SailPoint

Report

The Horizons of Identity Security

From frontier to fabric: Mastering the autonomous enterprise

2026-2027



Executive summary

The enterprise security landscape has crossed a critical threshold. In 2026, 79% of organizations run AI agents in production environments. These agents are making autonomous decisions, querying critical databases, and executing transactions on behalf of the business. Yet only 2% have deployed identity security tools purpose built to govern and secure them. This gap, between the breakneck speed of agentic deployment and the lagging maturity of the identity controls meant to govern it, is the central focus of the fifth annual Horizons of Identity Security report.

To quantify this shift, SailPoint surveyed 340 senior identity, IT, cybersecurity, and risk leaders across North America, Europe, Asia Pacific, and Latin America. Evaluating organizations across strategy, operating models, talent, and technical capabilities, the 2026 Horizons of Identity Security maturity model assesses where enterprises stand on a five-tier evolutionary scale:

- **Horizon 1 - No formal program:** Identity is fragmented, tactical, and siloed. Human access relies on manual provisioning, while machine and autonomous agent identities operate largely outside formal governance.
- **Horizon 2 - Manual, tool assisted:** Centralized IAM tooling is deployed for core business systems, but operations remain heavily ticket driven. A standard access process starts to cover some non-human identity types, though agentic maturity remains nascent.
- **Horizon 3 - Broadly deployed:** Identity controls are deployed broadly across the enterprise. A single, standard, repeatable process becomes the norm for human and non-human access. Basic non-human governance is established, with centralized inventory, explicit ownership, and distinct lifecycle management across environments.
- **Horizon 4 - AI strengthened:** Identity functions as a strategic business enabler. Dynamic, context aware, and just-in-time access controls span virtually all human and non-human environments, integrated directly with SecOps for automated threat response. More sophisticated controls are deployed, like prompt-injection monitoring and behavioral risk analysis.
- **Horizon 5 - Ecosystem integrated:** Identity governance is a primary pillar of enterprise innovation. It integrates seamlessly across the digital ecosystem, functioning as a self-governing identity fabric. In this advanced state, autonomous “guardian” agents monitor peer agents and processes self-optimize, while a federated agent trust model governs secure agent-to-agent collaboration across organizational boundaries.

Key findings

Five patterns stand out from the 2026 data.

1. **Agentic adoption has severely outpaced security controls.** Non-human identities now vastly outnumber the human workforce: AI agents represent 22% of all non-human accounts today. Yet most organizations continue to stretch traditional, human-speed tools across autonomous entities, with 85% relying on legacy general purpose identity tools or platform native controls (such as on-premises directories, manually maintained roles, or tools not designed for cloud, machine, and AI agent identities).
2. **Overall maturity remains anchored at foundational levels.** Roughly 60% of organizations remain in Horizons 1 and 2, while less than 1% have achieved Horizon 5. This lag is overwhelmingly driven by non-human identity deficits: 87% rate their human IAM capabilities as capable or better, while only 43% report mature processes for agentic access.
3. **A pronounced awareness gap masks operational vulnerability.** Broad leadership optimism sharply diverges from technical reality. 80% of respondents perceive only a minor gap in their tooling, yet only 15% can provision machine access in real time, and only 15% possess mature capabilities for detecting shadow AI. Similarly, 57% express confidence in meeting regulatory requirements, while only 43% feel prepared to produce verifiable evidence for an AI-related audit inquiry.
4. **Foundations remain the primary operational roadblock.** Before enterprises can realize dynamic, intent-based security, they must first master baseline hygiene and visibility. Credential lifecycle management, shadow AI discovery, and real-time monitoring are rated as severe operational hurdles across the market.
5. **The business value of non-human identity security is immediate and measurable.** Securing non-human identities delivers direct business velocity, not just defensive value. 62% of organizations report measurable workforce and developer productivity gains, and 46% report safer, faster AI deployment, once manual access bottlenecks are eliminated.

What this means for security leaders

Identity architectures built for humans cannot govern autonomous agents executing thousands of transactions per second. The data points to a coverage gap, not a competence gap: the same organizations that manage human identity well are the ones struggling to extend equivalent controls to machines and agents.

For CISOs and IT executives, the mandate is straightforward. Bring all identities, human, machine, and AI agents, into a unified identity fabric, replacing manual gates with continuous, context-aware automation built for the speed of machines. The chapters that follow examine where this gap originates, what it costs to leave unaddressed, and lay out a concrete, horizon-by-horizon roadmap toward autonomous security.

Chapter 1

The future of identity: Elements and latest trends



1 Access

Zero standing privilege at machine speed

- Zero standing privilege - no baseline access until task start
- Dynamic policy shifts with behavior, device, context
- Machine-speed authorization - access decided in seconds
- Agent-to-agent auth runs on Auth 2.1 / OIDC to MCP servers

Future signals

- **Workload access mgmt -** runtime, least-privilege for services & agents
- **Intent-based access -** access tied to declared task, not standing role

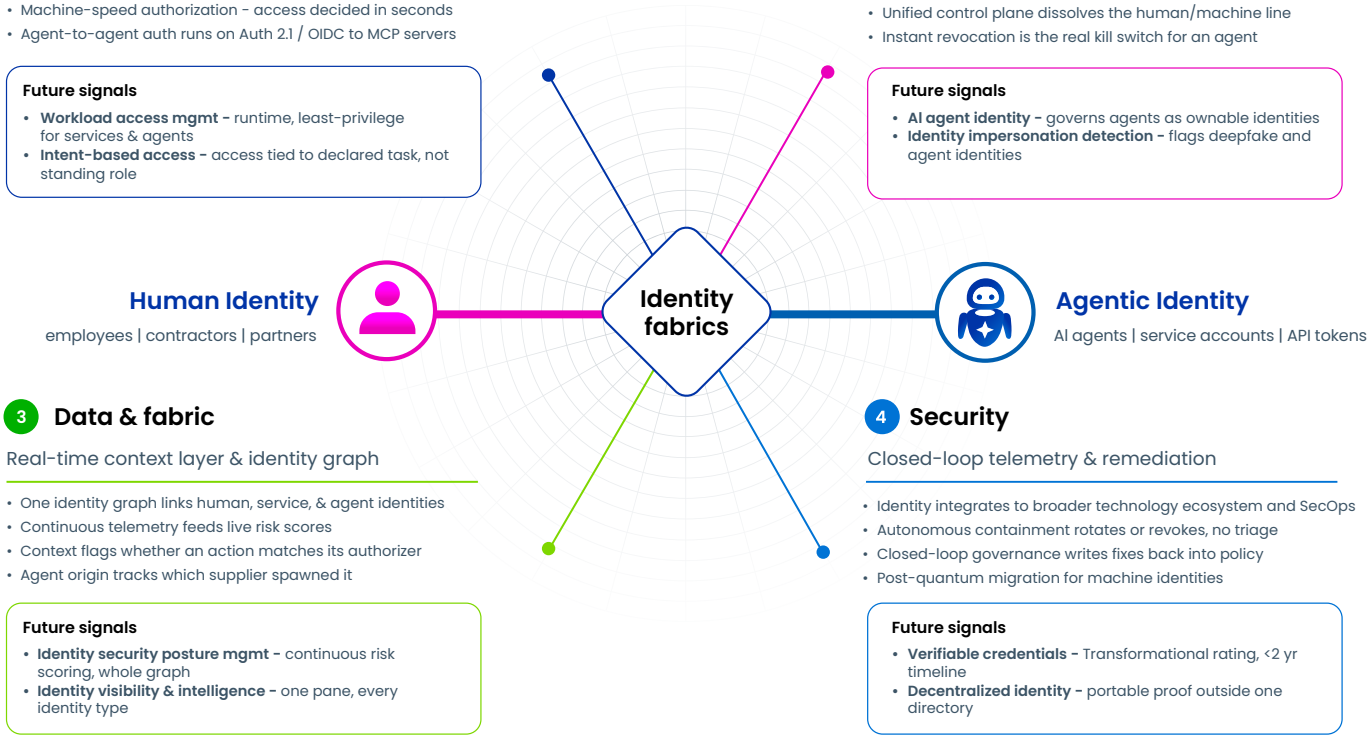
2 Governance

Unified oversight, human & machine

- Automated discovery - every agent cataloged with owner & scope
- Ephemeral governance - credentials live minutes, not months
- Unified control plane dissolves the human/machine line
- Instant revocation is the real kill switch for an agent

Future signals

- **AI agent identity -** governs agents as ownable identities
- **Identity impersonation detection -** flags deepfake and agent identities



3 Data & fabric

Real-time context layer & identity graph

- One identity graph links human, service, & agent identities
- Continuous telemetry feeds live risk scores
- Context flags whether an action matches its authorizer
- Agent origin tracks which supplier spawned it

Future signals

- **Identity security posture mgmt -** continuous risk scoring, whole graph
- **Identity visibility & intelligence -** one pane, every identity type

4 Security

Closed-loop telemetry & remediation

- Identity integrates to broader technology ecosystem and SecOps
- Autonomous containment rotates or revokes, no triage
- Closed-loop governance writes fixes back into policy
- Post-quantum migration for machine identities

Future signals

- **Verifiable credentials -** Transformational rating, <2 yr timeline
- **Decentralized identity -** portable proof outside one directory

Autonomous identity converges access, governance, data, and security into one self-directing fabric

For years, identity and access management functioned as a defensive checkpoint: A perimeter gate verifying human credentials before granting static privileges. In the autonomous enterprise, identity has evolved into the central nervous system connecting applications, data, cloud infrastructure, and autonomous agents. The future of identity security is defined by the convergence of four foundational pillars: Access, Governance, Data, and Security.

Access

Dynamic, zero standing privilege at machine speed

Governance

Unified oversight across human and machine identities

Data & fabric

Real-time context layer and identity graph mapping

Security

Closed-loop telemetry and remediation

Pillar 1: Privilege access is moving from standing grants to continuous, contextual trust

- **Zero standing privilege (ZSP).** Accounts hold zero baseline permissions until an authorized task is initiated.
- **Dynamic policy enforcement.** Permissions adjust in real time based on behavioral signals, device trust, session context, and data sensitivity.
- **Machine-speed authorization.** AI-driven policy engines evaluate access requests in seconds, eliminating "admin-speed" approval delays.

Pillar 2: Governance is expanding to cover every non-human and agent identity

- **Automated discovery and registration.** Every machine identity, API token, and AI agent is cataloged at inception with an explicit owner, purpose, and scope.
- **Ephemeral identity governance.** Credentials living for mere minutes or seconds, including across chains of sub-agents, are dynamically assigned, audited, and decommissioned without human intervention.
- **Unified control plane.** The artificial boundary separating human employees, service accounts, and AI agents dissolves into a single, cohesive governance fabric.

Pillar 3: Identity data is consolidating into an integrated context layer

- **Unified identity graph.** Human, service, and AI agent identities are linked via entity resolution into a single authoritative identity graph, no more separate context for who a person is and what an agent is doing on their behalf.
- **Continuous behavioral telemetry.** Access patterns feed dynamic risk scores continuously, rather than during quarterly audit cycles.
- **Context-rich intelligence.** Human and agent identity signals provide granular business context, answering not just who holds access, but what normal behavior looks like, and whether an agent's actions align with the human or workflow that authorized it.

"Identity data can no longer live in siloed repositories. Bringing signals together into a single context layer is what transforms static directory records into continuous, risk-aware intelligence."

Enterprise CISO & Advisory Panelist

Pillar 4: Identity telemetry is powering automated threat detection and response

- **Identity integrates with the broader technology ecosystem:** Real-time identity signals feed SecOps and the broader security ecosystem to instantly detect credential misuse and lateral movement.
- **Autonomous containment:** Suspicious behavior triggers automated credential rotation, step-up verification, or session termination, without waiting for human triage.
- **Closed-loop governance:** Incident remediation feeds directly back into access policies, preventing identical vulnerabilities from recurring across the environment.
- **Post-quantum migration for machine identities:** Protect the credentials protecting today's non-human accounts and AI agents from getting harvested by tomorrow's post quantum decryption.

"Integrating telemetry into detection is foundational today, but automated remediation remains further out—because a bad automated response can create an operational incident of its own."

Enterprise CISO & Advisory Panelist

Chapter 2

The state of identity security today: The Horizons of the invisible fabric



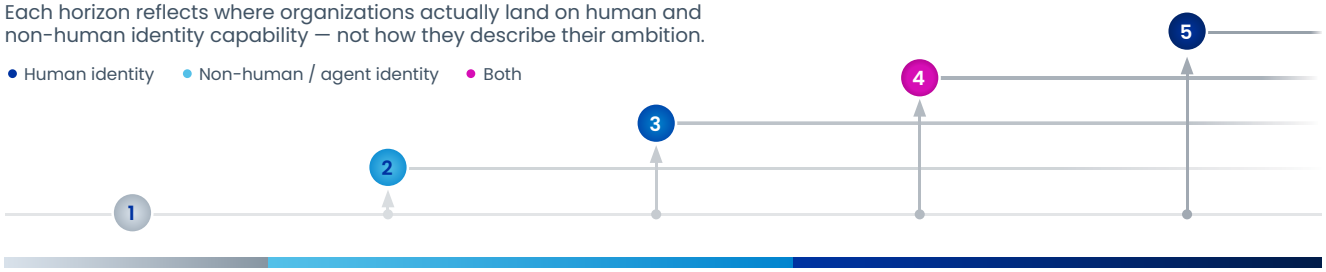
The maturity model: Evaluating capabilities in the autonomous era

The 2026 Horizons Maturity Index retains the standing identity-security framework (strategy, talent, operating model, and technology) while deepening non-human and agentic readiness across its scoring and separating technical confidence from technical capability. Last year's report established that non-human identities were outpacing humans in sheer volume. This year, the data reveals that the challenge has evolved from a crisis of scale to a crisis of agency. We are no longer just managing passive workloads; we are dealing with autonomous agents acting without oversight.

Exhibit 2.1:

The five Horizons of Identity Security, separated by real capability gaps in the 2026 data

Each horizon reflects where organizations actually land on human and non-human identity capability — not how they describe their ambition.



No formal program

Strategy and technology for a digital identity program not yet in place

- **Human capability lags** — human controls remain manual, fragmented, and ad hoc.
- **Additional identity coverage is limited** — machine, cloud, and SaaS identities are only partially covered by their IAM program.
- **Non-human strategy is still emerging** — they recognize the need for agent identity strategy but have no defined approach yet.
- **Standard access processes are rare** — no single standard process for non-human identity access.
- **Shadow AI ungoverned** — no visibility into unsanctioned AI tools or deployments
- **Team expertise is informal** — identity teams have only informal, untrained familiarity with AI agent security.

Manual, tool-assisted

Some identity tooling live, but still heavily dependent on manual work

- **IAM automation stays partial** — most human identity tasks are still manual alongside some that are automated.
- **Access processes cover some identity types** — have a standard process for only some non-human identity types, not all.
- **Strategy graduates to "defined"** — a documented non-human identity strategy starts to emerge.
- **Manual processes/limited tech to detect and govern Shadow AI** — limited emerging visibility to Shadow AI.
- **Team expertise is developing** — some staff trained specifically on AI agent security.
- **Agents start entering audit scope** — starting to include AI agents in identity audits

Broadly deployed

Identity capabilities deployed broadly across the business

- **Human identity capability turns strong** — core capabilities are established and have strong integration with SecOps.
- **A single standard process becomes the norm** — repeatable processes govern most non-human access.
- **Non-human strategy aligns with enterprise AI** — "aligned" becomes the leading goal for the first time.
- **NHI inventory established** — visibility across endpoint, SaaS, cloud; shadow AI controls emerging.
- **Ownership & lifecycle assigned** — explicit ownership for machine identities; distinct lifecycle mgmt begins.
- **Team expertise becomes capable** — the team is now broadly trained on AI agent controls and security.
- **Audit readiness turns a corner** — organizations are more prepared for AI-related audit inquiries.
- **Integration with SOC** — both human and non-human signals.

AI-strengthened

Runs automated identity programs and uses AI to strengthen them

- **Human identity capability is strong across the board** — controls are strong for human identities with automated, SecOps-integrated threat response.
- **A standard non-human access process is nearly universal** — consistent governance covers virtually all non-human identities.
- **Agentic processes reach maturity** — governance processes are mature or continuously self-optimizing.
- **Team expertise turns advanced** — possesses deep, specialized expertise that is continuously developed.
- **Every audit includes agents** — 100% of identity audits cover AI agent identities.
- **Prompt injection, intent controls and risk-based behavior monitoring** in place for non-human identities.
- **Real-time and just-in-time authorization is norm** — instantaneous access remains the norm at this stage.
- **Centralized policy engine** granting agent access.

Ecosystem-integrated

Dissolves the line between internal identity controls and the broader ecosystem

- **Human capability is fully optimized** — core controls are automated, intelligent, and seamless.
- **Agentic processes self-optimize** — continuously adapting to new use cases without manual rework.
- **Non-human strategy is fully aligned** — explicitly aligning with enterprise AI strategy.
- **Identity talent spans the ecosystem** — distributed beyond the enterprise, collaborating with partner companies and industry forums.
- **Audit readiness is complete** — organizations are extremely prepared for human and AI-related audit inquiries.
- **Self-governing identity fabric** — autonomous guardian agents monitor peers; processes self-optimize.
- **Federated agent trust** — agent-to-agent trust chains governed across organizational boundaries.

INSUFFICIENT	FOUNDATIONAL	ADVANCED

Horizon = H26_LEVELSr6 (overall 2026 horizon level). Human-identity items: H3, H4, H5, H6. Non-human/agent items: Q12, Q13, Q14, Q15, Q16A, Q16B, Q16C. Both: Q15B, Q20B. Horizon 5, n=2 — directional only, not statistically reliable.

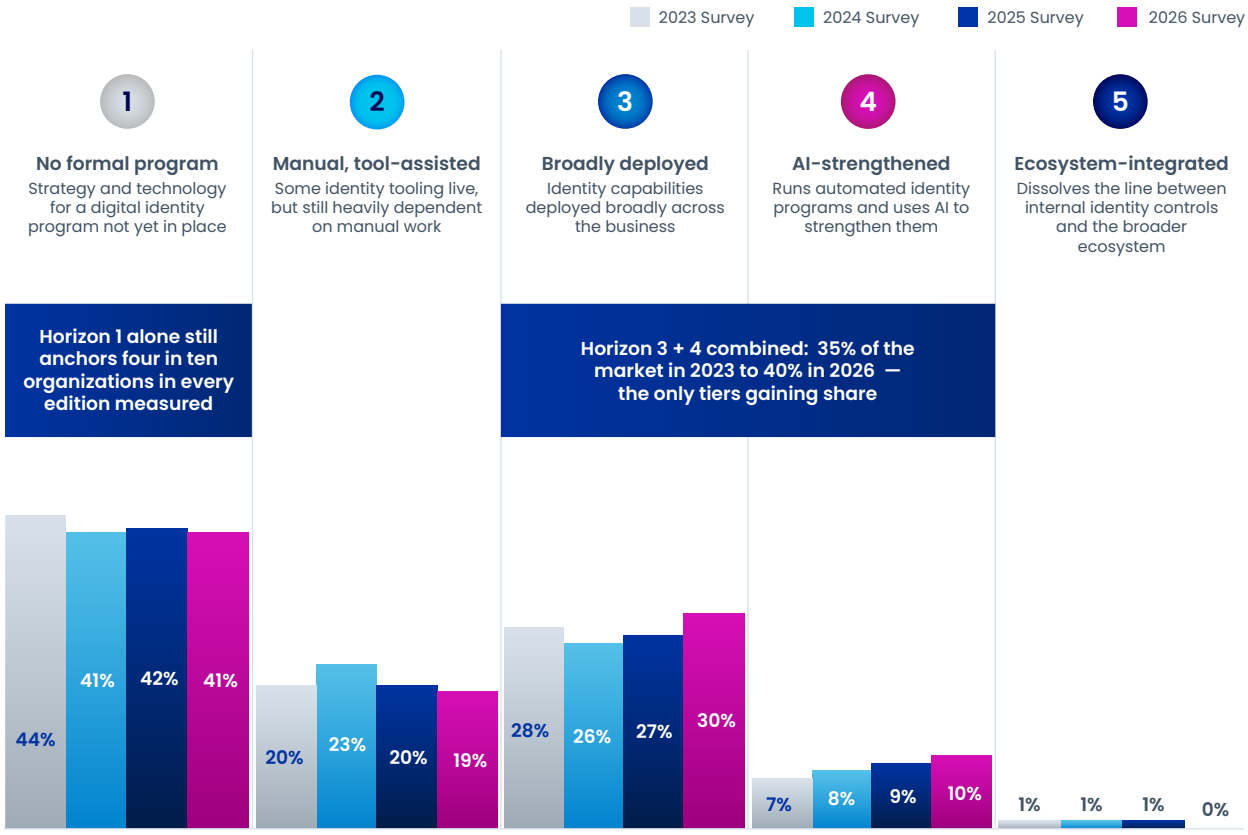
Source: ETR/SailPoint Horizons of Identity Security survey, N=340.

The five horizons in practice

Each Horizon describes a recognizable organizational state, not just a score band: Horizon 1 is loud and manual, while Horizon 5 is silent, continuous, and autonomous. Moving across these stages marks a fundamental transition from siloed, ticket-based administration to digitized policy governance, culminating in dynamic, AI-driven trust.

Exhibit 2.2:

Distribution of identity maturity over time (2023 to 2026)



Source: SailPoint Horizons of Identity Security survey, N=340. Overall horizon-maturity distribution, trended against the prior three editions of the survey; each edition is an independent cross-sectional sample, not a tracked panel. 2026, N=340. Horizon definitions condensed from the report's own five-level maturity scale.

A market anchored in the foundational stages

Despite years of investment and a clear understanding of the risks, the market's center of gravity remains firmly planted in the foundational stages of identity maturity. Our 2026 data shows that a combined 60% of organizations are still in Horizon 1 (41%) or Horizon 2 (19%). The multi-year persistence of the maturity curve reveals a critical structural reality: The market has hit an architectural ceiling, not an adoption problem. Most organizations have spent the last three years refining identity governance for human employees, only to find that the operational playbooks built for the workforce cannot govern autonomous agents executing thousands of transactions per minute.

Three structural dynamics explain why the market's center of gravity remains anchored in the foundational tiers:

- **The tool-sprawl illusion in early horizons:** Organizations in the foundational stages often have identity technology, but it is typically deployed as disconnected point solutions rather than a unified control plane. Point solutions are deployed tactically across SaaS, infrastructure, and DevOps silos, creating disconnected pockets of administration where non-human accounts and non-human secrets proliferate without centralized ownership or lifecycle tracking.
- **The "digitization trap" at the middle tier:** The middle horizon represents the upper limit of what manual and scheduled processes can achieve. Organizations at this stage have successfully digitized employee joiner-mover-leaver (JML) workflows and periodic access reviews. However, attempting to apply these same scheduled review cycles to ephemeral machines and AI agents creates severe operational drag, preventing teams from advancing toward dynamic security.
- **The fundamental pivot to machine-speed trust:** Breaking into the upper horizons requires an architectural paradigm shift rather than incremental process improvements. The few organizations operating at the advanced tier have decoupled access decisions from manual ticketing, replacing standing privileges with continuous, contextual evaluation and automated policy enforcement.

> **"Human access management always received the people, process and technology investment because that is where we believed the risk was… Agents now shine a bright light on that problem, because we are not tooled for non-human identity security."**
>
> CISO
> Large enterprise technology company

Until enterprises bridge this gap by bringing machine identities into the same unified governance fabric as their human workforce, overall maturity will remain tethered to the foundational levels.

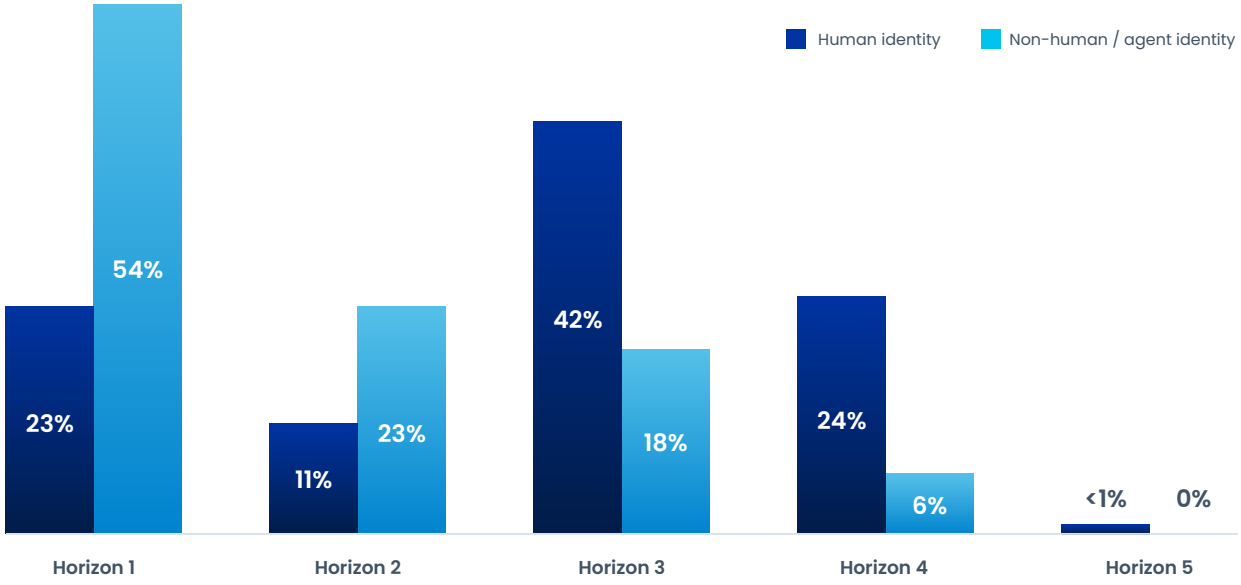
A tale of two maturities: Human governance outpaces non-human readiness

This year's data reveals a stark division: Organizations are progressing along two entirely different maturity timelines. When this study began in 2022, 45% of organizations sat at Horizon 1, the least mature tier, with identity programs still fragmented and manual. Five annual editions later, human identity maturity has been cut nearly in half: Just 23% remain at Horizon 1, and 67% have reached Horizon 3 or beyond.

Agent identity security is starting with a much steeper deficit. Today, 54% of organizations sit at Horizon 1 for agent identity maturity, surpassing the 45% starting point established by human programs in 2022 (Exhibit 2.3). If agent identity follows the same five-year curve, adequate maturity doesn't arrive until the early 2030s. With 79% of organizations already running agents in production today, that timeline is not viable. The five-year human curve has to become a one-year agent curve.

Exhibit 2.3:

The non-human view of the Index sits clearly to the left of the human view



Source: SailPoint Horizons of Identity Security survey, N=340. 2026 Horizons Maturity Index, human-identity and non-human identity component views, market-level construction (n=418; 78 screened respondents assigned to Horizon 1 in both views).

This disconnect is not a reflection of how well identity teams do their jobs. It reflects a skills and tooling gap for a fundamentally different problem. Nearly nine in ten leaders report strong technical capabilities for managing human access, yet nearly a third struggle to extend those same standards across cloud workloads and agentic environments. The specialized skills and purpose-built tools this requires are still early in development across the market.

Rather than a failure of competence, enterprises face an operational coverage gap spanning four key areas:

- **Strategy:** AI adoption is moving faster than identity planning; most organizations are only beginning to align their non-human identity frameworks with broader enterprise AI roadmaps.
- **Talent:** Identity teams possess deep expertise in traditional IAM, but specialized skills in securing autonomous agents remain in early development across the market.
- **Process:** While employee onboarding and access reviews follow mature, repeatable playbooks, non-human lifecycles are frequently managed through fragmented, ad hoc procedures.
- **Technology:** Foundational capabilities like compliance reporting are relatively stable, but dynamic, machine-speed controls, such as real-time authorization and shadow AI discovery, represent the least mature capabilities across the enterprise.

Securing the autonomous enterprise does not require rebuilding your identity program from scratch. The immediate priority is extending the proven governance disciplines already protecting the human workforce to cover the unmanaged non-human identities operating across your digital estate.

Inside the maturity index: Where organizations lead and lag

An anatomy of the maturity score reveals a market that is strong on ambition but constrained on execution. Organizations feel confident in their people and their plans, scoring highest in the Personnel and Strategy dimensions. However, these high scores mask a fundamental weakness in the underlying infrastructure, with Operating Model and Technical Capabilities lagging significantly behind. This creates a dangerous disconnect where strategy outpaces the technical ability to deliver, leaving a critical gap between plans and real-world, machine-speed enforcement. (Exhibit 2.4)

Exhibit 2.4:

People, strategy, and confidence score ahead of operating model and technical capability



Source: SailPoint Horizons of Identity Security survey, N=340. 2026 Horizons Maturity Index dimension scores, each scored 0–5: Strategy (Q16A), Personnel (Q16B), Operating Model (Q16C), Technical Confidence (Q17), Technical Capabilities (Q16D, 12-item average).

What sets advanced organizations apart?

Organizations operating at the upper horizons (Horizons 4 and 5) are not simply doing more of the same: They operate with a fundamentally different security philosophy. Rather than treating identity as a reactive IT support function, advanced organizations position identity security as an essential accelerator for business velocity, autonomous AI adoption, and enterprise resilience.

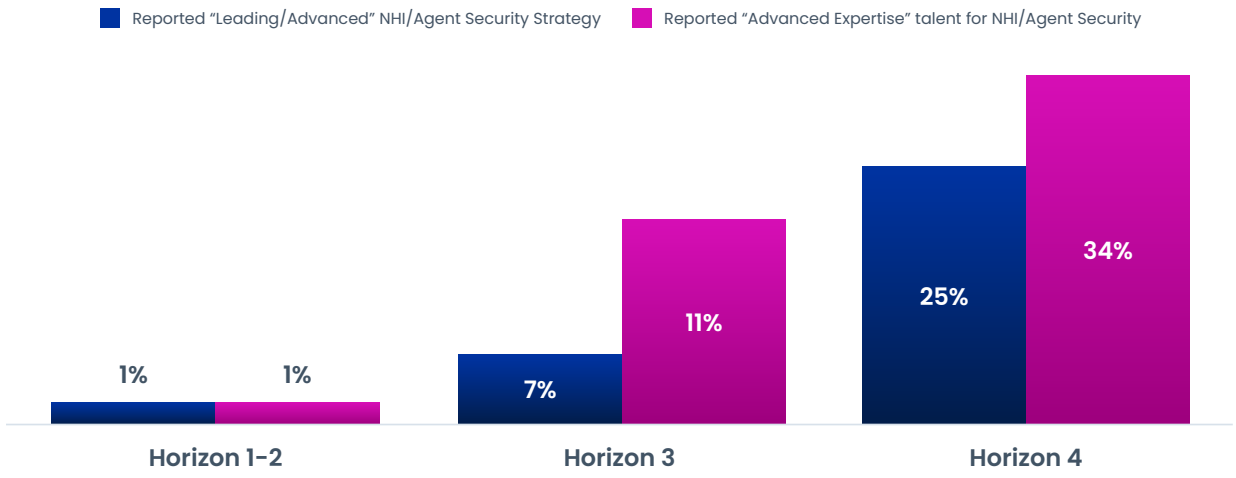
Our research reveals three core differentiators that set these leaders apart:

1. They treat non-human identity as a core pillar of enterprise AI strategy

At foundational levels, identity teams are kept at arm's length from AI initiatives, often learning about new autonomous agent deployments only after they are already in production. In contrast, advanced organizations embed identity governance directly into the foundation of their enterprise AI roadmap. Leaders at Horizons 4 and 5 invest heavily in developing specialized in-house expertise around agentic architectures, delegated authority, and machine credential lifecycles. By building this capability internally, they bridge the historical divide between security teams and development units, ensuring that security guardrails are designed into autonomous workflows from day one. (Exhibit 2.5)

Exhibit 2.5:

Self-reported scores for Strategy and Talent rise with each Horizon



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Selecting the top tier of each scale ("Leading" / "Advanced"), grouped by horizon. "Which option best describes how your organization's identity strategy addresses non-human and AI agent identities?" "Which option best describes your identity team's capabilities and understanding of AI agent controls and AI security?"

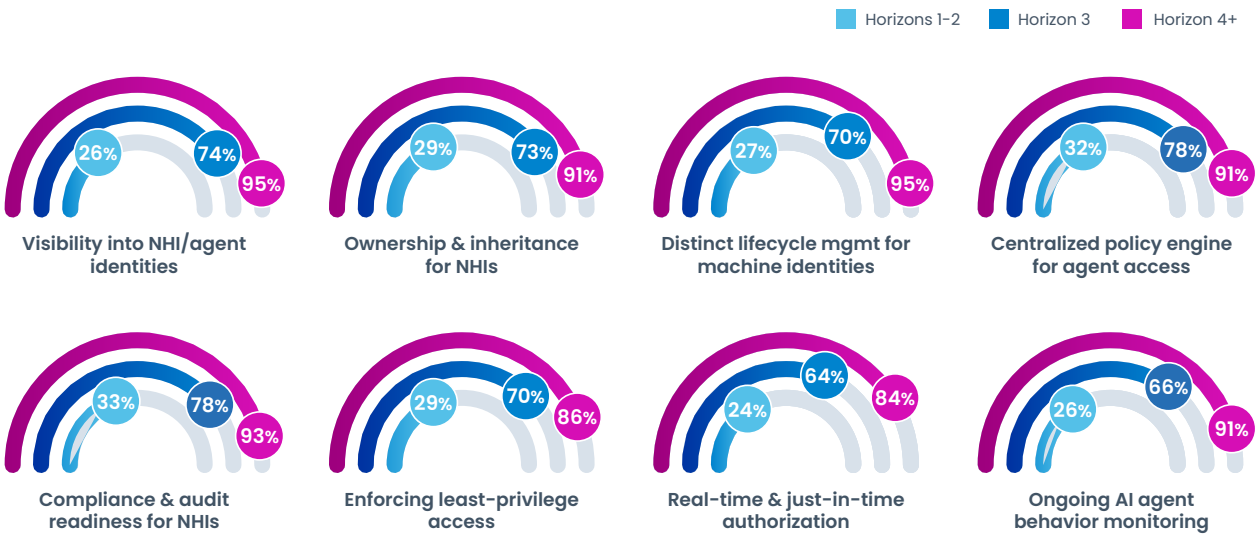
2. They master the entire visibility-to-accountability chain

While early-stage organizations struggle with baseline inventory and credential sprawl, advanced enterprises have established end-to-end governance across their entire non-human estate: They see every identity, assign explicit ownership, and continuously enforce policy. For agent-to-agent workflows, that ownership extends to delegated authority itself, tracing which human or system originally authorized an autonomous agent to act.

The single largest technical differentiator for Horizon 4 and 5 organizations is the transition from static, standing privileges to machine-speed dynamic authorization. Rather than granting persistent access that sits exposed for months, mature organizations enforce zero standing privilege (ZSP) and real-time, just-in-time access for both human engineers and autonomous agents. This drastically shrinks the attack surface while eliminating manual ticketing delays. (Exhibit 2.6)

Exhibit 2.6:

Organizations at Horizons 4–5 have far more mature technical capabilities than organizations at Horizons 1–2



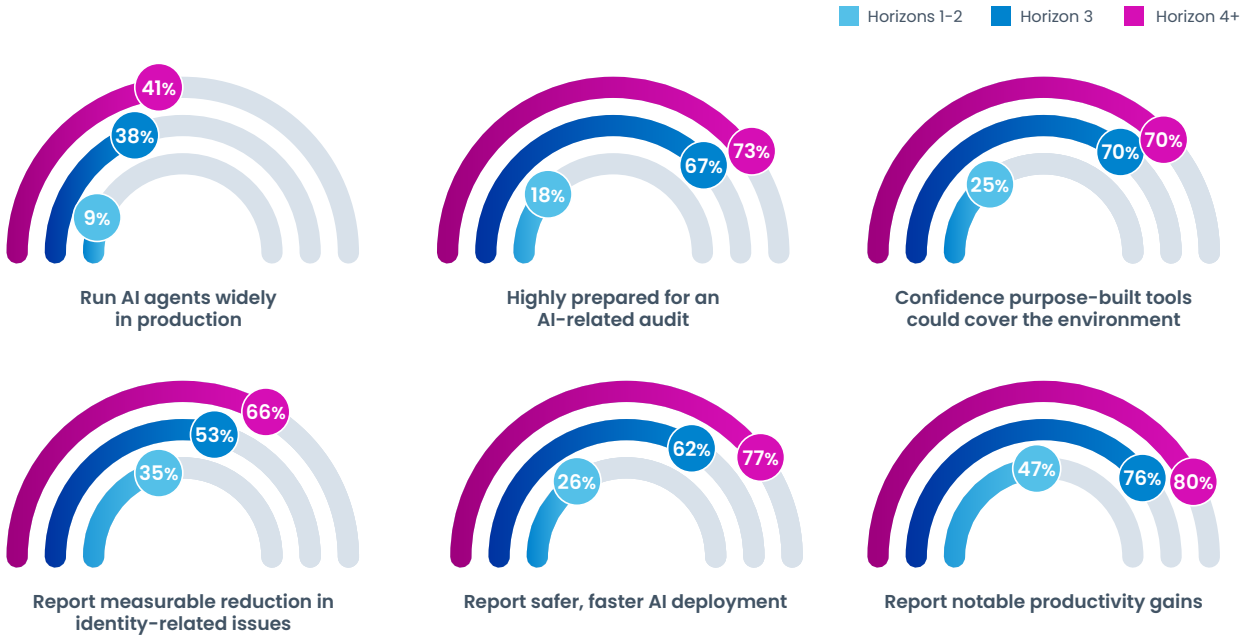
Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Q16D, 8 of 12 items with the largest spread from Horizon 1–2 to Horizon 4–5, top-2-box ("strong/very strong capability") by Horizon group. (Rate your organization's current capability (1–5) in each of 12 technical-control areas for non-human and AI-agent identities.)

3. They transform governance into a measurable business accelerator

The true hallmark of advanced maturity is business velocity. While lower-horizon organizations often view security as a bottleneck, Horizon 4 and 5 leaders turn identity into an innovation enabler. By replacing manual ticketing delays with continuous, policy-driven automation, they achieve faster and safer AI deployments, on-demand audit readiness, and measurable productivity gains across developer and security teams. (Exhibit 2.7)

Exhibit 2.7:

Advanced maturity travels with wide deployment, audit readiness, and reported value



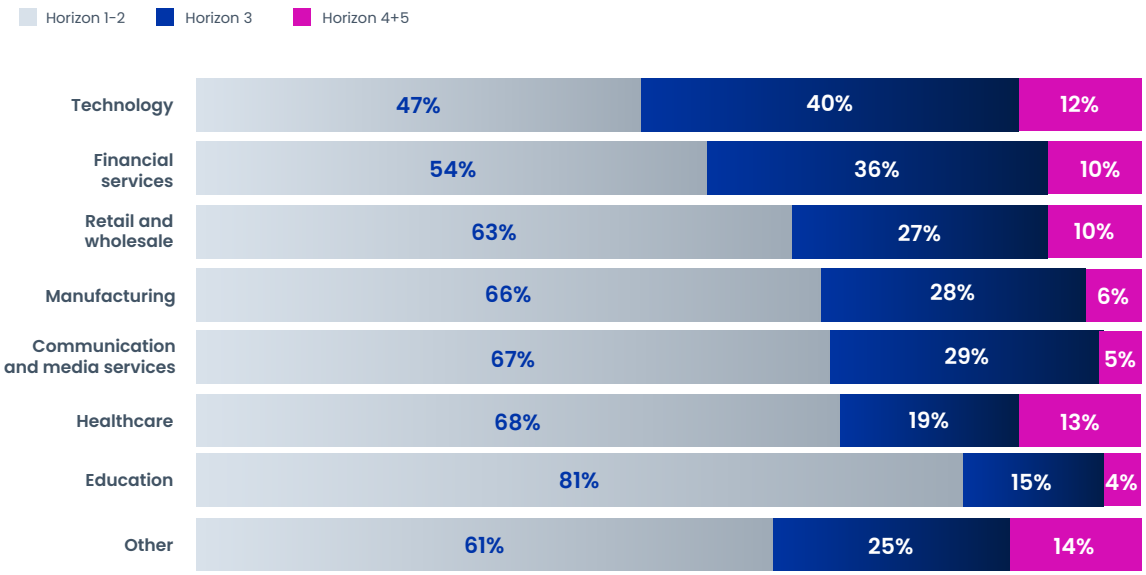
Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). S9A, Q20B, Q21, Q25A, Q25B, Q25C by 2026 Horizons Maturity Index band (n=340; Horizons 1–2 n=171, Horizon 3 n=125, Horizons 4–5 n=44).

Industry and regional dynamics

Maturity varies substantially across sectors and geographies. Technology and Financial Services lead the market, driven by heavy cloud adoption and strict compliance requirements. Healthcare and Education lag significantly, encumbered by legacy systems and resource constraints. (Exhibit 2.8)

Exhibit 2.8:

Maturity distribution across industries. Technology and Financial Services lead on Horizon level



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Horizon level split by industry; market-level base of n=418 assigns screened-out respondents to Horizon 1.

Regional dynamics: Global variance in identity maturity

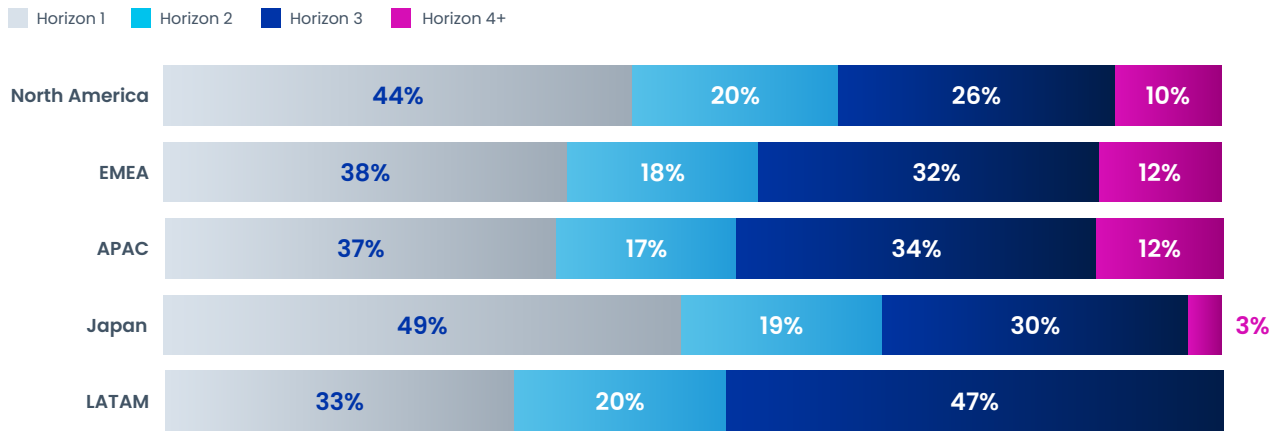
Global maturity levels reflect distinct regional priorities and architectural legacies rather than uniform progress (Exhibit 2.9):

- North America leads in advanced, AI-driven access controls, driven by aggressive public cloud adoption and agentic deployment velocity.
- EMEA shows the strongest mid-tier governance and audit readiness, propelled by strict regulatory pressures like DORA and the EU AI Act.
- APAC (ex-Japan) demonstrates rapid cloud-native adoption and high agility, steadily transitioning from foundational to digitized governance.
- Japan exhibits exceptional governance process discipline but remains structurally anchored in foundational tiers due to custom on-premises infrastructure and outsourced IT operating models.
- LATAM is actively consolidating identity tools as enterprises prioritize baseline access hygiene amid expanding cloud migrations, yet as a whole, the region shows the most opportunity for foundational maturity growth.

Across all geographies, the common tipping point for advancement remains identical: evolving from manual, regional stopgaps to a unified, machine-speed identity fabric.

Exhibit 2.9:

Horizon distribution varies sharply across all five regions



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Horizon level split by region; market-level base of n=418 assigns screened-out respondents to Horizon 1.

Chapter 3

Ambition, reality, and the path to true maturity



The awareness vs. reality gap

A pervasive theme in the 2026 data is the significant disconnect between leadership awareness and operational execution. Across multiple survey measures, broad sentiment questions yielded high optimism, while granular technical questions exposed critical gaps.

Consider the contrast.

- **The perception:** A majority of leaders (57%) express high confidence in their organization's ability to adapt and adopt the tools required to govern AI-agent identities.
- **The reality:** Over half of enterprises struggle with foundational hygiene, rating real-time monitoring (57%) and credential lifecycle management (58%) as difficult operational hurdles.

This is not simply a matter of leaders being overly optimistic; it reveals a deeper awareness gap. For many leaders, "identity security" still primarily means securing the human workforce, and their confidence reflects years of investment in maturing those legacy processes. The data shows this confidence is misplaced in the agentic era, exposing a fundamental lack of awareness of what 'good' now looks like. It is a strategy-execution gap where the ambition to secure the autonomous enterprise has outpaced the operational ability, and even the vocabulary, to do so.

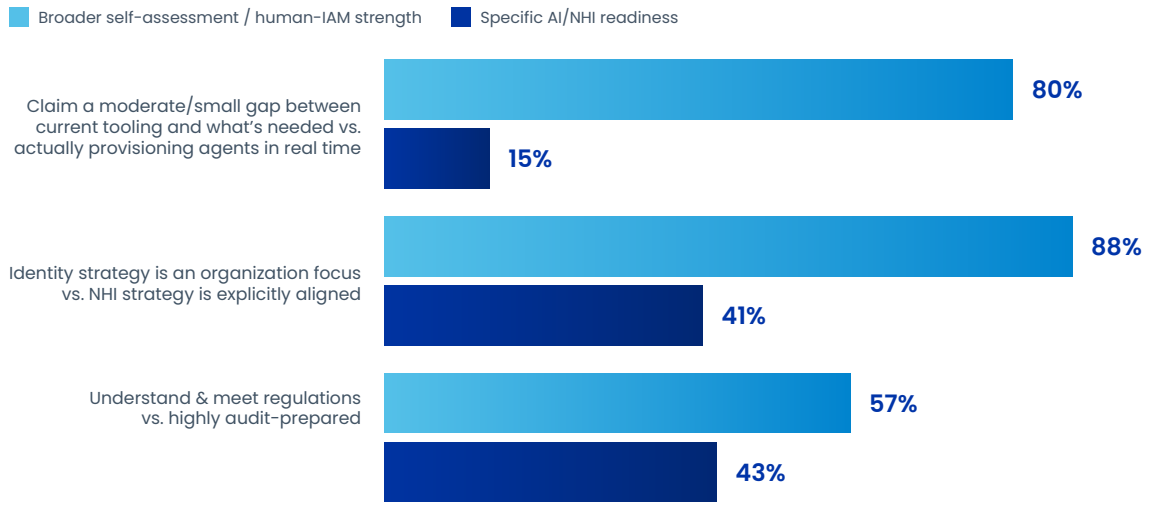
> "Leaders have already had the benefits of AI set as expectations before anyone knows how to deliver them, so the widest gap is between that overexuberant expectation and the actual ability to deliver… The more durable forcing functions are the compliance and regulatory obligations you have to meet to operate the business."
>
> Mitch Ashley
> Industry Analyst
> The Futurum Group

This gap manifests across three critical domains (Exhibit 3.1):

1. **The regulatory gap:** While a majority of leaders (57%) express confidence in their ability to meet regulatory requirements, less than half (43%) feel highly prepared to produce the specific evidence required for an AI-related audit.
2. **The strategy gap:** An overwhelming 88% of organizations rate their human identity strategy as a focus area or strategic enabler, showing how far human-identity has matured. Yet less than half (41%) have a non-human identity strategy that is explicitly aligned with their enterprise AI roadmap, evidence that this same strategic maturity has not yet transferred to agents. This suggests that "identity strategy" for many still primarily means the human workforce.
3. **The tooling and operating model gap:** Most strikingly, 80% of leaders believe the gap in their current tooling is moderate or smaller. The operational reality is that only 15% of organizations can provision non-human access in real time, the baseline requirement for securing autonomous systems operating at machine speed.

Exhibit 3.1:

Leaders self-assessment runs well ahead of specific readiness on regulation, strategy, tooling, and operating-model readiness



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Q19 and Q20B (4 or 5 on a 5-point scale); H1 (3 to 5) and Q16A (4 or 5, aligned or leading); Q20 (gap rated 1 to 3) and Q15 (effectively real time). Each pair sets a broad self-assessment against a specific measure; the measures differ in construct and the gaps are illustrative, not a single scale.

As one CISO from a large educational institution noted after a recent review, this gap has tangible consequences: "I personally reviewed the administrative console of one of our enterprise AI platforms and found approximately 30 licenses still assigned to accounts that should have been offboarded... That exposed a gap: Disabling the primary identity did not automatically trigger the complete lifecycle management of every downstream AI entitlement. Without an integrated identity fabric, offboarding remains incomplete, leaving unmonitored access paths wide open."

The path to maturity: Discover, Govern, Protect

Maturity does not move through Discover, Govern, and Protect as three sequential, self-contained stages. Each phase advances all three disciplines at once, just at different depths. Early phases plant the seeds of governance and protection long before they become the primary focus, and later phases still depend on the discovery and governance work done before them.

Phase	Discover	Govern	Protect
Phase 1 (Horizons 1 and 2)	**Primary focus.** Centralized machine inventory, shadow AI discovery, clear account ownership.	Emerging. Initial ownership assignments and ad hoc access rules begin to take shape.	Minimal. Manual, reactive incident response with no automation.
Phase 2 (Horizon 3)	Sustained. Inventory is continuously updated as new accounts and agents appear.	**Primary focus.** Policy engines, automated joiner-mover-leaver processes, credential lifecycle standards.	Emerging. Early just-in-time access for high-risk accounts and basic automated alerting.
Phase 3 (Horizons 4 and 5)	Sustained. Visibility becomes real time and self-updating.	Sustained. Policies enforced dynamically and continuously, with no manual upkeep required.	**Primary focus.** Real-time JIT access, continuous audit trail, autonomous threat remediation.

Phase 1: Discover (Horizons 1 and 2). Finding the blind spots

Organizations cannot protect what they cannot see. With 29% of enterprises reporting limited or absent non-human visibility, Phase 1 establishes an exhaustive, centralized inventory of every service account, bot, API token, and AI agent. This is primarily a discovery exercise, but it also lays the first bricks of governance, since every newly discovered identity needs an assigned owner, and the first bricks of protection, since visibility is what makes even basic incident response possible.

Phase 2: Govern (Horizon 3). Establishing control and standardization

Once visibility is established, organizations implement structured policy frameworks. The focus shifts to standardizing access models, assigning organizational owners to every machine identity, and formalizing lifecycle processes to eliminate orphaned accounts. Discovery work continues in the background to keep the inventory current, and the first real protective capabilities emerge here too, typically just-in-time access for the highest-risk accounts and automated alerting on policy violations, ahead of full machine-speed enforcement.

Phase 3: Protect (Horizons 4 and 5). Automating security at machine speed

At advanced maturity, governance transitions to real-time enforcement. Organizations deploy just-in-time provisioning, intent-based access policies, and automated threat containment to protect dynamic environments without introducing human latency. None of this works without the discovery and governance foundations built in the earlier phases: Real-time protection is only as good as the inventory and policies feeding it.

Lessons from the road: The power of a unified identity fabric

A common misconception in the market is that securing autonomous agents requires buying disjointed, standalone security point products. The data disproves this premise: Organizations that bring AI agents into their centralized enterprise identity platform achieve control maturity on par with or ahead of any siloed alternative.

The research shows that 64% of organizations already govern emerging machine and agent identities through their existing centralized identity platforms. Rather than creating new operational silos with point solutions, leading enterprises recognize that a single, unified identity fabric, one that governs employees, third parties, service accounts, and autonomous agents together, provides the comprehensive visibility and consistent policy enforcement required for machine-speed security.

The primary differentiator in security maturity is committing to a centralized, formal governance framework that covers every identity across the enterprise, not adopting disconnected niche tools.

Chapter 4

The AI velocity paradox: You cannot run an AI-speed enterprise with human-speed security



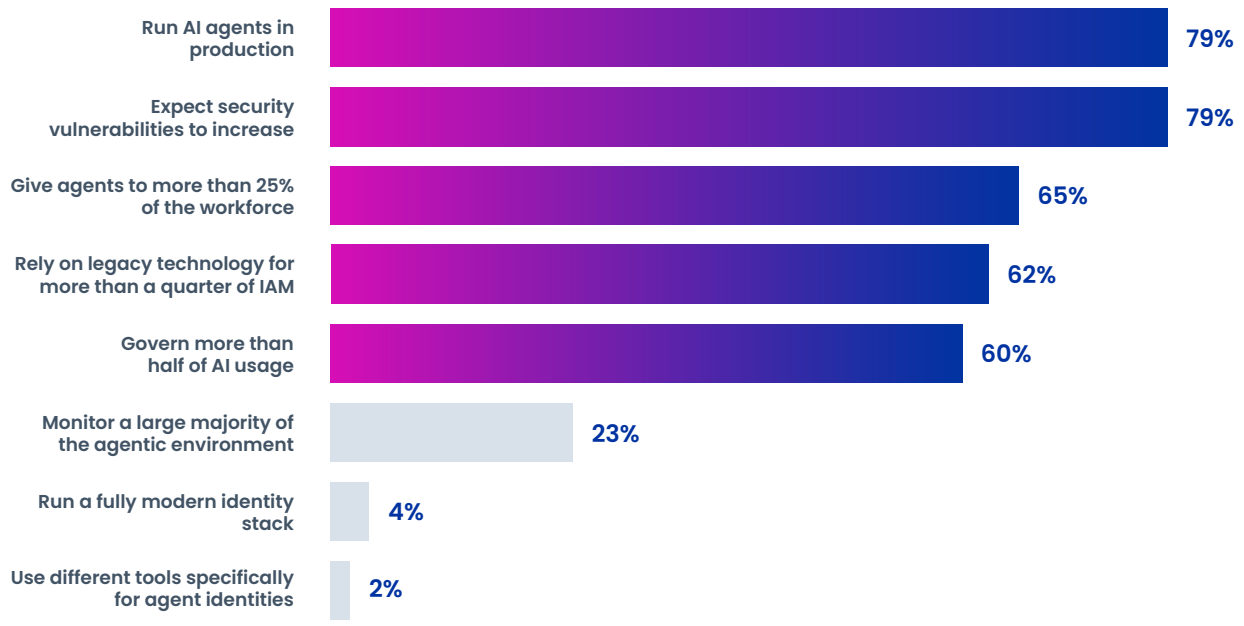
The collision of machine speed and human controls

Enterprise AI deployment represents a classic velocity paradox: Organizations are investing millions to accelerate business processes with autonomous agents, but those agents remain tethered to manual, human-speed identity controls.

> **"CIOs must urgently recognize that AI agents introduce novel and rapidly evolving risks that span security, compliance, and operational integrity, and these cannot be managed with legacy controls." "Only organizations that invest in new, adaptive risk management frameworks and AI-specific oversight will safely unlock agent-driven value."**
>
> Source: Gartner®, AI Agents: Novel Risks Will Proliferate and Require Novel AI Solutions, Rajesh Kandaswamy, Avivah Litan, 11 November 2025. GARTNER is a trademark of Gartner, Inc. and/or its affiliates.

The friction is no longer just an inconvenience; it is becoming the bottleneck that determines whether AI initiatives launch securely at all. Exhibit 4.1 tells this story with two starkly contrasting sets of numbers.

Exhibit 4.1:
AI agent adoption is running ahead of security foundations



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). S9A, Q6, S9B, Q5, Q13, Q14, S10. Independent survey measures shown together for comparison.

The insight: A 40x gap between adoption and modern security

At the top of the chart, the story is one of rapid, widespread adoption. A clear majority of enterprises are already running AI agents in production (79%), giving them to a significant portion of their workforce (65%), and governing at least half of their AI usage (60%).

At the bottom of the chart, the story is one of profound architectural lag. Only a tiny fraction of organizations are running a fully modern identity stack (4%) or using purpose-built tools to secure their agents (2%). The gap between the 79% who have adopted AI agents and the mere 2% who use specific tools to secure them is the AI velocity paradox in its purest form. It is a 40-to-1 mismatch between ambition and infrastructure. This is the single data point that explains the operational friction, security vulnerabilities, and audit gaps that define the modern enterprise security landscape.

> **"Most enterprises are still guarding the front door while the windows learn to open themselves… IDC's own numbers confirm it: non-human identity abuse now rivals stolen credentials as an attack path, and control has to move from 'approved at onboarding' to 'verified at the moment of action.' The old model trusted identity for months; agents only need milliseconds to prove that's no longer enough."**

Emanuel Figueroa
Senior Research Analyst, Identity and Access Management Security
IDC

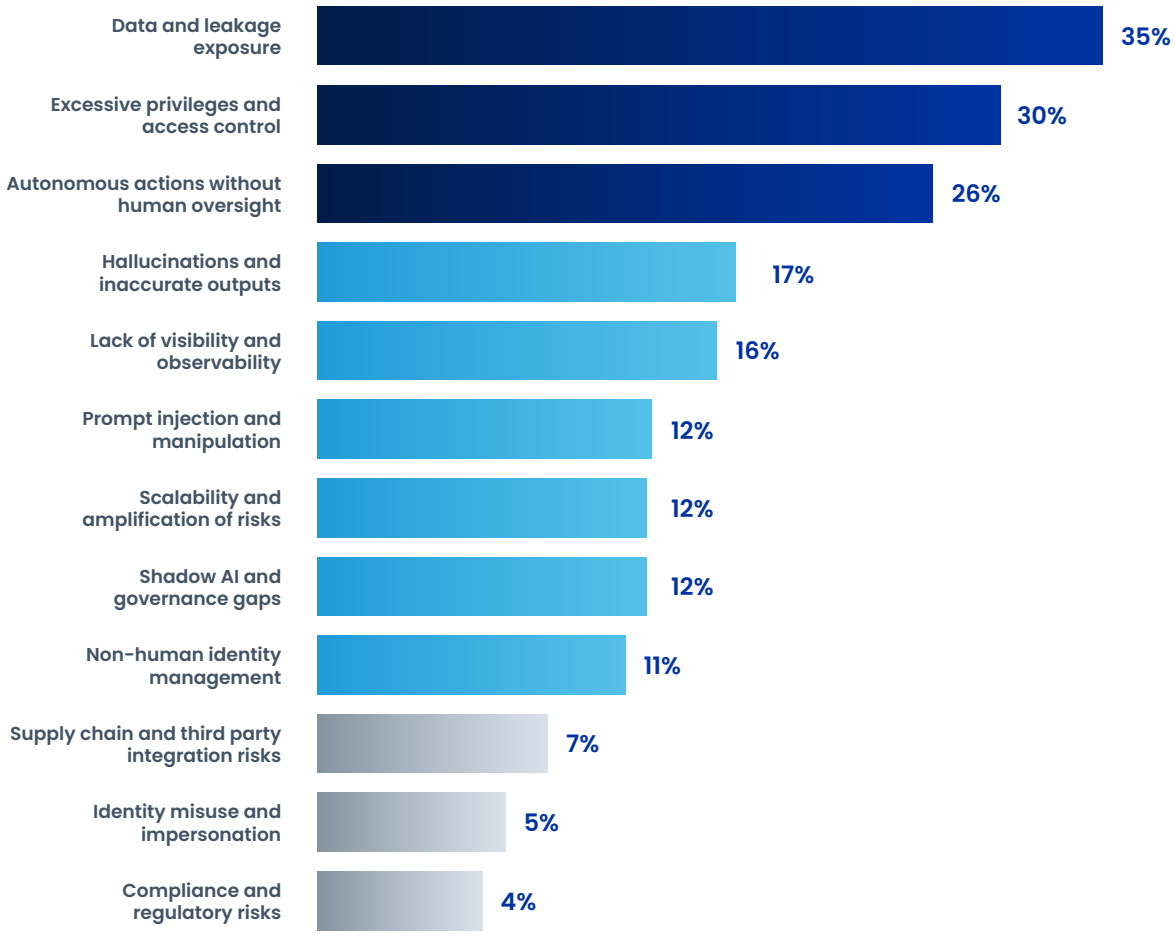
Rising threat exposure: Leaders' fears are focused on foundational access failures

The velocity paradox is not a theoretical risk; it is actively expanding the enterprise attack surface in real time. When asked what keeps them up at night, security leaders are not primarily concerned with exotic AI failure modes like model hallucinations. Their top-of-mind risks are classic identity security failures, amplified by the speed and scale of AI. As Exhibit 4.2 shows, data leakage and exposure (35%) and excessive permissions and privilege creep (30%) are the dominant concerns.

This reframes the entire AI security conversation. The primary threat isn't that agents will behave unpredictably; it's that they will behave exactly as configured: With too much access, for too long, over too much sensitive data.

Exhibit 4.2:

Data exposure and access-control risks dominate perceived AI-agent risk



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Q1: Top-of-mind AI agent risks, % of respondents citing each (coded).

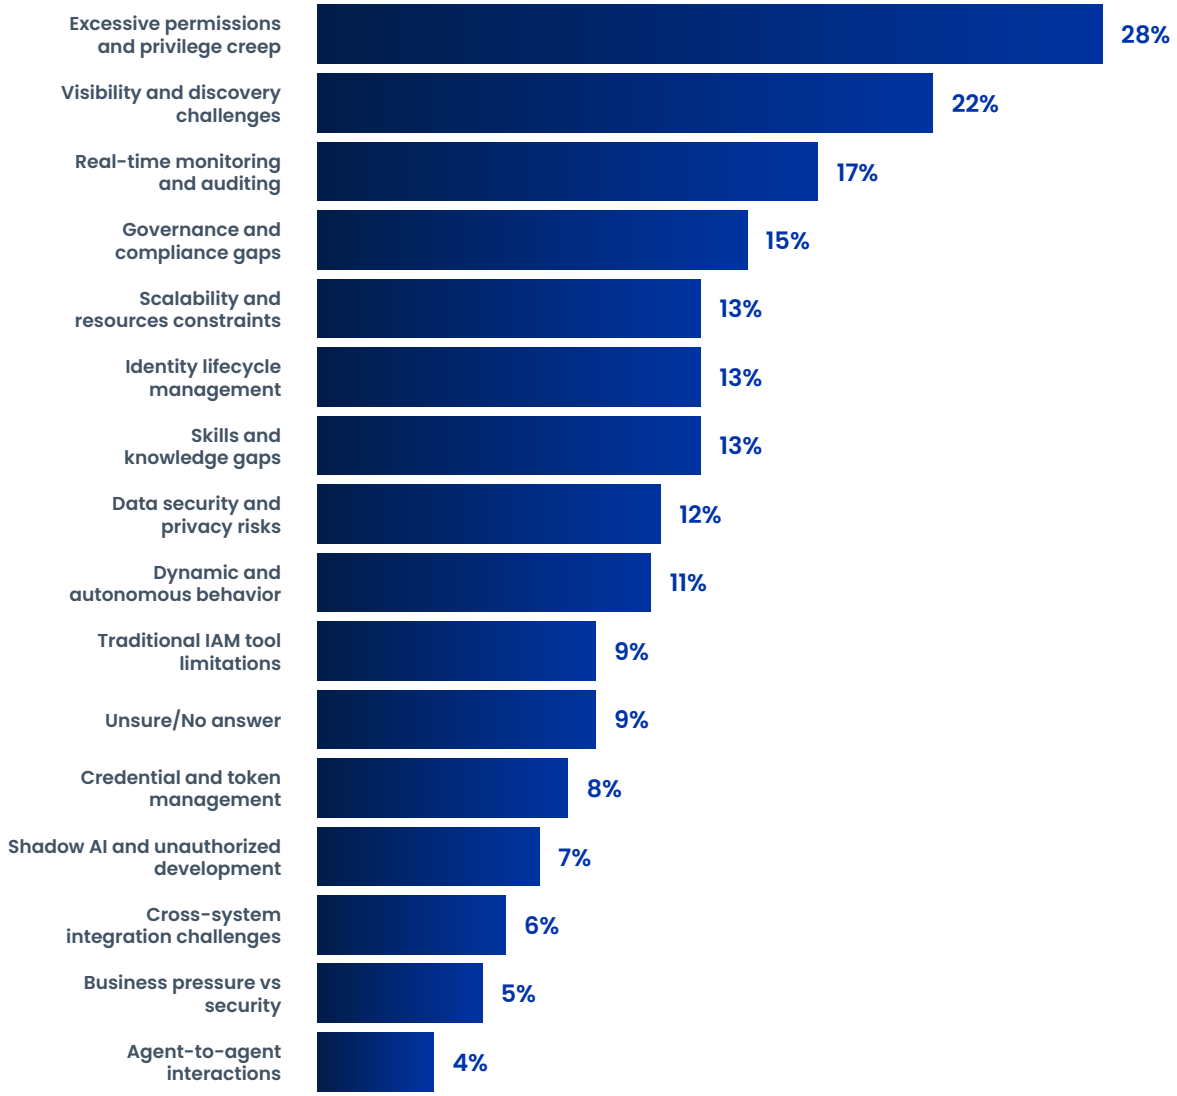
The operational reality: Alignment between what leaders fear and where teams struggle

If the primary fear is access control failure, the survey data reveals that security teams are acutely aware of the vulnerability. “Excessive permissions and privilege creep” is not only a top perceived risk; it is also the #1-cited challenge facing identity security teams today. (Exhibit 4.3)

This alignment between executive concern and practitioner difficulty is a critical insight. It proves that awareness was never the issue for organizations here. The operational model is. Security teams recognize the risk of over-permissioned agents, but legacy ticketing workflows prevent them from auditing and scoping non-human entitlements dynamically at scale.

Exhibit 4.3:

Biggest challenges facing identity security teams for managing AI-agent identities and access points



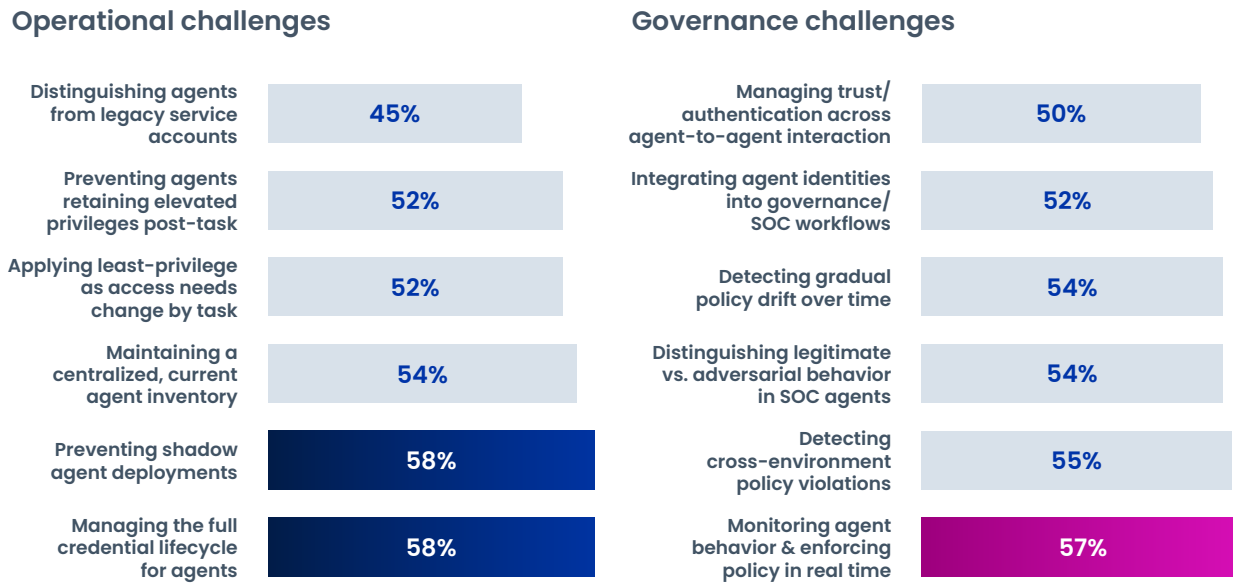
Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Q3: In your own words, what are the biggest challenges your identity security team faces when managing the identities and access points of AI agents? (Coded Open-Text Responses).

The root cause: Foundational controls remain the hardest challenges

When evaluating specific operational and governance tasks, the difficulty profile reinforces why privilege creep remains rampant (Exhibit 4.4). Core operational tasks like credential lifecycle management (58%) and shadow AI prevention (58%) are rated as severely difficult, while top-tier governance tasks like real-time behavioral monitoring (57%) prove equally challenging.

Exhibit 4.4:

Operational vs. governance challenges in managing AI agents



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Q4A/Q4B (% rating 4 or 5 on a 5-point difficulty scale).

The insight here is that you cannot separate operations from governance in the agentic era. A failure to manage the credential lifecycle is a governance failure. An inability to detect shadow AI is an operational blind spot. The data shows that the market is struggling with both sides of this equation simultaneously, indicating that point solutions for one area will not solve the systemic problem.

The false compromise: Speed vs. security

This widespread operational difficulty forces enterprises into a difficult position, creating what appears to be a compromise between speed and security (Exhibit 4.5). When faced with the choice, nearly half of the market (49%) claims to "balance both equally," while a cautious 38% lean toward governance. Only a small fraction (13%) prioritizes speed above all else.

Exhibit 4.5:

Few organizations are willing to accelerate past the risk

Stated posture when AI deployment speed conflicts with security and governance



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Q27 "Organizations often face a tension between deploying powerful AI capabilities quickly and maintaining effective security and governance over them. Thinking about how your organization currently operates, where does your organization tend to fall on the following spectrum?"

However, the data from the previous section reveals this to be a false compromise. A stated posture of "balance" without the underlying operational capability to enforce it is not a strategy; it is a stall, dressed up as a compromise. Organizations are caught in the velocity paradox: They have an AI-speed ambition but a human-speed foundation, leaving them unable to move decisively in either direction. This is where most of the market currently sits, waiting for an architectural shift that can resolve the paradox.

Chapter 5

Continuous governance: Turning compliance into a competitive advantage



Moving from periodic scramble to continuous assurance

As global regulators and internal auditors turn their attention to artificial intelligence, compliance can no longer operate as a periodic, spreadsheet-driven fire drill. The sheer velocity and volume of autonomous agent activity demand a fundamental shift: From reactive audit preparation to a state of continuous assurance, where compliance is an automated and persistent byproduct of a mature identity security program.

Our research reveals a significant "audit preparedness gap." While a majority of leaders (57%) are confident in their high-level ability to understand and meet regulatory requirements, far fewer (43%) feel truly prepared to produce the specific, granular evidence an auditor would demand for AI agents.

Exhibit 5.1:

Four in five expect compliance workload to rise, while fewer than half report high audit preparedness



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Q8 (increase slightly, moderately, or substantially), Q19 (4 or 5), Q20B (4 or 5). Q8 is a respondent expectation for the next 12–24 months.

This gap is especially concerning given the operational reality: 69% of organizations already include AI agents in their identity audits, and a staggering 81% expect that workload to increase in the next 12-24 months. The market is caught in a difficult position: Audit demands are accelerating, yet the operational machinery to meet those demands with confidence is lagging.

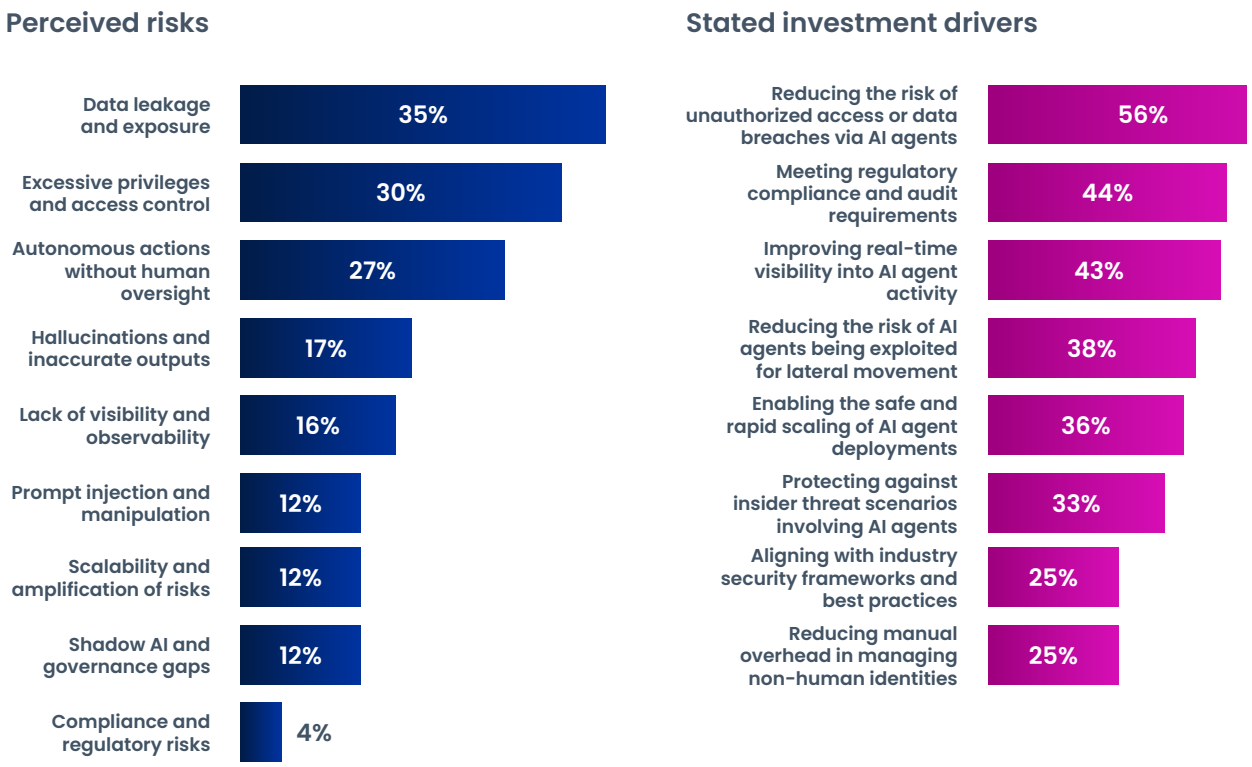
As one CISO on our advisory panel candidly noted, the bar for audits is getting higher: "AI is going to make the auditors superhuman too. Rather than sampling ten items for cost reasons, they will be able to sample everything and find every flaw… which makes me feel more exposed and tells me the audit pressure is rising even faster than the data alone suggests."

Compliance as a catalyst for investment, not a goal

If audit confidence is this thin, one might expect compliance to be the primary driver of new security investments. However, the data tells a more nuanced story. When asked to rank investment drivers, leaders place "Reducing the risk of unauthorized access or data breaches" (56%) ahead of "Meeting regulatory compliance and audit requirements" (44%).

Exhibit 5.2:

Compliance is the least-cited top-of-mind risk, the second-most-selected investment driver, and the third-most-selected future priority



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Q1 (open text, coded; top eight categories, plus bottom ranked compliance risks), Q23 (all eight items, share ranking each in top three).

This reveals a key insight into the mindset of mature security leaders: They don't invest in identity security to become compliant; they invest in it to become secure. Audit readiness is simply a valuable byproduct of a well-architected, risk-focused identity program. An advisory panel CISO confirmed this perspective: "I was not surprised that 69% say they already audit agent identities while only 10% feel extremely prepared. Those are two very different standards."

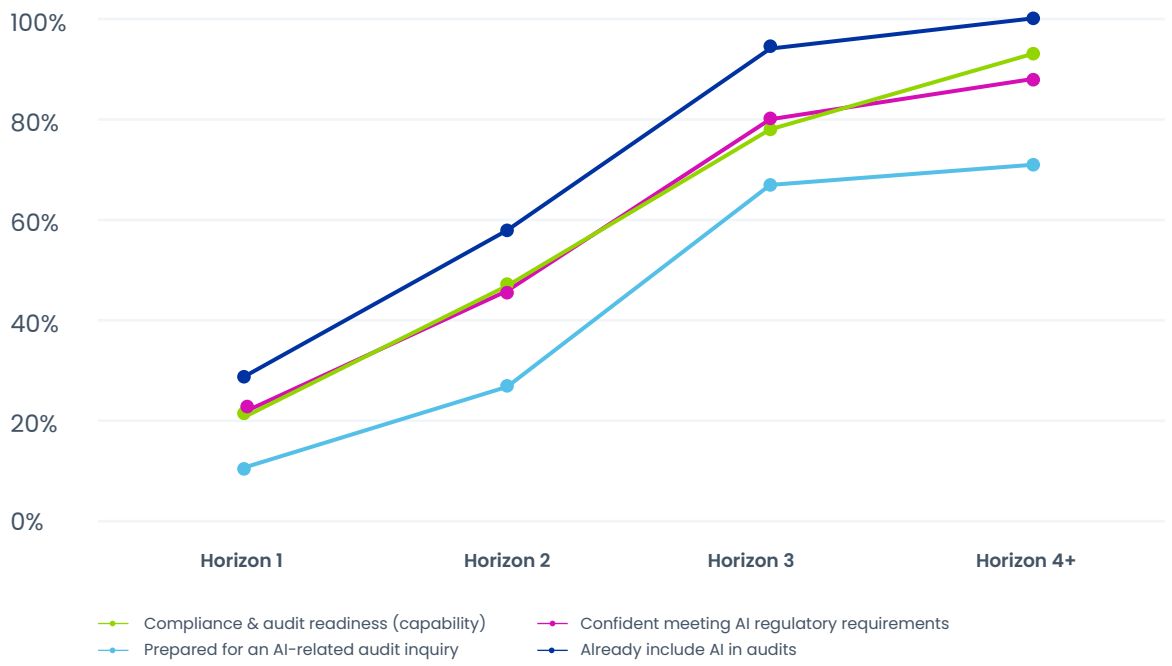
Organizations that chase compliance as a primary goal will always be in a reactive, fire-drill mode. Those that build a foundation of continuous governance achieve compliance as a natural outcome.

Audit readiness climbs with maturity

The data proves conclusively that audit readiness is not a standalone function but a direct outcome of overall identity maturity. As organizations advance through the Horizons, their composite governance readiness (a measure that combines audit preparedness with core technical controls) scales directly and impressively. (Exhibit 5.3)

Exhibit 5.3:

A composite of all four governance metrics climbs from 21% at Horizon 1 to 88% at Horizon 4 – every measure rises together, level by level.



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Q16D, Q19, Q20B, Q15B. Grouped by overall Horizon level (Levels 1–4); n=92, 79, 125, 42 respectively. Level 5 (n=2) excluded. Self-reported; correlation, not proof of causation.

As the exhibit shows, organizations in Horizons 1 and 2 struggle across the board, with only a small fraction demonstrating mature controls or audit readiness. By Horizon 3, a significant shift occurs as organizations begin to standardize processes. At Horizon 4, the capabilities converge, with a strong majority demonstrating high levels of readiness across all measures.

This powerful correlation delivers a clear verdict: The path out of the "preparedness gap" runs through a higher level of identity maturity, not a better audit process. By automating discovery, standardizing lifecycles, and enforcing policy consistently, mature organizations generate a continuous, verifiable audit trail by default, making the periodic audit a simple proof point rather than a major operational lift.

Chapter 6

The ultimate ROI: Securing the speed of business



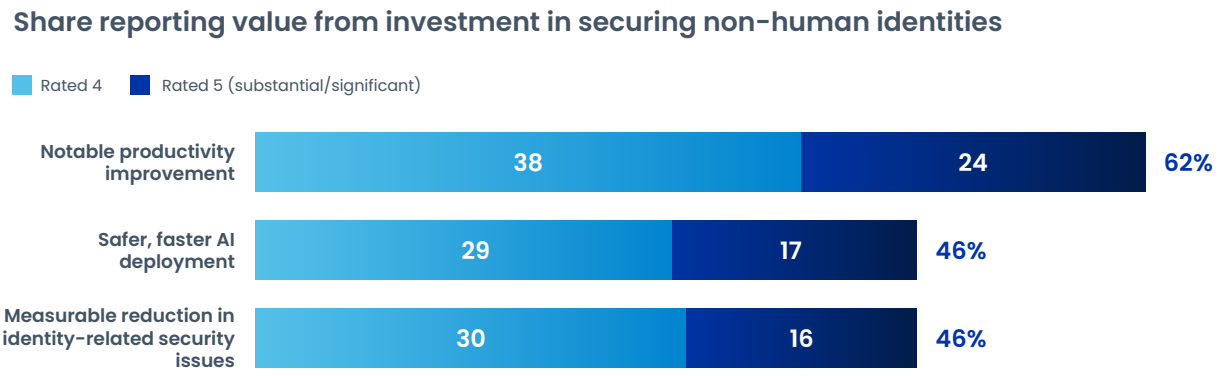
Reframing identity security from cost center to growth engine

For decades, the business case for identity security has rested on risk avoidance and mitigating regulatory penalties. The 2026 survey data reveals a far more compelling narrative: Modern identity security is a primary driver of business velocity and workforce productivity.

Securing non-human identities does more than reduce risk: It delivers direct, measurable business impact. When asked about the value realized from their investments, organizations consistently pointed to operational acceleration. 62% of organizations report a notable productivity improvement after implementing non-human identity security, leading the list of realized benefits (Exhibit 6.1). This demonstrates that when manual approval bottlenecks are eliminated, development and operations teams can innovate faster and more securely.

Exhibit 6.1:

Realized business value from securing non-human identities



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Q25A, Q25B, Q25C (4 or 5 on a 5-point scale). Midpoints: Q25A some reduction, Q25B somewhat accelerated, Q25C no notable effect. Not-yet-invested responses (about 3%) included in the base.

Tellingly, "Safer, faster AI deployment" (46%) and "Measurable reduction in identity-related security issues" (46%) are tied. This disproves the outdated notion that security and speed are in opposition. Mature identity governance is the engine that allows both to happen simultaneously. Securing agents is not the toll booth on the road to velocity. It is the road.

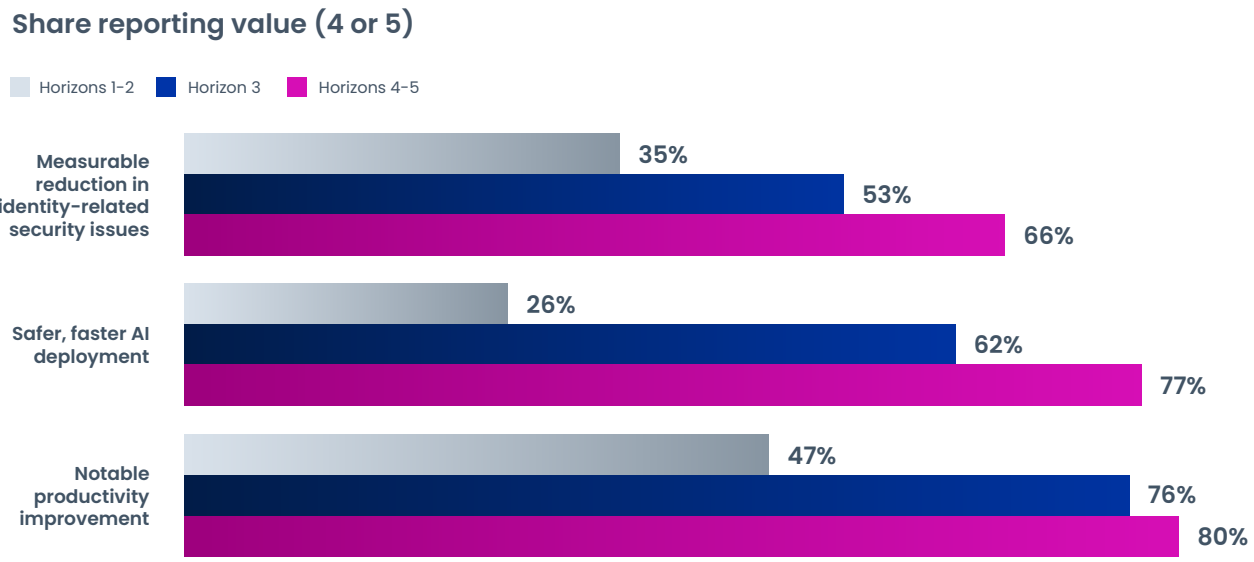
Maturity as a productivity multiplier

This value story becomes even more pronounced when viewed through the lens of maturity. The data shows a direct and powerful correlation between an organization's Horizon level and the business value it realizes.

As shown in Exhibit 6.2, the most mature organizations (Horizons 4-5) are roughly twice as likely to report significant productivity gains and reductions in security incidents compared to their peers in the foundational tiers. Most impressively, the benefit of "safer, faster AI deployment" is nearly three times higher for mature organizations. This proves that maturity pays off well beyond compliance, acting as a productivity multiplier that unlocks the full value of AI and automation initiatives.

Exhibit 6.2:

Horizon 4 and 5 organizations are far more likely to report realized value on every measure



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Q25A, Q25B, Q25C (4 or 5) by 2026 Horizons Maturity Index band (Horizons 1–2 n=171; Horizons 4–5 n=44). Not Index components; the Horizons 4–5 base is small.

Broad economic uplift across the enterprise

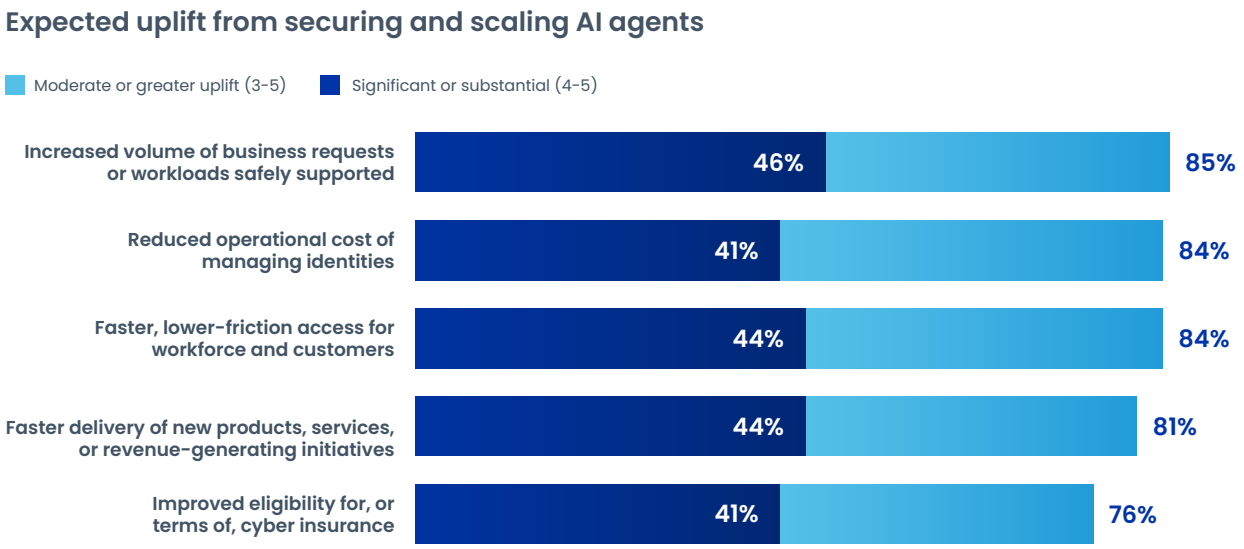
The gains from mature identity security are not confined to the identity team; they ripple outward across the entire enterprise. When asked about the future, leaders expect significant, measurable business value across a wide range of operational and financial dimensions.

As shown in Exhibit 6.3, the expected uplift is both broad and substantial. Leaders anticipate that mature identity security will allow them to safely support a higher volume of business requests and workloads (85%), enable faster, lower-friction access for employees and customers (84%), and ultimately lead to the faster delivery of new products and revenue-generating initiatives (81%).

Crucially, this value extends to the financial bottom line. A significant 76% of leaders expect their identity security posture to improve their eligibility for, or the terms of, their cyber insurance. In an era of rising premiums and increasingly stringent underwriting, the ability to demonstrate mature, automated identity governance is becoming a key lever for controlling insurance costs and proving a strong security posture.

Exhibit 6.3:

Three-quarters or more expect at least moderate uplift on every value dimension, including cyber insurance



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Q26, five dimensions (n=340). The 2026 realized- and expected-value measures update the prior Horizons TEI lens, adding explicit cyber-insurance and access-friction measures.

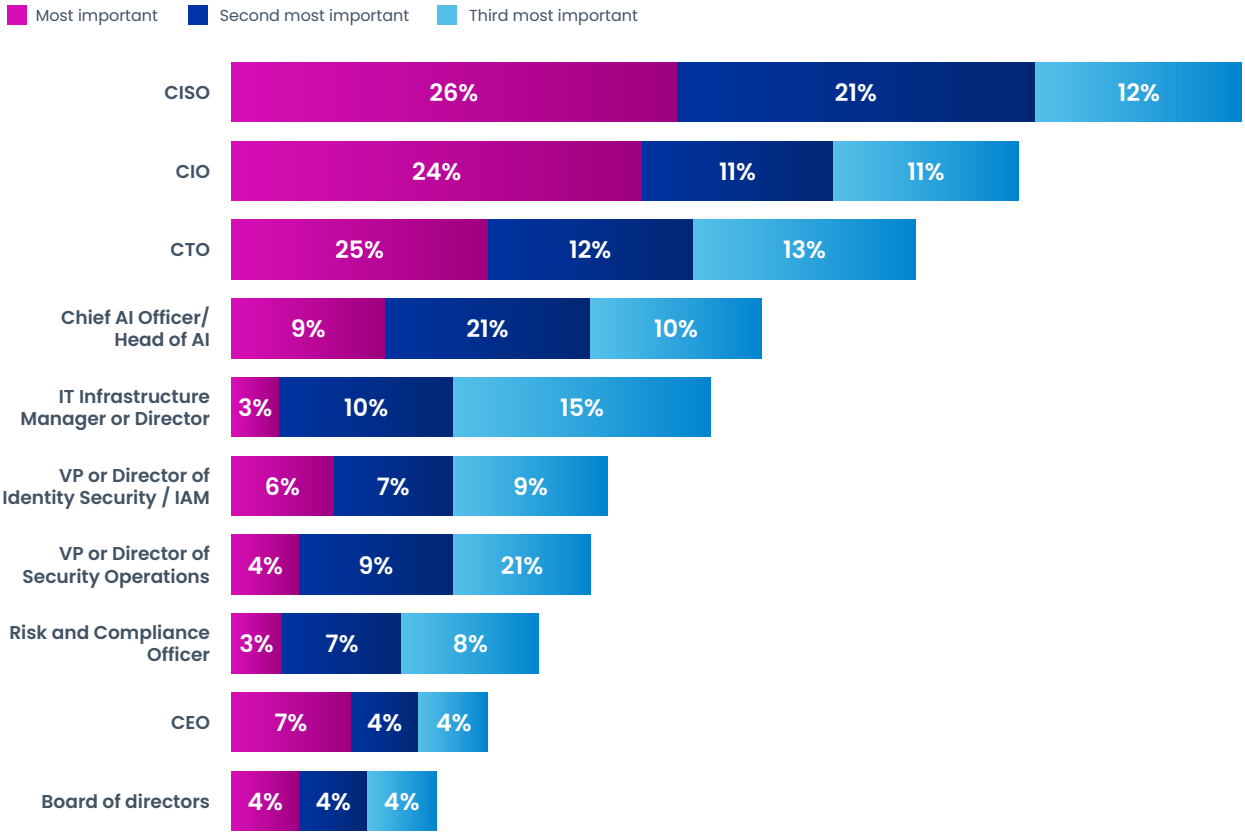
This forward-looking data completes the ROI picture. It reframes the investment from a defensive necessity to a strategic enabler that drives productivity, accelerates revenue, and delivers a tangible financial return in the form of optimized cyber insurance terms.

Making the business case: Ownership and alignment

This year's data provides a clear playbook for how to frame the business case for identity investment internally. The CISO is the most influential decision-maker (ranked #1 by 26% of respondents) for adopting new agent identity security technologies, ahead of the CIO and CTO.

Exhibit 6.4:

Key decision-makers for AI agent identity security adoption



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Q22, "Within your organization, who are the most important decision-makers when it comes to adopting new technologies related to securing AI agent identities?" Select and rank up to three (n=340).

Given the CISO's central role, the business case should be anchored in risk reduction and readiness, themes that resonate strongly with security leaders. However, the data also shows a clear opportunity to expand the conversation. As one CISO from a large HR services enterprise noted, "From discussions I've had with my peers, many security leaders are using agentic identity as the reason to finally upend their identity and access management technology stack."

The ultimate ROI of identity security is about enabling the business to move faster and more safely. By framing the conversation around business velocity, productivity, and enabling AI innovation, CISOs can build a far more compelling, enterprise-wide case for investment that resonates all the way to the boardroom, not just within the security department.

Real-world impact: Customer spotlights in identity value and scale

CRITEO

Criteo is a leading global commerce intelligence platform company. Identity management at Criteo began as a custom internal capability — a model that grew harder to sustain as the business and technology footprint expanded. Criteo started its transition to a standardized, cloud-first identity governance foundation in 2021 and went live in 2022, built on the principle that IT's role isn't to build an IGA platform, but to create business value while securing identities and data.

The evolution: From custom stack to identity-first foundation

Criteo's identity program evolved through three distinct phases:

- **Building a standardized foundation (2021–2022):** Replaced an eight-year-old custom identity-management stack, previously maintained by a team of eight, with SailPoint IdentityNow (now Identity Security Cloud). From day one, governance scope extended beyond employees to include technical accounts and machine identities, each with an accountable owner, defined purpose, lifecycle dates, approval workflows, access reviews, and an audit trail.
- **Maturing machine identity governance (ongoing):** Migrated from the legacy technical-account model to Machine Identity Security (MIS) as the capability matured, building a stronger foundation and practical experience for the next evolution.
- **Extending to AI agents (2025–present):** Applying the same identity-first principles established for humans and machines to AI agents, while explicitly building and validating the model rather than treating it as solved.

Business impact & scale

More than four years after going live, Criteo's identity platform is now operated by a team of four, down from the eight required to maintain the previous custom stack, while creating more business value in the process.

Criteo key impact metrics

84 sources, 320 applications: Governance at meaningful scale, with standardized controls across the enterprise application landscape.

5,500 access profiles, 80 workflows: Structured, repeatable access governance built for scale, not one-off exceptions.

2,000+ certification campaigns in 50 months: Identity governance operates as a continuous process, not a periodic project, with a repeatable model for reviewing and remediating access.

The next horizon: Governing the autonomous frontier

Criteo sees AI agents as fundamentally shifting the balance of power in the enterprise: Capabilities once concentrated in IT or expert teams are now in the hands of far more users. That democratization fuels innovation, but it also multiplies the identities, integrations, and access paths that need governing — and without a unified approach, risks leaving behind what Criteo describes as "an archipelago of disconnected services," where ownership and accountability become difficult to track.

> **"The objective is not to slow down AI adoption. It is to make AI adoption safe enough to scale."**
>
> Jérôme Robin
> Director of Corporate Security
> Criteo

Criteo's response is to extend the same governance discipline to every agent: Define its purpose and scope, apply least privilege, control its lifecycle, secure its authentication and secrets, and keep its actions traceable and auditable — connected to the same governance processes as every other non-human identity.

AT THE FOREFRONT
UChicago Medicine

University of Chicago Medical Center (UCMC) is a leading academic medical center navigating identity governance under the added weight of HIPAA regulation and the sensitivity of protected health information (PHI). UCMC migrated from IdentityIQ to SailPoint Identity Security Cloud roughly a year ago, part of a broader push to modernize identity governance beyond basic network-account and email provisioning.

The evolution: From basic provisioning to full lifecycle governance

UCMC's identity program is in the midst of a multi-stage modernization:

- **Foundational provisioning (before migration):** Identity management centered on basic network-account and email provisioning, with privileged-access ownership and role-based access control (RBAC) as persistent gaps — making it difficult to ensure the right people had the right access at the right time.
- **Platform modernization (2024–2025):** Migrated to Identity Security Cloud and began standing up Business Access Roles and automated application onboarding, extending governance from basic provisioning toward full lifecycle management.
- **Extending to agents (in progress):** Standing up an Agentic Security Fabric to identify UCMC's AI agent population, assigning ownership, and establishing governance controls — addressing a gap: The organization adopted AI quickly, but lacks a complete inventory of agents and who owns them.

> "It really comes down to establishing guardrails and understanding what information you want your agents to have access to. In a healthcare setting, there's a lot of HIPAA regulation, PII, PHI — so it's critical to define what data your agents can touch. Combine that with an AI advisory or governance board that sets the framework for how and where AI is used, and that's what lets you safely accelerate AI, especially in healthcare."
>
> Eric Burnett
> Director, Identity and Access Management
> University of Chicago Medical Center

The next horizon: Governing the autonomous frontier in a HIPAA-regulated environment

For UCMC, the path to trusting an agent with real access starts with visibility: Understanding and defining what information an agent should be allowed to touch, before that access is ever granted, to avoid blind spots. Every agent needs a human owner, someone accountable for monitoring its behavior and tying its activity back to a person, backed by a clear blueprint of the data and infrastructure each agent can reach.

That discipline matters more in healthcare than almost anywhere else. With personally identifiable information (PII) and protected health information (PHI) embedded throughout its systems, UCMC can't treat agent access as one-size-fits-all; access has to be constrained by the sensitivity and intended use of the underlying data. UCMC's approach pairs that technical guardrail with an institutional one: An AI advisory or governance board that sets the overarching framework for how and where AI can be used across the organization. Together, the two let UCMC move faster on AI adoption without loosening the safety controls a hospital can't compromise on.

RWE

RWE is already one of the leading companies in the field of renewable energy. Its identity program has developed in three chapters: A foundational phase building core processes and skills, a move to SaaS that laid the groundwork for AI-driven security and Zero Trust, and a current phase focused on delivering identity as a service. Throughout, RWE has scaled its identity operations with a small core team backed by a managed service.

The evolution: A three-chapter maturity journey

RWE's identity security program progressed through three distinct phases:

1. **Building the foundation (2013–2018):** Introduced IdentityIQ on-premise and migrated the deployment to AWS in 2018, formalizing core enterprise policies, standardized procedures, and internal IAM skills.
2. **Transitioning to cloud & AI (2021–2022):** Rebuilt application access around SailPoint Identity Security Cloud (SaaS), layering in AI-driven identity insights and laying the operational groundwork for Zero Trust.
3. **Delivering identity as a service (2024–2026):** Elevated identity into an agile business service: Analyzing more than 3,000 user tickets to eliminate fulfillment friction, optimize request journeys, and standardize governance for long-term scalability.

> "Every non-human identity needs someone accountable for it. Whether it supports an application, an automated process, or an AI agent, we need to understand what it can access, why it needs that access, and when that access should end."
>
> Nikk Gilbert
> Global CISO
> RWE

Business impact & scale

RWE has supported growing identity governance demand through centralized governance, standardized processes, and automated workflows.

RWE key impact metrics

~6% average monthly growth: Over a 24-month period, IGA demand grew at an average rate of roughly 6% per month across application onboarding, application changes, certification campaigns, and reporting.

600+ integrated applications: More than 5,000 self-service access requests submitted monthly, with 4,000+ access profiles available across all operating companies.

1,500+ annual certification campaigns: More than 40,000 individual accesses reviewed regularly through structured certification workflows.

The next horizon: Governing the autonomous frontier

As RWE prepares for greater use of automation, non-human identities, and AI agents, the organization is extending its governance framework to encompass machine actors and automated background processes.

RWE's core operating principle ensures that automation never outpaces accountability: Every non-human identity is cataloged with an accountable owner, an explicit business purpose, defined access boundaries, clear lifecycle limits, and the ability to review and withdraw access when it is no longer required.

Chapter 7

The mandate: Achieving radical trust



The data and trends explored in this report point to a clear and urgent mandate for security leaders. To survive and thrive in the age of AI, organizations must build an architecture where trust is radically continuous but entirely invisible to the user and the machine. Human-speed controls, periodic audits, and fragmented visibility are no longer viable. The future belongs to a continuous, autonomous, and self-healing identity fabric.

> **"Identity is foundational. It is where security starts, and if you do not have it you do not have security… Consider that an agent can now be written and executed by another agent, having never existed before, with no pre-existing code to call, because AI builds agents, it does not just run existing ones. That is your ephemeral identity problem in a sentence."**
>
> Mitch Ashley
> Industry Analyst
> The Futurum Group

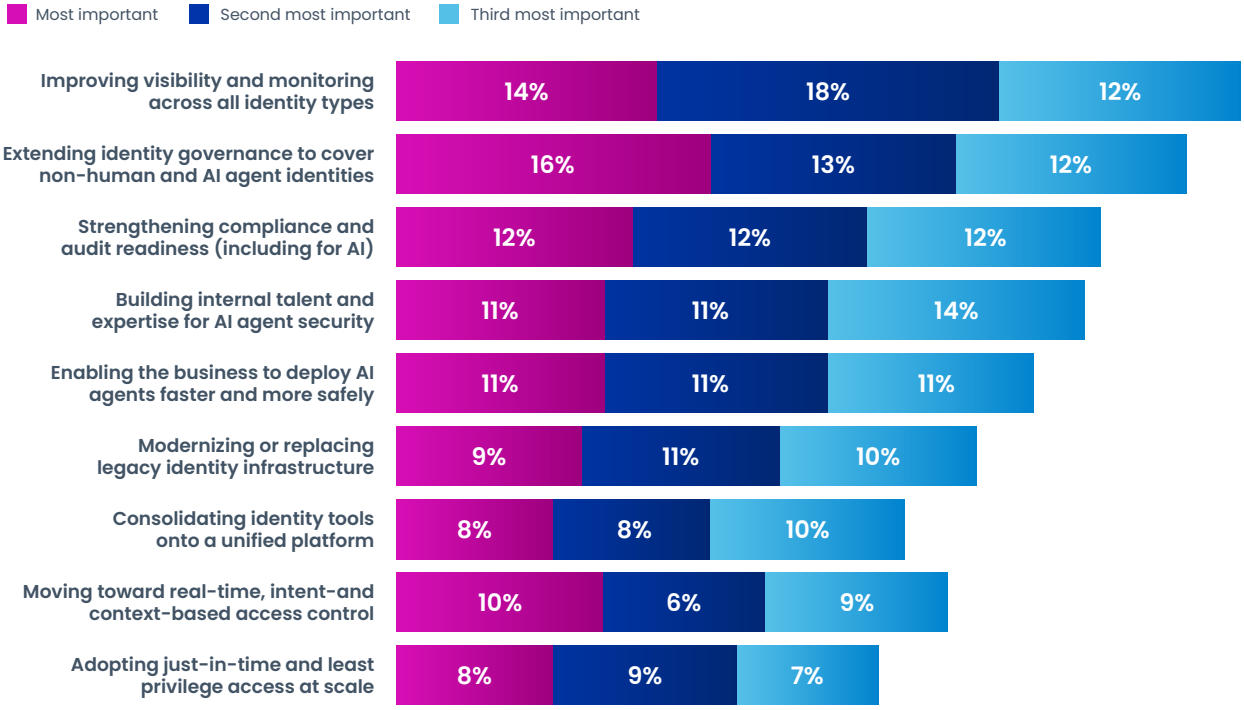
The path forward: Climbing the ladder from foundations to automation

The market is not there yet, but its own stated priorities show it is climbing toward this mandate by starting with the foundations. As Exhibit 7.1 shows, leaders are overwhelmingly prioritizing foundational work like improving visibility and monitoring (44%) and extending governance to non-human identities (41%).

The most advanced, machine-speed capabilities, such as real-time, intent-based access control (26%) and just-in-time access at scale (25%), rank at the bottom of the list. This is not a sign of neglect; it is a sign of sequencing. The market correctly understands that it cannot leap to the top of the ladder without first climbing the rungs.

Exhibit 7.1:

Foundational priorities like visibility and governance lead stated investment plans



Source: SailPoint Horizons of Identity Security survey, 2026 (n=340). Q24, "Looking ahead over the next 12-24 months, what are your organization's top priorities for identity security?". Rank up to three. Individual rank shares and top three totals are rounded independently from respondent counts; components may not sum to the displayed total.

A three-stage roadmap to continuous, autonomous identity

The journey from today's manual processes to tomorrow's autonomous security follows a clear, three-stage roadmap. The destination is continuous, invisible trust, but the path runs through these foundational stages.

> "Start with visibility, but don't assume technology will give you complete visibility. You cannot govern an agent you don't know exists, and sometimes the best discovery tool you have is the trust of your community. Build enough trust that people bring their AI use cases to you rather than hide them from you."
>
> AVP, CIO, & CISO
> Large Education Institution

Stage 1: Discover (Horizons 1–2)

The objective: Establish complete visibility across all human and machine identities.

You cannot govern what you cannot see. For organizations in the foundational tiers, the immediate priority is eliminating blind spots by creating a single, authoritative inventory of all identities across the digital estate.

Key milestones

- Catalog all non-human actors: Inventory every service account, API token, script, and autonomous agent across on-premise, cloud, and SaaS environments.
- Detect shadow AI: Identify unsanctioned agents or tools operating outside governance before they're brought into the estate.
- Assign explicit ownership: Map every machine identity and automated workflow to an accountable human owner and documented business purpose.
- Bring machines into the audit perimeter: Include non-human accounts in regular access reviews and audit readiness assessments.

Leadership guidance

Resist the temptation to deploy dynamic authorization before achieving baseline discovery. Build from your established identity capabilities and leverage a unified identity fabric to inventory the estate before attempting machine-speed enforcement.

Stage 2: Govern (Horizon 3)

The objective: Standardize controls and automate lifecycles across the identity fabric.

Once visibility is established, the focus transitions from discovery to structured governance. Organizations move away from ad hoc, ticket-based approvals toward consistent, policy-driven lifecycles.

Key milestones

- Standardize joiner-mover-leaver (JML) for machines: Automate the provisioning, privilege adjustments, and decommissioning of service accounts and AI agents at end of life.
- Enforce automated credential rotation: Eliminate hardcoded passwords and long-lived API keys in favor of managed vaulting and scheduled rotation.
- Align identity with enterprise AI roadmaps: Establish formal governance guardrails for development teams deploying agentic workflows.
- Unify human and agent context: Extend the same identity graph from discovery to link human and agent signals together, so no governance decision is made on agent context alone.

Leadership guidance

Turn maturing, ad hoc processes into codified enterprise standards. Focus on eliminating standing privileges across critical infrastructure and SaaS platforms to shrink the attack surface.

Stage 3: Protect (Horizons 4–5)

The objective: Achieve machine-speed enforcement and continuous assurance.

At advanced maturity levels, identity governance transitions from scheduled reviews to active, real-time protection. Security operates continuously and invisibly in the background.

Key milestones

- Combine human and agent context in real time: Enforcement decisions should draw on both agent behavior and the human or workflow that authorized it — you cannot secure an agent without that context, or vice versa.
- Implement zero standing privilege (ZSP): Transition to just-in-time (JIT) access that grants ephemeral privileges only when a specific task is initiated and revokes them immediately upon completion.
- Deploy dynamic, context-aware authorization: Adjust permissions in real time based on continuous behavioral signals, device posture, and data sensitivity.
- Integrate identity telemetry with SecOps: Feed real-time identity signals directly into SIEM/SOAR platforms to enable automated containment (e.g., session termination, token revocation) during active threats.

Leadership guidance

Treat identity as the central nervous system of your security operations. With continuous governance in place, compliance reporting becomes an automated byproduct of daily operations, freeing your security team to focus on strategic business enablement.

Beyond the horizon: The five levels of identity autonomy

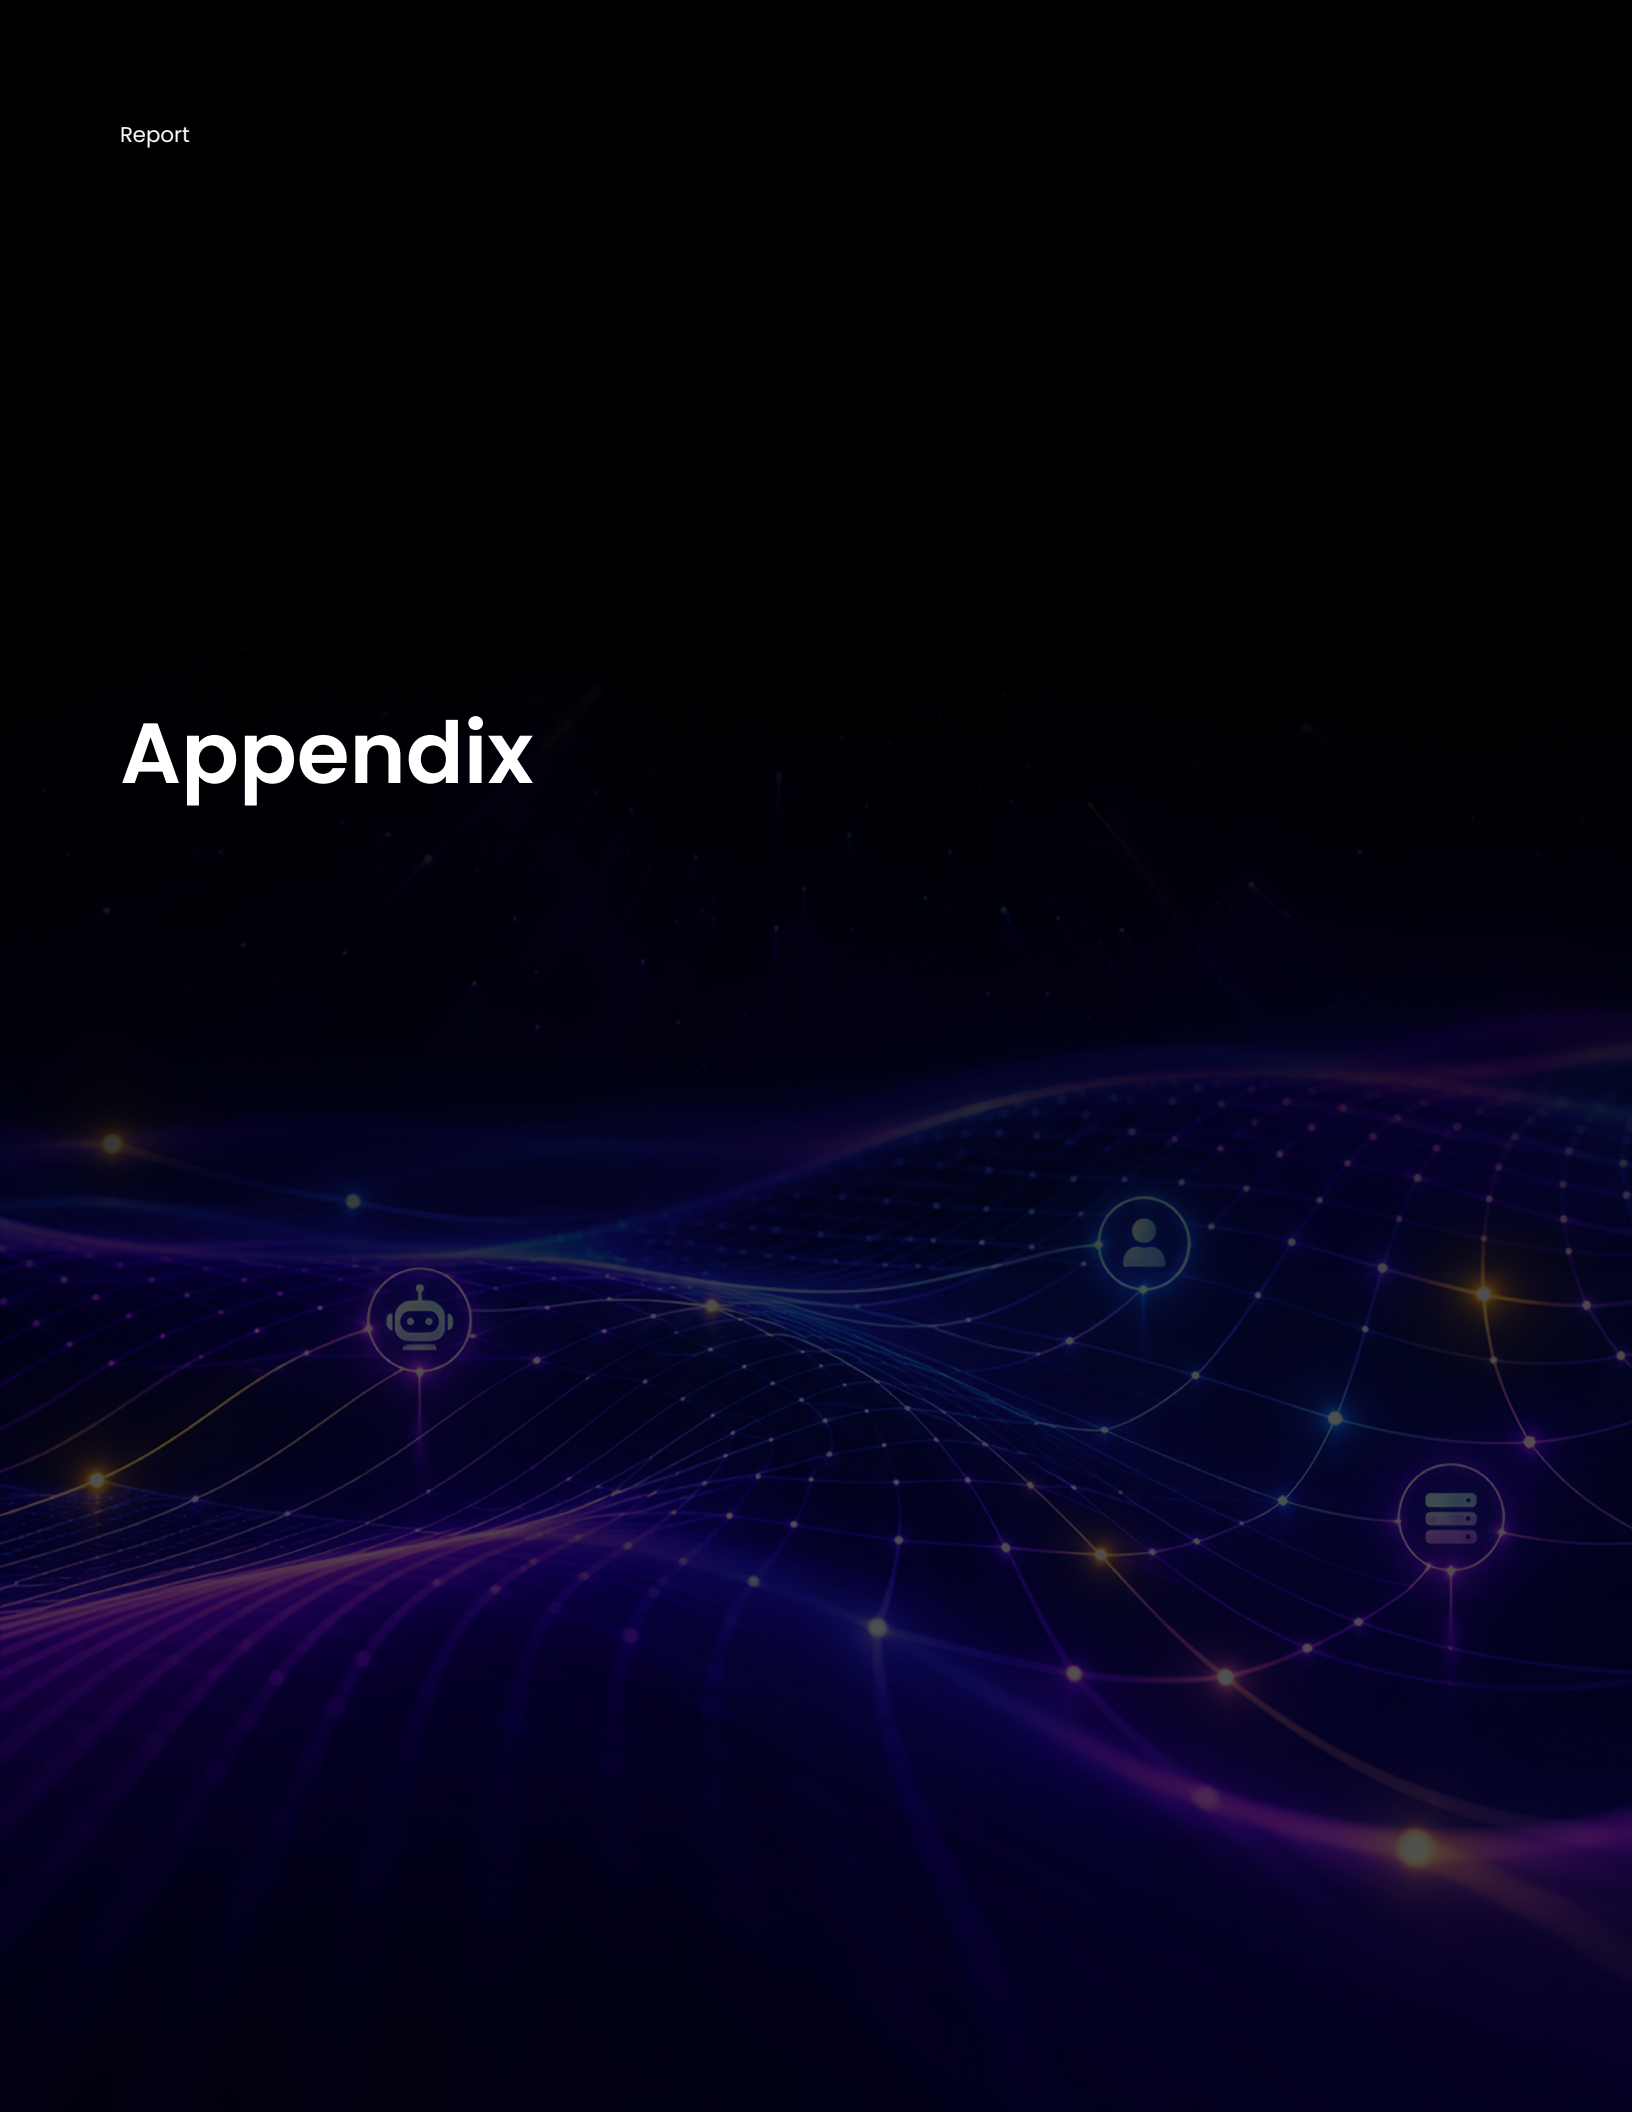
Achieving this three-stage mandate is the immediate goal. Looking further, however, a new paradigm emerges: A multi-layered model of identity autonomy that will define the next decade of security. This journey progresses from simple, human-led automation to a fully self-governing identity fabric where "guardian" AI agents monitor and manage peer agents without a human in the loop.

As this model matures, guardian agents extend past a single enterprise's walls. The same federated trust that lets agents collaborate across organizational boundaries also lets guardian agents coordinate with their counterparts at partner organizations, vouching for external agents before they're granted access to a shared workflow. And as confidence in these systems grows, guardian agents move from watching peer agents to acting on what they see — containing a compromised agent or revoking its credentials without waiting on a human in the loop.

This is the ultimate vision for identity security: A system so intelligent, adaptive, and seamless that it becomes the invisible, self-healing nervous system of the modern enterprise. The work to get there starts today: Use SailPoint's online assessment to see where your organization falls on the maturity curve, and get a prescriptive path for closing the gap.

Take the assessment →

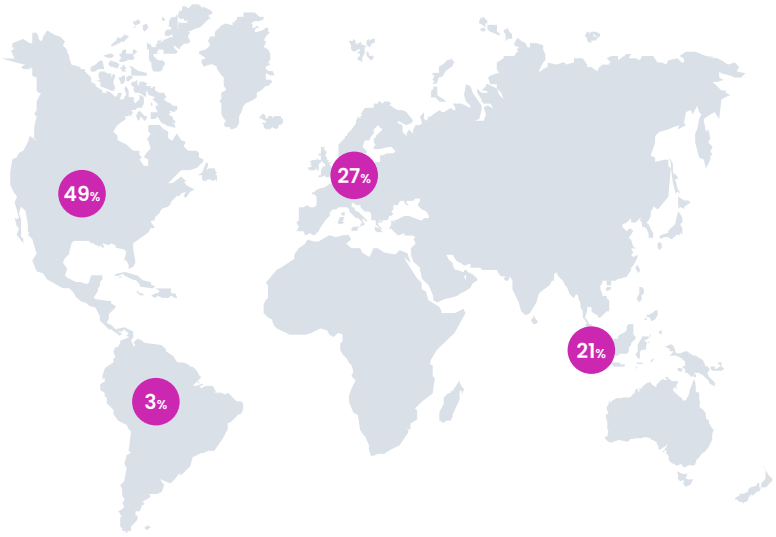
Appendix



Research approach

The 2026 to 2027 Horizons of Identity Security report is grounded in primary quantitative research fielded by SailPoint in 2026. The study reached 340 qualified identity security, IT, cybersecurity, and risk leaders across global enterprises.

Geo headquarter breakdown (n=340)



Enterprise scale

100%

Organizations with 1,000 or more employees.

44%

Large enterprises with over 10,000 employees

Executive decision-making authority

72%

Hold C-suite, VP, or Director titles, including CISO, CIO, CTO, and VP of IAM

81%

Hold direct budget and technology decision-making rights.

Respondents came from these industries

23%	21%	12%	12%	11%	5%	6%	10%
Technology	Financial services	Manufacturing	Retail & wholesale	Healthcare	Communications & media services	Education	Other verticals



About SailPoint

SailPoint (Nasdaq: SAIL) is defining the new era of adaptive identity security. In a world where non-human identities now significantly outnumber humans, our AI-powered platform unifies identity, security, and data intelligence to protect today's enterprise from advanced identity-based threats. We deliver the identity solution that spans both the breadth of identities and the depth of context needed to drive real-time access with confidence. Built on principles like zero-standing privilege and contextualized risk, our SailPoint platform transforms identity from a point of vulnerability into a powerful security advantage. Trusted by many of the world's leading organizations, SailPoint secures the enterprise with intelligent, autonomous identity security.

© 2026 SailPoint Technologies, Inc. All rights reserved. SailPoint, the SailPoint logo and all techniques are trademarks or registered trademarks of SailPoint Technologies, Inc. in the U.S. and/or other countries. All other products or services are trademarks of their respective companies.