SailPoint dimensional research

AI agents: The new attack surface

A global survey of security and IT professionals and executives

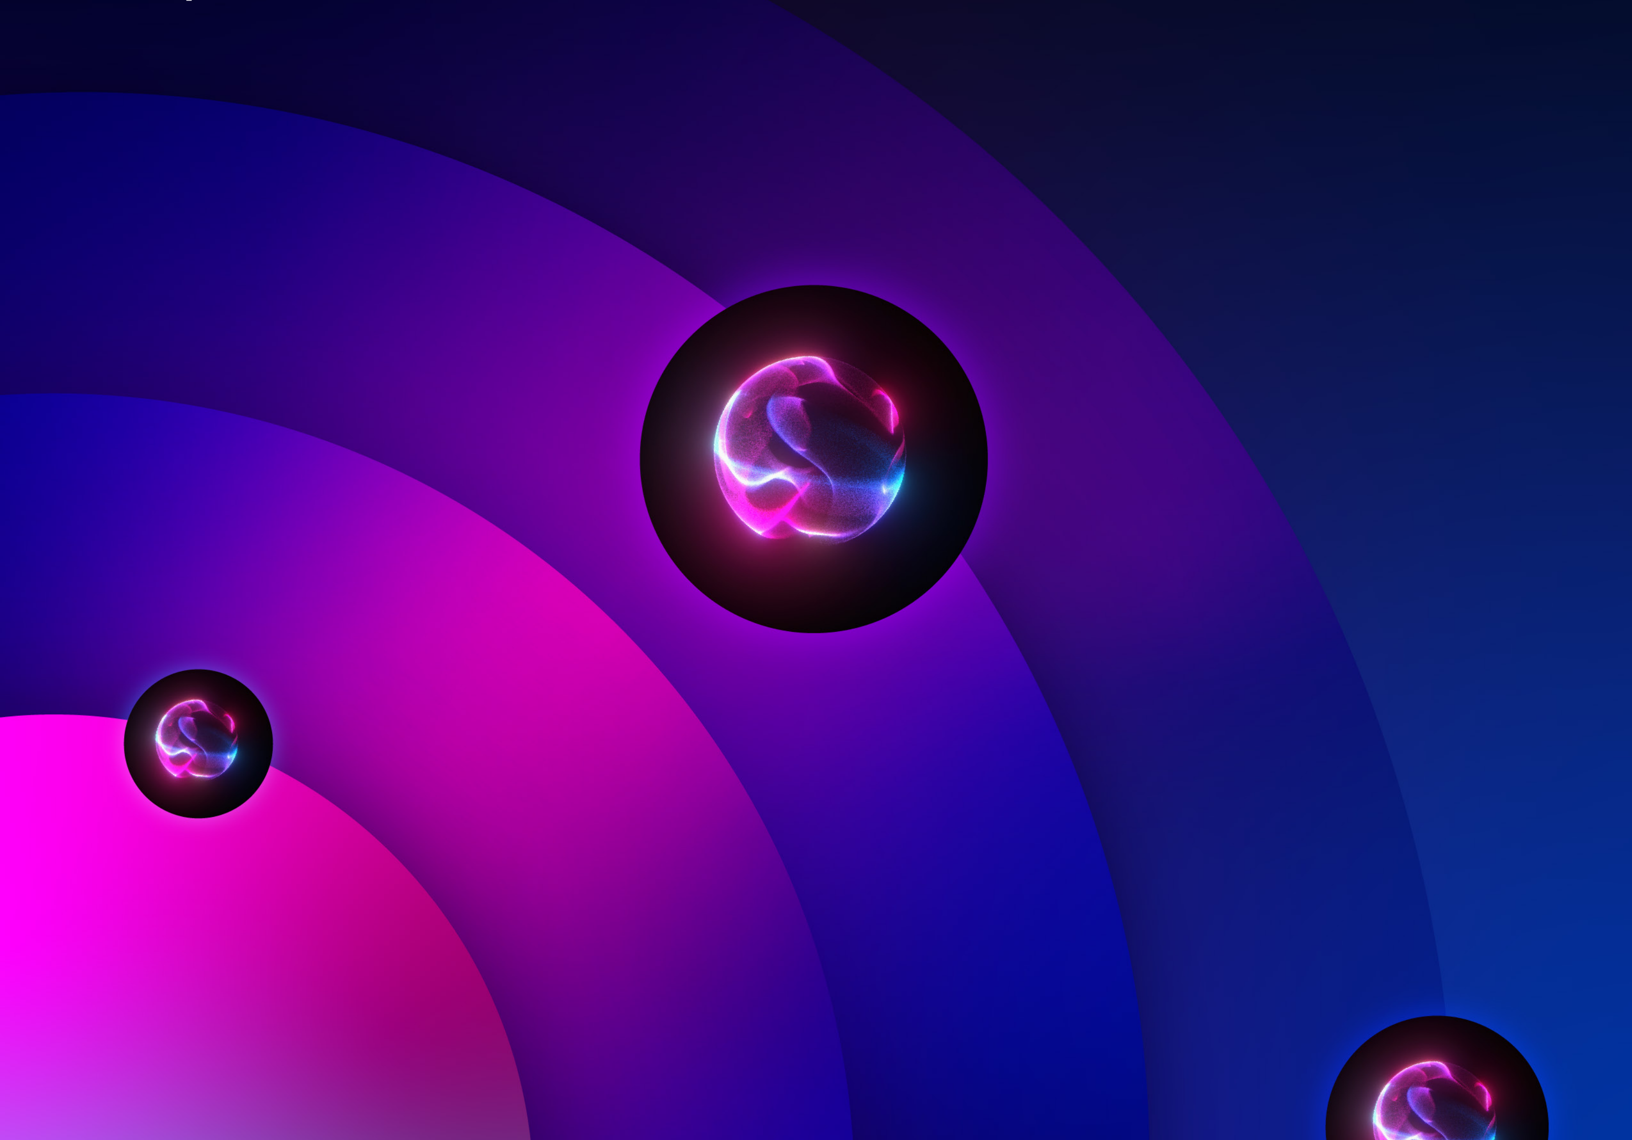

Introduction

This paper presents key findings from a global primary research survey conducted by independent firm Dimensional Research. Through this research, SailPoint aimed to examine the current use, adoption, and governance of AI agents, with a particular focus on the distinct risks their identities present compared to those of human and machine identities. It explores issues such as unintended actions, gaps in governance, and the underlying causes of AI agent risk, as well as the extent to which organizations are leveraging identity security tools to provision and manage these identities.

Executive summary

Research shows a concerning 82% of companies now utilize AI agents, with over half reporting these agents access sensitive data daily. Alarmingly, 80% of organizations have experienced unintended actions from their AI agents, including inappropriate data sharing and unauthorized system access. Some AI agents have even been coerced into revealing access credentials.

This lack of control has led 96% of technology professionals to identify AI agents as a growing security threat—66% believe this risk is immediate, while 30% see it emerging in the near future. The primary concerns include inadequate data access and data sharing controls and unpredictable AI agent behaviors. These agents handle diverse sensitive information including customer data, financial records, intellectual property, legal documents, and supply chain transactions.

While 92% of respondents recognize AI agent governance as crucial to enterprise security, only 44% have implemented relevant policies. Although 71% of IT departments claim awareness of AI agent data access, this knowledge extends to compliance, legal, or executive teams in less than half of the surveyed companies.

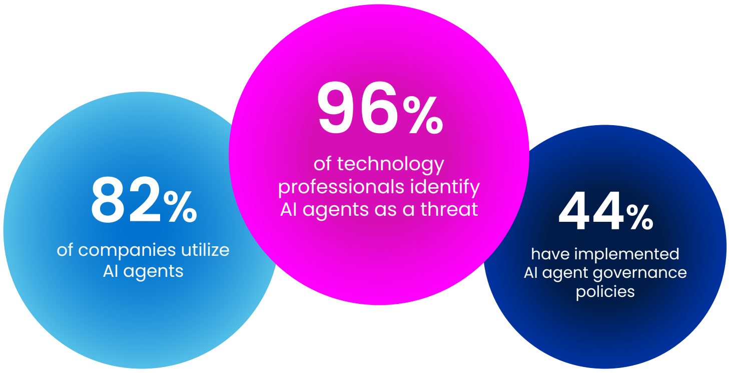

82% of companies utilize AI agents

96% of technology professionals identify AI agents as a threat

44% have implemented AI agent governance policies

Those surveyed indicated that AI agents pose a greater risk than both machine and human identities. Unlike traditional identities, AI agents often require broader privileges across more systems, data, and services. They are also more difficult to govern, with rapid access typically provisioned directly within IT. Despite these concerns, just over 60% of companies employ identity security solutions to manage access. With 98% of organizations planning to deploy new AI agents within the year, data exposure risks are escalating rapidly.

The business value of AI agents is undisputed, but the potential consequences of compromised sensitive data could be devastating. Companies urgently need comprehensive solutions to govern access permissions and monitor and control which systems and data AI agents are accessing.

Key findings

Things of note: In the survey the term “**AI agents**” (also known as **Agentic AI**) was defined as autonomous systems that perceive, make decisions, and take action to achieve specific goals within an environment. AI agents or Agentic AI often require several different machine identities to access needed data, applications and services.

AI agent use is already pervasive but unintended actions are exposing sensitive data

- 82% of companies are already using AI agents
- 53% acknowledge AI agents are accessing sensitive information
- 80% reveal AI agents have performed unintended actions of accessing and sharing inappropriate data

Growing security risk driven by diverse data access and lack of governance and auditability

- 66% state AI agents are a growing security risk
- Numerous data control issues are driving AI agent security risk
- Numerous teams already using AI agents
- 92% state governing AI agents is paramount to enterprise security
- Only 44% currently have any governance policies in place for AI agents

AI agents lead identity risks with broader access and truncated visibility and approval processes

- 72% state AI agents pose a greater risk than machine identities
- 64% confirm that AI agents require multiple identities to access necessary data, applications, and systems
- AI agents require broader privileges and are harder to govern, with faster access and limited approval processes

Detailed findings

AI agents with access to sensitive data are used daily by most companies

AI adoption is nearly a ubiquitous topic today among most organizations, along with generative AI, large language models (LLM), and AI-based analytics. However, this research finds that the use of AI agents, also known as Agentic AI has surged, as 82% of companies state they are using AI agents today.

Is your company using any applications that utilize AI agents?

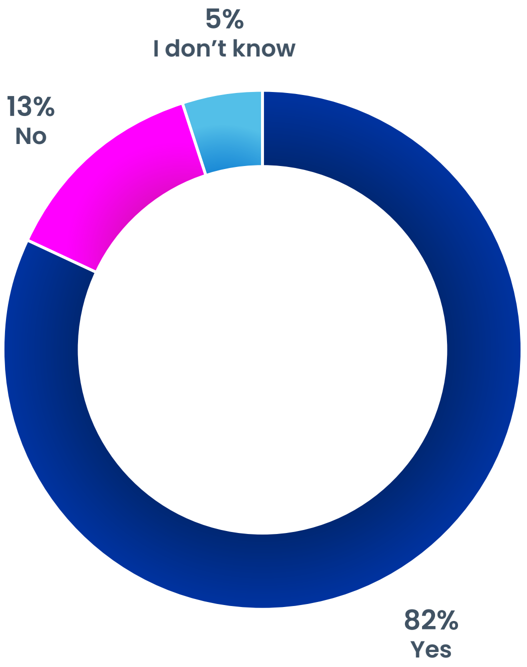

AI agents tend to be goal-based where a task is given to the AI agent and it must find the information and resources to satisfy that request. As such, technology professionals shared that more than half (53%) of AI agents are accessing sensitive information. And 58% of those AI agents are accessing that sensitive information daily. The chart below also reveals that 10% don't know if AI agents are accessing sensitive data, a concerning and reoccurring finding throughout this report.

In your experience, will AI agents have access to sensitive company information?

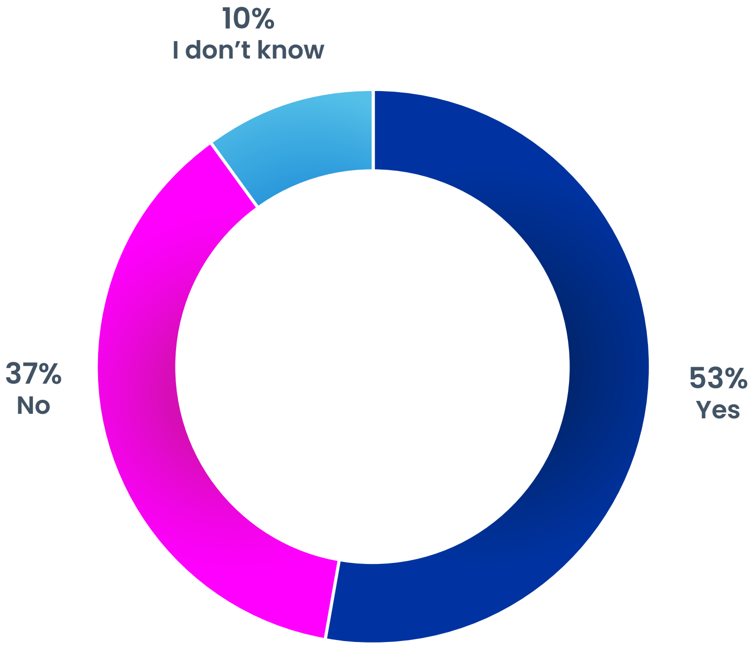

AI agents access and share inappropriate and sensitive data

While many studies have focused on the benefits of AI, this research study sought to understand whether AI agents are performing actions outside their intended scope. In the chart below, the key takeaway is only 20% of companies state that their AI agents have NOT performed unintended actions, which deductively means 80% of companies are experiencing AI agents performing unintended actions.

Leading the list of unintended actions, 39% of respondents reported AI agents accessed unauthorized systems, while 33% said agents accessed inappropriate or sensitive data. Although these behaviors may reflect attempts to fulfill a task, the next set of actions is more concerning: 32% noted that AI agents enabled the download of sensitive data, and 31% said the data was inappropriately shared. Additionally, AI agents have accessed the internet in search of information, introducing unverified data into their outputs. Perhaps most alarming, nearly one in four companies (23%) reported that AI agents were coaxed into revealing access credentials—potentially opening the door for cybercriminals.

What type of actions have your company's AI agents performed that were beyond its intended scope?

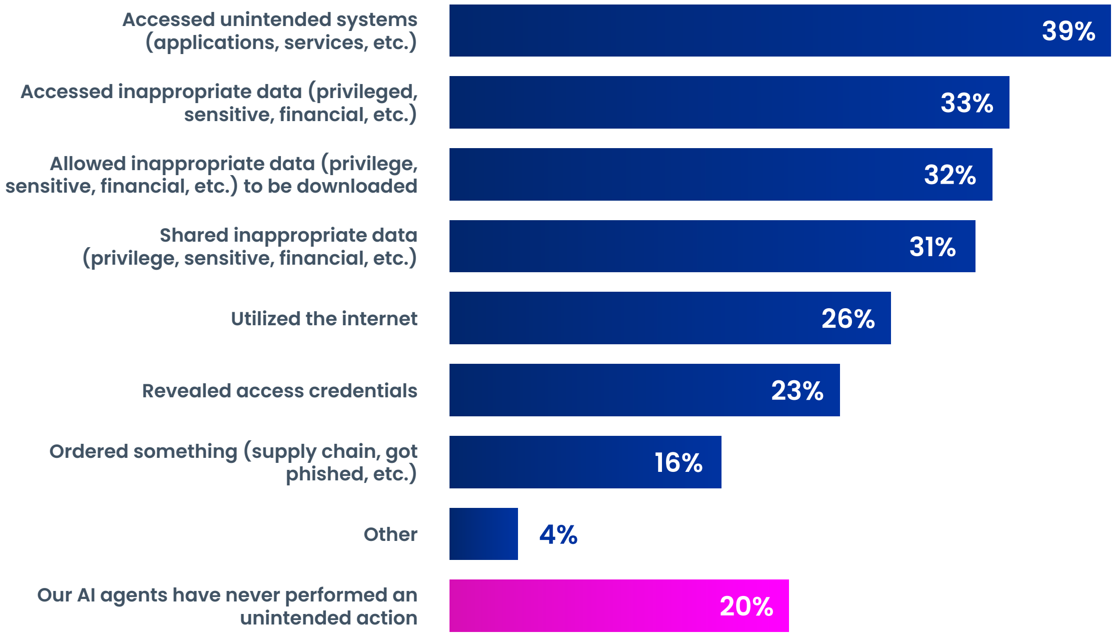

AI agents are a growing security risk

With Agentic AI performing unintended actions for 80% of the companies, it is not surprising that 96% state AI agents are a growing security risk, with 66% stating that risk is present today.

In your experience, are AI agents a growing security risk?

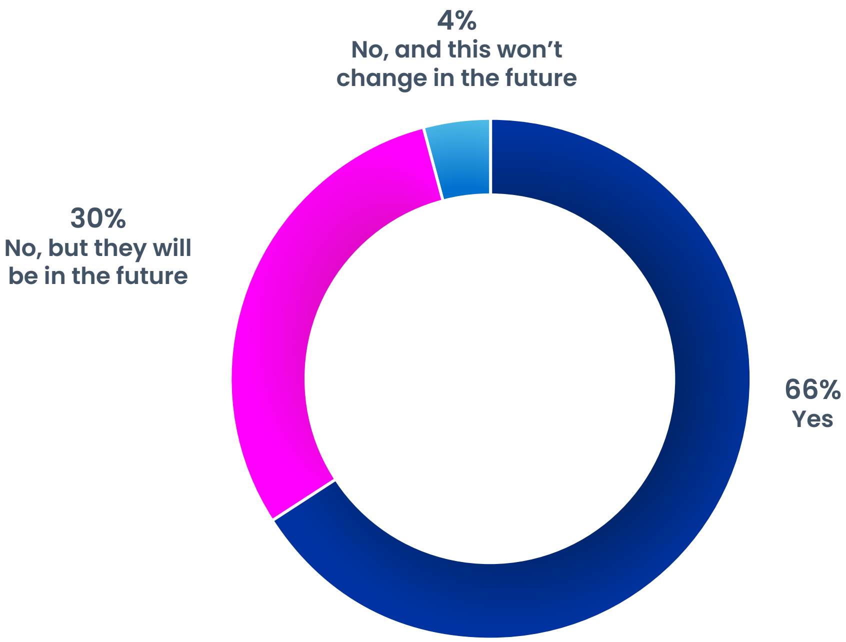

AI agents present numerous business risks and data control issues

The survey also explored the specific factors contributing to AI agents as a security risk. The top six concerns were closely ranked, with only a six-point spread. Leading the list is AI agents' ability to access privileged data (60%), followed by their potential to perform unintended actions (58%), as detailed earlier in this report. Other major concerns include sharing privileged data (57%), making decisions based on inaccurate or unverified data (55%), and both accessing and sharing inappropriate information (54%). Additionally, 49% of respondents cited the risk of AI agents generating inaccurate outputs for users. Notably, 38% reported incidents involving the disclosure of security information to internal bad actors, while 29% cited exposure to external threats. From a business standpoint, these risks span compliance failures, data privacy breaches, security vulnerabilities, and the dissemination of incorrect information to employees, partners, and customers.

In your experience, what makes an Agentic AI a security risk?

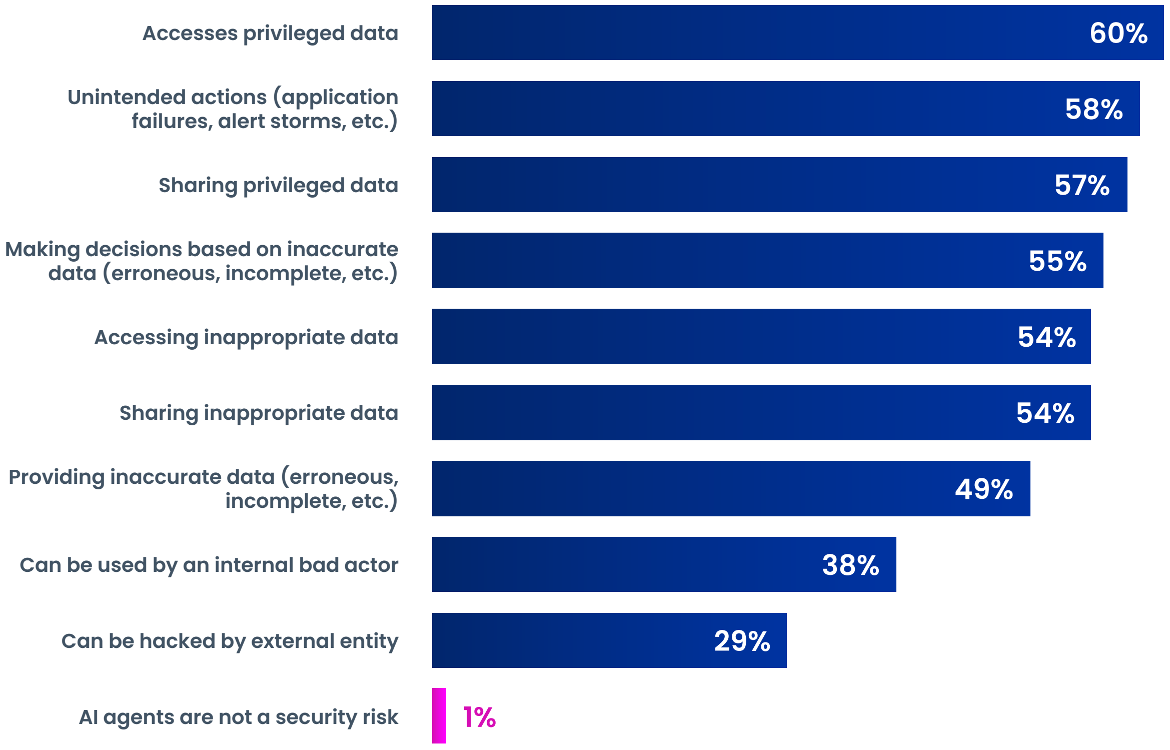

AI agents have access to key data across the enterprise

To better understand AI agent adoption and the intrinsic risks, the research identified which teams are currently using AI. As expected, IT leads the way at 52%, followed by customer service (46%), cybersecurity (44%), support desk (44%), and software development (39%). Notably, AI agents are also being utilized by product management (26%), sales (25%), compliance (24%), as well as marketing and HR (both at 23%). This data highlights not only the breadth of enterprise adoption but also the wide range of sensitive business data and information AI agents are now accessing.

At your company, what teams are currently using solutions that utilized AI agents?

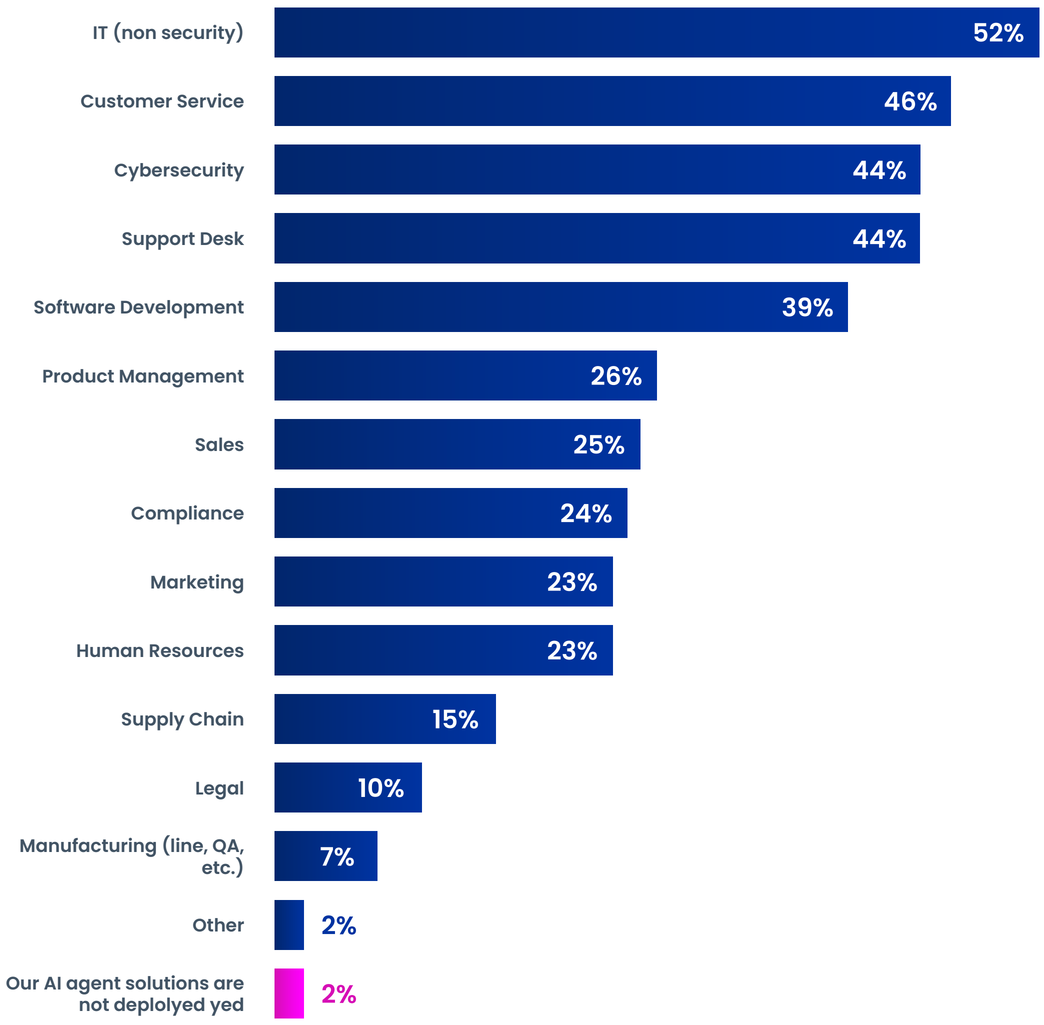

AI agent governance critical to enterprise security

Given the broad scope of data AI agents can access, the previously cited risks, and the frequency of unintended actions, an overwhelming 92% of technology professionals indicated that governing AI agents is critical to enterprise security.

In your opinion, is governing AI agents critical to ensuring enterprise security?

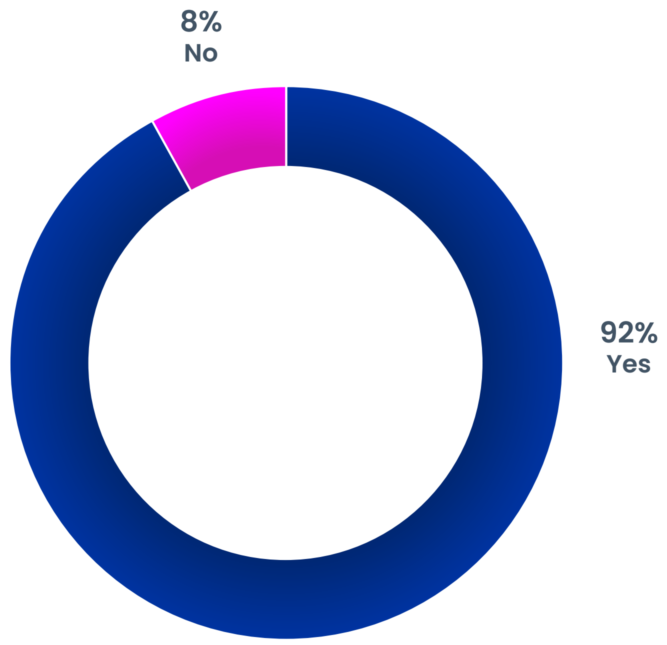

Despite widespread recognition of the risks among the survey respondents, only 44% of organizations currently have governance policies in place to manage AI agents and the data they access and share. While 53% are in the process of developing such policies, the reality is that most remain exposed today. Notably, just 3% of respondents reported having no plans to implement AI agent governance at all.

Does your company currently have governance policies specifically for AI agents?

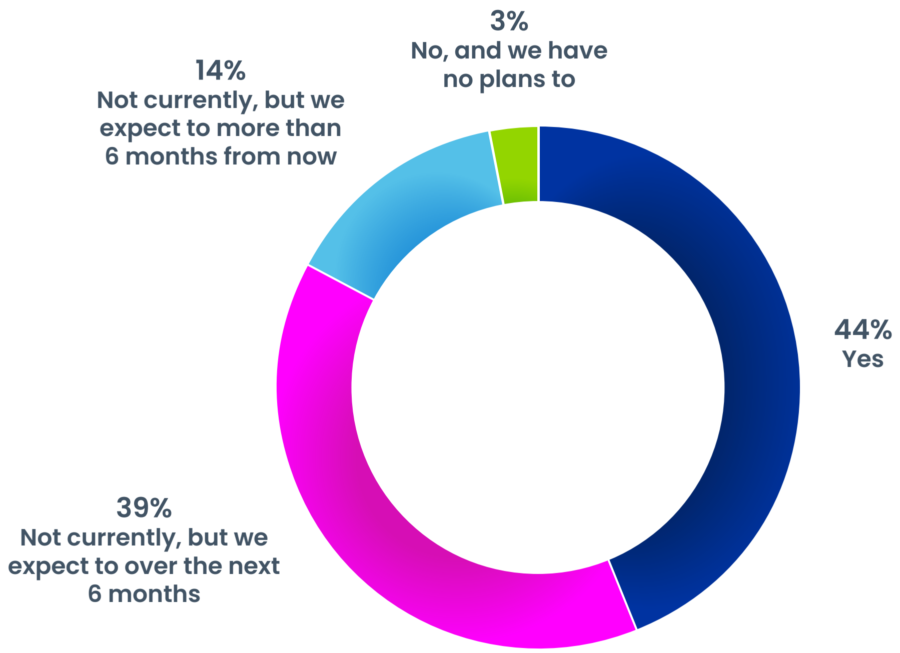

Executives, compliance, legal: Uninformed about data AI agents access

Effective governance of AI agents begins with a clear understanding of the data they can—and should—access. This information must be shared with teams responsible for compliance and data protection. As shown in the chart below, IT is the most informed team regarding AI agent data access, given their role in implementing the technology, managing configurations, and provisioning credentials. However, awareness drops significantly among other critical stakeholders: compliance (47%), legal (39%), executives (34%), and other departments—despite their essential role in identifying sensitive data, safeguarding the organization, and minimizing risk.

Which teams have been advised of all of the data that AI agents have access to?

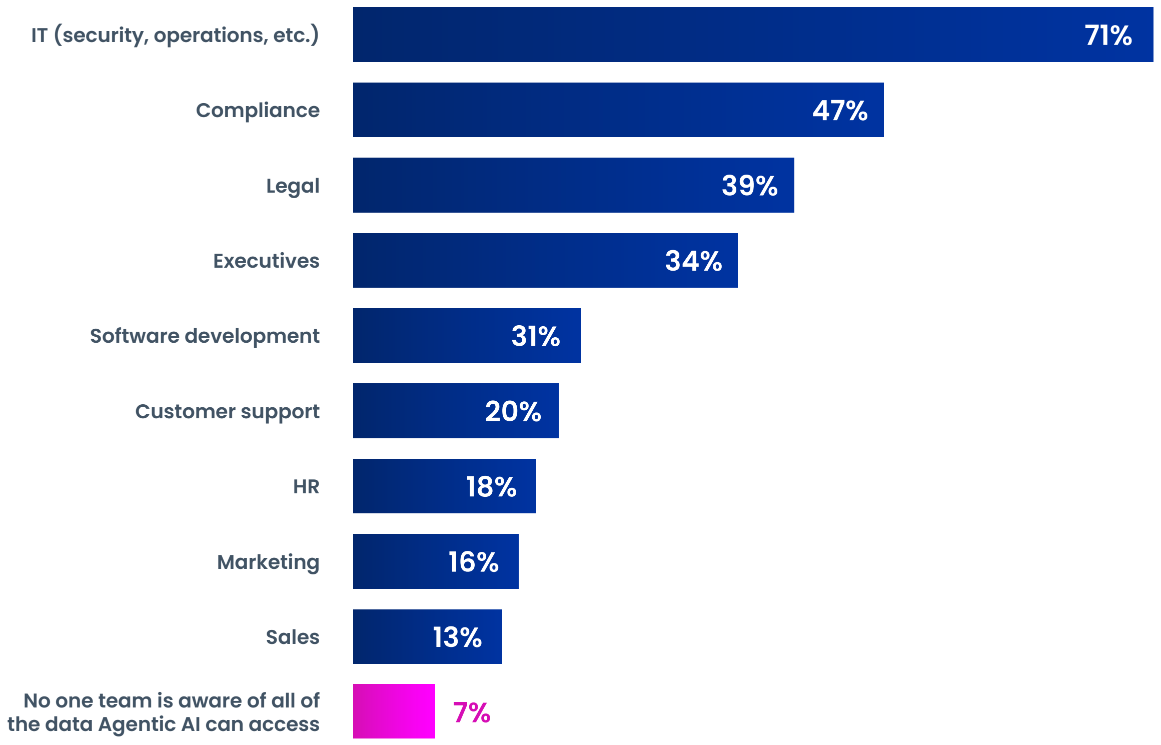

This lack of visibility into data access has resulted in only 52% of companies reporting that they can track and audit all data used or shared by AI agents. Consequently, nearly half of organizations remain unaware of what data is being accessed or exposed—often putting them at risk of violating data protection regulations.

Is your company able to track and audit every piece of data an AI agent accesses?

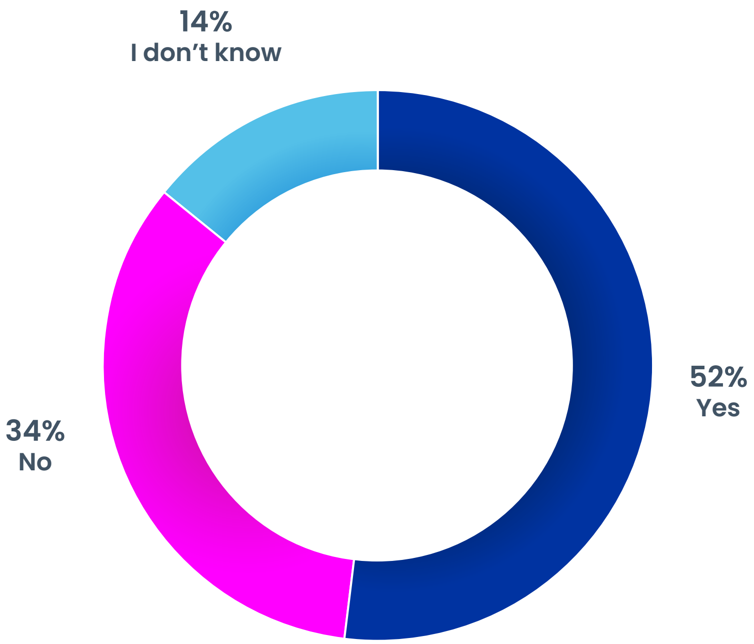

Access governance critically important to manage AI agent risk

The lack of control and visibility over the data AI agents access and share, as outlined in the preceding sections, appears to have prompted the 62% of respondents in the chart below to identify access governance for managing risk of AI agents as critically important. Notably, no respondents selected "low importance" or "not at all important," resulting in unanimous agreement—100% of participants view AI agent access governance as essential to managing associated risks.

In your experience, how important is access governance for managing the risks of AI agents?

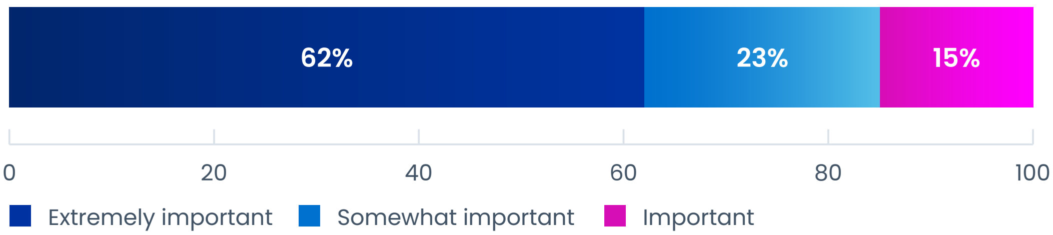

AI agent identities create more risk than machine and human identities

IT has long faced data and governance challenges with applications and services that access and share information—typically managed through machine identities. However, 72% of technology professionals now believe AI agents present a greater risk to the business than traditional machine identities.

In your opinion, do AI agents present a greater risk than other machine identities?

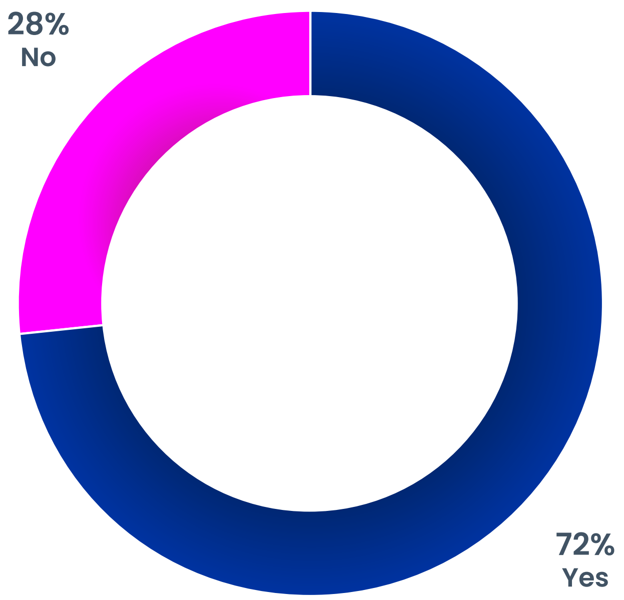

As shown in the chart below, 90% of participants indicated that AI agent identities differ significantly from human identities. To better understand the unique risks AI agents pose at the identity level, respondents were asked how these agents compare to what has traditionally been viewed as the highest-risk identity: humans. The top concern cited was that AI agents often have broader access to applications and data (54%) than typical human users. Additionally, 40% noted that AI agents are more difficult to govern—likely due to limited visibility and their potential for unpredictable actions. While human identities typically undergo structured access approvals involving managers or executives, AI agent access is often provisioned solely by IT (35%) and approved more quickly (34%). As a result, IT may lack full awareness of the specific types of data being accessed—such as customer information, intellectual property, or employee records—making it difficult to apply appropriate compliance or sensitivity controls.

How are AI agent identities different from human identities?

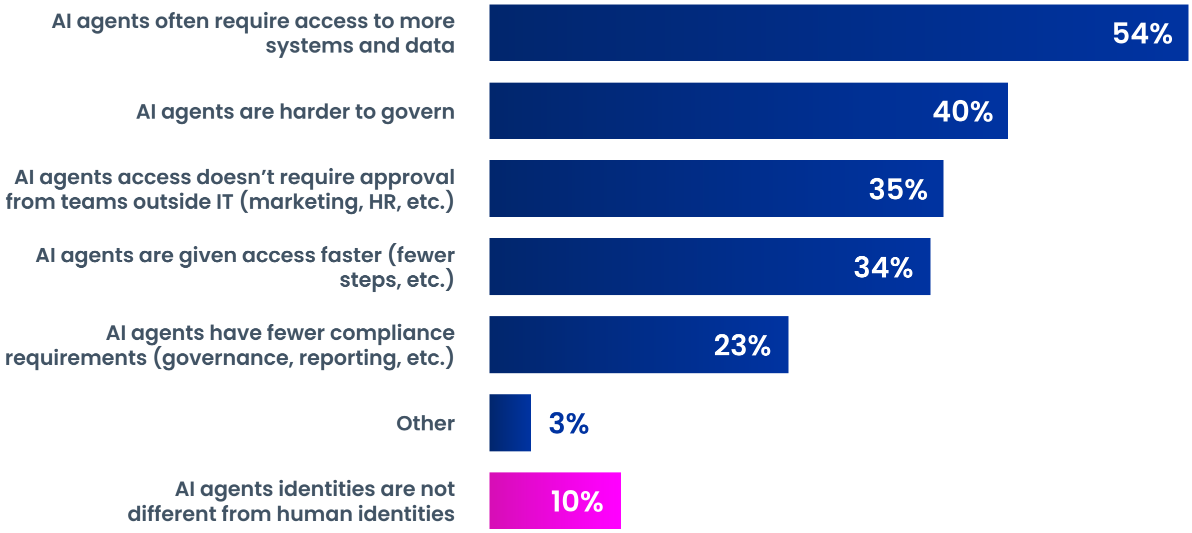

AI agents utilize multiple identities to access needed information

To explore why AI agents may be more difficult to govern, participants were asked whether a typical AI agent requires multiple identities to access the systems, applications, and data it needs. Sixty-four percent confirmed that AI agents often rely on several access identities, complicating efforts to track and correlate data usage and sharing. Despite this complexity, only 62% of organizations reported using an identity security solution to manage AI agents and their multiple identities.

At your company, do AI agents require multiple identities to access necessary systems, applications, and data?

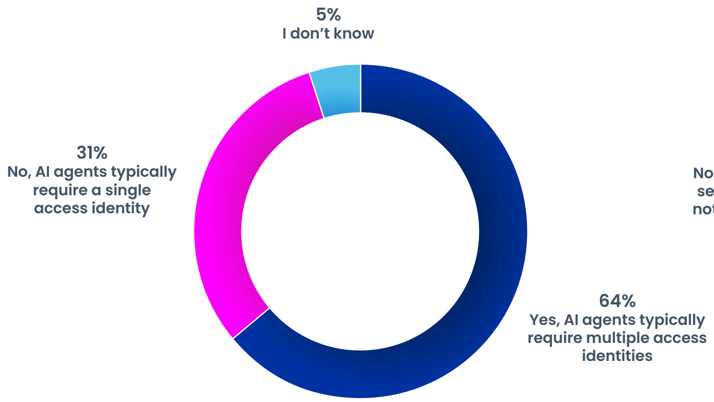

At your company, is access for AI agents provisioned and governed by identity security solutions?

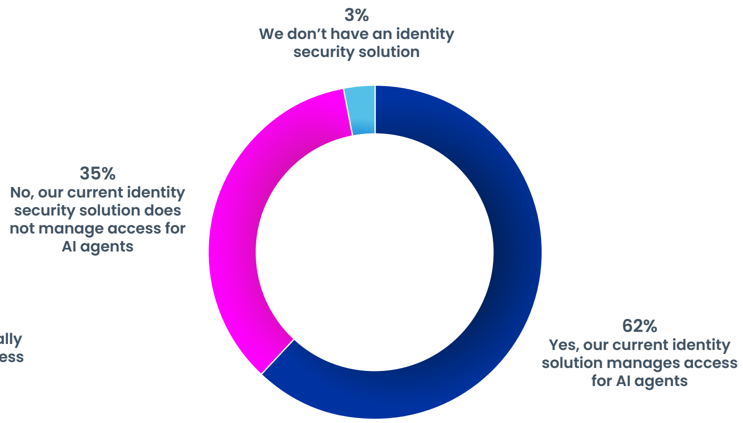

AI agents on a rapid enterprise-wide rollout plan

Given the widespread adoption of AI agents across teams and growing awareness of the associated risks, the research aimed to determine whether organizations are pausing deployments to strengthen security and identity controls for their AI agents. The findings suggest otherwise—an overwhelming 98% of companies plan to expand their use of AI agent-driven solutions within the next 12 months, spanning nearly every team across the enterprise. While AI holds the promise of greater value, it also significantly amplifies exposure and risk.

Over the next 12 months, which teams will deploy new solutions that utilize AI agents?

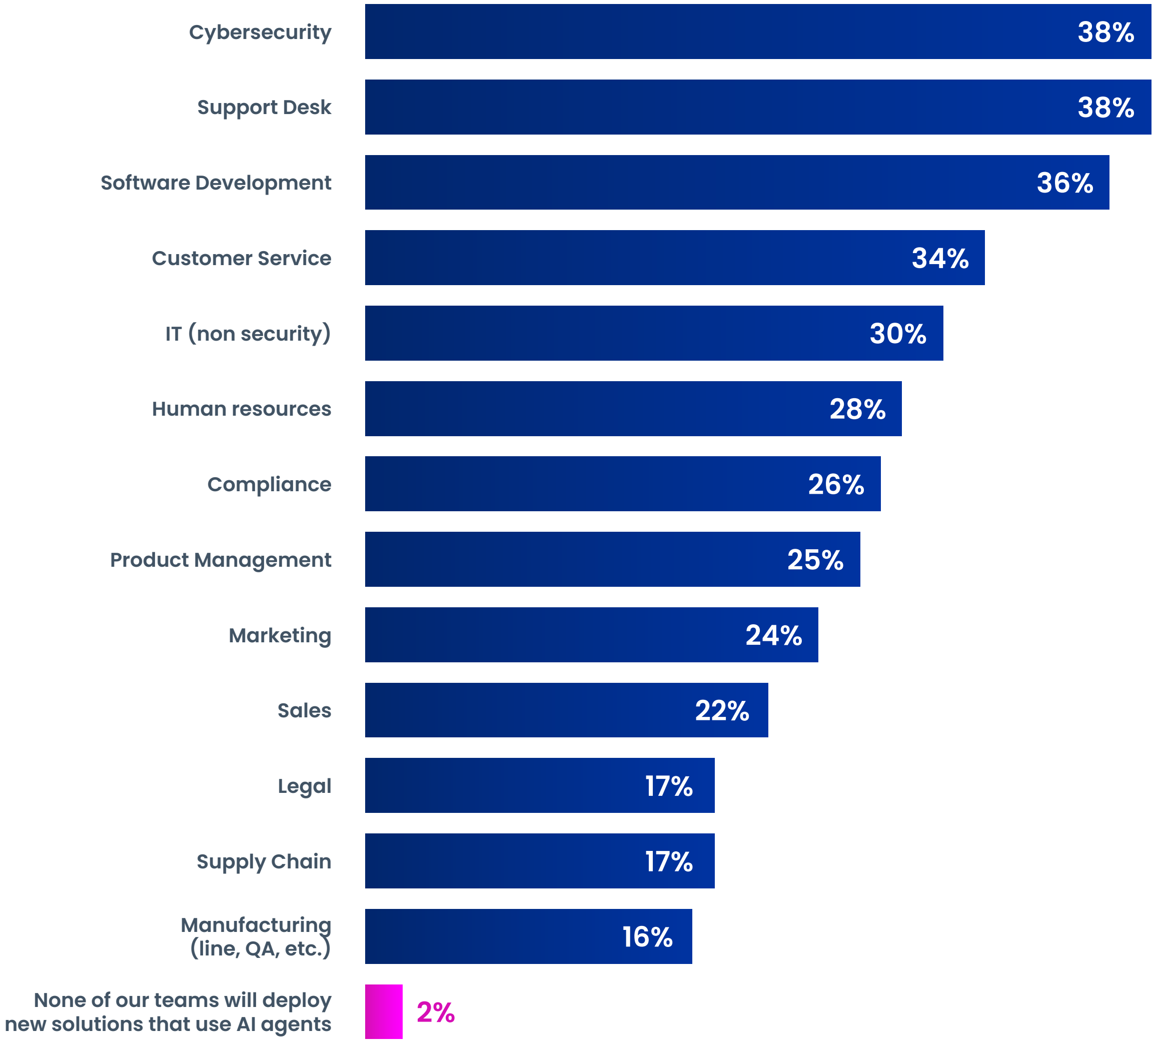

Conclusion

AI agents have rapidly become integral across organizations, with 98% of companies planning to expand their AI agent deployments in the next year. This widespread adoption promises efficiency and innovation as these agents access and process data throughout the enterprise.

However, this progress comes with significant risk—80% of organizations report their AI agents have already performed unauthorized actions, including accessing and sharing sensitive information. Beyond regulatory compliance issues, this creates vulnerabilities affecting employees, partners, and customers who may receive inaccurate information or, more dangerously, expose access credentials to malicious actors.

This reality explains why an overwhelming 96% of respondents identify AI agents as an escalating security threat. While establishing governance over AI agent access is widely recognized as essential, fewer than half of surveyed companies have implemented any governance policies. Most organizations fail to track or audit the data AI agents access, leaving legal teams, compliance officers, and executives without visibility into the information these systems can reach.

The security challenge is magnified because AI agent identities present greater risks than traditional machine identities and require more complex management than human users. Despite needing multiple permissions to perform their functions, AI agents typically receive expedited access through IT departments alone. By design, these agents will explore all accessible resources to fulfill requests—creating inherent security vulnerabilities.

To address these risks effectively, organizations must implement specialized identity security solutions with AI agent-specific controls that can restrict access to sensitive data, maintain comprehensive audit trails, and provide transparency to all stakeholders. With data breaches already causing significant financial and reputational damage, an unmanaged AI agent represents an even greater vulnerability—capable of compromising enterprise security with a single response to a cleverly crafted question.

Survey methodology

IT professionals responsible for AI, security, identity management, compliance, and operations at enterprise companies representing all seniority levels were invited to participate in a survey on their company's use of AI agents.

A total of **353 qualified participants** completed the survey which was conducted by Dimensional Research, an independent third-party. All participants had enterprise security responsibilities. Participants were from 5 continents providing a global perspective. The survey was administered electronically, and participants were offered token compensation for their participation.

Responsibilities

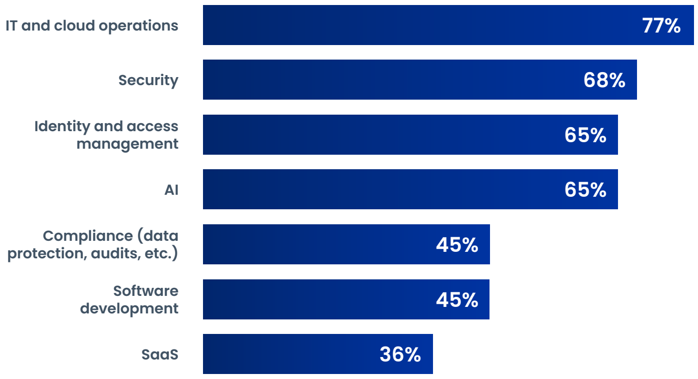

Size

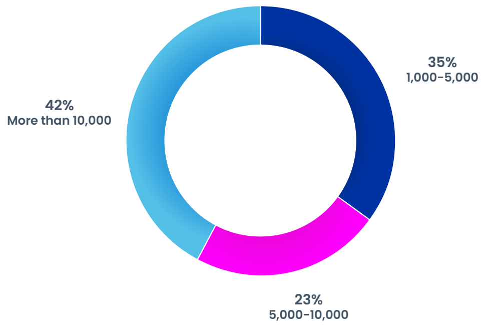

Industry

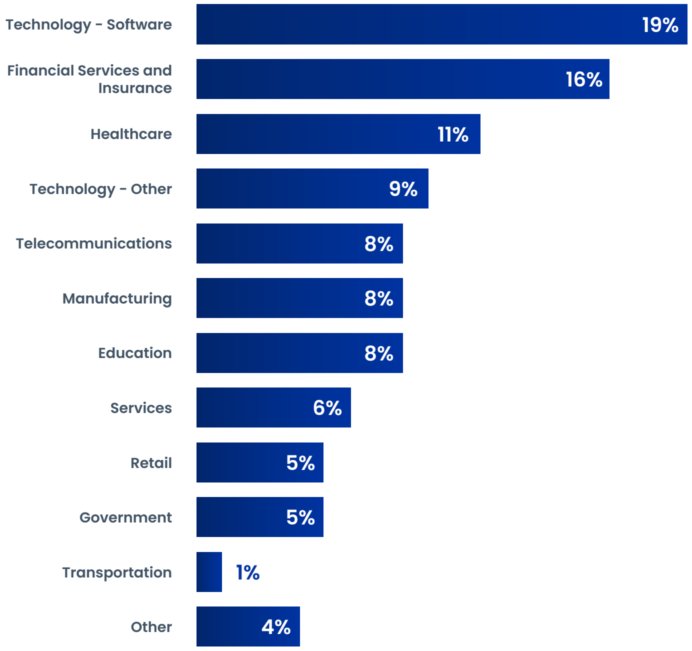

Senority

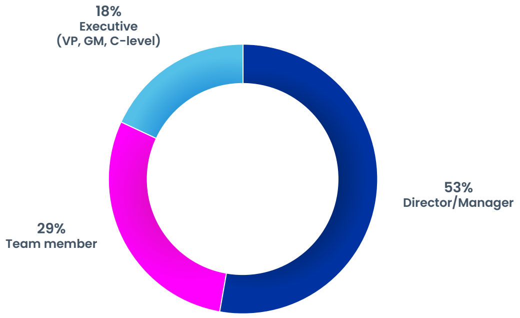

About Dimensional Research

Dimensional Research provides practical marketing research to help technology companies make their customers more successful. Our researchers are experts in the people, processes, and technology of corporate IT and understand how IT organizations operate. We partner with our clients to deliver actionable information that reduces risks, increases customer satisfaction, and grows the business.

For more information, visit **www.dimensionalresearch.com**.

Disclaimer: The information contained in this report is for informational purposes only, and nothing conveyed in this report is intended to constitute any form of legal advice. SailPoint cannot give such advice and recommends that you contact legal counsel regarding applicable legal issues.

SailPoint®

About SailPoint
At SailPoint, we believe enterprise security must start with identity at the foundation. Today's enterprise runs on a diverse workforce of not just human but also digital identities—and securing them all is critical. Through the lens of identity, SailPoint empowers organizations to seamlessly manage and secure access to applications and data at speed and scale. Our unified, intelligent, and extensible platform delivers identity-first security, helping enterprises defend against dynamic threats while driving productivity and transformation. Trusted by many of the world's most complex organizations, SailPoint secures the modern enterprise.