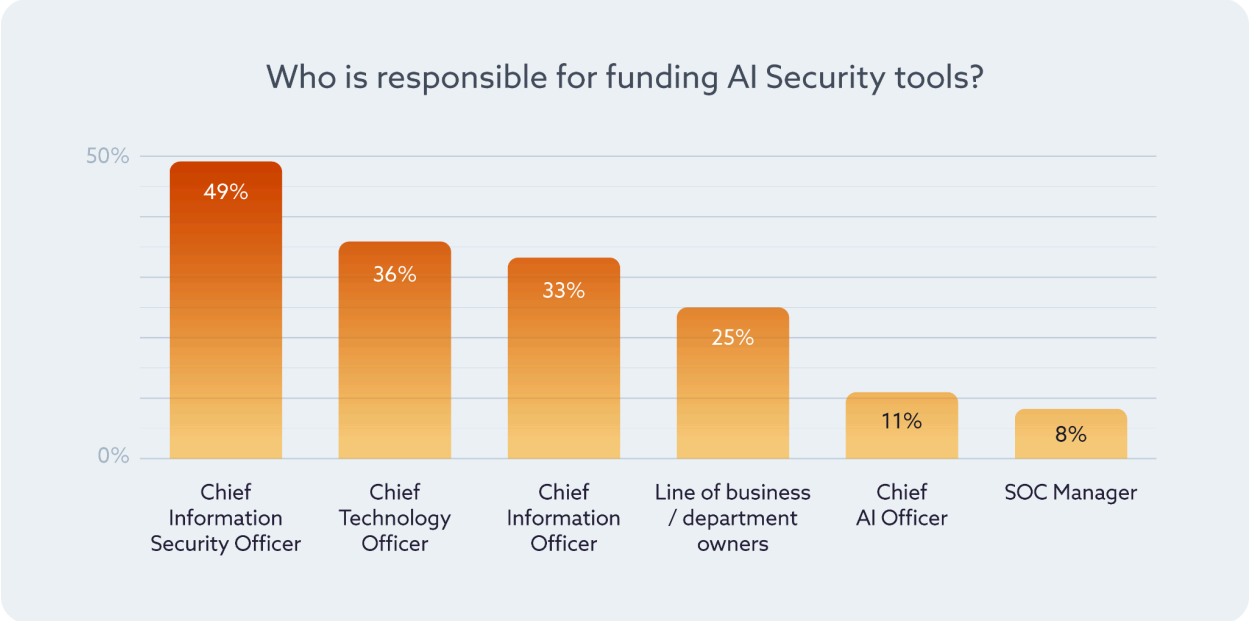

Taken together, these findings indicate that organizations are still refining how AI fits within existing operational and governance structures. Security ownership is solidifying under established teams, but deployment and funding responsibilities remain diffuse. The rise of dedicated AI/ML teams reflects progress from last year's intentions to create formal governance groups, but the current fragmentation points to an ongoing need for clearer accountability and coordination. The fact that security teams are assuming primary responsibility so early in AI's maturity may mark a notable shift—one where security itself is becoming an early adopter, shaping the guardrails for responsible AI innovation.

Key Finding 6:

Understanding AI Risk and Closing Skills Gaps Are the Biggest Challenges with Securing AI

Even with clearer ownership emerging, organizations continue to face significant challenges in building the skills and risk understanding required to secure AI effectively. Organizations cite **understanding AI risks (61%)**, **skill gaps (53%)**, and **lack of knowledge among current staff (52%) as the top hurdles** to getting started with security for AI implementations.

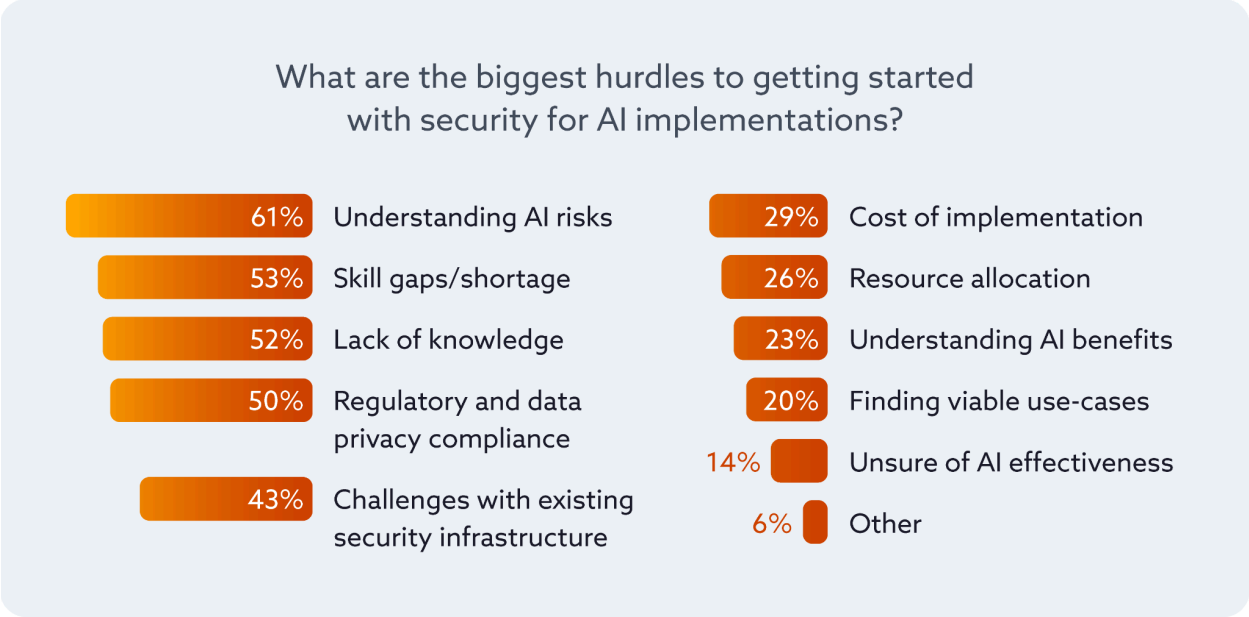

Additional barriers include regulatory and data privacy compliance (50%), integration with existing security infrastructure (43%), and practical constraints like cost (29%) and compute/resource allocation (26%). Lower on the list are understanding AI benefits (23%), finding viable use cases (20%), and unsure of AI effectiveness (14%)—suggesting most organizations see the value, but are constrained by risk comprehension, skills, and compliance execution rather than by lack of business demand.

When asked to rank their top security concerns, **organizations overwhelmingly point to sensitive data exposure as their primary risk, with 52% ranking it as their most concerning issue**. This far exceeds all other risks—compliance challenges were next at just 16%, followed by model integrity compromise (12%) and data poisoning (10%). Far fewer respondents view prompt injection (5%) or model theft (5%) as top-tier threats. On the opposite end of the

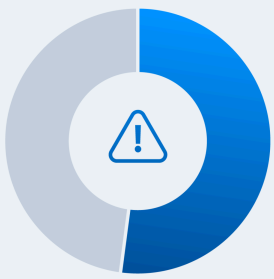

52%

rank **sensitive data exposure** as most concerning risk

spectrum, model theft was most frequently ranked least concerning (37%), underscoring that organizations are currently more focused on data leakage and regulatory exposure than on more technical or theoretical AI attack vectors. The prioritization of data and compliance risks suggests that many organizations are treating AI security as an extension of existing privacy and governance frameworks—reinforcing the perception that the most immediate danger lies not only in adversarial attacks, but in losing control of sensitive information through AI systems and integrations.

Privacy & Safety Insight

Fifty percent of respondents cite privacy and regulatory compliance as their top challenge, only 21% highlight risks that affect model reliability and integrity, including threats like data poisoning or prompt injections. This reveals a persistent gap between data protection and safety governance. Organizations should get an AI security assessment and determine the best approach to extend their privacy controls to include safety-by-design principles such as content integrity evaluation, hallucination mitigation, and bias testing within TEVV (Testing, Evaluation, Verification and Validation) workflows. These safeguards support both Google's Secure AI Framework (SAIF) and CSA's AI Controls Matrix (AICM).

While many practitioners still approach AI workloads as an extension of cloud environments, the underlying risk landscape is shifting. Traditional cloud-native issues—misconfiguration, network exposure, and access control weaknesses—now intersect with AI-specific threats such as prompt injection, model-output data leakage, and model drift. These risks introduce behavioral and data-flow uncertainties that exceed what conventional cloud controls were designed to manage. The findings show that organizations recognize their primary risks—especially around data exposure and compliance—even as they continue to build the skills, tools, and governance needed for effective AI security. Extending existing privacy and security controls remains necessary but insufficient; controls-based approaches alone cannot address the non-deterministic and behavior-driven nature of AI systems. Closing this tooling gap will require adaptive reasoning–based defenses and purpose-built operational practices capable of managing AI behavior at scale.

Conclusion

This year's survey provides a comprehensive view of how organizations are both securing AI systems and adopting AI within security operations—revealing an ecosystem evolving quickly yet still solidifying its foundations. The findings show a landscape advancing rapidly but deliberately: AI adoption has moved from experimentation to enterprise scale, governance structures are taking shape, security teams are emerging as early adopters, and leadership enthusiasm remains high even as understanding of AI risks and skills gaps continues to mature.

Across the key findings, governance maturity stands out as the strongest predictor of readiness and responsible innovation. Only a minority of organizations report comprehensive AI security governance today, but where unified frameworks are in place, outcomes consistently improve—earlier experimentation, higher board awareness, greater confidence in securing AI systems, and more robust staff training. Organizations must shift from fragmented policies to a unified governance model that spans all teams involved in AI. Frameworks such as the CSA AI Controls Matrix (AICM) and Google's Secure AI Framework (SAIF) provide the structure to do so, enabling organizations to align policy with emerging regulatory obligations while carrying forward critical lessons from cloud governance. Organizations may also benefit from engaging independent security consultants to evaluate AI governance readiness and implementation—drawing on expertise from both internal and external sources.

At the same time, security teams are becoming early adopters of AI, a marked shift from previous cycles where security typically lagged behind emerging technologies. Nearly all organizations are testing or planning to test AI for detection, investigation, and response, and agentic AI is gaining real momentum. Training programs are expanding, and confidence rises sharply when governance is established—further reinforcing the connection between policy maturity and operational capability.

LLM adoption is also accelerating, consolidating around a small number of major providers while organizations pursue multi-model strategies to balance innovation, control, and risk. As AI becomes foundational infrastructure, governance and vendor risk management become essential for ensuring resilience across increasingly complex and interconnected AI environments.

Despite this progress, leadership enthusiasm continues to outpace confidence in managing AI risks. Many executives are driving AI initiatives forward, but only a portion feel fully aware of or equipped to manage the associated security implications. This gap underscores the need for sustained executive engagement, deeper alignment with security leaders, and increased investment in specialized expertise as ownership structures evolve and responsibility for AI deployment and protection remains distributed across teams.

Finally, organizations continue to face persistent skills gaps and challenges in understanding AI risks. Risk comprehension, staff knowledge, and compliance burdens remain the top barriers to secure adoption. While data exposure and regulatory concerns dominate security priorities, model-level risks—such as data poisoning, prompt injection, and model manipulation—receive less attention. Extending traditional security and privacy controls remains necessary but insufficient; the non-deterministic, behavioral nature of AI requires adaptive reasoning–based defenses and purpose-built operational practices capable of managing AI behavior at scale.

Demographics

This survey gathered insights from 300 IT and security professionals across a diverse range of organizations, spanning different industries, sizes, and geographic regions. The demographic breakdown provides important context for understanding the findings, highlighting the varied experiences and challenges faced by organizations in different sectors and operational scales.

Location

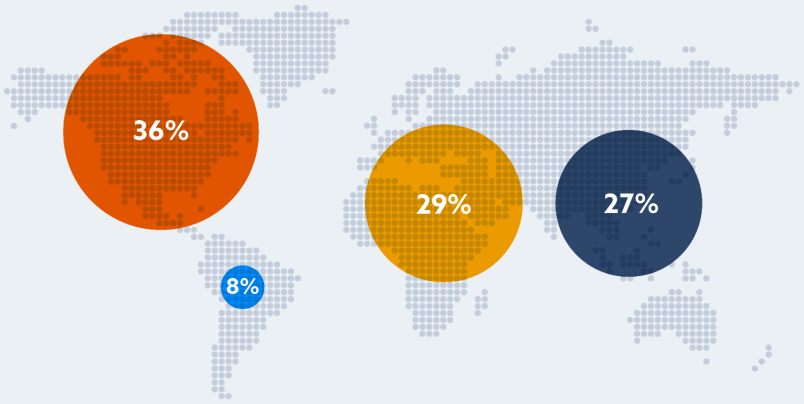

- North America
- Latin/South America
- Europe, Middle East, Africa (EMEA)
- Asia Pacific (APAC)

Organization Size

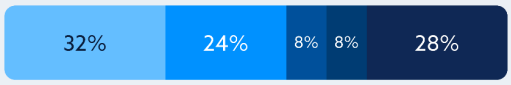

- 1–250
- 250–2,500
- 2,501–5,000
- 5,001–10,000
- 10,000+

Job Level

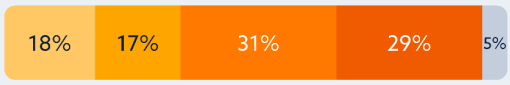

- C-level or Executive
- Director
- Manager
- Staff
- Other

Organization Industry

- **35%** Telecommunications, Technology, Internet & Electronics
- **20%** Finance & Financial Services
- **7%** Education
- **6%** Healthcare & Pharmaceuticals
- **5%** Government
- **3%** Business Support & Logistics
- **2%** Manufacturing
- **2%** Agriculture
- **2%** Retail & Consumer Durables
- **2%** Utilities, Energy, and Extraction
- **1%** Construction, Machinery, and Homes
- **1%** Entertainment & Leisure
- **1%** Insurance
- **1%** Nonprofit
- **1%** Advertising & Marketing
- **1%** Real Estate
- **0.5%** Health & Fitness
- **0.5%** Transportation & Delivery
- **6%** Prefer not to answer

Survey Methodology

The Cloud Security Alliance (CSA) is a not-for-profit organization with a mission to widely promote best practices and ensure cybersecurity in cloud computing and IT technologies. CSA also educates various stakeholders within these industries about security concerns in all other forms of computing. CSA's membership is a broad coalition of industry practitioners, corporations, and professional associations. One of CSA's primary goals is to conduct surveys that assess information security trends. These surveys provide information on organizations' current maturity, opinions, interests, and intentions regarding information security and technology.

Google commissioned CSA to develop a survey and report to better understand the industry's knowledge, attitudes, and opinions regarding AI security and governance. Google financed the project and co-developed the questionnaire with CSA research analysts. The survey was conducted online by CSA in Summer 2025 and received 300 responses from IT and security professionals from organizations of various sizes and locations. CSA's research analysts performed the data analysis and interpretation for this report.

Goals of the Study

The goal of this survey is to explore how enterprises address security challenges related to AI development, deployment, and use. It also investigates potential cybersecurity threats and vulnerabilities as observed from within the enterprise's cloud environment. Our goal is to gain a deeper understanding of:

- AI adoption trends
- AI security and governance
- Ownership and accountability in AI development, deployment, and use