PERMISO

PERMISO STATE OF IDENTITY SECURITY REPORT 2026

FROM FALSE CONFIDENCE TO TRUE VISIBILITY

TABLE OF CONTENTS

EXECUTIVE SUMMARY 2

INFRASTRUCTURE & COMPLEXITY 4
Multi-Cloud Identity Complexity 4
Identity Provider Fragmentation 5
Human Identity Management Trends 6
Non-Human Identities Surge 7
Risk Perception vs. Reality 8

THE VISIBILITY CRISIS 10
Identity: The Dominant Attack Vector 10
The Visibility Illusion 11
Unified Tracking Remains Elusive 12
Proactive Detection Fails 14
The Audit Trail Gap 16
False Confidence in NHI Inventory 17

NON-HUMAN IDENTITY MANAGEMENT 19
Discovery Methods Split Industry 19
Non-Human Identity Maturity Levels 20
The Credential Graveyard 22

DETECTION & RESPONSE 23
Detection Improves, Response Lags 23
The Blast Radius Bottleneck 24
Visibility Gaps Create Alert Fatigue 26

THE AI IDENTITY CHALLENGE 27
AI Transforms Identity Creation 27
AI Agents Access Sensitive Data 28
AI Identity Growth Expectations 30
AI Identity Crisis 30

OPERATIONAL COSTS & FRAGMENTATION 32
Tool Sprawl Fragments Visibility 32
Manual Correlation Costs 35

BUSINESS IMPACT & MARKET RESPONSE 37
Business Impact of Limited Visibility 37
What Security Teams Actually Want 39
Investment Surge in 2026 41

CONCLUSION 43

EXECUTIVE SUMMARY

The 2025 State of Identity Security Report reveals a watershed moment in cloud security. Drawing from 512 organizations worldwide, this study exposes a sobering reality: as identity infrastructure grows exponentially more complex, the gap between what organizations believe they can see and what they actually control has never been wider.

KEY FINDINGS INCLUDE

- **The Complexity Explosion:** Organizations now juggle an average of 2-3 cloud service providers alongside 2-3 identity providers, creating a fragmented authentication landscape where visibility becomes nearly impossible. Most organizations manage between 1,000 and 5,000 human identities, while non-human identities in the same range have surged to 44% of all organizations.

- **Identity Attacks Dominate:** For the first time, we quantified what security teams have long suspected: 77% of organizations report that between 26% and 75% of all security incidents are identity-related. This isn't a theoretical threat. This is the primary attack vector of 2025.

- **The Visibility Crisis:** While 46% of organizations claim comprehensive visibility into all identities, the data tells a darker story. Only 43% can detect risks proactively before incidents occur, just 29% can determine blast radius within minutes when compromise happens, and visibility gaps trigger security alerts frequently or occasionally in 82% of organizations.

- **The AI Identity Surge:** Organizations expect AI-generated identities to increase by 1-50% in the next 12 months (62% of respondents), yet 82% already have AI agents accessing production data, with most reporting between 1% and 50% of their sensitive data exposed to AI systems.

- **The AI Identity Surge:** Organizations expect AI-generated identities to increase by 1-50% in the next 12 months (62% of respondents), yet 82% already have AI agents accessing production data, with most reporting between 1% and 50% of their sensitive data exposed to AI systems.

- **The Tool Sprawl Tax:** 71% of teams use 3-10 separate tools to achieve identity visibility, burning 10-40 hours per week manually correlating identity data from different sources (60% of organizations). This isn't just inefficiency. It's a critical security gap masquerading as a productivity problem.

- **The ROI of Visibility:** 71% of organizations believe that 26-75% of security incidents could have been prevented with comprehensive identity visibility, with security breaches cited as the primary business impact of limited visibility by 44% of respondents.

- **Investment Signals Change:** Despite these challenges, 89% of organizations plan to increase their identity security investment in 2026, with 38% planning significant increases of over 30%. The market has recognized the problem. Now comes the hard work of solving it.

- **The message is clear:** identity security in 2025 is defined not by what organizations think they control, but by the exponential growth of what they cannot see. The organizations that survive the next wave of attacks will be those who close this visibility gap before attackers exploit it.

INFRASTRUCTURE & COMPLEXITY

MULTI-CLOUD IDENTITY COMPLEXITY

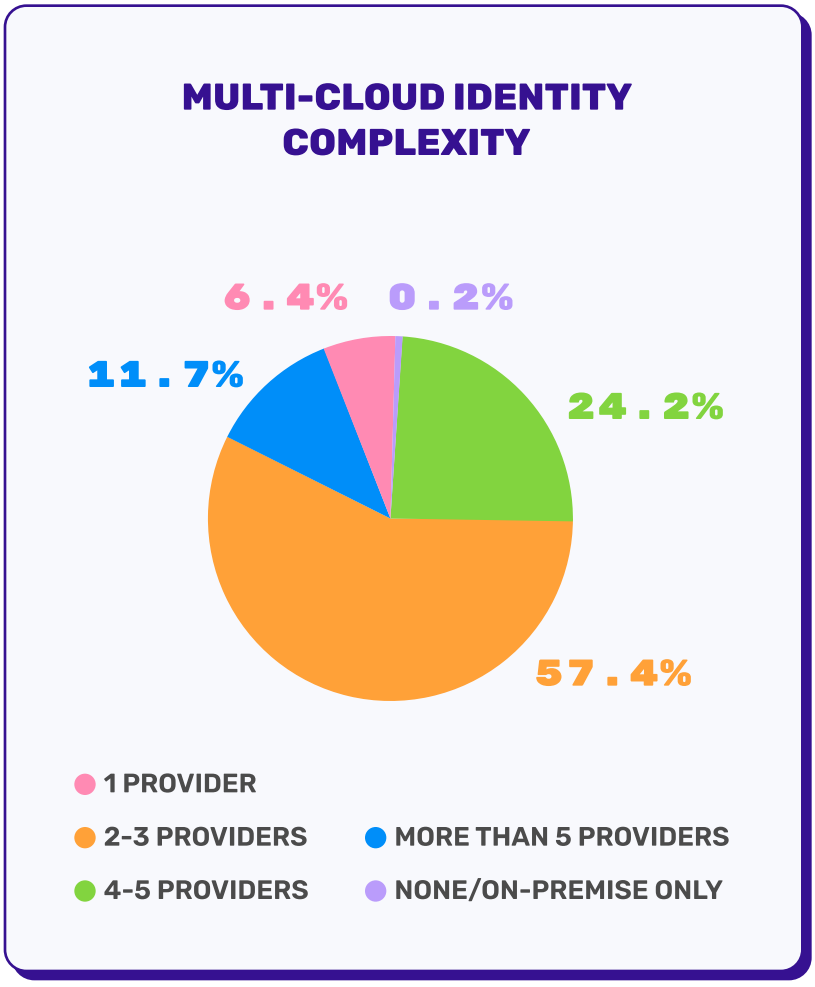

Question asked, "How many cloud service providers does your organization use?"

The multi-cloud reality has solidified. In 2025, 57% of organizations operate across 2-3 cloud providers, while 24% manage 4-5 providers, and 12% wrangle more than five. Only 6% maintain a single-provider strategy.

This distribution closely mirrors 2024 findings, where organizations reported using an average of 2.5 cloud service providers. The consistency suggests multi-cloud has moved from strategy to standard operating procedure, with AWS continuing its market dominance at 25%, followed by Azure at 22% and GCP at 7%.

The overwhelming dominance of the 2-3 provider bracket represents deliberate architectural decisions: a primary cloud provider, a secondary for redundancy or specific workloads, and occasionally a third for specialized services. The quarter of organizations operating 4-5 providers likely represent larger enterprises with complex M&A histories or global operations requiring regional providers.

IDENTITY PROVIDER FRAGMENTATION

The Identity Security Implication

Every additional cloud provider multiplies identity complexity exponentially. Each brings its own IAM framework, identity primitives, authentication flows, and security model. Layer in 2-3 identity providers, and you create a fragmented landscape where comprehensive visibility becomes nearly impossible. The question isn't whether multi-cloud is here to stay (it is). The question is whether organizations can build identity security architectures that span these environments without creating blind spots attackers can exploit.

For the first time in our survey, we quantified identity provider fragmentation. The results: 57% use 2-3 identity providers, 22% rely on a single provider, 16% manage 4-5 providers, and 6% juggle more than five.

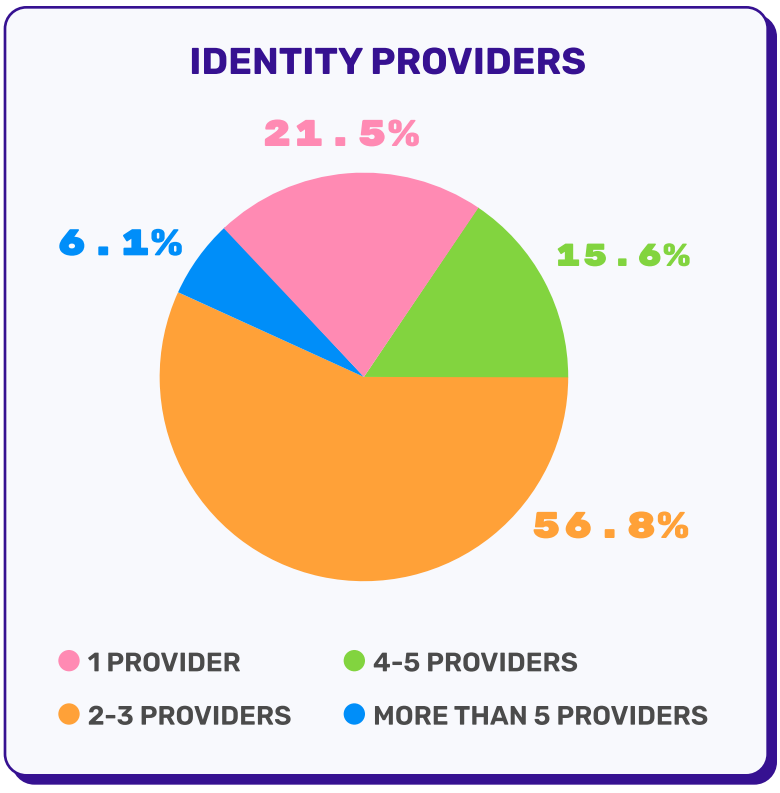

Question asked, “How many identity providers (IdPs) like Okta, Entra ID, Ping Identity, or others does your organization use across all environments?

More than half of all organizations operate across 2-3 identity providers: think Okta for SaaS, Entra ID for Microsoft workloads, Ping Identity for legacy systems. Unlike cloud infrastructure where workloads can be isolated, identities must flow across every system, creating a complex web of federation relationships, trust boundaries, and authentication handoffs. Each handoff is an opportunity for misconfiguration. Each trust relationship is a potential pivot point for attackers.

HUMAN IDENTITY MANAGEMENT TRENDS

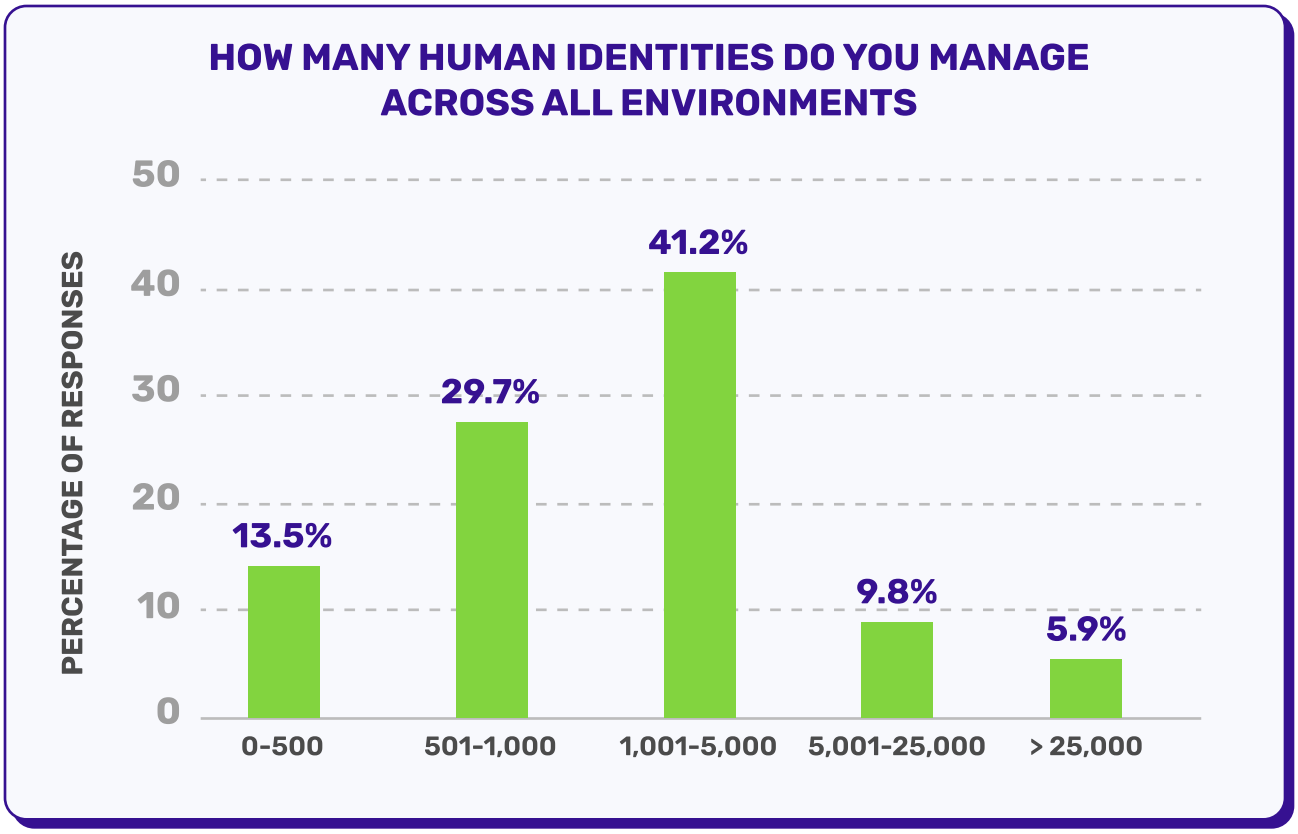

Question asked, "How many human identities do you manage across all environments?

Organizations continue clustering in the mid-range: 41% manage 1,001-5,000 identities, while 30% manage 501-1,000. This represents a 7-point decline from 2024's 48%, suggesting organizations are either consolidating identities through aggressive deprovisioning of dormant accounts, implementing more disciplined identity lifecycle management, or simply growing more slowly than their non-human identity populations.

The clustering in the 500-5,000 range (71% of all organizations) reveals the universal challenge: organizations are large enough to need sophisticated identity management but not so large they can afford dedicated identity security teams. This is the danger zone where complexity exceeds capability, where visibility gaps emerge, and where attackers find their opportunities.

NON-HUMAN IDENTITIES SURGE

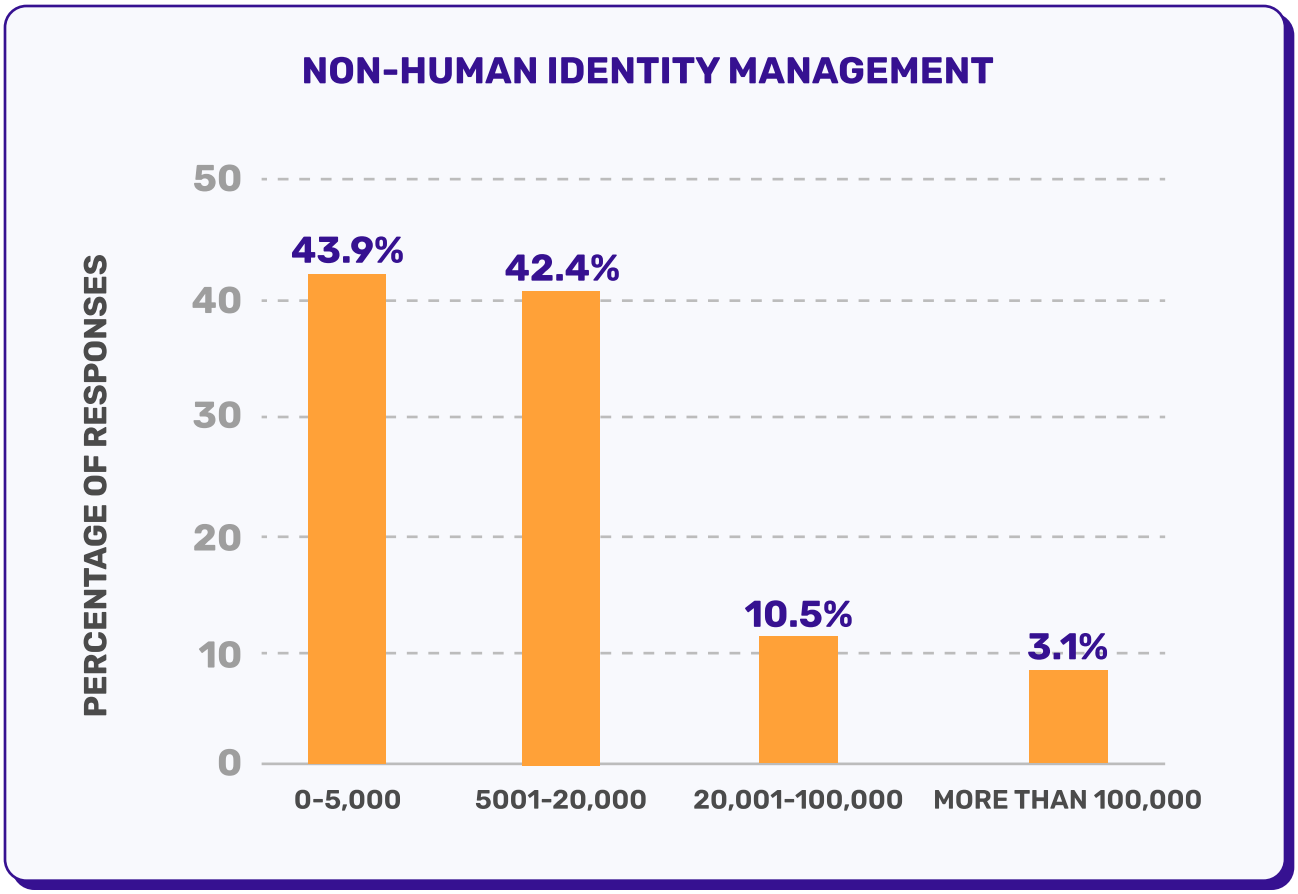

Question asked, "How many Non-human identities do you manage across all environments?"

If human identities have found equilibrium, non-human identities are experiencing explosive growth. While 41% of organizations manage 1,001-5,000 human identities, 44% manage up to 5,000 non-human identities and 42% manage between 5,001 and 20,000 non-human identities (a bracket that barely exists for humans).

In 2024, we documented 42% managing 1,000-5,000 non-human identities, calling them the "silent workhorses" of cloud environments. The 2025 data confirms this trend has intensified, with the concentration shifting toward even higher volumes.

The Great Inversion: When Machines Outnumber Humans

For a substantial portion of organizations, non-human identities outnumber human identities by factors of 3x, 5x, or even 10x. The majority of identities in your environment are no longer people. They're service accounts, API keys, access tokens, certificates, and increasingly, AI agents.

The 14% managing more than 20,000 non-human identities operate at bleeding-edge scale. Consider an organization with 100,000 non-human identities: if each has an average 90-day lifespan before rotation, you're creating, managing, and retiring over 1,000 identities per day. No human team can track this manually. No traditional IAM system was built for this scale.

Most organizations can barely track their human identities. When you add 5,000 to 20,000 non-human identities, each with different lifecycles and access patterns, comprehensive visibility becomes a pipe dream.

RISK PERCEPTION VS. REALITY

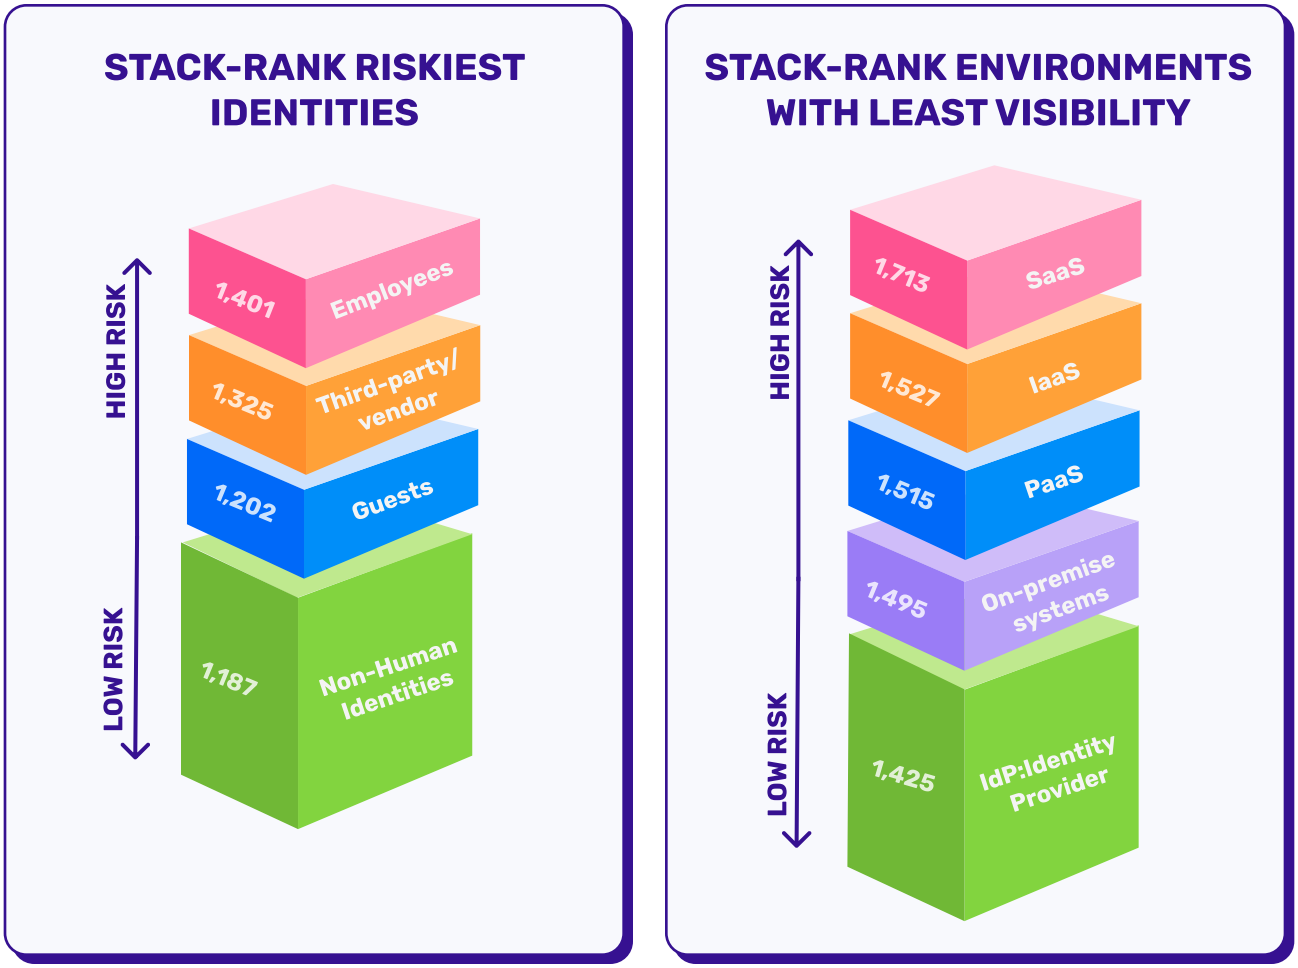

Two questions reveal a fascinating tension: what organizations fear most versus where they have the least visibility to detect those threats.

The Risk Hierarchy

1. Employees (most risky)
2. Third-party/Vendor
3. Guests
4. Non-Human Identities (least risky perception)

The Visibility Crisis

1. SaaS (worst visibility)
2. IaaS
3. PaaS
4. On-premise systems
5. IdP: Identity Provider (best visibility)

Employees remain the top perceived risk, consistent with 2024, but the gap between employees and third-party/vendors has narrowed considerably. Organizations are waking up to the reality that vendor access represents an enormous attack surface. Every major breach story (SolarWinds, Okta, MOVEit) has involved compromised vendor access. The visibility hierarchy also remains unchanged from 2024, with SaaS continuing to have the worst visibility despite its growing dominance in enterprise IT.

Despite managing thousands of non-human identities, organizations rank them as the least risky identity type. This is the risk perception gap in action. Non-human identities don't click phishing links, but they do get hardcoded in GitHub repos and left with overly permissive access for years.

THE VISIBILITY PARADOX

SaaS environments have the worst visibility, followed by IaaS and PaaS. Organizations have moved their most critical applications to SaaS while simultaneously losing visibility into who has access and what they're doing. If employees are your biggest risk and SaaS is your biggest visibility gap, then the intersection represents your highest-risk, lowest-visibility attack surface. That's precisely where attackers operate.

THE VISIBILITY CRISIS

IDENTITY: THE DOMINANT ATTACK VECTOR

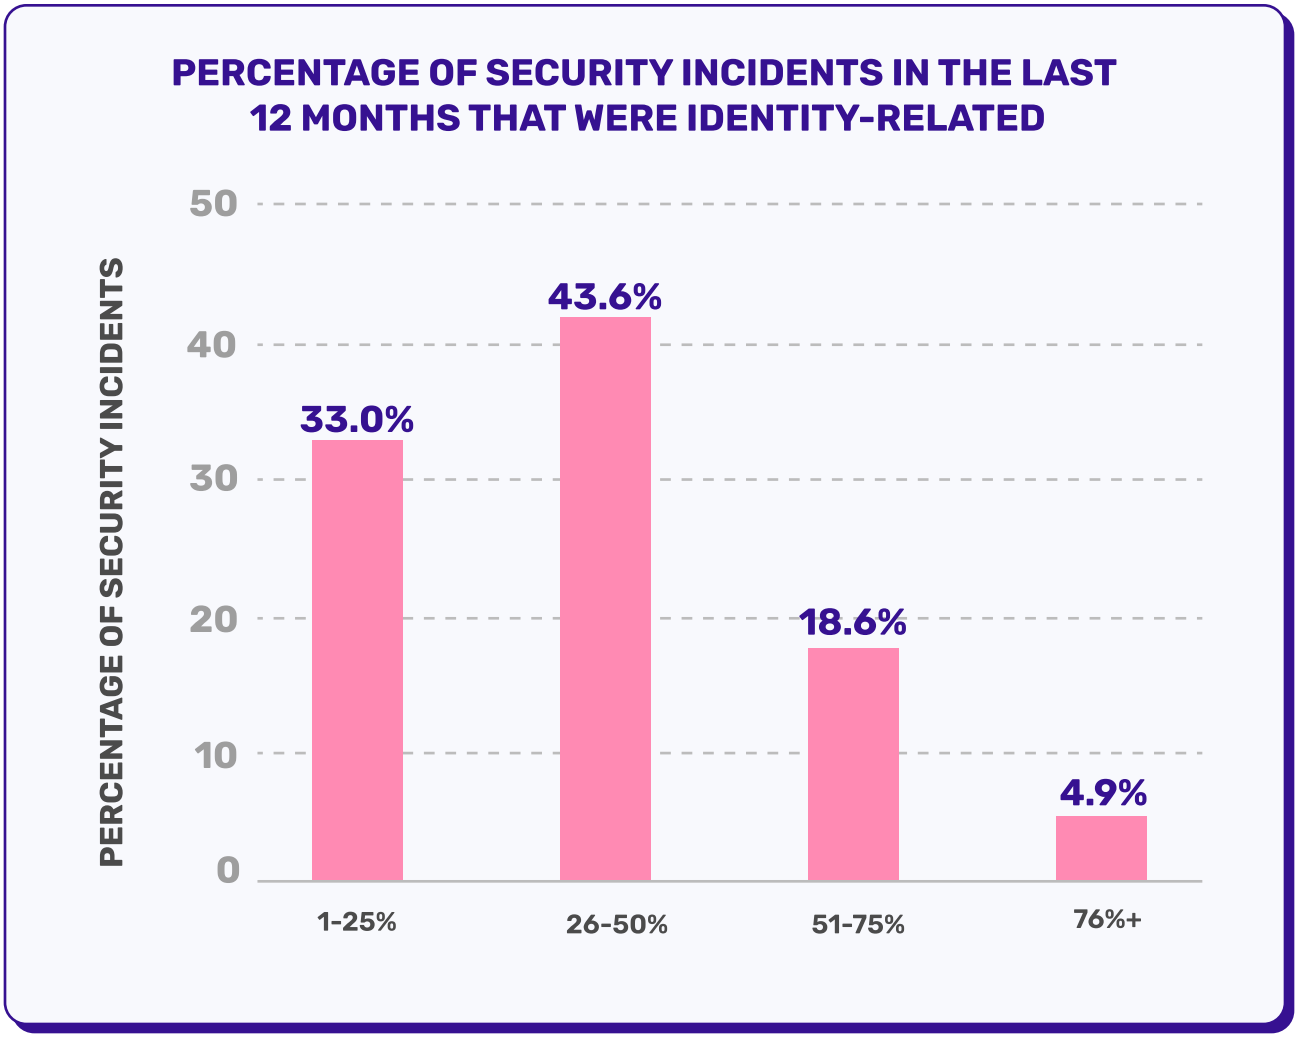

Question asked, "What percentage of security incidents your organization experienced in the past 12 months were identity-related?"

For the first time, we quantified what security teams have long suspected: identity compromise is the dominant attack vector. 77% of organizations report that between 26% and 75% of all security incidents involve identity compromise, with 44% reporting identity attacks constitute 26-50% of their incident volume.

When identity-related incidents represent a quarter to half of your security events, you're no longer primarily fighting network attackers. You're fighting attackers who bypass your perimeter entirely by logging in with valid credentials. Organizations spending more on network security than identity security in 2025 are optimizing for yesterday's threat landscape.

THE VISIBILITY ILLUSION

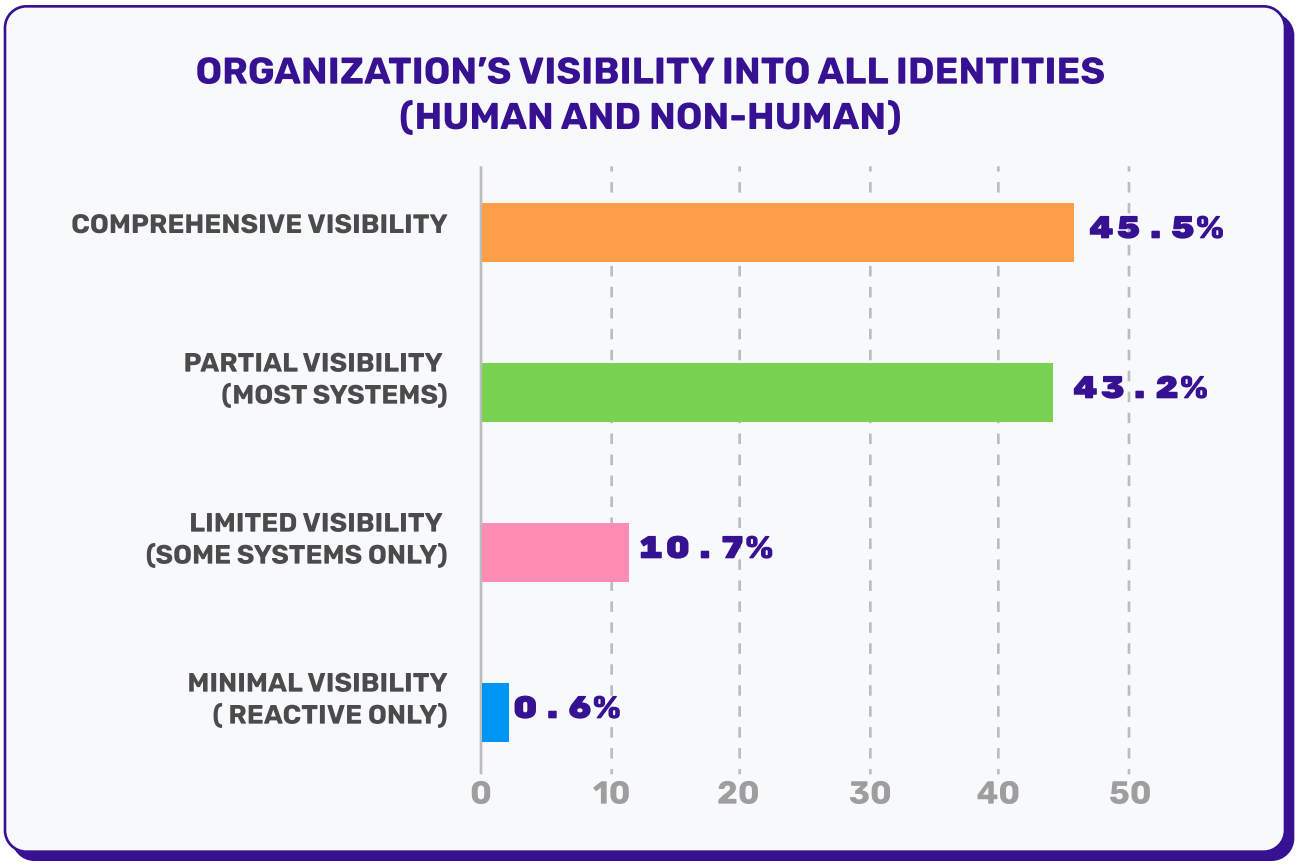

Question asked, "Which best describes your organization's visibility into all identities (human and non-human)?"

Ask organizations if they have comprehensive visibility into all identities, and 46% say yes. This confidence seems reasonable until you compare it to 2024, when 93% claimed comprehensive inventory. That's a 47-percentage-point confidence drop (the most dramatic shift in our entire survey).

But this isn't a collapse in capability. It's a correction in understanding. The 2024 question asked about maintaining an inventory (a static list you generate and file away). The 2025 question asks about visibility (dynamic, real-time tracking and continuous monitoring of what identities actually do). The 47-point plunge exposes the gap between knowing identities exist and truly seeing how they behave.

This is Level 1 of what we call **The Visibility Illusion Cascade,** where each deeper question about visibility reveals progressively lower capability.

UNIFIED TRACKING REMAINS ELUSIVE

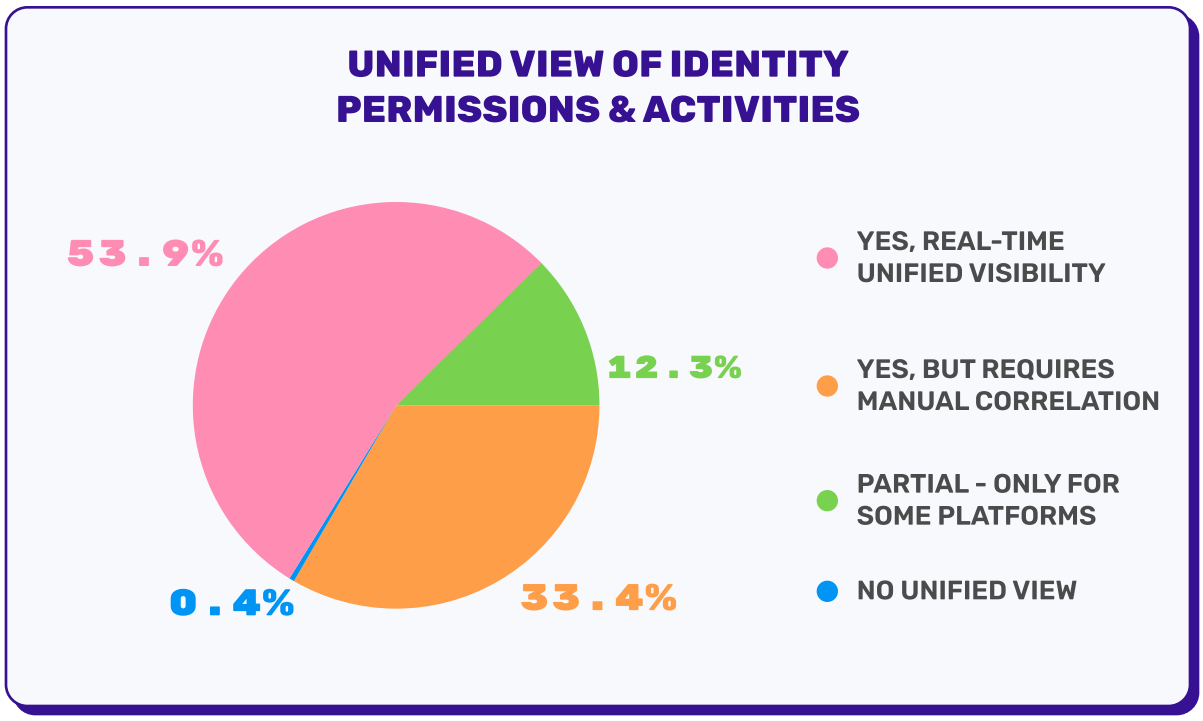

Question asked, "Can you track identity permissions and activities across all platforms in a unified view?"

While 54% claim real-time unified visibility, another 33% admit they have tracking capability but require manual correlation. This is where we introduce **The Manual Correlation Tax**, the hidden cost organizations pay when they have data everywhere but insight nowhere.

In 2024, 93% claimed they could monitor services and resources accessed by identities in real-time. The 2025 question probes deeper (not just monitoring, but unified tracking across all platforms). The result: only 54% can do this without manual effort, revealing that most "real-time monitoring" was actually fragmented across multiple tools.

These organizations have SIEM logs, IdP dashboards, cloud IAM reports, and CSPM findings. They have all the data. What they don't have is the ability to automatically connect the dots. When an incident happens, they spend hours manually reconstructing attack paths, pulling logs from different tools, and piecing together what happened.

By the time they've manually correlated the data, the attacker has moved three steps ahead. Having the capability to eventually figure out what happened is not the same as having real-time visibility to stop attacks in progress.

The Visibility Illusion Cascade accelerates: from 46% claiming comprehensive visibility to 54% claiming unified tracking, but a third are paying The Manual Correlation Tax to maintain even that illusion.

PROACTIVE DETECTION FAILS

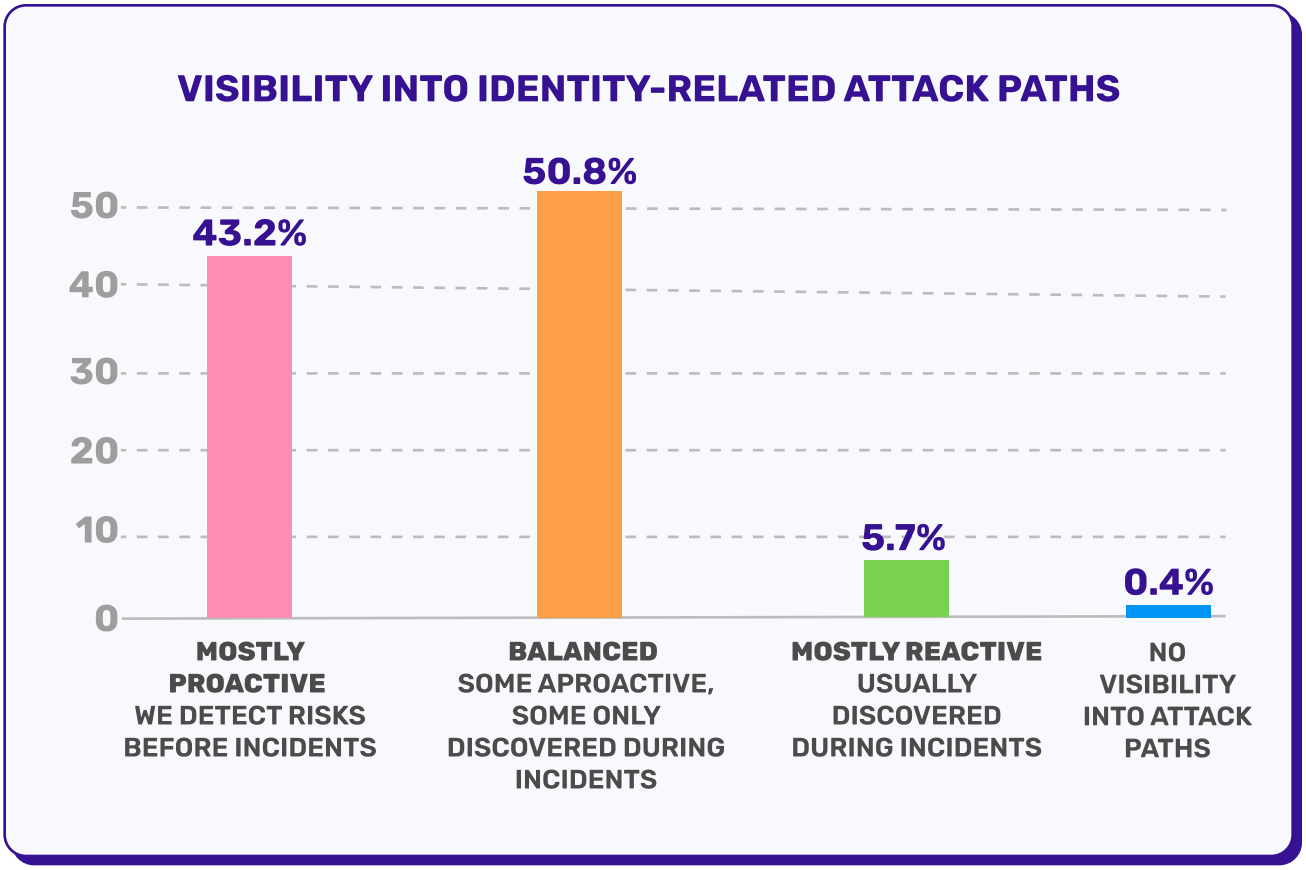

Question asked, "When it comes to attack paths involving identities, how would you describe your visilibily?"

Only 43% of organizations can detect identity-based risks before incidents occur. This is the gold standard: knowing not just what identities are doing, but what they could do if compromised. Notice how this capability drops from the 54% claiming unified visibility (that 11-point drop represents organizations that can see activity but not risk).

The 51% describing themselves as "balanced" (some proactive, some discovered during incidents) deserve scrutiny. In practice, this means they catch obvious risks proactively while discovering complex attack paths only when attackers exploit them. Being "balanced" sounds reasonable until you realize attackers don't use obvious paths.

This is The Proactive Security Myth: the belief that deploying security tools makes you proactive, when most tools only detect incidents after they've begun. Real proactive security requires complete visibility into all identity permissions, continuous analysis of how they could be chained together, and automated blast radius modeling.

The Visibility Illusion Cascade revealed:

46% claim comprehensive visibility (down from 93% in 2024)

54% claim unified tracking, but 33% require manual correlation

Only 43% can proactively detect risks before incidents

The Visibility Cascade Score (VCS)

To measure how organizational confidence erodes when tested against actual capabilities, we introduce the Visibility Cascade Score:

VCS = (Comprehensive Visibility % + Unified Tracking % + Proactive Detection %) ÷ 3

2025 Score: (46% + 54% + 43%) ÷ 3 = 47.7%

- **What This Reveals:** The VCS exposes a critical gap: while organizations average 48% effective visibility across the cascade, this masks dramatic capability erosion at each level. The 11-point drop from Stage 2 to Stage 3 is particularly revealing. Organizations that can unify their data still struggle to analyze it proactively. This is the difference between having visibility and using it effectively.

- **Industry Benchmark:** A VCS below 50% indicates compromised visibility at multiple levels. Organizations above 60% demonstrate mature capabilities across detection, tracking, and analysis. The 2025 industry average of 47.7% suggests most organizations are barely maintaining functional visibility.

- **Year-over-Year Comparison:** In 2024, if we apply the same cascade effect observed in 2025 (an average 8-point drop per stage), the theoretical score would have been approximately 77%. The 29-point decline represents not just improved measurement, but genuine recognition that static inventory is insufficient for dynamic security.

THE AUDIT TRAIL GAP

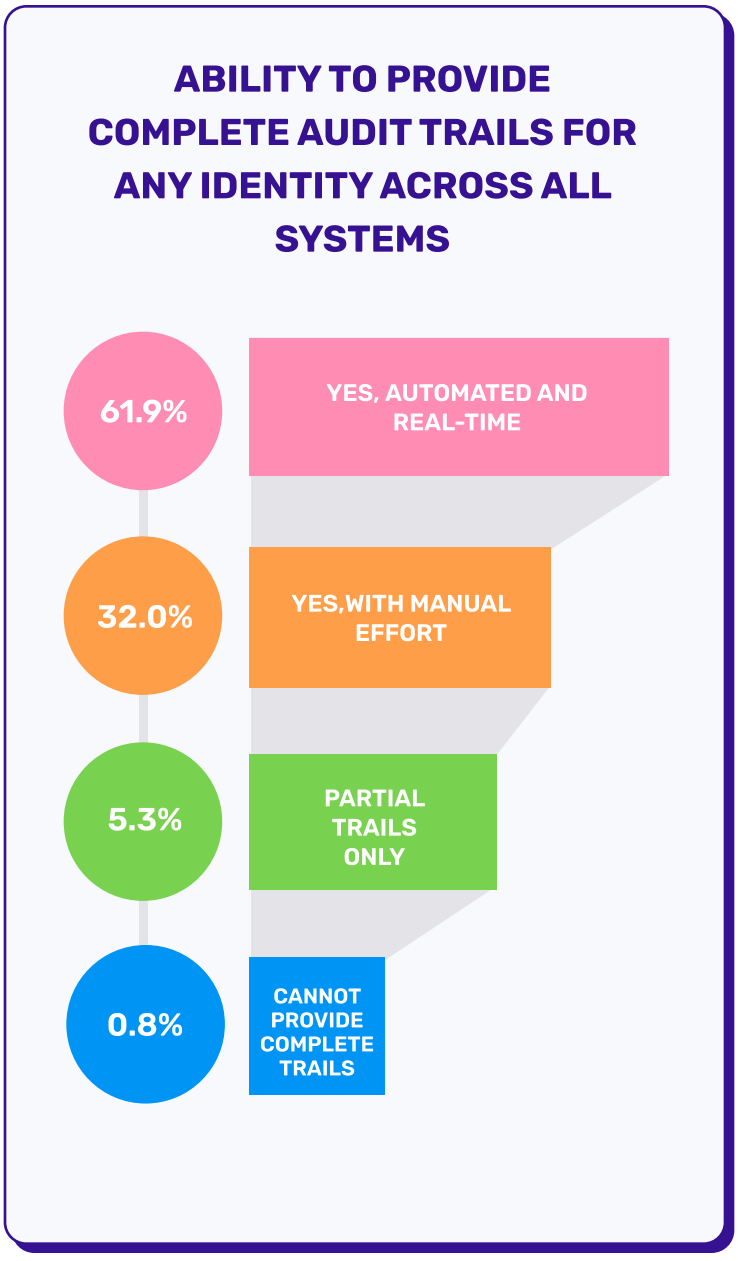

Question asked, "Can you provide complete audit trails for any identity's activities across all systems?"

When incidents happen, 62% claim they can provide automated, real-time audit trails. Another 32% can do it with manual effort.
Here's where The Manual Correlation Tax strikes at its most painful moment: during active incident response.

Organizations paying this tax spend 4-8 hours reconstructing what a single compromised identity did, pulling logs from multiple systems, correlating timestamps, and piecing together attack paths. By the time they present findings to leadership, the attacker has established persistence or moved laterally.

The 62% with "automated" audit trails may be overstating capability. Having automated collection doesn't mean automated analysis, instant attack path visualization, or quick blast radius determination.

FALSE CONFIDENCE IN NHI INVENTORY

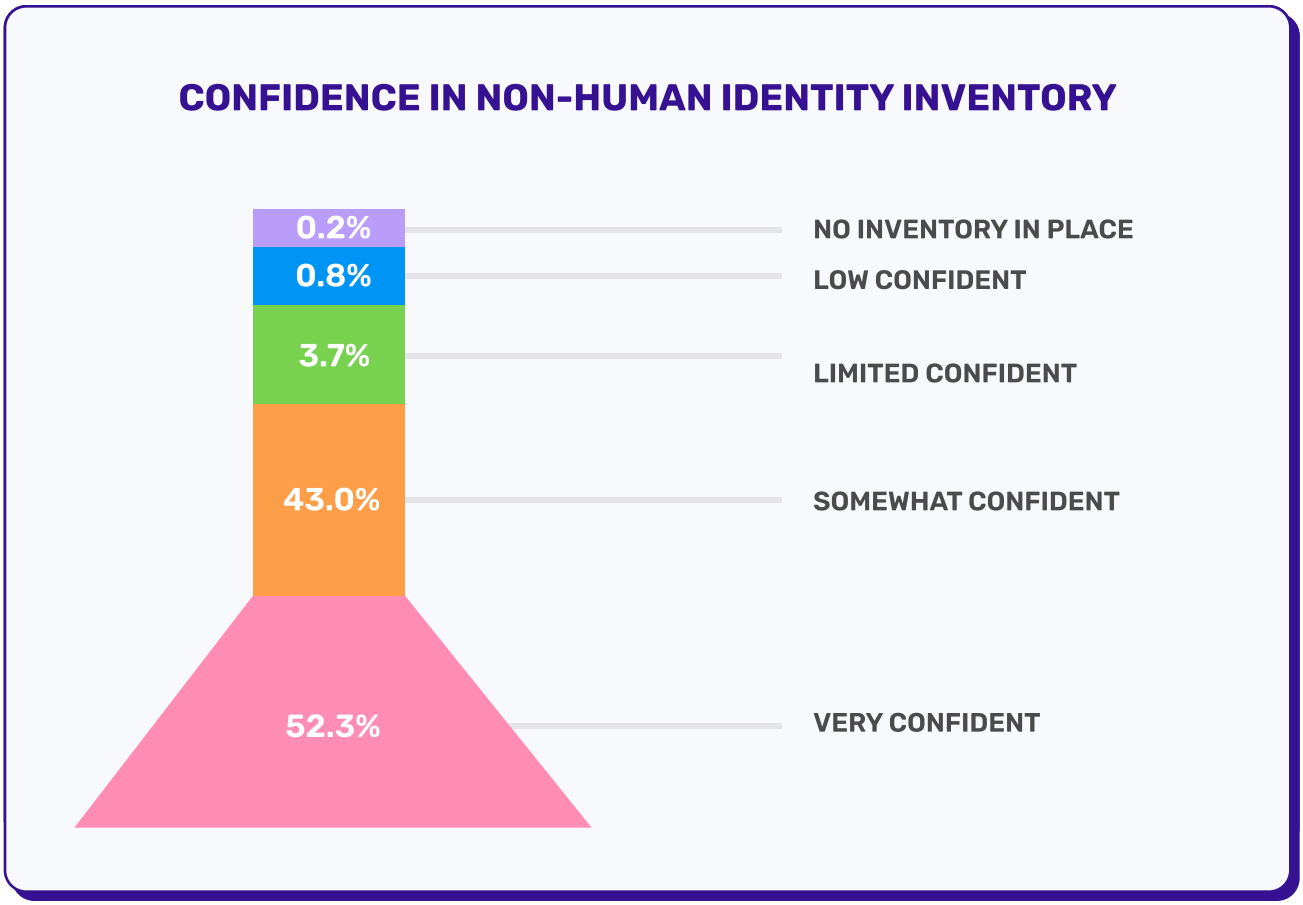

Question asked, "How confident are you in your inventory of ALL non-human identities with access to critical systems?"

After watching organizations struggle with visibility, tracking, and detection, we arrive at the most surprising finding: 95% express confidence (52% very confident, 43% somewhat confident) in their inventory of all non-human identities with access to critical systems.

This is dramatically higher than any other metric. Only 46% claimed comprehensive visibility into all identities, only 43% can proactively detect risks, yet 95% are confident about non-human identity inventory.

This is The Confidence-Reality Inversion: organizations express highest confidence in the area of greatest complexity. Organizations manage anywhere from 5,000 to 20,000 non-human identities across 2-3 cloud providers and 2-3 identity providers, with SaaS environments (the worst visibility) proliferating with service accounts for integrations and APIs.

What organizations call "inventory" is often a best-effort list that's immediately outdated. Developers create new service accounts daily, CI/CD pipelines generate credentials dynamically, and shadow IT deploys applications with their own credentials.

Non-human identities are skeleton keys to modern infrastructure: long-lived credentials, broad permissions, no MFA, minimal monitoring. When organizations express 95% confidence in tracking these while managing thousands across fragmented environments, they're demonstrating the most dangerous form of security posture: **false confidence.**

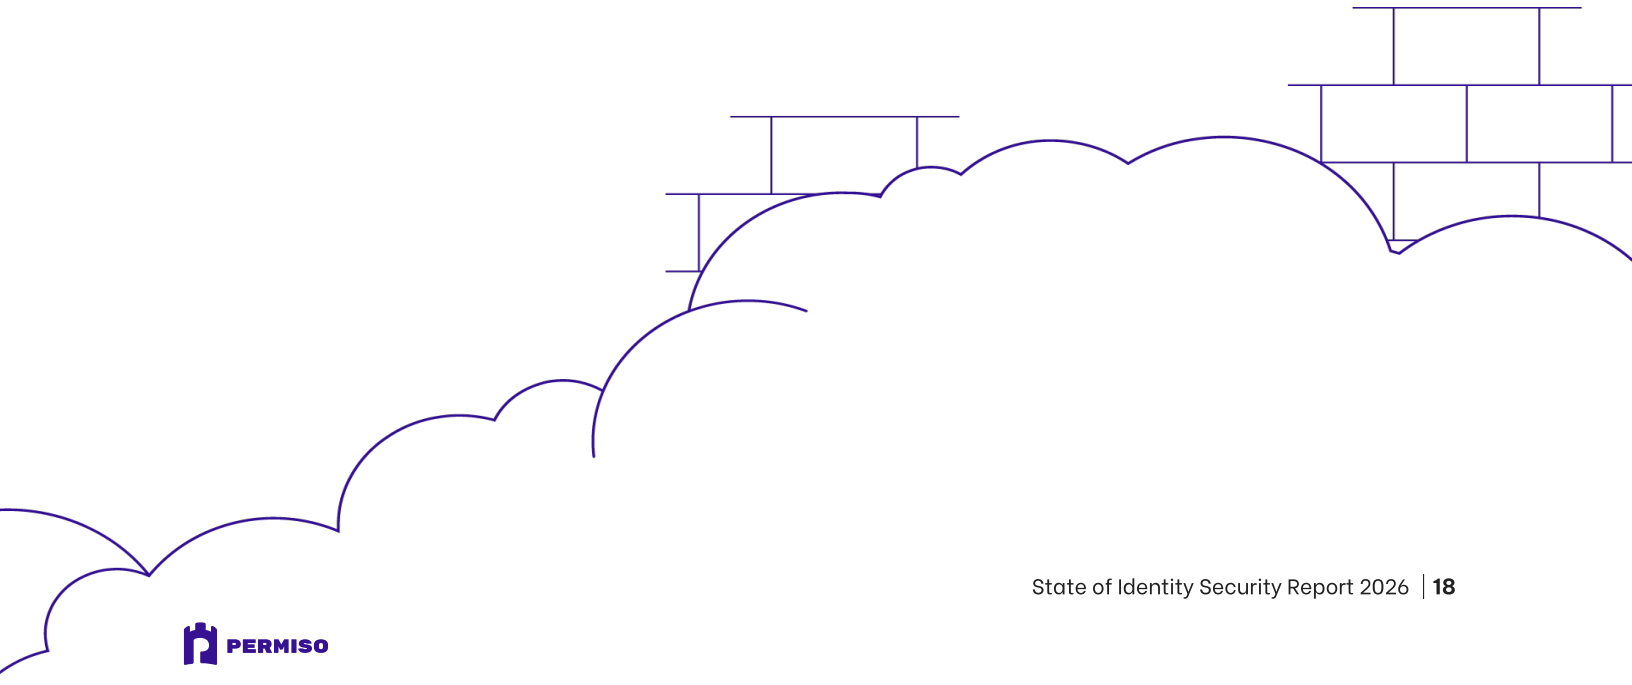

NON-HUMAN IDENTITY MANAGEMENT

DISCOVERY METHODS SPLIT INDUSTRY

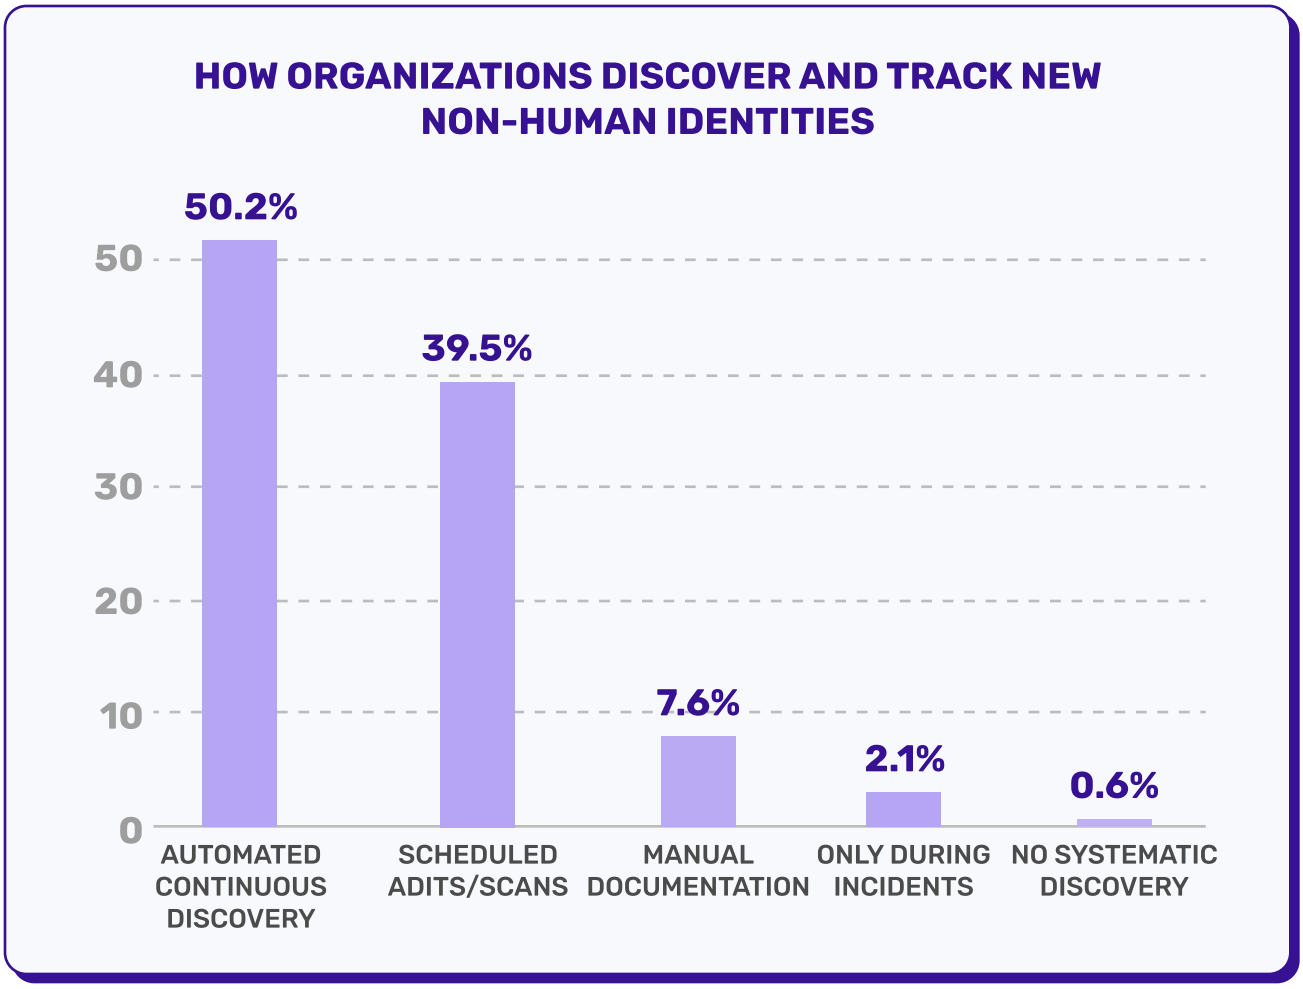

Question asked, "How do you discover and track new non-human identities?"

Despite 95% confidence in non-human identity inventory, the methods organizations use reveal a troubling divide. Half the industry (50%) has automated continuous discovery, while the other half relies on scheduled audits (40%), manual documentation (8%), or discovers identities only during incidents (2%).

This 50-50 split matters because non-human identities don't wait for quarterly audits. Developers spin up service accounts daily, CI/CD pipelines generate credentials automatically, and containers create identities on the fly. If you're discovering these through scheduled scans, you're always operating with an outdated inventory.

NON-HUMAN IDENTITY MATURITY LEVELS

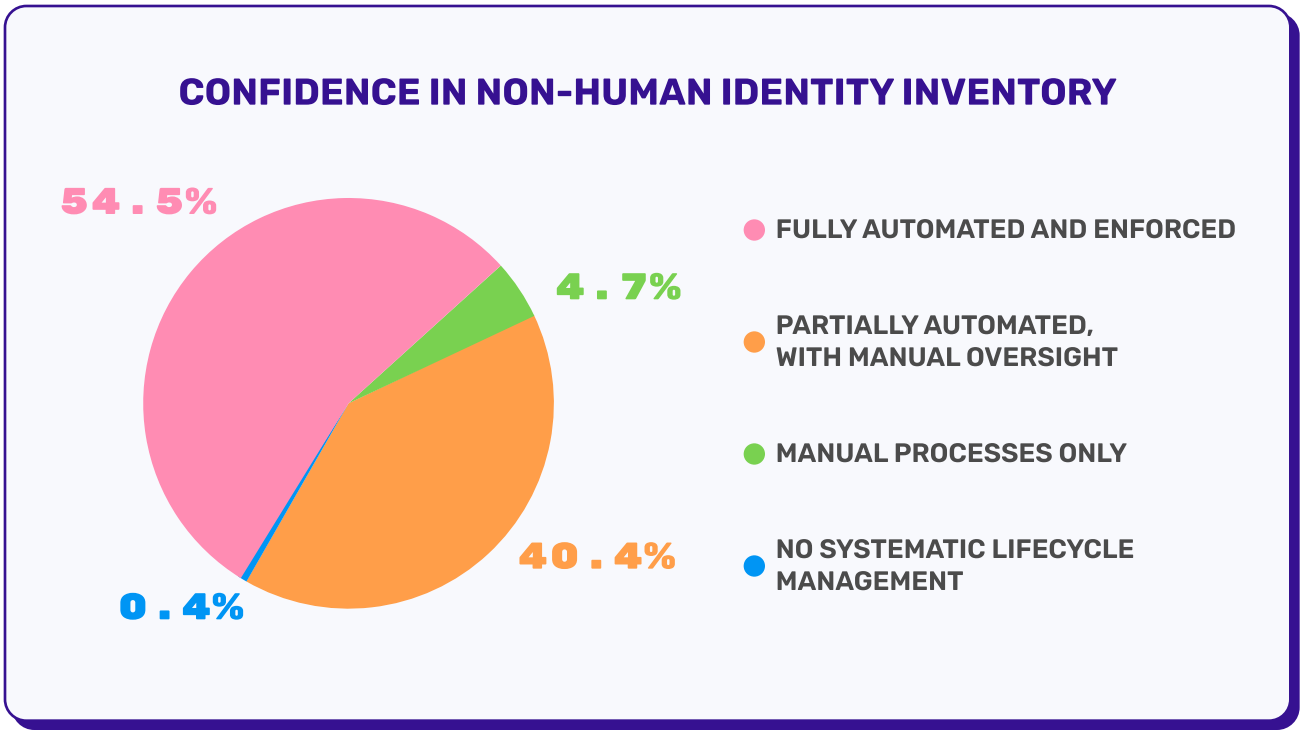

Question asked, "Which best describes how your organization manages non-human identities across their lifecycle (creation → rotation → retirement)?"

Combining discovery methods with lifecycle management approaches reveals distinct maturity levels. On lifecycle management, 55% report fully automated and enforced processes, while 40% have partial automation with manual oversight.

THE MATURITY MATRIX:

Level 4 - Automated Excellence (27% of organizations):
Automated continuous discovery + fully automated lifecycle management. These organizations know when identities are created, monitor their usage, rotate credentials automatically, and retire them when no longer needed.

Level 3 - Hybrid Approach (23%):
Automated discovery + partial automation with oversight. Strong discovery but lifecycle management requires human intervention for key decisions.

Level 2 - Periodic Management (20%):
Scheduled audits/scans + partial automation. These organizations play catch-up, discovering identities in batches and managing them reactively.

Level 1 - Manual Processes (30%):
Manual documentation or incident-based discovery + manual processes or no systematic management. These organizations have neither automated discovery nor automated lifecycle controls.

Only 27% of organizations operate at Level 4 maturity, where both discovery and lifecycle are fully automated. This explains the Confidence-Reality Inversion. Organizations are confident about inventory they're managing with immature processes.

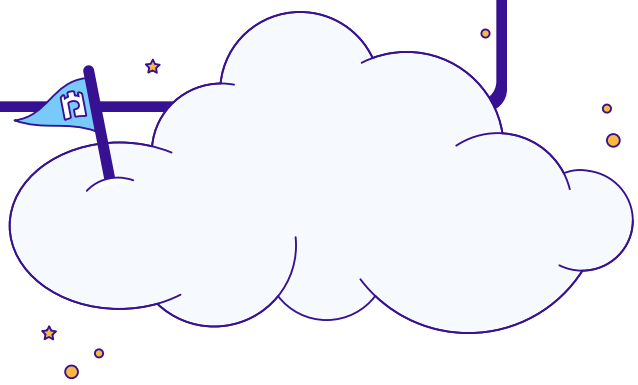

THE CREDENTIAL GRAVEYARD

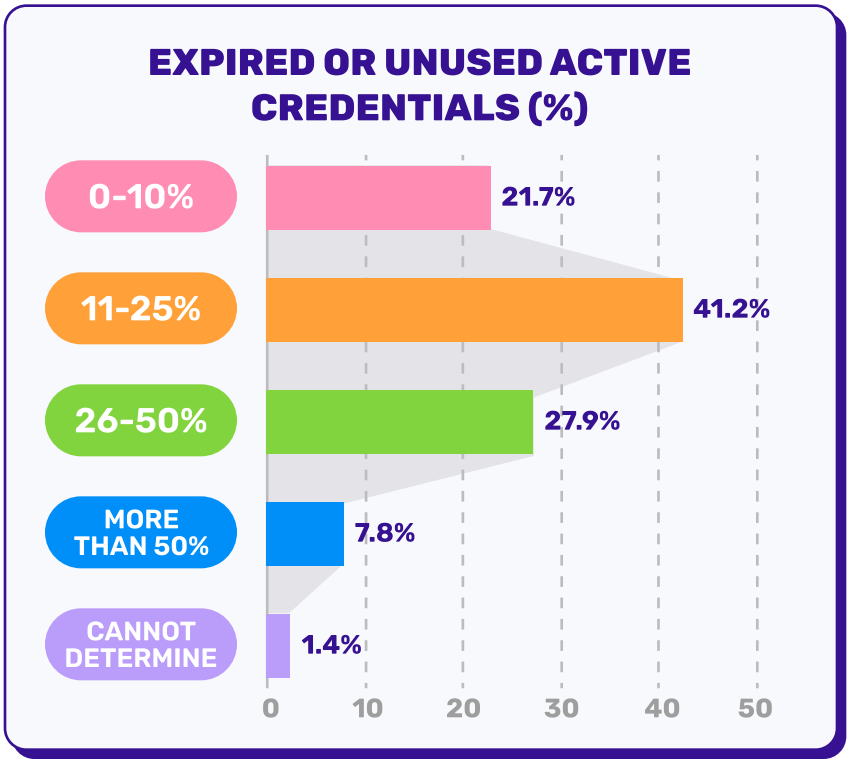

When asked about expired or unused credentials that remain active, 69% of organizations admit that 11-50% of their credentials fall into this category. Another 8% report more than 50% of credentials are expired but still active.

Question asked, "What percentage of credentials (keys, tokens, certificates) are expired or unused but still active?"

This is **The Credential Graveyard Problem**: credentials that should be dead but continue to haunt your environment with active access. These are service accounts from decommissioned applications, API keys for tools no longer in use, certificates that expired months ago but still authenticate successfully, and access tokens generated for one-time tasks that never got revoked.

The largest cluster sits at 11-25% (41%), meaning the typical organization has roughly one in five credentials that are expired or unused but still grant access to systems. For an organization managing 10,000 non-human identities, that's 2,000+ zombie credentials waiting to be discovered by attackers.

*"Expired credentials represent some of the lowest-hanging fruit for attackers," notes **Jason Martin**, **Co-CEO** at **Permiso Security**. "When organizations don't have automated lifecycle management, credentials pile up like unpaid technical debt. Eventually, that debt comes due when an attacker finds a three-year-old service account with admin access that nobody remembered existed."*

DETECTION & RESPONSE

DETECTION IMPROVES, RESPONSE LAGS

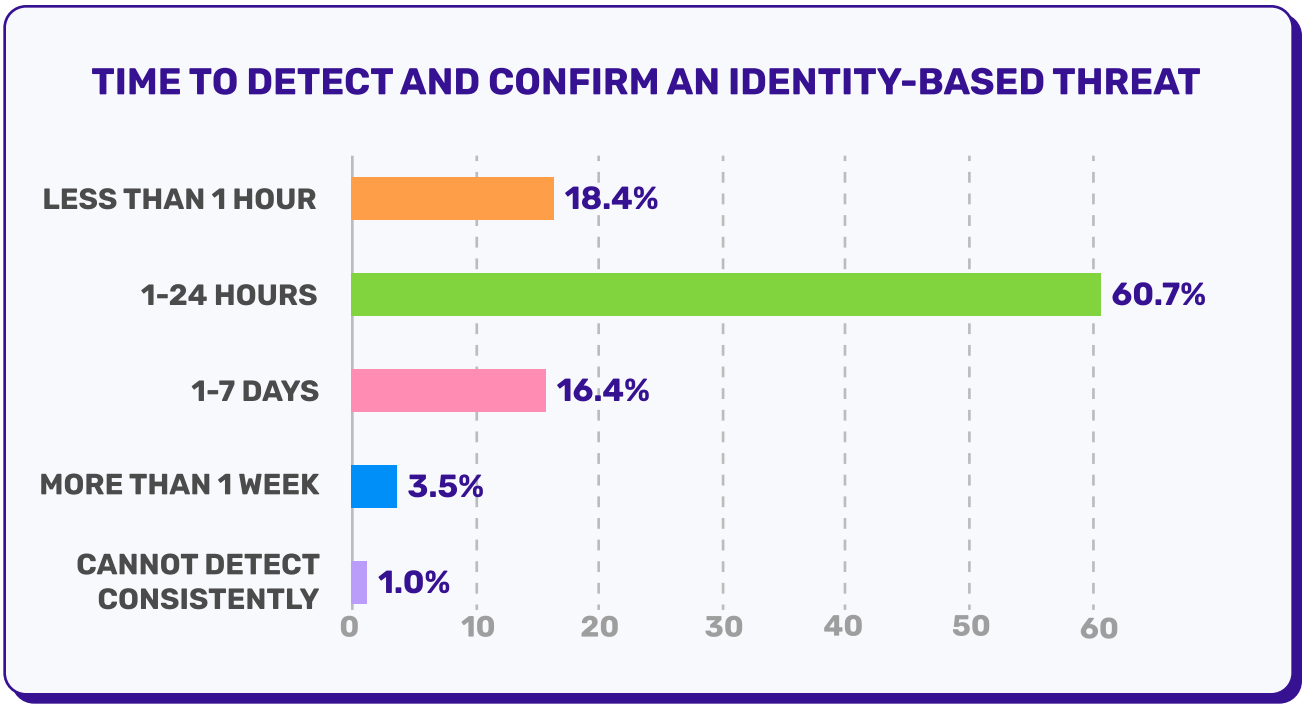

Question asked, "How long does it typically take to detect and confirm an identity-based threat?"

When identity credentials are compromised, every hour matters. The 2025 data shows 79% of organizations can detect and confirm identity-based threats within 24 hours, with 18% achieving sub-hour detection.

This represents dramatic improvement over 2024, when only 61% claimed 24-hour detection (an 18-point increase that's one of the most encouraging findings in our report). Organizations have invested heavily in faster detection capabilities, likely responding to the rising tide of identity attacks.

However, 16% still require 1-7 days for detection, and 4% need more than a week. For these organizations, attackers have essentially unlimited dwell time. The 1% who cannot detect consistently represent organizations with such poor visibility that identity compromise can go undetected indefinitely.

THE BLAST RADIUS BOTTLENECK

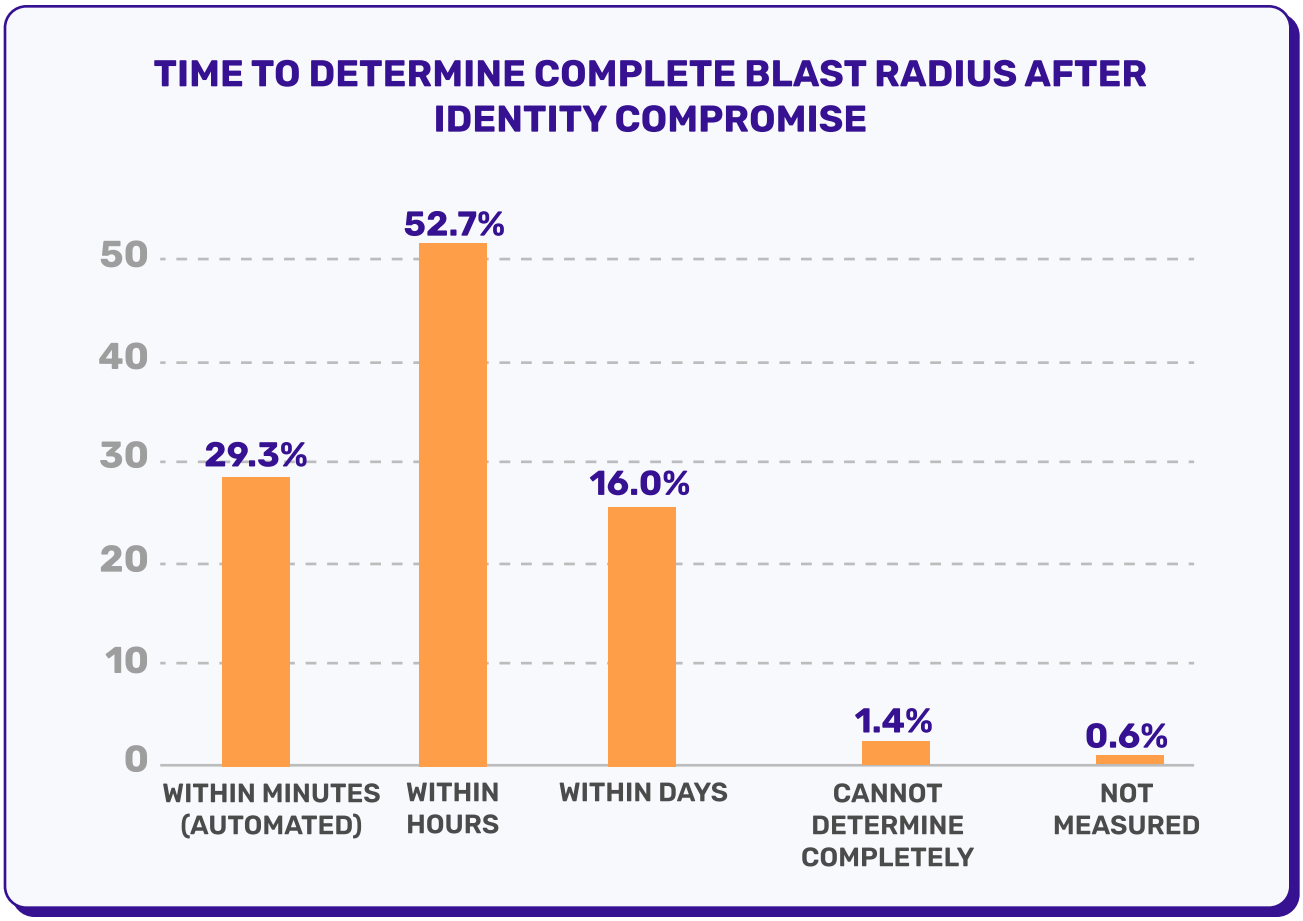

Question asked, "When an identity compromise is detected, how quickly can you determine its complete blast radius?"

Detecting a compromise is only half the battle. Understanding the blast radius (what the compromised identity could access and what damage might have occurred) is equally critical. This is where we see **The Response Speed Gap emerge.**
While 79% can detect threats within 24 hours, only 29% can determine complete blast radius within minutes when compromise is detected. Another 53% need hours, and 16% require days.

The gap matters because detection without blast radius understanding leaves leadership with impossible questions: *"Should we shut down systems? Should we notify customers? How bad is this?"* While analysts spend hours manually correlating access permissions and activity logs, the incident response team makes decisions in the dark.

THE DETECTION-RESPONSE GAP INDEX (DRGI)

Detection without response is theater. The Detection-Response Gap Index measures the critical gap between knowing a compromise occurred and understanding its scope.

DRGI = (% Detect within 24hrs) × (% Determine blast radius within minutes)

2025 Score: 79% × 29% = 22.9%

Zone	% of Organizations	Capability Profile
Response Ready	23%	Detect within 24 hours AND determine blast radius within minutes. Can make immediate containment decisions. Likely have automated correlation and attack path mapping.
Detection Bottleneck	56%	Detect within 24 hours BUT need hours/days for blast radius. Know they're compromised but can't act decisively. Pay the highest Manual Correlation Tax during incidents.
Dual Deficit	21%	Slow detection (over 24 hours) AND slow blast radius determination. Attackers operate with impunity. Likely discover breaches through external notification.

The Time-to-Response Reality:

For the 56% in the detection bottleneck zone, the response timeline looks like:

- **Hour 0-12:** Detect anomaly, confirm compromise
- **Hour 12-36:** Manual correlation across tools to map identity access
- **Hour 36-48:** Determine blast radius and lateral movement
- **Hour 48+:** Begin containment

During those 48 hours, attackers are moving laterally, exfiltrating data, and establishing persistence. By the time blast radius is understood, the initial compromise is the least of your problems.

The Improvement Path: Comparing this to the 18-point improvement in detection speed (from 61% to 79% achieving 24-hour detection), it's clear organizations have invested in detection. The DRGI reveals where the next investment must go: blast radius determination and attack path mapping.

VISIBILITY GAPS CREATE ALERT FATIGUE

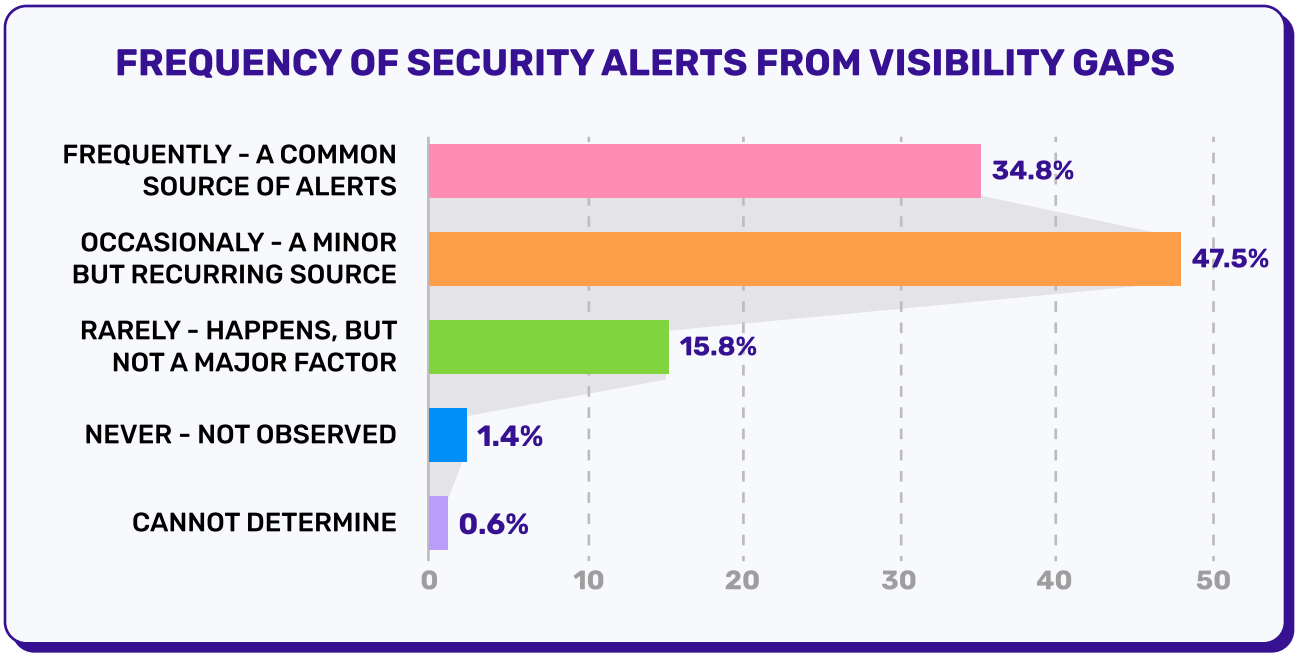

Question asked, "How often do visibility gaps (e.g., unmanaged accounts, shadow identities, misconfigured permissions) result in security alerts in your environment?"

When visibility gaps (unmanaged accounts, shadow identities, misconfigured permissions) trigger security alerts, 82% of organizations report this happens frequently or occasionally. Specifically, 35% report visibility gaps are a common source of alerts, while 48% experience them occasionally.

These alerts represent **The Alert Fatigue Multiplier**: every visibility gap creates false positives, requires manual investigation, and distracts teams from real threats. When your monitoring systems fire alerts about unmanaged accounts you didn't know existed, you're not doing proactive security. You're doing damage control.

*"The organizations that tell us visibility gaps rarely or never trigger alerts are usually the ones with the worst visibility," observes **Paul Nguyen**, **Co-CEO** at **Permiso Security**. "If you can't see the gap, you can't alert on it. The fact that 82% are getting these alerts actually suggests a growing awareness of the problem, even if they haven't solved it yet."*

THE AI IDENTITY CHALLENGE

AI TRANSFORMS IDENTITY CREATION

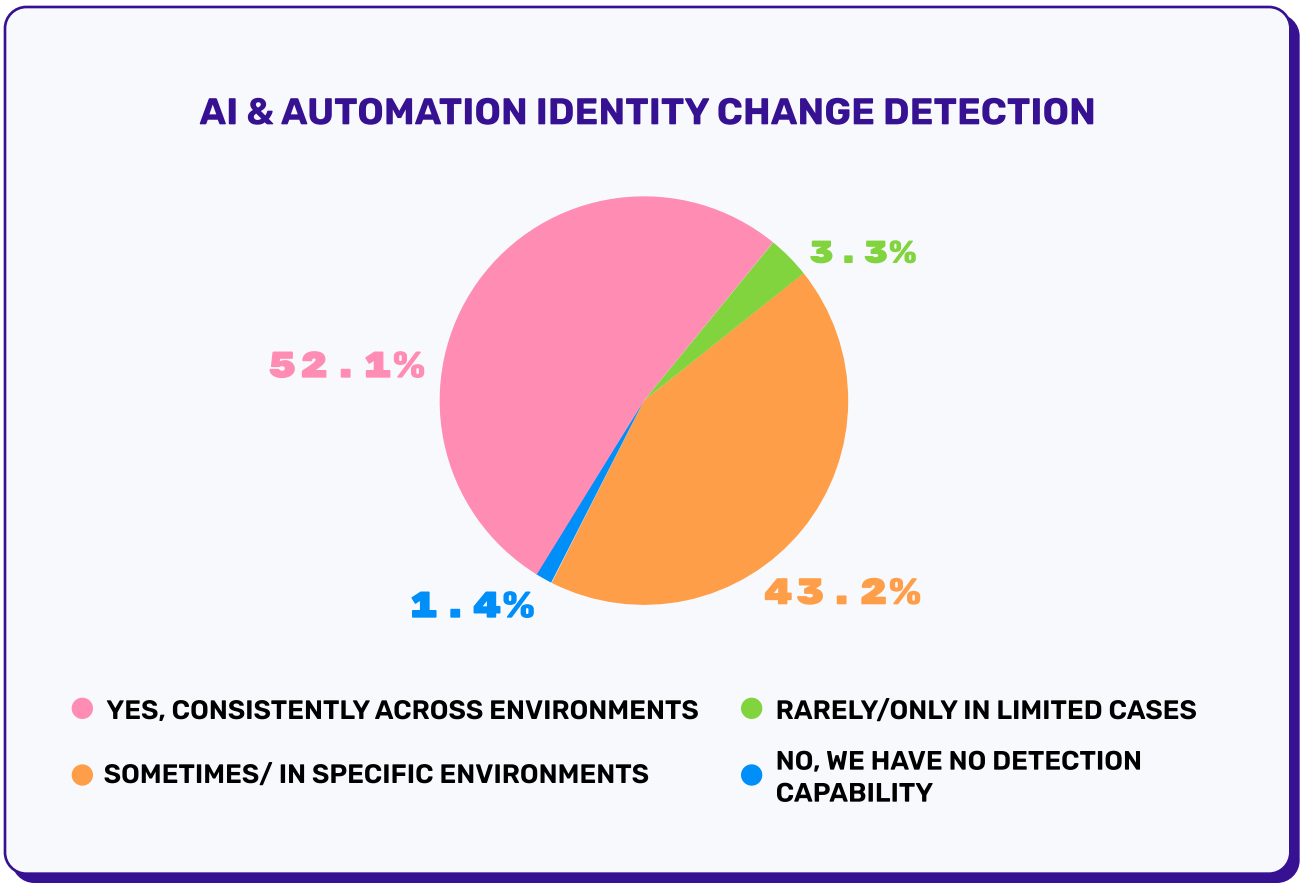

Question asked, "Can you detect when AI systems or automation tools create/modify identities or permissions?"

The rise of AI isn't just changing how we work. It's fundamentally altering who and what has access to systems. In 2025, 95% of organizations report that AI systems or automation tools can create or modify identities and permissions in their environments, with 52% reporting this happens consistently across all environments.

For decades, humans created identities through ticketing systems, approval workflows, and manual provisioning. Now, AI agents and automation platforms are generating identities dynamically, modifying permissions based on algorithmic decisions, and creating access patterns that no human ever reviewed.

The security implications are staggering. When humans create identities, you can audit the decision, question the business justification, and trace accountability. When AI creates identities, who's responsible? The developer who wrote the automation? The business owner who approved the AI deployment? The AI model itself?

AI AGENTS ACCESS SENSITIVE DATA

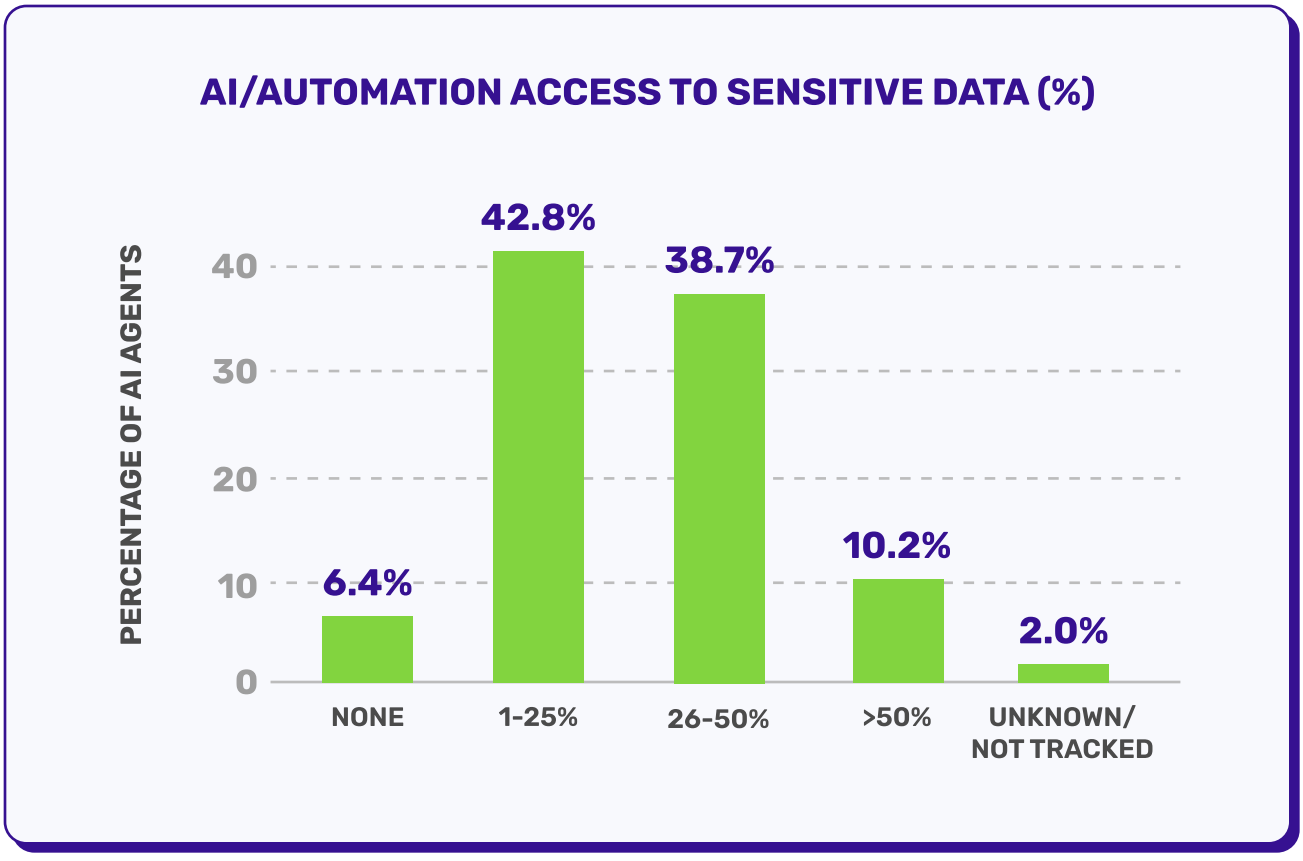

Question asked, "What percentage of AI agents or automated systems have access to production/sensitive data?"

While 95% acknowledge AI systems can create identities, 82% admit AI agents or automated systems already have direct access to production or sensitive data right now. The distribution: 43% report 1-25% of their data is exposed to AI systems, 39% report 26-50%, and 10% report more than 50%.

*"The challenge with AI agents isn't that they're accessing data," explains **Paul Nguyen**, **Co-CEO** at **Permiso Security**. "The challenge is that most organizations don't have visibility into which AI systems have access, what permissions they hold, or what they're doing with the data. These are non-human identities on steroids, with access patterns that traditional monitoring can't detect."*

Consider what 39% reporting 26-50% AI access means in practice. If you have 10TB of sensitive data, AI systems have unfettered access to 2.5-5TB of it. That's customer records, financial data, intellectual property, and trade secrets being processed by systems operating outside your traditional security controls.

AI IDENTITY GROWTH EXPECTATIONS

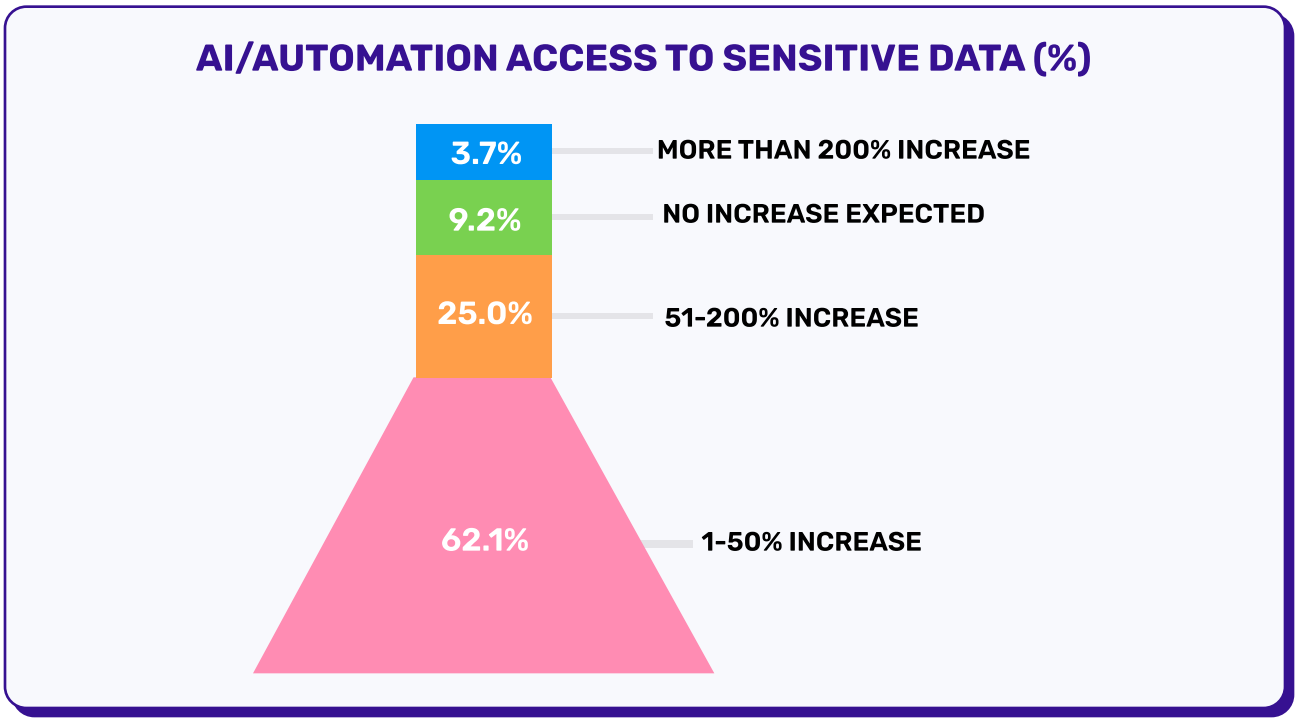

Question asked, "What percentage increase in AI-generated identities do you expect in the next 12 months?"

The AI identity explosion isn't slowing down. Organizations expect AI-generated identities to surge in the next 12 months, with 91% anticipating increases. The majority (62%) expect 1-50% growth, while 25% expect 51-200% growth, and 4% expect increases beyond 200%. Only 9% expect no increase, a figure that seems disconnected from market reality. Every major cloud provider is pushing AI services, every SaaS vendor is adding AI features, and every enterprise is experimenting with AI agents.

AI IDENTITY CRISIS

Truth #1:

AI CREATES IDENTITIES FASTER THAN YOU CAN TRACK THEM

95% of organizations report AI systems can create or modify identities and permissions. Unlike human-requested identities with approval workflows, AI-generated identities appear instantly (no ticketing system, no audit trail, no business justification documented).

Truth #2:

YOU'RE GRANTING DATA ACCESS YOU DON'T TRACK

82% have AI agents accessing production or sensitive data right now. For 39% of organizations, AI systems have access to 26-50% of their data. That's customer records, financial data, intellectual property, and trade secrets being processed by systems that don't appear in your IAM console, don't trigger your DLP policies, and operate outside your traditional security controls.

Truth #3:

THE GROWTH IS EXPONENTIAL, NOT LINEAR

91% expect AI-generated identities to increase in the next 12 months:

- 62% expect 1-50% growth (steady)
- 25% expect 51-200% growth (doubling or tripling)
- 4% expect over 200% growth (exponential)

Your current identity management tools weren't designed for this scale.

THE QUESTION EVERY CISO SHOULD ASK:

"If an AI agent created 500 service accounts last month, can you name even 10 of them?"

THE AI IDENTITY TSUNAMI

Organizations are deploying AI systems faster than they can secure them, granting access faster than they can track it, and generating identities faster than they can manage them. This tsunami collides directly with the visibility crisis we've documented throughout this report. Organizations already struggle to maintain comprehensive visibility into human and traditional non-human identities.

Now they're adding thousands of AI-generated identities with dynamic permissions and unpredictable access patterns. The visibility gap isn't closing. It's accelerating.

OPERATIONAL COSTS & FRAGMENTATION

TOOL SPRAWL FRAGMENTS VISIBILITY

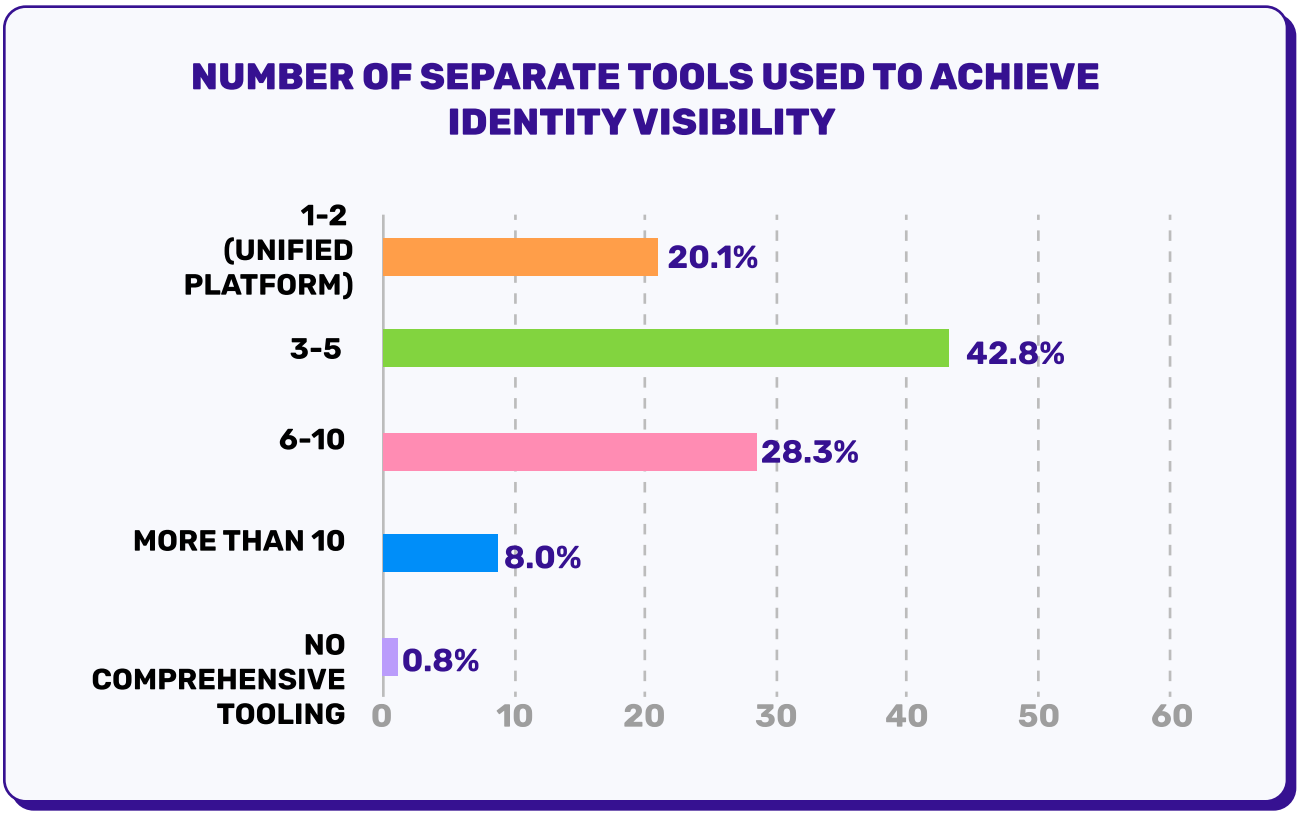

Question asked, "How many separate tools does your team use to achieve identity visibility?"

Faced with identity complexity, organizations have responded predictably: they've bought more tools. The 2025 data shows 71% of organizations use between 3 and 10 separate tools to achieve identity visibility, with 43% using 3-5 tools and 28% using 6-10 tools.

In 2024, organizations were using an average of 2.61 security tools (up 16% from 2.25 in 2023). The 2025 data reveals the situation has worsened specifically for identity visibility, with 71% now using 3-10 separate tools just for identity management. This isn't security strategy. This is desperation manifested as procurement.

The Tool Sprawl Paradox emerges clearly: organizations believe more tools equal better visibility, when in reality, more tools create more fragmentation. Each tool provides a piece of the puzzle (your IdP shows authentication, your CSPM shows cloud identities, your SIEM shows access logs) but no single tool shows the complete picture.

The 8% using more than 10 tools face an identity visibility crisis masquerading as comprehensive security. At this scale, you don't have a security architecture. You have a collection of overlapping capabilities that require full-time staff just to maintain.

THE TOOL BURDEN MATRIX

The relationship between tool count and manual correlation hours isn't linear (it's multiplicative). The Tool Burden Matrix quantifies this correlation and identifies four distinct burden zones:

Burden Zone	Tool Count	Manual Hours/Week	% of Organizations	Annual Cost
Efficient	1-2 tools	<10 hours	20%	<$31K
Moderate	3-5 tools	10-40 hours	43%	$31K-$125K
High Burden	6-10 tools	40-80 hours	28%	$125K-$250K
Critical	10+ tools	80+ hours	8%	$250K+

KEY FINDINGS INCLUDE

- **The Correlation Factor:** Organizations using 6-10 tools spend 4-8x more time on manual correlation than those using 1-2 tools. This isn't additive complexity (each additional tool creates exponential correlation overhead).

- **The Moderate Majority:** 43% of organizations sit in the moderate burden zone with 3-5 tools, spending 10-40 hours weekly (25-50% of two FTEs) just connecting data.

- **The Critical Zone:** The 8% in the critical zone (10+ tools, 80+ hours/week) are spending more than two full-time employees' worth of effort just on manual correlation.

- **The Efficiency Gap:** Only 20% of organizations operate in the efficient zone, having achieved unified platforms or tight tool integration that frees analysts for actual security work.

- **The Industry Cost:** The industry is collectively spending hundreds of millions of dollars annually on a problem that unified visibility would eliminate. This tool sprawl directly enables The Manual Correlation Tax. When identity data lives in 3-10 different tools, someone has to correlate it manually.

MANUAL CORRELATION COSTS

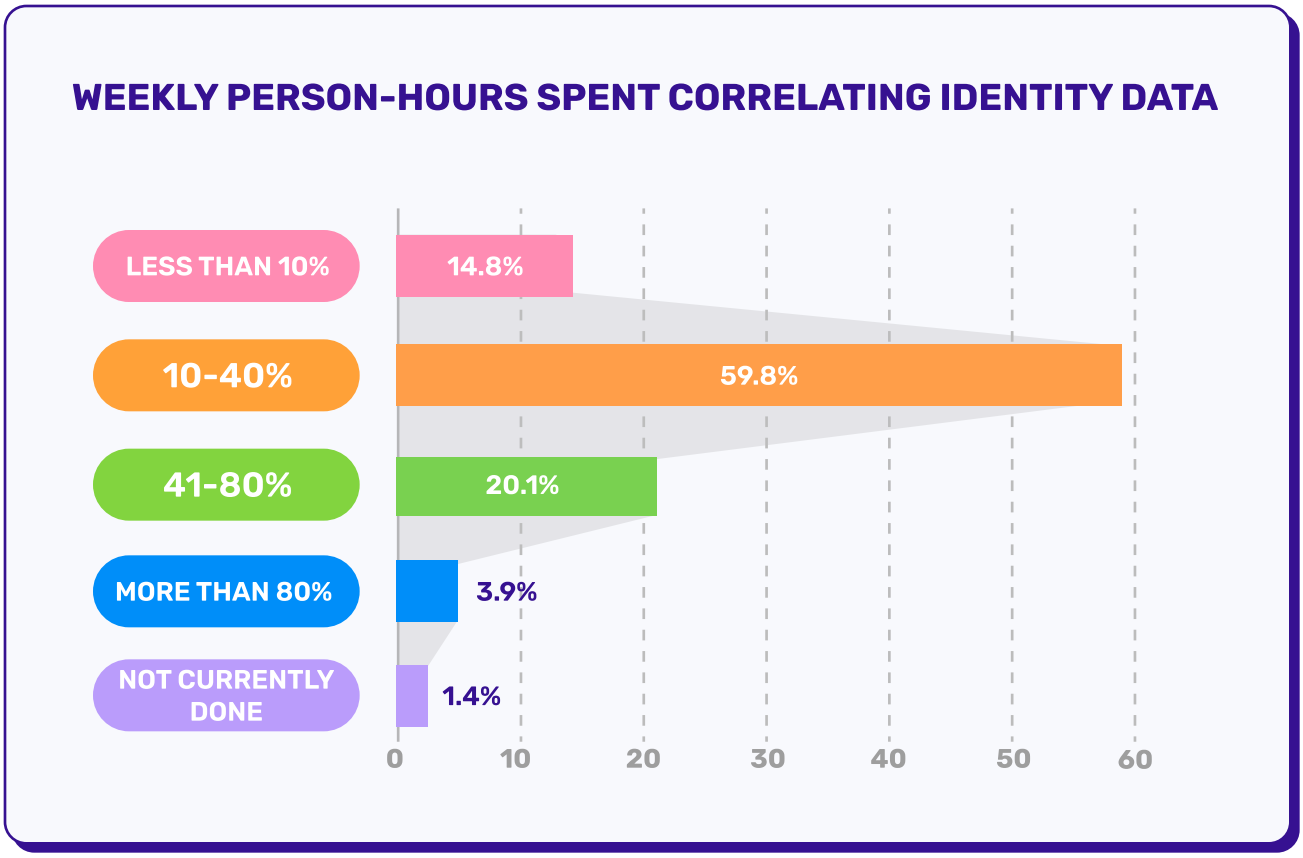

Question asked, "How many person-hours per week does your team spend correlating identity data from different sources?"

We've referenced The Manual Correlation Tax throughout this report. This question finally quantifies it, and the numbers are brutal. Organizations spend 10-40 hours per week (60%) or even 41-80+ hours per week (20%) manually correlating identity data from different sources.

A security analyst earning $120K annually costs roughly $60/hour. For the 60% spending 10-40 hours weekly on manual correlation, that's $600-$2,400 per week ($31K-$125K annually) in direct labor costs. For the 20% spending 41-80+ hours, the costs approach or exceed $200K annually when you account for multiple analysts.

But the real cost isn't just money. It's opportunity cost. Every hour spent manually correlating identity data is an hour not spent hunting threats, improving defenses, or responding to incidents.

> *"When we talk to security teams drowning in manual correlation, they all describe the same pattern," notes* ***Ian Ahl, CTO*** *at* ***Permiso Security****. "They know which identities to investigate, but by the time they've pulled logs from five different systems, mapped the identity across three different formats, and reconstructed the timeline, the incident has evolved. They're always responding to yesterday's attack."*

The connection between tool sprawl and correlation hours is direct: more tools equal more correlation burden. The industry's response to identity complexity has been to add more point solutions, which has only deepened the problem.

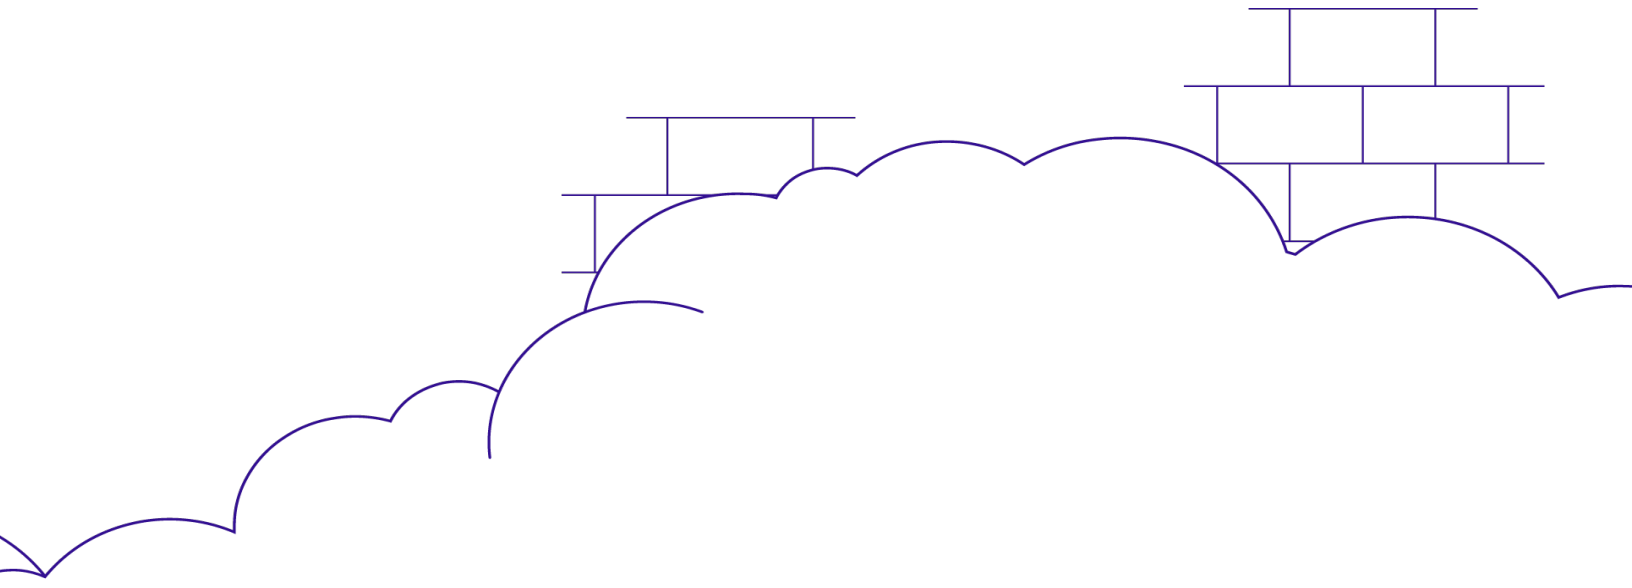

BUSINESS IMPACT & MARKET RESPONSE

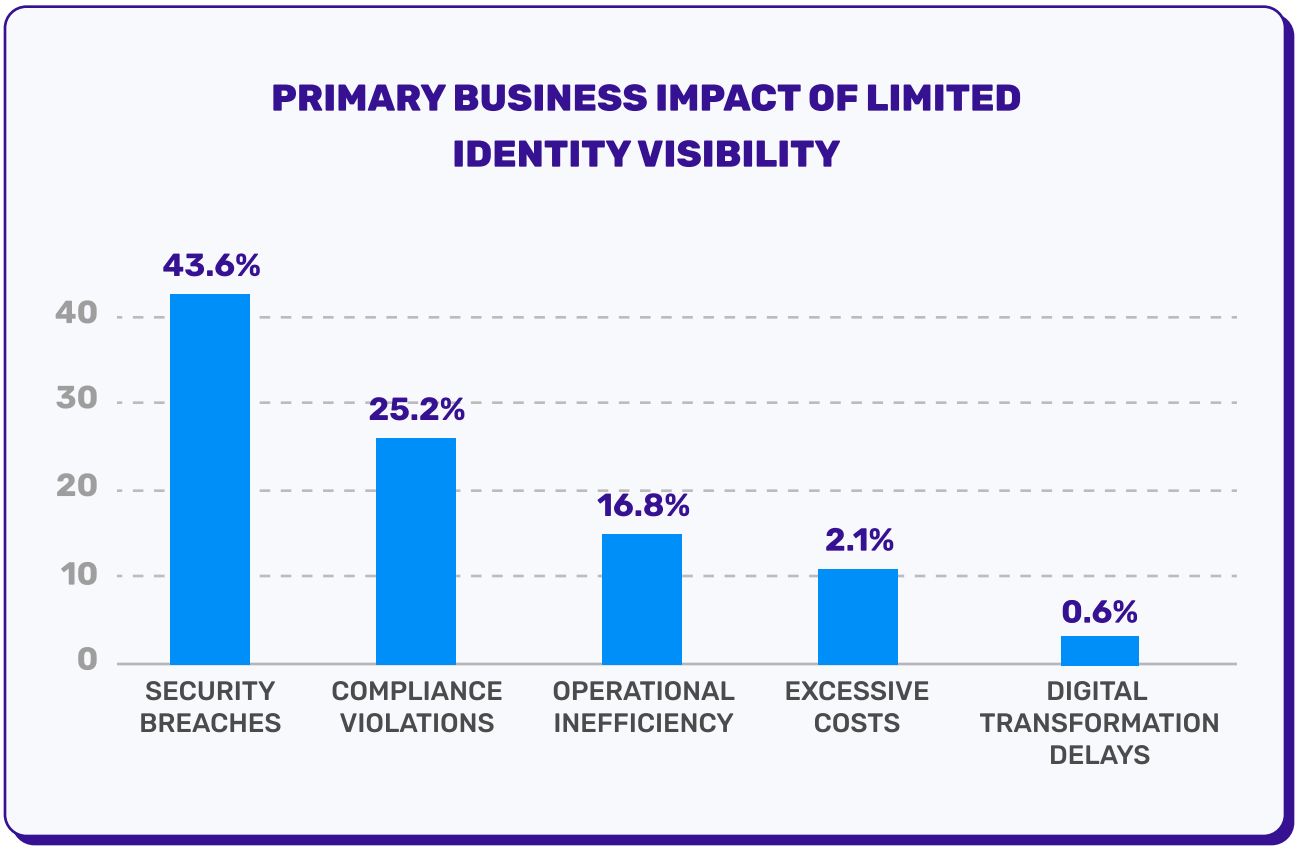

Question asked, “What is the primary business impact of limited identity visibility?”

Security breaches dominate the business impact discussion, with 44% citing them as the primary consequence of limited identity visibility. But the complete picture reveals that visibility gaps touch every corner of the business:

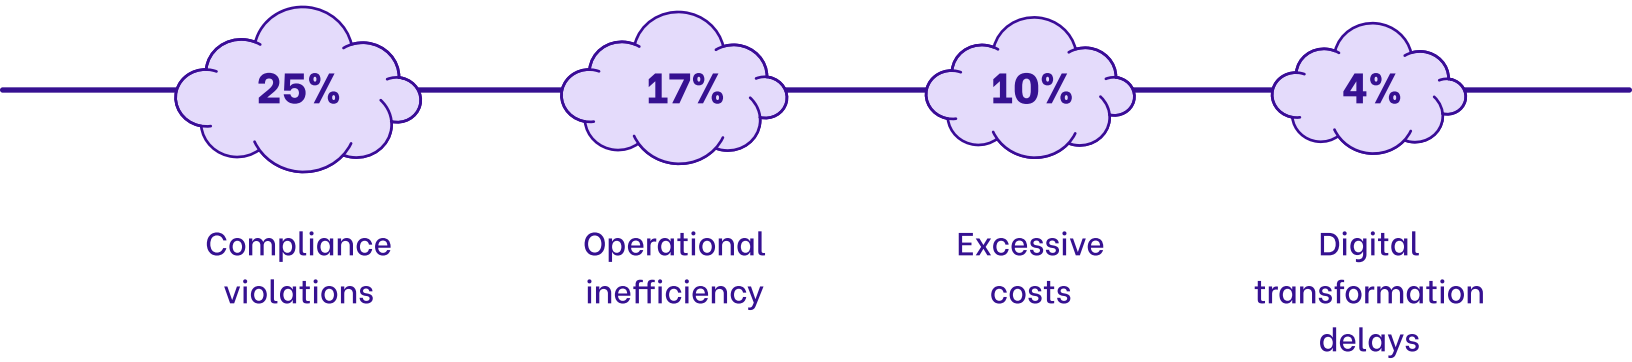

Compliance violations

Operational inefficiency

Excessive costs

Digital transformation delays

The 44% focused on security breaches understand the direct line between visibility gaps and successful attacks. The 25% citing compliance violations face a different but equally serious concern (regulations from SOC 2 to GDPR require organizations to know who has access to what data). Limited identity visibility makes compliance impossible to prove.

Perhaps most interesting is the 17% identifying operational inefficiency. These organizations have recognized that **The Manual Correlation Tax** isn't just a security burden. It's operational overhead that slows down incident response, delays access provisioning, and consumes resources that could be deployed elsewhere.

WHAT SECURITY TEAMS ACTUALLY WANT

MOST IMPACTFUL CAPABILITY

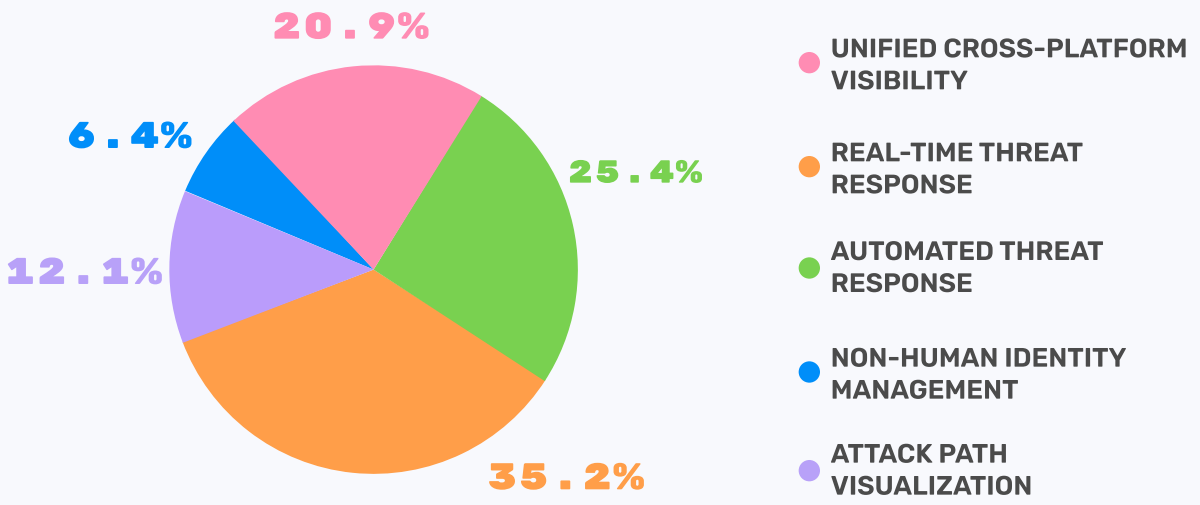

Question asked, "Which identity visibility capability would most improve your security posture?"

OTHER MEANINGFUL CAPABILITIES

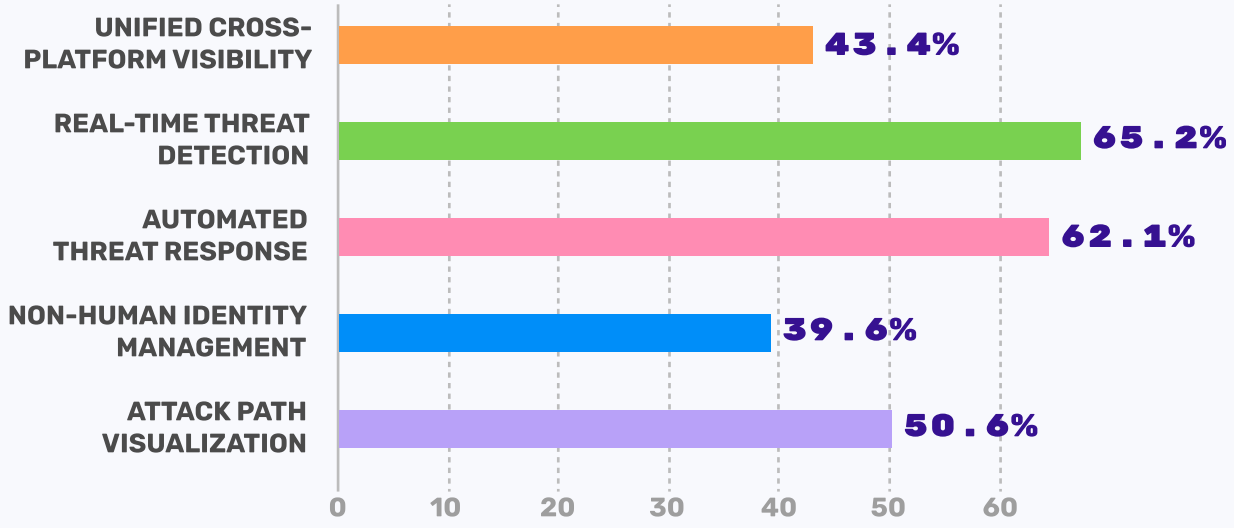

Question asked, "Which other capabilities would also meaningfully improve your posture?"

The gap between what teams want and what they can achieve reveals the next frontier of identity security. Real-time threat detection stands out as the top capability, with 35% selecting it as their first choice and 65% naming it as a meaningful improvement overall.

Unified cross-platform visibility also ranks high, chosen by 21% as their single most important capability and by 43% overall. Automated threat response (25% first choice, 62% overall), non-human identity management (6% first choice, 40% overall), and attack path visualization (12% first choice, 51% overall) round out the top priorities.

THE GAP BETWEEN WHAT TEAMS WANT AND WHAT THEY CAN SEE

The results show a clear pattern: security teams are not asking for more tooling. They're asking for faster answers and fuller visibility. Real-time threat detection ranks highest because teams are tired of discovering breaches after the fact. Automated threat response comes right behind it because detection without action still leaves a gap attackers can exploit.

But here's where the real insight emerges. Even though real-time detection and automated response are the top picks, unified cross-platform visibility still shows up strongly across both charts. This tells us that teams know detection and response only work if they're built on complete identity awareness. If your visibility is fragmented, your detection will be incomplete and your automation will be unreliable.

THE TAKEAWAY

Security teams don't just want data. They want context, correlation, and confidence. The emphasis on unified visibility makes it clear that the next frontier of identity security isn't more alerts or faster responses. It's building a foundation where teams can see everything happening with every identity, in real time, across every platform.

*"Organizations keep asking us for faster threat detection," observes **Jason Martin**, **Co-CEO** at **Permiso Security**.*

"But when we dig into what's slowing them down, it's always the same answer: fragmented visibility. You can't detect what you can't see, and you can't respond quickly when you're spending hours correlating data manually. The fastest path to better detection isn't better detection tools. It's unified visibility."

INVESTMENT SURGE IN 2026

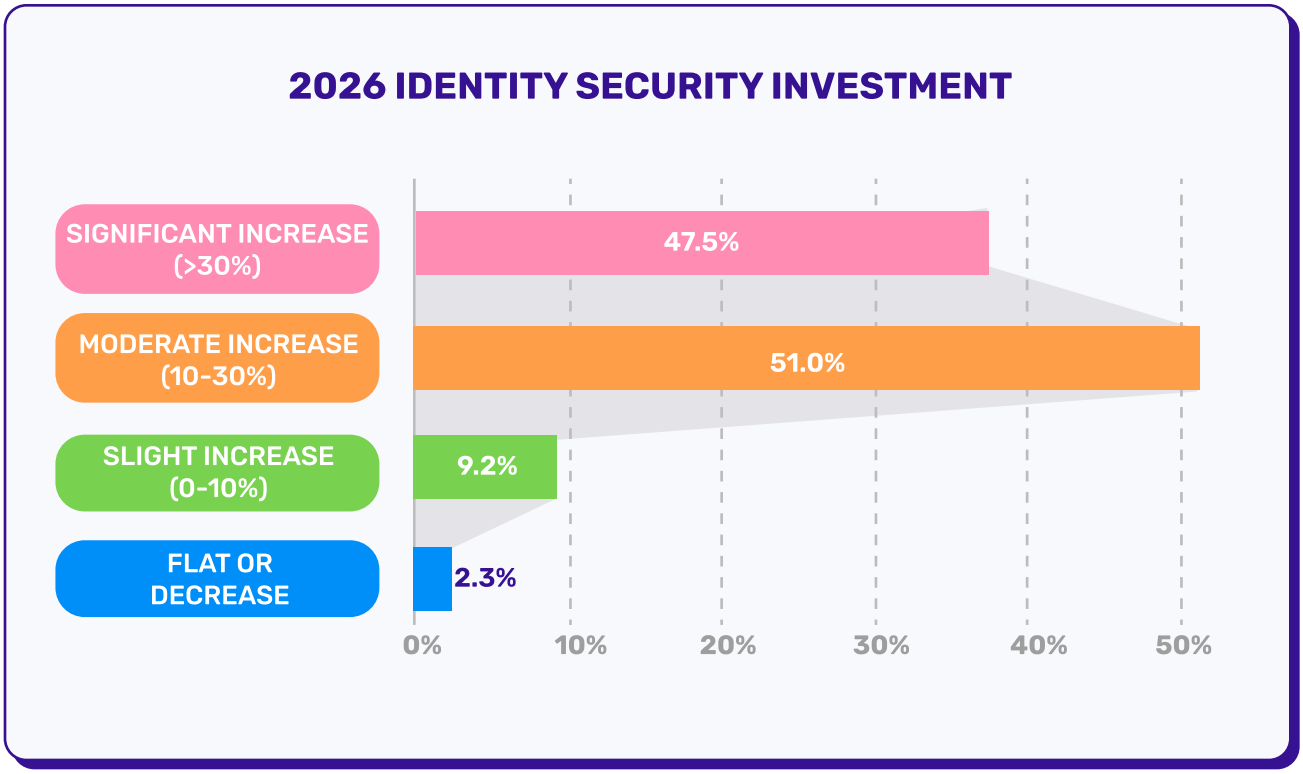

Question asked, "What is your planned investment change in identity security for 2026?"

Despite all the challenges documented in this report (or perhaps because of them), organizations are committing to dramatic increases in identity security investment. 89% plan to increase spending in 2026, with 38% planning significant increases over 30% and 51% planning moderate increases of 10-30%.

Only 9% plan slight increases, and just 2% expect flat or decreased budgets. This represents one of the strongest investment signals we've ever seen in identity security.

THE INVESTMENT SURGE DRIVERS

Multiple factors are driving this surge. First, the identity incident data provides clear justification. When 77% of organizations report that 26-75% of security incidents are identity-related, CFOs can't argue against identity security investment.

Second, the visibility gaps documented throughout this report (The Visibility Illusion Cascade, The Manual Correlation Tax, The Confidence-Reality Inversion) have created operational pain that demands solutions. Organizations are burning 10-80 hours per week on manual correlation while missing threats that comprehensive visibility would catch.

Third, The AI Identity Tsunami is forcing investment. Organizations cannot deploy AI systems at scale while managing AI-generated identities with tools built for human users. The 91% expecting AI identity growth recognize they need new capabilities.

The organizations planning 30%+ budget increases represent the vanguard (organizations that have experienced the consequences of limited visibility firsthand, whether through breaches, compliance failures, or operational inefficiency). They've moved identity security from a component of their security budget to a strategic priority.

CONCLUSION

The 2025 State of Identity Security Report reveals an industry at an inflection point. Organizations face unprecedented identity complexity: multi-cloud infrastructure, fragmented identity providers, exploding non-human identity populations, and The AI Identity Tsunami adding thousands of new identities with unpredictable access patterns.

Against this complexity, we've documented The Visibility Illusion Cascade. Organizations claim comprehensive visibility (46%, down 47 points from 2024's 93%), but when pressed about specific capabilities, confidence systematically declines. Only 43% can proactively detect risks before incidents, despite 95% expressing confidence in their non-human identity inventory. This is The Confidence-Reality Inversion: highest confidence in the area of greatest complexity and lowest actual capability.

The operational cost of this visibility gap manifests in The Manual Correlation Tax. Organizations spend 10-80 hours per week manually correlating identity data across 3-10 separate tools, burning analyst time and delaying incident response at the exact moments when speed matters most. This tax is measured in successful attacks, preventable breaches, and incidents that organizations estimate 26-75% could have been avoided with comprehensive visibility.

The identity threat landscape has fundamentally shifted. With 77% reporting that 26-75% of incidents are identity-related, attackers have clearly recognized what many organizations haven't yet accepted: identity is the new perimeter.

Yet hope emerges in the investment data and improved detection speeds. The 89% of organizations planning identity security budget increases in 2026 signals market recognition of the problem. The 18-point improvement in 24-hour detection rates (from 61% in 2024 to 79% in 2025) shows that when organizations invest in identity security, results follow.

Organizations understand they cannot secure what they cannot see, cannot detect threats across fragmented tools, and cannot respond effectively while paying The Manual Correlation Tax.

The path forward requires moving from point solutions to unified platforms, from manual correlation to automated analysis, and from reactive detection to proactive risk identification. Security teams don't just want more tools or more data. They want context, correlation, and confidence.

The organizations that achieve unified visibility will operate at a fundamentally different security posture than their peers. They'll detect threats before incidents occur, respond in minutes instead of hours, and prevent the 26-75% of incidents that comprehensive visibility makes preventable.

The organizations that don't will continue paying The Manual Correlation Tax, operating with The Confidence-Reality Inversion, and discovering breaches after attackers have already achieved their objectives.

Identity security in 2025 is defined not by what organizations think they control, but by what they can actually see. The future belongs to organizations that close the visibility gap before attackers exploit it.

PERMISO

THANK YOU

QUESTIONS?

CONTACT US AT HELLO@PERMISO.IO