GALLOS

GALLOS Pulse Survey of Alliance Security Leaders

AI Deployment in the Enterprise

December 2025

AI Deployment in the Enterprise

GALLOS Pulse Survey, July 2025

The Survey

In July 2025, the GALLOS Pulse Survey asked Alliance members about their organisations' AI deployment. The objectives were to understand:

I. The extent to which AI is being deployed in the enterprise;

II. Where AI is used – and where it is not;

III. Who is managing its deployment;

IV. What is being used to secure and govern AI; and

V. What is preventing organisations from deploying AI across the board.

This document explores the themes above, giving Alliance members a window into where their organisation sits relative to others, and where they might look for tools to enhance, secure and govern adoption.

Executive Summary

AI adoption has reached ubiquity: every GALLOS Pulse Survey respondent confirmed their organisation is actively deploying AI. What differs is the depth of adoption and the contexts in which AI is most applied. Here is a brief summary of our findings:

- **Efficiency-focused implementation strategy** – Enterprises are prioritising workflow optimisation and productivity enhancement over workforce replacement, maintaining a "human-in-the-loop" approach that emphasises pragmatic, incremental improvements
- **Security and governance as primary barriers** – Model governance, application oversight, and data security remain the most significant obstacles to scaling AI initiatives, with business value often unclear to stakeholders
- **Strong demand for third-party security tools** – Security leaders plan widespread deployment of external AI TRiSM (trust, risk, security and management) solutions across all categories, including usage monitoring, observability, red teaming, and auditability, reflecting the inseparable link between AI adoption and risk management
- **Fragmented ownership structure** – AI innovation responsibility is distributed across CIO offices, decentralised business units, Heads of Data, and emerging centralised AI teams, creating significant challenges for AI TRiSM vendors in identifying primary enterprise buyers
- **Tool proliferation creating new challenges** – The rapid emergence of specialised AI governance and security products is generating procurement bottlenecks and uncertainty about whether the market will consolidate into integrated platforms or remain fragmented with point solutions

Survey results are based on the responses of 22 GALLOS Alliance members.

AI Deployment in the Enterprise

GALLOS Pulse Survey, July 2025

What are the primary cybersecurity challenges you are facing over the next 6-12 months?

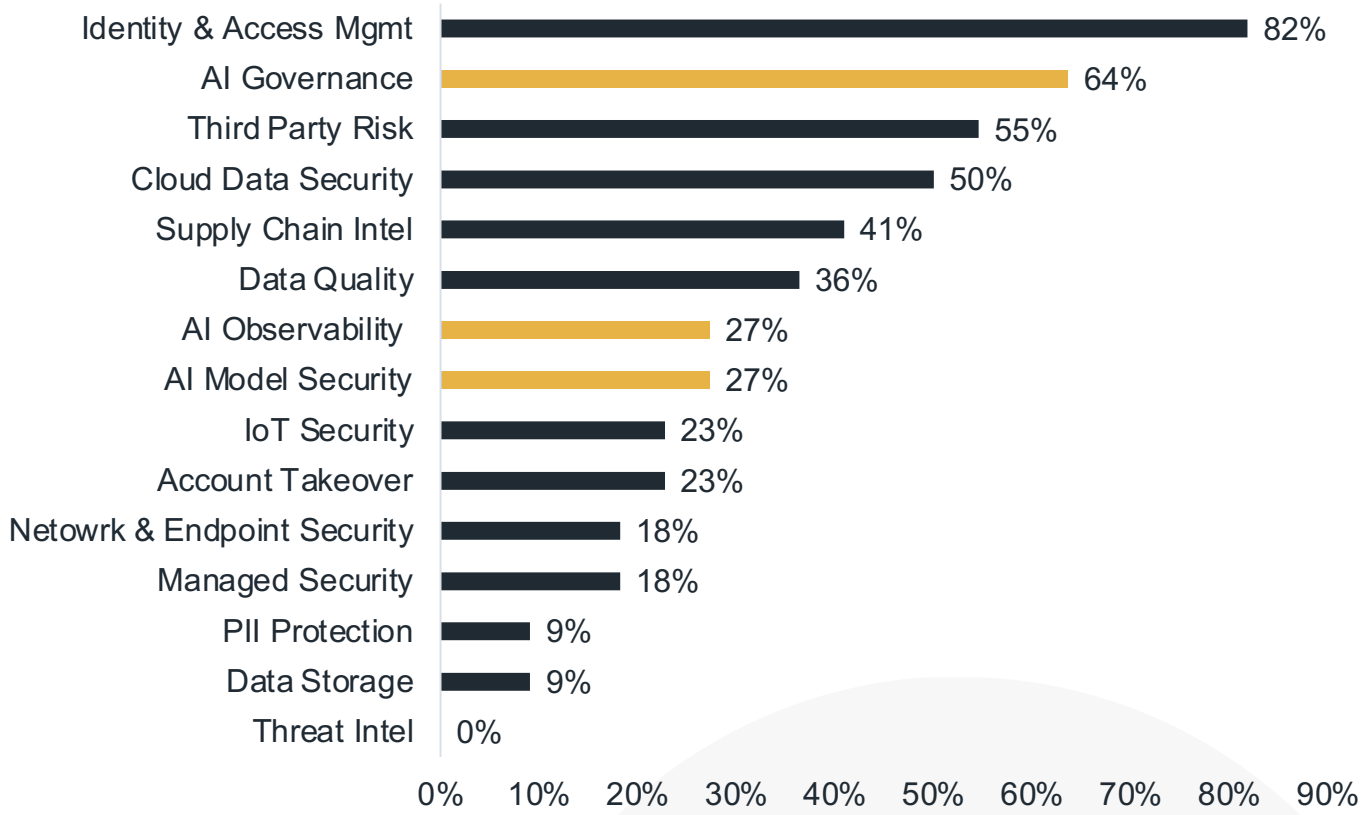

Whilst AI-related priorities are prominent, they do not yet dominate the top tier of security leaders' agendas. Categories such as employee AI usage monitoring (86%), model and data observability (73%), and AI vulnerability testing (68%) are being pursued at high levels, but they still trail the very highest priorities typically seen in broader cyber programs, such as identity security, cloud posture, and ransomware resilience. AI risk is recognised as critical, but is still being absorbed into established frameworks rather than displacing legacy concerns.

100% of respondents believe AI will shape their cyber strategy within the next one to three years, highlighting an inflection point. Today, AI-specific controls are treated as adjacent to core security functions; within a short horizon, they are expected to become integral. The current emphasis on monitoring, observability, and red teaming for AI is likely just the starting phase. Over time, the weight of these AI-specific categories will rise to match the most entrenched cyber priorities, particularly as regulatory scrutiny and enterprise dependency on AI systems increase.

Security leaders are currently aligning AI governance with existing priorities rather than elevating it above them. But the universal consensus on AI's strategic impact indicates a trajectory where AI-related security and governance controls are poised to become co-equal with the highest current priorities, shifting from "emerging concern" to "core mandate" in the near term.

AI Deployment in the Enterprise

GALLOS Pulse Survey, July 2025

Over half of organisations have adopted AI at scale

To what extent is your organisation currently deploying AI?

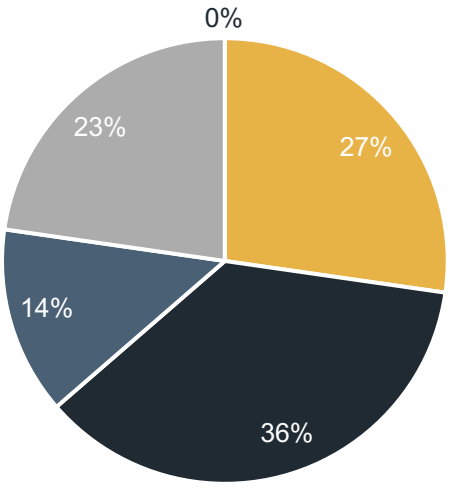

36%	In many use cases and departments
27%	In major initiatives/aggressively
23%	Cautiously/limited in scope
14%	Moderately
0%	Not at all

AI has crossed the threshold from experimental pilots to mainstream enterprise adoption. More than 60% of organisations deploying AI either across multiple departments or through aggressive, large-scale initiatives; it is clear that AI is no longer being treated as a fringe innovation but as a foundational capability. The fact that only a minority are deploying AI “moderately” highlights a shift in competitive dynamics: opting out of AI is becoming increasingly rare, and a marker of falling behind.

Most organisations are applying AI to enhance efficiency and workflows, reflecting a pragmatic focus on augmentation rather than disruption. Ownership remains fragmented between CIO offices and decentralised business units, mirroring early cloud adoption patterns where governance trailed enthusiasm. At the same time, the surge in AI TRiSM tooling adoption shows that enterprises are expanding use while still building trust and assurance mechanisms.

Overall, AI has become embedded but not yet fully institutionalised—scaled in practice, still maturing in governance and clarity of value.

77%

of respondents’ organisations are at least moderately deploying AI, confirming that the enterprise has moved beyond experimental use

100%

of respondents’ organisations are deploying AI to some extent

AI Deployment in the Enterprise

GALLOS Pulse Survey, July 2025

Business impact and governance are hurting adoption

What are the primary roadblocks, if any, holding up AI deployment in your organisation?

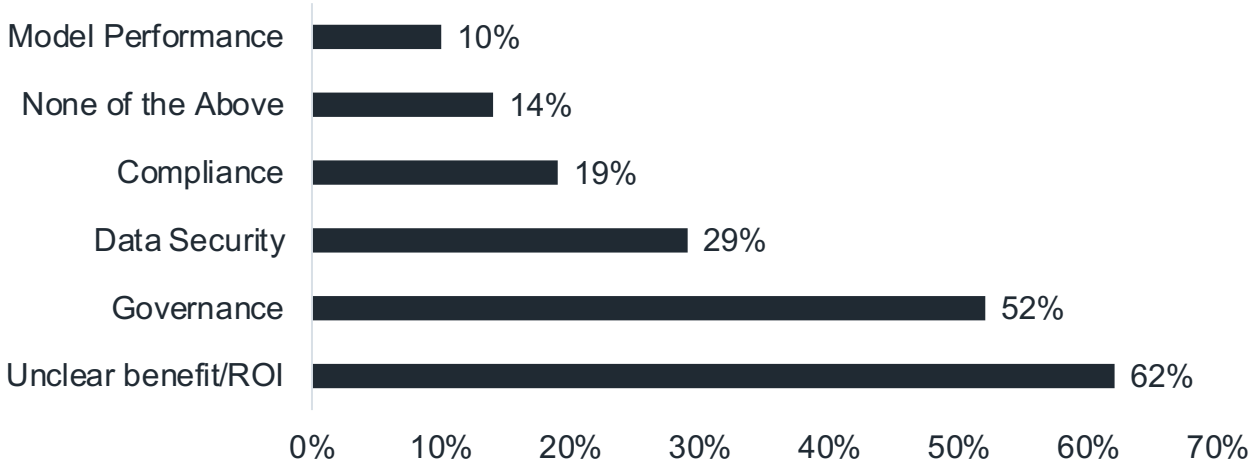

The greatest barriers to AI deployment are no longer technical, but strategic and organizational. The top two roadblocks – unclear benefit or ROI (62%) and governance (52%) – indicate that enterprises are struggling less with whether AI can work, and more with how to do so responsibly. This marks a turning point in AI maturity: performance and feasibility are largely solved problems, but aligning AI with measurable business outcomes and trustworthy governance remains elusive.

The prominence of unclear ROI underscores a pervasive challenge in the enterprise AI landscape: many projects demonstrate technical success without clear economic justification. This often stems from scattered experimentation, lack of baseline metrics, and the difficulty of quantifying the value of efficiency or augmentation compared to direct revenue gains. As a result, business leaders hesitate to scale AI programs without defensible value narratives – particularly in regulated or cost-constrained environments.

Meanwhile, governance (52%) and data security (29%) reveal that trust and control continue to be major constraints. As AI systems become embedded in decision-making and customer-facing functions, questions around accountability, explainability, and policy compliance grow sharper. Compliance (19%), though lower, still reflects the uncertainty surrounding evolving global regulatory frameworks, from the EU AI Act to sector-specific guidance.

Only 10% cite model performance as a primary obstacle, suggesting that technical capability is no longer the limiting factor. The challenge has shifted from building working models to building models that the organisation – and its regulators, customers, and boards – can trust. In essence, the data paints a picture of enterprises moving from an era of experimentation to one of justification and governance: AI can deliver, but proving and managing that delivery has become the new frontier.

AI Deployment in the Enterprise

GALLOS Pulse Survey, July 2025

Employee efficiency and workflow enhancements are the most common uses

Where is your organisation most active in applying AI?

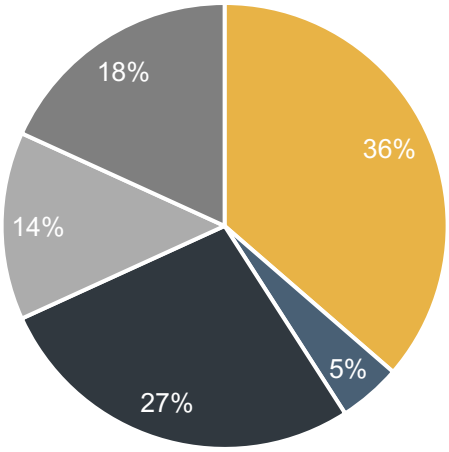

36%	Employee efficiency
27%	Enhancing existing workflows
14%	Product improvements
5%	Employee replacement
18%	Other

Employee efficiency and workflow enhancements as the most common use cases reflect a pragmatic approach to value creation. Rather than prioritising disruptive applications such as employee replacement or radical product transformation, enterprises are channelling AI toward augmenting existing processes. This implies two things: first, organisations view AI less as a tool for wholesale reinvention and more as a lever for productivity in knowledge and operational work; second, the cultural and organisational barriers to using AI for workforce substitution remain strong, even in an era of automation obsession.

Another implication is that enterprises are treating AI as a horizontal enabler, rather than a vertical point solution. By embedding AI into workflows, firms are effectively integrating it into their operating fabric. However, the relatively lower emphasis on product improvement suggests that AI's role in shaping competitive differentiation is still underdeveloped. For now, enterprises appear to be pursuing low-risk, incremental gains, while the more transformative potential of AI in reshaping business models remains a future horizon.

5%

of respondents' organisations are actively replacing employees with AI – "human-in-the-loop" prevails for now

14%

of respondents' are most actively using AI for product improvements – surprisingly low in the era of AI coding tools

AI Deployment in the Enterprise

GALLOS Pulse Survey, July 2025

Security leaders expect to use tooling across AI TRiSM

Where are you using or considering using third party tools for AI governance and security?

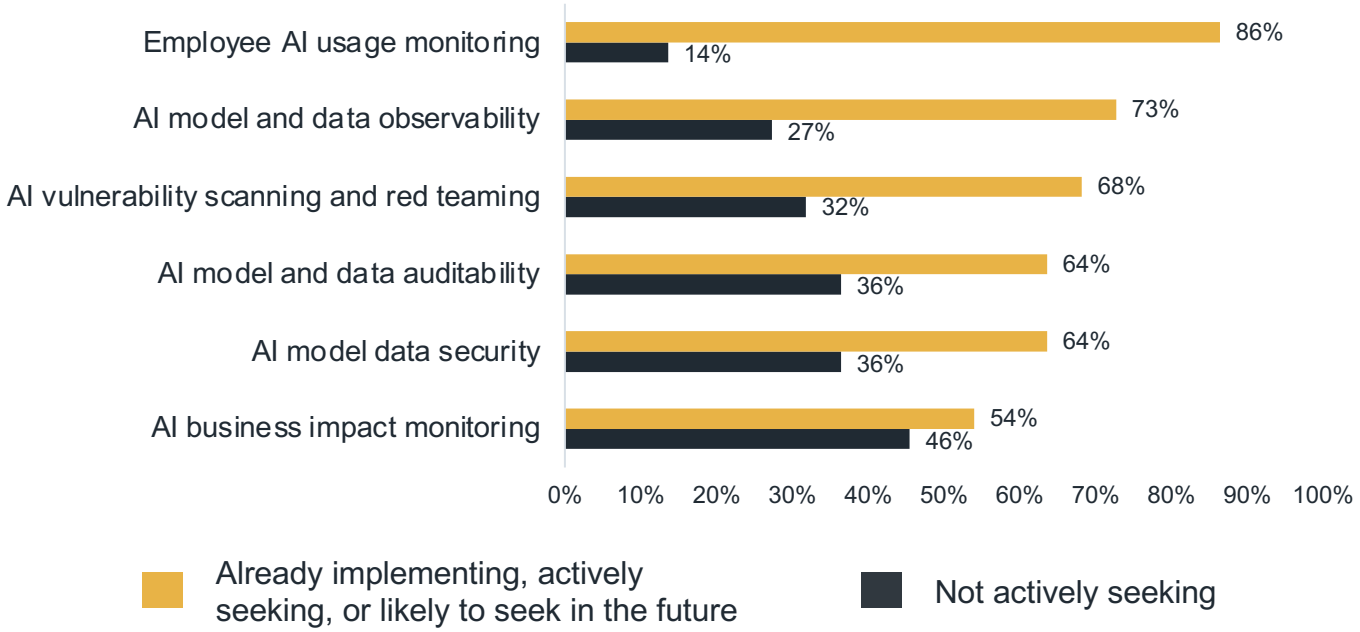

Security leaders are not approaching AI as an uncontrolled innovation, but as a domain that requires the same rigour and tooling as any other critical enterprise system. The emphasis on employee AI usage monitoring (86%) signals that organisations are acutely aware of the human layer of risk: shadow use of generative AI, data leakage through prompts, and policy non-compliance are seen as the most immediate threats. This reflects a shift in security posture, where visibility into how staff interact with AI is treated as foundational.

The strong demand for AI model and data observability (73%) and vulnerability scanning/red teaming (68%) suggests that enterprises are beginning to adapt familiar DevSecOps practices to the AI context. Much as continuous monitoring and penetration testing became table stakes for cloud adoption, similar mechanisms are now emerging as hygiene requirements for AI. Notably, interest in auditability and model security (64%) indicates that leaders are anticipating external pressures – from regulators, auditors, and customers – who will expect traceability and assurance around AI decisions.

The lower (surprisingly, given respondents' earlier emphasis on unclear ROI), but still significant, interest in business impact monitoring (54%) implies that many organisations remain focused on controlling risk inputs (usage, data, vulnerabilities) before turning their attention to measuring downstream business consequences. This may point to a maturity gap: AI risk is being addressed tactically in the near term, while systematic integration into enterprise risk management frameworks is still in development.

Overall, the picture is one of convergence: enterprises are treating AI security not as a novel category, but as an extension of established governance and assurance patterns. The prioritisation of monitoring and observability over replacement or transformation mirrors the broader adoption trend – AI is being operationalised incrementally, with security leaders ensuring the guardrails are in place before innovation accelerates further.

AI Deployment in the Enterprise

GALLOS Pulse Survey, July 2025

Innovation and deployment responsibility is fragmented

Which group within your organisation is managing AI innovation?

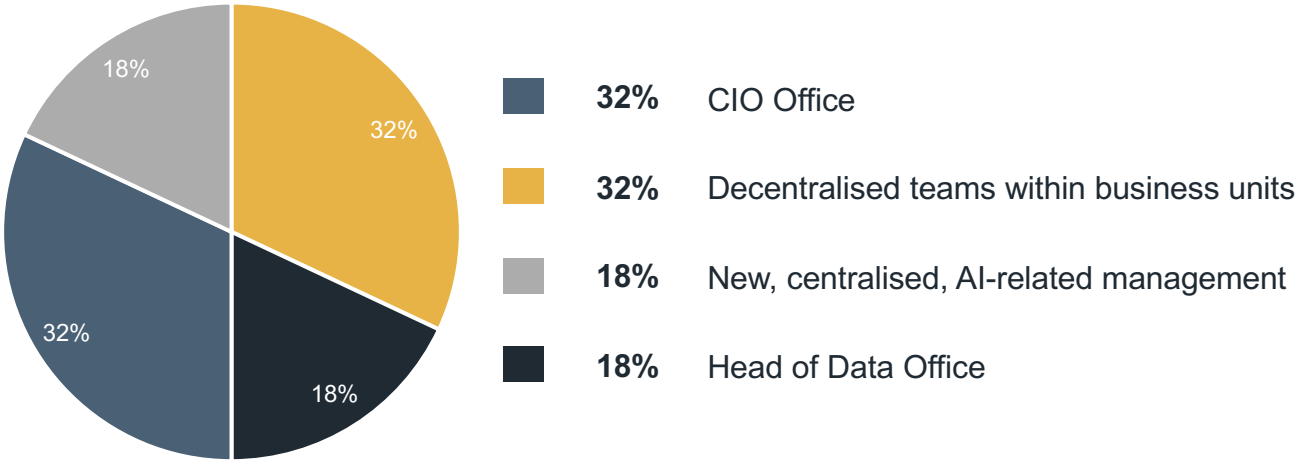

The distribution of responsibility for AI innovation reveals a fragmentation in enterprise governance. With CIO offices (32%) and decentralised business-unit teams (32%) equally likely to be leading AI efforts, organisations are pursuing two parallel models: centralised oversight for strategic alignment and distributed ownership for speed and domain-specific innovation. This duality reflects both the cross-cutting nature of AI and the uncertainty over where ultimate accountability should sit.

The emergence of new, centralised AI-related management groups (18%) indicates that some enterprises are beginning to formalise AI as a first-class discipline, on par with security or data governance. The fact that this number is still relatively low suggests the field is in transition – most companies are experimenting with structures rather than locking into a permanent model. Similarly, the Head of Data offices (18%) playing a leadership role underscores the strong data foundation required for AI, but also signals that data teams alone cannot absorb the full strategic and operational complexity of enterprise AI.

AI ownership is still fluid, lacking the institutional clarity that cloud adoption or cybersecurity eventually acquired. The equal weighting between centralised (CIO or AI-specific units) and decentralised (business units, data leaders) approaches implies a tension: enterprises are trying to capture innovation at the edges while avoiding fragmentation of governance and risk oversight. Over time, this balance is likely to resolve into more formalised AI centres of excellence or enterprise-wide AI governance frameworks, but for now the picture is one of experimentation – both in technology and in management models.

GALLOS

Survey designed and administrated by Jack Pearson, Principal, Gallos Ventures

Contact: ventures@gallostech.io