STIFEL | VENTURE BANKING

SINET

AI Deployment and Security in the Enterprise

A 2025 Survey of Risk Executives

February 2026 • Report

Inside...

A View from the Field

How Risk Leaders Are Governing AI at Scale

AI adoption is no longer confined to experimentation. Across a 2025 survey of Risk Executives conducted by SINET in collaboration with Stifel Bank, GenAI is now universal with first-party Agentic AI rising in prominence and increasingly embedded in production environments. As deployments expand, security and governance have become the gating factor for scale. This report benchmarks how organizations are managing employee AI use, implementing governance and observability tooling, maturing red teaming practices, and prioritizing near-term security investments.

Methodology (at a glance): Based on a 2025 survey of 102 Risk Executives conducted by SINET in collaboration with Stifel Bank, supplemented by follow-up qualitative interviews.

Five Themes Shaping Enterprise AI in 2026

1. AI is now a production reality.

Nearly all respondents report using AI in their organizations today, with adoption spanning industries and functions. Risk executives are increasingly focused on safe expansion into Agentic AI, and leading-edge organizations report the highest concentration of production deployments.

2. Perceived maturity varies widely and correlates with governance investment.

While a subset of respondents identify as “Leading Edge,” most place themselves as Middle of the Pack or Behind. Leaders start with Governance and fit AI into existing frameworks and strategies. Responsible adoption combines (i) well-articulated policy (policy as code), (ii) administrative controls (GRC/shiftleft), and (iii) technical controls.

3. Risk executives can enable the business and experimentation.

The security organization should use this moment to positively impact the business and educate other departments around responsible adoption. Organizations are increasingly addressing employee AI use through approved enterprise platforms, proxy and gateway controls, and self-hosted environments rather than relying on blocking alone. The leaders allow all employees (not just engineers) to experiment quickly with increasing levels of friction and quality assurance as you move toward production. This shifts the conversation to use cases and ROI, rather than simply "we need AI." We heard that more than 95% of proposed pilots drop off once teams move from ideas to building, especially for business-led projects using low-code and no-code tools.

4. Despite enthusiasm for software development and creator roles, skepticism around AI ROI remains for many.

Concerns around ROI and cost are slowing AI adoption outside leading-edge organizations, despite top down pressure to act. Many respondents cite developer productivity as a key use case, but an equal number worry about a generational gap as junior engineers lean on Assistant Agents and miss foundational learning. Others flag long-term economics, noting costs may rise as dependency grows and providers subsidize usage with low gross margins. LLM-based assistants are now widespread, but those who identify as Behind remain skeptical that GenAI alone will deliver the next productivity boom. Leaders are increasingly focused on Tasker Agents and Orchestration Agents to unlock that value.

5. Red teaming remains a lagging capability, even as risk escalates with Agentic AI.

Organizations increasingly recognize the need for red teaming, but consistent practices are not yet standard outside of the leaders. As agentic use expands into workflows that make changes to systems of record rather than simply synthesizing information, red teaming maturity will likely become a decisive differentiator for secure scaling. Intuitively, you cannot practice red teaming without real use cases. We heard many cases where Middle of the Pack organizations evaluate red teaming tools, only to pause procurement until more meaningful use cases emerge. The need for red teaming is proportional to the risk and scope of the use cases. If you’re deploying an LLM on an internal knowledge base, the need is considerably less than having Tasker Agents and Orchestration Agents making real business decisions and changing systems of record.

Enterprise AI is Broad and Accelerating

AI is now near universal among surveyed organizations. Respondents report active AI use alongside continued investment in machine learning and NLP, with growing uptake of hosted AI services and early momentum in Agentic AI.

What We're Seeing

- AI adoption is no longer niche. It is embedded across enterprise environments.
- Adoption spans industries, though leaders were largely concentrated in software. These businesses face an existential threat if their product and user experience falls behind competitors, which requires them to push the boundaries in development.
- Leaders are already recognizing ROI from Agentic AI use cases, and we 100% believe it is enduring. Almost all our interviewees are building their AI Security policy to support agentic use cases in the future.

AI Modalities in Use Today

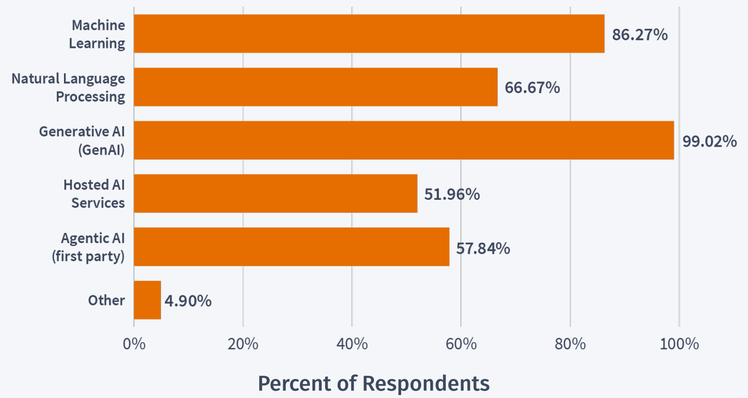

Nearly all respondents are using GenAI, alongside strong adoption of ML and NLP and growing use of agentic and hosted AI services.

Survey Respondents by Industry

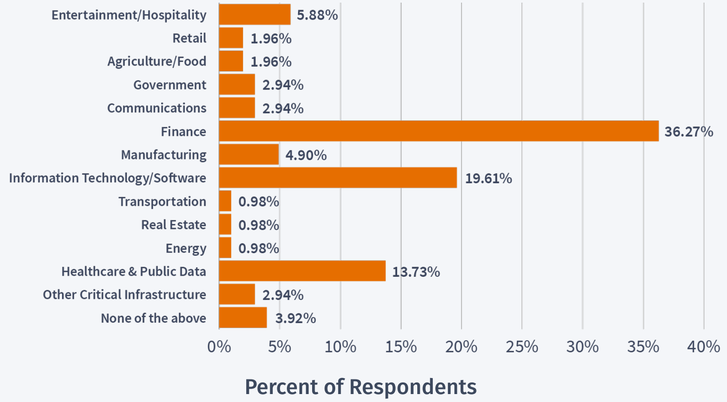

Finance and IT/software lead the sample, with substantial representation from entertainment and healthcare.

Most Organizations See Themselves as “Middle of the Pack”

Organizations differ significantly in how they perceive their AI maturity relative to peers, and the majority do not identify as Leading Edge. For most, AI maturity remains a work in progress.

How Organizations Perceive Their AI Maturity Relative to Peers

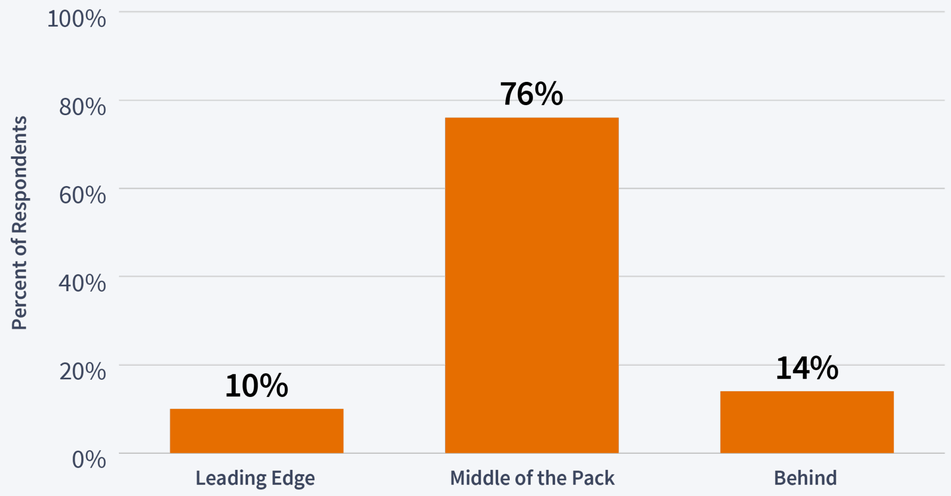

Middle of the Pack combines respondents who identified as "slightly ahead" or "about average."

Why This Matters

Self reported maturity is more than perception. It correlates with investment in governance, guardrails, and red teaming. This self perception shows up in practice: most organizations remain concentrated in lower risk AI deployments, with higher impact use cases gated by governance readiness.

> We’ve got internal-facing chatbots and intake processes, but putting a common framework in place… has its own challenges.
>
> — Risk Executive at global services firm

Live Use Cases Skew Toward Lower Risk Categories

While AI is broadly adopted, many deployments remain concentrated in lower-risk areas such as productivity, research, and internal knowledge workflows. More risk-intensive categories are expanding but remain less consistently deployed.

AI Use Cases in Use or Planned Across Organizations

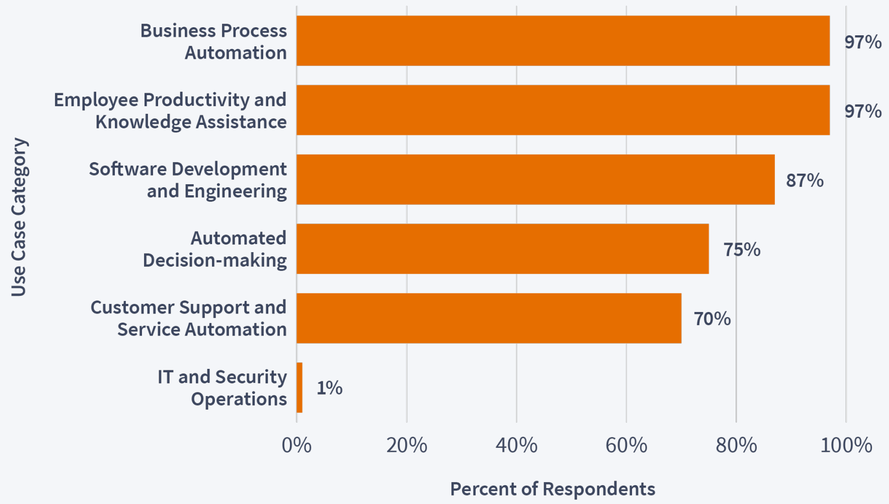

Respondents selected all that apply. Percentages reflect organizations reporting each use case as live or planned.

Automated decision-making includes traditional machine learning use cases such as approving or rejecting transactions, fraud detection, refunds, and trading.

Most organizations are deploying AI first in workforce productivity and internal-facing workflows. Higher risk categories (customer-facing applications, agentic automation tied to systems of record) show expanding interest but more measured deployment.

This reflects the broader governance pattern: organizations are moving forward, but often within boundaries of controllable risk.

> You don't find multi-agent, client-facing deployments as heavily in financial services… because of the risk of it going wrong and causing client harm or consumer damage. That's a big bright line for us. So a lot of the focus is internal optimization of processes—internal employee productivity things, saving time from drudgery.
>
> — Risk Executive at U.S. financial institution

AI Security and Governance Investment Scales With Maturity

Organizations that identify as Leading Edge consistently invest across guardrails, governance tooling, and red teaming. Middle of the Pack and Behind organizations show more fragmented adoption, reinforcing that maturity is driven less by access to AI models and more by investment in control infrastructure.

What We're Seeing

Maturity shows up in control investment. The largest gap between Leading Edge and other organizations is in governance and observability.

Guardrails, Red Teaming, and Governance Tooling by Maturity Group

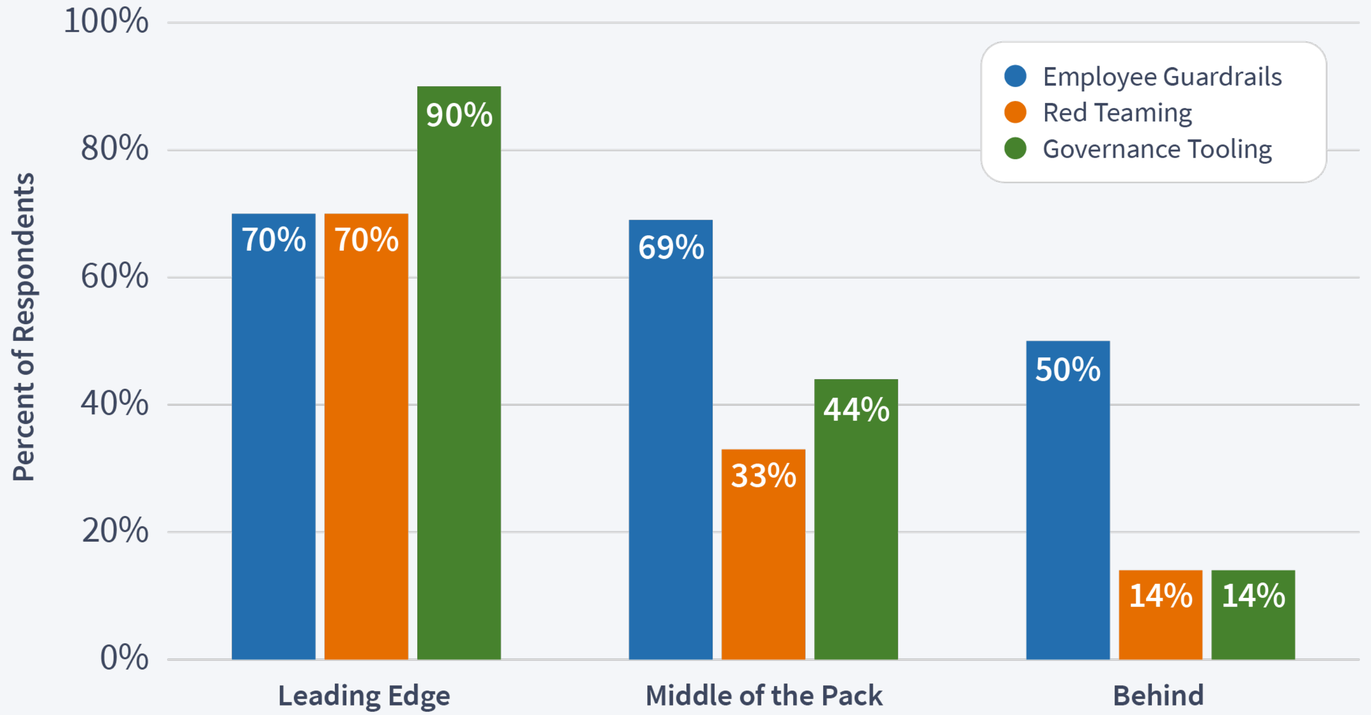

Guardrails Captured Initial Spend

Organizations are increasingly managing employee AI usage through enterprise-approved platforms, proxy and gateway controls, and internal deployment environments. Blocking public tools is less common than enabling controlled access.

What We're Seeing

- Enterprise platforms are the dominant baseline approach (Microsoft CoPilot, Gemini, ChatGPT Enterprise, etc.).
- Structural controls (proxy/gateway, self-hosting) are meaningful and increasingly common.
- Blocking alone is less common, suggesting a shift toward enabling use under guardrails. Anecdotally, most practitioners we've spoken to have some mechanism to ensure DeepSeek is not running wild across their organizations.

Approaches to Managing Employee GenAI Use

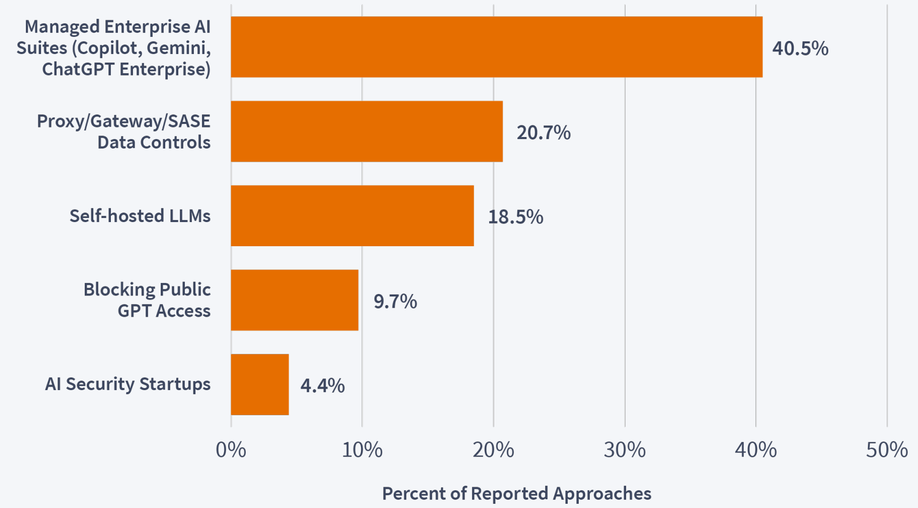

Respondents selected all that apply. Percentages reflect the share of respondents reporting each approach and do not sum to 100%.

> We do not prohibit our users from using generative AI bots. We want them to use it because it improves productivity. At the same time, we put guardrails in place through policy and tooling… The goal is not lockdown—it is governed use aligned to company policy.
>
> — Risk Executive at global logistics company

Governance and Observability Tooling Remains Uneven

Most teams are still early on governance and observability. Production deployment is concentrated among those who identify as Leading Edge.

As AI expands, governance and observability tooling are emerging as a practical requirement. Organizations increasingly need to answer:

- What is running in our environment, and what data can it access?
- How are these models configured?
- How are we compliant with our AI framework requirements?
- Who has access?
- How do we ensure agents are not overprovisioned, and what actions can they take?
- How are we testing these controls?
- Can we reconstruct an end-to-end AI driven workflow after the fact?
- Do we retain audit ready evidence of which agents, models, and systems were involved?

Status of AI Governance and Observability Tooling

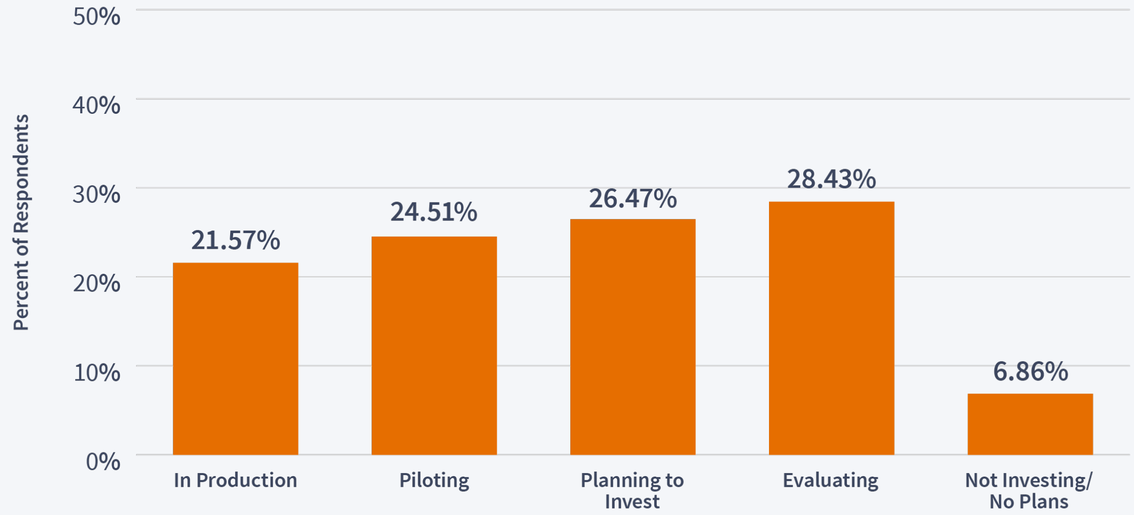

> As we start relying on these bots, algorithms, and agents, that's where observability is extremely important in governance, and being able to see what agents are running, where they're running, and what data they're hitting.
>
> — Risk Executive at global asset manager

Red Teaming Follows Production Use Cases

Red teaming is widely recognized as necessary, but approaches vary significantly across organizations. Many respondents report early-stage or ad hoc red teaming practices. We believe there will be widespread adoption as we use higher risk agents (Tasker and Orchestration Agents), but it is critical for many internal AI use cases.

What We're Seeing

- Red teaming maturity is lagging governance intent.
- Even organizations with live use cases often report inconsistent or early-stage practices.
- As customer-facing and agentic applications expand, red teaming will need to become continuous, repeatable, and integrated into deployment pipelines.
- While it would be unfortunate if an internal LLM exposed sensitive information such as executive compensation, the consequences are much less severe than an Orchestration Agent hallucinating and deleting everything in your GitHub repositories.

Current Approach to Red Teaming for GenAI

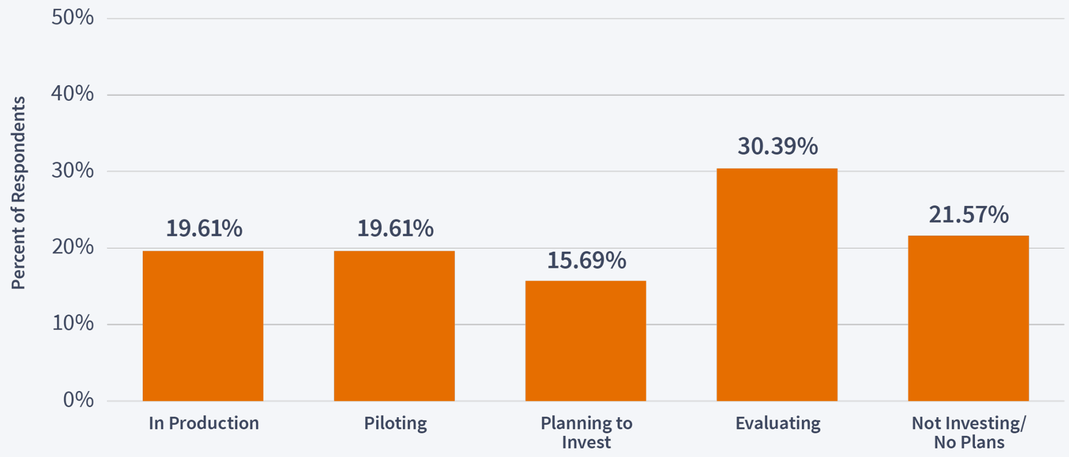

> We were working on a chatbot for our homepage... and discovered it was providing medical advice, which is not acceptable. That was something we uncovered through red teaming. We had to implement guardrails so that if someone asked a medical question, the chatbot would not attempt to triage or diagnose. Instead, it would limit responses to information about cell phone plans and TV channels.
>
> — Risk Executive at global communications provider

Security Investment Priorities Signal Where Scale Will Be Won or Lost

Respondents' planned security investments reinforce that AI governance is now a top security priority, alongside identity and SOC modernization.

Why This Matters

AI will not scale without stronger identity controls and data safeguards. These investments reflect the security foundation required to move beyond bounded use cases into fully operational deployment. There's a widespread belief (which we share) that the SOC is the perfect place to leverage AI to make the security team more efficient. As adversaries embrace AI tooling to accelerate the volume of attacks, defenders need to fight fire with fire to keep pace.

Top New Security Budget Priorities

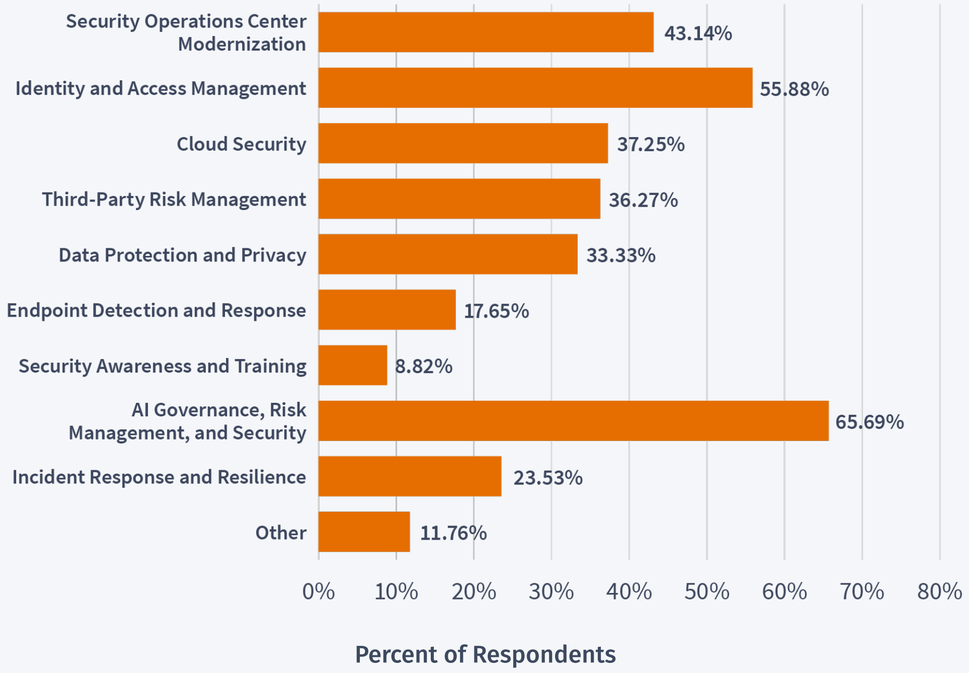

AI governance leads new security investment plans, with identity, cloud, and SOC modernization close behind.

Buyers Are Paying for Promise While Tooling Matures

Even organizations with budget, urgency, and mature security programs report that the AI governance tooling landscape is still catching up. Leaders largely built their tools before the wave of emerging startups came to market. While AI is already in production across nearly all respondents, only a minority report observability and governance tooling in pilot or production, creating a widening execution gap between adoption and oversight.

> If a tool can meet 85% of the requirement, we will buy it. If it cannot, we build. In this space, a lot of solutions still do not meet the bar, so you end up filling the gaps yourself.
>
> — CISO at a global technology company

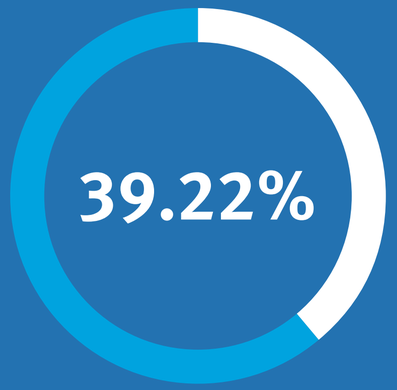

Only **39.22%** of respondents report that red teaming for customer-facing AI is currently **in production or piloting**, underscoring how far testing and validation practices still have to mature.

Implication

Platforms that integrate governance, observability, and red teaming will increasingly displace point solutions focused on a single risk layer. We believe enduring companies will be built around enabling and securing Agentic AI with early use cases for GenAI becoming commoditized and offered by legacy platforms.

From Insight to Execution: Enterprise AI in Practice

As AI moves from experimentation to execution, the limiting factors are increasingly organizational rather than technical—governance, ownership, operational readiness, and the controls required to scale safely. Survey results show broad adoption, but uneven maturity in the mechanisms that keep deployment accountable and observable.

To complement the quantitative findings, we conducted in-depth interviews with risk executives across industries and maturity levels. These perspectives illustrate how organizations translate intent into practice—pairing administrative (GRC/shift-left) controls that define accountability with technical (AI security) controls that enforce guardrails in real systems.

Featured perspectives include leaders from:

- A global communications provider operating AI at scale
- A regulated financial institution architecting governance-first AI programs
- A mature enterprise buyer navigating market and vendor constraints
- A pragmatic enterprise scaler balancing adoption velocity with operational readiness
- A governance-first organization prioritizing ownership, controls, and oversight early

The perspectives that follow are not presented as prescriptive models but as grounded examples of how enterprise AI programs are being shaped in the field today.

THE TOP 1% OPERATOR

Governance as Infrastructure, Not Policy

Global communications provider

Risk Executive

This organization represents the most advanced end of the maturity spectrum, operating at a scale where AI adoption cannot be managed through policy and employee awareness alone. Rather than allowing teams to access public tools ad hoc, leadership prioritized building a controlled access layer that could meet employee needs while enforcing security and privacy requirements by default.

The Risk Executive describes this as a deliberate effort to avoid the inevitable “shadow AI” pattern, where business units independently adopt public tools and governance becomes a retroactive cleanup exercise. Instead, the organization built an internal LLM gateway that routes employee usage through a governed environment, enforcing corporate guardrails and preventing sensitive data inputs.

Once governed access was in place, adoption accelerated quickly. What began as dozens of internal chatbots expanded into hundreds and then thousands of embedded AI-enabled use cases across operations, while governance remained centralized. In parallel, the organization formalized cross-functional testing, including security, privacy, legal, IT, and customer stakeholders, to identify failure modes early and build durable guardrails for customer-facing deployments.

> We built an LLM proxy. Everything routes through our own environment with guardrails and corporate guidelines. We didn’t want shadow AI, so we gave people an internal alternative that actually works, while preventing customer or PII data from being entered.
>
> Governance cannot be a policy document. If you want AI to scale, controls have to be built into the system so people default into a safe environment. And as you move into agents, the risks become clearer: overly permissive access to private data, exposure to untrusted content, and the ability to exfiltrate data. If all three exist, you can get disastrous outcomes, which is why guardrails must be in place from the beginning.

KEY TAKEAWAY

Governance maturity is not a brake on adoption. When guardrails are embedded into the access layer and reinforced through continuous testing, AI can scale safely across the enterprise.

THE TOP 10% ARCHITECT

Scaling AI in Regulated Environments

Mid-sized financial institution

Risk Executive

This security leader describes an adoption posture that is neither experimental nor “open exploration.” Instead, the organization is pursuing a structured internal maturity curve that aligns with risk management expectations and regulatory oversight. Most peers, he notes, are still concentrated in an “enterprise productivity” phase, with early wins tied to embedded platforms such as Microsoft Copilot and adjacent automation capabilities already present in the enterprise stack.

He emphasizes that the gating factor is not technical capability, but the ability to scale AI using the same core disciplines that govern cloud and other emerging technologies: policy, access control, change management, and auditable oversight. AI may introduce novel risk categories, but the control model should remain familiar. The priority, he argues, is to build a governance foundation that can withstand scrutiny before moving into client-facing or high-impact applications.

> “AI is a new layer of technology, but the control disciplines are familiar. Access management, change management, configuration management. Those requirements don’t go away. Copilot-style productivity use cases are relatively low risk. But once you move into agents that connect to systems of record and take action, expectations scale quickly. At that point, you need controls in place and the ability to explain the workflow clearly to an auditor or regulator. If security leadership is not pushing for that structure, the organization will build at velocity and governance will fall behind. And when the exam or audit comes, the organization will be on the back foot.

KEY TAKEAWAY

AI introduces new capabilities, but not new governance fundamentals. As organizations move from productivity copilots to agents that take action in systems of record, the requirements for access controls, change management, and audit-ready oversight increase sharply. The risk is not adoption. It is governance falling behind.

THE MATURE BUYER

Paying for Promise in an Immature Market

Large U.S. financial institution

Risk Executive

A mature enterprise buyer with an active pipeline of AI use cases across business lines is already investing time and budget into governance and security tooling. Yet across proof of concept evaluations, a consistent market maturity gap shows up. Many solutions demonstrate strength in individual domains, but few deliver reliable end to end coverage at the breadth and depth required for a scaled enterprise environment.

In practice, this creates an uncomfortable tradeoff. Organizations can wait for vendors to mature, or they can move forward with partial coverage and compensate with manual processes and internal controls. Compounding the challenge, many vendors are positioning enterprise buyers as “design partners” who help productize capabilities in real-world environments, but with pricing that already assumes maturity. For sophisticated buyers, the question is not whether AI security tooling is needed. It is whether the market can deliver what enterprises are being asked to pay for today.

> The POCs we have run show spotty coverage at best. We have use cases across business lines, and there is no single tool that supports them end-to-end, so gaps will remain. Either we wait for vendors to mature, or we operate with a blend of vendor capability and manual support.
>
> What we are seeing often is vendors asking customers to help mature their products through design partnerships. Design partnerships are not uncommon, but it is a design partnership with a sticker price that already reflects AI. In effect, these are immature products at very high price points.

KEY TAKEAWAY

Even mature security organizations with budget and readiness to deploy are constrained by a tooling landscape that has not yet caught up to enterprise requirements. Until vendor solutions provide consistent coverage across use cases, scaling AI will require a hybrid model of partial tooling plus internal controls.

THE PRAGMATIC SCALER

Operationalizing AI Without Overengineering

Global asset manager

Risk Executive

This organization represents a pragmatic adoption posture shaped by an external catalyst: a recent acquisition by a much larger firm with more mature AI capabilities. Historically, the organization had been a laggard in AI adoption. The acquisition changed the urgency. As integration accelerated, the business push for AI moved from “nice to have” to time-sensitive, and security leadership shifted from reactive evaluation to building repeatable governance mechanisms.

Rather than overengineering a complex program, the team focused on practical controls that can scale: an AI council, a standardized vendor intake process, and training that establishes a clear acceptable-use baseline for employees. The interviewee describes governance as a process of disciplined enablement. It is designed to reduce the most common risk scenarios and create a shared decision model that brings IT, security, legal, and compliance into the same operating rhythm. For use cases already delivering value, the focus is not on blocking. It is on creating guardrails and accountability that keep adoption defensible under scrutiny.

> “AI is different because the question is how you limit the blast radius. It comes down to training and end users knowing what they can and cannot do. For us, the business came to security early and said we will fall behind if we do not move, so we built a more regimented process. We use about 15 vetting questions and review tools across security, IT, and legal. Often we go back to vendors with follow-up questions before we greenlight anything.
>
> We also rely heavily on third parties. If we become over-dependent on one vendor and they go down, the business is exposed. The other issue is complexity. If, in five years, we have 2,000 agents in the environment, are they creating inefficiencies or acting outside our visibility? Identity also becomes critical because those agents will have privileges, some of which may be overentitled. At the end of the day, can we prove we have a firm grip on our AI tools and privileges, with maker-checker controls and real human oversight?

KEY TAKEAWAY

Pragmatic organizations scale AI by building repeatable governance mechanisms that enable the business while reducing avoidable risk. Standardized intake, shared decision-making, and employee training create a defensible foundation for adoption, especially when external catalysts increase speed and demand.

THE GOVERNANCE-FIRST ENTERPRISE

Enablement Requires Friction

Global asset manager

Risk Executive

This organization has taken a deliberate governance-first approach to AI adoption, resisting the common “experiment first, govern later” pattern. The executive describes widespread enthusiasm across the market, but views much of it as aspirational. In contrast, the firm’s approach centers on practical outcomes, clear decision rights, and a lightweight but structured governance process that enables adoption while applying the right level of friction.

Governance ownership sits outside the security function, led by a data science team supported by a cross-functional governance group that includes legal, compliance, audit, business teams, and security. Rather than trying to block AI use broadly, the organization prioritizes licensing and approved tool selection, requiring teams to register AI ideas so they can be screened for legal, privacy, contract, and data agency risks. The model is rooted in culture and accountability. The executive emphasizes that no technical control can anticipate every use case, so the organization focuses on training risk-aware teams, monitoring usage, and enforcing integrity when people bypass agreed processes.

Looking ahead, the organization sees Agentic AI as an immediate governance challenge. Its approach is to engineer the core agent ecosystem in advance, minimize ad hoc adoption, and create an onboarding process similar to API governance, where security becomes the outcome of strong architecture and disciplined engineering.

> “You can have a million policies and a million tools, but you will never outline every use case people will encounter. You need educated, risk aware people, and you need accountability. We license specific models and AI enabled software, and we tell people these are the tools you are required to use.
>
> If someone thinks there is a better tool, we expect them to let us know so we can review it. We encourage that. And we do monitor. If someone is operating without integrity, that is not a technical issue. That is an integrity issue. We have already given you tools, a process, and management sign off.
>
> The governance process is the most important thing we put in place. It creates the right amount of friction. And over time, these controls should be viewed as business enablement and business process enforcement, not security tooling.”

KEY TAKEAWAY

Governance-first organizations scale AI by treating governance as an operating model, not a security overlay. By establishing ownership, approved tools, lightweight screening, and an accountability regime early, they enable adoption that is faster, more repeatable, and more defensible as use cases move into higher-impact agentic workflows.

Governance Enables Accelerated Adoption

As AI moves beyond copilots, organizations are formalizing governance across security, IT, risk, legal, and business leadership to support faster adoption with clearer accountability. As shown earlier, organizations that scale adoption fastest are those that invest early in governance and control infrastructure.

Who Owns AI Governance and Oversight Today

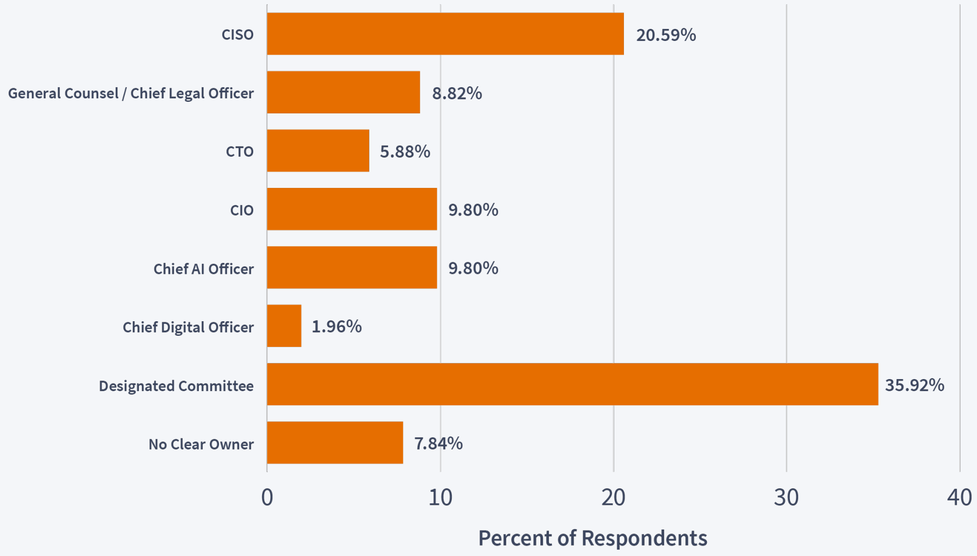

Cross functional governance improves alignment and speeds adoption. It only breaks down when ownership and escalation paths are unclear.

Why This Matters

Governance is not just policy. To scale from copilots to agents, it must translate into clear decision rights plus enforceable controls, monitoring, and escalation paths. The practical goal is to enable experimentation while keeping risk legible to leadership. Runtime detection and response is emerging as a promising control layer, but most organizations are still early.

This is an opportunistic moment for CISOs to educate non technical constituents and enable the business. Setting up a process for safe experimentation eliminates a lot of the noise from the business saying “we need AI, now.” This allows the committees to spend their time pushing higher ROI use cases into production.

Looking Ahead: The Next Maturity Divide Will Be Agent Governance

As organizations move from copilots to agentic workflows, enterprise leaders should be able to answer:

- [ ] Am I seeing all the agents running in our environment?
- [ ] Are the agents in our environment configured correctly?
- [ ] How is privilege and identity assigned, monitored, and revoked?
- [ ] How do we detect, respond to, and later reconstruct an agent deviating from its task?
- [ ] What fallback plans exist if a dominant AI vendor or platform fails?

AI adoption is now the default across enterprises. The differentiator in 2026 will not be access to models, but the ability to operationalize governance, controls, and oversight at scale. Organizations that build defensible guardrails and audit-ready observability early will be positioned to move faster into higher-impact deployments, while those that delay will face growing risk, regulatory friction, and operational uncertainty as agentic systems expand.

Board Readiness: Speaking Plain English About AI Risk

It is common that traditional board members do not understand security. We see the best-in-class risk executives educating their constituents and meeting them where they're at—with a language they can speak. We propose framing your AI journey in context of the four agents outlined below, with increasing governance and oversight as you move up the maturity curve. In the near term, we can control the blast radius of an incident by avoiding the Lethal Trifecta. We'll want to enable this over time, but it is a helpful tool to articulate how we're enabling innovation without taking unnecessary risk.

Chatbot Agent

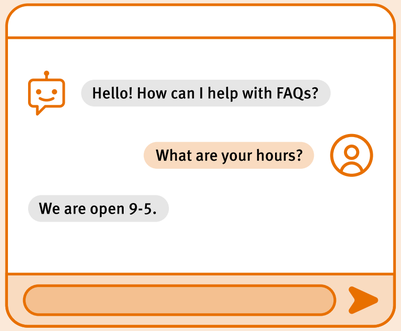

- Conversational Interface
- Reactive
- Limited Scope

LOW RISK

Assistant Agent

- Instruction-Driven
- Integrates Knowledge
- Supports Advanced Capabilities

MEDIUM RISK

Tasker Agent

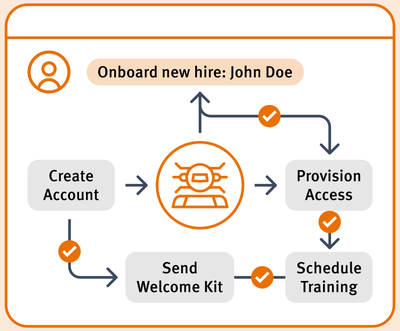

- Autonomous Execution
- Multi-Step Orchestration
- Workflow Automation

HIGH RISK

Orchestration

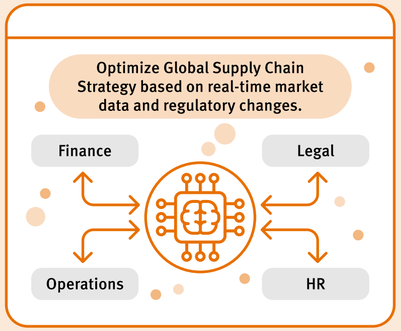

- High-level Coordination
- Complex Decision-Making
- End-to-End Process Management

CRITICAL RISK

INCREASING RISK

Framework adapted from Brian Fricke, LinkedIn, 2025.

To help boards and executive teams reason about risk as capabilities evolve, we frame agent adoption across capability and consequence. As organizations move from Chatbot Agents and Assistant Agents toward Tasker Agents and Orchestration Agents, business value increases and so does the risk surface.

To make this risk tangible, we use the “Lethal Trifecta,” originally articulated by technologist Simon Willison. The model highlights three conditions that, together, create disproportionate risk:

- Access to sensitive data
- Exposure to untrusted inputs
- The ability to take action or exfiltrate information.

Each can be managed in isolation, but combined they can produce outcomes that are hard to predict, detect, or reverse.

The Lethal Trifecta

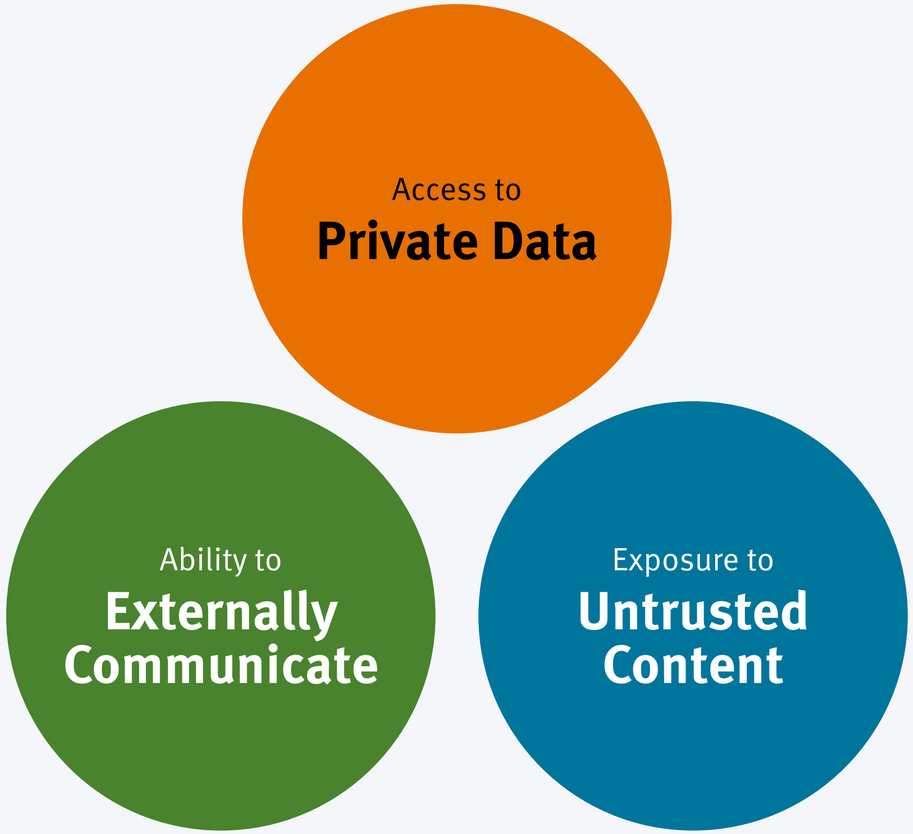

Framework adapted from Simon Willison, The Lethal Trifecta, 2025

Survey Methodology

The findings in this report are based on a survey of 102 risk executives conducted in 2025 by SINET in collaboration with Stifel Bank. Respondents include CISOs, CIOs, CSOs, and other senior leaders with direct responsibility for security, risk, and technology governance. Titles and affiliations are regularly validated to ensure relevance.

Survey results are complemented by in-depth interviews with a subset of respondents to provide qualitative context and real-world perspective.

Data Interpretation

- “Investment” reflects initiatives that respondents report as in pilot or production.
- Planning, evaluation, or policy-only approaches are excluded from investment figures.
- AI maturity is self-reported relative to industry peers.

Authors

Danny Hatfield
Email • LinkedIn

Managing Director, Software & Security
at Stifel Venture Banking

Robert Rodriguez
LinkedIn

Chairman and Founder of SINET
Venture Partner at SYN Ventures

Contributors

Brian Fricke
LinkedIn

CISO and IT Risk Head at City National Bank

Erik Naugle
LinkedIn

AI Security and Strategy Officer at Cranium AI

STIFEL | VENTURE BANKING SINET