SANS | Research Program

2025 SURVEY

2026 SANS State of Identity Threats & Defenses Survey Insights Event:

How Identity Became the New Security Perimeter—and What's Next

Written by **Rich Greene**

March 2026

P1 SECURITY

2026 SANS STATE OF IDENTITY THREATS & DEFENSES SURVEY INSIGHTS EVENT

Key Findings

The "Deployment vs. Resilience" Gap

High tool adoption has not correlated with low breach rates. While **68%** of organizations detect identity attacks within 24 hours, only **55%** contain them in that same window. This lag allows attackers to escalate privileges before the SOC intervenes.

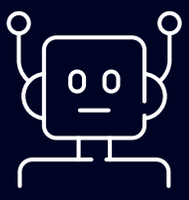

The Non-Human Identity Crisis

Non-human identities (NHIs) are the fastest-growing identity category, with **76%** of organizations reporting growth. Yet they are the least governed with **92%** of organizations failing to rotate the majority of their NHI credentials on a 90-day cycle, creating a landscape of "forever access."

Authentication Is Not Enough

Credential phishing accounts for only **35%** of attacks. The majority of incidents now involve techniques that bypass the login prompt entirely, such as session token hijacking (**23%**) and compromised browsers (**27%**).

The AI Shadow

73% of organizations are deploying agentic AI or automations that require credentials. This introduces a new class of identity that is not just automated but *autonomous*, entering production faster than governance frameworks can be built.

Business Impact

Identity failure is a brand crisis. **44%** of organizations cite reputation damage as a primary outcome of identity attacks, ranking it above many technical impacts.

Survey Author

Rich Greene
SANS Certified Instructor

CURRENTLY TEACHING

SEC301: Introduction to Cyber Security

VIEW PROFILE

Rich Greene has spent his career on the front lines of cyber defense, translating high-stakes operational experience into practical, accessible training for today's defenders. A SANS Instructor and author of *SEC301: Introduction to Cyber Security*, one of the organization's most widely attended foundational courses, he is known for turning complex security concepts into clear, engaging lessons that resonate with learners at every level. His approach is shaped by more than two decades in military, intelligence, and cybersecurity operations (including service in U.S. Army Special Forces and cyber roles supporting missions across 17 countries) where secure communications and rapid decision-making were mission-critical. Those experiences forged his belief that effective defense begins with understanding how adversaries think, move, and exploit weaknesses.

After transitioning from military service, Rich built a career focused on strengthening cyber readiness across organizations. Through his company, SITH2, LLC, he helps teams improve detection, response, and defensive decision-making, insights that directly inform the hands-on labs and real-world scenarios in SEC301. He holds more than a dozen GIAC certifications across incident handling, threat intelligence, forensics, and leadership, along with the CISSP and advanced degrees in cybersecurity from the SANS Technology Institute. Whether teaching in the classroom, speaking on podcasts and conference stages, or designing immersive learning experiences, Rich brings energy, clarity, and a mission-driven commitment to developing the next generation of defenders. His students consistently describe him as engaging, inspiring, and transformative, an instructor who makes cybersecurity not just understandable, but exciting and achievable.

Executive Summary

For a decade, the cybersecurity industry operated on a tacit assumption that modernizing authentication stacks with single sign-on (SSO) and multi-factor authentication (MFA) would suffocate the breach epidemic. The 2026 SANS Identity Threat Detection and Response (ITDR) Survey proves this assumption is dangerous. The defining finding is what we call the "deployment vs. resilience" gap: Although 68% of organizations detect identity attacks within 24 hours, only 55% contain them in that same window. Organizations have invested in the sensors to hear the alarm but have not yet built the operational muscle to put out the fire.

The survey reveals a defining paradox. Identity security tools are widely deployed, with 85% of organizations reporting active use of ITDR capabilities. Yet 55% of these same organizations admit to experiencing an identity-related compromise in the past 12 months. This gap signals that the industry has entered a new phase of conflict. The challenge is no longer about purchasing technology; it is about operationalizing it against a threat landscape that has changed shape.

The survey exposes a core paradox: Identity security tools are everywhere, but identity breaches are still common.

Although security teams successfully hardened human access, the attack surface shifted to the side windows: non-human identities (NHIs), post-authentication sessions, and autonomous AI agents. Critically, the majority of attack techniques identified in this survey—session token hijacking, compromised browsers, and OAuth abuse—occur after authentication succeeds. The threat landscape has moved beyond the login prompt. Organizations are fighting a modern war with outdated maps, heavily invested in protecting logins while attackers exploit what happens after the front door opens: ungoverned sessions, static service account credentials, and trusted integrations.

This analysis serves as a strategic roadmap. It moves beyond adoption statistics to analyze *why* identity defenses are failing in practice and provides actionable guidance for closing the gap between perceived readiness and actual risk.

Demographics

To understand the state of the industry, this analysis draws upon data collected from the 2026 SANS Identity Threat Detection and Response (ITDR) Survey. The respondent base represents a diverse global cohort of cybersecurity professionals.

Most respondents were based in the United States (85%) with participants also from Europe, Canada, Africa, Asia, and Australia. The top industries represented were the usual mix of respondents from technology, banking and finance, government, and telecommunications. There was a diverse representation of organization size with the largest showing from organizations ranging from 100-1,000 employees. The survey captures insights from across the organizational hierarchy. Respondents range from technical practitioners (security architects, SOC analysts) to strategic decision-makers (CISOs, Directors of Identity). This blend highlights the frequent misalignment between executive perception of risk and practitioner reality Figure 1 shows the survey demographics in detail.

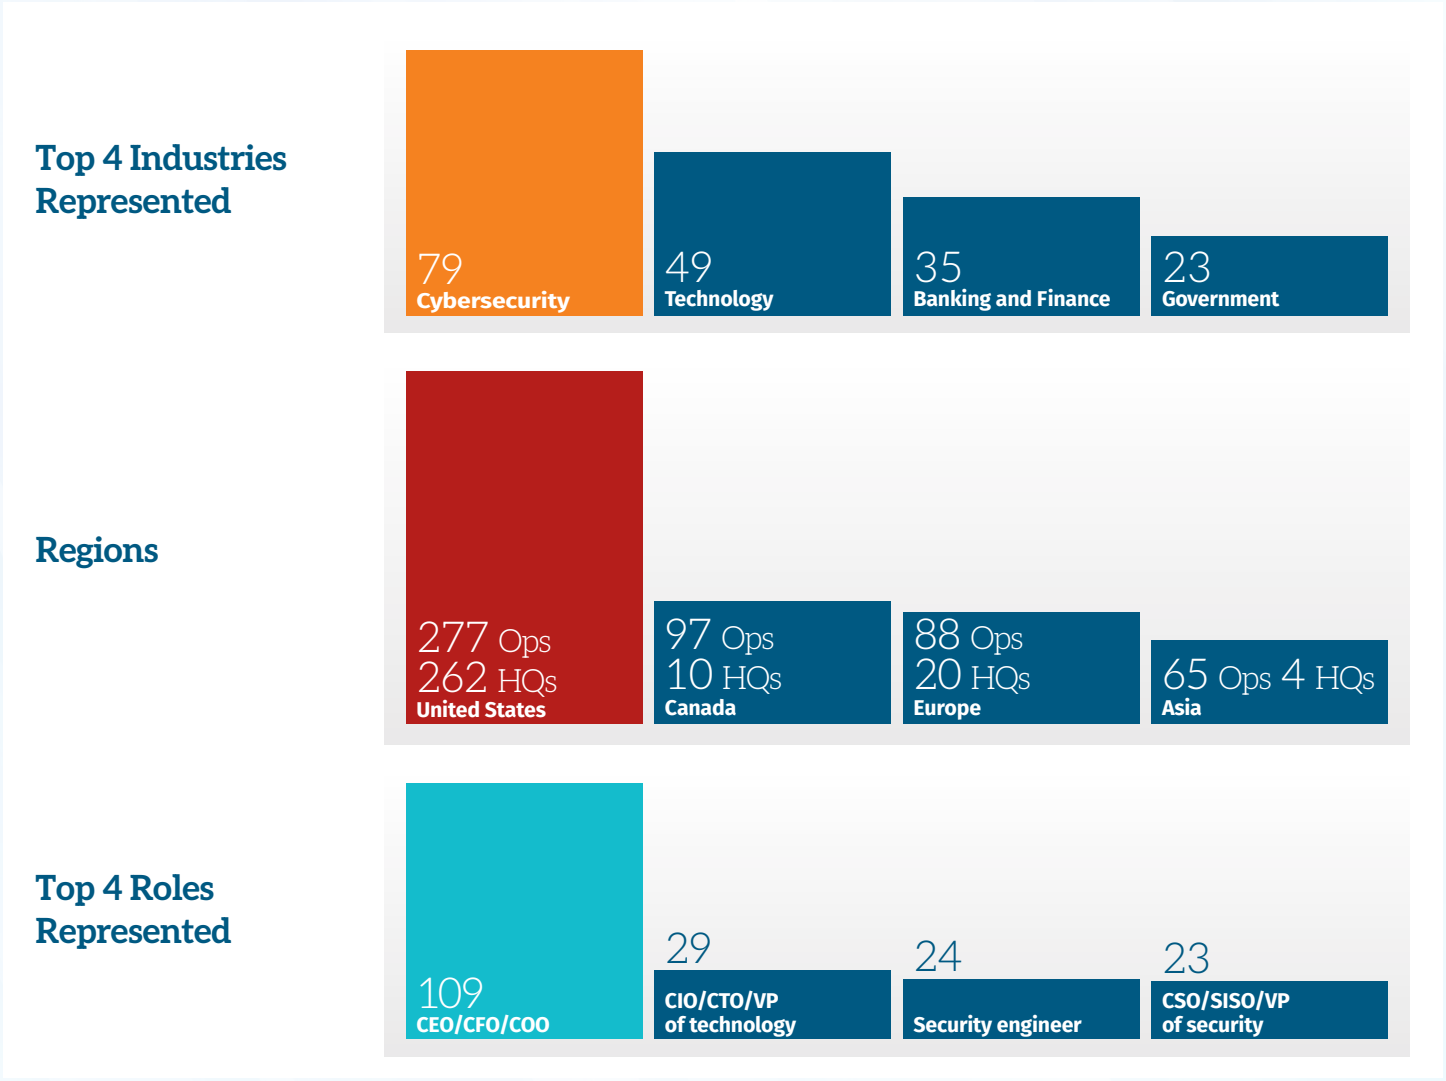

Figure 1. Demographics of Survey Respondents

Environment Scope

The data reflects the complexity of modern infrastructure. Respondents largely manage hybrid identity environments utilizing a complex mix of:

- **On-premises infrastructure—**Traditional Active Directory remains critical for 45% of respondents.
- **Cloud identity providers (IDPs)—**Platforms like Microsoft Entra ID and Okta serve as the primary control plane.
- **Multicloud infrastructure—**AWS, Azure, and GCP environments where workload identities proliferate.

It is important to note that the majority of respondents do not operate purely on-premises or purely in the cloud. The norm is hybrid with overlapping environments managed simultaneously, where a single identity transaction may traverse on-premises Active Directory, a cloud IDP, and a SaaS application in a single session. This hybrid reality is a defining characteristic of the modern identity attack surface and a recurring theme throughout this analysis.

Methodological Note

We pay special attention to "negative space" in the data. Where respondents indicated "Unknown," we interpret that lack of visibility as a critical finding. In identity security, what you do not know usually represents your highest risk.

The ITDR Effectiveness Paradox

The survey data presents a striking contradiction. With 85% of respondents reporting that their organizations currently deploy ITDR tools, one would think that widespread adoption should correlate with a decline in intrusions. Instead, identity-related compromises remain stubbornly prevalent.

Among the organizations surveyed, 55% experienced at least one identity-related cyberattack leading to unauthorized access within the past 12 months. Even more concerning, 21% reported multiple incidents (see Figure 2).

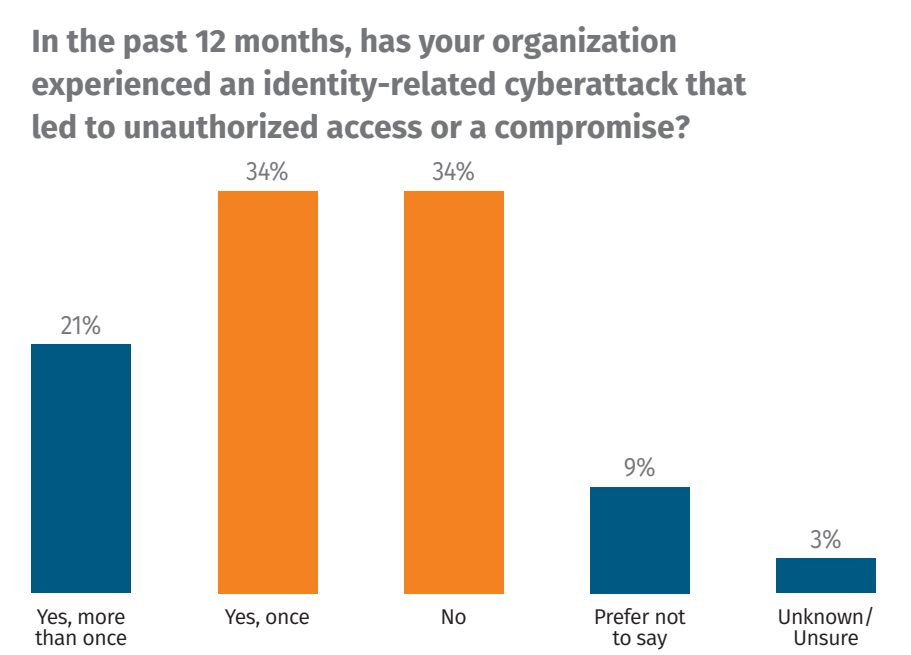

Figure 2. Identity-Related Cyberattacks Leading to Unauthorized Access or Compromise

The Operational Gap

This disconnect does not suggest that ITDR technology is flawed. Organizations are correct to invest here as attackers are clearly targeting identity. The problem lies in the distance between procuring a solution and operationalizing it. Many ITDR deployments exist in name rather than in practice. These may be pilot programs that never expanded, implementations limited to narrow use cases, or alert-generating systems that organizations lack the resources to act on.

We are witnessing a "containment lag." The survey indicates that while most organizations detect identity attacks within 24 hours, significantly fewer contain them in that same window. This gap allows attackers time to escalate privileges, move laterally from on-premise Active Directory to cloud infrastructure, or exfiltrate data before the SOC can intervene.

In short, organizations have built the sensors to hear the alarm, but they have not yet built the muscle memory to put out the fire. This paper examines that gap, exploring how deployment patterns and organizational constraints shape whether these tools protect organizations or merely provide a false sense of coverage.

What's Working: Detection and Investment Momentum—Despite the containment gap, the data reveals genuine forward motion. 41% of breached organizations detect identity attacks within four hours. 45% of organizations report ITDR at broad deployment or mature stages. And 79% plan to increase ITDR investment over the next 12 months, with 44% planning increases of 10–25%. The sensors are working; the operational muscle to act on them is what needs to catch up.

Key Findings and Thematic Analysis

The survey results coalesce around five major themes that define the state of identity security in 2026. These themes reveal an industry struggling to adapt legacy governance models to a hyper-connected reality.

Identity isn't just about people anymore. It's about machines, bots, and AI, and security teams are losing the race to govern them.

Theme 1: Identity Has Changed Shape (from Humans to Systems)

The vocabulary of identity security still carries the weight of an earlier era when "identity" meant a human employee. That mental model no longer matches reality.

The Explosion of Non-Human Identities

75% of respondents report growth in their organization's use of NHIs over the past year. Only 13% report no change. The identity perimeter has quietly doubled or tripled in size, but the new population consists of service accounts, API keys, and automation bots (see Figure 3).

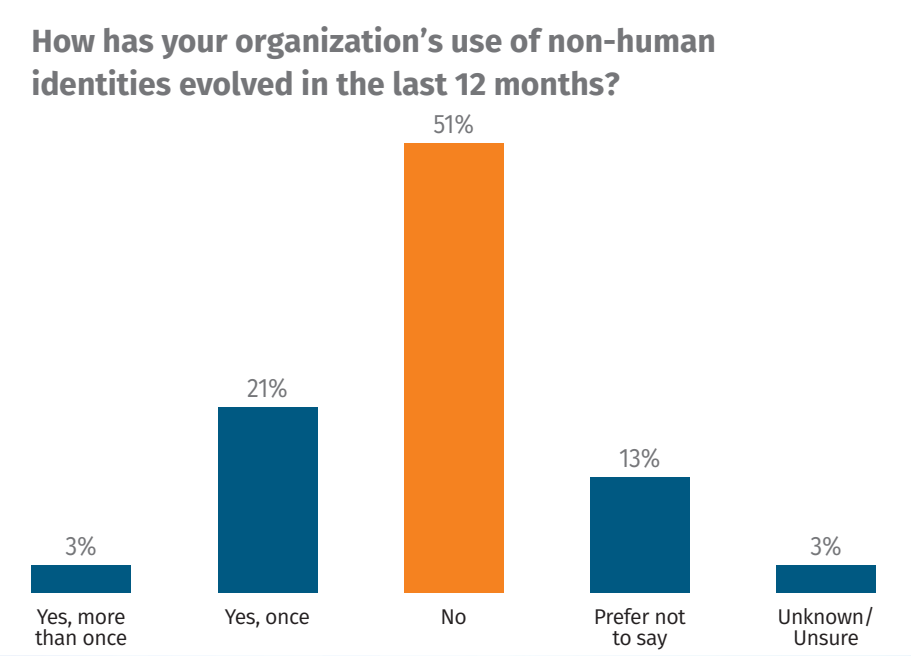

Figure 3. Growth in Non-Human Identities

These identities operate according to different rules. They authenticate continuously, do not use MFA, and often hold elevated privileges because they automate critical infrastructure tasks.

The Dispersion of Risk

Critical identities no longer reside in a single directory. When asked where their most critical identities live, respondents revealed a fractured landscape:

- 47% cite legacy/on-prem AD.
- 43% cite cloud infrastructure (AWS IAM, Azure RBAC).
- 38% cite DevOps/CI/CD pipelines.
- 35% cite third-party integrations (OAuth/OIDC).

Analysis: Security teams are trying to defend a fragmented empire. Attackers understand this fragmentation better than defenders do, targeting the gaps between these silos. The high ranking of DevOps environments (38%) is particularly alarming as these environments often prioritize speed over access control.

Theme 2: The Hygiene Crisis in NHIs

If NHIs are the fastest-growing attack surface, they are also the most neglected. The survey exposes a systemic failure in basic hygiene for machine credentials.

The Rotation Gap

Credential rotation is the most basic defense against long-term compromise. Yet the data shows that rotation is the exception:

- Only 8% of organizations rotate most (>75%) of their NHI credentials every 90 days.
- 15% admit they *don't know* their rotation rate.
- The majority (59%) rotate fewer than half of their NHI credentials quarterly.

The Implications of Static Secrets

This means the vast majority of service account keys and API tokens are static. They persist for months or years, often hardcoded in scripts or buried in config files.

The reason for this failure is rarely negligence; it is operational fear. Unlike a human user who notices a password reset, a service account simply breaks when its key is changed. This causes downtime and discourages operations teams from touching them. Organizations are effectively prioritizing operational uptime over security hygiene, leaving the keys to the kingdom exposed for years.

However, the rotation gap is a symptom of a deeper structural problem: The governance model for NHIs is inherited from human identity management. Periodic access reviews, manual rotation schedules, and ticket-based provisioning were designed for populations of hundreds or thousands of human users. They do not scale to NHI populations that number in the tens of thousands, authenticate continuously, and span cloud platforms, CI/CD pipelines, and SaaS integrations. Until organizations recognize that NHI governance requires a fundamentally different program model—one built for machine-scale life cycle management rather than human-scale periodic review—the rotation gap will persist regardless of policy mandates.

What's Working: NHI Controls Gaining Traction. **The picture is not entirely bleak. 34% of organizations have deployed secrets vaults to eliminate hardcoded keys. 33% report automated key and secret rotation in place. 33% enforce scoped, least-privilege roles for NHIs. And 31% have implemented CI/CD signing and provenance controls. Additionally, 37% of organizations report NHI Management programs at broad deployment or mature stages, with another 26% in early adoption. The foundation is being laid. The challenge is scaling it to match NHI population growth.**

Theme 3: Authentication Reality (The Hierarchy of Protection)

We have achieved consensus on MFA for employees, but universal enforcement remains a myth. The data reveals a distinct "Hierarchy of Protection" where security degrades rapidly as you move away from the core workforce. To understand where organizations are truly protected, and where they are exposed, it helps to break this hierarchy down:

- **Employees: High coverage—**Security teams have successfully mandated MFA here.
- **Contractors: A major gap—**73% of organizations fail to enforce MFA on all partner/contractor accounts. This is a critical vulnerability as contractors often have elevated access to internal systems (IT support, development) but operate outside the corporate device trust boundary.
- **Customers: 38% enforce MFA on less than half of customer accounts—**This reflects the tension between security and user experience (UX). Businesses fear friction will drive customers away, leaving data vulnerable to credential stuffing.
- **Non-human identities: The bottom of the hierarchy—**46% require MFA for fewer than half of NHIs (often 0).

The Phishing-Resistant Gap

While organizations know that legacy MFA (SMS, OTP) is bypassable, adoption of phishing-resistant tech (FIDO2/Passkeys) is critically low:

- Only 7% of organizations report broad passkey support across business-critical apps.
- 36% report no or almost no support.

We're investing heavily in securing the front door (employees) while leaving the side doors (contractors) and back windows (NHIs) exposed. Attackers simply take the easiest path in (see Figure 4).

Figure 4. Business-Critical Apps that Support Passkeys or FIDO2

Theme 4: Attackers Are Multidimensional

Security teams often optimize for "preventing the login" by blocking brute force attacks. The data proves attackers have moved on.

Although credential phishing (35%) is a top factor, it is not the *only* factor. Attackers use a diverse menu of techniques to bypass authentication controls:

- **Compromised browsers (27%)—**Attackers steal the session token directly from the browser cookies, bypassing MFA entirely.
- **MFA fatigue (26%)—**Attackers spam users with push notifications until they accept.
- **Token hijacking (23%)—**Replaying valid session tokens to gain access without credentials.

The Failure of Event-Based Detection

The rise of session hijacking and privilege escalation highlights the failure of "event-based" detection. A single log entry (an authentication event) might look innocent. A marketing user accessing a finance app is valid if they have permission. A session token being used from a new IP might be travel.

However, the *sequence* reveals the attack. Traditional SIEMs often struggle to correlate these disparate events across cloud, on-prem, and SaaS environments. Real detection requires monitoring post-authentication behavior, meaning what the identity does after it gets in.

This points to an architectural limitation in how identity security has been built. Traditional detection architectures concentrate visibility at the authentication boundary: login events, failed attempts, MFA challenges. The post-authentication layer—what identities do with their access once inside—remains largely unmonitored. SSO validates the login. MFA strengthens the login. SIEM logs the login. But the session that follows, where privilege escalation, lateral movement, and data exfiltration actually occur, often falls into a visibility gap. Closing this gap requires extending detection and governance into the post-authentication layer, not just strengthening the perimeter around it.

What's Working: Post-Authentication Visibility Is Emerging—Organizations are beginning to extend monitoring beyond the login prompt. 45% report continuous authentication or session risk scoring at broad deployment or mature stages. 41% enforce conditional access policies based on geo, device, or risk signals. 37% use risk-based step-up or continuous re-authentication during active sessions. And the top signals improving identity detections are endpoint/EDR telemetry (46%) and IdP sign-in and risk logs (34%)—both of which provide behavioral context beyond the initial authentication event.

Theme 5: The Next Crisis (Agentic AI)

The survey identifies a governance challenge that mirrors the NHI sprawl of the cloud era but introduces a new variable: autonomy.

74% of organizations are already using AI agents or automations that require credentials. These are not passive chatbots or deterministic scripts. They are autonomous actors executing transactions, modifying infrastructure, and accessing customer data based on nondeterministic decision-making. At their core, however, they are still identities—and they require the same governance fundamentals as any other identity: inventory, least privilege, time-bounded access, and auditability.

The Governance Lag

Despite this rapid adoption, governance is nonexistent:

- No single control (approvals, audit trails, sandboxing) is used by more than 40% of respondents.
- 5% of security leaders *don't know* if agentic AI is running in their environment.

This matters because traditional NHIs follow fixed logic; if you check the code, you know what they will do. AI agents interpret instructions. An ungoverned AI agent is essentially an overprivileged insider threat that operates at machine speed with the ability to hallucinate actions. Organizations are deploying "decision-making" capabilities without "decision-constraining" governance. The industry failed to govern NHIs early in the cloud era and is still paying the price. Agentic AI presents the same governance challenge on a compressed timeline, and the window to get ahead of it is narrowing.

What's Working: Early Governance of AI Agents—Unlike the NHI story, organizations appear to be building governance earlier in the adoption cycle for agentic AI (see Figure 5). 37% have implemented human-in-the-loop approvals for AI agent actions. 31% require mandatory audit trails and watermarking. 30% deploy runtime sandboxes or policy guardrails. And 28% use secrets brokers to avoid giving AI agents long-lived credentials. These are meaningful controls. The question is whether they will scale as agentic AI moves from pilot programs to production workloads.

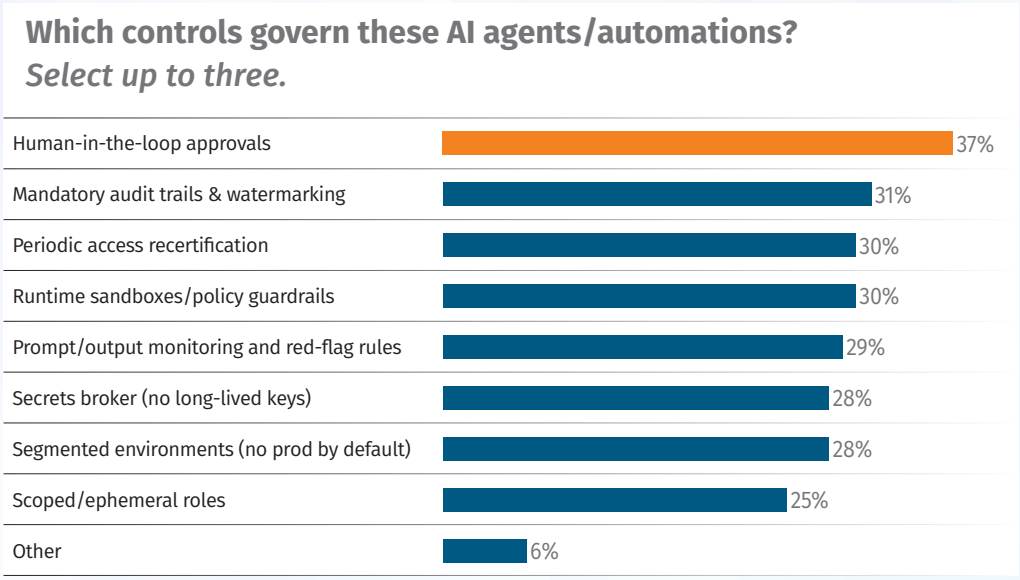

Figure 5. Controls Used to Govern AI Agents/Automations

Anatomy of a Modern Attack

To see why today's identity defenses are failing, it helps to follow how the survey findings play out across a modern attack chain. The era of "single point of failure" breaches is over, today's attacks unfold in phases, with each step exploiting a different identity weakness. The following is a representative, four-phase sequence that shows how separate control gaps compound into a full-scale compromise:

- **Phase 1: Initial access via the weakest link—**The attacker does not target the CEO's corporate laptop. Instead, they target a contractor (protected by SMS MFA, per the 66% gap) or a service account with hardcoded credentials found in a public repository (a symptom of the 92% rotation failure).
- **Phase 2: Evasion via session abuse—**Once inside, the attacker utilizes token hijacking (23%). They export the session cookies from the compromised contractor's browser. Now they *are* the contractor. They access the SaaS portal without triggering a new MFA prompt. The identity provider sees a valid token and assumes all is well.
- **Phase 3: Lateral movement via NHIs—**Inside a code repository, the attacker finds a hardcoded cloud API key (NHI) that has not been rotated in two years. This key has "Admin" privileges because the developer wanted to avoid permission errors. This is the pivot point where the attack moves from "User Identity" to "Workload Identity."
- **Phase 4: Impact via integration abuse—**Using the API key, the attacker creates a new OAuth application (16% of attacks involved OAuth abuse). They grant this malicious app "Read All Email" permissions. The attack is now persistent. Even if the contractor resets their password, the malicious app retains access.

Modern identity attacks don't break in, they chain together, exploiting contractors, sessions, NHIs, and integrations to turn small gaps into full-scale compromise.

The Failure of Traditional Detection

Throughout this chain, no "failed login" alerts were generated. The attacker used valid credentials and valid tokens. To a traditional SIEM relying on event-based logs, this looks like business as usual. This illustrates why 55% of organizations are breached despite having tools: The tools are looking for *events* (failed logins), but the threat is *behavior* (unusual access).

The failure here is not just one of detection methodology—event-based vs. behavioral—but of detection scope. Every phase of this attack chain succeeded using valid credentials and valid sessions. Controls that concentrate visibility at the authentication boundary cannot see the lateral movement, privilege escalation, and integration abuse that follow. This gap between authentication-layer security and post-authentication activity is where modern identity attacks live, and it is where the next generation of ITDR capabilities must operate.

Barriers to Maturity and Business Impact

Why does the industry struggle to close these gaps? The survey points to structural and financial barriers rather than technical impossibilities:

- **The budget barrier—**38% of respondents cite budget constraints as the top barrier to effective ITDR deployment. Identity security is often viewed as "plumbing" rather than a frontline defense. It struggles to compete for budget against more visible threats like ransomware, even though identity is the primary vector *for* ransomware.

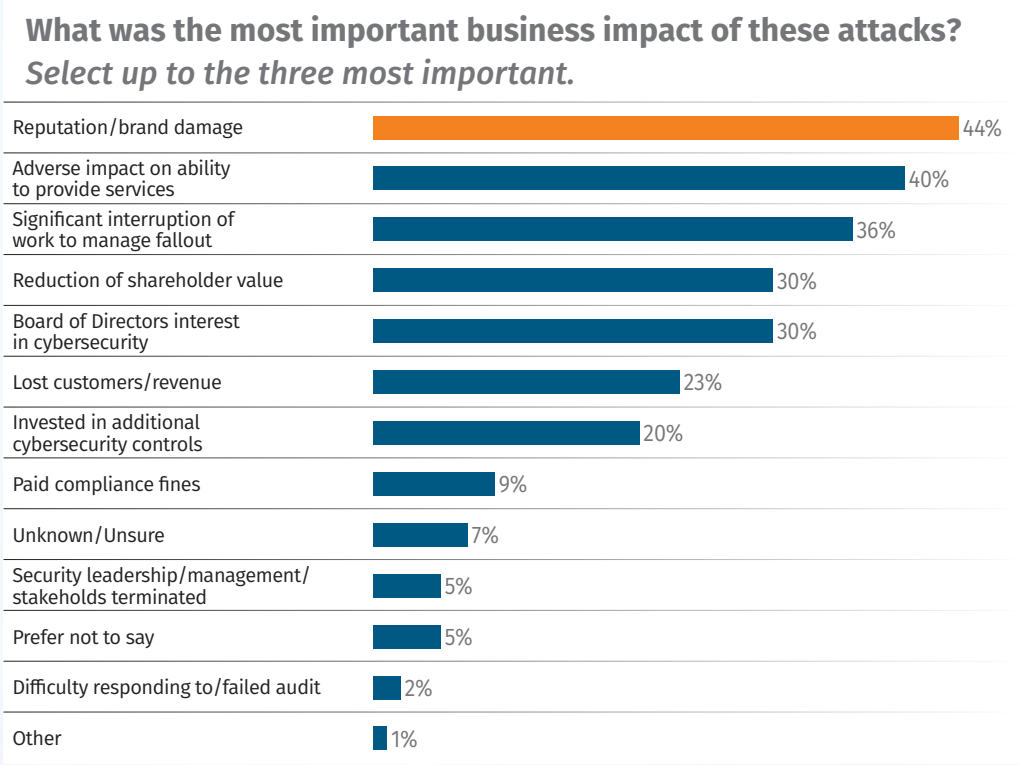

Figure 6. Business Impact of Identity-Related Cyberattacks

- **Legacy anchors—**Reliance on legacy protocols (cited in 19% of attacks) persists because organizations cannot easily retire critical business apps that rely on NTLM or LDAP. This forces security teams to maintain "downgrade" paths that attackers exploit. Organizations are effectively paying a "security tax" for every piece of legacy infrastructure they refuse to retire.
- **The business consequences—**The cost of this under-investment is high and visible. When asked about the primary impact of identity-related attacks, respondents cited (see Figure 6):
  - **Reputation or brand damage (44%)—**Identity breaches are visible public events that erode trust.
  - **Service disruption (40%)—**Identity is the control plane. When it goes down, business operations halt.
  - **Loss of shareholder value (30%)—**Markets punish identity failure swiftly.

This data provides a clear narrative for security leaders seeking budget. Identity security is not an IT ticket; it is a brand preservation strategy.

Strategic Implications for Organizations

The data dictates five immediate shifts for security leadership. The time for "identity awareness" is over and the time for "identity operations" has begun:

1. **Rebalance the portfolio from prevention to detection.** Prevention (MFA, SSO) is necessary but insufficient. With a 55% compromise rate, organizations must assume the perimeter will be breached. Investment must shift toward ITDR capabilities that can spot:
   - Lateral movement via service accounts
   - Abnormal usage of valid session tokens
   - Privilege escalation in SaaS platforms

Equally important, organizations should invest in proactive governance that enforces policy at the time of action—at the moment of an access request or privilege change—rather than relying solely on retroactive detection after a violation has occurred. Proactive governance reduces the threat surface that ITDR must monitor, making detection programs more focused and effective.

2. **Treat NHIs as privileged users.** Stop treating service accounts as "infrastructure configuration" and start treating them as VIP users.
   - **Immediate action—**Inventory all NHIs. If you can't see them, you can't rotate them. Use discovery tools to scan code repositories and cloud environments for static keys.
   - **Goal—**Move from static long-lived secrets to short-lived, federated credentials (e.g., OIDC) to kill the value of a stolen key. Short-lived credentials limit the attacker's window of exploitation. A stolen token that expires in 60 minutes cannot be used for persistent access or lateral movement days later, dramatically reducing the blast radius of any single compromise.
3. **Operationalize the SOC for identity.** Identity is now a Big Data problem. A SOC analyst looking at a firewall log cannot see an identity attack. Identity telemetry must be integrated into SOC workflows.
   - **The metric that matters: *Mean Time to Contain*—**If you detect in 24 hours but contain
     in 48, you have lost the battle.
   - **Automation—**Detection logic must be coupled with automated response actions. For example, if a high-fidelity alert triggers for "impossible travel," the system should automatically suspend the user or revoke their session token without

waiting for a human analyst.

4. **Bridge the cultural divide (IAM vs. SecOps).** A subtle but critical implication is the friction between IAM teams (who own the directory) and SecOps teams (who own the response). Successful ITDR requires these teams to fuse their operations. IAM must care about logs. SecOps must understand authentication flows.

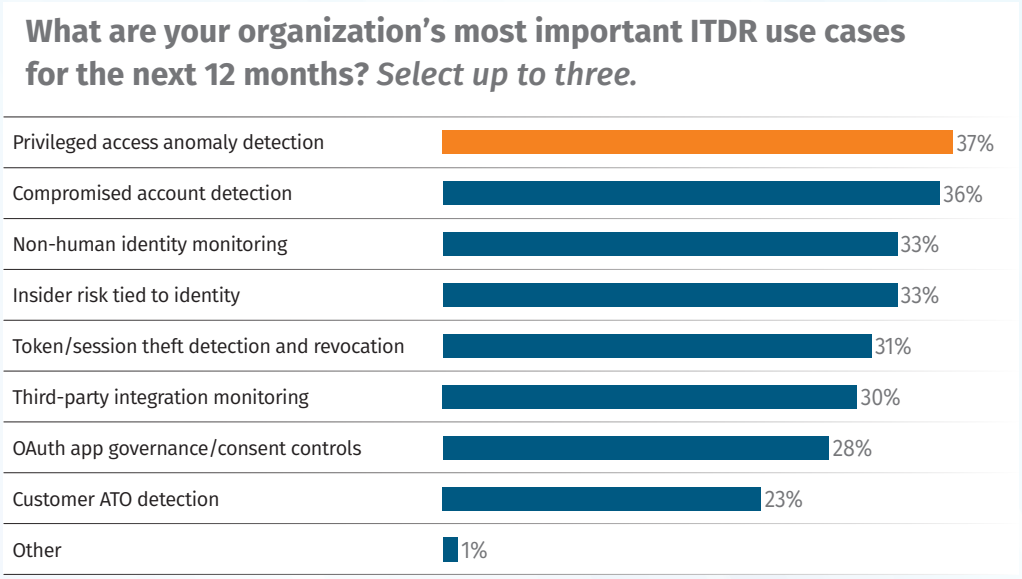

Figure 7. Expected ITDR Use Cases for the Next 12 Months

5. **Looking forward: The future of ITDR—**The next 12 to 36 months will see ITDR move from a "tool category" to an operational discipline (see Figure 7). The market is maturing, and expectations are rising.

   - **Identity security will become non-human by default—**As AI and automation scale, human logins will become the minority of authentication events. Defenses designed for human behavior (biometrics, 9-to-5 patterns) will fail. Behavioral baselining for machine identities will become the primary requirement.
   - **The rise of "immune system" response—**Manual response is too slow for machine-speed attacks. We will see a shift toward automated containment where identity systems self-heal by automatically rotating compromised keys or stepping up authentication requirements in real-time when risk is detected.
   - **Consolidation of visibility—**The fragmentation of identity (cloud vs. on-prem vs. SaaS) is unsustainable. Organizations will demand a "unified identity control plane" that provides a single view of risk across the entire ecosystem.

Conclusion

Most organizations have successfully deployed the *technologies* of identity security. They have bought the tools. But they have not yet mastered the *operations* of identity resilience. Attackers no longer need to break in. They simply log in. They use what already exists: legitimate credentials, valid sessions, and trusted integrations. To defend against this, organizations must stop looking at identity as a static list of users and start treating it as a dynamic, living attack surface.

The future belongs to organizations that can answer four questions:

1. Can we see every identity (human and machine)?
2. Can we detect when a valid identity behaves badly?
3. Can we see and govern what an identity does after it authenticates?
4. Can we stop it before it leaves the blast radius?

This report is not a verdict on failure; it is a call for recalibration.

The data shows we are not there yet. But it also shows us exactly where we need to go. For organizations looking to act on these findings immediately, we recommend five concrete starting points:

1. **Measure your containment gap.** Run a tabletop exercise or review your last identity incident and compare your Mean Time to Detect against your Mean Time to Contain. If there is a gap, identify whether the bottleneck is technical (lack of automated response), procedural (unclear escalation paths), or organizational (IAM and SecOps operating in separate workflows). This single metric will tell you more about your ITDR maturity than any deployment dashboard.
2. **Inventory your NHIs.** You cannot govern what you cannot see. Conduct a discovery scan across cloud environments, code repositories, CI/CD pipelines, and SaaS integrations to catalog every service account, API key, and automation credential. Assign an owner to each one. The 92% rotation failure starts with the fact that most organizations do not know what NHIs they have, let alone who is responsible for them.
3. **Extend MFA to contractors and third parties.** The data shows that 73% of organizations do not enforce MFA across all contractor and partner accounts. These identities often have elevated access but operate outside the corporate device trust boundary, making them the path of least resistance for attackers. Prioritize phishing-resistant methods (FIDO2/passkeys) where possible, as legacy MFA (SMS, OTP) remains bypassable.

4. **Add post-authentication monitoring to your detection stack.** Evaluate whether your current detection capabilities can see what happens after a successful login. Specifically, ensure you have visibility into session token usage patterns, privilege escalation within SaaS platforms, and anomalous access to resources that do not match an identity's historical baseline. If your SIEM only alerts on failed logins, you are monitoring the wrong layer.
5. **Establish governance for agentic AI before it scales.** If your organization is among the 73% deploying AI agents or automations with credentials, implement foundational controls now: Require human-in-the-loop approvals for high-risk actions, enforce least-privilege and time-bounded access for every AI agent identity, maintain audit trails for all autonomous actions, and use secrets brokers rather than long-lived credentials. The cost of building governance early is a fraction of the cost of retrofitting it after an incident.

Sponsor

SANS would like to thank this survey's sponsor:

PO SECURITY

SANS | Research Program

About the SANS Research Program

The SANS Research Program is a key initiative by the SANS Institute and a premier global provider of cybersecurity research and information. SANS Research Program is designed to provide cybersecurity practitioners and leaders with data-driven insights, thought leadership, and solutions that help them better understand and respond to evolving security challenges. All content is authored by SANS instructor experts from around the world who apply their years of experience from hands-on practitioner work in the field, advisory roles, and the classroom to provide education, guidance, and actionable insights that help make the cyber world a safer place.

To learn about sponsorship opportunities for research, content, and in-person or virtual events, email us at **Sponsorships@sans.org** or go to **www.sans.org/sponsorship**.