#shifthappens
powered by AvePoint

The State of AI:

Go Beyond the Hype to Navigate Trust, Security and Value

AvePoint

Table of Contents

01

Introduction

INTRODUCTION

AI is no longer a future disruptor – it's a present reality.

But as adoption accelerates, trust in AI outputs is eroding. Organizations realize that without strong governance, resilient data strategies, and a commitment to quality, AI can just as easily become a liability as a competitive advantage.

According to data from 775 global business leaders across financial services, government, and healthcare, AvePoint found:

#1

AI usage fuels data security incidents.

75.1% of organizations reported at least one data security incident where oversharing sensitive information negatively impacted them.

#2

As AI becomes more popular, scrutiny increases.

Microsoft 365 Copilot has become the top sanctioned generative AI tool, and over 50% of employees have integrated the use of generative AI tools into daily and weekly work. Still, 88.3% of organizations have only rolled this out to some employees, with pilot programs focused on data security and readiness.

#3

Security and accuracy concerns stall AI rollouts up to a year.

The two main reasons AI is not more widespread are inaccurate output (68.7%) and data security concerns (68.5%). These challenges are delaying AI rollouts up to 12 months for over three-quarters of organizations.

#4

Traditional information management frameworks fall short in the AI era.

While 90.6% of organizations have what they perceive to be an effective information management program or framework in place, few have data classification and incident prevention in place. Their overconfidence will create challenges as data increases at a rapid clip.

#5

Data growth isn't just a scale issue – it's a security risk.

79.2% of organizations now manage 1 PB or more of data, up 25% from last year, and nearly 20% of organizations believe that in 12 months more than half of their data will be created by generative AI. As data sprawls across cloud platforms, organizations must deploy true information lifecycle management to reduce their risk.

#6

Inaccurate AI output threatens rollouts and erodes employee judgment.

68.7% of organizations have slowed the rollout of generative AI assistants due to inaccurate AI output, and more than two-thirds of respondents express significant concern for the impact of inaccurate outputs on employee judgment.

#7

Cultivating AI literacy is critical to the success of any AI rollout.

99.5% of organizations have used a range of interventions to strengthen AI literacy among employees.

#8

Public AI use is finally being governed.

AI Acceptable Use policies have increased by 79.8% year-over-year, with 84.5% of organizations now enforcing or developing clear AI usage guidelines.

#9

AI success is measured by customer impact – but that's where most fall short.

Enhancing customer insights and personalization is what organizations want to improve when implementing AI, yet there is a 5.8% gap between what they hope to achieve and what they actually do.

#10

AI adoption is measured by usage – not by impact.

89.6% of organizations use built-in or third-party reports to track how generative AI assistants are used and how well they perform. But far fewer look at the human or business impact using interviews or other qualitative methods.

Data & Methodology:

Key findings and data in this report comes from a 2025 global study by Osterman Research for AvePoint. Building on AvePoint's 2024 AI & Information Management Report, this 2025 study contrasts year-over-year data, revealing how enterprises are moving from AI experimentation to enterprise-wide enablement, emphasizing governance, risk mitigation, and measurable outcomes. 775 respondents with responsibility for information management, data security, or AI programs at their organization were surveyed.

JOB ROLE

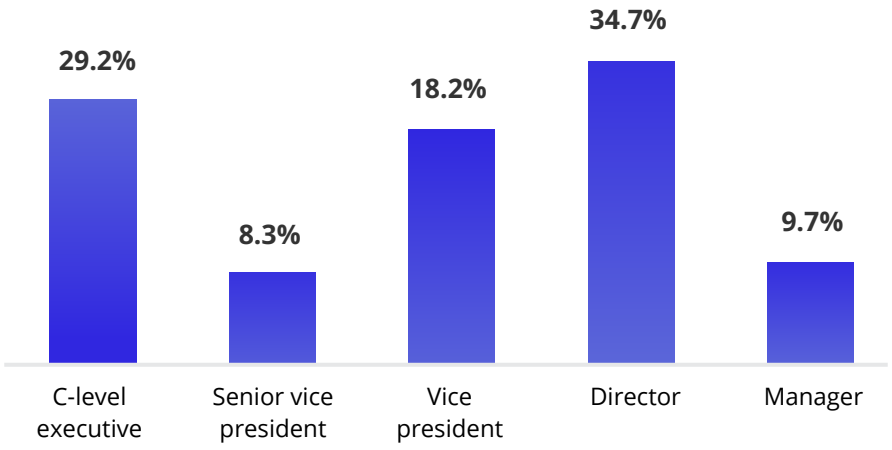

INDUSTRY

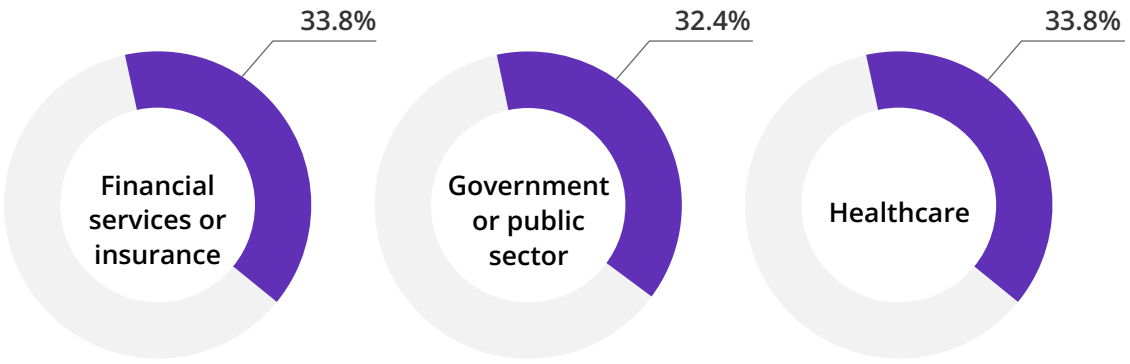

ORGANIZATION SIZE

GEOGRAPHY

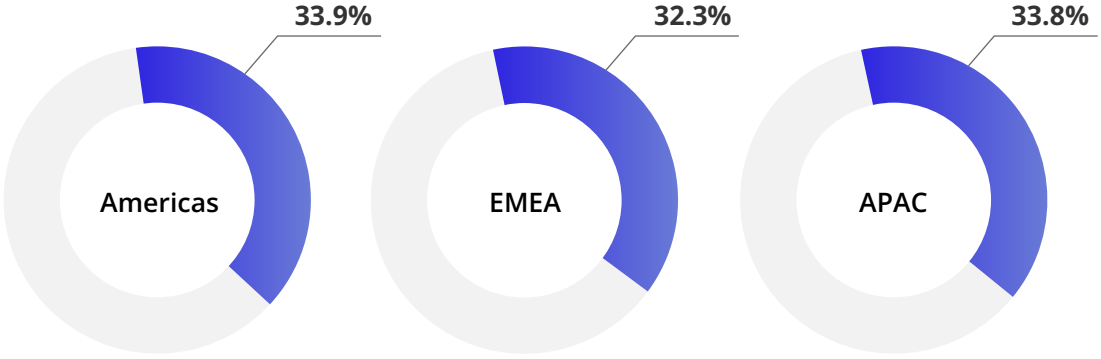

02

Current State of Enterprise AI

CURRENT STATE OF ENTERPRISE AI

Enterprise AI adoption is accelerating

Over the past year, more employees have gained access to approved generative AI tools and most now use them daily or weekly.

Microsoft 365 Copilot and ChatGPT are used at similar rates, with only a 0.4% difference. However, Copilot has seen faster growth – rising from 47% to surpass ChatGPT, which grew slightly from 65%. Approval for Google Gemini also increased, growing faster than ChatGPT but slower than Copilot. Most organizations approve two (44.3%) or three (25.1%) generative AI tools for employee use. See Figure 2.1.

Figure 2.1

Top generative AI tools sanctioned for use by organizations

Percentage of respondents

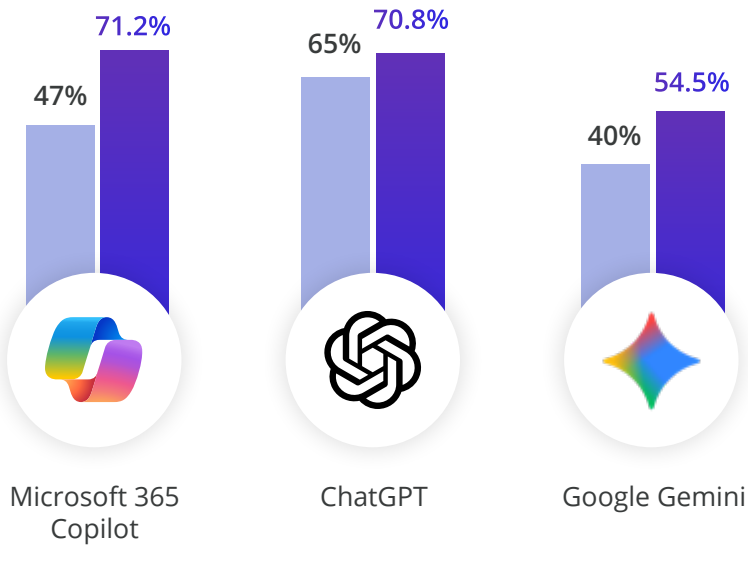

More than half
of employees have integrated generative AI tools in their work one or more times per week.

The remainder have made less progress in doing so, with an average of 18.1% of employees using generative AI tools only one or more times per month, and 14% less frequently than monthly. This means there is still significant opportunity for employees to integrate generative AI tools more deeply into their daily workflows. See Figure 2.2.

Figure 2.2

Integration of generative AI tools into regular work processes

Average percentage of employees using generative AI tools per cadence

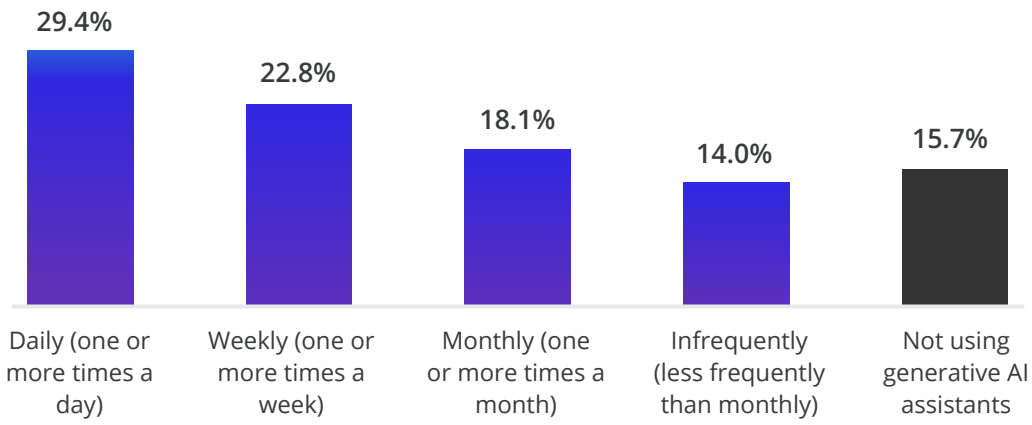

Most started with a pilot program to assess the real benefits of AI

88.3% of organizations started their AI journey with a pilot – a small-scale, short-term experiment to assess the feasibility of a wider rollout. Only 11.7% of organizations jumped immediately into production rollout without assessing feasibility in advance. On average, 37.4% of employees were included in a pilot program – although there was a wide range of approaches taken among the respondents to this research. 62.6% conducted a pre-assessment of where AI could assist employees in their work and assigned pilot licenses accordingly, while 52.9% had senior leadership specify participants. See Figures 2.3 and 2.4.

Figure 2.3

Percentage of employees involved in a pilot program for AI

Percentage of respondents

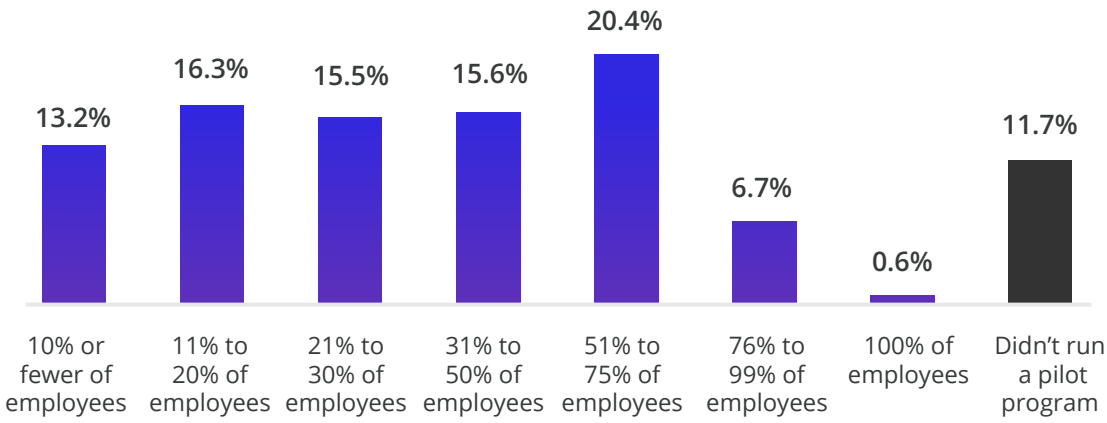

Figure 2.4

Methods for deciding which employees should participate in an AI pilot program

Percentage of respondents

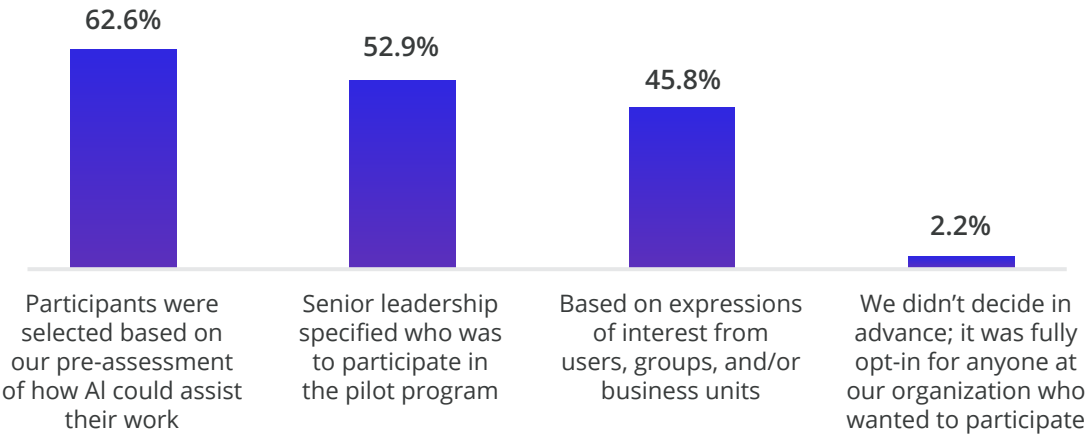

Most pilot programs included the IT department, with data teams also commonly involved. That makes sense, since these groups are responsible for evaluating data security and readiness – two key goals of AI pilot programs. Differences across industries can be seen in Figure 2.5.

Figure 2.5

Industry	Most common	Second most common	Third most common	Fourth most common	Fifth most common
Financial services and insurance organizations	IT (department, professionals)	Data analysts, scientists, administrators	Finance team or department Customer support or service	Operations (frontline staff)	Senior leaders
Government and public sector organizations	IT (department, professionals)	Data analysts, scientists, administrators	Senior leaders	Operations (frontline staff) Business unit leaders	Developers
Healthcare organizations	IT (department, professionals)	Clinical professionals (doctor, nurse, researcher, pharmacist)	Data analysts, scientists, administrators Admin staff Customer support or service	Marketing	Developers Operations (frontline staff) Finance team or department

Data concerns drive pilot program decisions

The two highest rated reasons for running a pilot program focus squarely on data:

84.4%
said testing data security risks was very or extremely important

81.3% said assessing the readiness of data. See Figure 2.6.

Figure 2.6

Reasons for running a pilot program for AI

Percentage of respondents indicating "very" or "extremely" important

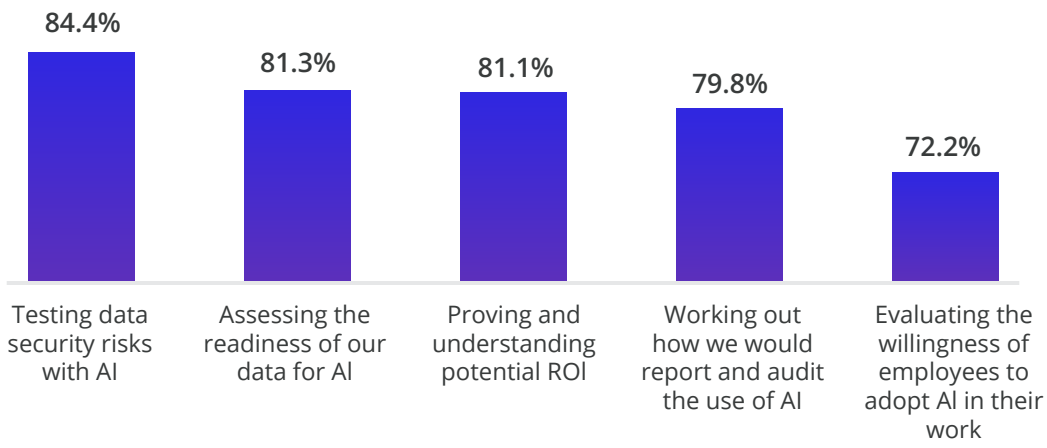

In addition to sanctioned AI, shadow AI tools are widely used

Sanctioned options such as Microsoft 365 Copilot, ChatGPT, and Google Gemini aren't the only generative AI assistants being used by employees. The percentage of employees using unsanctioned AI is growing year-over-year. See Figure 2.7.

Figure 2.7

Use of unsanctioned generative AI tools

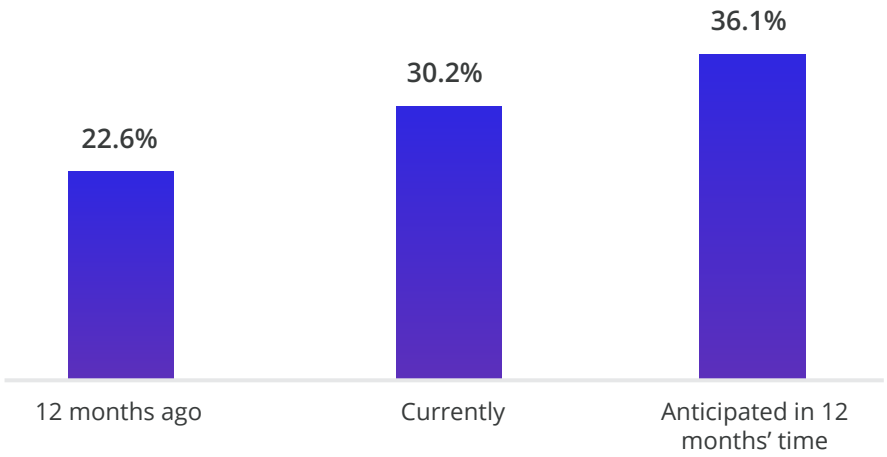

When correlating the use of unsanctioned generative AI tools by employees with the number of sanctioned generative AI tools:

- Organizations that approve only one generative AI tool are the least aware of what other tools employees are using – 18.7% admit they don't know. In contrast, for organizations that approve of 2-3 tools, only 10-13% don't know what their employees are using.
- Among organizations only sanctioning a single generative AI tool, use of unsanctioned generative AI tools by employees is about 15% lower compared to those sanctioning two or more. This is likely due to greater organizational support for the single sanctioned toolset, covering adoption and process change investments, although it could also be due to lower visibility into unsanctioned usage.

Takeaways

01

Data security is the top reason why organizations are not rolling out AI faster or to more people in their organization. It's not a question of cost or ROI.

02

AI adoption is widespread, but shadow AI use is growing which highlights a governance gap.

03

Pilots are the norm, but data readiness – not the technology – is the real test.

EXPERT PERSPECTIVE

Jeremy Thake

Chief Architect, AvePoint

Many organizations use pilot programs to find and fix data security and governance gaps, which helps them mitigate risk and adjust their policies. However, the rise of task-specific AI agents introduces new, and often unanticipated, challenges. We are no longer just concerned about oversharing data; we are now facing the potential risk of oversharing agents themselves. In fact, Gartner predicts 40% of enterprise apps will feature task-specific AI agents by 2026, up from less than 5% in 2025. Organizations need to move quicker to match the pace of innovation and the speed at which agents will transform our workplace as we know it.

EXPERT PERSPECTIVE

03

Signs of strengthening guardrails for AI

SIGNS OF STRENGTHENING GUARDRAILS FOR AI

AI Acceptable Use Policies are becoming standard practice

An AI Acceptable Use policy provides guidelines and rules for the ethical and responsible use of AI systems within an organization - 98% of organizations will have such a policy in place within the next 12 months.

When we asked this question in 2024, it was a yes or no question - with 47% of respondents saying "yes" to having a policy and 53% saying "no." In 2025, we assessed for greater nuance via five options - including a fully developed and implemented policy and a policy in active development. See Figure 3.1.

Figure 3.1

Implementation status of an AI Acceptable Use policy

Percentage of respondents

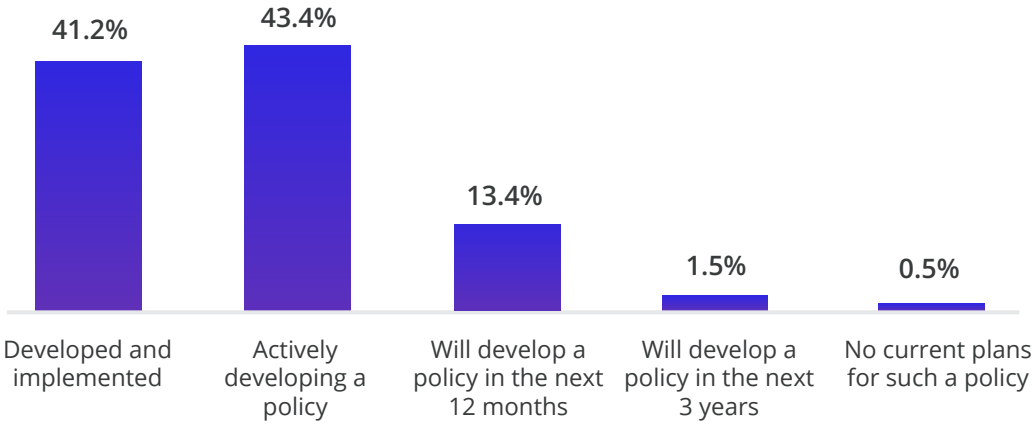

It is possible to interpret the year-over-year data in one of two ways. First, the safeguard of an acceptable use policy for AI has declined, assuming that the 47% in 2024 is the same as the 41.2% in 2025 that say they have a fully developed and implemented policy. From the perspective of comparing two different assessment methodologies, this is unlikely to be the case because the hard lines of a "yes" or "no" answer push respondents in the middle (e.g., currently developing a policy) to one or other extreme.

This leads to the second and likely more accurate interpretation, that the safeguard of an acceptable use policy for AI has increased, assuming that the 47% in 2024 aligns with the 84.5% that have a policy already (41.2%) or are currently developing one (43.4%). This is a year-over-year increase of 79.8%.

Over time, organizations are likely to cycle continuously between having a developed and implemented policy for AI and actively developing one. Especially now with the growth of agentic AI, AI policies will need to change in lockstep - and technology can help. Tools can monitor usage and raise risks because we all know just because a documented policy exists does not mean everyone follows it (as evidenced by the use of unsanctioned AI tools, for example).

AI guidelines are a team sport

Most believe in shared responsibility for developing guidelines for the safe and effective usage of AI. The highest responsibility sits with the organization itself (93.2% indicated "high" or "extreme" responsibility), trailed by industry coalitions (77.8%) and government (68.4%). See Figure 3.2.

Figure 3.2

Responsibility for developing guidelines for safe and effective usage of AI

Percentage of respondents indicating "high" or "extreme" responsibility

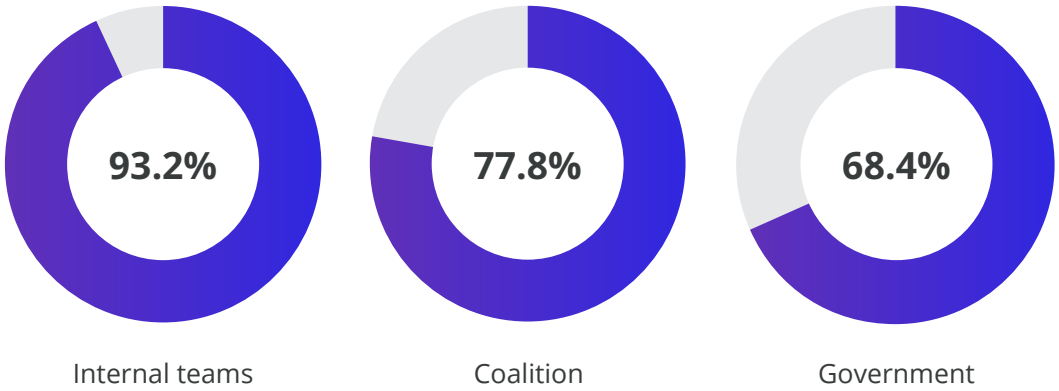

How AI responsibilities vary by region

Across the three regions represented in this survey, internal teams are always viewed as holding primary responsibility for developing guidelines for safe and effective usage of AI. Both EMEA and APAC view an industry coalition as having greater weight in the development process than the Americas, and the importance of government input in the APAC region is higher than in the other two. See Figure 3.3.

Figure 3.3

Responsibility for developing guidelines for safe and effective usage of AI - by region

Percentage of respondents indicating "high" or "extreme" responsibility

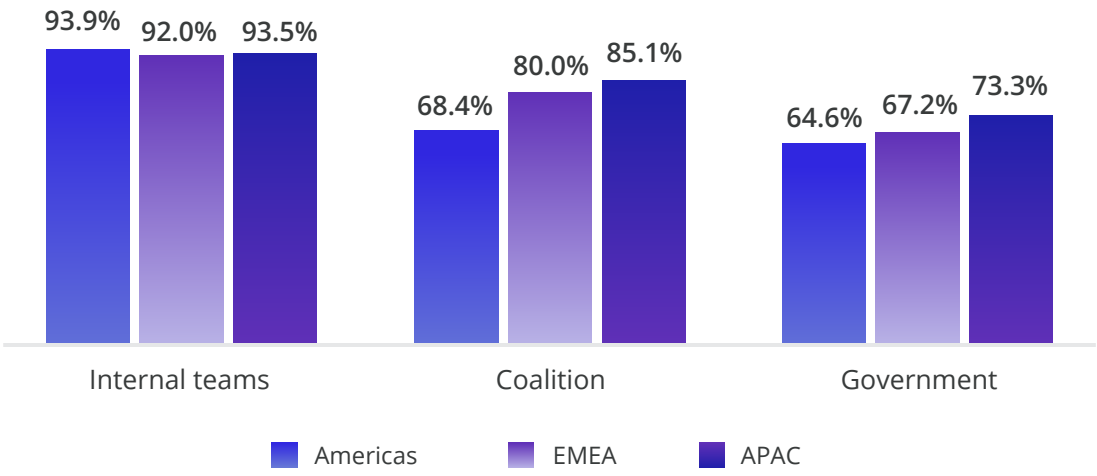

Responsibility for AI guidelines vary by industry

When correlated by the three industries we investigated in this research, internal teams still come out with primary responsibility for developing safe and effective usage guidelines for AI. Across the other two possible groups – an industry coalition and government – respondents from government or public sector organizations attribute higher responsibility to these two groups. Financial services or insurance organizations are the least likely to want responsibility sitting with an industry coalition or government bodies. See Figure 3.4.

Figure 3.4

Responsibility by industry for developing guidelines for safe and effective AI usage

Percentage of respondents indicating "high" or "extreme" responsibility

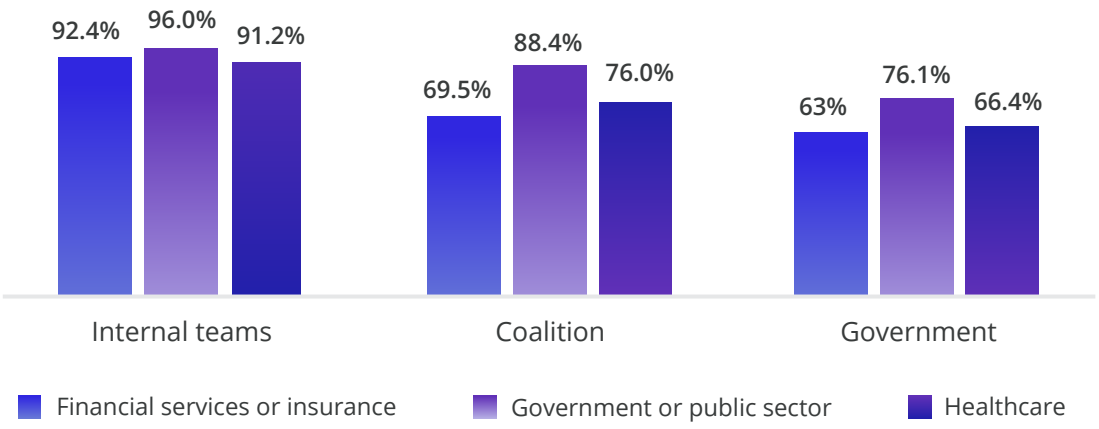

Takeaways

01

AI policies must evolve continuously: Organizations recognize that AI policies are not static documents. They require ongoing iteration to keep pace with technological change.

02

Governance is a shared responsibility, but organizations lead the way. Internal teams hold the greatest accountability, though industry coalitions and governments play rising roles, especially in the public sector and APAC

04

Implementation status and impediments

IMPLEMENTATION STATUS AND IMPEDIMENTS

Access to corporate-approved generative AI assistants is increasing year-over-year

More employees year-over-year are being given access to corporate-approved generative AI assistants, with 40% currently having access and the expectation of this increasing to 54.6% in 12 months' time. See Figure 4.1.

Figure 4.1

Access to corporate-approved generative AI assistants

Average percentage of employees having access

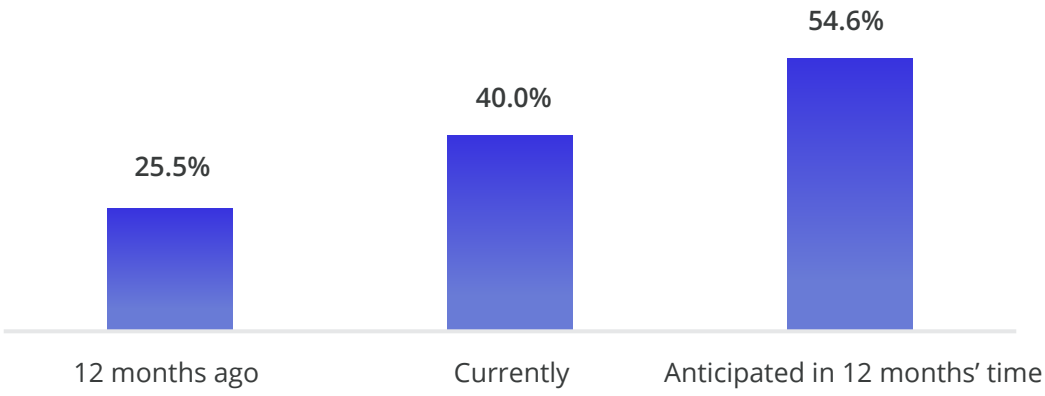

In looking at the data:

- 81.3% of organizations see year-over-year increases in the percentage of employees having access to corporate-approved generative AI assistants.
- Among organizations that have completed their production rollout of AI (10% of organizations surveyed), the average percentage of employees with access to corporate-approved generative AI assistants is 43.7% this year. Only one organization in this group has made generative AI assistants available to 100% of employees, indicating that for most deployments, a "completed rollout" does not currently mean 100% employee coverage.
- Among organizations currently rolling out AI to everyone in their organization, the average percentage of employees with access to generative AI assistants at this point is 45.1%, although this is anticipated to grow to 61% in 12 months' time. While less than 1% of this group have made generative AI assistants available to 100% of employees already, this is anticipated to grow to 8.7% in 12 months.

Irrespective of increasing adoption, significant data concerns remain

Most organizations (85.7%) are hitting the brakes on rolling out generative AI tools because of two main problems: **bad data** and **data security worries**.

1. Inaccurate AI output due to outdated data, irrelevant data, and hallucinations is biggest risk (68.7%). This reinforces the fundamental argument we made in our 2024 research – that without strong governance and information management disciplines, AI would be hampered from the start.
2. Data security concerns like unauthorized exposure of sensitive data due to AI are the second highest concern (68.5%).

Of least concern is the lack of a framework for change management with AI (61.4%), which means that more organizations think they know how to manage the people side of AI-enabled change but are more significantly hindered by the technical considerations of outdated and ill-protected data. See Figure 4.2.

Figure 4.2

Reasons for slowing down the rollout of generative AI assistants (n=664)

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

In comparison to the list of concerns with AI implementation in our 2024 research, concerns around data privacy and security have declined slightly (from 71% last year), while concerns due to the quality and categorization of internal data has increased significantly (from 61%).

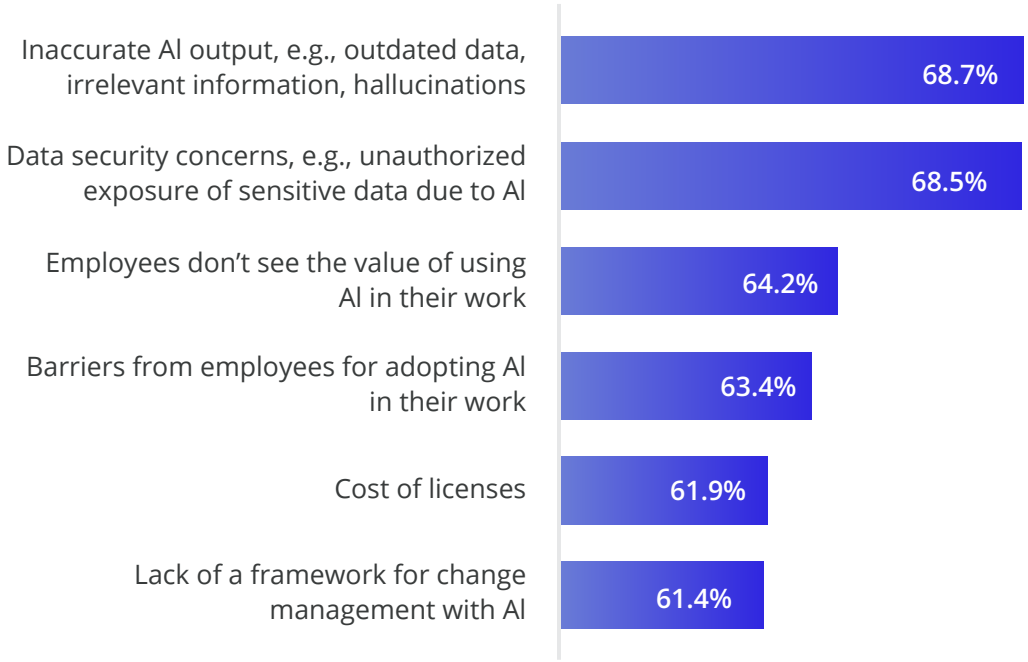

Takeaways

01

Most organizations are in production with generative AI, but few have full employee coverage.

02

Data quality and security concerns – not change management – are the biggest blockers to rolling out generative AI assistants.

03

85.7% of organizations slowed generative AI rollouts due to data quality and data security concerns.

05

Addressing data security and data management issues with generative AI

ADDRESSING DATA SECURITY AND DATA MANAGEMENT ISSUES WITH GENERATIVE AI

Most rollouts delayed up to a year due to data security and data management issues

81.3% of organizations delayed their deployment of generative AI assistants due to data security and/or data management issues. Three to six months was the most common time frame for a delay to address such issues (34.8%), followed by an almost equal split between three months or less (20.5%) and seven to 12 months (21%). The average delay was 5.8 months. See Figure 5.1.

Figure 5.1

Delay time frame for generative AI assistants

Percentage of respondents

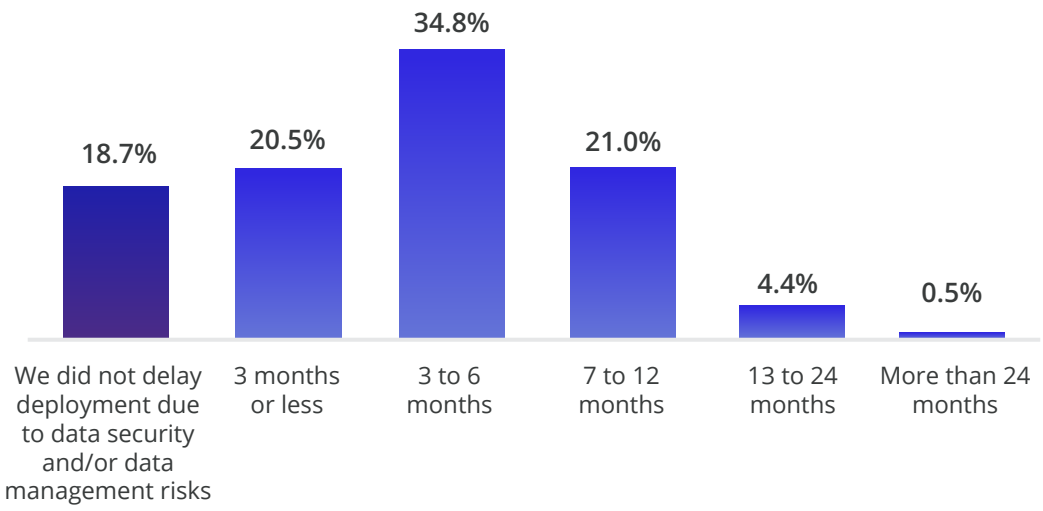

For organizations that delayed rollouts, 67% provided training for employees on how to safely use generative AI assistants in their work and 55% deployed third-party governance tools to assess generative AI assistants' output for accuracy and alignment with data governance policies.

AI-related security breaches

75.1%
of organizations experienced one or more AI-related security breaches over the previous 12 months.

Breaches attributed to external threat actors were most common, affecting both sensitive employee data (occurring at 65.8% of organizations that were breached) and sensitive customer data (58.2%). Following close behind were these same breach types caused by employees. See Figure 5.2.

Figure 5.2

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

Percentage of respondents who experienced one or more breaches (n=582)

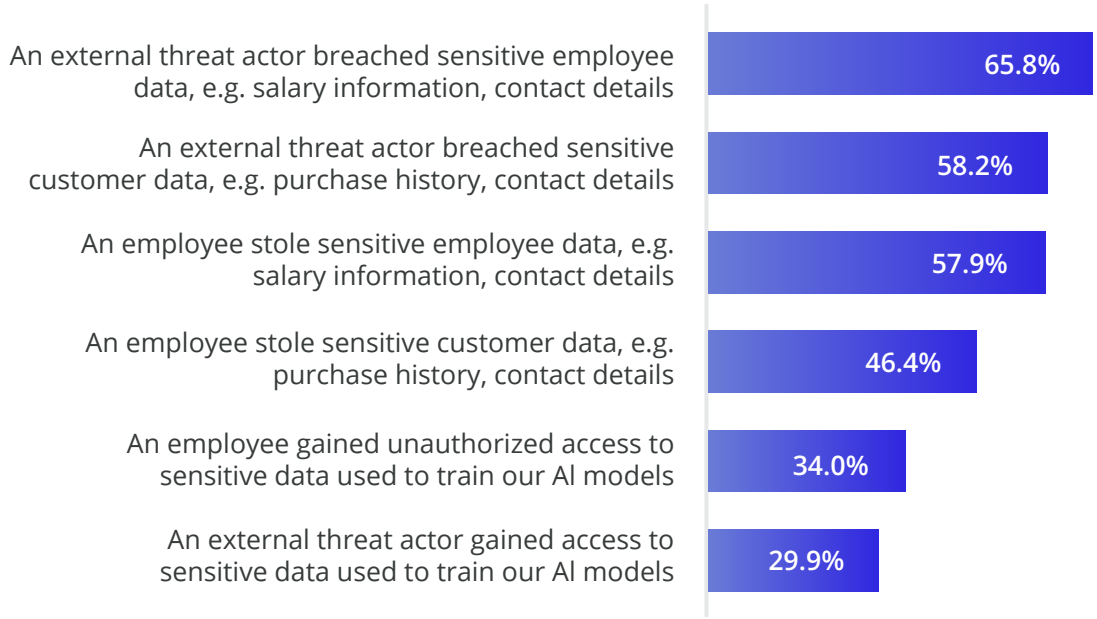

Increased investments in third-party tools planned

An average of 52.7% of organizations plan on increasing their investment level across a range of third-party tools, with third-party governance tools to assess AI output for accuracy and alignment with data governance policies in leading the way (64.4%). Third-party data security tools are in second place (54.5%). See Figure 5.3.

Figure 5.3

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

Percentage of respondents

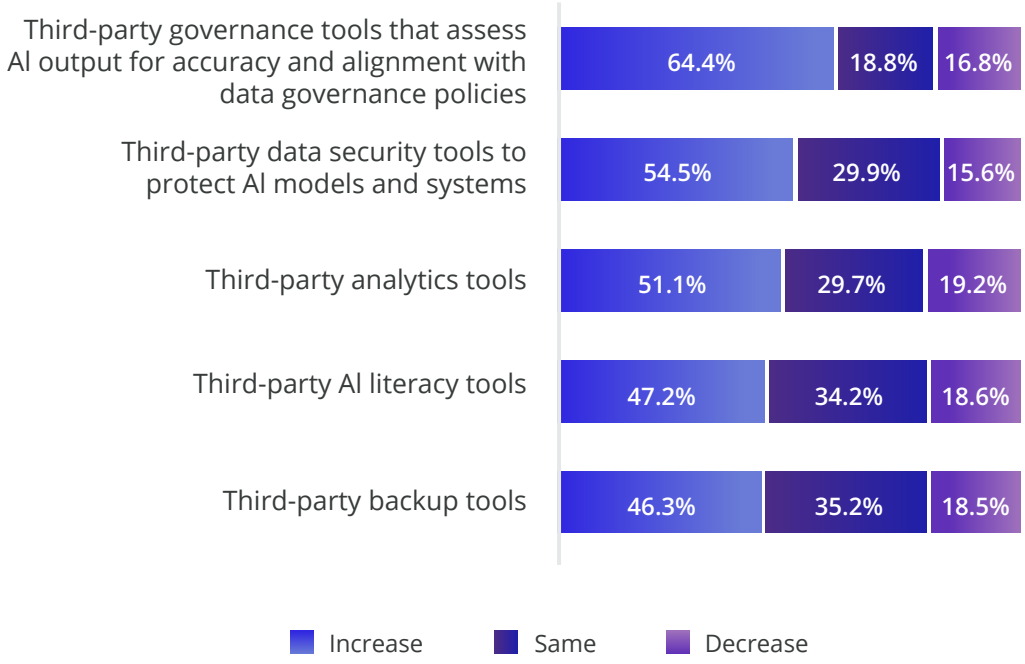

Takeaways

01

Security concerns are delaying AI rollouts. 81.3% of organizations delayed deployment due to data risks, by an average of nearly six months.

02

AI-related breaches are widespread: 75.1% of organizations experienced at least one AI-related security breach in the past year, underscoring the urgency of proactive governance.

03

With 85% relying solely on native security tools, many organizations are vulnerable to rising data incidents – prompting a surge in investment in third-party governance and security tools.

EXPERT PERSPECTIVE

Dana Simberkoff

Chief Risk, Privacy and Information Security Officer, AvePoint

Organizations today face new data security challenges that native tools alone can't address. It's crucial to move beyond theoretical security policies and focus on operational implementation. This means every policy must be backed by clear procedures, technical controls, continuous monitoring, and strict enforcement. Moreover, the effectiveness of these controls must be constantly evaluated and updated as the AI landscape evolves. Our research states that 75% have had security incidents from AI oversharing, and IBM reports that 97% of organizations affected by these types of incidents lack internal access controls. Security incidents like data breaches risk both internal and client data, leading to serious financial, reputational, and legal issues. Everyone must take responsibility for strengthening AI-related data security and governance.

EXPERT PERSPECTIVE

06

Information management strategies

INFORMATION MANAGEMENT STRATEGIES

Most have an information management strategy they believe is delivering value

90.6% of the organizations in this research indicated they have an information management program or framework in place, up slightly from 88% in our 2024 research. When asked about the capability of their current program or framework to deliver against three key outcomes, most gave high marks across the board. See Figure 6.1.

Figure 6.1

Effectiveness of information management programs at achieving outcomes

Percentage of respondents

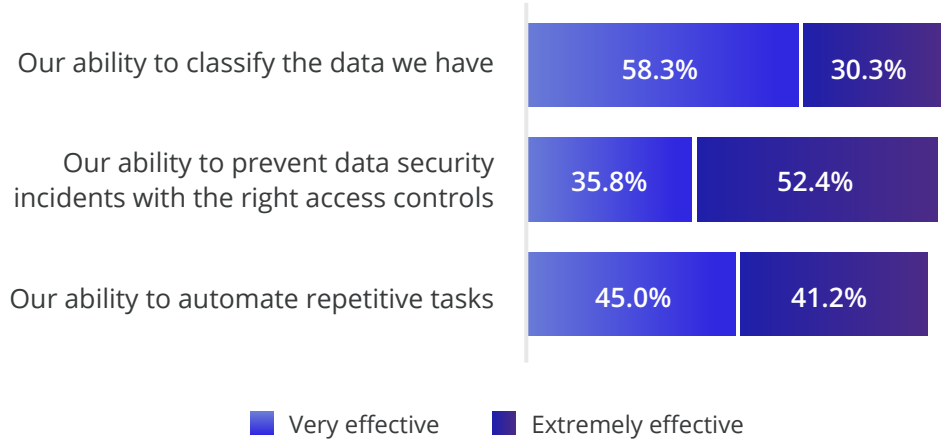

For the purposes of this study, information management (IM) strategy refers to a comprehensive approach to managing all aspects of information within an organization. This involves collecting, organizing, storing, preserving, retaining, and disposing information while maintaining control over its structure, processing, and delivery, both in electronic and physical formats, throughout its lifecycle.

Two significant numbers in Figure 6.1:

1. **For the ability to classify current data, only 30.3% claimed the highest level of effectiveness.** Data classification is a fundamental data discipline on which all other aspects of information management depend, meaning their confidence in their programs is overstated.

2. **For the ability to prevent data security incidents, 52.4% claimed the highest level of effectiveness for their program.** However, when correlated with data security incidents experienced during the previous 12 months, 77.2% of those claiming the highest level of maturity experienced internal (4.6%), external (7.6%), or both internal and external (64.9%) types of data security incidents, such as the theft of sensitive customer or employee data. There is an overconfidence in how access controls are currently being used to prevent data security incidents among the organizations in this research.

Opportunity to improve the maturity of information management tools and technologies

Underpinning the presence of an information management program or framework is a set of tools and technologies including areas such as data discovery, archiving, retention, lifecycle data management, and automation. Among the organizations in this research, an average of 41.8% claim the highest maturity rating across the five tools and technologies we asked about. See Figure 6.2.

Figure 6.2

Overconfidence in information management qualities

Percentage of respondents

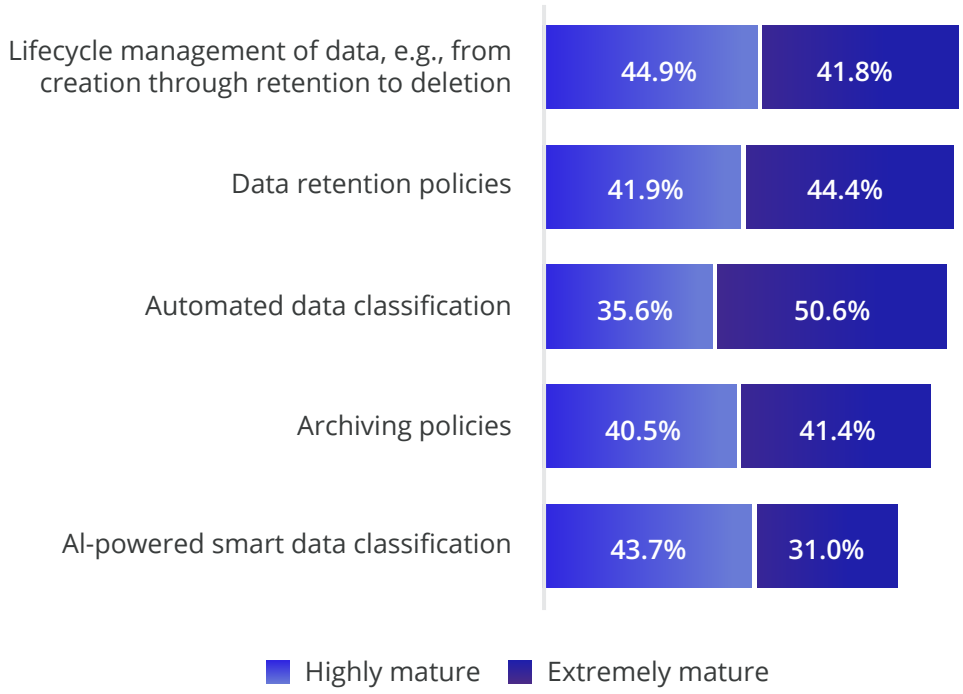

The highest isolated rating for maturity in Figure 6.2 was awarded to automated data classification (50.6%). A newer related technology – AI-powered smart data classification – is already being given a 31% extremely mature rating, despite its newness to market. The dichotomy, however, is that despite the high rating for the maturity of automated data classification, only 30.3% of respondents are fully happy with how their current approach actually classifies their data. Being able to understand what data an organization has is the fundamental data discipline on which all other aspects of information management rest, such as archiving, retention, and data security. Ultimately, there is an overconfidence in information management strategies that are underdelivering on core elements.

Higher data volume means higher risk

Most organizations (97%) with 6 or more exabytes of data under management have an information management strategy. In contrast, organizations with less than 1 petabyte of data are less likely to have an information management strategy (87%). See Figure 6.3.

Figure 6.3

Frequency and count of types of data security incidents - by data under management

Percentage of respondents

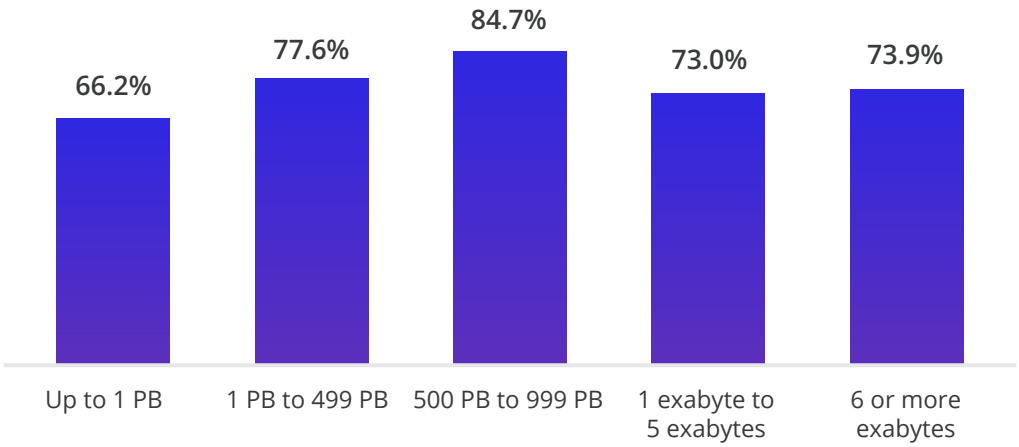

Takeaways

01

Most organizations have an IM strategy, but few excel at it. While 90.6% have an IM strategy, only 30.3% rate their data classification as highly effective.

02

Tool maturity and performance aren't the same. For example, despite high maturity ratings for automated classification tools, many organizations still struggle with actual data classification outcomes.

03

Data volume correlates with risk. Organizations managing 500-999 PB of data had the highest rate of security incidents, highlighting the need for scalable IM frameworks.

EXPERT PERSPECTIVE

Alyssa Blackburn

Program Manager, Information Management, AvePoint

Most information management strategies haven't kept pace with today's complex challenges. In fact, fewer than one-third of organizations believe their data classification is effective. This says to me that the 90% of organizations with an IM strategy in place have policies and frameworks that aren't delivering real value. To address this, organizations should adopt AI-driven tactical strategies that are smarter and more efficient. For example, as data volumes continue to grow, there is simply no choice but to employ AI driven classification systems. An AI-powered IM strategy enables better decision-making at scale - and those that fail to adapt risk falling behind.

EXPERT PERSPECTIVE

07

That’s a whole lot of data

THAT'S A WHOLE LOT OF DATA

The growth in data volumes isn't changing, and neither are the consequences

The story is the same in this year's research as last year: the sheer volume of data managed by organizations continues to grow. In last year's research, 64% managed at least 1 PB of data and 41% had 500 PB or more. Those numbers have dialed up this year, with 79.2% managing at least 1 PB and 51.6% more than 500 PB. **Year-over-year, that's a 25% increase in data volume.** See Figure 7.1.

Figure 7.1

Amount of data organizations manage

Percentage of respondents

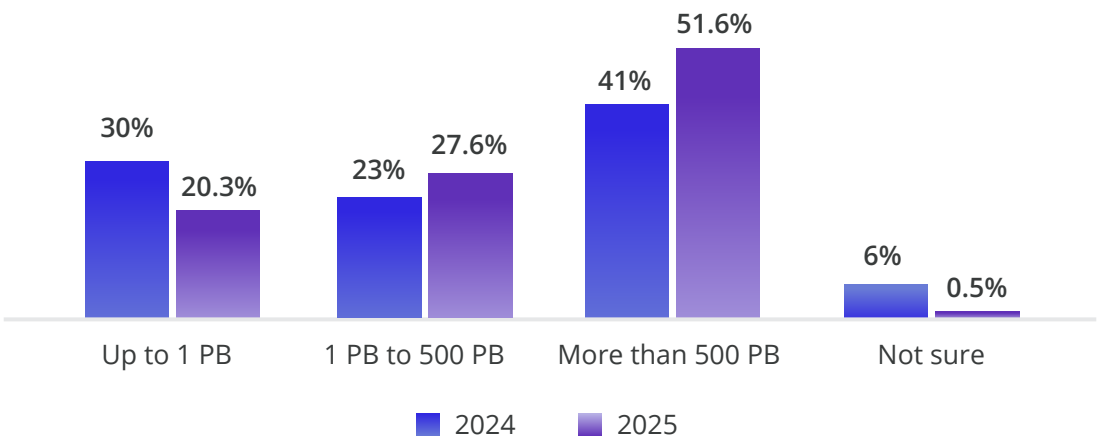

For the purposes of this study, data refers to any information that is relevant to respondents' businesses and its operations. This could include – though is not limited to – customer data, financial records, analytics data, product information, marketing data, inventory or supply chain data, employee information, or other types of relevant data. Data may be structured or unstructured and could be stored in a company's physical or digital information system in various formats, including physical documents, spreadsheets, databases, or cloud-based storage systems.

Much of this increase is driven by the creation of new data but without the corresponding removal of old data. In our 2024 research, respondents said that 50% of their organizational data was more than 5 years old.

> *This year, 70.7% of respondents say that at least half of their data is more than 5 years old.*

More is being added without removing what is increasingly redundant, obsolete, or trivial.

From whatever perspective you look at this data reality, significant consequences flow:

- For information management, the relentless increase in data volumes reinforces the criticality of underlying disciplines to manage that data, with data discovery and classification essential for retention, archiving, and lifecycle management. Data classification should also drive the automated removal of redundant, obsolete, and trivial data.
- For data security, understanding the scope and breadth of data being produced and stored is just as essential to ensure that only the right people have access, unprotected data is rapidly mitigated, and data breaches rendered unlikely.
- For AI, training internal models on outdated, redundant, or low quality data will result in outputs, AI-driven insights, and recommended decisions that likely diverge from the best of what the organization has to offer. If this happens for too long, the core competitive dynamics of the organization will be destroyed by the very technology that was supposed to lift it to the next level.

Complexity of storage location

- Only 15.4% of the organizations in this research store data in a single location, with 90.8% of these using one of the three public cloud services we asked about (i.e., Microsoft Azure, Amazon Web Services, or Google Cloud Platform).
- The remainder – the vast majority (84.6%) - use multiple platforms in parallel. Whether for reasons of redundancy, resilience, feature access - or due to legacy tech debt - the average number of storage locations is 2.8, with 35.7% using two and 31.2% using three storage locations. See Figure 7.2

Figure 7.2

Frequency and count of types of data security incidents - by data under management

Percentage of respondents

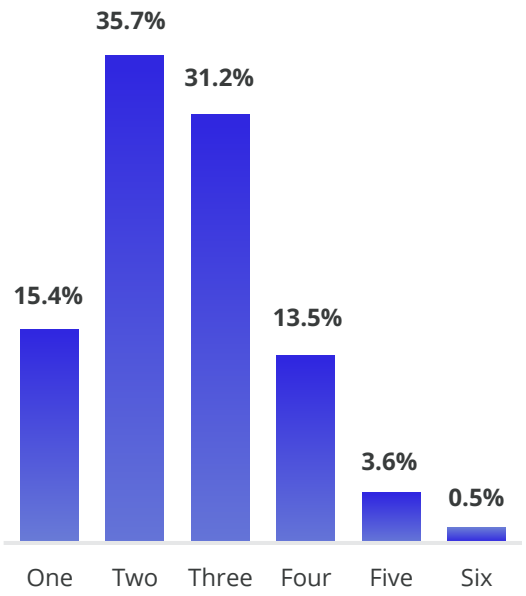

As with growing data volumes, organizations face an array of consequences of using multiple platforms. For instance:

- How to ensure multi-cloud and hybrid strategies for data discovery, collection, archiving, and retention are managed effectively.
- How to apply consistent security policies to data to minimize breaches.
- How to use data for AI when it's stored across a multi-cloud and hybrid storage environments.

Organizations anticipate even higher data growth in the next year

We calculated a 25% year-over-year growth rate in data over the past 12 months, based on how data storage volumes changed between our research study in 2024 and this year (see Figure 7.3). Based solely on the data we collected this year, that's very close to what respondents said they've actually seen – with an average data volume growth rate of 23.8% over the past 12 months lifting to an average of 31.6% anticipated over the upcoming 12 months.

Figure 7.3

Actual and anticipated growth in data created or received per year

Percentage of respondents

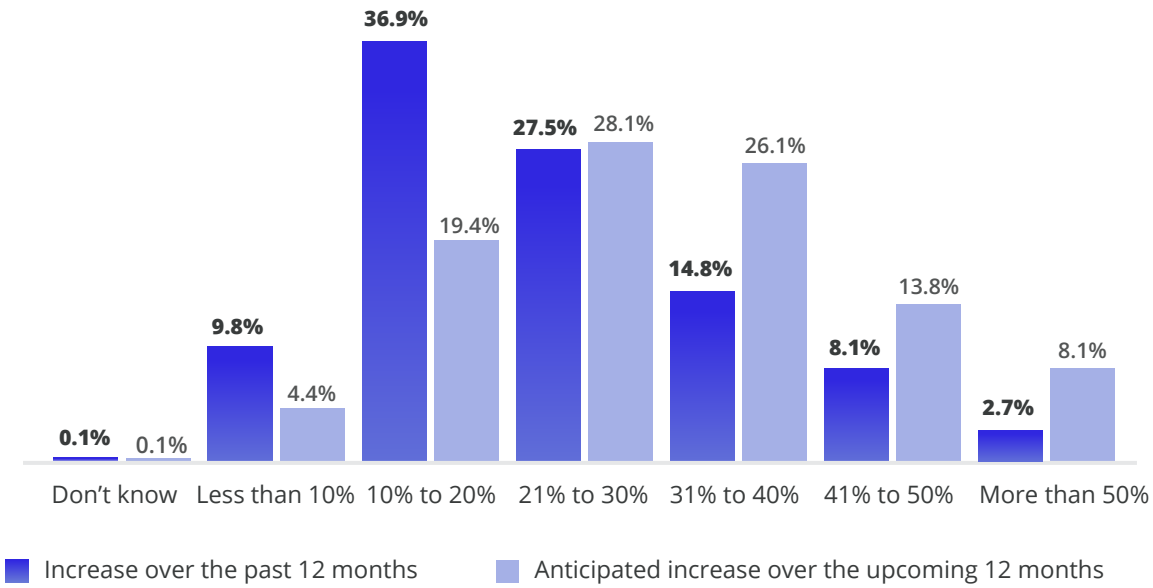

Generative AI assistants contribute a significant share of growing data volumes

Data growth due to generative AI assistants is expected to nearly double over a two-year period, increasing from 22.5% to 40%. The organizations in this research say that it will increase from an average of 22.5% of data created 12 months ago to an anticipated 40% of data created in 12 months. See Figure 7.4.

Figure 7.4

Data created by generative AI assistants

Percentage of respondents

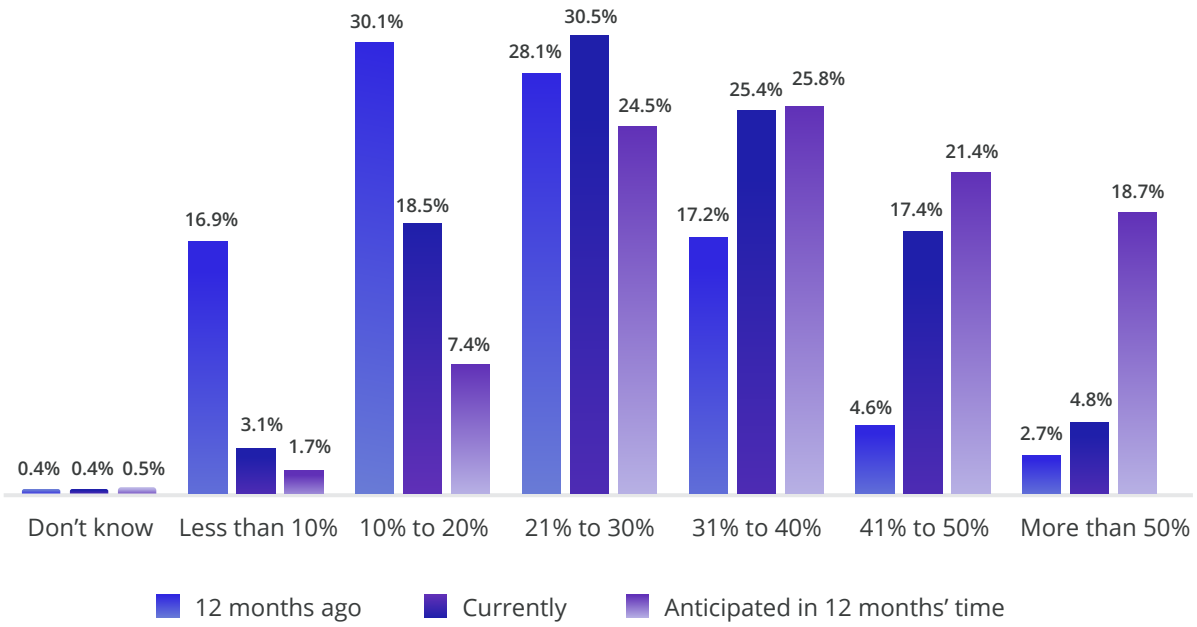

Key generative AI concerns can't be addressed by merely increasing cloud spend

This substantial growth in data volumes due to generative AI assistants comes with many implementation concerns indicating that more is not necessarily better. In first place is how to ensure data security and privacy, followed by integration complexity, meeting regulatory compliance obligations, and implementation costs. Of much lower concern is the ability of the data processing

infrastructure to handle additional processing, but this is a problem in the cloud space that is "solved" easily by increasing cloud spend. The higher ranked concerns require more intentional engagement and strategic interventions. See Figure 7.5.

Figure 7.5

Concerns about implementing generative AI assistants

Percentage of respondents

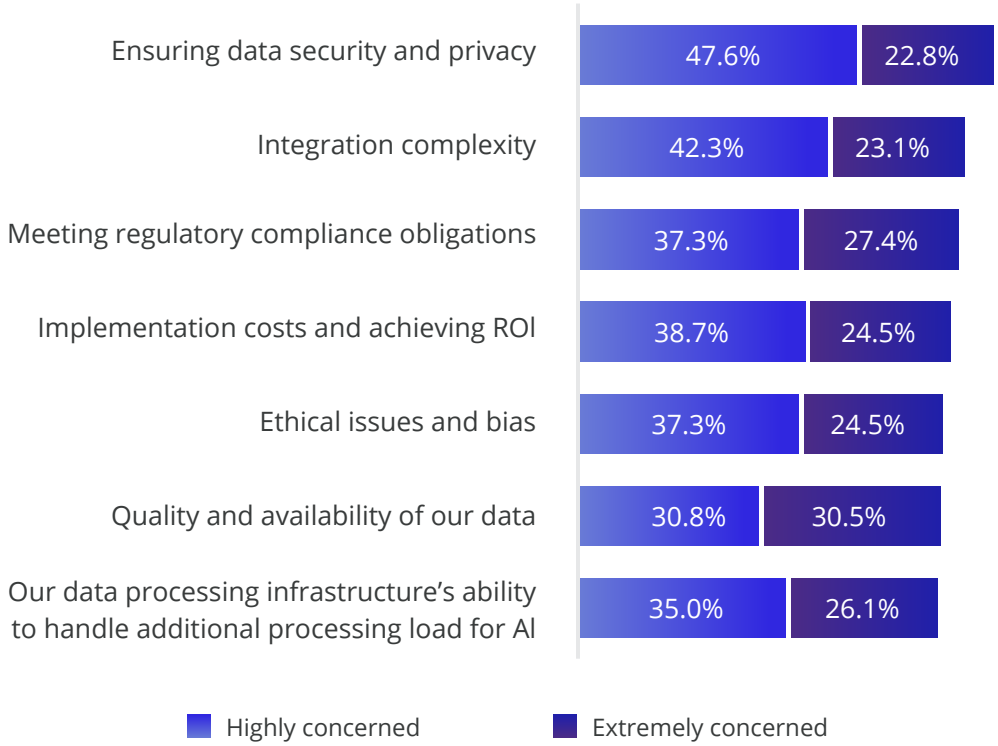

Takeaways

01

AI is a major driver of data growth, but increasing cloud spend alone can't solve data governance challenges.

02

Multi-cloud complexity demands unified information management and security strategies.

03

Data volumes are up 25% YoY, with generative AI expected to account for 40% of new data next year.

08

Data quality concerns with generative AI hinder value

DATA QUALITY CONCERNS WITH GENERATIVE AI HINDER VALUE

AI rollouts stall and incorrect AI outputs erode human judgment

Earlier, we shared that most organizations (85.7%) are hitting the brakes on rolling out generative AI tools because of two main problems: **bad data** and **data security worries**.

This same red flag is echoed in how respondents weigh their concern with actually using generative AI assistants once implemented.

The highest concern overall (66.8%) is employees losing the ability to differentiate truth from fiction, in other words, the loss of judgment and the hard-won ability to sense when data is not right due to deep engagement.

When it comes to extreme concern, 32.5% of respondents are worried about generative AI assistants hallucinating and producing incorrect outputs. This is more than the 17.3% who are extremely concerned about irrelevant information, illuminating the perception that it is easier to tell when something is not relevant but harder to confirm when it's wrong. See Figure 8.1.

Figure 8.1

Concern about threats from using generative AI assistants

Percentage of employees

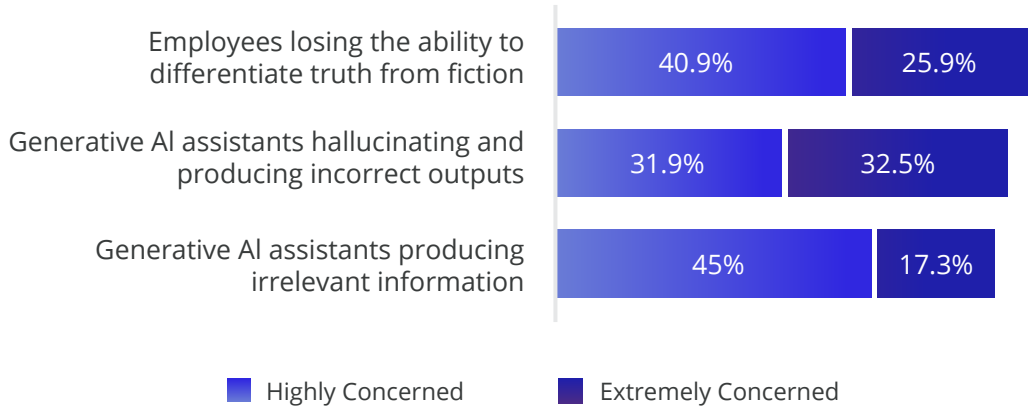

If organizations implement generative AI within a culture of "more and faster" with no counterbalancing emphasis on "better quality," then employees will be incentivized to use generative AI outputs to complete discrete work tasks faster. Given the growing time and effort gap between using the suggested output from a generative AI tool versus undertaking further manual analysis or research to confirm accuracy, there will be increasing reliance on the quick answer – irrespective of quality and accuracy. As this cycle of dependence on generative AI to complete work tasks faster deepens, employees will become unquestionably reliant on it for their compensation and job security, leading to the top-rated concern in Figure 8.1 becoming a reality.

LLMs are primarily trained using retrieval augmented generation (RAG), which uses an organization's internal data to augment prompt engineering or fine tuning, which retrains a model on a focused set of external data to improve performance. Consequently, there is inherent risk of them having bad data or not enough data which leads to inaccuracy and hallucination. What can help?

- Continual assessment of the quality of data being used for AI models. Especially with such a high percentage of current data being more than five years old, information management is critical.
- Increased employee literacy to differentiate between fact and fiction and to improve AI prompting.

Takeaways

01

The highest concern overall (66.8%) from organizations is that their employees cannot differentiate between factual and fictitious AI outputs.

02

45% of organizations are highly concerned with irrelevant information which can be improved with better data quality and better prompt construction by users.

EXPERT PERSPECTIVE

John Peluso

Chief Technology Officer, AvePoint

Consistent inaccuracies in AI output might feel like a superficial glitch, but in fact, such issues are usually a symptom of a deeper problem with the way your organization's data is structured, governed, and secured. If you fail to correct these problems before rolling out AI, you're simply going to amplify the risks posed by these defects, while also creating a mass disinformation machine that poses novel risks of its own. To avoid this, create a strong data foundation for your AI tools by implementing automated access controls, and eliminating redundant, outdated, and trivial data (ROT). These steps might sound straightforward, but many AI-adopting organizations haven't implemented these or other basic data governance procedures. Subsequently, there's a very real cost to pay for that, both in the near-term and further down the road, especially with agentic AI expected to exacerbate these issues even more.

EXPERT PERSPECTIVE

09

Capturing a return on AI investments

CAPTURING A RETURN ON AI INVESTMENTS

Customer outcomes define AI success

The two highest ranked success measures for AI implementations are all about customers. Organizations seek enhanced customer insights and personalization (40.9%), followed by increased customer engagement and satisfaction (39.7%). A set of enabling and internal factors – including increased data quality, better decision-making, and strengthened cybersecurity posture – follow close behind, albeit with declining ratings at the highest level and increasing ratings at the second to highest level. See Figure 9.1.

Figure 9.1

Measures of success for AI implementations: What's important

Percentage of respondents

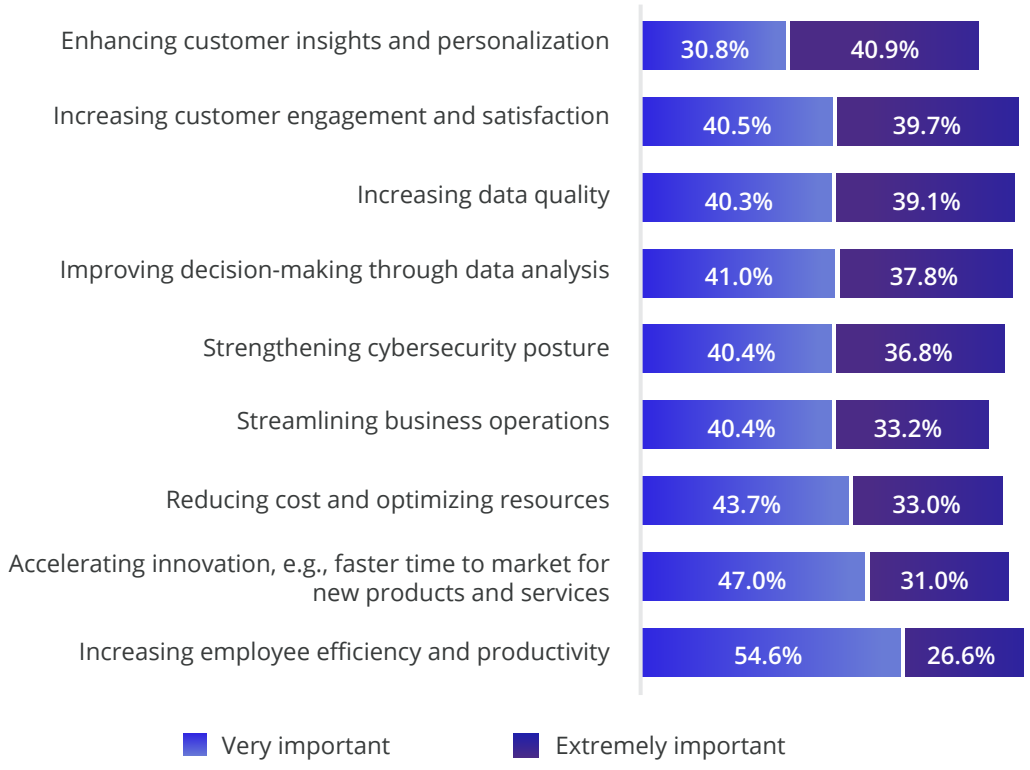

AI results fall short of expectations

Organizations say their biggest returns from AI investments have come from improving employee productivity, boosting data quality, and enhancing customer engagement.

But when we compare how important these goals are to how much value AI has actually delivered, there's still a gap. See Figure 9.2.

For example: LLMs are primarily trained using retrieval augmented generation (RAG), which uses an organization's internal data to augment prompt engineering or fine tuning, which retrains a model on a focused set of external data to improve performance. Consequently, there is inherent risk of them having bad data or not enough data which leads to inaccuracy and hallucination. What can help?

- Among those who said "enhancing customer insights and personalization" is very important, 1.5% more said they've achieved high value in that area
- But among those who rated it extremely important, 5.8% fewer said they've achieved extreme value – showing a clear shortfall between expectations and results.

Figure 9.2

Measures of success for AI implementation: Importance versus achieved

Percentage difference between the importance of the success measure and the achievement of the success measure

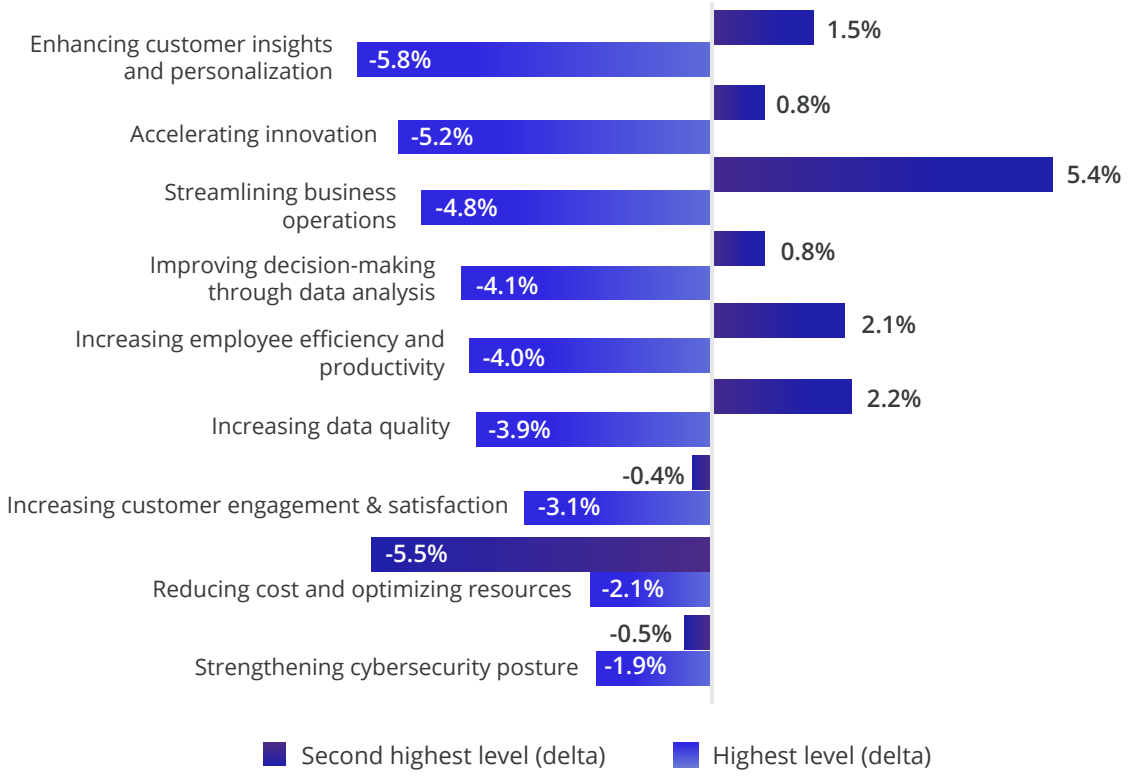

From Figure 9.2:

- Organizations are generally seeing a greater likelihood of achieving value if expectations are lower. Among organizations taking the extreme view, actually landing that value is proving more difficult.
- The most important success measure from Figure 9.1 - enhancing customer insights and personalization - has the largest gap in importance versus achievement.
- Despite vendor claims, AI hasn't delivered cost savings or resource optimization. In fact, both measures are trending negatively. Organizations appear to be finding it difficult to track net financial benefits to their AI investments.

Unmet expectations haven't diminished AI ambition

An average of 79% of respondents have a strong vision for next steps with their AI journey. This vision rests primarily on strengthening the data security and information governance underpinnings of using AI within the organization, followed by developing new products and services – a customer-focused strategy – and expanding the use of AI to additional departments. See Figure 9.3.

Figure 9.3

Importance of plans and strategies for future AI investments

Percentage of respondents

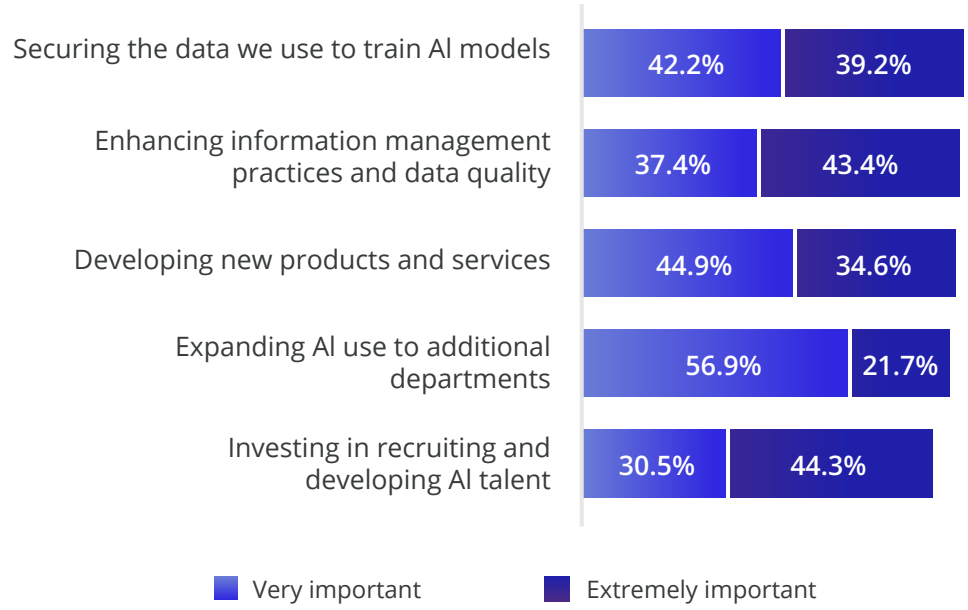

Patience for AI ROI is measured in months – not years

81.9% of organizations expect a return from their AI investments within 12 months, although only 9.4% want to see this return in less than three months. Most are willing to wait for a somewhat longer period of three to six months or seven to 12 months. Very few are willing to wait for more than 24 months. See Figure 9.4.

Figure 9.4

Timeframe for returns on AI investments

Percentage of respondents

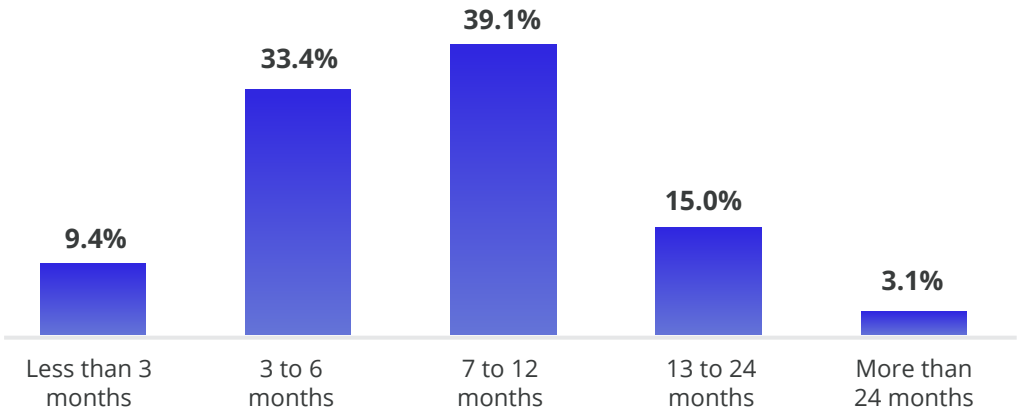

Demands for a quick return in the short-term (one year or less) are driven by the easy availability of the technology and less by the challenge of training employees, assuring high quality data inputs for model training, redesigning work processes, revamping operations, and ensuring that suggested outputs reflect the appropriate quality level.

Takeaways

01

ROI from AI investments is perceived as attainable in the short term and it's expected to make positive change in less than 12 months on average. This is often proving unrealistic, however, as many struggle to meet expectations due to data and process readiness gaps.

02

Customer-centric outcomes are the leading reasons for implementing AI. Enhancing customer insights and customer engagement are the two top success metrics, but they also show the largest gap between importance and achievement.

03

The vision for AI is strong. 79% of organizations have a clear roadmap for expanding AI use, especially in governance, security, and product innovation.

10

Collecting feedback on AI tools from employees

COLLECTING FEEDBACK ON AI TOOLS FROM EMPLOYEES

Organizations measure AI usage, not AI impact

Usage reports and employee surveys are the most used ways that organizations collect feedback on the use and effectiveness of generative AI assistants. Only 1.9% of organizations indicate that they do not collect feedback on AI usage and effectiveness from employees at all. See Figure 10.1.

Organizations typically rely on two sources for AI usage reports. The most common is the built-in reporting from the platform itself – like Microsoft 365 for Copilot, or Google for Gemini. The other option is using a third-party reporting tool. Overall, 89.6% of respondents rely on usage reports from one or both of these two sources, with 49.8% relying on their underlying platform only and 27.8% relying on both. The remainder (22.5%) use a third-party add-on reporting tool only.

Figure 10.1

Methods of collecting feedback on the use and efficacy of generative AI assistants from employees

Percentage of respondents

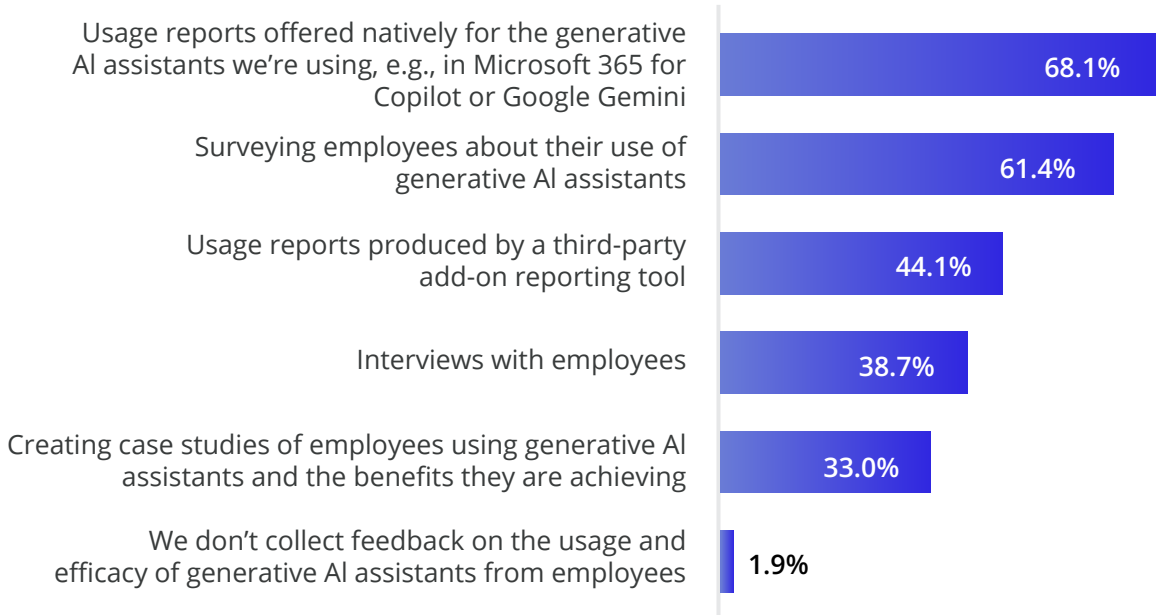

Usage reports assemble quantitative data based on actual usage patterns, which is good for assessing counts, averages, and frequency. Such reports are commonly used by IT departments to assess adoption technically. Usage reports, however, lack the ability to portray the human and business impact of embracing generative AI assistants, including how employees are finding success and overcoming roadblocks in the change and adoption journey. They also lack the signals of leadership support, engagement, and empowerment that other types of assessment approaches – such as surveys, interviews, and case studies – portray inherently.

Most organizations use two or three methods for collecting feedback. Of the organizations we surveyed for this research, 73.9% are using both a quantitative method (one or both types of usage reports) and a qualitative one (at least one of surveys, interviews, and case studies). 13.9% rely solely on quantitative methods, and 10.2% rely solely on qualitative methods.

Takeaways

01

Collecting employee feedback on AI success takes a blend of data sources. In this research, most rely on quantitative sources (89.6%), but underutilize qualitative insights (e.g., interviews, case studies).

02

Multi-method feedback is best. 73.9% use both quantitative and qualitative methods, which provides a more holistic view of AI adoption and impact.

03

Platform-native tools dominate. Most organizations rely on built-in reporting from platforms like Microsoft 365 and Google, with fewer using third-party analytics.

11

Interventions are strengthening AI literacy

INTERVENTIONS ARE STRENGTHENING AI LITERACY

AI fluency starts with role-based guidance

Where and how to integrate AI into core work processes across the organization relies on having a cadre of employees who are skilled and trained in the use of AI. This includes where AI makes sense, how to use AI without compromising data security or compliance mandates, and how to identify hallucinations in AI output. Almost all organizations (99.5%) have used a range of interventions to strengthen AI literacy among employees.

Role-based guidelines and/or use cases were rated as the intervention type with the highest impact on employees, with 79.4% of respondents giving a rating of "highly" or "extremely" impactful. Personal experimentation, webinars, and conferences were in a second band of impactful interventions, although these are more about leveraging what employees as individuals can learn from wider perspectives to bring into the organization. The high rating given to these second band options indicates that, for many organizations, incorporating AI into work practices is still in its early days, because the referential ideals are from outside the organization, not within it. See Figure 11.1.

Figure 11.1

Impact of interventions on improving AI literacy among employees

Percentage of respondents

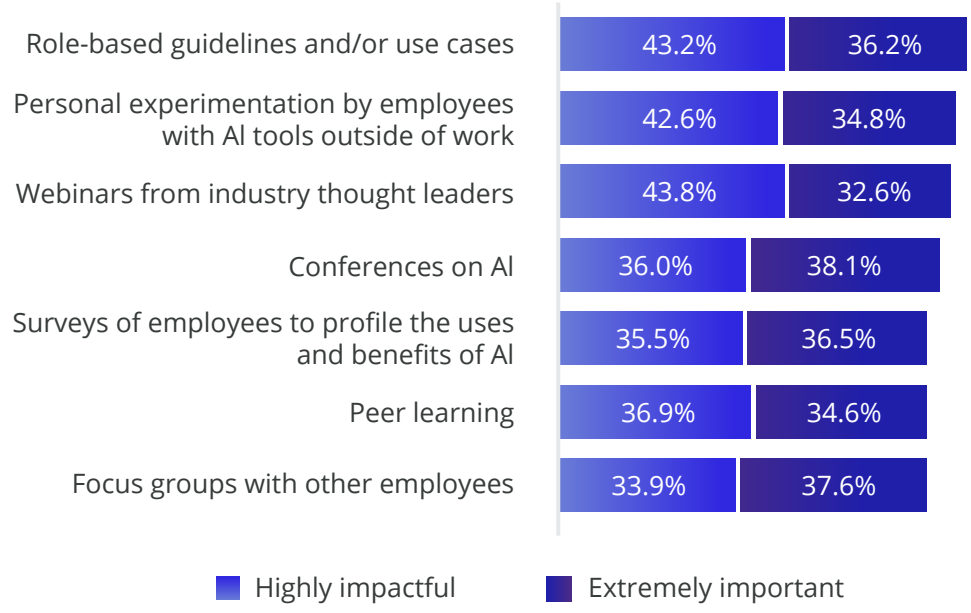

All seven interventions were rated as being highly or extremely impactful by an average of 74.6% of respondents. In other words, doing something intentional – anything, even – to increase AI literacy will pay high dividends among employees.

AI literacy demands stronger human judgment

Developing AI literacy is essential due to the growing list of challenges for humans when AI is used. This includes:

- Model collapse, where training data for large language models becomes less effective over time, resulting in degraded outputs. This often happens when models are trained on synthetic, AI-generated output. Employees need judgment and expertise to identify outputs that contradict reality, and a way of escalating such outputs for remediation.
- Employees who use generative AI to confirm their positions – rather than to challenge and extend them – can get caught in the grip of extreme confirmation bias. The end result is degraded decision-making skills among employees, with significant negatives for organizations as a consequence.
- Reduction in deep work, where employees rely on AI for quick answers and stop engaging critically with their tasks. Taken too far, this undermines critical thinking abilities and de-skills the workforce.

While it is unlikely that the use of AI will be stopped, working to amplify its positives and mitigate its negatives will be essential for organizations.

Takeaways

01

Investing in improving AI literacy is nearly universal: 99.5% of organizations have implemented interventions to improve employee understanding and safe use of AI.

02

Role-based training is most impactful: Tailored guidelines and use cases are rated highest in effectiveness, suggesting a need for more contextual learning.

03

AI literacy is a resilience strategy: As risks like hallucinations and confirmation bias grow, literacy becomes essential to preserve human judgment and critical thinking.

12

Recommendations

RECOMMENDATIONS

1 Build trust before scale.

AI adoption without trust is a risk multiplier. Before expanding access or investing in more licenses, organizations must address the root causes of mistrust – namely, inaccurate outputs, shadow AI usage, and weak data foundations. Prioritize data quality, intentional governance, and employee literacy to ensure AI is a force multiplier, not a liability. Trust is not a milestone – it's a practice that must evolve alongside your data and workforce.

2 Govern AI like a critical business system.

AI is not a side project – it's a core operational capability. Treat it with the same rigor as financial systems or cybersecurity infrastructure. That means implementing and continuously evolving AI Acceptable Use policies, operationalizing governance into day-to-day workflows, and ensuring visibility across sanctioned and unsanctioned tools. As AI adoption grows, so must oversight across business units, platforms, and data sources.

3 Secure the data pipeline, not just the endpoint.

Security breaches and hallucinated outputs often stem from poor data hygiene, not just model flaws. Invest in upstream controls: robust information management, smart data classification, and third-party governance tools that validate AI outputs against policy and context. Resilience starts with the data, not the dashboard.

4 Redefine AI success around resilience and impact.

Usage metrics alone don't tell the full story. Shift from measuring AI adoption to measuring AI impact – on employees, customers, and business outcomes. Combine quantitative usage data with qualitative feedback to understand what's working, what's not, and where governance or training gaps are holding you back. True success also demands accountability – clear ownership of AI systems, visibility into outcomes, and mechanisms for course-correction as adoption scales.

13

Conclusion

CONCLUSION

> *The age of AI is here - but trust must be earned.*

This year's research confirms what many organizations are beginning to realize: the real differentiator in AI success isn't speed – it's stewardship. The organizations seeing the greatest returns from AI are not those who adopted first, but those who governed best. They've invested in data quality, embraced continuous policy evolution, and prioritized human judgment alongside machine intelligence.

As AI becomes more deeply embedded in business operations, the stakes are rising. Inaccurate outputs, data breaches, and unchecked automation are no longer theoretical risks. They are real-world consequences that can erode trust, damage reputations, and stall innovation. The path forward demands more than enthusiasm; it requires intentionality.

At AvePoint, we believe that secure, resilient, and governed AI is not just a best practice – it's a business imperative. Our mission is to help organizations build the foundations that make AI trustworthy, scalable, and sustainable. That means strengthening data governance, elevating employee literacy, and embedding ethical guardrails into every layer of AI deployment.

As you reflect on the insights in this report, we invite you to assess your organization's AI maturity. Where are your strengths? Where are the gaps? And how can we help you close them?

Because in the age of intelligent automation, leadership isn't just about what you build with AI. It's about how responsibly you build it.

www.AvePoint.com | @AvePoint