451 Research
by S&P Global

Vanguard Report

August 2026

A new access model for the agentic enterprise

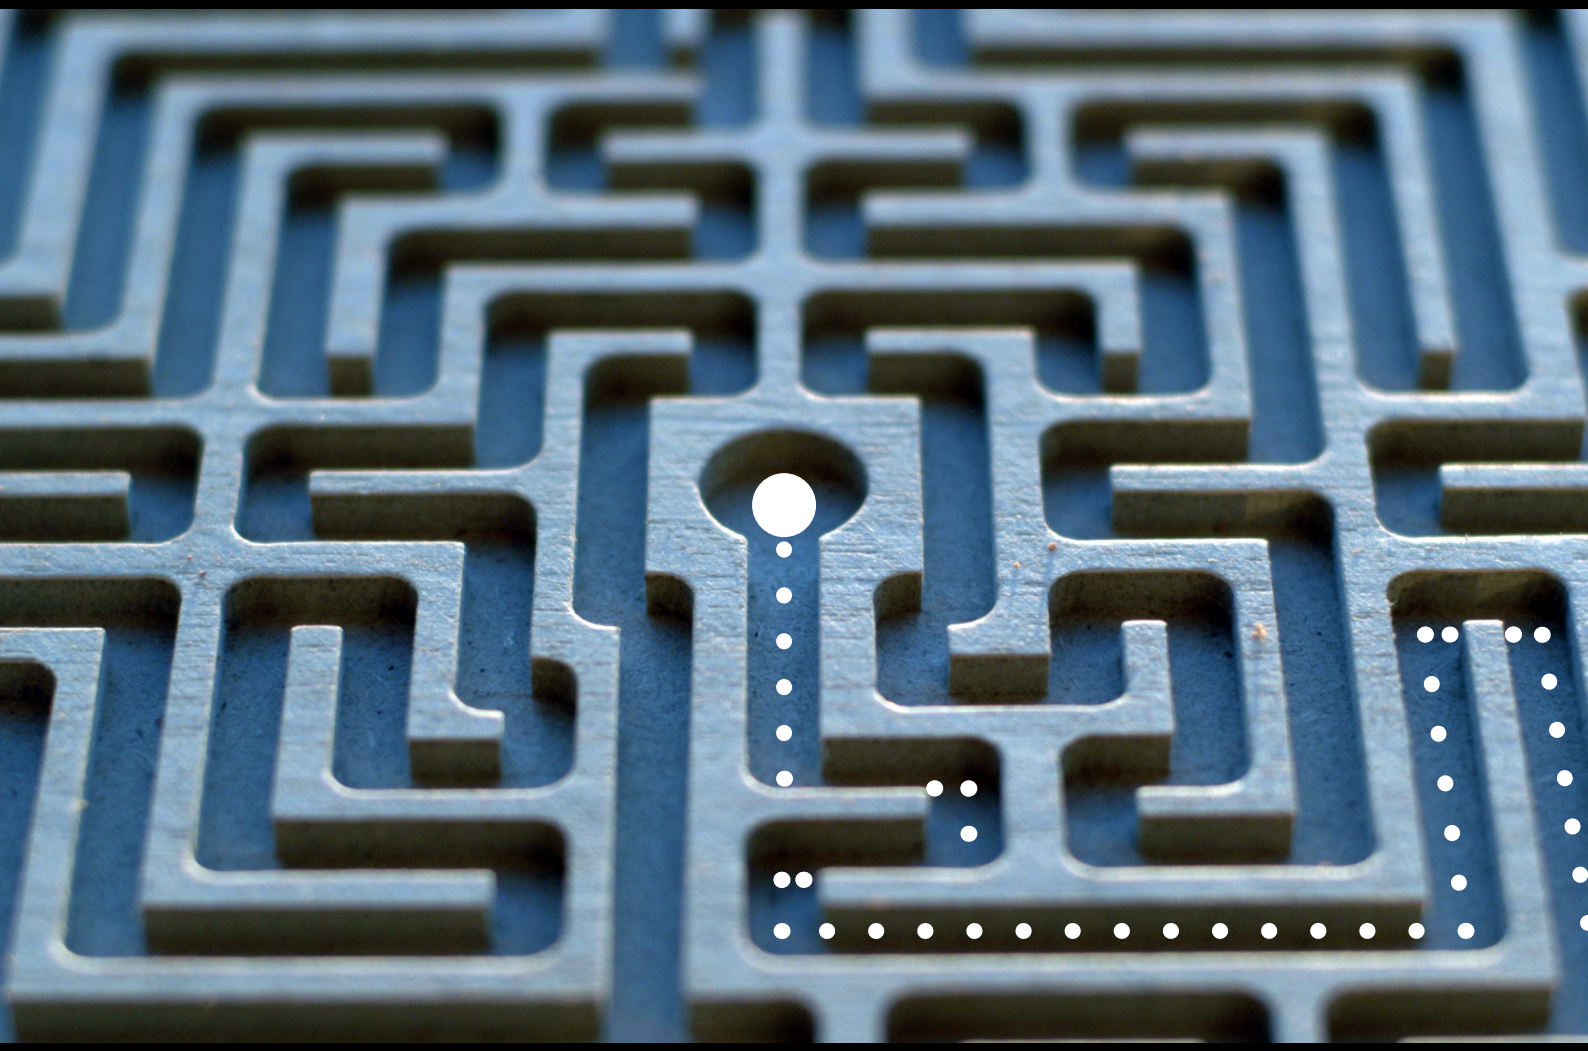

Commissioned by 1Password

Table of contents

Introduction

The identity perimeter is expanding

For decades, identity and access management (IAM) was primarily a human problem. Enterprises built security architectures around people who were onboarded by HR, issued static credentials and governed by policy. That model was not designed for the world that is taking shape.

Traditional IAM and privileged access management (PAM) platforms were designed for human users at human scale. However, non-human identities (NHIs) such as service accounts, API keys, tokens and machine credentials already vastly outnumber human identities in most organizations: Multiple independent studies suggest the ratio ranges from 50:1 to as high as 144:1. Unlike human identities, NHIs are created inside code repositories, CI/CD pipelines and DevOps scripts by developers that are mainly focused on shipping software, not enforcing security policy. Agentic AI is widening that gap: AI agents can log into services, make decisions and act across enterprise environments, all without human involvement.

The Take

The identity perimeter is expanding, but most organizations are not keeping pace. Security strategies are lagging agentic AI adoption — a pattern that has repeated with every major technology advance of the past two decades. What makes agentic AI distinctly challenging to secure is not its scale but its unpredictability. Agents' non-determinism breaks the core assumption of traditional access policy — that administrators can define in advance what a given identity should and should not do.

The agentic AI opportunity is substantial for IAM vendors, but only for those whose platforms are built to govern identities that behave in unanticipated ways. Modern access control needs to address credentials and permissions for humans, machines and agents in a single control plane that unifies governance, policies and auditing. Those organizations that are best positioned to navigate this transition will treat it as an architectural reset rather than an incremental upgrade. Organizations that act now will help to define what secure agentic AI looks like in the future.

The agentic frontier: How AI is redefining the identity challenge

The rise of agentic AI represents a transformational shift in the NHI challenge, not an incremental one. Traditional NHIs such as service accounts and API keys performed narrowly scoped, repetitive tasks on predictable schedules. Their behavior, while often poorly governed, was at least deterministic. Security teams could, in principle, define in advance what a given credential should and should not do.

AI agents break that assumption. Unlike a service account executing a fixed script, an agent can make contextual decisions and act across enterprise environments, often without requiring human-in-the-loop involvement for critical activities and content. That unpredictability is a fundamental property of how LLM-based systems operate. For that reason, traditional static access policies built around known behaviors and unchanging scopes are structurally insufficient for the probabilistic nature of agents.

Figure 1: Agentic AI use within organizations

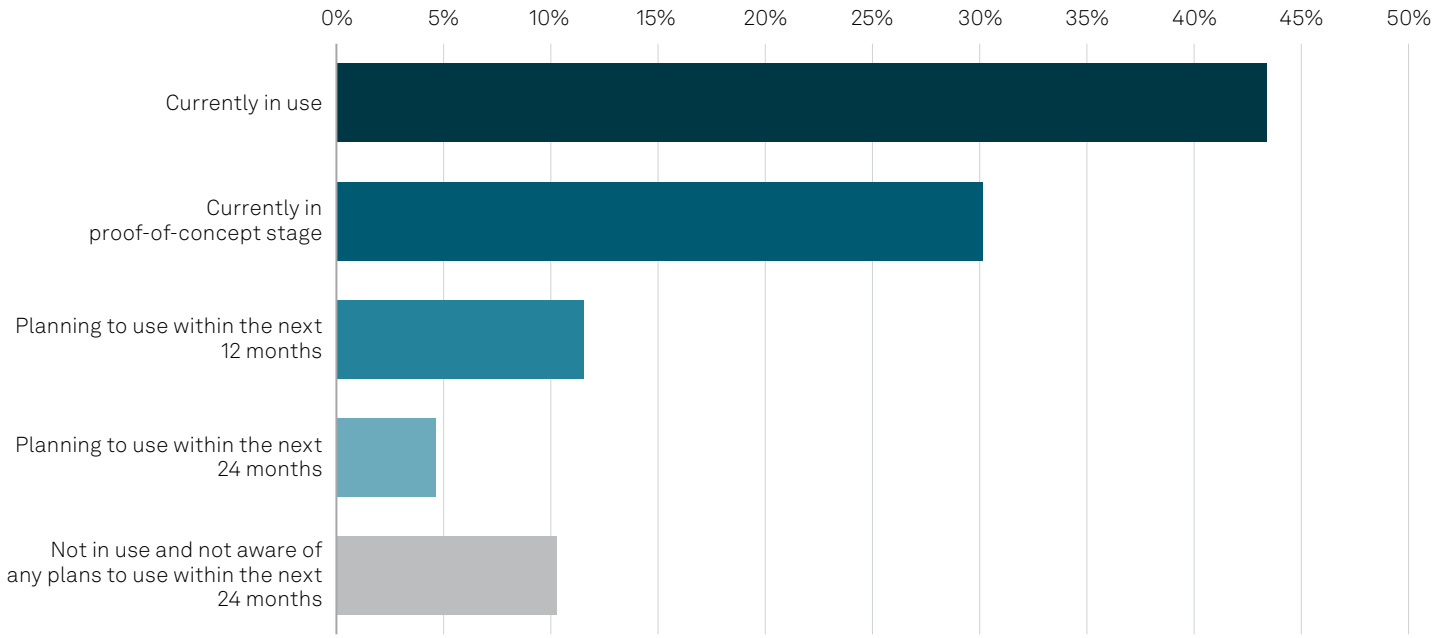

Q. Does your organization currently use Agentic AI within its business and/or IT operations?
Base: All respondents (n=408).
Voice of the Enterprise: DevOps, Agentic AI 2026.
Source: 451 Research by S&P Global.

Agentic AI adoption is outpacing security

451 Research's Voice of the Enterprise (VotE) DevOps survey data shows that 69% of enterprise respondents have already deployed some form of agentic AI, and roughly 90% of enterprises are either using agentic AI for business or IT operations or plan to in the next two years (see Figure 1). Moreover, nearly two-thirds of those organizations expect to adopt agentic AI enterprise-wide or across multiple divisions.

As we have seen with virtually every other technological paradigm shift in recent years, CISOs and CIOs are being pushed to accelerate AI integration at the cost of following basic security processes and principles. Not surprisingly, security concerns related to agentic AI have not kept pace with AI adoption rates. For example, less than 20% of respondents to a recent 451 Research study expressed concern or skepticism about the autonomous nature of agentic AI, while the majority are cautiously optimistic (34%).

Figure 2: Sentiment toward the autonomous behavior of AI agents

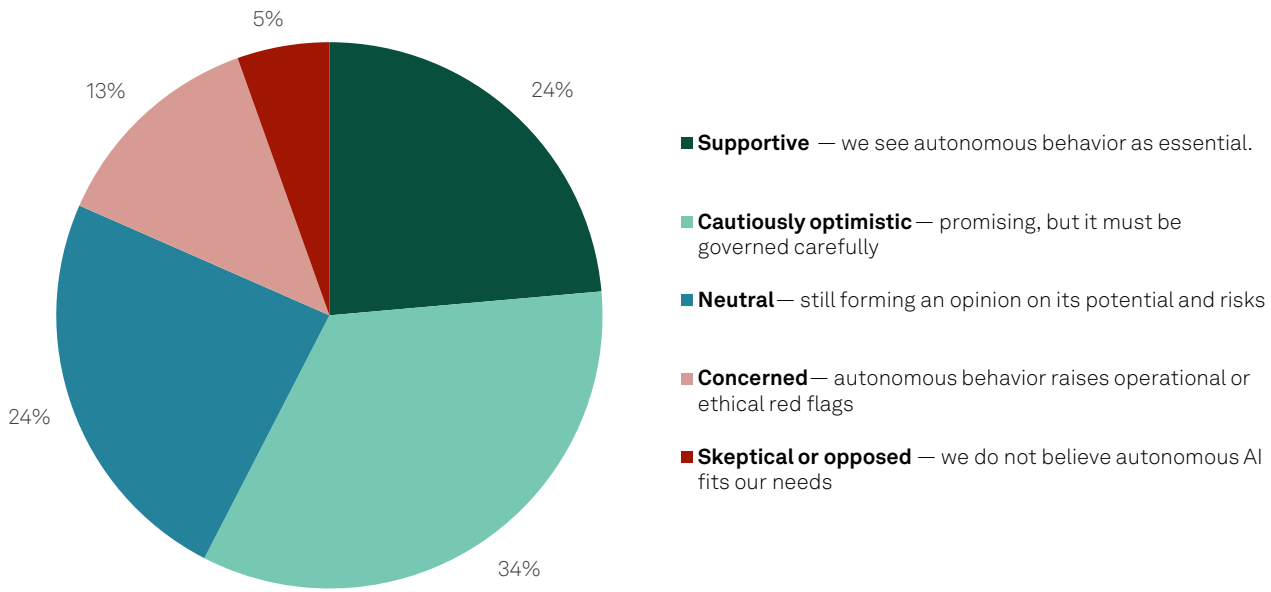

Q. How would you describe your organization's current sentiment toward the autonomous behavior of agentic AI?
Base: All respondents (n=733).
Voice of the Enterprise: AI & Machine Learning, Agentic AI 2025.
Source: 451 Research by S&P Global.

The access trust gap

When web browsers, SaaS apps and mobile devices took hold in the enterprise, a gap emerged between how legacy identity stacks were designed to operate and how access happened: Work shifted to the edge, shadow IT became endemic and security lagged for years. Agentic AI is accelerating on that same arc, but at machine scale and with less margin for error. Agents can be spun up instantly, granted credentials and deployed across an organization in ways that centralized IAM teams have no visibility into, performing actions that humans could not foresee and thus that policies cannot effectively govern. Shadow AI has become the new shadow IT.

The developer dimension

Most identity security risks related to agents originate in developer workflows rather than the security organization itself. Developers are the primary creators of NHIs — provisioning service accounts, embedding secrets and configuring agent access as part of software delivery. Critically, most developers are not IAM experts. They want to build securely, but they lack enabling tools and workflows. Until remediation fits naturally into developer workflows, security debt will accumulate where agents are created.

Key agentic security challenges

The elements contributing to the "access trust gap" in an agentic world can be grouped into five categories:

- **Lack of visibility.** Unlike directory services for humans, there is no unified inventory of existing agents, their credentials, the systems they can reach and how they may behave at runtime. Organizations that cannot answer basic discovery questions cannot begin to govern agent access.
- **Excessive permissions.** Developers provisioning agents may grant overly broad or persistent access to ensure functionality. Unlike excessive permissions for humans, which are often detected during access reviews, over-permissioned agents go largely undetected, and their permissions are rarely revoked.
- **Attribution and accountability.** When something goes wrong, security teams must be able to explain what went wrong, which agent used which credential, under whose authority and when. In environments with large populations of agents operating across shared credentials, those questions are nearly unanswerable.
- **Long-lived credentials and static secrets.** Credentials are routinely embedded in environment files, configuration scripts and code repositories, where they persist long after they are needed. There is no widely adopted model for securely delivering credentials to agents according to just-in-time principles and revoking them once their task is completed.
- **Auditing and remediation gaps.** Even where visibility exists, organizations lack automated remediation workflows to respond to anomalous agent behavior or revoke over-permissioned credentials in a timely way.

A new access model for the agentic enterprise

Traditional IAM and PAM platforms are structurally inadequate for a world where most active identities are non-human, agents behave non-deterministically and credentials are scattered throughout developer toolchains that are often beyond the reach of centralized governance.

The organizations that successfully navigate this transition will treat it as an architectural reset —rethinking identity security from the ground up to govern people, machines and agents in a unified way, with a single control plane that integrates governance, policy management and auditing. Enterprises are already moving in this direction. According to survey data from 451 Research's VotE: Information Security, Identity Management, 41% are already using third-party tools to manage NHIs, 18% are running proofs of concept and another 15% plan to implement NHI security tools within 6 months. The momentum is real, but the gap between tooling adoption and comprehensive governance remains wide.

The fundamentals: What every organization must get right

- **Start with discovery and visibility.** Before organizations can govern agent access, they must know what exists. Discovery means answering basic inventory questions: What agents are operating? What credentials do they hold? What plain-text secrets are embedded in config files or repositories? As the adage goes, you cannot secure what you cannot see, and visibility must be continuous, not a one-time snapshot, to catch shadow AI emerging outside of centralized IT. However, discovery is just a starting point, and it is rapidly becoming table stakes in a comprehensive agentic identity security strategy.
- **Establish a single system of record for credentials.** Today, credentials are scattered across vaults, environment files, repositories and local tooling with no single source of truth. Organizations need a single, authoritative system of record for all credentials and secrets — spanning human users, service accounts, machine identities and AI agents. Every credential, regardless of type, must have a known owner, defined scope, expiration policy and audit trail.
- **Move to just-in-time credential delivery.** Static credentials — stored in vaults, embedded in .env files or hardcoded in repositories — persist far longer than needed, creating broad attack surfaces, and should thus be reserved for non-critical systems. In an agentic world, credentials for critical workloads should be brokered and delivered at runtime using just-in-time methods, scoped precisely to the task and revoked immediately on completion, helping organizations move toward a zero-standing-privilege strategy.
- **Implement guided remediation for developers.** Remediating NHI and agentic risk must fit naturally into developer workflows and must be supported by tooling that makes the secure path the default.
- **Ground identity security in runtime authority.** Authentication is necessary but insufficient. The critical question is whether a given identity should be trusted to take a specific action, in a specific context, at a specific moment. For AI agents, this distinction is especially important: An agent may authenticate legitimately and still act outside its intended scope. Runtime authority means continuously evaluating agent behavior against expected parameters and enforcing access boundaries dynamically, not just at session initiation.
- **Ensure full auditability and clear attribution.** When an incident occurs, security teams must reconstruct exactly what happened: which human or agent used which credential, when and under whose authority. Without clear attribution, incident response is guesswork, and regulatory exposure is substantial. Every agent action should generate a tamper-evident record tied to a specific identity, credential and authorization context.

Conclusion

C-level guidance: Owning the agentic security transition

The fundamental points above describe technical requirements, but this transition is also an organizational challenge, with three imperatives for C-level leaders:

- **Involve cross-functional teams.** Securing agentic AI requires participation from IAM, infrastructure, engineering, AI enablement and — critically — the developer community.
- **Enable developers, don't police them.** Governance policies built without developer input will be ignored or bypassed. Invest in tooling and workflow integrations that make secure agent provisioning the default, not an additional burden on top of delivery pressure.
- **Adopt a framework, not just tools.** The vendor landscape is fragmented: PAM vendors are expanding into identity governance, IGA platforms are extending to machine identities, and a wave of NHI startups has entered the market. Organizations without a coherent framework risk accumulating point solutions that fail to address broad organizational needs. Use a maturity model to sequence investments.

Unified access platform

1Password Unified Access platform governs human, machine, and AI agent identities from a single trusted system. It's built on the foundation of Enterprise Password Manager, the credential vault and system of record for over 180,000 businesses, and secures over 1.5B credentials. The new 1Password Credential Broker extends that same vault and governance model to machine workloads and AI agents at runtime, issuing credentials scoped to the job, with full auditability for every access event. 1Password Privileged Access eliminates standing privileges across cloud and hybrid infrastructure. Access is created at request time, scoped precisely to the task, and deprovisioned automatically when the work is done.

About the author

Garrett Bekker

Principal Research Analyst

Garrett Bekker is a principal research analyst at 451 Research by S&P Global, leading the identity and access management (IAM) vertical within the Information Security channel. Prior to his coverage of IAM, Garrett also covered cloud security data security and governance, risk and compliance. Within IAM, Garrett's current research specializations include passwordless authentication, cloud-native authorization, identity governance and administration (IGA), privileged access management (PAM), identity security and non-human identities (NHIs).

He arrived at S&P Global through its 2019 acquisition of 451 Research, which he joined in 2014. Garrett has viewed security from a variety of perspectives over the past 25 years. He started his career in security as an equity research analyst at several investment banking firms, most recently Merrill Lynch, where he covered information security, infrastructure software and networking companies. Garrett has also worked with early-stage enterprise security vendors, including Bat Blue (acquired by OPAQ Networks), in sales and marketing roles.

Garrett holds a bachelor's degree in international studies from the University of Buffalo. He has completed all coursework for a doctorate in economics from The New School in New York. He also completed undergraduate studies at McGill University in Montreal, and graduate work at Cambridge University in the UK.

About 451 Research

451 Research by S&P Global provides essential insight into key trends driving digital transformation across the entire technology landscape. By offering a combination of expert analyst insight and differentiated data, 451 Research provides client organizations with the critical information and perspectives they require to make more effective business decisions.

451 Research offerings are organized around the prevailing themes driving digital transformation today, including data centers and energy, AI, cybersecurity, digital industries, and resilient digital infrastructure.

For more information about 451 Research, please go to: spglobal.com/451research.

CONTACTS

www.spglobal.com
www.spglobal.com/en/enterprise/about/contact-us.html