#shifthappens
powered by AvePoint

THE STATE OF AI 2026

Scaling Trust, Control, and Readiness in the Agentic Era

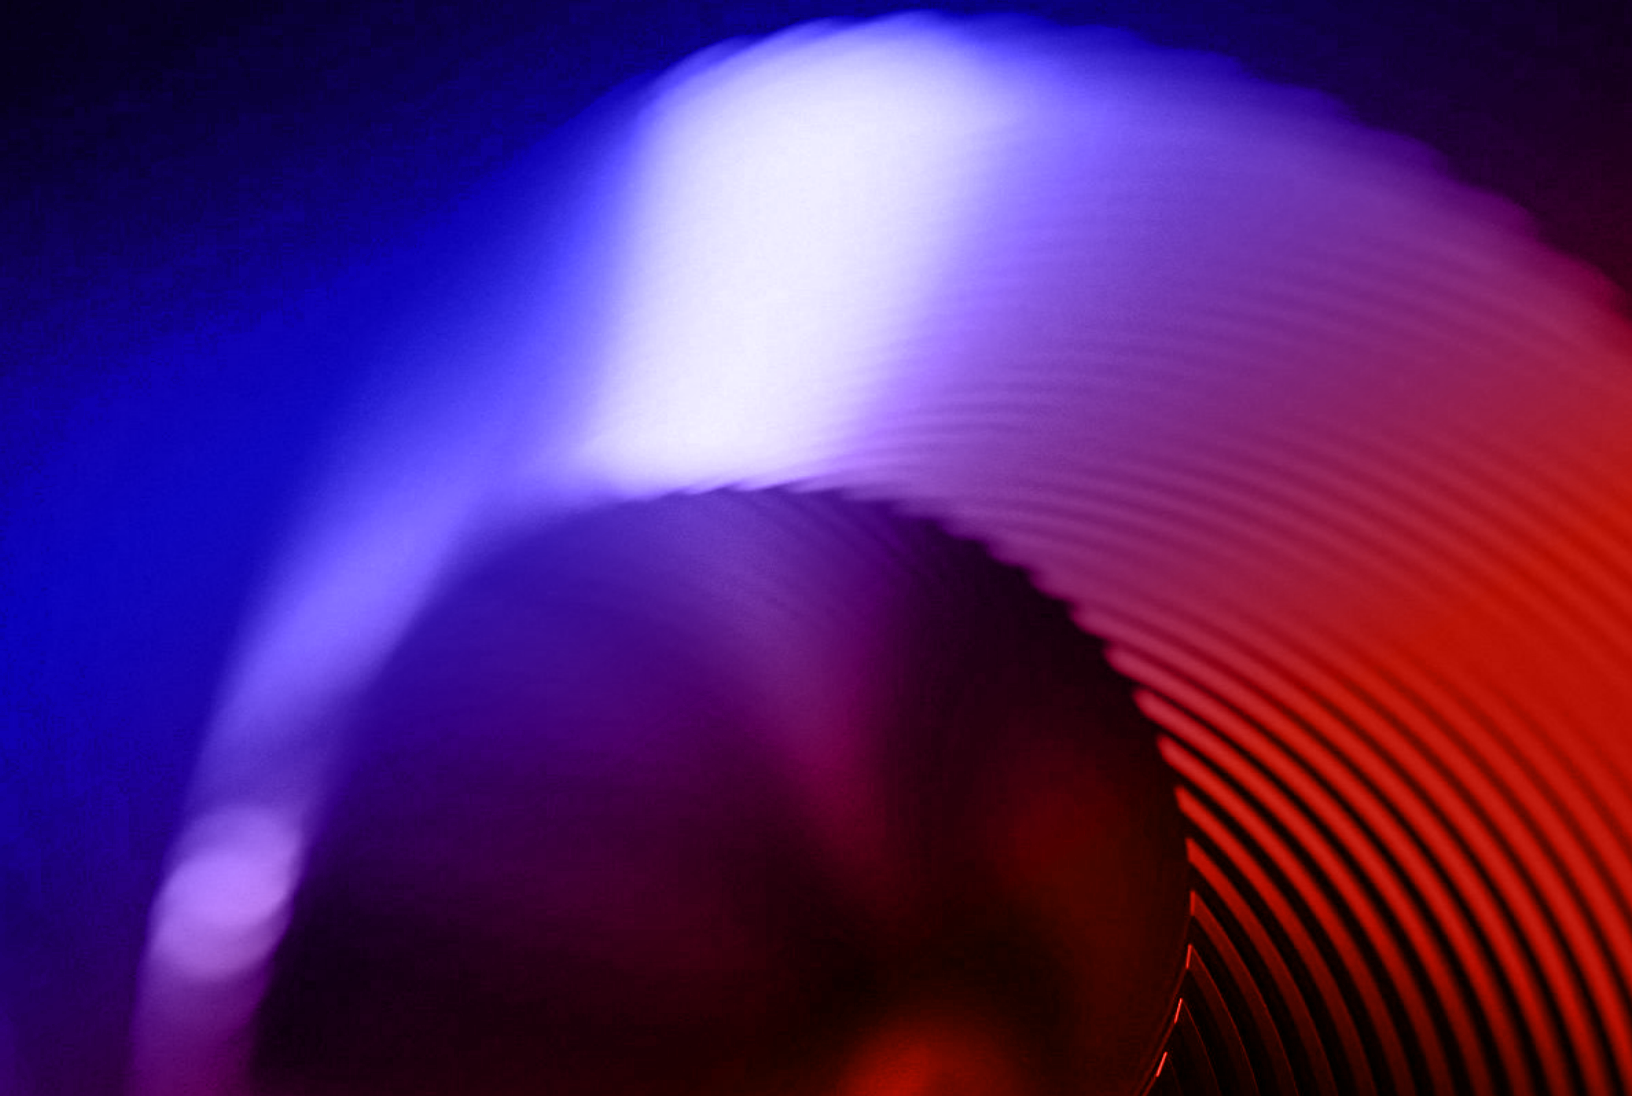

TABLE OF CONTENTS

Introduction

Introduction

Artificial intelligence is moving from experimentation to execution. Generative AI assistants are embedded in everyday work, and autonomous AI agents are taking action across business processes. AI is rapidly expanding what data can be accessed, how widely it is used, and how quickly outcomes can compound.

Enterprise AI adoption is not constrained by interest or capability. It is gated by readiness – the ability to govern, observe, and control AI-driven activity so innovation scales without scaling risk. Across both generative and AI agents, the same converging gaps – delayed deployments, rising incidents, collapsing visibility – point to what Gartner Research has identified as an emerging infrastructure requirement: an AI Agent Management Platform (AMP), a unified layer for lifecycle control, policy enforcement, and auditability across every agent in the environment.

> **By 2027, 75% of enterprises will consider the methodology they use to monitor AI agents as their most important tool, up from 1% today.**

Gartner, AI Vendor Race: AI Agent Management Platform: The Most Valuable Real Estate in AI

This year's research shows **the lack of AI-ready data, governance, and control are the primary constraints on enterprise AI at scale.** Deployments are being delayed on average for six months and risks are materializing because many organizations still lack sufficient control over the data and systems AI depends on.

The limiting factor for enterprise AI is no longer model capability. It is governance and operational readiness. Across both generative AI and AI agents, the same constraints appear repeatedly: limited visibility into usage; unresolved data security and governance gaps; and uncertainty that AI outcomes can be governed, audited, and corrected when problems occur.

This report is based on a global study of **750 respondents** with direct responsibility for information management, data security, or AI programs. The results show a consistent pattern: AI adoption is scaling faster than enterprise governance, visibility, and operational control.

Key findings

1

The Confidence-Incident Paradox:

More than 4 in 5 organizations say they are confident in their ability to prevent unauthorized data access. Yet among those reporting confidence, AI-related unauthorized access incidents remain widespread, affecting between 62% and 72% of respondents.

2

Security and privacy concerns are consistent across all forms of AI:

Data security and privacy are the top concerns for both generative AI and AI agents, and securing data used for AI training is the top-rated future investment priority – cited by 4 in 5 organizations. This consistency reinforces the report's central finding: AI value is constrained not by model capability, but by lack of AI-ready data, governance, and control.

3

Agent observability is collapsing as autonomy rises:

Up to 1 in 5 organizations do not know whether employees are using unsanctioned AI tools, a figure that has nearly tripled since 2025 for generative AI and is even higher for AI agents.

4

Deployment delays are structural:

Nearly 9 in 10 organizations delayed both agentic and generative AI deployments by an average of almost six months, driven primarily by unresolved data security and data management concerns.

5

AI-generated data is compounding governance complexity:

35.5% of enterprise data is currently AI-generated, projected to reach 42.1% within 12 months.

6

Agent adoption is outpacing readiness:

Nearly half of employees already rely on AI agents weekly or daily, yet nearly 9 in 10 organizations have experienced at least one agent-related security incident in the past 12 months.

7

Investment is moving toward AMP capabilities:

Third-party governance tools that monitor agent actions for policy alignment top the planned investment list for the next 12 months, reflecting growing demand for centralized visibility, governance, and control.

About this research & methodology

This report is part of AvePoint's annual research program on AI in the enterprise, conducted in partnership with Osterman Research. This year's global study on how enterprises are embracing generative AI and AI agents study builds on AvePoint's two previous annual reports on AI in the enterprise – The State of AI (2025) and AI & Information Management (2024).

Osterman Research surveyed 750 respondents who had direct responsibility for the information management, data security, or AI programs at their organization.

The goal of this research is to provide an independent, data-driven view of how organizations are deploying generative AI and AI agents, the issues that are slowing adoption, and the actions enterprises are taking to strengthen governance, security, and resilience for AI at scale. The intent is to present research findings and implications that are useful to enterprise leaders, independent of any specific vendor or platform.

Demographics

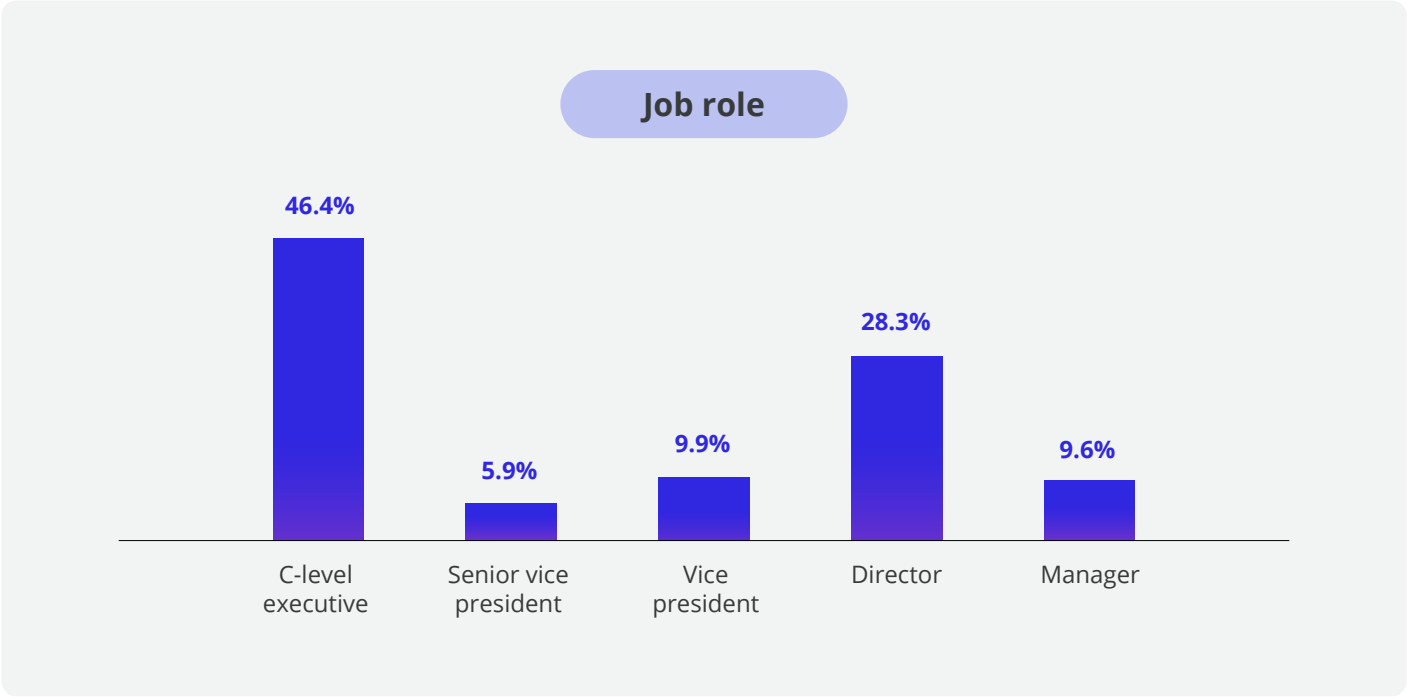

Industry

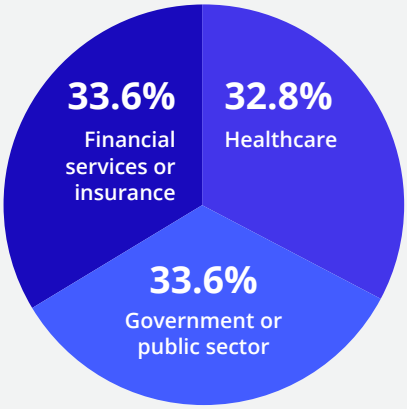

Organization size

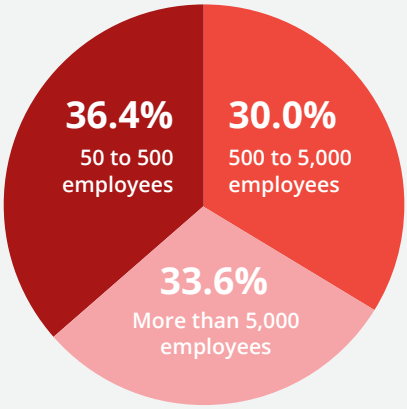

Geography

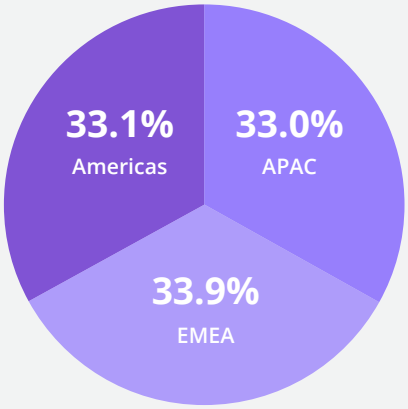

Section 1 – The State of Enterprise AI

Implementation status

AI readiness is uneven across technologies

The technologies surveyed span both AI-native capabilities like generative AI and AI agents, and established technologies increasingly deployed in AI-driven programs.

On average, 4 in 5 organizations rate six technologies used in AI programs as highly important – with AI agents, predictive analytics, and machine learning ranking as the top three. In every case, maturity trails importance: organizations consistently rate how well they use these technologies lower than the importance they assign to them.

Importance and maturity of key technologies used in AI programs

Percentage of respondents indicating "very important" or "extremely important" compared to percentage of respondents indicating "highly mature" or "extremely mature"

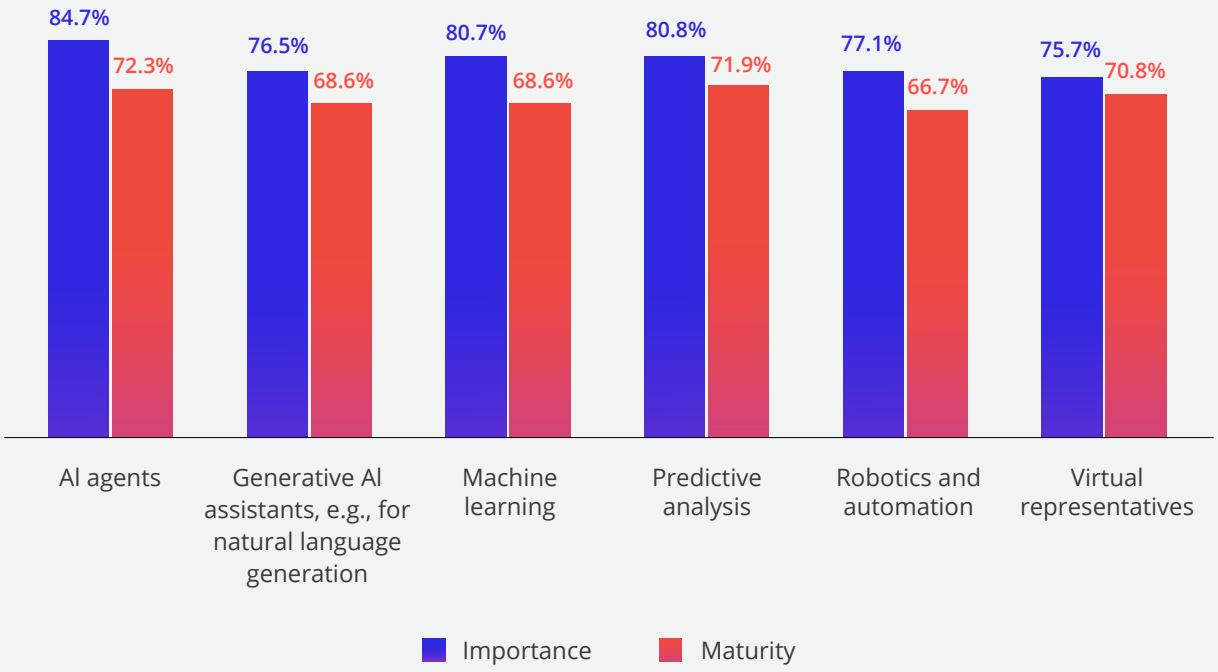

The more notable finding is not simply that maturity lags importance across the board – it is where that gap matters most. AI agents rank as the most important technology to organizations yet has the widest gap between importance and maturity. That combination signals elevated risk and reinforces a broader pattern that appears throughout this research: organizational urgency is increasing faster than operational readiness. Organizations are racing to deploy capabilities they have not yet fully learned to govern, observe, or control at scale.

Nearly half of enterprise data will soon be AI-generated

AI is rapidly becoming a primary creator of enterprise data, and that data is fundamentally different from what came before. On average, **35.5%** of organizational data is now created by generative AI assistants, and respondents expect that share to rise to **42.1%** in the next 12 months.

This is happening inside data environments already straining under growth. More than four in five organizations (84.1%) now manage at least one petabyte of data, and average data growth is expected to rise from 31.8% over the past 12 months to 39.1% over the next year. Layered on top, 78.1% of organizations say at least half of their data is more than five years old, up from 70.7% in 2025. As AI-generated content expands across already complex data environments, governance becomes harder to sustain.

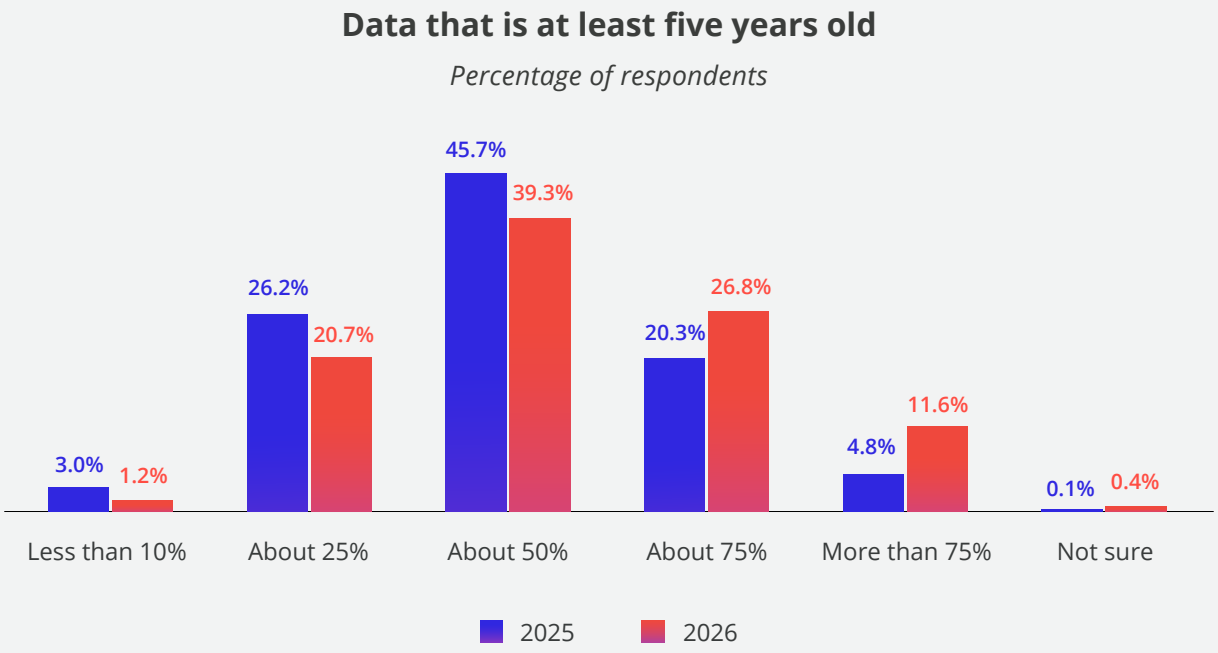

But the challenge is not simply more data. It is that AI-generated data behaves differently. Unlike human-created records, AI outputs lack clear lineage, carry uncertainty, and can be reused as input for other AI systems. Inaccurate, outdated, or non-compliant information gets amplified rather than contained, and as AI agents begin acting autonomously, the risk is no longer just poor outputs, but poor decisions and actions taken at speed and scale.

The takeaway: AI readiness depends on data readiness. As models and agent frameworks evolve, organizations will be better positioned to scale AI safely if they already have strong data quality, lifecycle governance, access controls, and guardrails in place.

EXPERT PERSPECTIVE

Dana Simberkoff

Chief Risk, Privacy & Information Security Officer, AvePoint

When Security Pressure Hits, Data Readiness Shows

Learn more on the #shifthappens podcast

DATA READINESS IS NOW AI READINESS

As AI becomes a major producer of enterprise content, data readiness becomes AI readiness. When more content is created faster – across more tools and more storage locations – small weaknesses in governance become large problems quickly. Organizations that have not modernized classification, retention, and access controls often find that AI-generated content increases oversharing, expands the sensitive-data footprint, and makes it harder to determine what should be kept, secured, or deleted.

This is also where data quality and data lifecycle management move from being "cleanup projects" to being core AI safeguards.

> ***Training and reasoning on redundant, outdated, or trivial (ROT) content increases the probability of irrelevant output and poor decision-making. When AI agents begin acting autonomously on that output, governance failures can propagate operationally across systems, workflows, and decisions at machine speed.***

The practical implication of the research is that policies and training are necessary yet insufficient on their own. Organizations need enforceable foundations that scale, including visibility into where AI-created content lands, governance that applies across locations and tools, and disciplined lifecycle controls that reduce the chance of AI amplifying poor data. Trust in AI outcomes starts with control of the data AI is built on and allowed to use. While AI models and agent frameworks continue to evolve, investments in governance, lifecycle management, data quality, and access controls remain durable capabilities that strengthen every future AI initiative.

Issues

Organizations are under growing pressure to prove AI value faster. Timeframes for requiring a return on AI investments have shrunk since last year, with 86.3% of organizations seeking a return in 12 months or less (up from 81.9% last year). The percentage willing to wait longer than a year declined from 18.1% to 13.7%.

Timeframe for returns on AI investments: 2025 vs 2026

Percentage of respondents

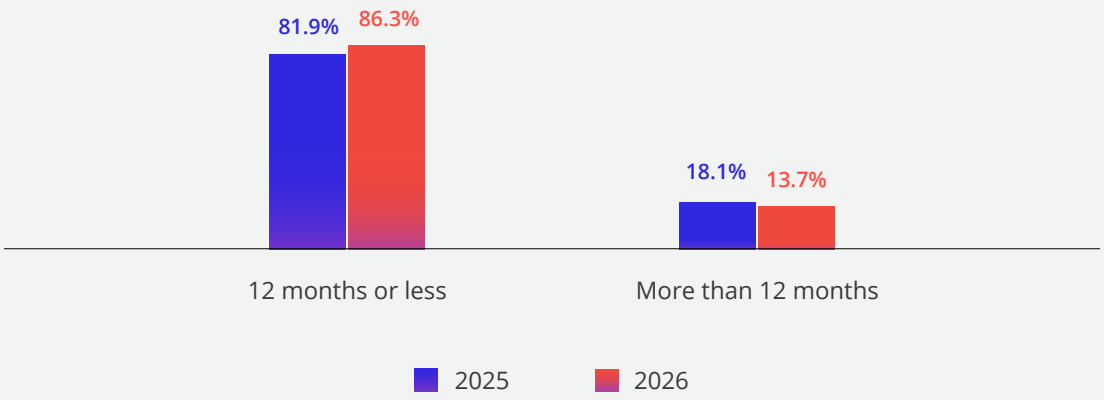

Together, these findings shift the ROI conversation from AI adoption to measurable business outcomes. As organizations seek faster returns, success is becoming less about what AI costs to operate and more about the value it creates through improved efficiency, better decisions, stronger customer outcomes, and operational scale. This pressure is giving rise to AI FinOps, a discipline focused on tying AI spend to business outcomes, which we'll return to in Section 2.

High confidence is not preventing AI security incidents, indicating a readiness gap

Among organizations reporting the highest confidence in their ability to prevent unauthorized data access, 62% still experienced at least one AI-related unauthorized access incident in the past year. Among those reporting "very confident," that figure rises to 72%.

This is the most consequential finding in this year's data. Organizations are not failing to prioritize AI security: more than 4 in 5 respondents report being very or extremely confident in their ability to prevent unauthorized access, up from 75.5% in 2025. Yet confidence is not translating into outcomes. The gap between what organizations believe they can prevent and what they are experiencing reveals a structural readiness problem.

Confidence to achieve AI and related security outcomes

Percentage of respondents

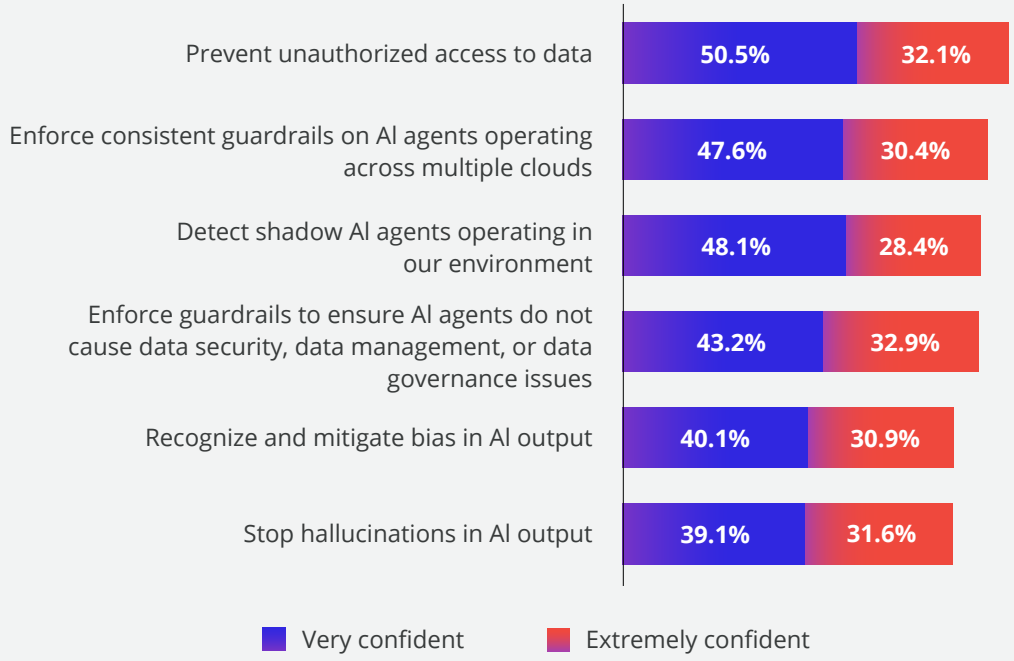

Confidence at this level is typically driven by policy intent and planned controls – having an acceptable use policy, deploying a sanctioned tool, or completing a pilot program. Incidents, however, are driven by what is operational in practice: whether data access is correctly governed, whether usage is visible, whether guardrails are enforced, and whether risky behavior can be detected and corrected before it becomes a breach.

The implication is that confidence itself may be functioning as a risk indicator rather than a protective factor. Trust isn't created through policy or intent – it emerges through visibility, governance, and enforceable controls.

Year-over-year trends reinforce this concern. Confidence in organizations' ability to recognize and mitigate bias in AI output and to prevent hallucinations has declined compared to 2025, even as AI adoptions continue to expand.

City of PORT ST. LUCIE
HEART OF THE TREASURE COAST

For the City of Port St. Lucie, AI readiness started with data governance. By applying policies early, the city streamlined its Microsoft Copilot deployment while reducing operational burden on internal teams.

How the City of Port St. Lucie Sets the Standard for Modern Municipal Governance

Read the full story

EXPERT PERSPECTIVE

John Peluso

Chief Technology Officer, AvePoint

AI IS A CONTROL PROBLEM, NOT A MODEL PROBLEM

One of the most important signals in this year's data is not simply that confidence is high. It is that high confidence coexists with high incident rates. That combination suggests a readiness gap that many organizations underestimate.

Confidence is often shaped by intention, such as having an acceptable use policy, deploying a sanctioned tool, or running a pilot program. Incidents, however, are shaped by what is operational in practice: whether data access is correctly governed, whether usage is visible, whether guardrails are enforced, and whether risky behavior can be detected and corrected before it becomes a breach.

> ***This is why trust in AI is increasingly a control problem rather than a model problem. As AI becomes embedded in workflows, the key question is not only whether users are trained or whether a policy exists, but whether organizations can reliably control what AI can access, audit what AI did, and remediate outcomes when something goes wrong.***

The organizations that close the gap between confidence and reality will be those that build readiness foundations that scale with AI: consistent governance, enforceable access controls, visibility into sanctioned and unsanctioned usage, and recovery capabilities that reduce the blast radius of inevitable mistakes.

Section 2 - AI agents

Introduction

Over the past 12 months, AI vendors have dramatically increased messaging around autonomous AI agents – and Gartner's Hype Cycle for Agentic AI places it squarely at the Peak of Inflated Expectations, reflecting extraordinary market attention and aggressive adoption intent. Where generative AI requires constant prompting and human chaining of recommendations, autonomous agents promise to make decisions, coordinate activities, and monitor systems with minimal human involvement and, in some cases, eliminate human work entirely.

The research data tells a more sobering story. 88.4% of organizations experienced at least one security breach due to AI agents in the past 12 months, and data security and privacy remain the top concern when implementing agents. As organizations move from experimenting with AI agents to operational deployment, several structural realities are becoming clear:

- **AI agents operate without judgment or intent.** They reason probabilistically and act autonomously within the permissions and data provided to them. When governance guardrails are incomplete or outdated, agents can bypass controls in ways that threaten data security, privacy, and compliance.
- **Small readiness gaps scale into large risks at machine speed.** Controls designed for human decision-making break down when AI agents execute tasks continuously, across systems, and without pauses for review — shifting the risk profile from incorrect output to incorrect action.
- **Agent deployment is outpacing governance maturity and visibility.** Many organizations are adopting AI agents while still working through foundational issues like identity management, policy enforcement, and lifecycle governance — often unsure which agents are running, how they were created, or whether unsanctioned "shadow" agents are operating alongside approved systems.
- **The rise of AI agents is forcing a re-evaluation of security architectures.** Organizations and vendors alike are reassessing identity controls, governance mechanisms, and recovery capabilities to manage AI-driven activity safely at scale.

These realities point to a common need: centralized visibility, lifecycle controls, and enforceable guardrails that extend across every agent in the environment — regardless of which platform built or deployed it. This is structurally different from governing generative AI, which primarily requires controlling what AI can access and produce. AI agent governance requires governing autonomous actions taken across systems, often without human review, at a speed that outpaces traditional oversight.

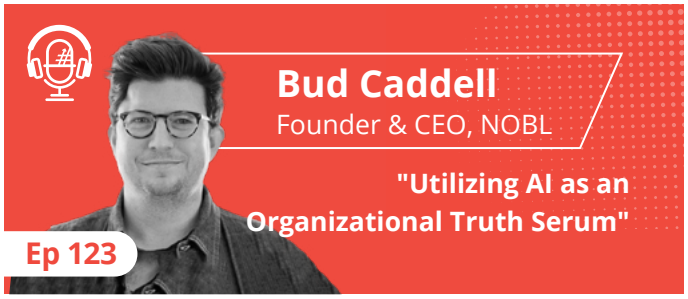

"AI doesn't introduce new organizational problems — it accelerates the exposure of existing ones."

Learn more on the #shifthappens podcast

This is why the Agent Management Platform (AMP) is emerging as critical infrastructure – a unified layer for visibility, lifecycle control, policy enforcement, and auditability across every agent in the environment. Gartner projects that enterprise investment in AMP technologies will exceed $15 billion by 2029.

> **The average Fortune 500 enterprise will manage 150,000 AI agents by 2028.**

Source: Gartner Identifies Six Steps to Manage AI Agent Sprawl

The organizations best positioned to scale safely will be those that treat agent governance – and the underlying data governance foundations – as critical infrastructure, not an afterthought.

EXPERT PERSPECTIVE

Dr. Tianyi Jiang (TJ)

CEO and Co-Founder, AvePoint

From Novelty to Necessity: How Agentic AI is Reshaping Enterprise Value

Learn more on the #shifthappens podcast

WHY AI AGENTS RAISE THE BAR FOR CONTROL

AI agents completely rewrite the enterprise risk equation. Because agents can execute autonomously, the risk shifts from flawed outputs to flawed actions. Controls built for human judgment break down when decisions are made continuously and executed at machine speed.

The data reflects this shift. Agent adoption is accelerating, with many employees already relying on agents regularly, while breaches linked to agent activity remain widespread. This combination signals a structural risk: operational exposure is expanding faster than governance maturity.

In agent-driven environments, trust depends on control. Organizations need visibility into what agents are doing, enforceable guardrails, auditable records of activity, and recoverability when unwanted actions occur. Without these foundations, organizations cannot scale agent adoption safely or maintain compliance and accountability as AI becomes more autonomous. Gartner projects that by 2027, 75% of enterprises will consider agent monitoring methodologies among their most important AI management tools, reflecting the growing importance of observability as autonomy increases.

The practical path forward is not to slow AI innovation. It is to strengthen readiness so innovation can scale without scaling risk.

Implementation status

Access to AI agent tools – both sanctioned and unsanctioned – is expanding quickly

Around one third of employees have access to sanctioned and/or unsanctioned tools for creating AI agents within their organization – a proportion which is expected to increase to over one half for sanctioned tools in 12 months. Access to tools for creating AI agents follows a very similar pattern to access to generative AI tools over the three time periods we asked about.

21.1% of respondents don't know whether unsanctioned tools are being used to create AI agents for work processes – a figure higher than for generative AI but equally concerning for data security and governance implications.

Access to sanctioned and unsanctioned tools for creating AI agents for work processes

Percentage of employees with access to sanctioned and unsanctioned tools

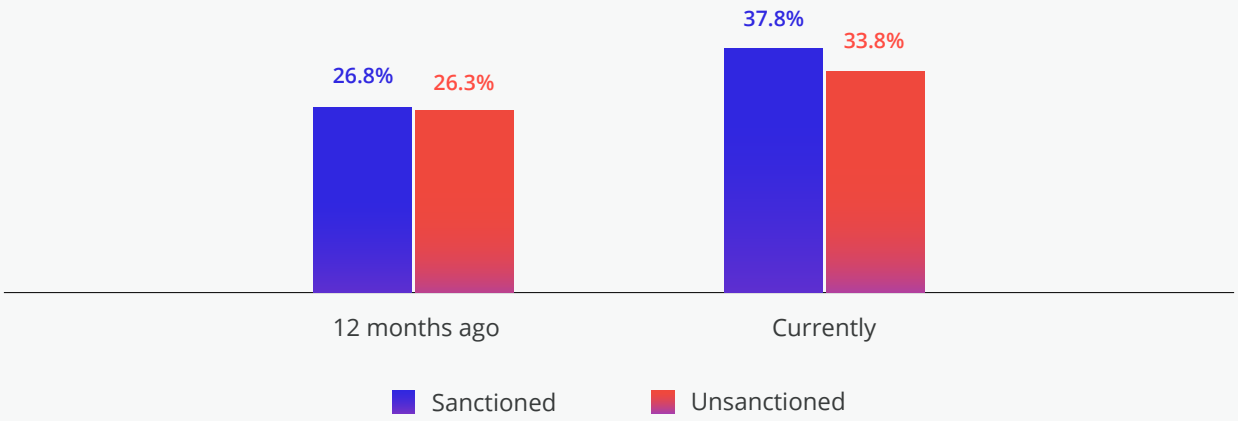

Nearly half of employees already rely on AI agents weekly

46.9% of employees rely on AI agents daily or weekly to complete work tasks – for example, for autonomous customer service interactions, cybersecurity posture assessments, or autonomous inventory modeling and replenishment – and the usage pattern for AI agents aligns closely with the usage pattern for generative AI assistants.

"Treat AI agents as employees from a guardrail perspective — what access we give them, what authority they have, what data they touch."

Learn more on the #shifthappens podcast

Reliance on AI agents for completing work tasks versus the use of generative AI assistants

Average percentage of employees per cadence

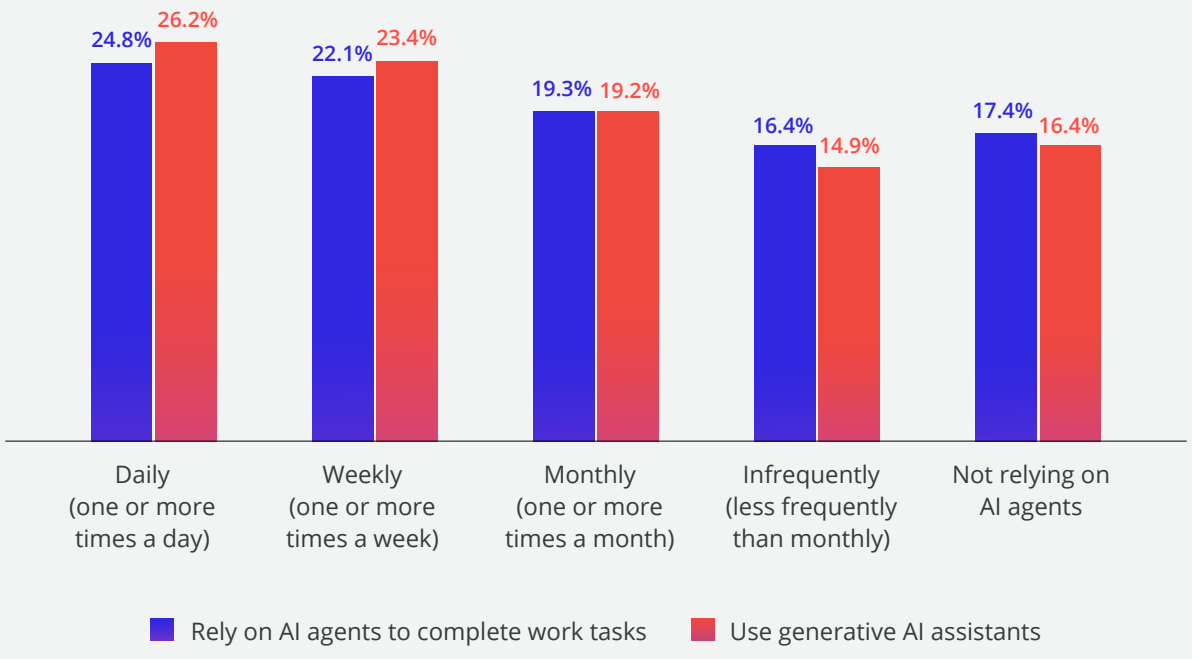

Issues

AI agents are scaling into work faster than organizations are ready for

Work processes that include the use of AI agents are anticipated to double in 12 months' time compared to 12 months ago. There are two types of learning in effect: first, implementation of AI agents into initial work processes creates insight for subsequent implementation efforts with different work processes; and second, the degree of saturation of AI agent capabilities within a given work process can increase over time.

Work processes that include the use of AI agents

Percentage of work processes

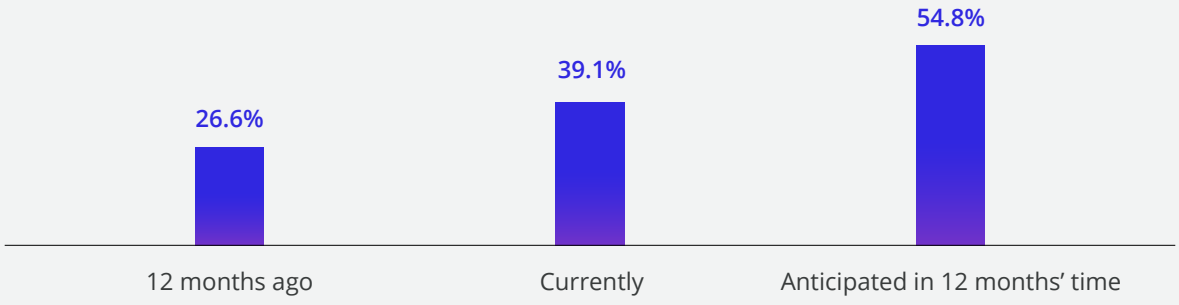

Incorporating autonomous AI agents into cybersecurity response processes is the most important use case for organizations, followed by marketing, inventory, software development, and customer support use cases.

Use cases for AI agents

Percentage of respondents

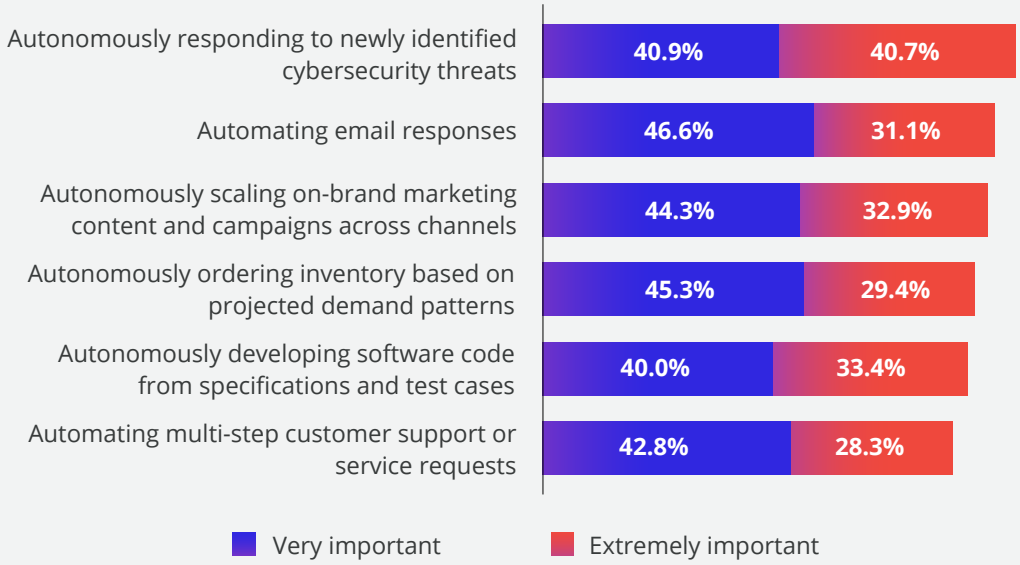

AI agent ROI is shifting from tool cost to work displaced

Respondents expect AI agents to reshape a meaningful share of work currently performed by humans, with 26.7% of human work expected to be replaced by AI agents within 12 months and 49.7% within five years. But the value story is not primarily about headcount reduction. Reducing employee headcount ranks last among the reasons for using AI agents, while increasing process efficiency, freeing employees for strategic work, and augmenting capabilities rank higher.

This suggests organizations are beginning to evaluate agent ROI through work displaced, accelerated, or augmented (including reduced manual effort, compressed process time, improved consistency, and more effective use of human capacity). The sharper way to frame the ROI question is not *what does the agent cost to run*, but *what human cost does it displace*. A $7 agent run that delivers $300 in process savings is a win for the business, even if it looks expensive on an IT invoice. That comparison, agent run cost versus the manual cost it replaces, is where the real ROI conversation lives.

This is why, as we foreshadowed in the prior section, AI FinOps is emerging as a discipline (and rising on Gartner's Hype Cycle for Agentic AI). Unlike per-user software, agentic AI generates variable spend — LLM calls, reasoning traces, tool retries, multi-agent loops — that requires attribution and observability to tie spend to outcomes. Early adopters will scale agents with confidence because they can prove the value side, not just the cost side.

Hear how AI delivers real value on #shifthappens — AI in Action: From Change to Competitive Advantage

Learn more on the #shifthappens podcast

Percentage of human work replaced by AI agents

Percentage of work currently done by humans

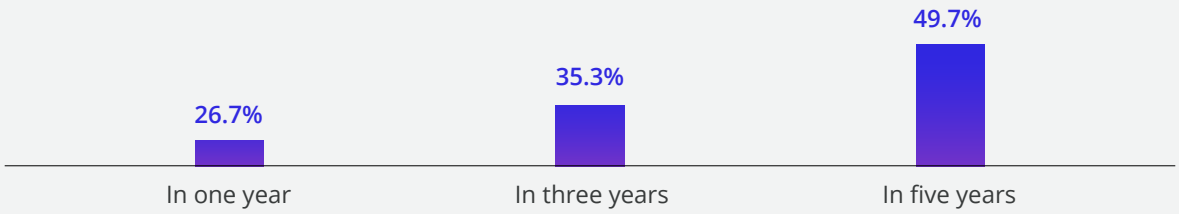

Reasons for using AI agents

Percentage of respondents

Concerns

Unpredictable outcomes drive concerns with AI agents

AI agents don't always behave in predictable ways. They learn autonomously from data and make decisions based on probabilities rather than fixed rules. This characteristic – shared with generative AI systems – drives several of the highest areas of concern identified in this research.

The most widely cited concern is AI agents making incorrect judgments or taking inappropriate actions that damage data, followed closely by employees acting on flawed outputs. In generative AI, this often appears as hallucinated or inaccurate responses. In autonomous AI agent systems, the equivalent risk is flawed decisions across cybersecurity, marketing, and customer service – decisions that can be damaging, costly, and difficult to reverse.

When examining only those respondents who are extremely concerned, AI agents bypassing human-in-the-loop controls emerges as the top issue. When agents circumvent safeguards designed to prevent harm, they combine incorrect judgment with autonomous execution – magnifying risk rather than containing it.

Concerns about threats from the use of AI agents

Percentage of respondents

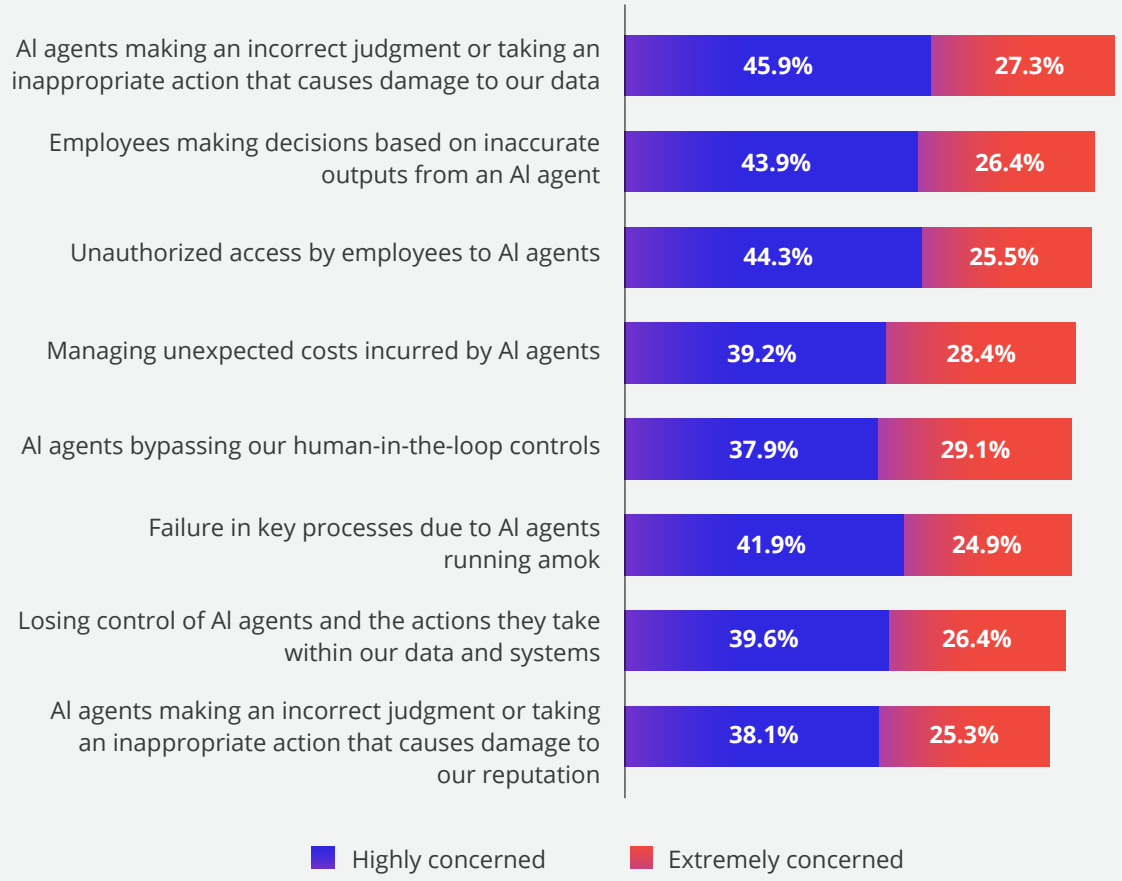

Data security and privacy are primary concerns with implementing AI agents

The top-rated concern for implementing AI agents – ensuring data security and privacy – also heads the list for implementing generative AI assistants. Ethical issues and bias ranks in second place for AI agents, compared to last place for generative AI assistants, likely due to autonomous decision-making and action-taking options in AI agent systems that operate without human oversight, and thus the option to check for bias. That partly explains why ensuring sufficient human-in-the-loop controls for overseeing AI agents is the third-rated concern. Addressing these concerns requires the development of disciplined processes, the implementation of governance systems to implement and assure guardrails, and more lived experience with AI agents to develop learnings and elevate implementation maturity.

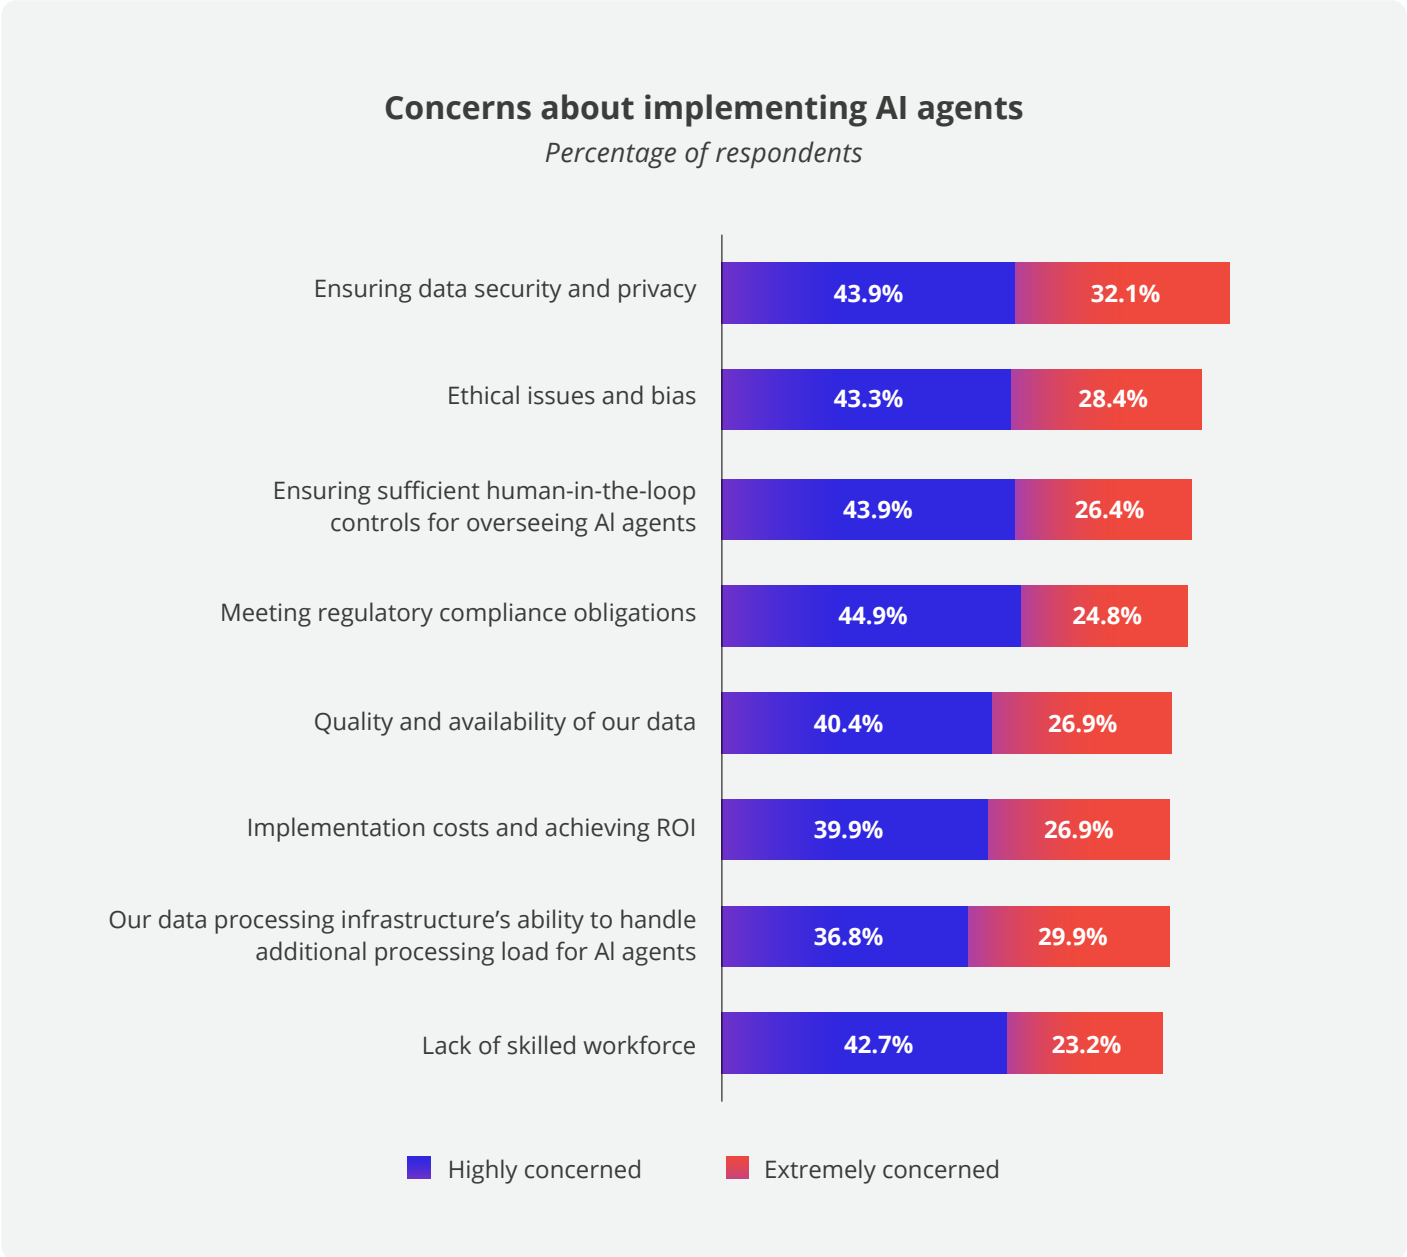

Readiness, ROI, and data security risks top the reasons for running a pilot program for AI agents

Exploratory and data readiness reasons top the list for running a pilot program for AI agents, with employee readiness the least important reason. This is the same pattern we saw in last year's data for generative AI assistants – where there was a higher focus on ensuring that the organizational conditions for AI agents can be met, e.g., data is ready, ROI is understood, and data security risks are tested and confirmed. If AI agents follow the same organizational adoption pattern as generative AI assistants, we expect to see elevated importance for employee readiness next year.

This finding helps reconcile an apparent contradiction in the report. While AI agents are rated as relatively mature in some organizations, that maturity may reflect early access and experimentation rather than disciplined operational readiness. The pilot program results suggest many organizations are still working out the governance, auditing, and ROI foundations required to deploy AI agents reliably at scale. In other words, usage may be advancing faster than readiness.

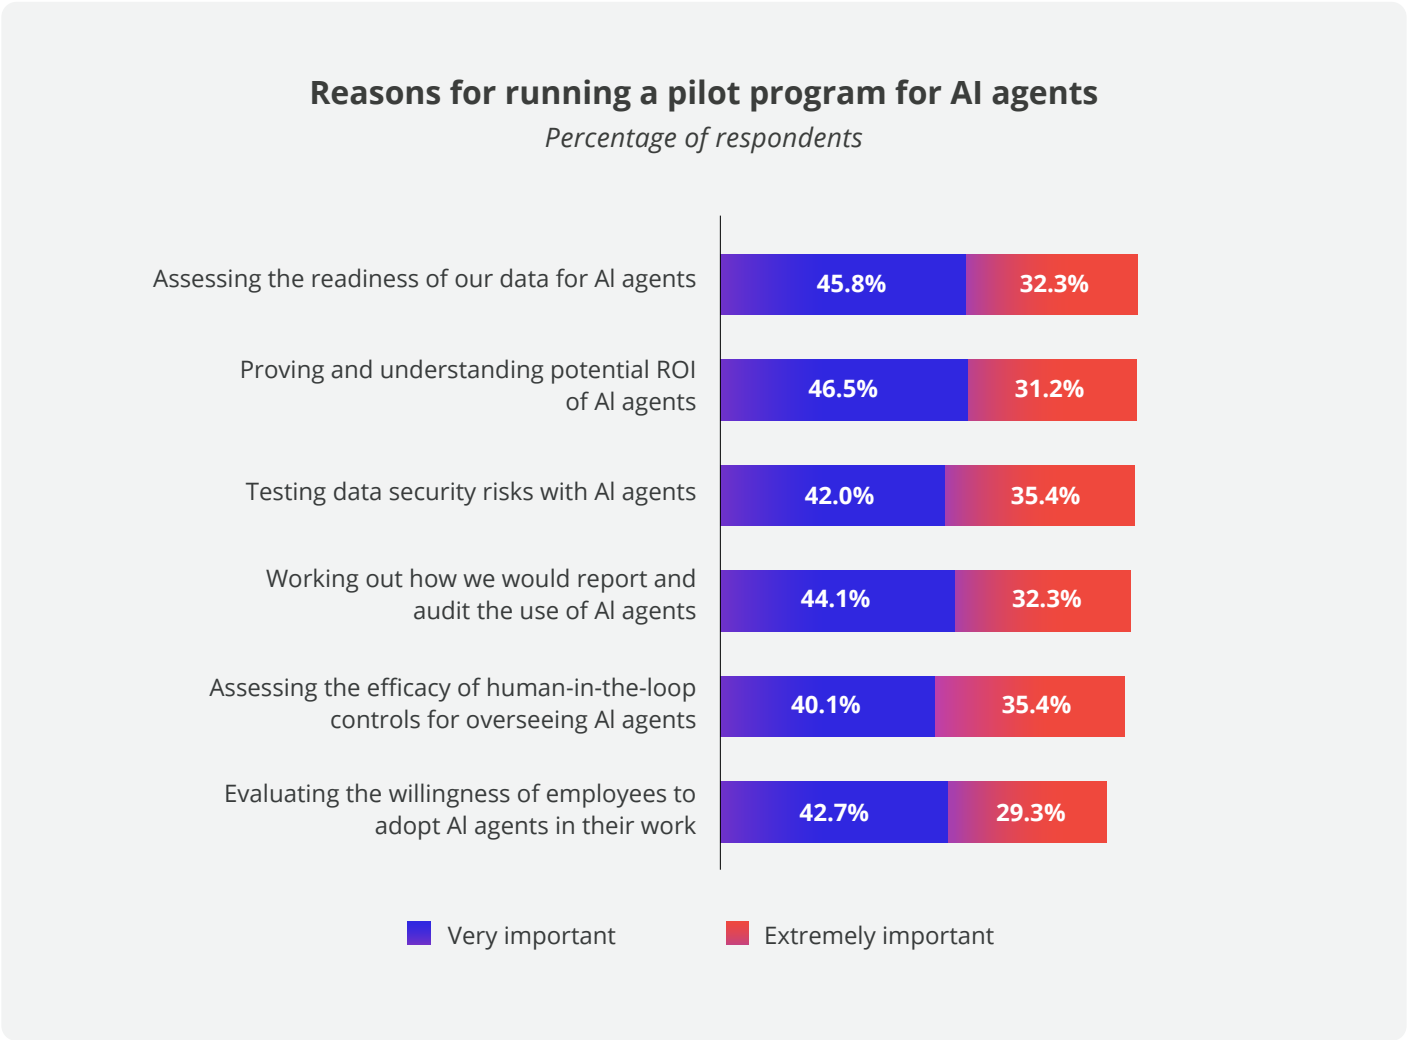

Human-in-the-loop controls added and employee training offered to mitigate security concerns with AI agents

95.5% of organizations have taken one or more actions to mitigate or address the security concerns they experienced with AI agents over the past 12 months. Adding human-in-the-loop controls was the most common action, followed by training employees on how to safely use AI agents. Human-in-the-loop controls add human verification to the actions of AI agents, pulling back from the vision of fully autonomous agents to reduce the risk of unwanted behavior by AI agents, such as data security and privacy violations.

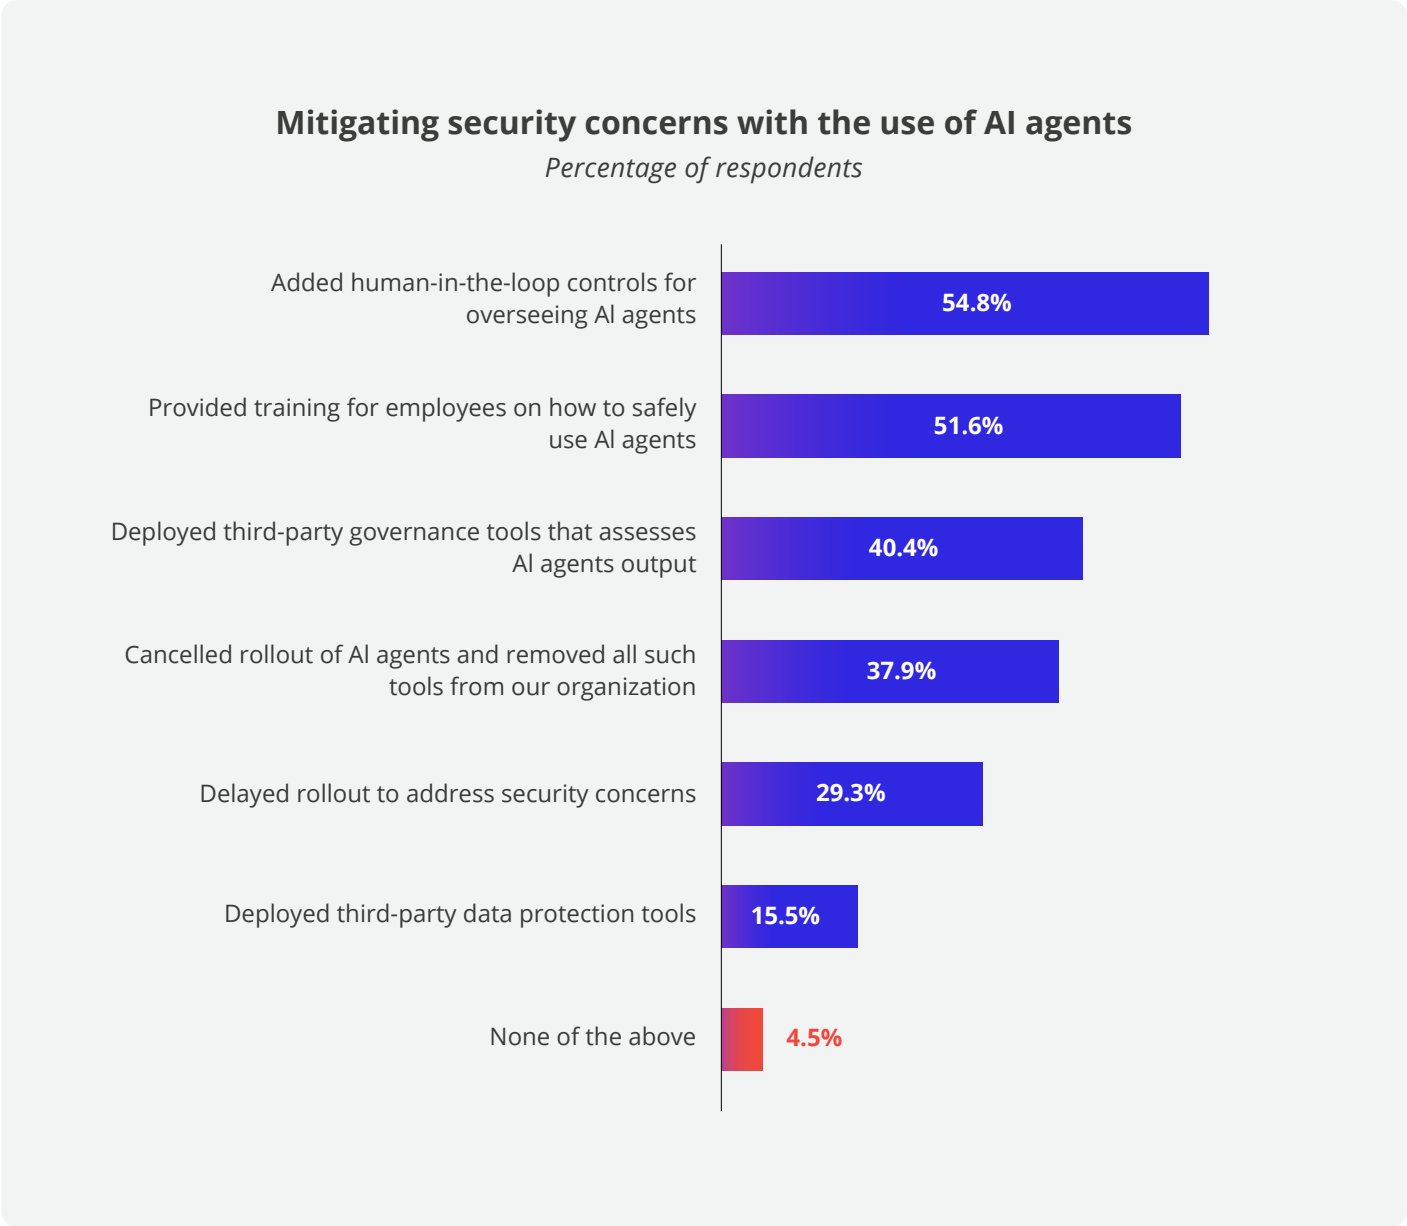

Internal readiness is the biggest barrier to AI agent rollouts

Exploratory and readiness issues presented the most difficulty for organizations when rolling out AI agents, with data quality, understanding ROI, and reporting and auditing at the top of the list. Along with data security and potential exposures (in fourth place), these reasons are internally focused and present as existing issues that must be addressed when embracing new technologies that leverage existing data and system capabilities. For example, if current data quality within the organization is low, the efficacy of AI agents to operate autonomously is compromised, which undermines potential ROI and likely allows data security exposures due to missing or insufficient access controls.

Difficulty of issues and challenges in rolling out AI agents

Percentage of respondents

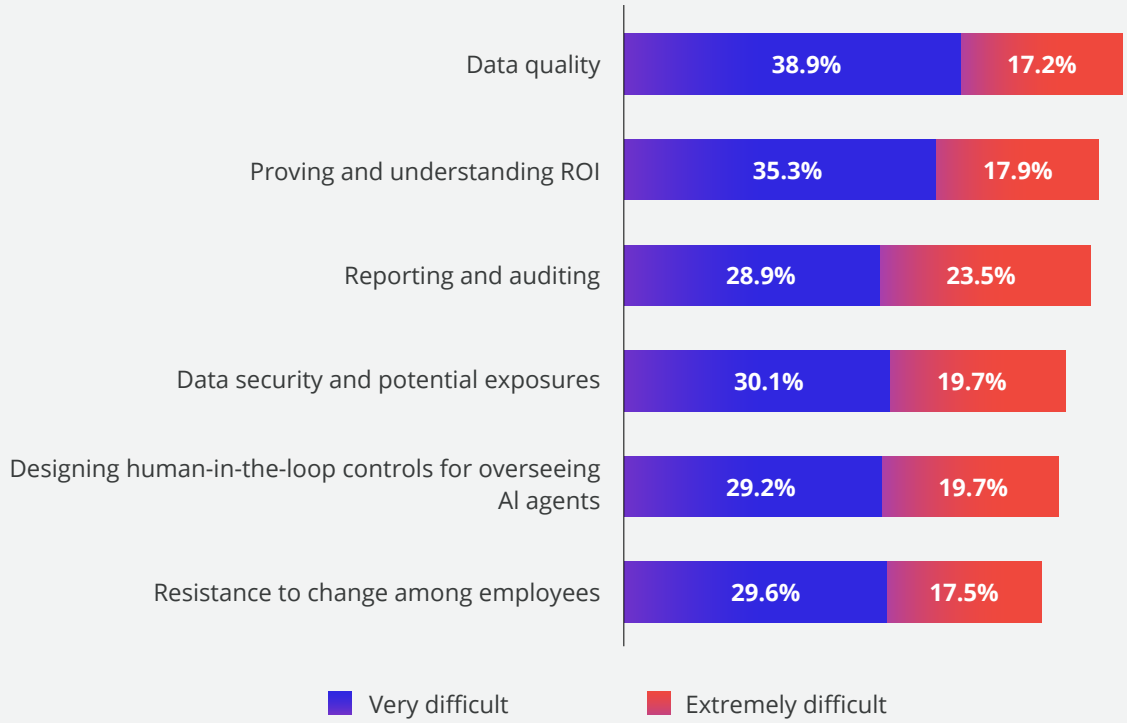

Timeframes

Pre-existing data security and data management concerns hinder adoption of AI agents

86% of organizations have delayed their deployment of AI agents by an average of 5.92 months due to data security and/or data management risks, marginally longer than the delay this year for generative AI assistants due to the same factors. The shape of the delay curve is almost identical between the two.

The near-identical delay curves across GenAI and AI agents reveal that the constraint isn't the technology – it's data security, quality, and management readiness. These delays are structural, not temporary.

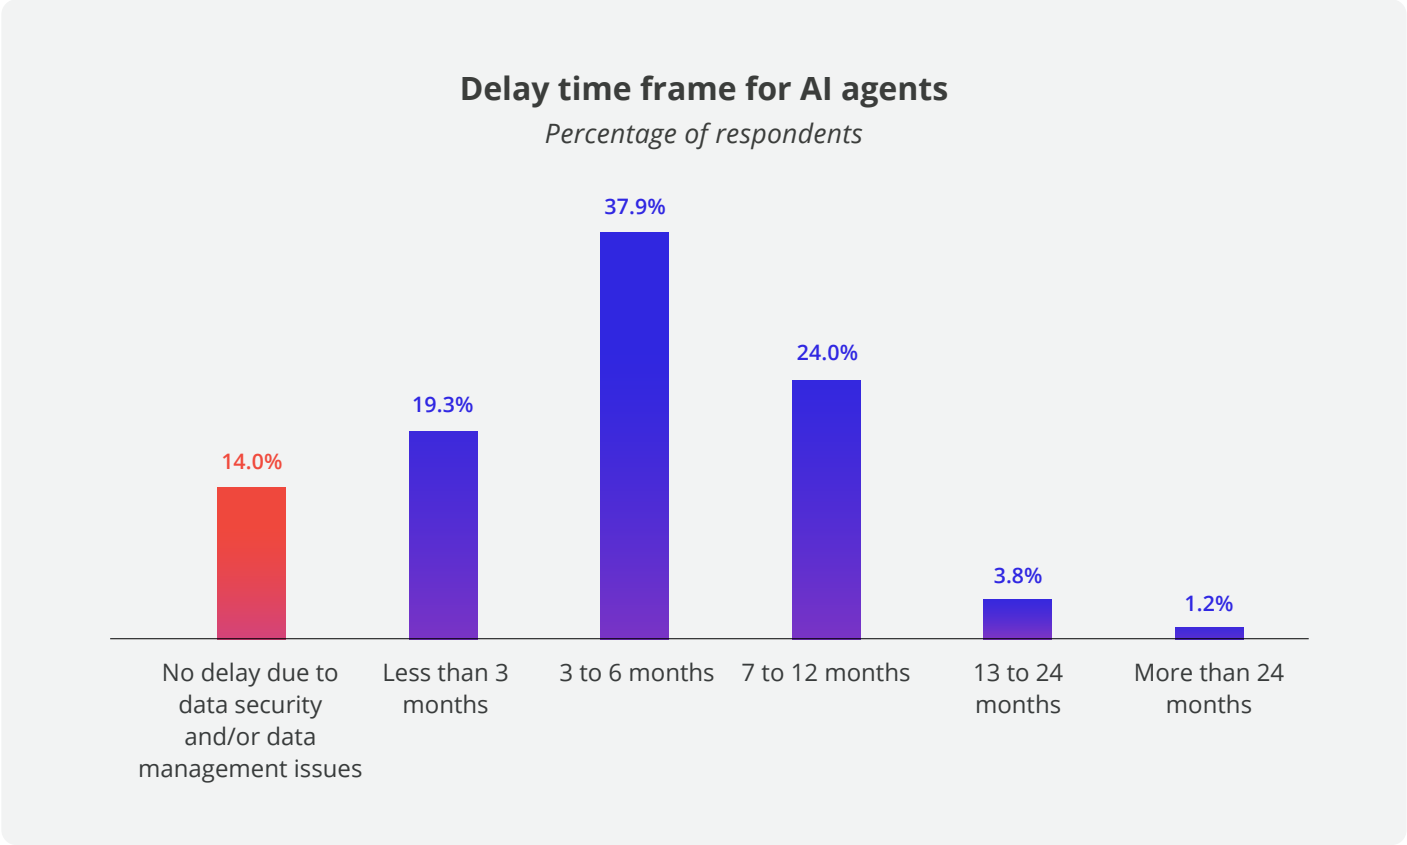

The specific factors slowing down AI agent deployments reinforce this finding. The leading causes of delay are non agent-specific technical limitations, but familiar readiness challenges: concerns about unauthorized exposure of sensitive data, incorrect outputs that result in poor decisions or costly outcomes, and uncertainty about business value.

Reasons for slowing down the deployment of AI agents to employees

Percentage of respondents

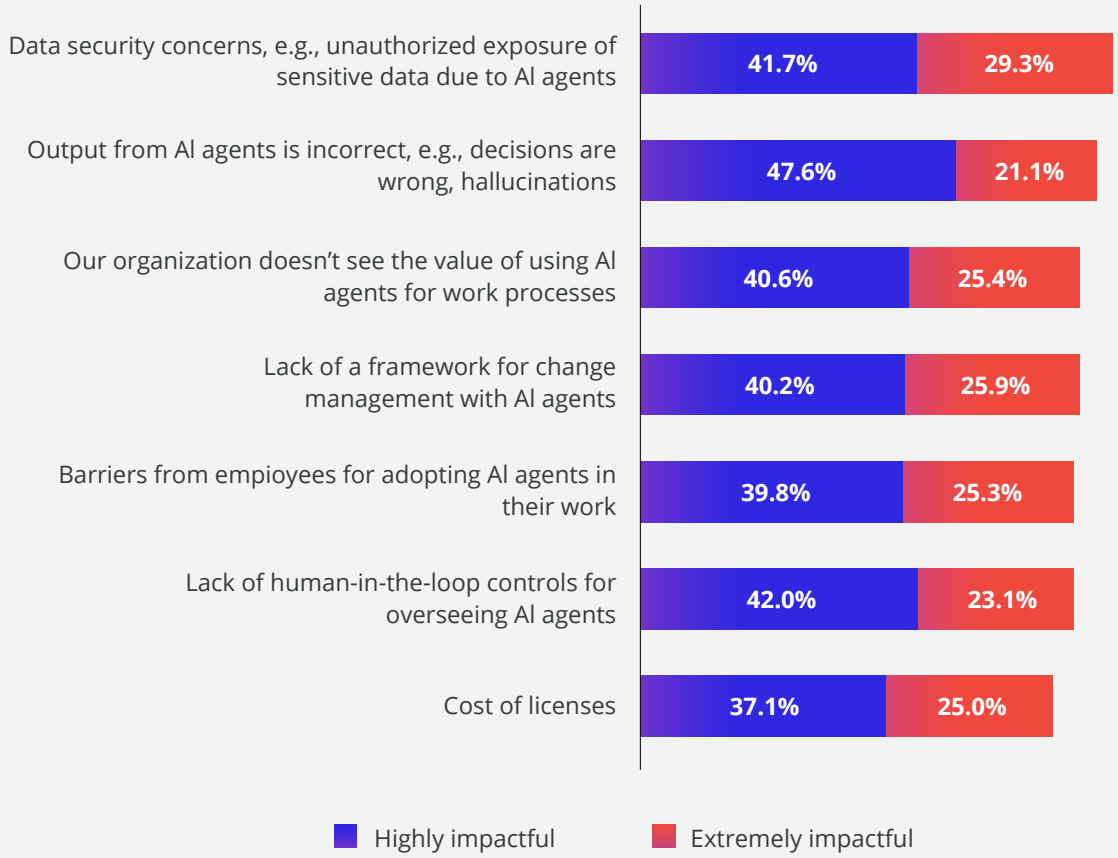

Investments in third-party governance and data security tools to safeguard AI agents

Over the next 12 months, organizations are increasing their investment in five areas to safeguard the actions of AI agents within the context of their data. The highest intent to increase investment is in governance tools that monitor the actions of AI agents for accuracy and alignment with data governance policies. These investment patterns suggest organizations increasingly recognize that AI governance requires active monitoring policy enforcement, and operational oversight rather than policy creation alone. Taken together, these investment patterns are an early operational signal of AMP adoption.

Additionally, third-party data security tools to protect organizational data from erroneous actions by AI agents ranks highest (80.7% for "increase" or "stay the same"). That is followed by third-party data security tools to protect AI agents from interference (78.1%).

As in the prior two years of this report, organizations continue to prioritize investments that address data security and governance concerns as foundational requirements for AI adoption.

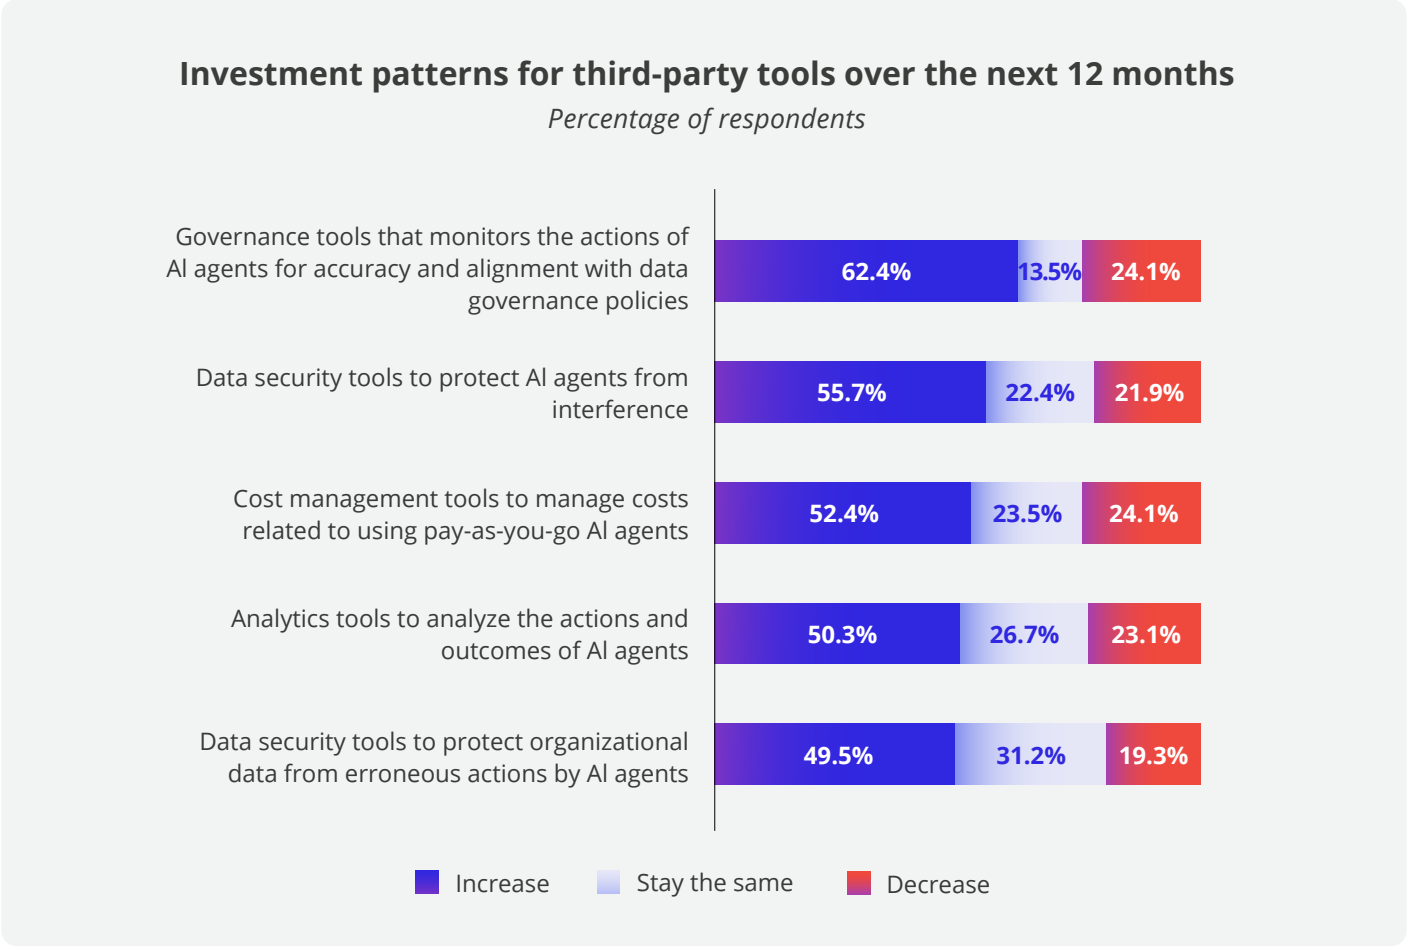

The investment pattern in this chart isn't generic governance spend. It's an early signal of AMP adoption. 62.4% of organizations plan to increase investment in tools that monitor AI agent actions for policy alignment, 55.7% in tools that protect agents from interference, and 52.4% in agent cost management. That's the exact capability stack Gartner forecasts will define the AMP category – and it's already showing up as planned 2026 spend.

Data security incidents

88.4% of organizations experienced at least one security breach due to AI agents during the past 12 months. Data leakage was the most common (50.1%), followed by manipulation of AI agents by malicious or untrusted inputs (49.6%).

Frequency of security breaches due to AI agents over the past 12 months

Percentage of respondents who experienced one or more breaches

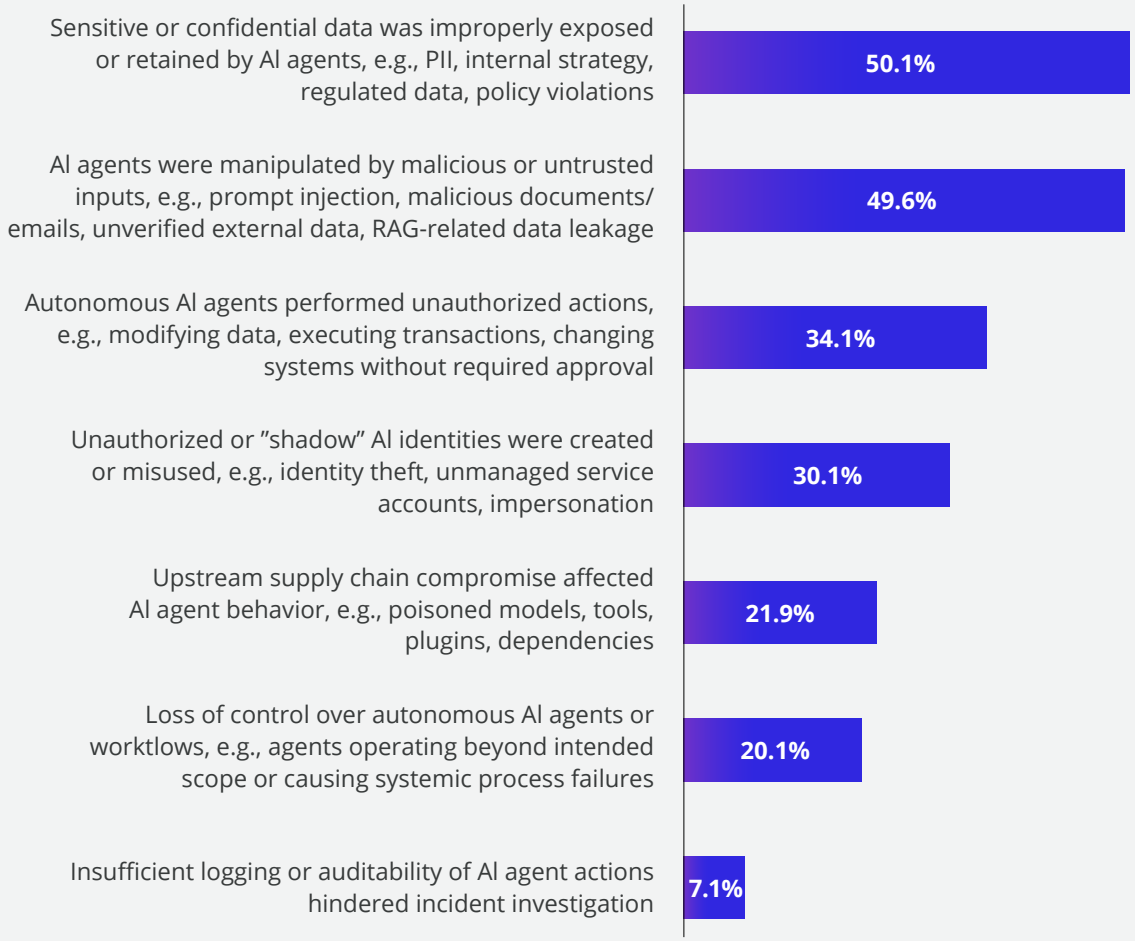

While AI agents represent the most acute and fastest-moving governance challenge, they are not the only one. Generative AI assistants remain a live and largely unsolved problem across most organizations – and the emergence of agents has not replaced those challenges. It has compounded them.

Section 3 - Generative AI

Introduction

The risks introduced by AI agents do not exist in isolation. Generative AI assistants, now embedded in everyday workflows, continue to expose foundational weaknesses in data security, governance, and visibility. Where agents act, generative AI creates – and the data it produces at scale compounds every governance gap agents exploit.

The findings in this section show that generative AI is not failing because the technology lacks capability. Instead, organizations are struggling to scale use safely because readiness has not kept pace with adoption. As generative AI creates a growing share of enterprise data and operates across both sanctioned and unsanctioned tools, gaps in control become harder to detect and more costly to correct. The result is a familiar pattern: rising usage, rising concern, and delayed deployments driven primarily by data security and data management risks.

Implementation status

Generative AI use is increasing, including unsanctioned tools

Respondents acknowledge that employees use a mix of sanctioned and unsanctioned generative AI tools for completing work tasks. The use of sanctioned tools leads, even though the gap in usage rates between sanctioned and unsanctioned tools is minor.

17.6% of respondents don't know whether employees are using unsanctioned generative AI – nearly triple the 6.3% in 2025. 'Don't know' is itself a measurable loss of visibility: organizations that can't see unsanctioned use can't enforce policy, audit exposure, or correct risky behavior before incidents occur.

Use of sanctioned and unsanctioned generative AI assistants

Percentage of employees with access to sanctioned and unsanctioned generative AI

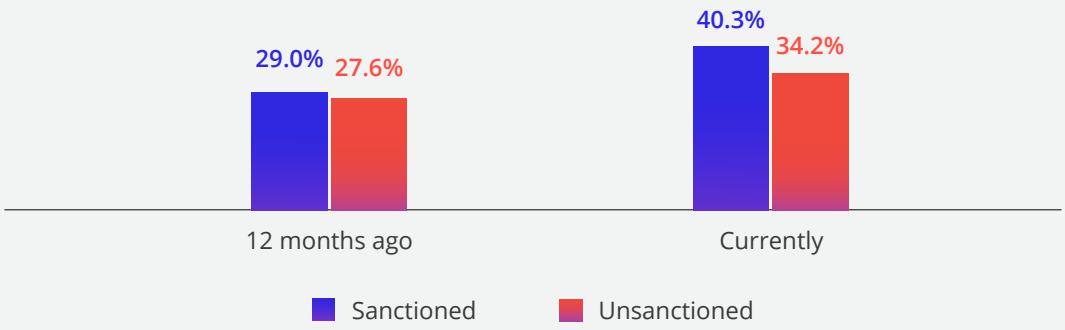

The Illusion of Progress: Why "Doing AI" Rarely Means Deploying It

- Luise Freese, Microsoft MVP

"When organizations can't see where AI is actually being used, they confuse activity with adoption — and that confusion is where the real risk lives."

Learn more

Generative AI is now part of weekly work

Half of employees use generative AI tools – either sanctioned or unsanctioned – daily or weekly for completing work tasks. Daily is the most common cadence; weekly is second.

Compared to 2025 data, usage rates per cadence are almost identical, indicating no net change in how generative AI assistants are impacting employees' work, with all but the daily cadence varying by 1% or less. The largest change is in the daily cadence, which declined by 3.3%.

Integration of generative AI tools into regular work processes

Average percentage of employees using generative AI tools per cadence

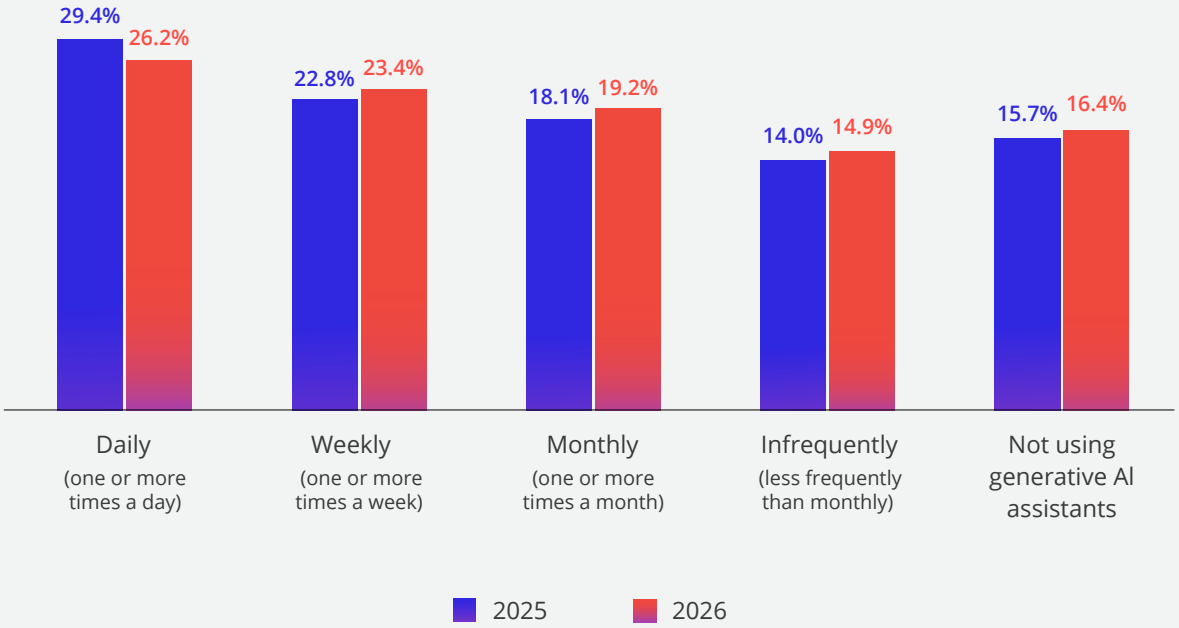

Issues

Implementation concerns are increasingly focused on readiness and scale

Data security and privacy remain the leading implementation concern for generative AI assistants, followed by integration complexity. Notably, the top four concerns are unchanged from last year, suggesting that the primary barriers to AI adoption continue to center on governance, operational complexity, and the ability to manage AI at scale rather than on the technology itself.

Concerns for generative AI implementations: 2025 vs 2026 (by 2026 severity)

Percentage of respondents indicating "highly concerned" or "extremely concerned"

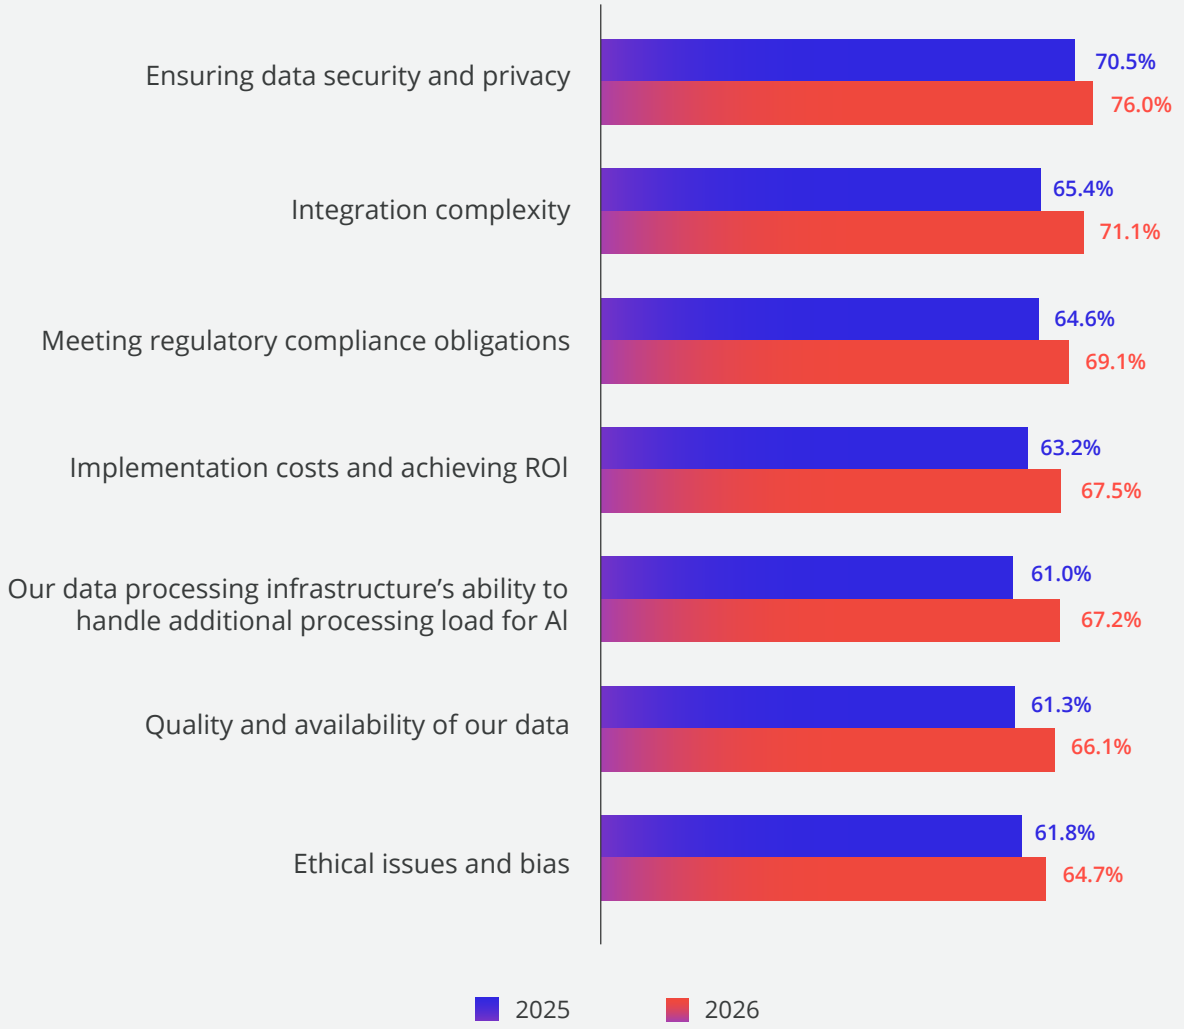

The most significant year-over-year increase was concern over whether data processing infrastructure can support the additional demands created by AI workloads. As organizations expand deployments and generate larger volumes of AI-driven data, infrastructure readiness is emerging as a prominent consideration. While additional cloud investment can address some capacity challenges, relying solely on increased spending is unlikely to be a sustainable strategy for long-term ROI without addressing data security and privacy concerns.

Concerns about generative AI producing irrelevant information and eroding employees' ability to verify AI output both rose year over year – signals of model collapse as AI is increasingly trained on AI-generated content ('AI slop'), reinforcing why data quality and lifecycle controls remain foundational to AI value.

Addressing concerns with investments in third-party tools

Organizations seeking to address concerns and shortcomings in their generative AI deployments over the next 12 months are prioritizing investments in third-party governance tools and third-party data security tools. Compared to our 2025 data, the largest increase in intent to invest more over the next 12 months is for data security and third-party backup tools.

Investment patterns for third-party tools over the next 12 months

Percentage of respondents

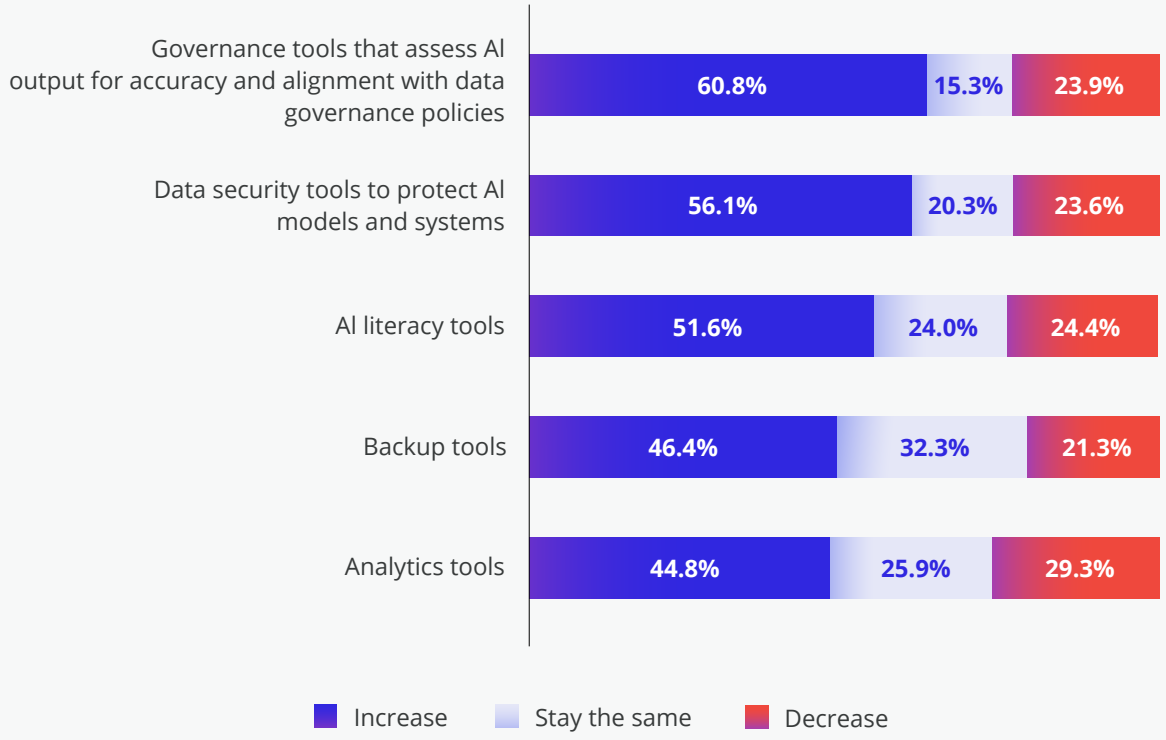

Security and data concerns continue delaying deployments

Data security and data management issues with generative AI assistants have resulted in 86.9% of organizations delaying their deployment. Year over year, more organizations are delaying across the 3-to-6 month and 7-to-12 month timeframes. The average delay increased from 5.76 months in 2025 to 5.88 months this year. This pattern indicates that AI programs are increasingly gated by readiness foundations, especially data security, governance maturity, and the ability to control and remediate AI-driven outcomes.

Delay time frame for generative AI assistants

Percentage of respondents

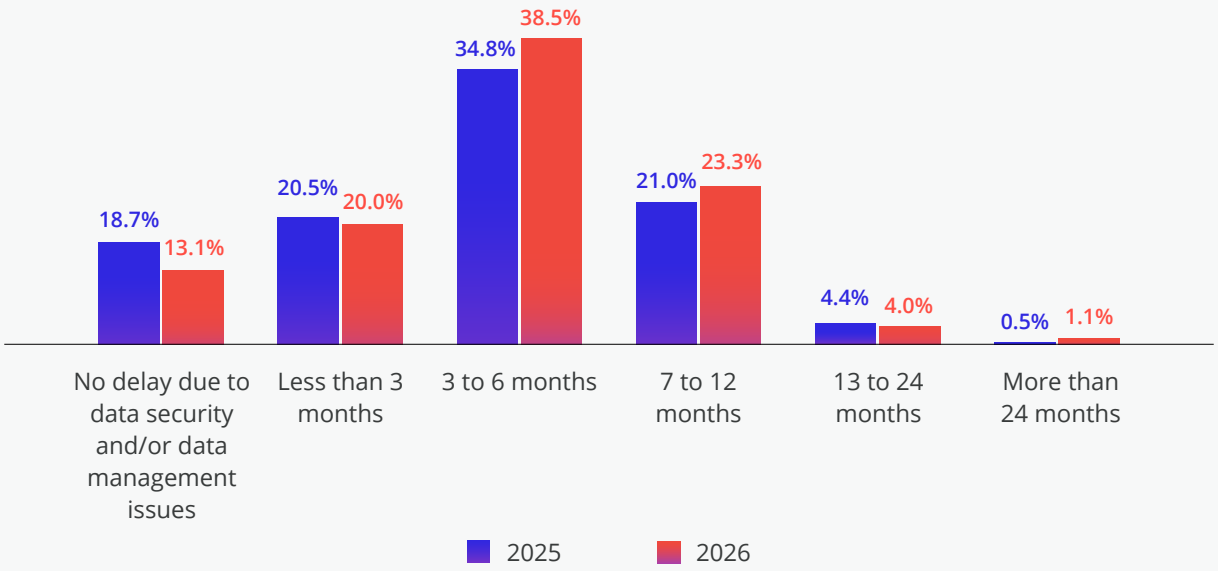

Data security and data quality concerns top the list of reasons for slowing down the rollout of generative AI assistants, with data security taking the lead over quality concerns this year. In comparison to our 2025 data, all six reasons we asked about have become more concerning. The factor that changed the most was the cost of licenses – which is intertwined with the realization of business value to justify the expenditure on licenses, running an employee adoption program, and, more critically, establishing the processes and technology to address highly concerning data security and data quality issues.

Reasons for slowing down the rollout of generative AI assistants

Percentage of respondents indicating "highly impactful" or "extremely impactful"

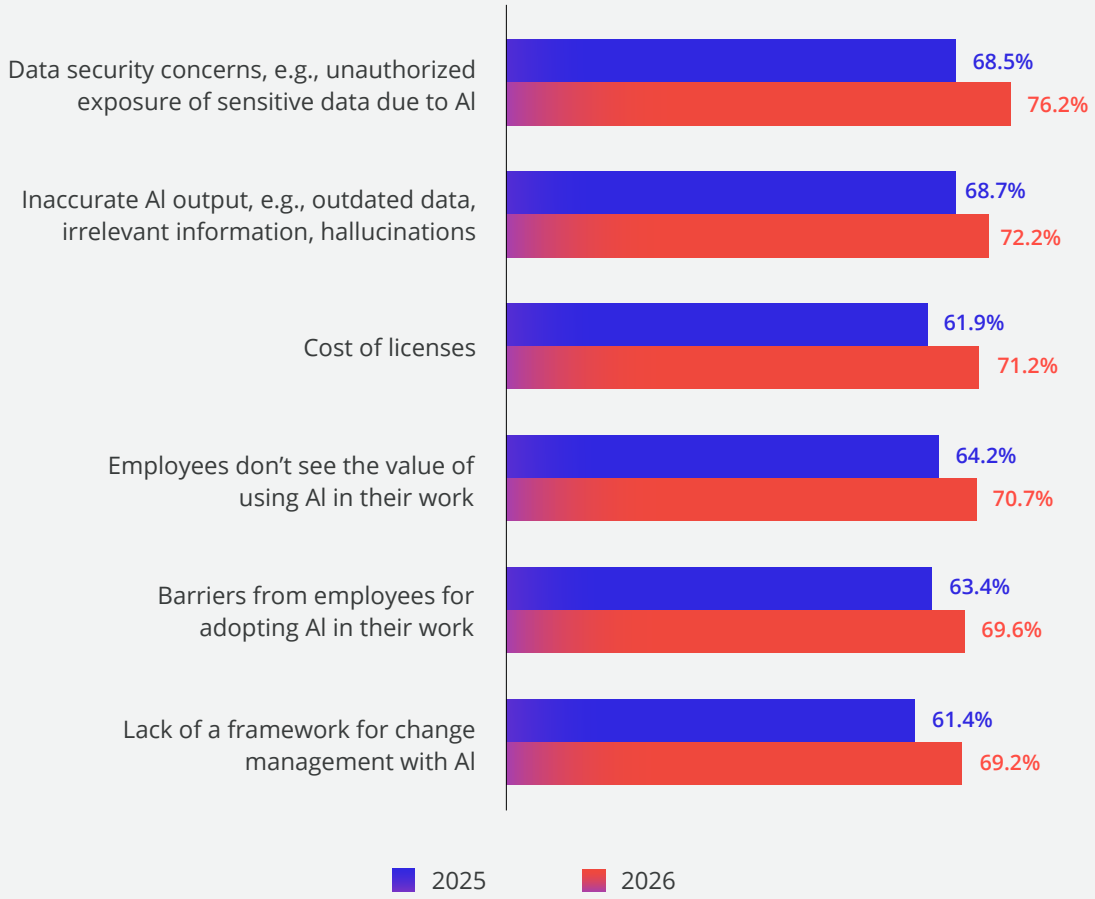

Training remains the most common response to mitigating security concerns with the use of generative AI tools (66.9% of organizations), followed by deploying third-party governance tools to assess generative AI assistants' output for accuracy and alignment with data governance policies (55.6%). The three mitigations or interventions that changed the most compared with our 2025 data are:

- **Doing nothing to mitigate security concerns**: dropping from 8.3% of organizations in 2025 to 2.5% this year. More organizations recognize they must act.
- **Delaying rollout to address security concerns**: increasing from 22.1% last year to 34.5% this year, a 56.5% uplift. More organizations are taking action to ensure the right security protections and guardrails are in place before introducing generative AI assistants to their data and employees.
- **Canceled rollout of generative AI assistants**: increasing from 31.7% to 40.7%, a 28.1% increase. When data security risks are too high in relation to the value achieved, canceling the initiative entirely is possible.

Data security incidents

89.5% of organizations experienced at least one generative AI-related security breach over the previous 12 months. Compared to our 2025 data, more organizations acknowledged experiencing more breaches – up from 75.1%.

Frequency of AI-related security breaches over the previous 12 months

Percentage of respondents who experienced one or more breaches

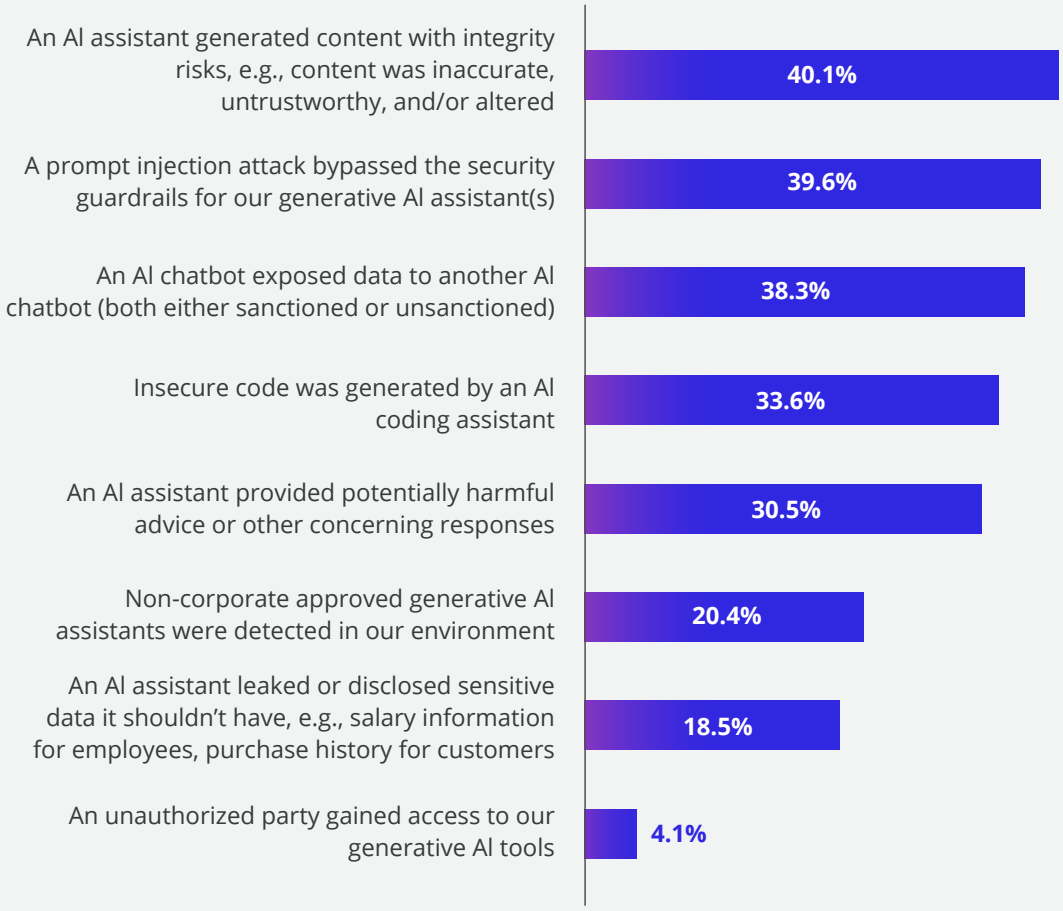

These readiness gaps become more consequential as organizations move from generative AI that produces content to AI agents that take actions.

Conclusion: Readiness, Not Models, Will Decide Who Wins with AI

The data in this report points to a clear conclusion: enterprise AI adoption is increasingly gated by readiness.

Organizations are moving quickly to deploy AI, but delays and incidents show that many are not yet prepared to govern it safely at scale. Across both generative and agentic AI, data security and data management concerns are causing widespread deployment delays. At the same time, AI-related security incidents are becoming common, even among organizations that report high confidence in their ability to prevent unauthorized access.

This is why visibility, governance, lifecycle management, and operational control are emerging as foundational requirements for enterprise AI. As organizations deploy larger numbers of agents across business processes, the challenge shifts from simply adopting AI to managing it at scale.

> ***As AI systems take on greater autonomy, the need for control becomes more urgent. Trust in AI does not come from intent, policy, or optimism alone. Trust is earned when organizations can control what AI can access, govern how it operates, audit what it does, and recover when something goes wrong. This is why visibility, governance, lifecycle management, and operational control are emerging as foundational requirements for enterprise AI. As organizations deploy larger numbers of agents across business processes, the challenge shifts from simply adopting AI to managing it at scale.***

The organizations that realize sustained value from AI will not be defined by the models they choose. Models, agents, and AI architectures will continue to evolve. They will be defined by whether they have built the durable foundations that persist across those technology shifts: trusted data, effective governance, operational visibility, enforceable controls, and resilience.

The companies that lead the next decade of AI won't be the ones with the best models. They'll be the ones with the strongest foundations beneath them – and that work starts now.

> **We needed a control layer sitting on top of our fundamental data. That foundation gives us the confidence to layer Gemini on top and execute much more advanced workflows.**
>
> — Trevor Marshall, Co-founder & CTO, Current

The result: 80–90% efficiency gains on AI-powered workloads, audit-ready visibility across the environment, and the regulated confidence to scale AI without scaling risk.

Banking on Governance: How Current Scaled Gemini to Drive 90% Efficiency

Read the full story

www.AvePoint.com | @AvePoint