Omdia
by informa techtarget

AI Agents and Identity Security: How Enterprises Are Rewriting the Rules

Todd Thiemann | *Principal Analyst*

March 2026

This Omdia research and eBook was commissioned by BeyondTrust and is distributed under license from TechTarget, Inc.

Research objectives

Organizations are eagerly investing in AI and AI agents; however, that investment poses new identity security challenges due to an increased cybersecurity attack surface, new compliance concerns, and novel AI agent management challenges. Identity security is foundational to AI security, and enterprise teams are coming to terms with the identity implications of agentic AI technology. AI agents can streamline enterprise operations, but over-permissive access and ungoverned agents can put data at risk of compromise and exfiltration.

To understand these trends and market dynamics, Omdia executed a survey of 400 IT and cybersecurity professionals at organizations in North America involved with or responsible for identity security technology products and services.

THIS STUDY SOUGHT TO:

- **Assess** the state of AI agent deployments across the enterprise.
- **Identify** AI agent identity security processes and key pain points.
- **Uncover** the perceived business risks and identity security risks for AI agents.
- **Establish** the key constituents driving strategy and purchasing decisions for AI agent management and governance.

Note: Totals in figures and tables throughout this eBook may not add up to 100% due to rounding or organizations choosing more than one answer to select questions.

Key findings

AI agents are proliferating and permeating the enterprise

PAGE 4

Identity teams strive to gain visibility into the AI agent population, with mixed results

PAGE 10

Identity security drivers for AI agents include mitigating risk, ensuring compliance, and accelerating projects

PAGE 13

Achieving AI agent identity governance and management involves a changing toolset

PAGE 17

AI agents are the dominant enterprise priority, and the CIO, CISO, and CTO drive identity security strategy and investments

PAGE 20


AI agents are proliferating and permeating the enterprise

AI agents emerge as a strategic priority

While AI generally has captured the imagination of business, IT, and cybersecurity decision-makers, eight in 10 organizations report that AI agents are the top (24%) or a high (58%) priority relative to other AI initiatives. AI agents provide an optimal way to improve productivity and efficiency and achieve the desired return on investment (ROI) for AI initiatives, but teams will be under scrutiny to accelerate AI projects while ensuring secure and well-governed AI agent deployments.

Prioritization of AI agents compared with other AI initiatives.

24%	Top priority (i.e., AI agents are prioritized above all other AI priorities)
58%	High priority (i.e., AI agents are one of several high AI priorities with significant efforts directed toward AI agent integration into several areas of business)
16%	Moderate priority (i.e., AI agents are prioritized with ongoing efforts to integrate where sensible, but other AI investment areas hold higher priority)
2%	Low priority (i.e., AI agents receive some consideration but are not a significant focus compared with other organizational AI priorities)

Back to Contents

AI agents are widely deployed

Organizations are embracing AI agents across various functions, with IT operations providing the most frequently cited functional area for AI agents.

Identity security and IAM leaders typically have visibility into projects across their environments and are tasked with ensuring that agent deployments are both secure and well-governed to meet security, corporate governance, compliance, and ROI goals.

Identity and access management (IAM) teams will need to establish common AI agent governance and management for diverse AI agent fleets to scale and accelerate secure AI agent deployments.

Departments or functional areas currently using, piloting, or planning to use AI agents.

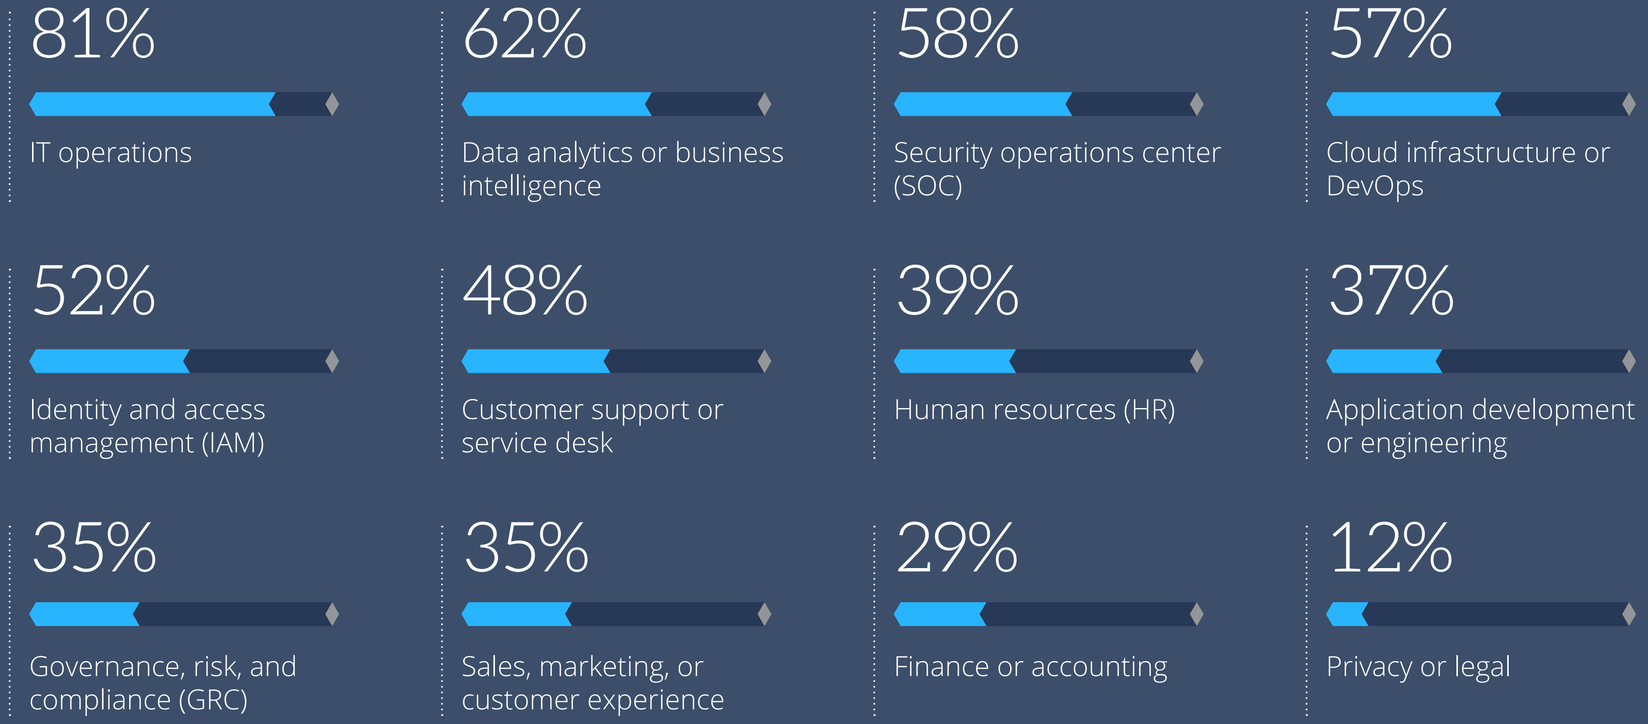

Back to Contents

AI agents are supporting an array of use cases

AI agent use cases are diverse and distributed across corporate IT ecosystems. Identity teams need to engage inside and outside of their organizations, whether through an “AI steering committee” or by collaborating with other AI agent constituents, to help streamline secure development and deployment for their AI agent fleets. Identity teams have frequently been perceived as an obstacle to rapid project rollouts in the past; however, AI agents provide an opportunity to change that perception. Identity teams can facilitate and accelerate projects through common identity infrastructure, policies, and management.

Highest priority use cases for AI agents.

Back to Contents

AI agent projects proliferate

Overall, organizations are in the middle of 22 distinct AI agent projects on average, each of which can involve multiple agents. Ensuring consistent, scalable management and governance for these projects will drive identity teams to establish common identity security processes.

Number of distinct AI agent projects, workflows, or deployments.

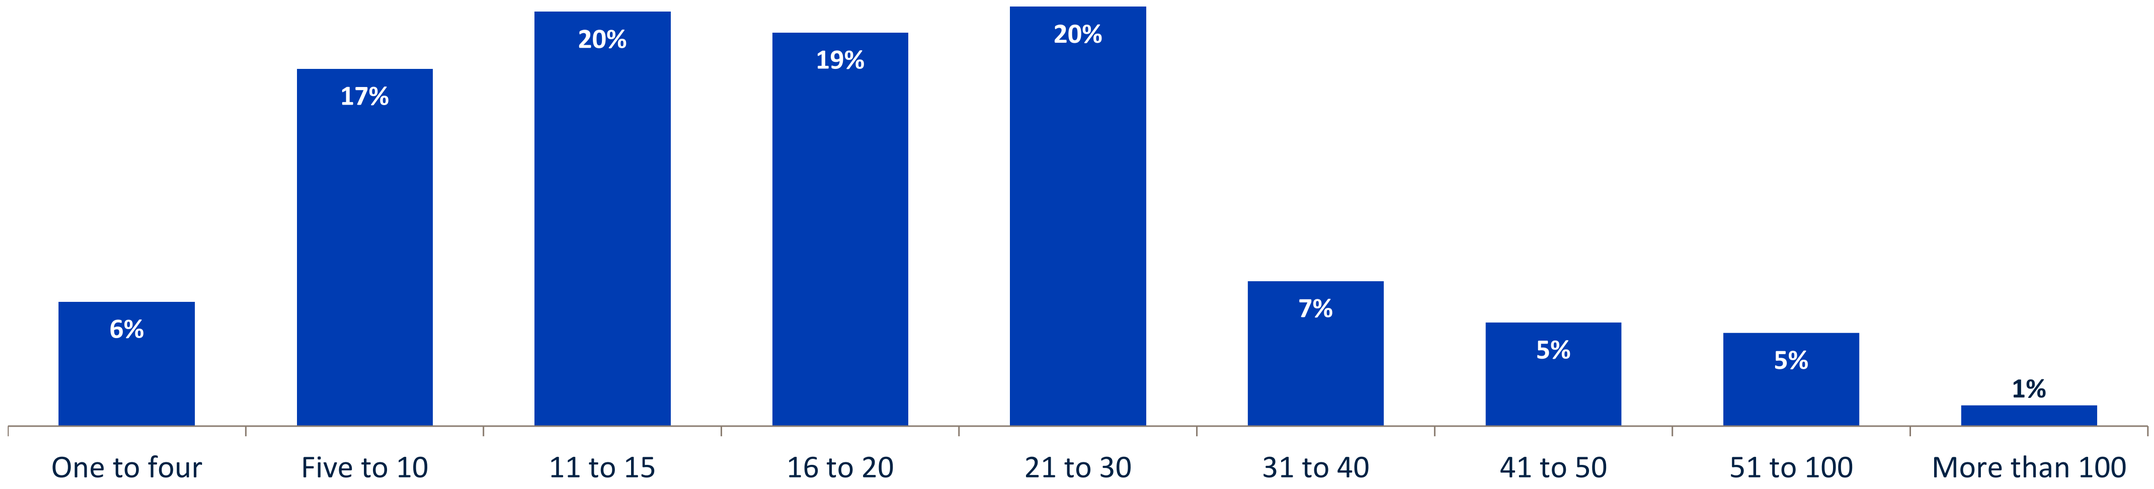

Back to Contents

AI agent deployment locations often include public cloud and SaaS

Cloud service provider AI agents and SaaS provider AI agents are the most common locations prioritized for deployments. SaaS provider agents represent easy wins for organizations that can deploy a ready-made agent in the SaaS "walled garden" to improve their SaaS application productivity.

AI agents are a recent phenomenon. As organizations get more experience with AI agents, expect the public cloud AI agent utilization to grow relative to most other locations given significant opportunity to deliver greater business value. However, that value also comes with greater security risk given the sensitive data stores being accessed.

Identity teams need holistic solutions for AI agent governance and management that can operate across the many locations where AI agents are deployed. This may require new identity technology tools for new AI agent use cases to supplement existing tooling.

Deployment locations or types prioritized for AI agents.

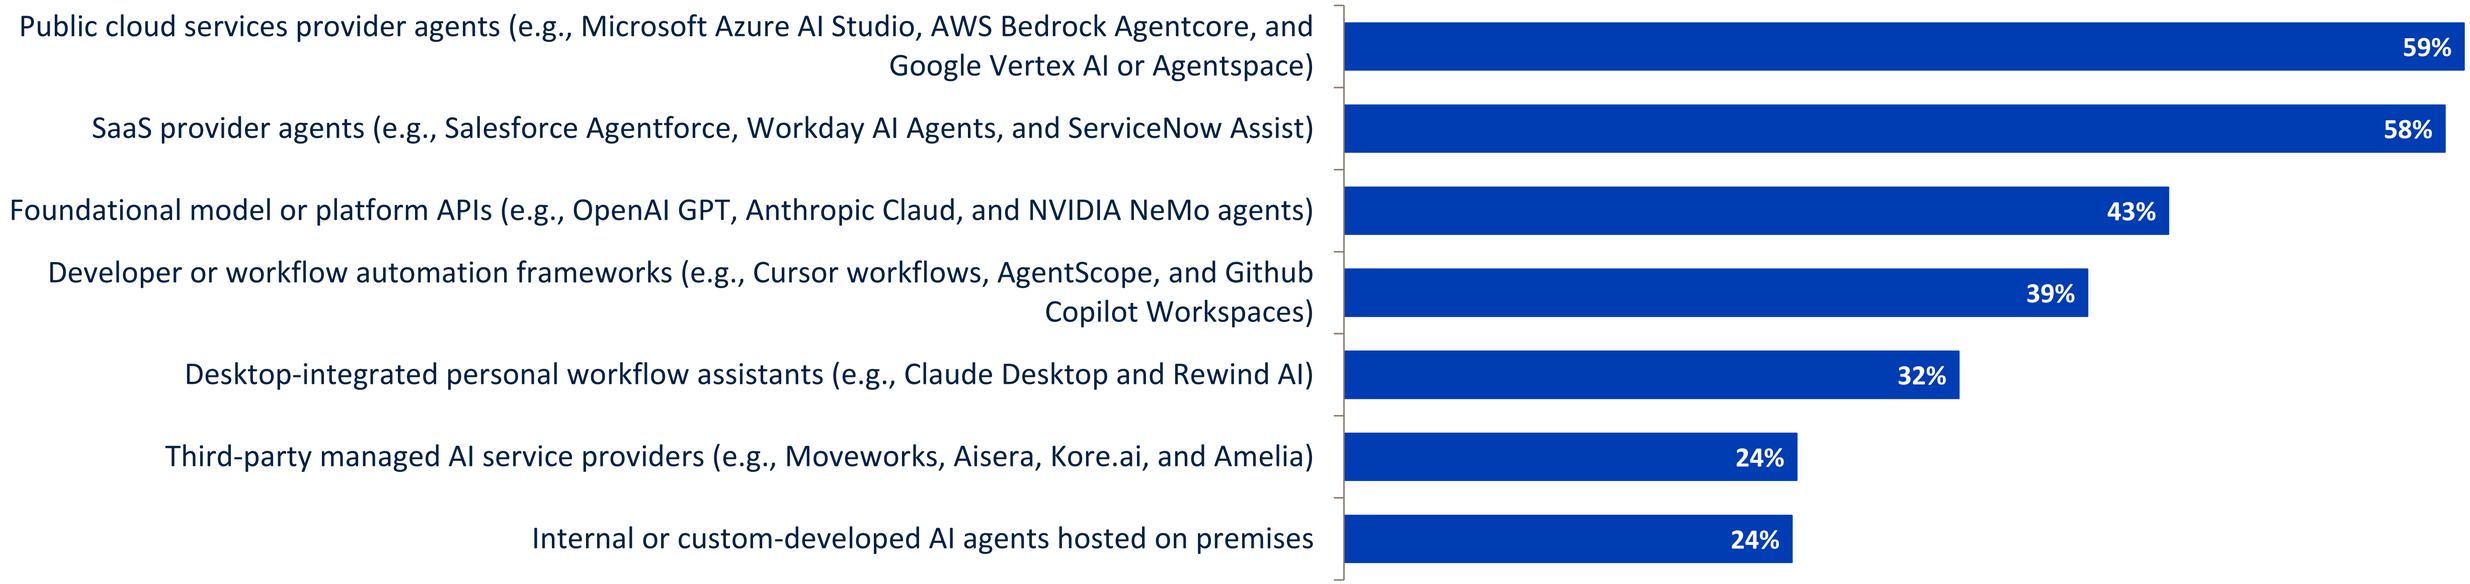

Back to Contents

Identity teams strive to gain visibility into the AI agent population, with mixed results

AI agent visibility varies, with gaps emerging

While more than half (51%) of respondents believe they have complete insight into their agents and their interconnectedness across the environment (agent-to-agent communication and MCP server communication), 44% indicate they have complete visibility to just a subset of their AI agent population. Given how diffused and decentralized AI agent utilization can be, these perceptions will change as the "shadow AI" phenomenon becomes better understood.

Visibility organizations have into the interconnectedness of AI agents with other agents or with MCP servers.

Complete visibility to all AI agents

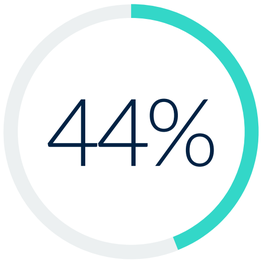

Complete visibility to some AI agents

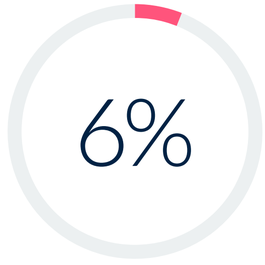

Limited visibility to all AI agents

Back to Contents

Discovering model context protocol (MCP) servers often involves existing tools

Organizations are using existing tools to discover MCP servers across their environments, with CSPM and EASM tools being the most common approaches, followed by network management tools, endpoint or container management tools, and API cataloging tools. Leveraging existing tooling provides an effective way to deliver visibility and facilitate control for MCP infrastructure.

Tools and processes for discovering MCP servers.

65% Cloud security posture management (CSPM) and external attack surface management (EASM) tools

59% Network discovery/surface management (tools discovering MCP servers exposed on internet-facing or internal networks)

58% Endpoint/container security or management tools (EDR or other tools that can enumerate MCP tooling)

46% Package/API cataloging tools (tools that track AI libraries, models, and external API calls)

29% Manual registration in MCP registry or API inventory

28% MCP gateway with discovery functionality

2% We don't currently use MCP

Back to Contents

Identity security drivers for AI agents include mitigating risk, ensuring compliance, and accelerating projects

Data privacy and security vulnerabilities top the list of AI agent business risks

AI agents represent a significant expansion of the enterprise attack surface, and the nature and implications of that expansion are still emerging. Data privacy and security vulnerabilities are the most prominent risks cited by around one-third of respondents.

The other business risks are recognized but not as pressingly. As AI agents get deployed and the threat landscape becomes clearer, expect some shifting in risk priorities.

Compliance and regulatory risks historically take time to catch up with new technology, and that will probably be the case with fast-evolving AI agent deployments that raise novel regulatory issues.

General risks related to AI agents that cause the most concern for organizations.

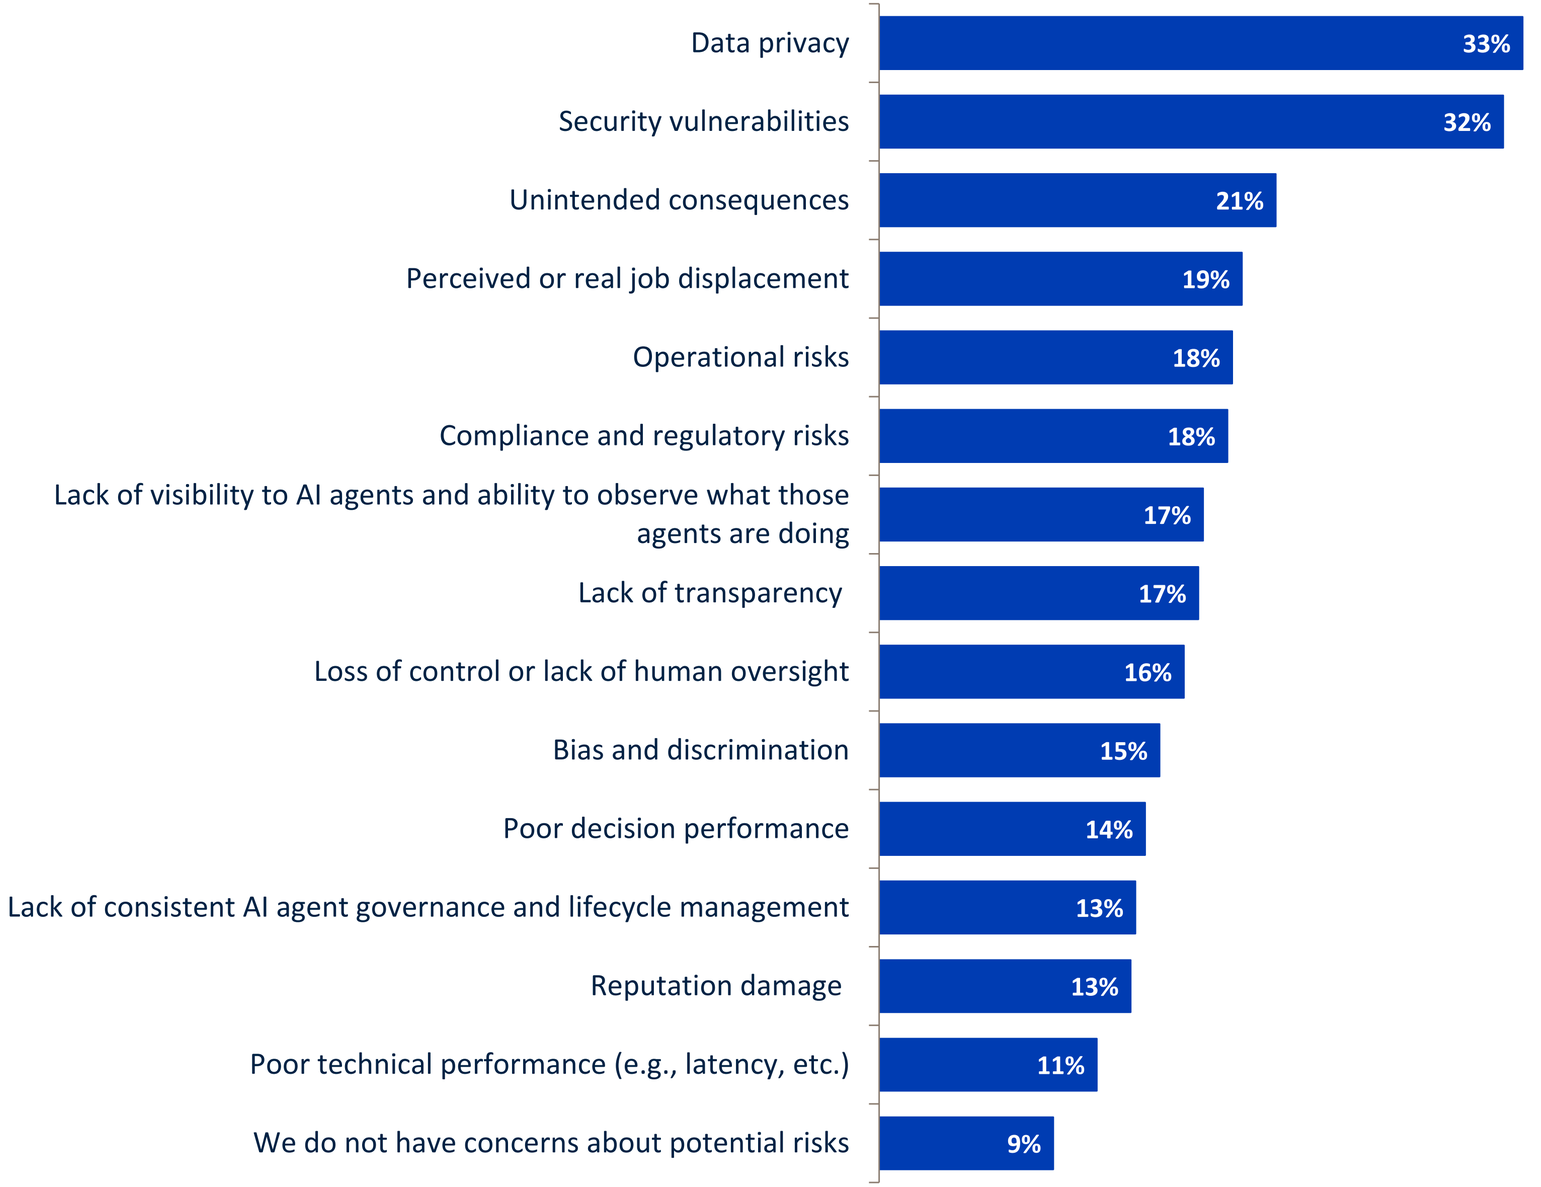

Back to Contents

AI agents pose identity security risks

Identity security is a pillar of any AI agent security strategy, and it counters risks that are different than other pillars of AI agent security like data security and AI security posture management. Identity teams share concerns around security vulnerabilities and data privacy but also identity-specific risks such as AI agent visibility (i.e., inventorying agents and observing their behavior), the potential for cross-protocol confusion, and the loss of control or lack of human oversight.

As identity security vendors innovate, they will need to solve for the gamut of identity issues posed by AI agents. Identity teams will need to be ruthless in meeting their objectives and choose the optimal tools to get the job done.

Identity security risks related to AI agents that cause the most concern for identity security teams.

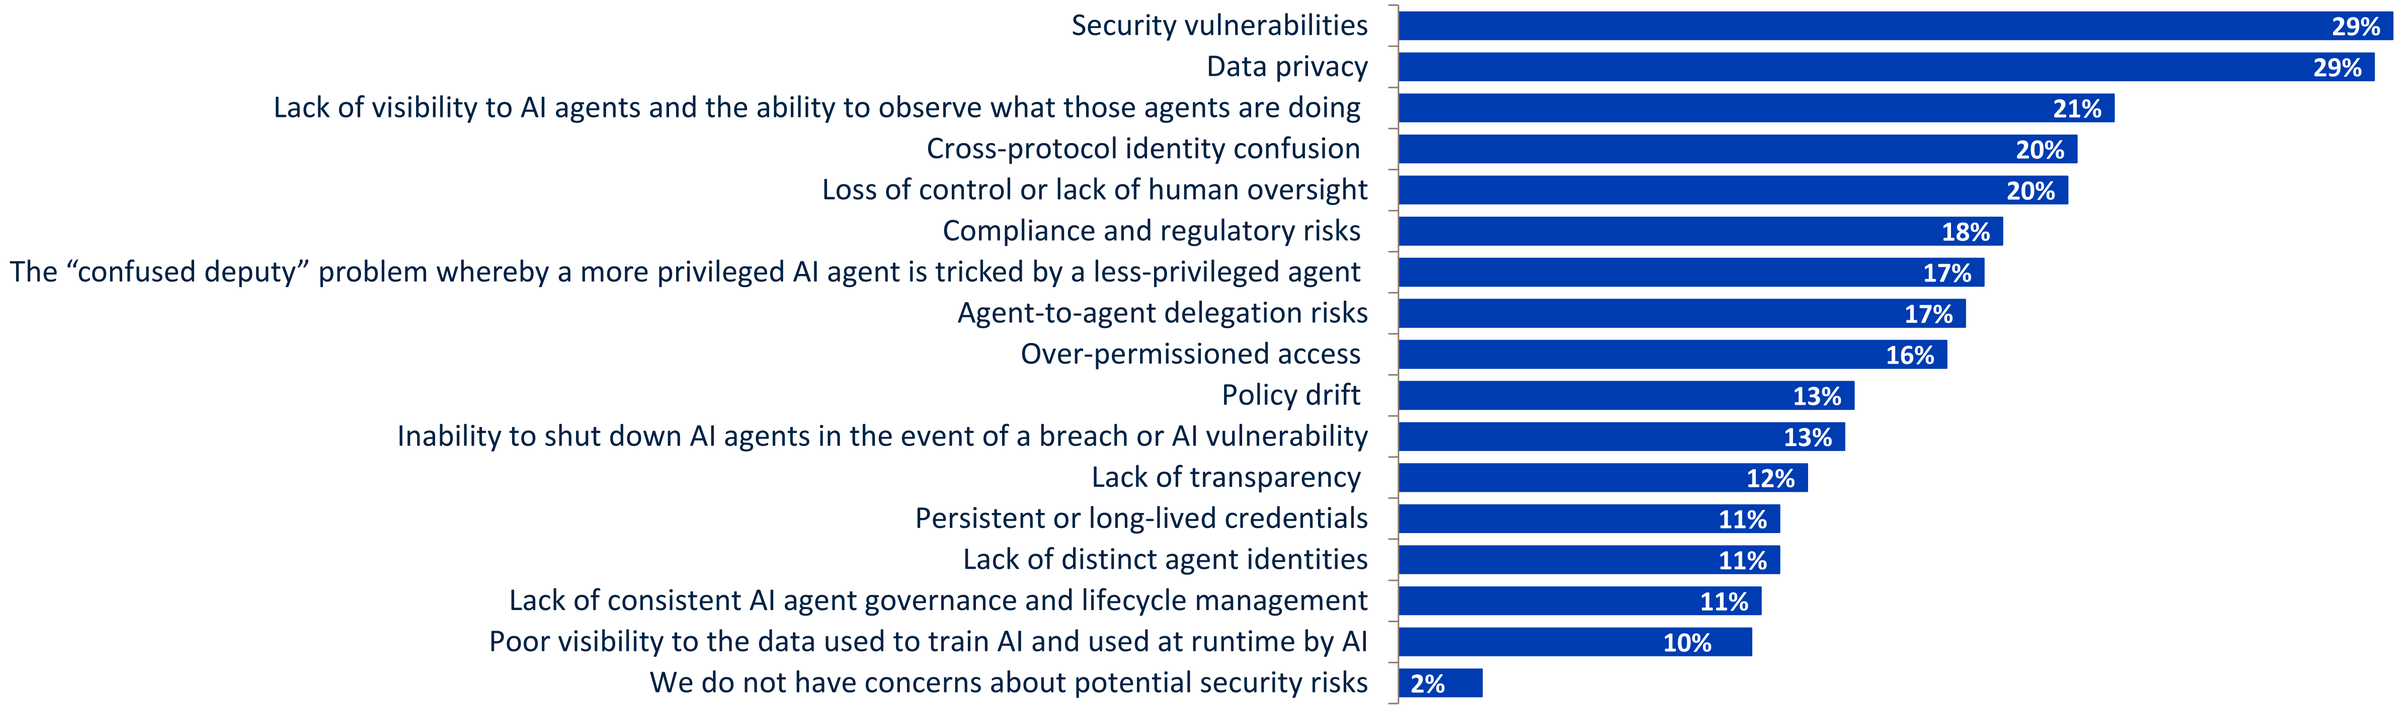

Back to Contents

A range of considerations drive identity security for AI agents

Organizations see compliance, mitigating cybersecurity risk, and operational efficiency as primary requirements driving identity security for AI agents, but they also expect IT governance benefits from having consistent management for their AI agent fleets.

Identity teams have an opportunity to accelerate AI agent development and deployment through consistent tooling and processes across the variety of AI agent projects.

Considerations influencing strategies to provide identity security for AI agents.

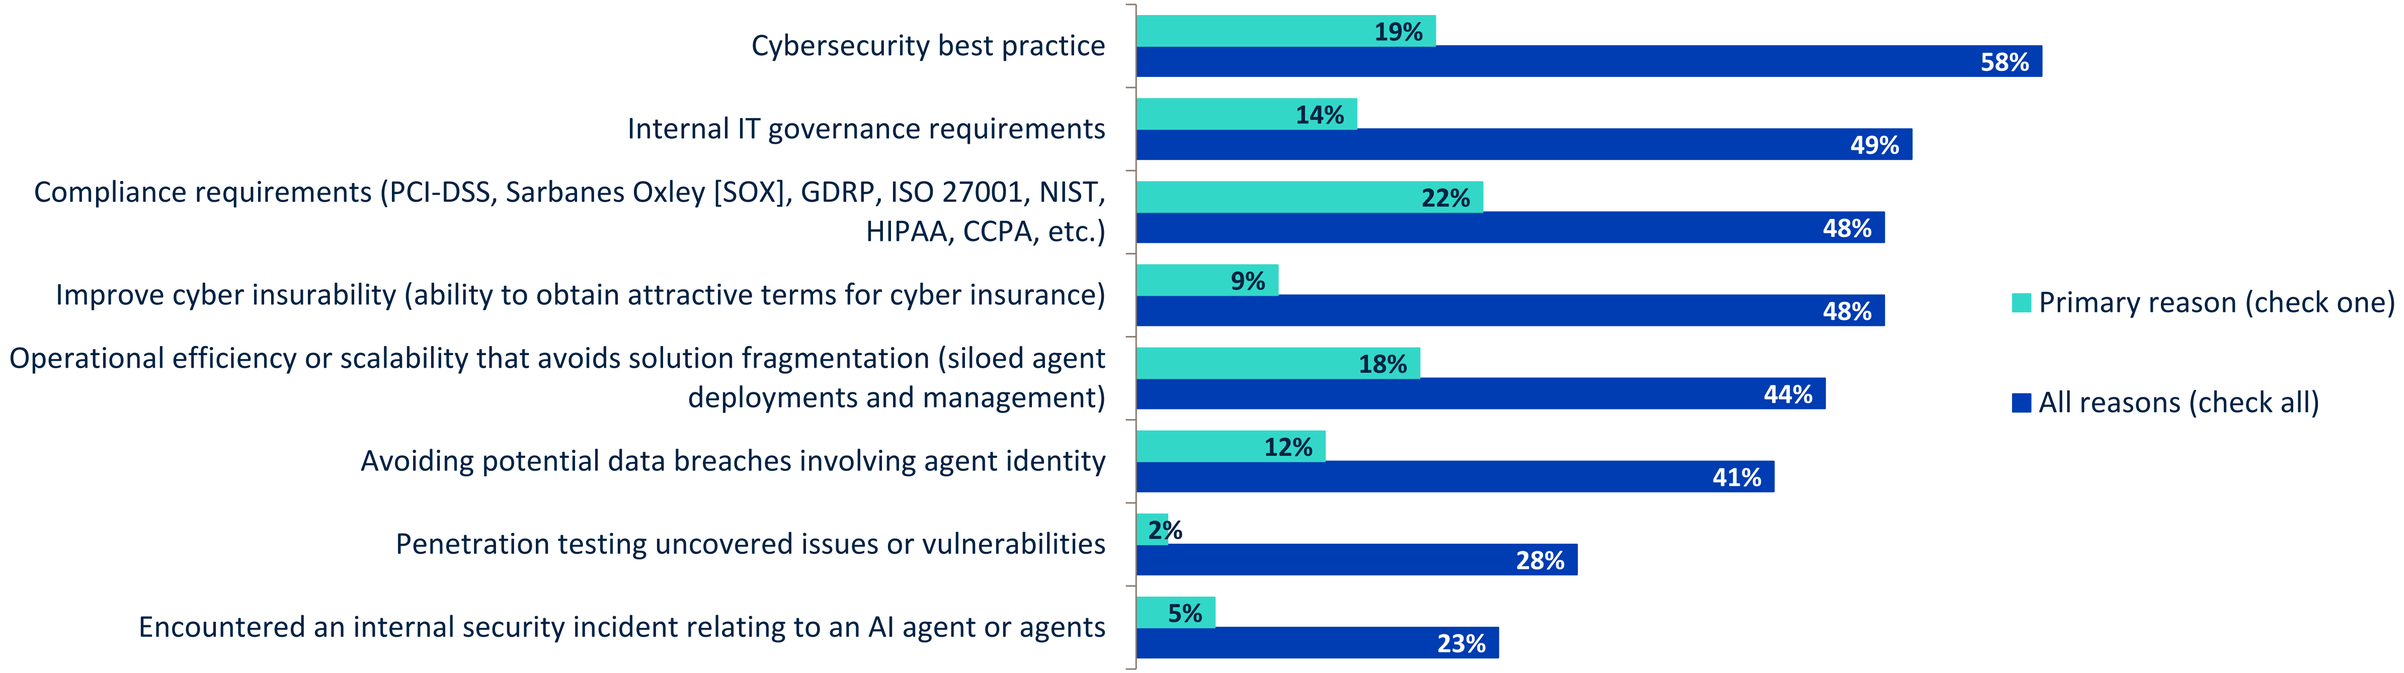

Back to Contents

Achieving AI agent identity governance and management involves a changing toolset

The perceived optimal AI agent solution approach is in flux

Organizations want both efficiency and effectiveness in any technology solution they deploy, and teams typically prefer extending an existing solution to cover a new use case rather than deploying a new tool.

Since AI agent security is a new issue, there is confusion and flux around data security versus identity security as well as what comprises a complete stack for AI agent identity security, governance, and management.

Security issues like prompt injection attacks and the potential for data loss in agentic infrastructure have preoccupied organizations and have probably resulted in data security platforms being selected as an optimal solution type. While a solution stack will have multiple elements, expect more identity-centric solutions to gain prominence as enterprises recognize the centrality of identity in securing AI agents.

Perceived optimal solution type to provide AI agent visibility, observability, access management, governance, and lifecycle management.

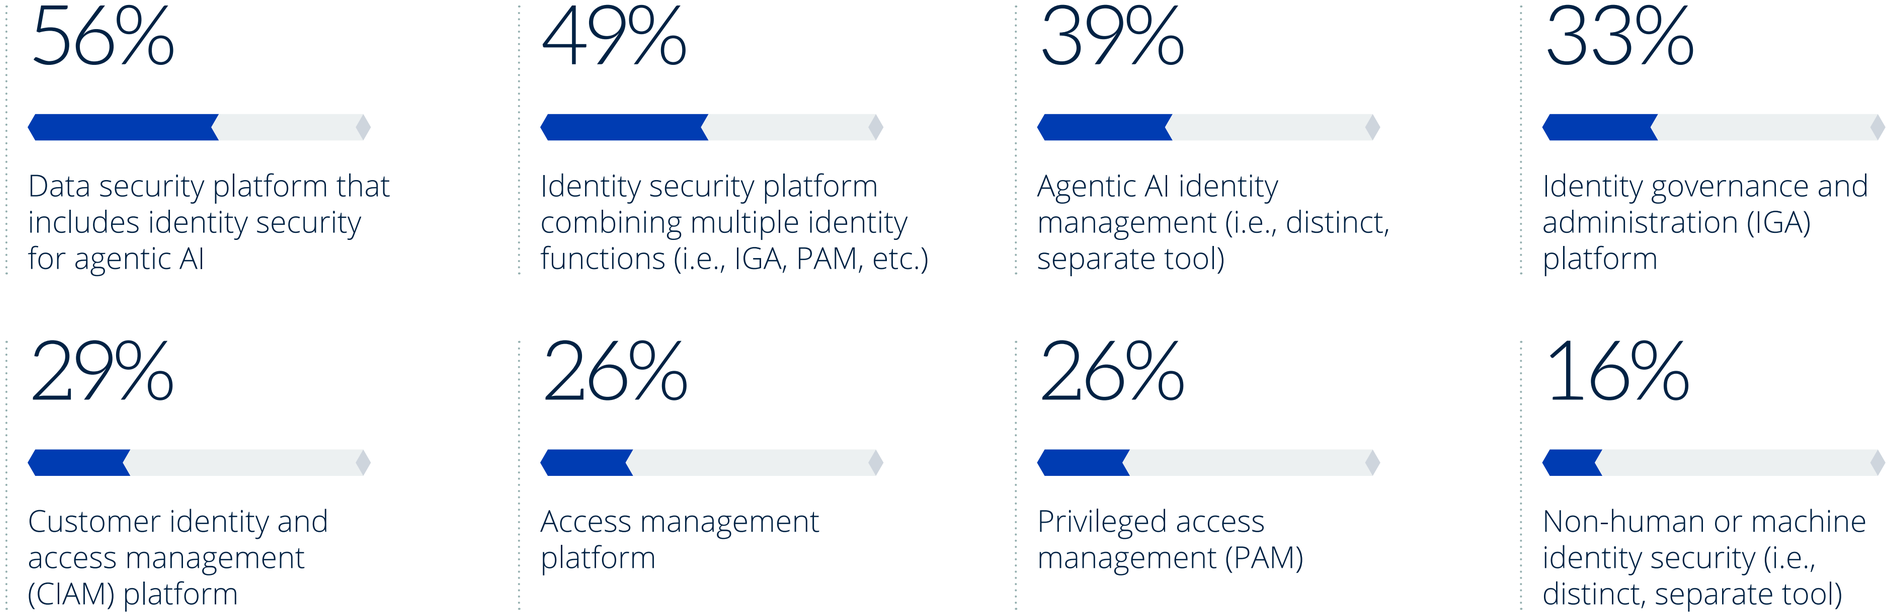

Back to Contents

Capabilities desired for AI agent management

AI agents represent a new type of identity that requires management and governance. While security will start with visibility to the AI agent population, organizations also need lifecycle management, automated remediation, and continuous authorization for ephemeral AI agent identities deployed at scale. And this needs to be achieved across an environment that includes cloud, SaaS, on-premises infrastructure, and endpoints. AI agents are new terrain for many identity teams and their AI infrastructure peers, and it can be expected that this capability priority list will shift in response to security threats, data compromises, and compliance requirements that emerge.

Capabilities organizations are prioritizing to secure and govern AI agent identities.

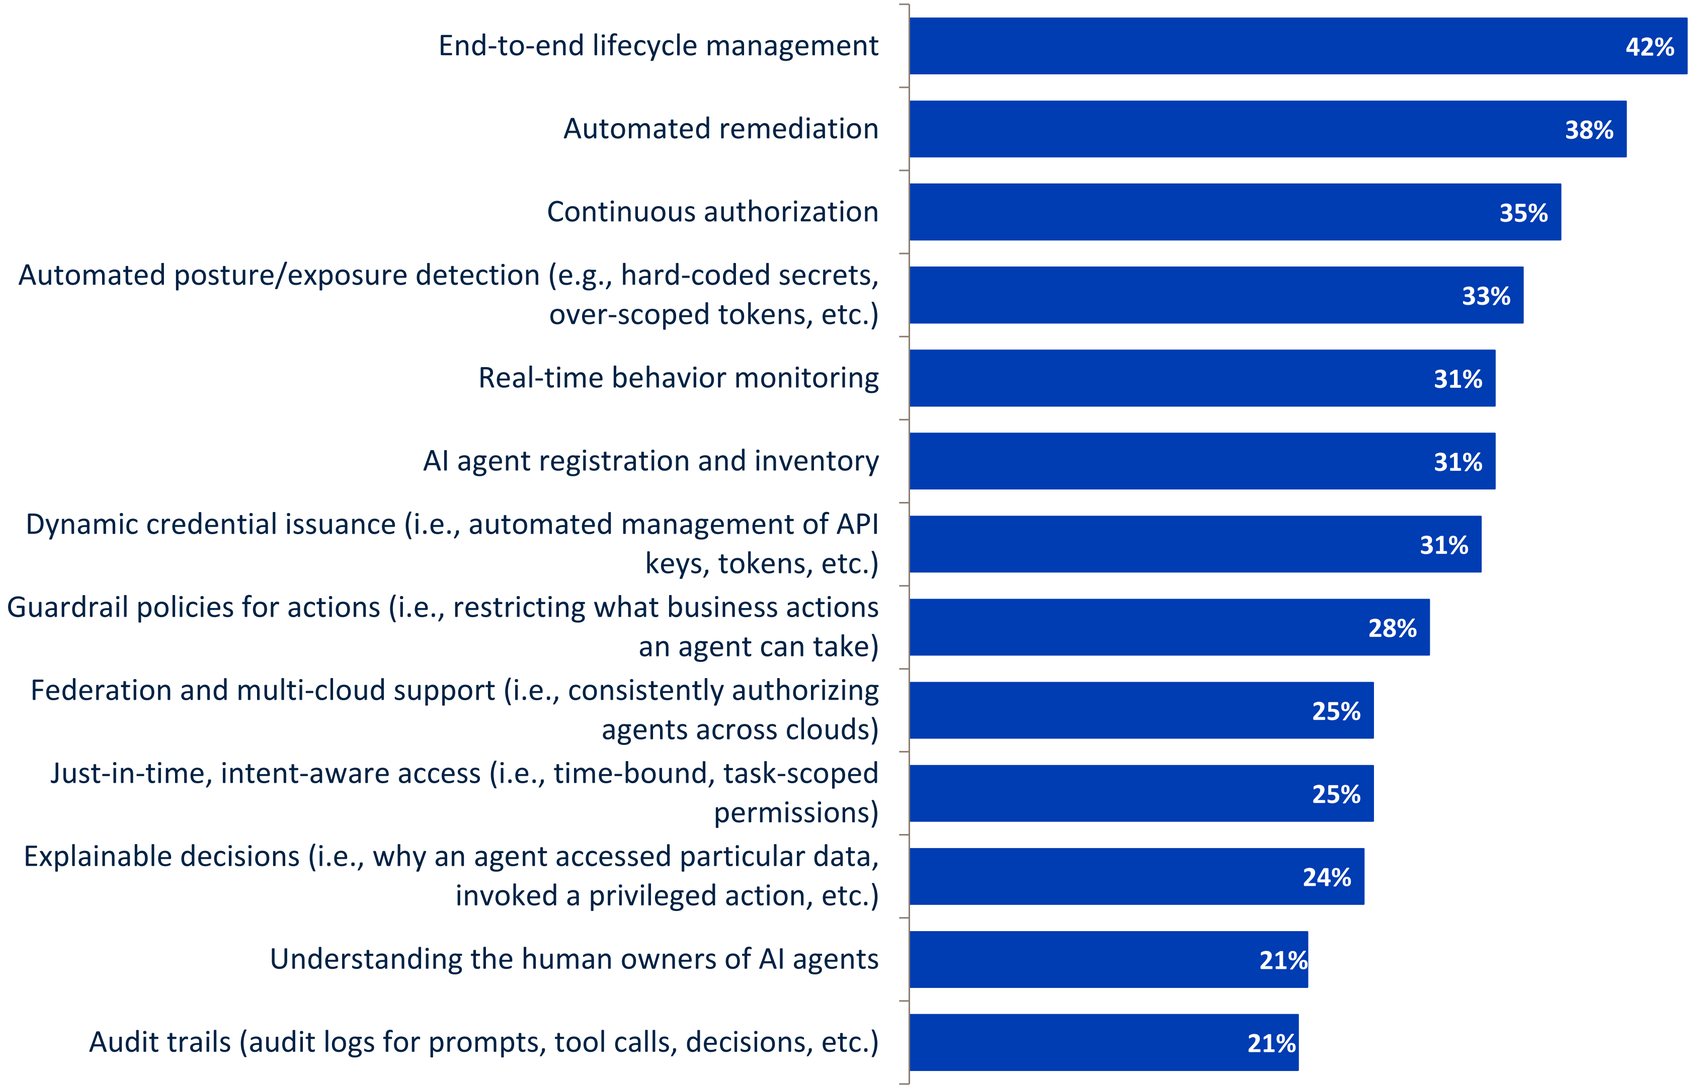

Back to Contents

AI agents are the dominant enterprise priority, and the CIO, CISO, and CTO drive identity security strategy and investments

CISOs are the most common AI agent security decision-makers

Organizations are frequently looking to security and identity leaders like the CISO and CIO for AI agent security decisions; however, the CTO is also commonly the primary decision-maker for AI agent strategy and planning generally, including security decisions.

AI agents are a recent phenomenon with new constituents, and the CTO or VP driving AI agent initiatives has become a significant persona with dedicated budget to facilitate all aspects of AI agent projects, including security and identity. Security and identity vendors that can understand and speak to this persona's concerns and preferences will be best positioned to add value.

Primary ownership of security for AI agents.

25%	CISO or cybersecurity team (including cloud security architects)
23%	Chief information officer (CIO, including cloud architects, application development, and platform engineering)
22%	Chief technology officer (CTO)
18%	An AI committee with cross-organization representation that establishes plans and strategy across the organization
9%	Regulatory compliance and risk management
2%	Line-of-business and application owners
2%	A disparate organization approach with no clear owner

Back to Contents

Security investment inhibitors are few and far between

Organizations are actively investing to secure their AI agent fleets and have few major inhibitors to those investment efforts. No investment concern garnered an inordinate amount of interest, with any issue garnering a quarter or less of respondents. While the threat landscape needs to become clearer for some respondents, an equal number of respondents indicated that nothing is holding back investments around identity security for AI agents.

As security threats to AI agents emerge and compliance mandates around AI agents become clearer, expect organizations to increase identity security investments for AI agents.

Issues holding organizations back from investing in identity security for AI agents.

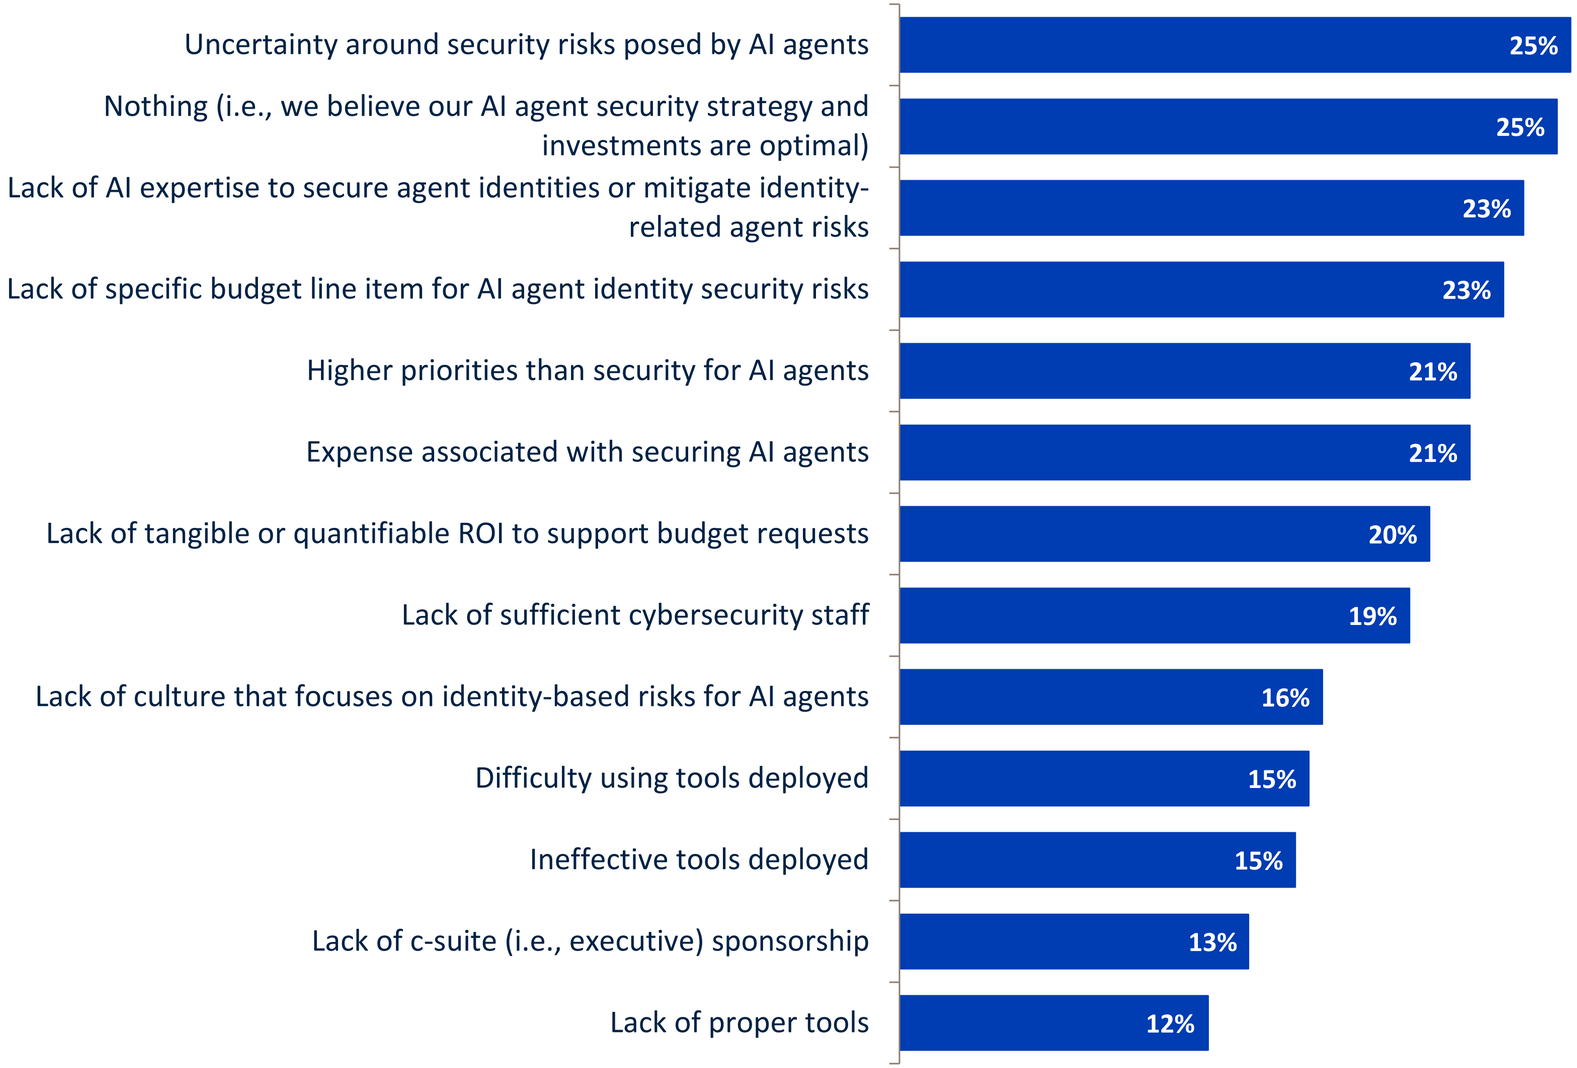

Back to Contents

Security for AI agents is a key priority

Security teams have finite resources and can take on a limited number of security projects in a given year. Securing AI agents is one of those projects in the coming 12-24 months. Indeed, securing AI agents is a top five priority for 93% of organizations and the top priority for 41%.

Priority level organizations will assign to security for AI agents over the next 12 to 24 months.

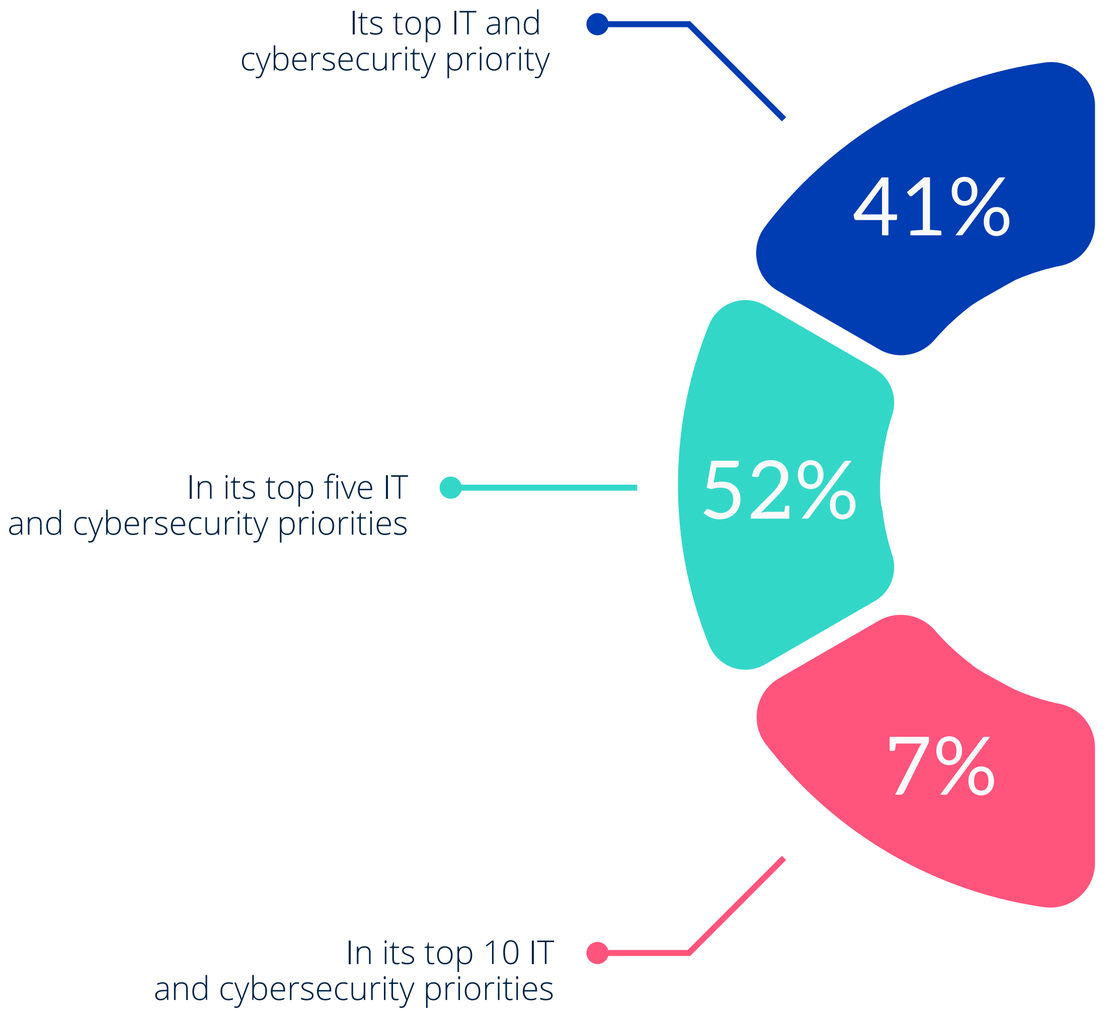

Back to Contents

BeyondTrust

About

BeyondTrust is the global identity security leader protecting Paths to Privilege™. Our identity-centric approach goes beyond securing privileges and access, empowering organizations with the most effective solution to manage the entire identity attack surface and neutralize threats, whether from external attacks or insiders.

BeyondTrust is leading the charge in transforming identity security to prevent breaches and limit the blast radius of attacks, while creating a superior customer experience and operational efficiencies. We are trusted by 20,000 customers, including 75 of the Fortune 100, and our global ecosystem of partners.

We believe AI should be a part of a company's identity program—not treated as an exception managed by yet another siloed tool. Learn how BeyondTrust applies visibility, intelligence, and protection to agentic AI as part of a holistic, single-platform approach to identity security: www.beyondtrust.com/solutions/ai-security.

Learn More

Research methodology and demographics

To gather data for this report, Omdia conducted a comprehensive online survey of IT and cybersecurity professionals from private- and public-sector organizations in North America between January 7, 2026, and January 15, 2026. To qualify for this survey, respondents were required to be involved with or responsible for identity security technology products and services. All respondents were provided an incentive to complete the survey in the form of cash awards and/or cash equivalents.

After filtering out unqualified respondents, removing duplicate responses, and screening the remaining completed responses (on several criteria) for data integrity, we were left with a final total sample of 400 of IT and cybersecurity professionals.

Respondents' organizations by number of employees.

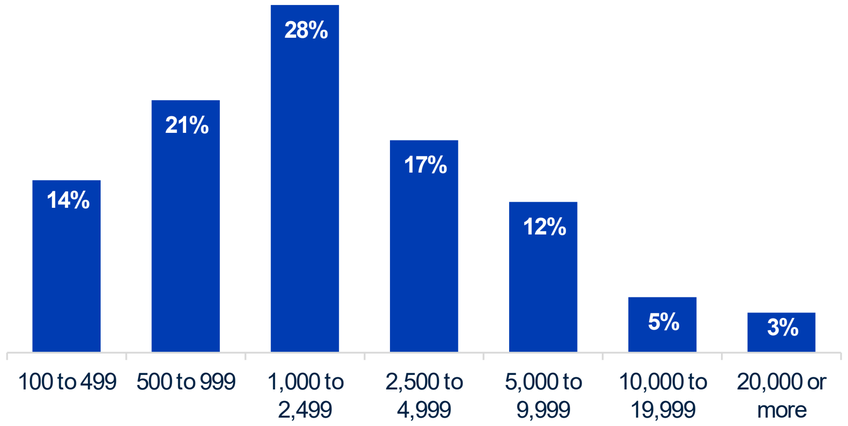

Respondents' organizations by years in operation.

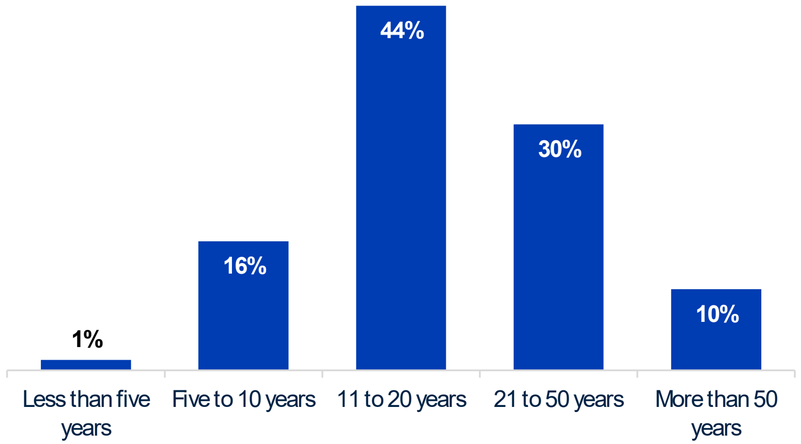

Respondents' organizations by industry.

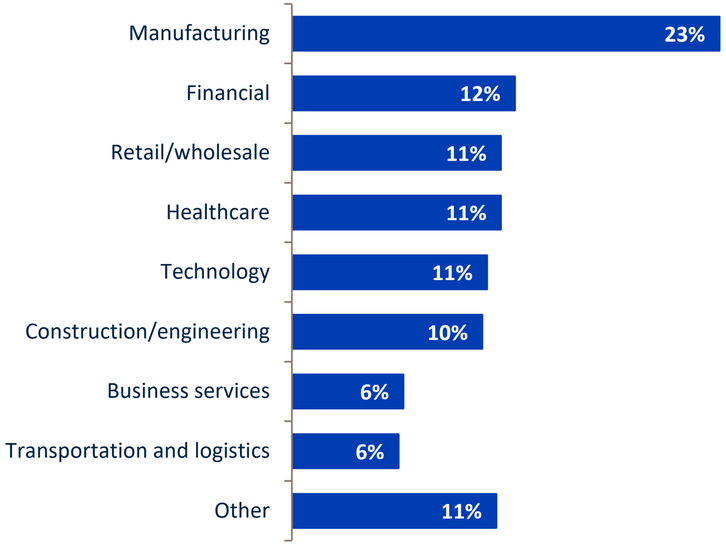

Back to Contents

Omdia provides focused and actionable market intelligence, demand-side research, analyst advisory services, GTM strategy guidance, solution validations, and custom content supporting enterprise technology buying and selling.