cloud security alliance®

aembit

Survey Report

Identity and Access Gaps **in the Age of Autonomous AI**

Acknowledgments

Lead Author

Hillary Baron

Contributors

Marina Bregkou
Josh Buker
Ryan Gifford

Graphic Design

Stephen Lumpe

About the Sponsor

Aembit is the identity and access management platform for agentic AI and workloads. It enforces access based on identity, context, and centrally managed policies, giving organizations a singular place to control access risk from AI agents, automate credential management, and accelerate AI adoption. With Aembit, enterprises can confidently control access to sensitive resources across all the workloads that power their business. You can learn more at aembit.io.

Table of Contents

Executive Summary

AI agents are rapidly becoming embedded across enterprise technology environments, interacting with applications, infrastructure, data platforms, and development pipelines. As these systems assume more autonomous operational roles, they introduce new challenges for identity management, access control, and governance. The report examines how organizations are managing AI agent identities and permissions—and whether existing identity and access management (IAM) models are keeping pace with their expanding operational role.

This survey reveals six critical insights:

1. AI Agents Are Already Operating Across Enterprise Systems

AI agents are widely deployed across enterprise workflows. Task automation agents are used by 67% of organizations, while roughly half report using research agents (52%), developer-assist agents (50%), and security or monitoring agents (50%). Only 15% report that AI agents are not used in production environments, and 73% expect agents to become very important or critical within the next year.

2. Most AI Agents Do Not Operate as Distinct Identities

AI agents often exist in an identity gray area—neither fully treated as human users nor consistently managed as first-class machine identities. While 52% use workload identities, 43% rely on shared service accounts, and 31% allow agents to operate under human user identities. As a result, 68% of organizations cannot clearly distinguish between human and AI agent activity.

3. No Single Team Clearly Owns AI Agent Identity and Access

Ownership of AI agent identity and access is fragmented. Security leads in 28% of organizations, followed by development or engineering (21%) and IT (19%), while only 9% identify IAM teams as the primary owner. This distributed ownership can lead to inconsistent controls and slower coordination when issues arise.

4. Confidence in AI Agent Access Exceeds Control Maturity

Organizations report moderate confidence in managing agent access, but operational controls reveal gaps. While 57% report moderate or high confidence in identity scoping, one-third of organizations are unsure how often AI agent credentials are rotated and 9% report they are rarely or never rotated. Only 22% report that access frameworks are applied very consistently to AI agents.

5. AI Agents Often Inherit Access and Expand the Attack Surface

AI agent access frequently derives from existing human permissions or automation logic rather than agent-specific entitlements. Nearly three-quarters agree that AI agents often receive more access than necessary (74%), and 79% believe agents introduce new access pathways that are difficult to monitor. More than half report that agents inherit access originally intended for humans or other systems at least sometimes.

6. Governance Mechanisms Are Compensating for Missing Identity Controls

Organizations frequently rely on procedural safeguards to control AI agent activity. Disabling identities or revoking tokens (49%) are the most common containment actions, while 42% report terminating the compute environment in which an agent runs. Looking ahead, organizations prioritize real-time visibility into agent actions (52%) and clearer identity separation between agents and humans (45%) to support safer scaling.

Takeaway

AI agents are already operating across enterprise environments, yet identity and access practices have not fully adapted to manage them. Many organizations rely on inherited permissions, fragmented ownership, and governance safeguards rather than identity-centric controls. As AI agents gain autonomy and operational scope, extending core IAM principles—clear identity separation, least privilege access, and continuous visibility—will be critical to managing risk and enabling secure adoption.

Key Findings

AI agents are rapidly becoming a functional layer within enterprise technology environments, interacting with applications, infrastructure, data platforms, and development pipelines. As these systems take on more autonomous tasks and integrate more deeply into operational workflows, they introduce new considerations for identity, access control, accountability, and governance. Understanding how organizations are managing these dynamics is essential for assessing whether existing security and IAM models are equipped to support the next phase of AI-enabled operations.

Key Finding 1:

AI Agents Are Operating Across Core Enterprise Systems

Agentic AI is already embedded in enterprise environments and expanding across critical workflows. Task-automation agents are reported by 67% of organizations, making automation the most common entry point. Adoption, however, extends well beyond efficiency use cases. More than half report data-retrieval or research AI agents (52%), and half report both code-generation or developer-assist AI agents (50%) and security or monitoring AI agents (50%). Infrastructure or IT operations AI agents are in use by 41%, while only 18% indicate no use of AI agents at all. This distribution indicates that AI agents are already influencing development workflows, security monitoring, and infrastructure operations. These areas typically require permissioned interaction with production systems. AI agent activity is therefore occurring in environments where identity scoping and access control have operational consequences.

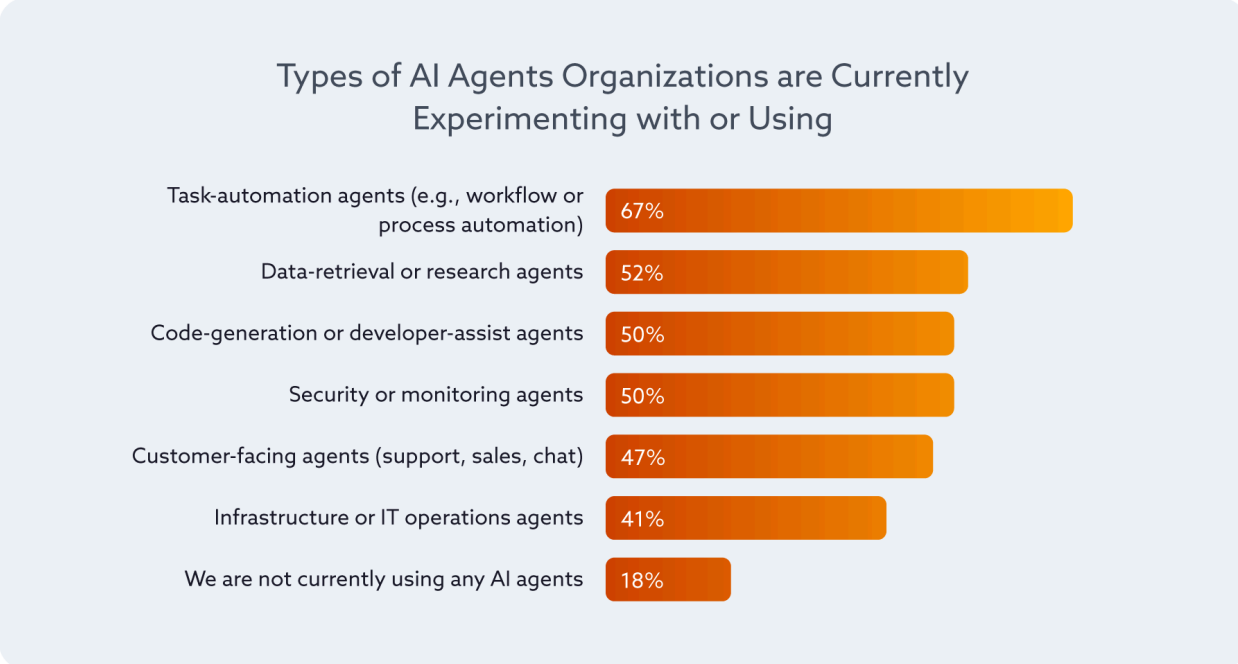

The environments in which AI agents interact further reinforce this operational footprint. A majority report interaction with internal applications or APIs (56%), and nearly half report activity within SaaS applications (49%) and cloud infrastructure (44%).

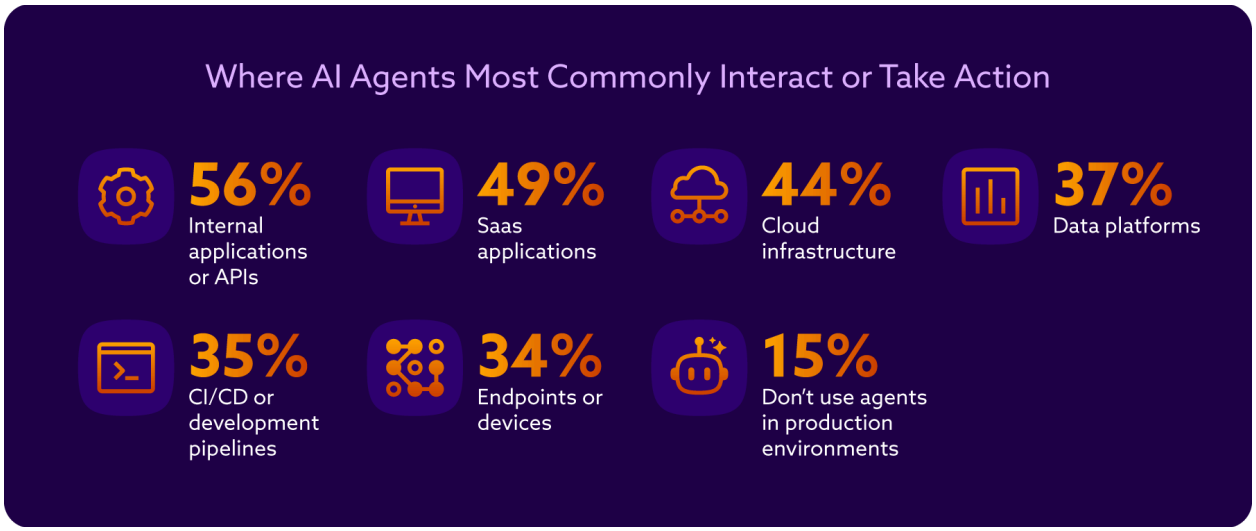

More than one-third indicate interaction with data platforms (37%), CI/CD or development pipelines (35%), and endpoints or devices (34%). Only 15% report that AI agents are not used in production environments, suggesting that most deployments extend beyond isolated test settings. As AI agents operate across internal systems, external platforms, and infrastructure layers, the number of integration points and access pathways expands. This cross-environment interaction increases the complexity of maintaining consistent identity governance and permission boundaries.

Forward-looking expectations suggest that this footprint is likely to deepen. A majority expect AI agents to become "very important" (54%) or "critical" (19%) within the next 12 months, and only 8% view them as "slightly important" or "not important." The concentration of responses in the upper tiers of importance indicates that agentic systems are expected to assume a more central role in business and technical operations. As strategic importance increases, integration across workflows and environments is likely to expand, further elevating the importance of identity and access governance.

Active deployment across multiple AI agent types, interaction with production systems, and rising strategic importance together indicate that agentic AI is no longer peripheral. It is operating within core enterprise environments today and positioned to become more central in the near term, increasing the scope and significance of how AI agent identities and permissions are managed.

Key Finding 2:

Most AI Agents Do Not Operate as Distinct Identities—They Borrow One

AI agents do not fit cleanly into existing human or traditional machine identity categories, and current security models reflect that ambiguity. It appears they aren't treated as human users nor first-class machine identities with any consistency. When asked how an AI agent's identity is typically represented when it performs an action, 52% of organizations report using an application or workload identity, 43% report using a shared or generic service account, and 36% report assigning the AI agent its own dedicated identity. At the same time, 31% indicate that AI agents operate under a human user's identity, while 12% are unsure how the identity is represented. No single model dominates, and multiple approaches often coexist within the same environment. Rather than reflecting a clear taxonomy, this distribution suggests a patchwork of identity treatments applied to AI agents.

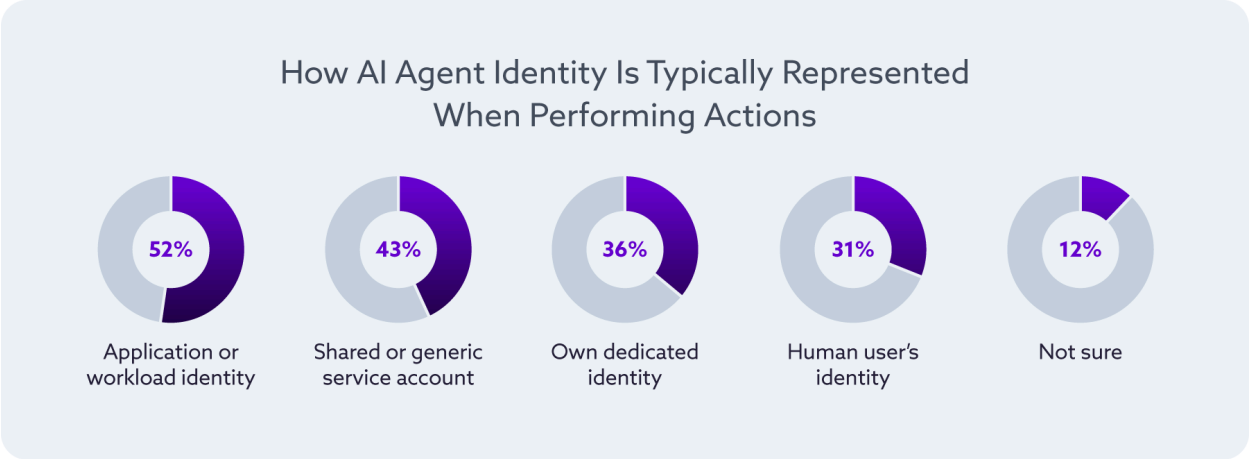

This fragmentation is not merely a naming or categorization issue. The way an AI agent is represented directly shapes how permissions are granted and enforced. When AI agents operate under human identities or shared service accounts, they inherit the full set of permissions attached to those accounts, whether or not those permissions align with the AI agent's intended role. Over time, this inheritance can unintentionally expand access scope and introduce policy drift, particularly when the underlying accounts were originally provisioned for interactive users or broader operational responsibilities.

Even in environments that use workload identities, the coexistence of multiple identity models complicates consistent policy application and attribution. When AI agents are sometimes treated as users, sometimes as services, and sometimes as standalone identities, enforcing uniform access boundaries becomes more difficult. This structural inconsistency does not remain abstract; it directly influences how clearly organizations can trace AI agent behavior in practice.

The impact becomes more visible when examining attribution. Sixty-eight percent of organizations report not being able to very clearly distinguish between actions performed by AI agents and those performed by humans. In contrast, only 32% are able to very clearly distinguish between actions performed by AI agents and humans. Strong, unambiguous attribution is therefore not universal. In environments where AI agents operate under human or shared identities, separating AI agent-initiated activity from human activity becomes more challenging, introducing friction for monitoring, forensic analysis, and accountability.

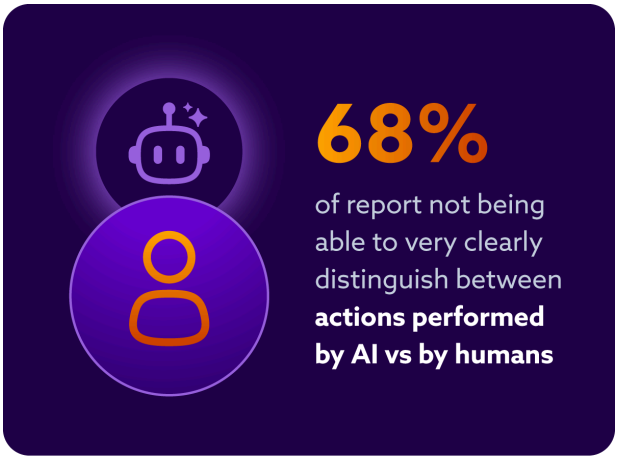

The ambiguity around AI agent attribution is compounded by inconsistent definitions across teams. Sixty-three percent agree that different teams describe AI agents in inconsistent ways (43% somewhat agree and 20% strongly agree). At the same time, 64% agree that policies explicitly differentiate between AI agents and human users (34% somewhat agree and 30% strongly agree), and 69% agree that they can trace an AI agent's actions back to the underlying context or initiator (44% somewhat agree and 25% strongly agree).

These results indicate that formal policies and traceability mechanisms often exist, yet conceptual alignment and operational clarity are not uniform. Agreement with policy differentiation coexists with acknowledgment of inconsistent terminology and interpretation, suggesting that documentation alone does not guarantee consistent implementation.

Across identity representation, attribution clarity, and definitional consistency, a clear pattern emerges: AI agents occupy an ambiguous position within existing identity frameworks. They are neither fully treated as human users nor consistently managed as first-class machine identities. This structural ambiguity shapes how permissions are assigned, how actions are traced, and how policies are enforced. Governance structure further reinforces this pattern.

Key Finding 3:

No Single Team Owns AI Agent Identity and Access

Ambiguity extends beyond taxonomy into organizational accountability. Responsibility for determining how AI agents authenticate and access systems is dispersed across multiple functions: security is identified as the primary owner by 28%, development or engineering by 21%, and IT by 19%, while only 9% point to identity or IAM teams as primarily responsible and another 9% report no clear owner. No single department emerges as the center of control.

When governance spans multiple domains without clear central accountability, consistent implementation and control maturity can be difficult to achieve. In practice, this means that how AI agents are provisioned and managed may differ from one team or environment to another. When no single function clearly owns agent identity and access, teams may rely on existing human or shared accounts, and controls may not be applied uniformly. Over time, this can lead to inherited permissions, uneven enforcement, and slower coordination when an agent's behavior needs to be reviewed or corrected. When something goes wrong, such as an agent acting outside its intended scope or exposing sensitive data, unclear ownership can delay decision-making, complicate remediation efforts, and make it harder to determine who is responsible for containment and follow-up.

Diffuse governance structures do not inherently signal immaturity. However, they do introduce variability. Without clear centralized accountability for AI agent identity and access, implementation practices may evolve unevenly. This could result in differences with credential hygiene, enforcement, revocation, and other controls being implemented.

Key Finding 4:

Confidence in AI Agent Access Exceeds the Controls Behind It

Organizations report moderate confidence in how AI agents are identified and controlled, yet underlying control practices reveal uneven visibility, operational friction, and distributed accountability. When asked how confident they are that AI agents have clearly defined identities and appropriately scoped access to systems and data, 40% report being moderately confident and only 17% very confident. Nineteen percent indicate slight confidence, while 21% report not being confident and 3% are unsure. A majority therefore express at least moderate confidence in identity scoping and access management. However, the relatively small share reporting very high confidence suggests that assurance may be qualified rather than absolute. Confidence may reflect familiarity with existing workflows rather than consistent evidence of mature, systematically enforced controls.

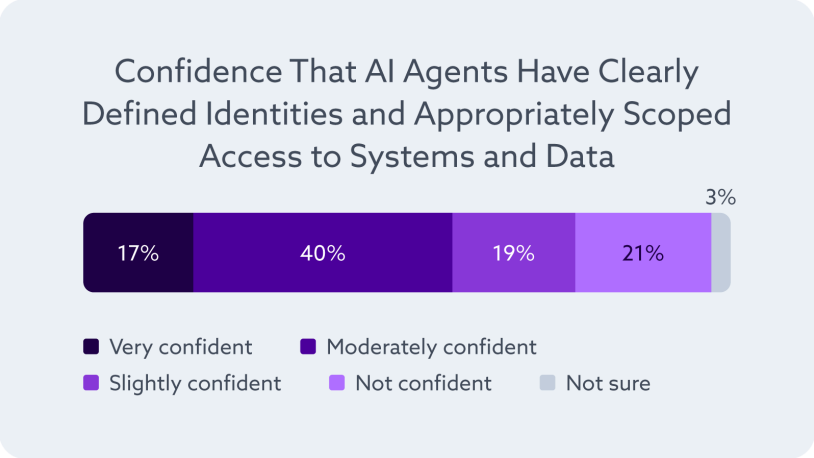

A similar pattern appears in revocation confidence. Thirty-five percent report being moderately confident that they could rapidly revoke an AI agent's access if it behaved unexpectedly or maliciously, and 29% report being very confident. Twenty percent indicate slight confidence, 10% not confident, and 6% unsure. While many believe revocation is achievable, strong assurance is not universal. Revocation may depend on manual intervention or coarse-grained mechanisms that are not consistently tested under time pressure.

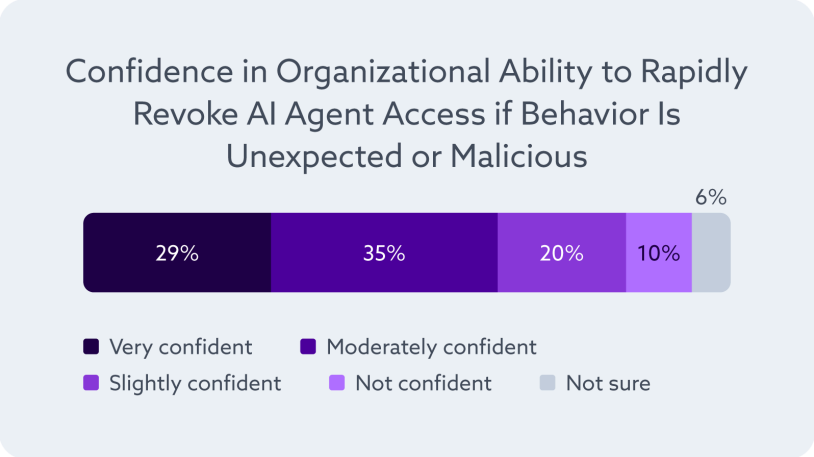

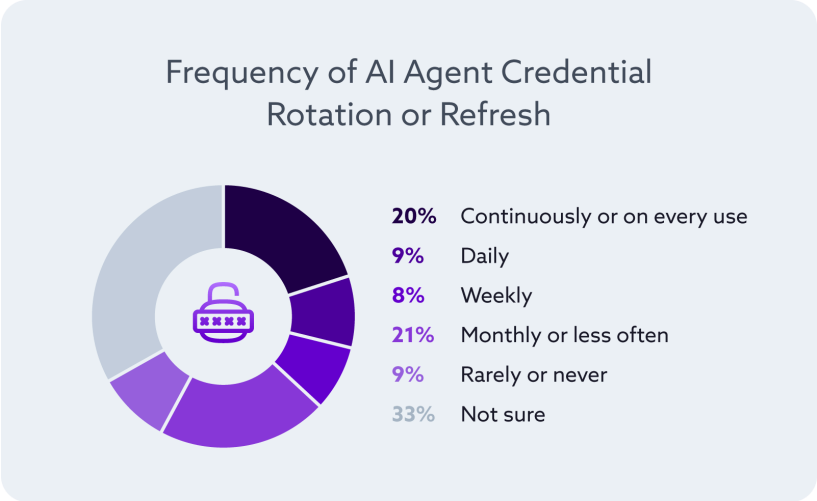

Credential rotation practices reveal one of the clearest control maturity gaps in the data. When asked how frequently credentials used by AI agents are rotated or refreshed, 33% report being unsure. Twenty-one percent indicate rotation occurs monthly or less often, 20% report continuous or per-use rotation, 17% report daily or weekly rotation, and 9% indicate credentials are rarely or never rotated. A full third of respondents lacking visibility into rotation cadence suggests that credential hygiene is not consistently tracked or owned. Where rotation practices are unclear, assurances around access durability and secret management may rely more on assumption than on verifiable process.

Engineering effort associated with AI agent authentication further illustrates operational complexity. Thirty-two percent report being unsure how much time is required to implement and maintain authentication or credential handling for a typical AI agent. Among those providing an estimate, 23% report one to three days, 20% report four to ten days, 17% report less than one day, and 7% report more than ten days.

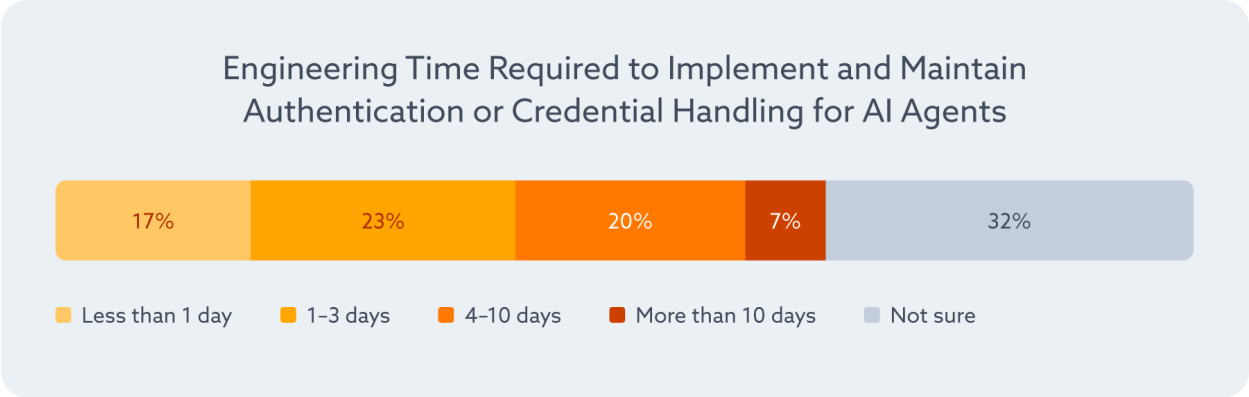

These results suggest that authentication and credential handling are not trivial tasks. When integration requires multiple days of engineering effort, reliance on shared accounts, inherited permissions, or long-lived credentials may become pragmatic shortcuts, even if they introduce governance trade-offs over time.

Consistency in applying access control frameworks to AI agents also varies. Thirty-two percent report applying their access control framework mostly consistently and 22% very consistently. However, 18% indicate only somewhat consistent application, 14% report not at all consistent application, and 14% are unsure. Even where formal access frameworks exist, their application to AI agent identities is not uniform across teams and environments. This uneven enforcement creates a gap between policy intent and operational reality.

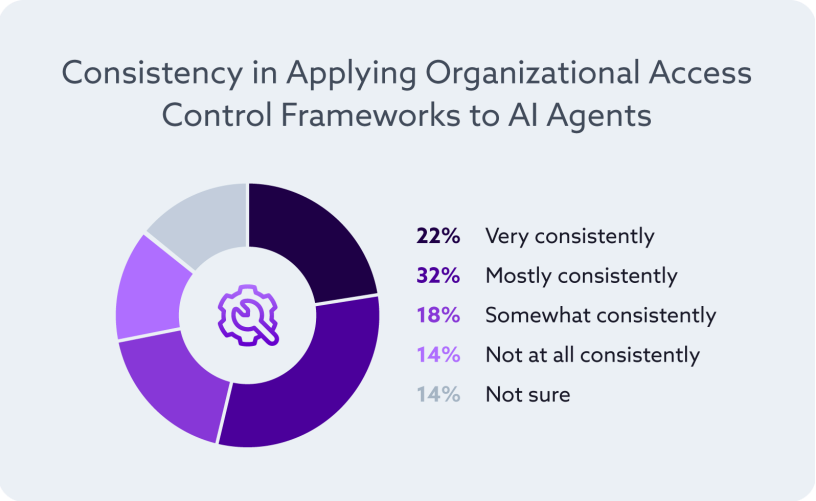

The data suggests that comfort with current workflows may outpace the maturity of the controls underlying them. Organizations often feel reasonably confident in managing AI agent access, yet the underlying operational mechanics reveal uneven implementation and limited transparency. This gap between perceived readiness and demonstrable control maturity reflects the broader challenge of adapting non-human identity practices to rapidly expanding AI agent deployments. These underlying control dynamics shape how access risk ultimately materializes in practice.

Key Finding 5:

AI Agents Inherit Access, Expanding an Organization's Attack Surface

AI agents introduce access patterns that expand privilege, create indirect pathways, and increase exposure to prompt-driven manipulation. Seventy-four percent of organizations agree that AI agents often receive more access than necessary to complete their tasks (53% somewhat agree and 21% strongly agree). Even higher agreement appears around prompt manipulation risk, with **81% agreeing that prompt manipulation could cause an AI agent to reveal sensitive credentials or tokens** (48% somewhat agree and 33% strongly agree). Similarly, 79% agree that AI agents introduce new access pathways that are difficult to monitor (51% somewhat agree and 28% strongly agree).

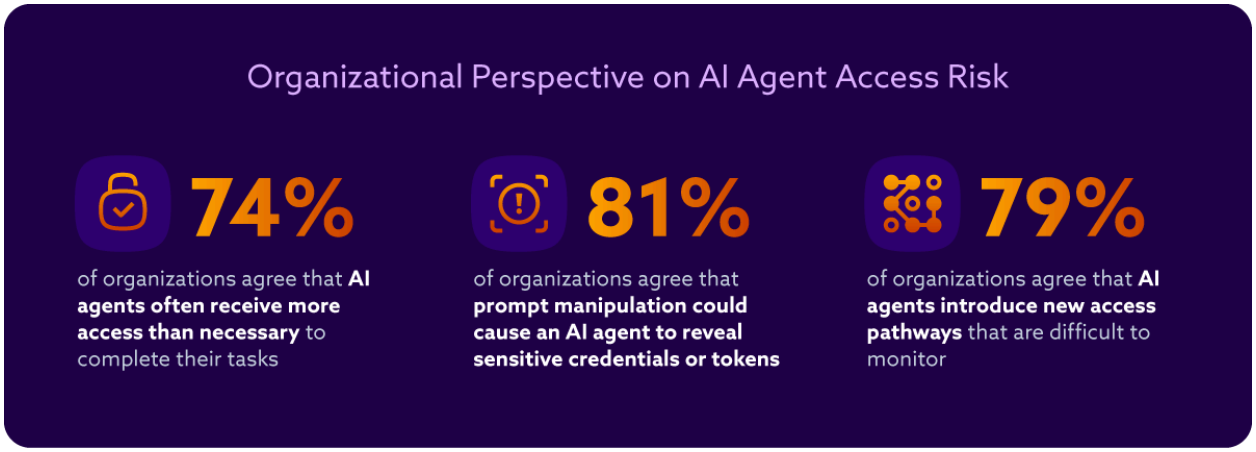

These results indicate broad consensus on the nature of AI agent-related access risk. The concern is not isolated to a narrow segment of respondents; it reflects a shared understanding that AI agents expand the effective access surface within enterprise environments.

These risk perceptions align closely with how AI agent access is determined in practice. When asked how an AI agent's access is typically determined, 25% report predefined rules or automation logic, and 23% report permissions of the human requesting the action. Only 18% report that access is based on the AI agent's own permissions, while 11% indicate reliance on a shared or generic account and 4% report that access is not explicitly scoped. Sixteen percent are unsure. This distribution suggests that access decisions are frequently anchored in human context or pre-existing automation logic rather than in narrowly defined, agent-specific permissions.

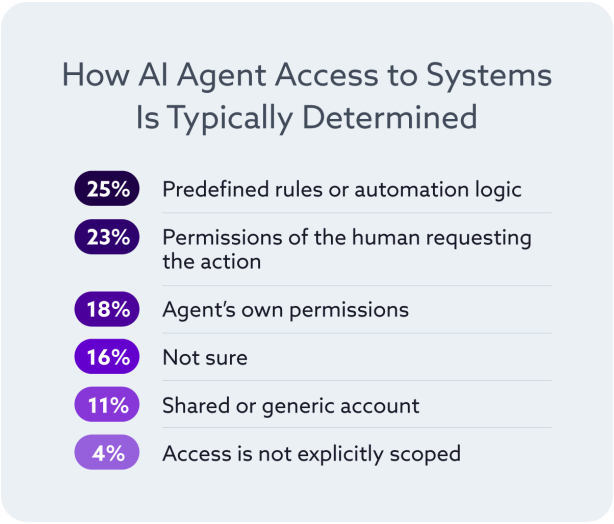

Where access derives from human entitlements or shared identities, the conditions for privilege inheritance are built into the access model itself.

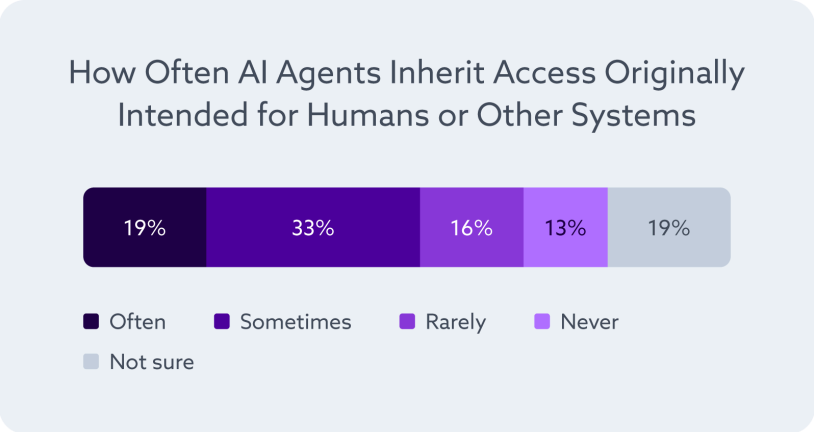

This structural reliance on human and shared contexts provides a plausible explanation for the over-privilege concerns identified earlier. When asked how often AI agents inherit access originally intended for a human or system rather than for the AI agent itself, 33% report sometimes and 19% report often. Twenty-nine percent report rarely or never, and 19% are unsure. More than half, therefore, acknowledge that AI agents inherit access at least some of the time. When access is derived from pre-existing human entitlements or shared accounts, permissions may extend beyond the AI agent's defined role. In this way, excess access becomes embedded in the identity design itself rather than arising solely from isolated misconfigurations.

This pattern creates an inherent asymmetry in how access is managed. When agents operate under humans or shared identities, associated permissions are applied automatically as a function of that identity. Refinding those permissions, however, often depends on downstream governance processes, policy updates, or manual intervention. Where enforcement mechanisms are uneven or inconsistent, inherited access can persist longer than intended, increasing exposure.

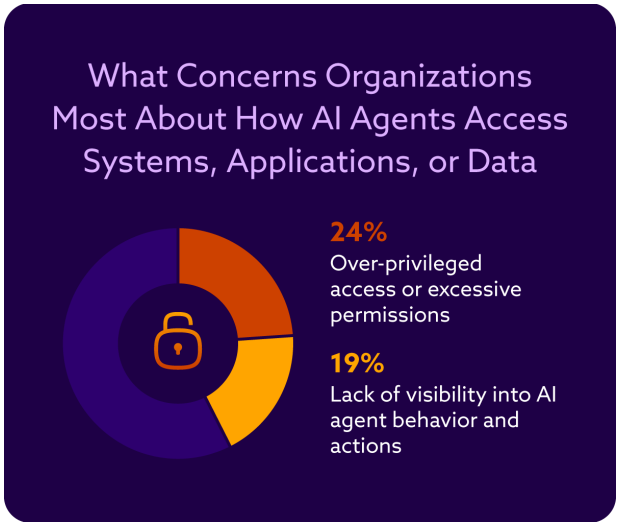

These structural patterns are also reflected in practitioners' open-ended concerns. Among responses describing what concerns organizations most about how AI agents access systems, applications, or data, over-privileged access or excessive permissions emerges as the most frequently cited theme (24%), followed by lack of visibility into AI agent behavior and actions (19%). The convergence between prior quantitative data and these qualitative responses indicates that these risks are being considered within active deployments rather than viewed as abstract possibilities.

As AI agents gain more standardized ways to invoke external tools and services (such as through agent-to-tool interaction models including Model Context Protocol (MCP)), the importance of precise identity scoping and monitoring increases. While the survey does not measure specific protocol adoption, architectural patterns that allow AI agents to interact directly with systems can amplify the consequences

of inherited permissions and limited visibility. In environments where identity governance is uneven or decentralized, these dynamics may further expand the effective access surface.

The overall pattern indicates that traditional access control assumptions do not translate cleanly to agentic behavior. Access is frequently inherited, influenced by human context, or governed through shared identities, creating conditions in which excess privilege and limited visibility can coexist. Addressing these dynamics requires enforcement models that account for how agent access is provisioned, inherited, and constrained in practice.

Key Finding 6:

Teams Are Using Governance as a Stopgap for Missing AI Agent Identity Controls

As AI agents assume greater operational responsibility, many organizations are relying on governance mechanisms—policy restrictions, human approvals, and post-action monitoring—to manage risk where identity-level IAM controls are not yet consistently embedded for AI agents. Nearly half of organizations report disabling the identity used by the AI agent (49%) or revoking active session tokens (49%) to limit access. Forty-two percent report terminating the compute environment in which the AI agent runs, and 37% report workflow automation to block further activity. Only 33% report removing or modifying access policies in real time, while 21% indicate manual review and intervention only and 7% report no clear revocation mechanisms.

The reliance on infrastructure termination as a control mechanism suggests that identity-layer enforcement may not yet be consistently embedded for AI agents. In environments where dynamic, per-identity access controls are limited, organizations appear more likely to intervene at the system or

session level rather than adjusting permissions directly at the identity layer. Such actions can halt activity, but they do not necessarily refine or recalibrate underlying entitlements, making them blunt and reactive.

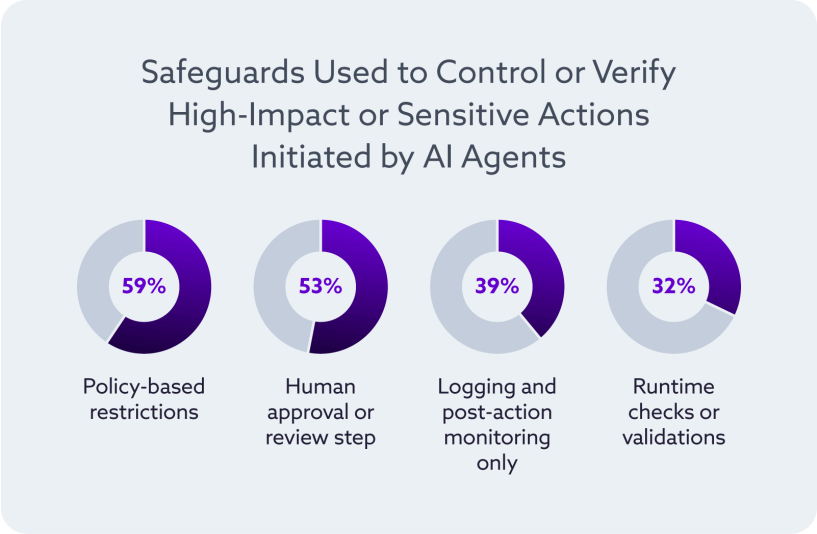

Controls applied to high-impact or sensitive AI agent actions show a similar pattern. Policy-based restrictions are the most commonly reported safeguard (59%), followed by human approval or review steps (53%). Logging and post-action monitoring only are reported by 39%, and runtime checks or validations by 32%. While human-in-the-loop mechanisms can provide an important checkpoint, their effectiveness depends on consistent and deliberate use. Where safeguards rely heavily on approval workflows or retrospective logging, risk management becomes procedural rather than continuously enforced at the identity layer. This reliance on governance and procedural safeguards suggests that dynamic, context-aware enforcement models are not yet uniformly embedded in AI agent deployments.

Accountability structures reinforce this picture. When an AI agent takes an unintended or undesired action, responsibility is most often assigned to security or IT teams (28%) or development or engineering teams (25%). Business or product owners are cited by 18%, 9% report shared responsibility across multiple teams, 6% assign responsibility to the initiating human user, and 15% are unsure. This diffusion of accountability indicates that governance remains distributed and situational. As AI agent autonomy increases, unclear ownership can complicate consistent enforcement and coordinated response.

Looking ahead, respondents signal awareness that governance-heavy approaches may not be sufficient as adoption scales. When asked which capabilities would most improve the ability to safely scale AI agents, the most frequently selected option is real-time visibility into AI agent actions (52%), followed by clear identity separation between AI agents and humans (45%). The ability to grant per-task, short-lived access is selected by 32%, standardized ways for AI agents to authenticate by 30%, and better guardrails to prevent misuse through prompts also by 30%. Immediate kill-switch or revocation control ranks lower at 24%. This prioritization suggests that sustainable control is associated more with continuous visibility and identity clarity than with emergency shutdown mechanisms.

Future identity and access concerns align with these preferences. Preventing over-privileged AI agents is cited by 59% of organizations, distinguishing between human and AI agent activity by 46%, and managing secrets and credentials and monitoring access patterns across multiple environments each by 36%. Applying consistent policies across teams and tools is cited by 30%. These anticipated challenges mirror the structural

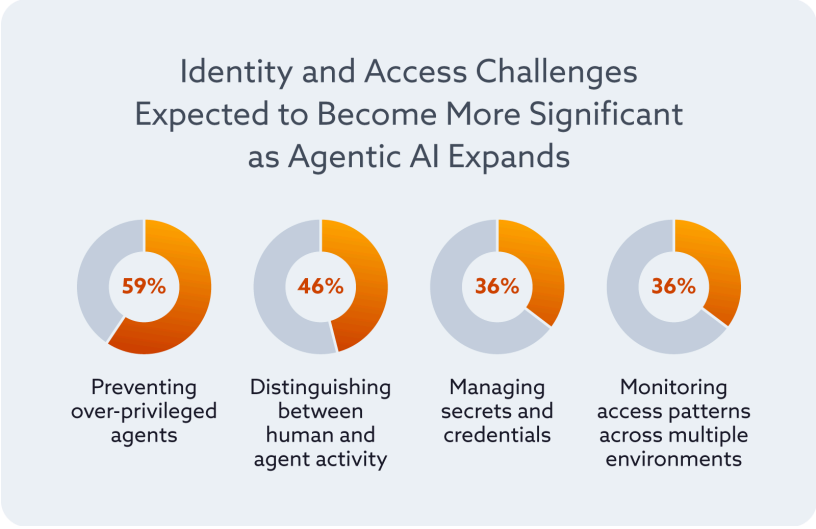

and operational gaps identified in earlier findings, particularly around over-privilege, attribution, and cross-environment enforcement.

As architectural patterns evolve to support more direct agent-to-tool interaction, including MCP, the operational burden of managing AI agent access increases. Expanding AI agent capabilities can place pressure on traditional identity enforcement models that were not originally designed for autonomous, tool-invoking systems. In environments where dynamic, per-identity controls are not consistently applied, organizations appear to rely more heavily on governance mechanisms, human approvals, and procedural safeguards to manage this complexity.

The overall pattern suggests that current control strategies emphasize oversight and containment more than embedded, identity-bound, real-time enforcement. Governance mechanisms are playing a central role in managing AI agent risk. Yet respondents' forward-looking priorities indicate recognition that visibility, clear identity separation, and short-lived access models are necessary to extend IAM principles more systematically to AI agents. The findings collectively point toward a shift from procedural safeguards to more systematic, identity-centric enforcement as AI agent autonomy continues to expand.

Conclusion

AI agents are becoming embedded across enterprise environments, interacting with applications, infrastructure, and data systems in ways that increasingly resemble operational actors rather than experimental tools. As these systems take on greater autonomy and operational responsibility, the identity and access models used to manage them are still evolving.

The findings suggest that many organizations are adapting existing identity frameworks rather than introducing new ones purpose-built for AI agents. As a result, agents are frequently embedded within human or shared identities, inherit permissions not originally designed for their tasks, and operate across environments where attribution and policy enforcement may vary. These structural dynamics influence how access is granted, how behavior is monitored, and how quickly organizations can respond when unexpected actions occur.

At the same time, organizations are developing practical safeguards to manage emerging risks. Policy restrictions, approval workflows, and monitoring mechanisms are commonly used to oversee AI agent activity, while capabilities such as real-time visibility, clear identity separation, and short-lived access are increasingly seen as priorities for the future. These signals suggest that practitioners recognize the need to evolve identity and access practices as AI agent deployments expand.

As architectural patterns continue to evolve and agents gain more direct pathways to systems and tools, the ability to manage identity, privilege, and attribution will become increasingly central to secure adoption. Strengthening identity-centric controls for AI agents—alongside visibility and consistent policy enforcement—will be an important step toward ensuring that agent autonomy scales in a controlled and accountable way.

Full Results

Current Landscape

Which of the following types of AI agents is your organization experimenting with or using today?

Where do the AI agents you use most commonly interact or take action?

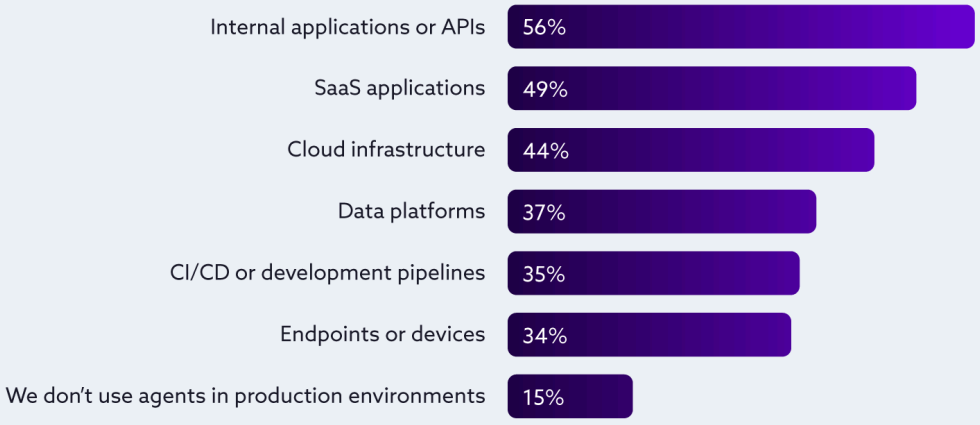

How strategically important do you expect AI agents to become in your organization over the next 12 months?

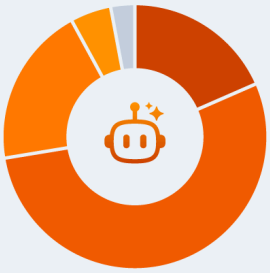

19%	Critical
54%	Very important
20%	Moderately important
5%	Slightly important
3%	Not important

Identity Models, Classification, and Attribution

When an AI agent performs an action in your organization's environment, how is that identity typically represented?

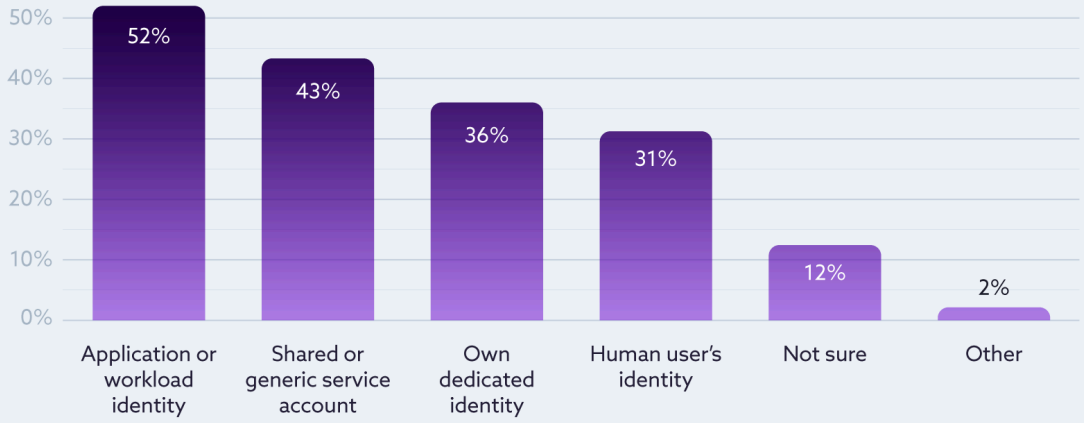

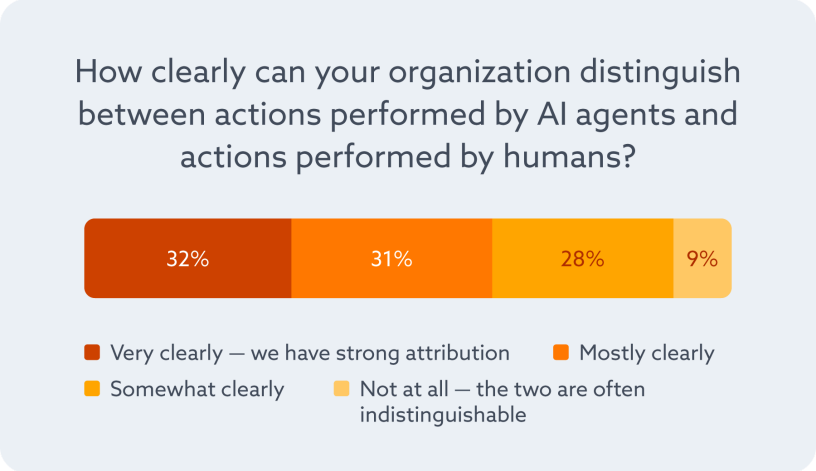

Who is primarily responsible today for determining how AI agents authenticate and access systems?

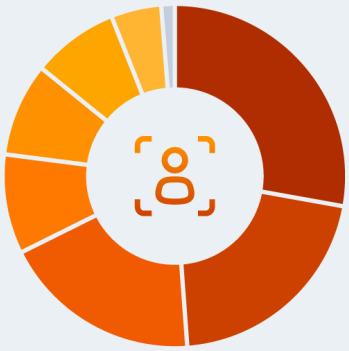

28%	Security	9%	No clear owner
21%	Development or engineering	8%	Not sure
19%	IT	5%	AI/ML platform
9%	Identity	2%	Other

How confident are you that AI agents in your organization have clearly defined identities and appropriately scoped access to systems and data?

Secrets, Access, and Exposure

Where do AI agents most commonly obtain the credentials or tokens needed to access systems today?

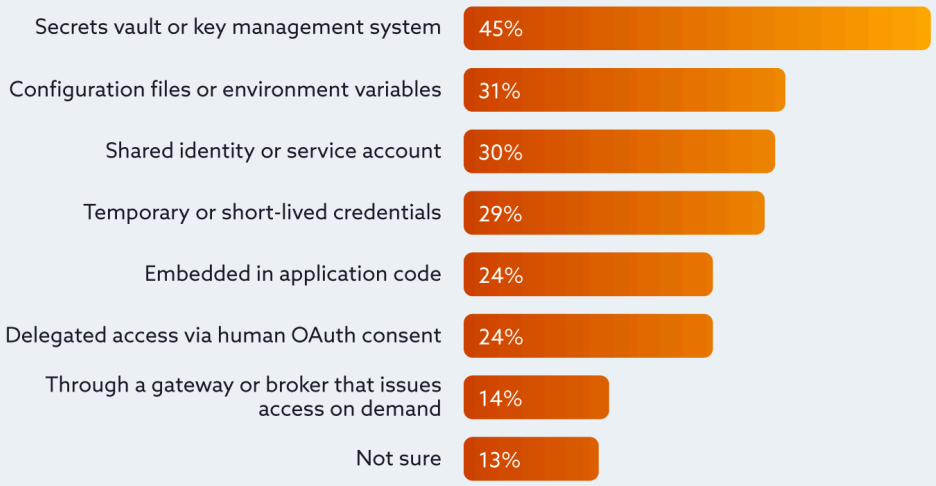

How frequently are the credentials used by AI agents rotated or refreshed?

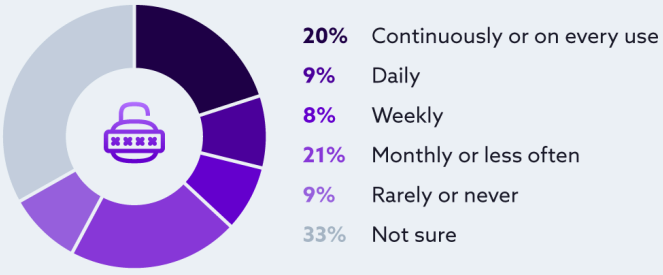

To what extent do you agree with the following statements about access risk?

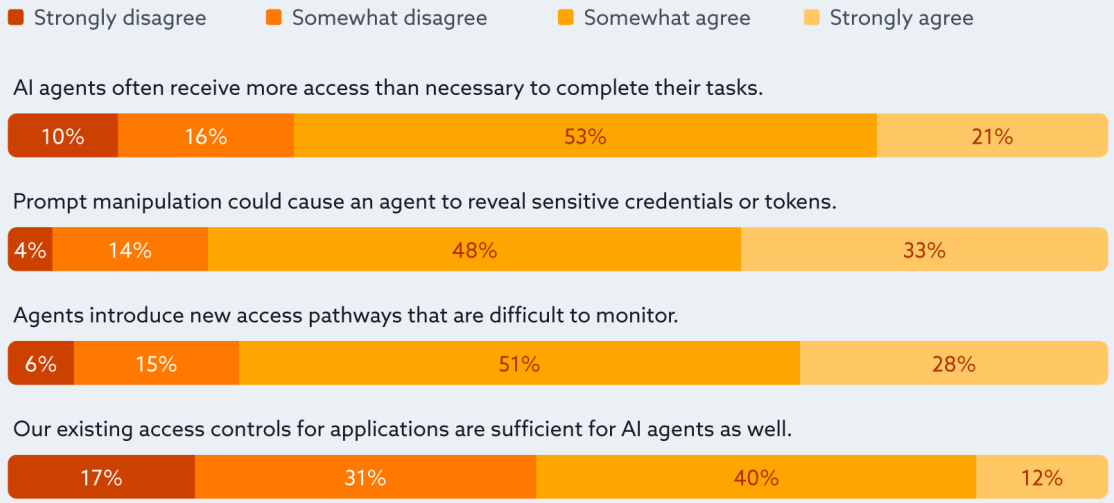

When an AI agent accesses a system, how is its access typically determined?

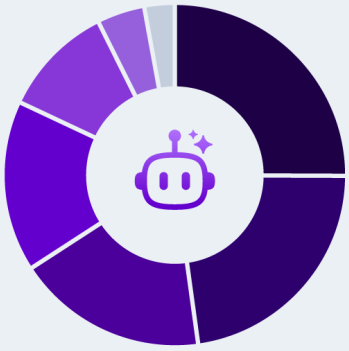

- **25%** Predefined rules or automation logic
- **23%** Permissions of the human requesting the action
- **18%** Agent's own permissions
- **16%** Not sure
- **11%** Shared or generic account
- **4%** Access is not explicitly scoped
- **2%** Other

Approximately how much engineering time is required to implement and maintain authentication or credential handling for a typical AI agent?

Revocation and Real-Time Control

How confident are you that your organization could rapidly revoke an AI agent's access if it behaved unexpectedly or maliciously?

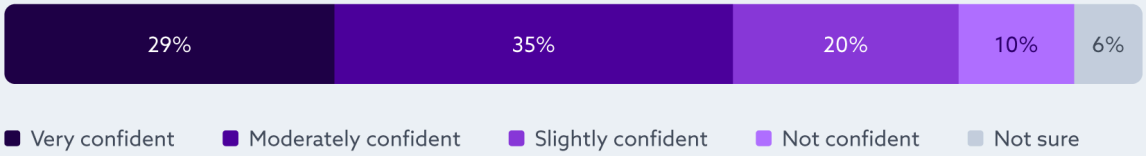

Which of the following mechanisms does your organization actually have for revoking or limiting an agent's access?

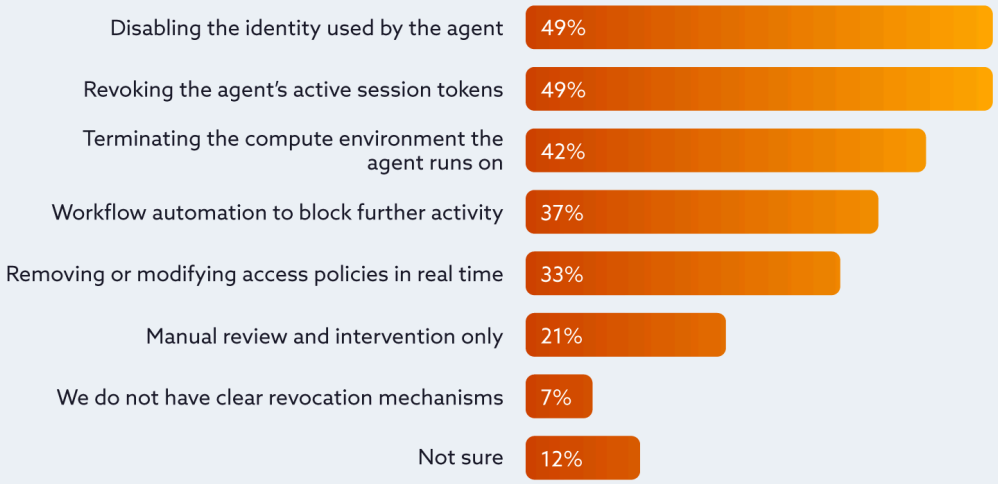

When an AI agent initiates a high-impact or sensitive action, which safeguards exist today to control or verify that action?

How consistently does your organization apply its access control framework to AI agents?

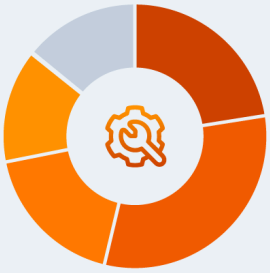

- 22% Very consistently
- 32% Mostly consistently
- 18% Somewhat consistently
- 14% Not at all consistently
- 14% Not sure

Architecture and Future Capabilities

How often do AI agents inherit access that was originally intended for a human or system rather than for the agent itself?

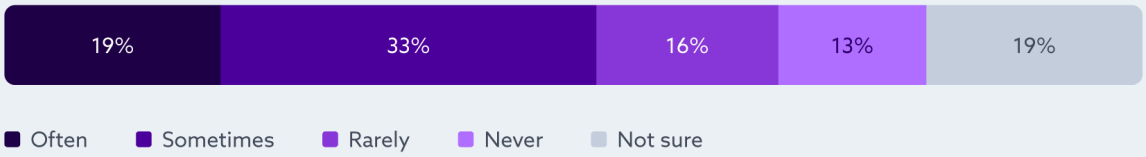

Which teams are most influential in shaping how AI agents access systems and data today?

- **61%** Security
- **44%** IAM/Identity
- **39%** IT operations
- **35%** Application development
- **31%** AI/ML or data platform teams
- **25%** Business units
- **8%** No single group has clear influence

Which of the following capabilities would most improve your organization's ability to safely scale AI agents?

Which best describes your organization's current approach to agent-to-tool authentication?

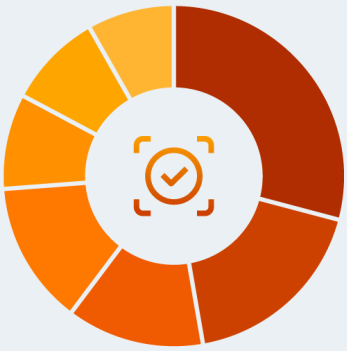

- **29%** Token-based access using established identity standards
- **18%** Short-lived or task-based credentials
- **13%** User-delegated access
- **13%** No consistent approach
- **9%** Not sure
- **9%** Custom or application-specific approaches
- **8%** New or agent-specific access models (e.g., MCP-style tool access)

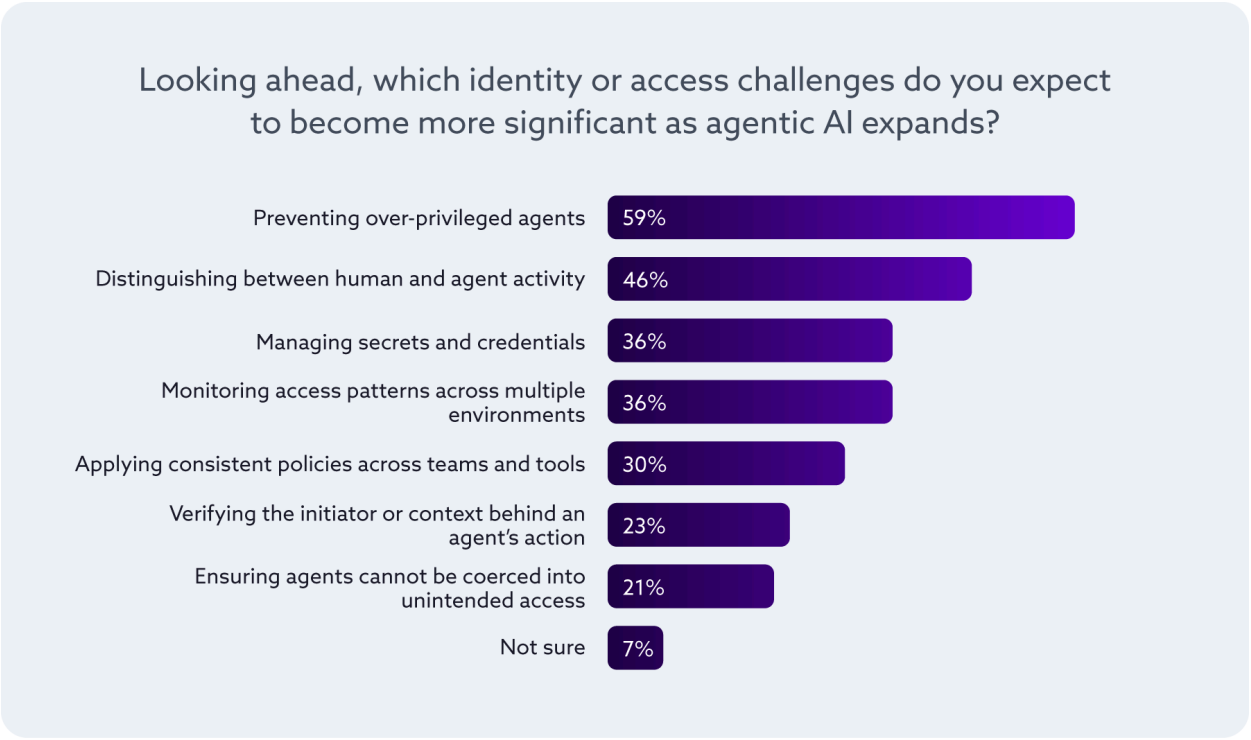

Demographics

Location

- Americas
- Europe, Middle East, Africa (EMEA)
- Asia Pacific (APAC)

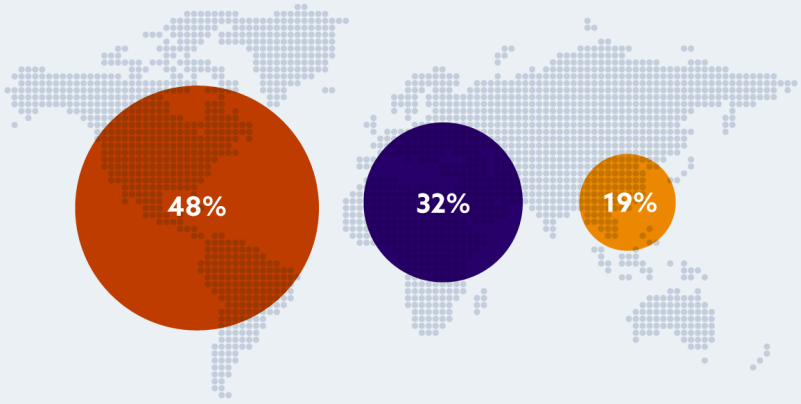

Organization Size

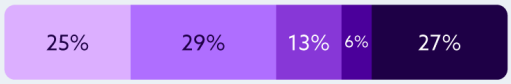

- <100
- 101–1,000
- 1,000–5,000
- 5,001–10,000
- 10,001+

Job Level

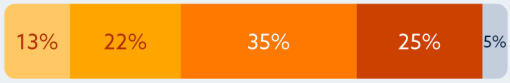

- C-level or Executive
- Director
- Manager
- Staff
- Other

Organization Industry

- **32%** Telecommunications, Technology, Internet & Electronics
- **13%** Finance & Financial Services
- **7%** Government
- **6%** Business Support & Logistics
- **5%** Construction, Machinery, and Homes
- **5%** Education
- **5%** Manufacturing
- **5%** Transportation & Delivery
- **3%** Healthcare & Pharmaceuticals
- **3%** Retail & Consumer Durables
- **3%** Utilities, Energy, and Extraction
- **2%** Advertising & Marketing
- **2%** Airlines & Aerospace
- **2%** Insurance
- **1%** Agriculture
- **1%** Entertainment & Leisure
- **1%** Food & Beverages
- **1%** Real Estate
- **6%** Prefer not to answer

Survey Methodology

The Cloud Security Alliance (CSA) is a not-for-profit organization with a mission to widely promote best practices and ensure cybersecurity in cloud computing and IT technologies. CSA also educates various stakeholders within these industries about security concerns in all other forms of computing. CSA's membership is a broad coalition of industry practitioners, corporations, and professional associations. One of CSA's primary goals is to conduct surveys that assess information security trends. These surveys provide information on organizations' current maturity, opinions, interests, and intentions regarding information security and technology.

Aembit commissioned CSA to develop a survey and report to better understand the industry's knowledge, attitudes, and opinions regarding autonomous AI agents. Aembit financed the project and co-developed the questionnaire with CSA research analysts. The survey was conducted online by CSA in January 2026, and it received 228 responses from IT and security professionals from organizations of various sizes and locations. CSA's research analysts performed the data analysis and interpretation for this report.

Goals of the Study

This research examines how identity, authentication, and access controls are evolving as autonomous systems become more embedded in enterprise environments. The goal is to establish a clear view of current access models, emerging risks, and organizational readiness as these systems move from experimentation into operational use.

Specifically, the survey aims to:

- Measure how organizations define and assign identity for autonomous systems across applications, infrastructure, and data environments
- Evaluate authentication and access approaches, including delegated access, token-based standards, short-lived credentials, and emerging models
- Assess governance and policy enforcement, including how consistently access controls are applied and who is accountable for system actions
- Explore oversight and intervention capabilities, including monitoring, revocation, and real-time control mechanisms
- Map readiness and capability gaps as organizations prepare to scale autonomous systems more broadly