Lumos

2026

AI, Automation, & Risk in 2026

IDENTITY AT A BREAKING POINT

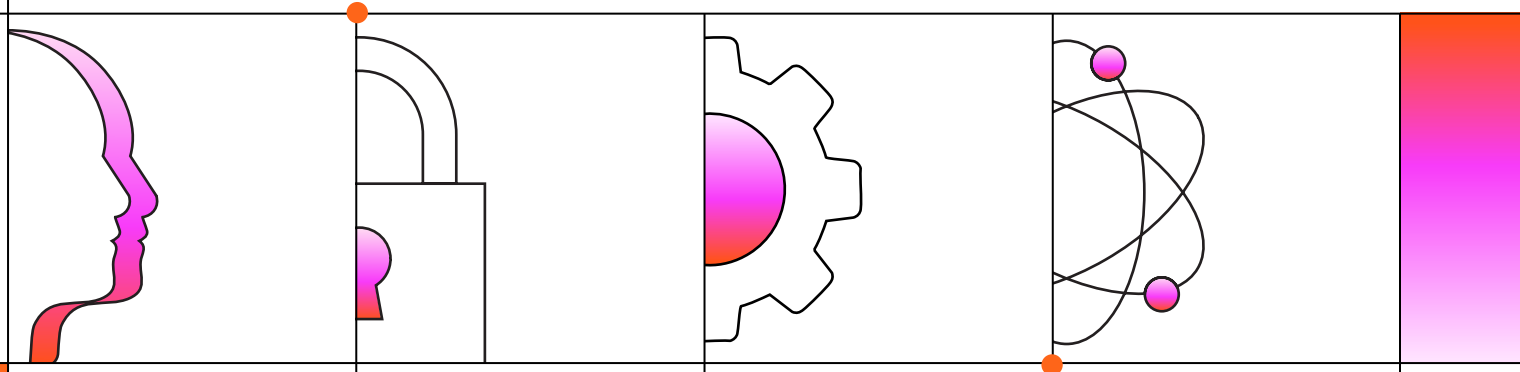

Table of Contents

FOREWORD

Identity is Your Greatest Exposure

Confidence is a funny thing. In the right hands, it's a superpower. In the wrong environment, it's a blindfold. Right now, there is a massive gap between how safe companies feel and how safe they actually are. It's a mirage of preparedness that looks good on a compliance spreadsheet but dissolves the moment it hits the real world.

The numbers tell a jarring story. Over 90% of security leaders believe they are ready for an identity-based attack. Yet, look at the rearview mirror: only 3.8% of organizations managed to navigate the last year without a significant identity-related incident.

This isn't because of a lack of effort, it's because of the operational chaos that lives in the shadows of every large enterprise. While we've been trying to keep up with the volume of requirements to satisfy auditors, the dynamic nature of the business and the sheer volume of tools means that inevitably something gets missed or overlooked. And in the background, attackers have been busy exploiting the messy, forgotten accounts and excessive permissions we left behind. We are no longer in an era where hackers are breaking in with sophisticated malware, groups like Volt Typhoon and Scattered Spider simply log in. They use legitimate credentials and quiet, lateral movement to hide in plain sight, evading the very detections we rely on to keep them out. Identity can no longer be treated as a technical control–it's a business risk that determines how much damage an incident can have.

We've reached a tipping point where technical debt and the explosion of AI- and SaaS-based identities have outpaced human-speed oversight. Organizations are in a state of paralysis, hesitating to embrace full automation because of a fundamental distrust in the results of automated tools.

Unfortunately, the reality is that while compliance overlaps with good security, it does not equal good security.

Risk has emerged as the ultimate catalyst for action in the field of identity and access management, far outweighing simple productivity or operational ease. We must move toward a world where we can assess risk at scale and take action with the same machine-speed as our adversaries. The path forward isn't just about more tools; it's about adopting AI-native processes that turn visibility into accountability.

What follows is an analysis of the specific vulnerabilities, from MFA fatigue to dormant access, that are being weaponized against us today. These insights are designed to help you move past the mirage of confidence and toward a state of actual, verifiable security.

Executive Summary

The AI, Automation, and Risk in 2026 report marks a definitive breaking point in enterprise security. There is a dangerous disconnect between identity leaders' perceived security and the operational realities of the current landscape.

Right now, identity is the main way hackers break into businesses: CrowdStrike's latest threat report found that 79% of cyberattacks involved no malware at all. Most modern attacks do not use complicated malware. Instead, they use stolen passwords or trick people into clicking "approve" on login notifications. This has become the primary battlefield for risk.

The mess is growing faster than anyone can keep up with due to the sheer number of accounts that now exist. We have reached a point where machine identities like bots, service accounts, and apps can outnumber human workers by as many as 20 to 1.

This massive sprawl of accounts, combined with people and machines having way more access than they actually need, has created a management nightmare that is impossible to handle manually.

There is a massive gap between how safe companies think they are and how safe they actually are. While over 90% of security leaders say they feel prepared to defend themselves, only 3.8% of companies actually made it through the last year without a significant incident. This gap exists because most identity tools are built to satisfy auditors and check boxes rather than stopping the operational chaos that hackers actually use to get inside.

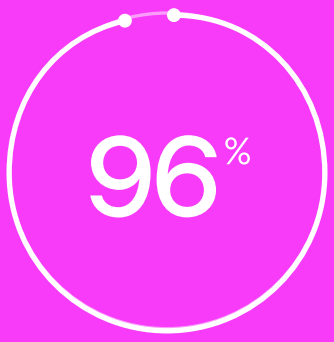

Have already experienced a significant identity-related problem

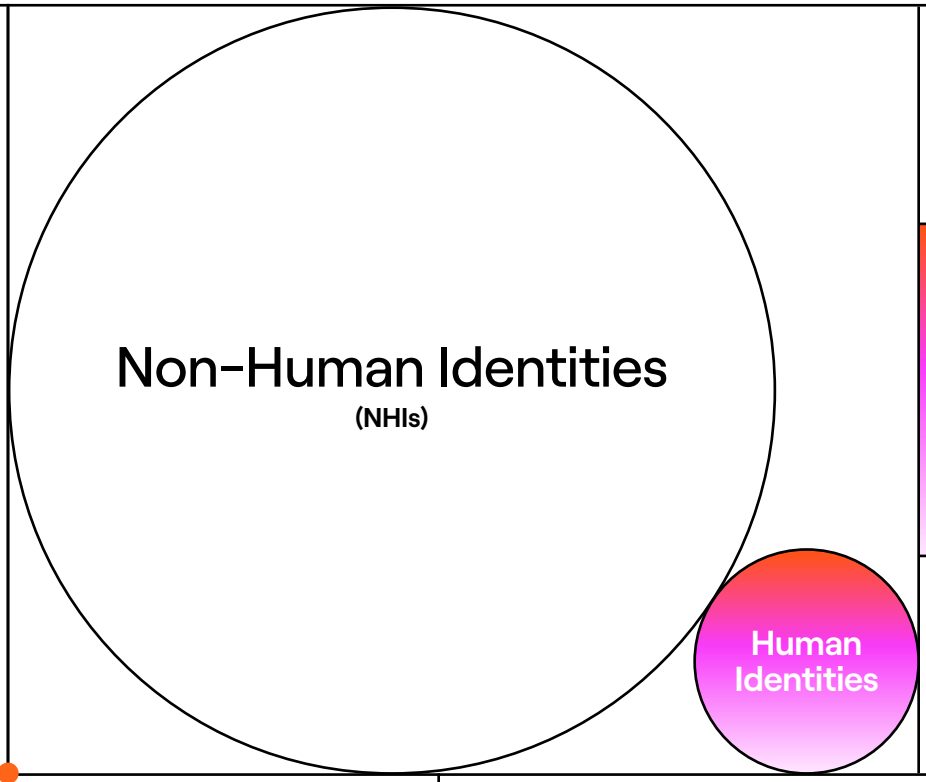

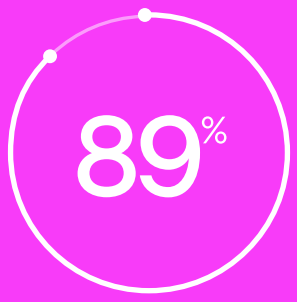

of leaders view AI as a necessity to manage administrative burden

Because of this, companies are realizing that managing risk is now more important than employee productivity or operational speed. We can no longer afford to rely on old ways of working that leave us vulnerable to real-world threats like dormant accounts or hackers moving through our systems. The need to change is no longer just an option. It is a fundamental necessity for an organization to survive.

88.7% of leaders view AI as a necessity to manage the overwhelming administrative burden of identity programs, yet the industry is still struggling to utilize these tools correctly. However, many teams are still stuck because they do not trust the results of a "black box" process that they cannot understand or audit.

The solution is a new approach that uses agentic, contextual intelligence to provide the transparency and visibility that teams need to trust automation. By moving from manual spreadsheets to smart, automated remediation, organizations can finally close the Confidence Gap and secure their systems at machine speed.

Methodolgy

The **AI, Automation, and Risk in 2026: Identity at a Breaking Point** report is based on primary research conducted among a cross-section of global enterprises to quantify the gap between perceived identity security and operational reality.

Survey Scope and Dataset

Qualified Respondents: Analysis is based on 133 qualified respondents.

Organizational Scale: Participating organizations ranged from 500 to over 10,000 employees.

Timeframe: Data was collected in November 2025 and reflects the security posture and incidents occurring throughout the preceding year.

Participant Demographics

Roles: The survey targeted high-level decision-makers and influencers, including CISOs, CIOs, CTOs, VP/Directors of IT, and Security/Identity Managers.

Industries: Data represents a diverse set of sectors, including Education, Energy & Utilities, Financial Services, Healthcare, Manufacturing, and Technology & Software.

Geography: 95% of respondents are based in North America, with additional representation from Latin America, Asia-Pacific, and the Middle East.

Key Definitions:

Dormant Access: Any account or entitlement that has remained inactive for greater than 90 days.

Non-Human Identities (NHIs)/Machine Identities: Accounts not tied to an individual human user, encompassing service accounts, API keys, tokens, secrets, certificates, and automation credentials.

Identity-Related Incident: Any security event or "near miss" stemming from the challenges of identity and access management, such as credential theft or MFA fatigue.

Introduction

Identity: The Modern Battlefield

The traditional network perimeter is gone, leaving identity as one of the last and most critical lines of defense. Despite heavy investment in MFA and IAM frameworks, 96% of organizations have already faced significant identity-related problems. This report explores why current defenses are falling short and how modern adversaries, like the collective **Scattered Spider** or **Volt Typhoon**, are aggressively manipulating human and identity systems to achieve their goals.

3 CORE AREAS OF CONCERN SHAPING THE 2026 THREAT LANDSCAPE:

The Accumulation of Excessive Privilege

55% of leaders cite the unchecked growth of permissions as their top hurdle, leading to "permission creep" where accounts hold far more access than required.

The Invisibility of Non-Human Identities (NHIs)

20:1 Machine identities now outnumber human users by ratios as high as 20:1 yet governance for these automated actors remains the area where organizations feel least prepared.

Real-Time Detection Gaps

48% of teams struggle to detect identity misuse in real time, leaving them blind to attacks as they happen.

Adversaries no longer need complex technical exploits; they use stolen credentials (43.6%) and MFA fatigue (48.1%) to walk through the front door. This report provides a roadmap for closing these gaps by embracing AI-driven automation and zero-trust architectures to manage the overwhelming administrative burden of modern identity programs.

43.6% Stolen Credentials

48.1% MFA Fatigue

Identity Isn't as Secure as Leaders Think

Modern enterprises are grappling with a fundamental disconnect between how they perceive their identity security and the actual, operational reality.

Leaders know that identity is a cornerstone of their organization's security. In fact, over 90% of leaders say their organization is "prepared" or "very prepared" to defend against identity-based attacks, but the actual data presents a different picture. Our research found that only 3.8% of organizations made it through the last year with no identity-related incidents, an alarming success rate for attackers.

of leaders who say their organization is "prepared" to defend against identity-based attacks

of organizations who made it through the last year with no identity-related incidents.

This disconnect isn't just a statistical anomaly; it stems from increasingly clever tactics that take advantage of the human element and administrative mistakes. When we examine the specific methods being used, the weaknesses in current defenses become apparent:

Credential Exploitation

The most common point of entry is also the simplest. Stolen or reused credentials (43.6%) continue to slip past basic protections, often paired with MFA Fatigue Attacks (48.1%), which exploit user frustration to gain unauthorized access.

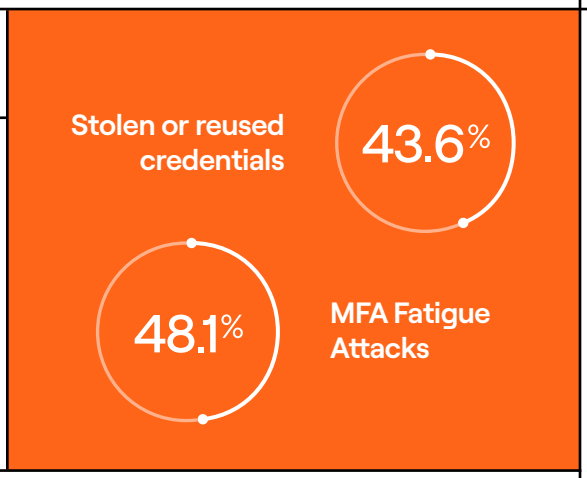

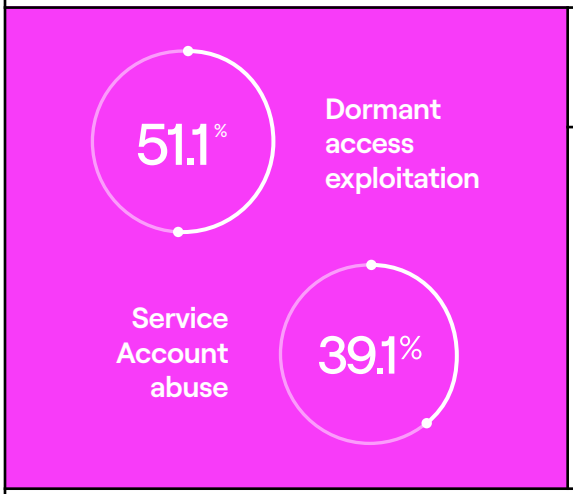

Wide Attack Surface

Perhaps most alarming is the exploitation of what IT teams cannot see or have forgotten. Dormant access exploitation (51.1%) and Service Account abuse (39.1%) create a large, unmonitored attack surface that is always open for bad actors.

Internal Vulnerabilities

Security is also failing from within. Insider Access misuse (46.6%) and Lateral Movement (37.6%) show that once a perimeter is breached—or if the threat is already inside—businesses lack the detailed visibility needed to prevent further damage.

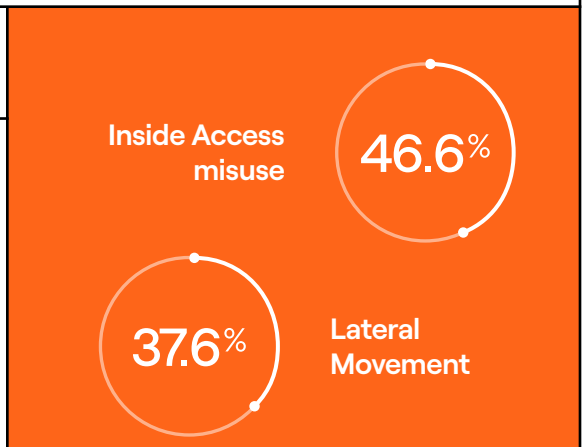

The data indicates that "preparedness" is often based on having tools rather than achieving effective outcomes. Organizations are checking off boxes for security compliance but remain at risk from the most common identity-based tactics.

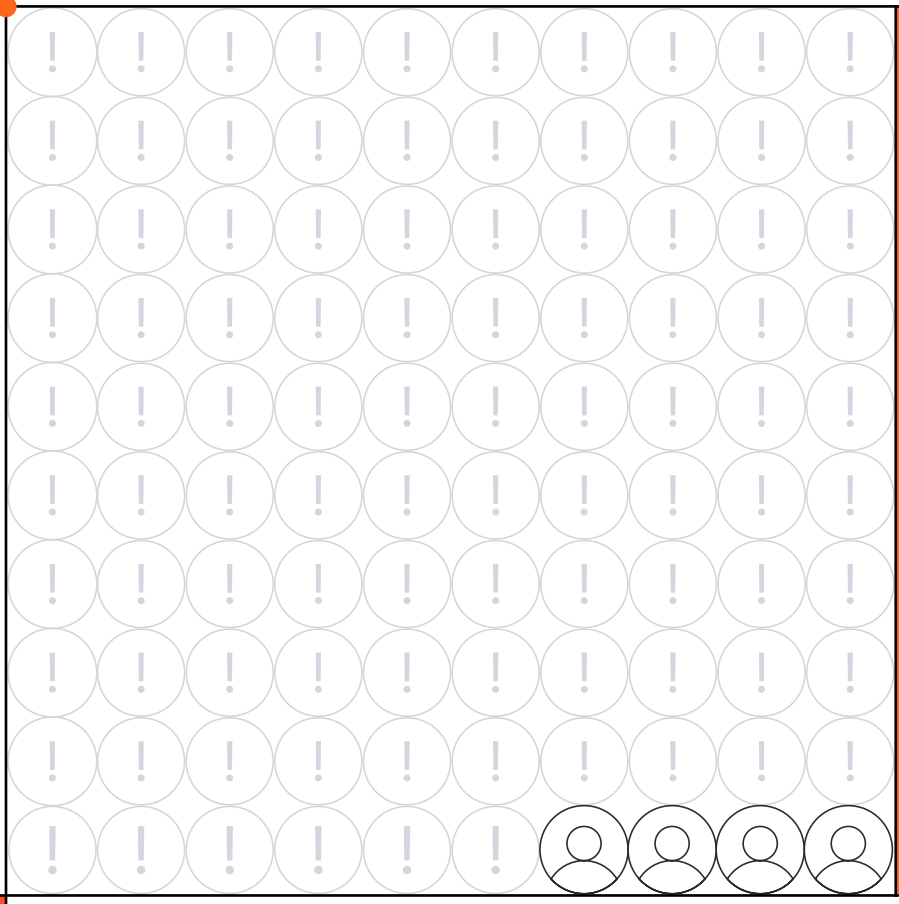

Only

3.8%

of organizations made it through the last year with no identity-related incidents

This gap exists because identity security is an ongoing operational challenge that can't be solved with set-and-forget tools. Until organizations move beyond the general perception of safety and start addressing specific, frequent threats like dormant accounts and MFA fatigue, the 3.8% of organizations that have remained "untouched" will only continue to decrease.

Risk is the First Priority

Properly managing Identity is difficult and time consuming, and operational efficiency is often one of the largest drivers of Identity conversations. But in the end, identity leaders are almost twice as concerned about managing risk as they are about employee productivity, compliance gaps, or operational efficiencies. Justifiably so, when **96.2% of organizations have experienced identity-related security incidents.**

The Two Most Pressing Identity Challenges

When forced to identify which specific identity hurdles cause the most friction, two primary challenges emerged as the most critical:

Almost 55% of respondents identified excessive privilege accumulation as one of their top concerns, and it isn't surprising that inability to detect misuse in real-time was the second-ranked issue at 48%. The gap in visibility prevents security teams from identifying and stopping identity-based attacks as they happen.

Excessive Privilege Accumulation

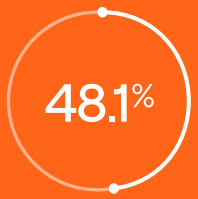

Inability to Detect Misuse in Real-Time

Combined, these two issues create a concerning blind spot for identity leaders. These issues are both a management as well as a visibility problem. Proper controls over access and privilege have to be combined with the ability to identify problem areas that need further control. It isn't enough to simply know an identity has excess privilege, but to understand why that privilege is an anomaly - without having to spend time digging for that context.

Otherwise, permissions gradually creep over time, leaving accounts with far more access than their roles require. And when that access is exploited, the same lack of visibility that led to the problem in the first place will leave organizations unable to see that there's been a compromise.

Identity Challenges Ranked

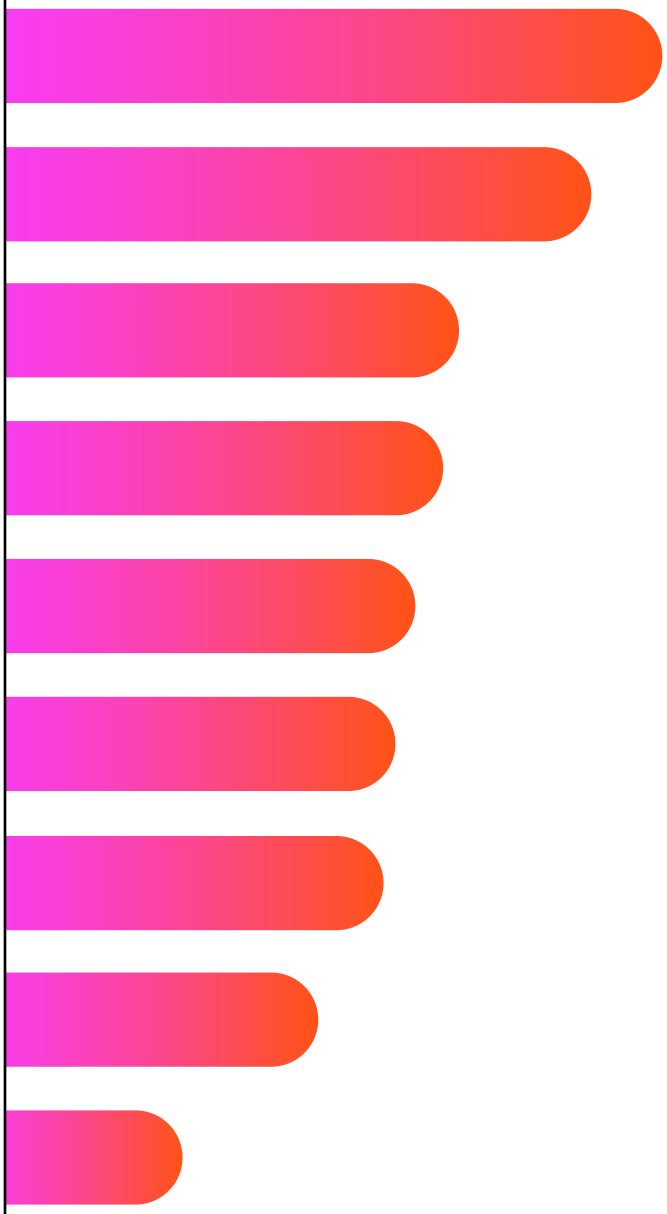

54.9%	Excessive privilege accumulation
48.1%	Inability to detect identity misuse in real time
34.6%	Gaps in access reviews/ stale entitlements
33.1%	Machine identity proliferation
30.1%	Manual provisioning/ deprovisioning delays
28.6%	Integration complexity across cloud + on-prem
27.1%	Shadow SaaS adoption
21.1%	Too many directories/ fragmented identity stores
7.5%	Lack of clear ownership between IT, Security, and HR

The Real-World Threat Landscape

One of the most significant hurdles in modern security is the realization that identity risk isn't a single, monolithic problem to be solved with a single tool. Instead, it is a plethora of interconnected issues that require a comprehensive strategy for both visibility and posture management. These aren't theoretical concerns; they are the exact vulnerabilities being weaponized in the wild today.

Our research found that nearly 65% of organizations have experienced 3 or more security incidents or near misses as a consequence of the challenges posed by identity and access management.

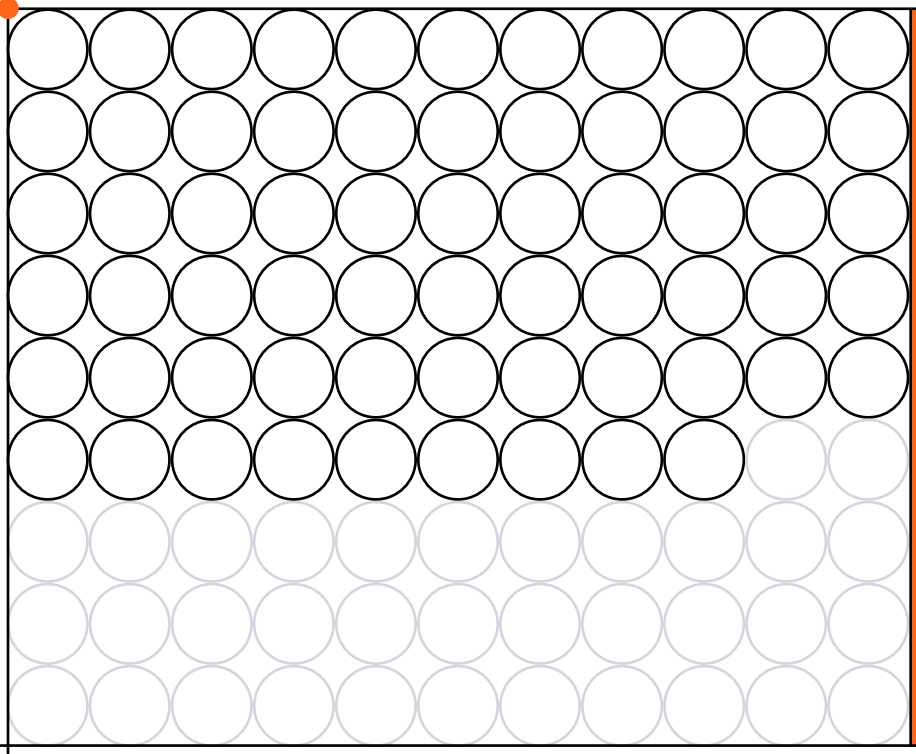

65%

of organizations have experienced 3 or more security incidents or near misses as a consequence of the challenges posed by identity and access management.

Unfortunately, there is no "silver bullet" to solve these issues, because the threat is not a single point of failure. Rather, there are a multitude of vectors driving identity threats, including MFA fatigue attacks, stolen or reused credentials, insider access misuse, and dormant account exploitation.

Identity-Driven Threats Experienced by Organizations

Dormant access exploitation

MFA fatigue attacks

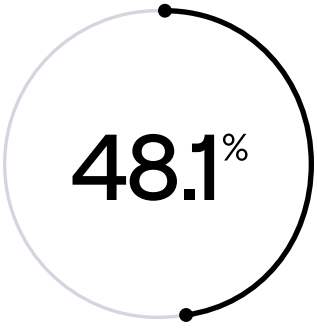

Insider Access Misuse

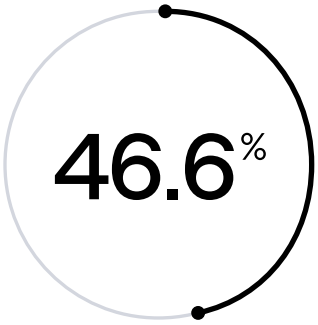

Stolen or reused credentials

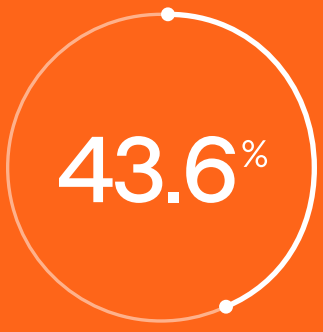

Service account abuse

Lateral movement via identity compromise

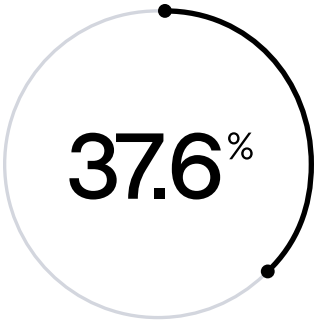

These concerns are driven by real-world threats and existing abuse, and the groups **Scattered Spider** and **Volt Typhoon** may be some of the best examples.

Threat Profile: Scattered Spider

The Modern Adversary Weaponizing Identity

Scattered Spider (also tracked as **UNC3944**, **Octo Tempest**, and **LUCR-3**) represents a shift in the cybercrime landscape, moving away from complex technical exploits toward the aggressive manipulation of human and identity systems.

Active since May 2022, this collective of young, native English-speaking individuals has successfully breached large enterprises in finance, hospitality, and retail, resulting in losses totaling hundreds of millions of dollars.

Scattered Spider's success relies on techniques like MFA fatigue attacks, dormant access abuse, and lateral movement through identity compromise.

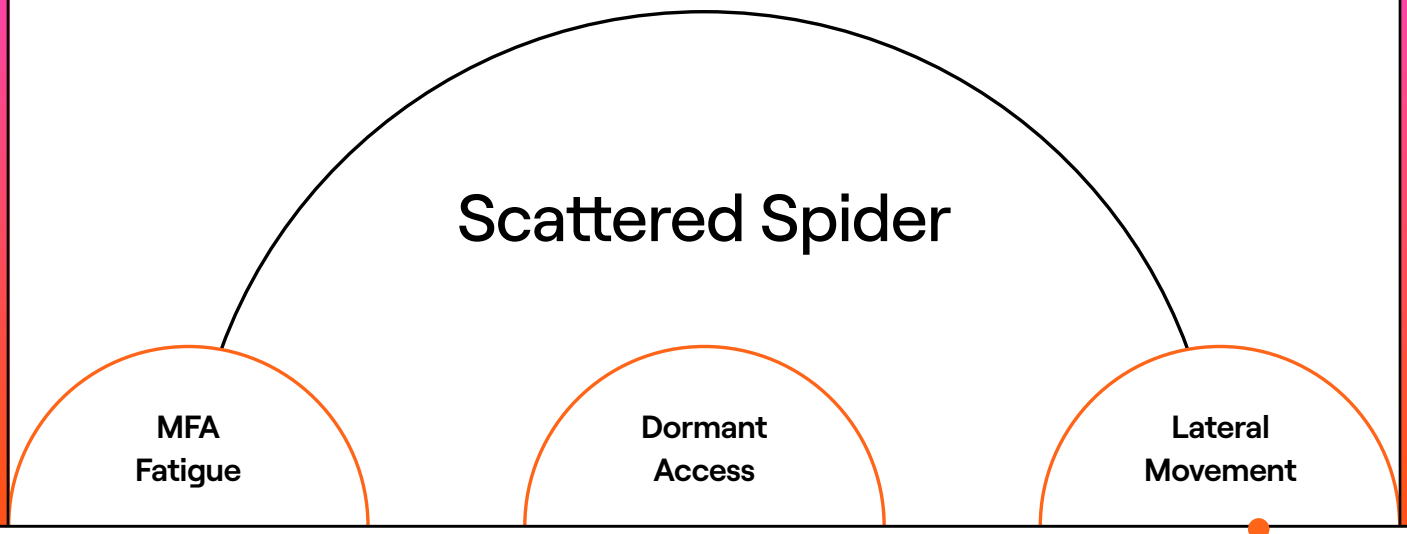

Identity-Centric Attack Vector

While many threat actors rely on software vulnerabilities, Scattered Spider's core strength is **social engineering mastery**. Their primary method of initial access involves:

Vishing (Voice Phishing)

Attackers call IT help desks impersonating high-value employees, such as the CFO, to manipulate staff into resetting passwords.

MFA Bypass

They specialize in re-enrolling Multi-Factor Authentication (MFA) to an attacker-controlled device or using "MFA push fatigue" to overwhelm users into granting access.

Real-Time Adaptation

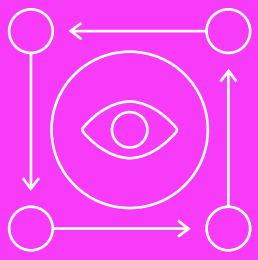

The group monitors internal corporate communications (e.g., Slack or Microsoft Teams) once inside to watch security teams' responses and adjust their tactics in real-time.

Living off the Land (LOTL)

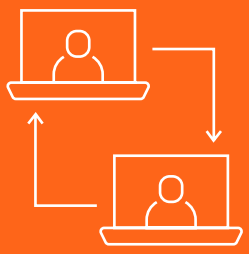

They utilize legitimate security tools, remote monitoring software (RMM), and endpoint detection (EDR) tools for persistence and lateral movement.

Rapid Operational Tempo and Pivoting

Once initial access is achieved, the group demonstrates a "fast and furious" operational tempo, often requiring defenders to respond in under **48 hours** before data encryption or exfiltration occurs.

Infrastructure Shifting

To evade detection, the group routinely changes their infrastructure and domain patterns every one to two months.

Multi-Lateral Pivots

They strategically shift alliances between major ransomware affiliates like ALPHV/BlackCat and Qilin to maximize their extortion efforts.

Threat Profile: Volt Typhoon

The Stealth Adversary Exploiting Trusted Identity

Active since at least 2021, this China aligned threat group has targeted enterprises and critical infrastructure environments by avoiding malware entirely and instead abusing legitimate credentials, trusted tools, and existing access paths.

Volt Typhoon represents a distinct evolution in identity based threats, prioritizing stealth and long term access over speed or immediate monetization.

Public reporting and joint government advisories have attributed Volt Typhoon's activity to actors associated with the People's Republic of China. Volt Typhoon's operations demonstrate how identity compromise can enable persistent, low noise intrusion in environments that otherwise appear secure.

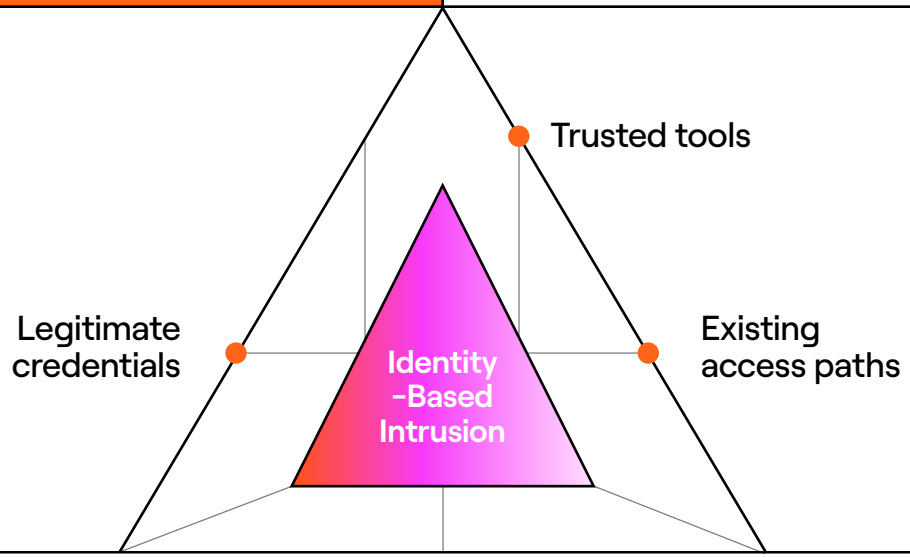

Identity-Centric Attack Vector

Unlike financially motivated groups, Volt Typhoon does not rely on phishing campaigns or overt user manipulation. Its operations focus on exploiting weaknesses in identity hygiene and access governance.

Initial access and persistence commonly involve the use of stolen or weakly protected credentials to authenticate as legitimate users or administrators. Once inside, the group relies heavily on built in operating system tools and administrative utilities to blend into normal activity. Excessive privilege and broadly scoped access allow attackers to move laterally without triggering suspicion or alerts.

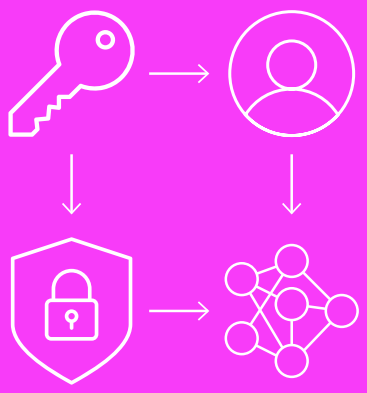

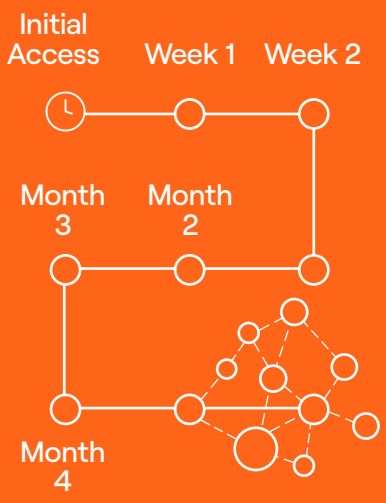

Slow and Persistent Operational Model

Volt Typhoon operates with patience rather than urgency, emphasizing persistence and resilience over rapid impact. Lateral movement is gradual and deliberate, often occurring over weeks or months.

Limited identity telemetry and insufficient real time monitoring enable long dwell times, while the absence of custom malware leaves few forensic artifacts for defenders to investigate.

Volt Typhoon highlights a critical reality for identity leaders. When attackers use valid credentials and trusted tools, traditional detection approaches struggle to distinguish malicious behavior from legitimate operations. Defending against this class of threat requires strong identity governance, least privilege enforcement, and continuous visibility into how access is actually used rather than how it is merely configured.

Machine Identity Risk is on the Rise

While human identities have historically been the focus of IAM programs, data shows that we have reached a tipping point where the human is rapidly becoming a secondary identity actor. The landscape of how our organizations are structured has shifted. Moves towards distributed, diverse workforces have led to account and tool sprawl, and the adoption of AI agents, SaaS platforms, cloud systems, and automated processes have created a growing pool of identities that aren't attached to any singular human in the workforce.

The rapid shift toward cloud-native architectures, microservices, and automated workflows has led to an explosion of machine identities, including service accounts, bots, APIs, and automated scripts.

A New Scale of Sprawl

The sheer volume of NHIs has fundamentally outpaced human growth, creating a massive management challenge.

Machine identities are now appearing at a staggering ratio, with 45% of orgs reporting that they have at least 1 nonhuman identity for every 3 humans. But some organizations reported blowing past half of the identities in their environment being NHI, reporting **10 or more non-human identities for every 1 human identity**, and in rare cases 20 or more.

This sprawl has made managing NHIs an essential component of modern identity governance.

45% of orgs report having at least 1 nonhuman identity for every 3 humans.

Despite the massive volume of machine identities, organizations feel a sense of confidence in their current capabilities. And yet, our research revealed another paradoxical gap between organizational assessment of maturity and the realities of identity security.

High Perceived Maturity: 73.7% of respondents reported feeling confident about their ability to inventory their NHI, with nearly half claiming to be "absolutely confident". An even higher 78.2% report confidence in their ability to govern and secure NHI, and 43.6% of those claiming the highest confidence. This suggests that many leaders view machine identities as a high-maturity area where they feel they have sufficient control, even as the sprawl continues to accelerate.

of respondents reported feeling confident about their ability to inventory their NHI

report confidence in their ability to govern and secure NHI

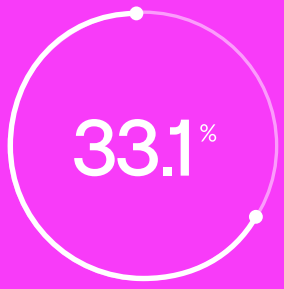

report NHI proliferation as one of their most pressing identity risks

However, this confidence may be misplaced when it comes to long-term governance and automation. Regardless of control of machine identities, human identities still currently pose a greater risk, and as NHI continues to proliferate those identities can't be left as an exposure point.

NHI proliferation has already become a material risk for 33.1% of organizations, and as non-human identities continue to outpace human users, the resulting access sprawl will accelerate faster than most security teams leveraging current frameworks and tooling can respond.Platforms that can handle the future scale and complexity of identity will be important, but even more vital is to implement tooling to handle today's realities in order to prepare for tomorrow.

AI is the Answer to both Risk and Operational Efficiency

As threats become more sophisticated and rapid, time to incident resolution stretches, and machines outpace human identities, manual oversight of identity governance is a less viable strategy with every passing day. Very few organizations have fully embraced automated workflows and are hindered by psychological and technical barriers to adopting AI. Leaders expressed feeling that they must get control of their data and governance before they can utilize AI, but the transition from manual workflows to autonomous, agentic identity requires different thinking. Moving into fully AI-assisted operations and automated governance starts small with AI-based intelligence and assessments, leveraging agentic capabilities to begin to automate, and finally transforming identity management with AI to deliver governance, enhance security, and drive productivity that evolves as an organization does.

Identity Scale & Complexity

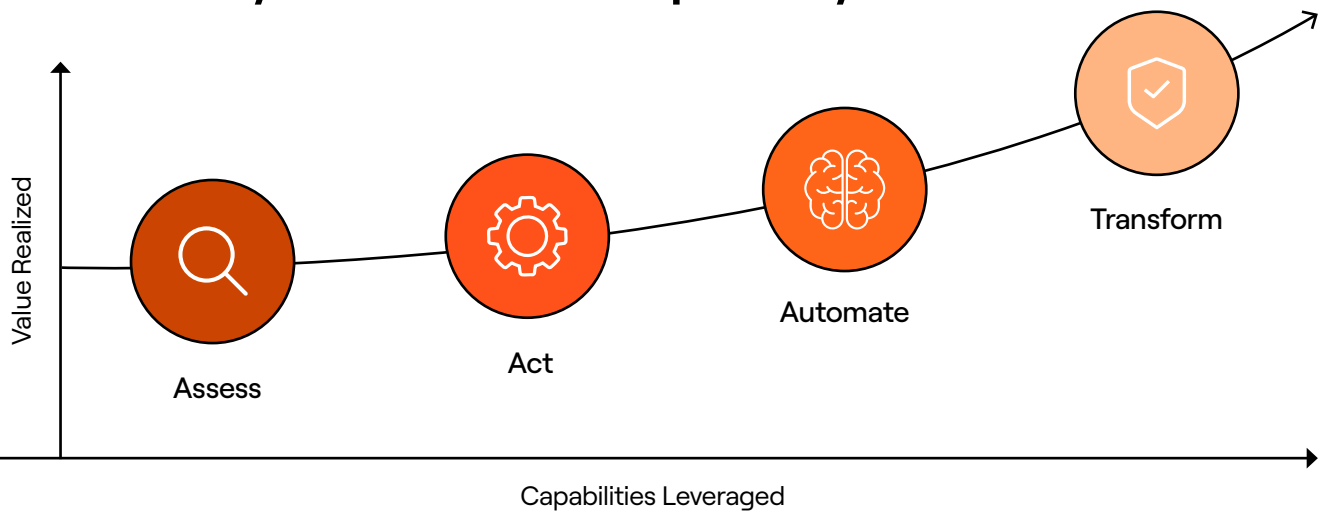

The current landscape

The growing technical debt and rapidly increasing number of identities to be managed have made manual management impossible. AI is becoming a fundamental necessity for organizations to close the gap between their current posture and a resilient security state. Administrative complexity in implementing mature identity programs and the risk involved in getting it wrong require systems that can evaluate and recommend action at scale.

Operational Priorities and Automation

The pressing nature of identity risk is forcing organizations to pivot toward several key pillars for improving their identity and access management programs:

Automating User Access Reviews

Moving away from "rubber-stamp" spreadsheets toward intelligent, automated review cycles.

Least-Privilege & Zero Trust

Implementing automated policies that ensure users and machines have only the access they need, exactly when they need it.

Governance & Analytics

Enhancing compliance reporting and improving identity analytics to spot anomalies that a human eye would miss.

Optimizing Velocity & Detection

Shrink critical windows of exposure, specifically improving **MTTP (Mean Time to Provision)** across apps and **MTTD (Mean Time to Detect)** identity misuse.

Organizational priorities for improvement

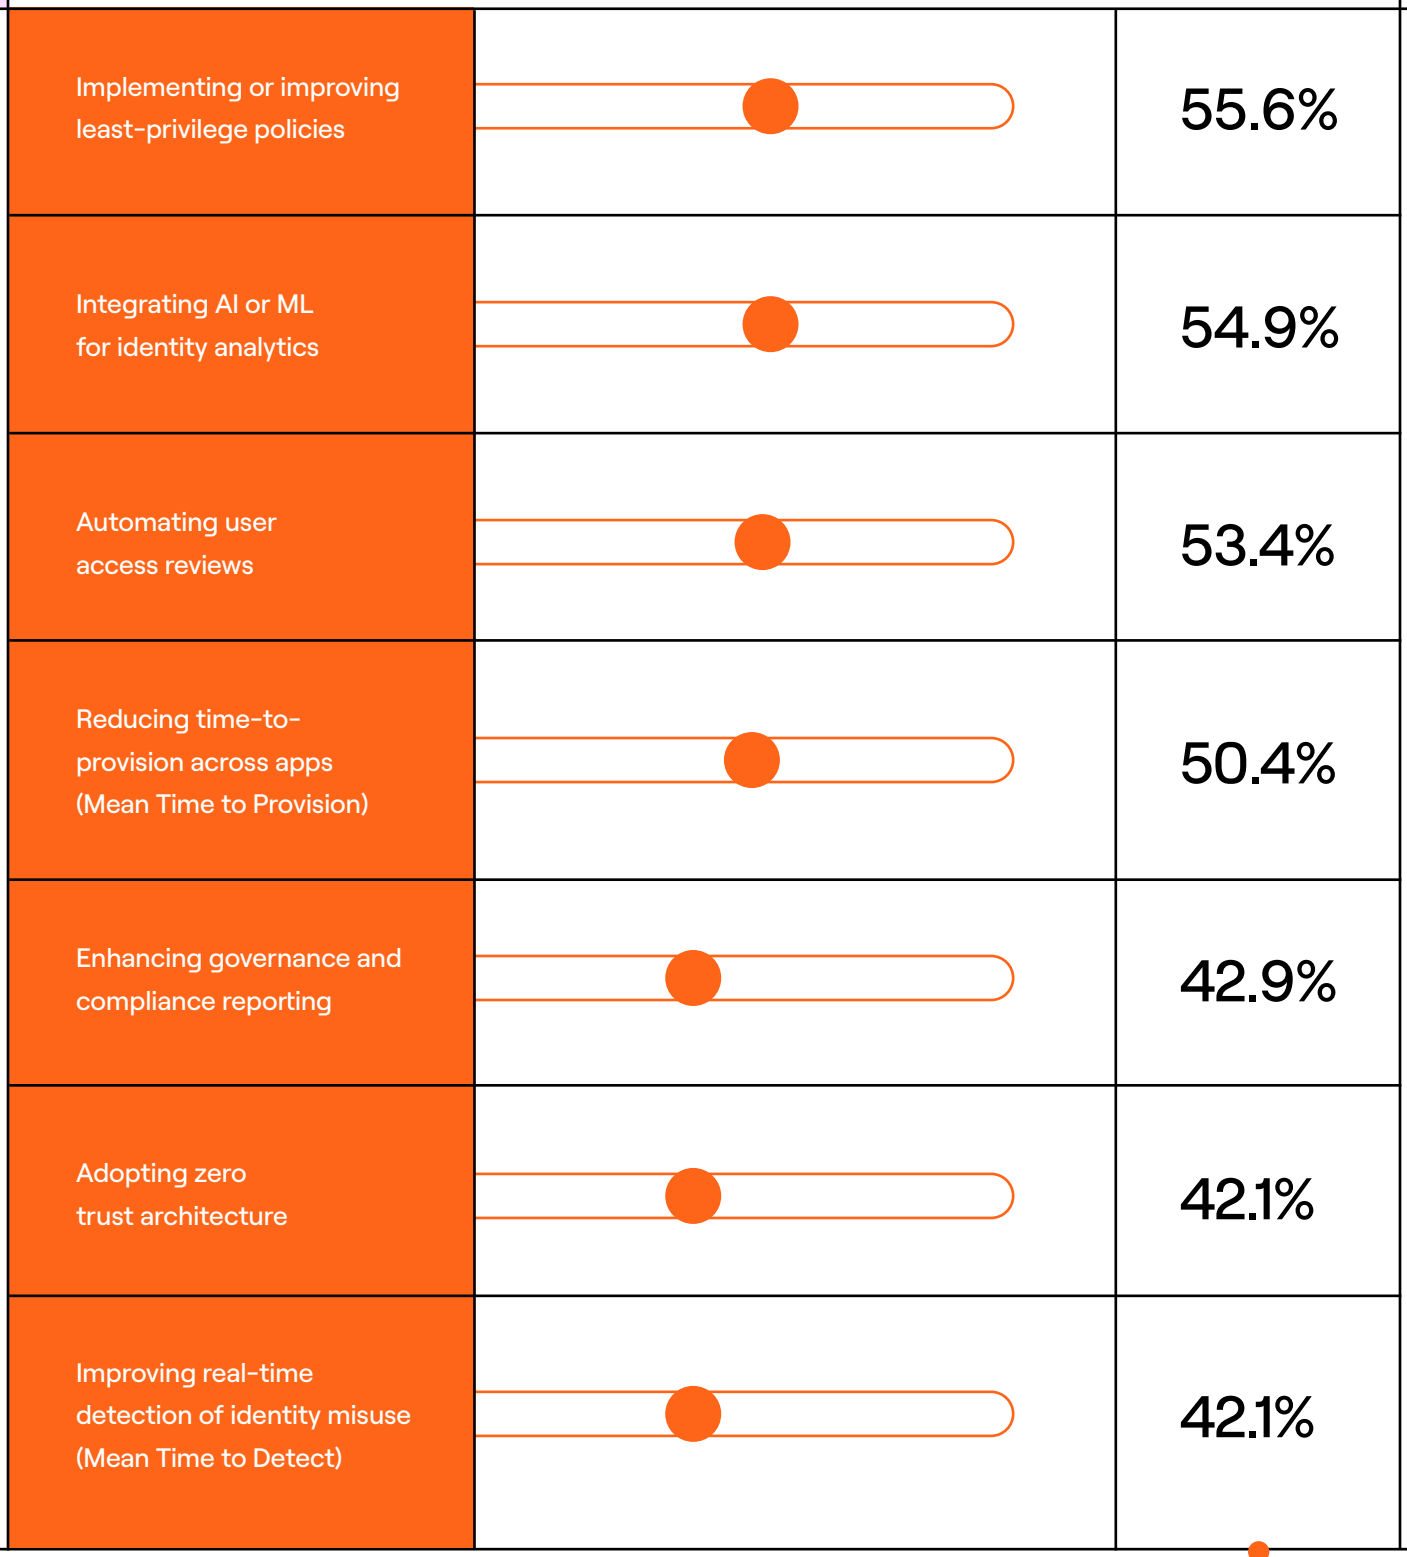

Despite hesitancy to fully embrace AI and automation, there is a large overlap between where organizations believe they would see the most benefit from AI and where they are prioritizing improvement.

Threat Detection and Risk Scoring

Leveraging AI to detect anomalies in access and reduce risk through recommendations and automated actions.

Audit and Compliance Reporting

Improving governance and analytics with full documentation for audits and compliance.

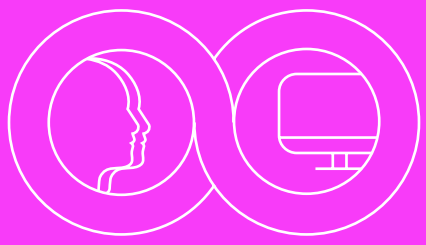

Identity Lifecycle Management for Human and Non Human Identities

Removing work for the Joiner/Mover/Leaver lifecycle and keeping up with NHI sprawl through automated provisioning and controls.

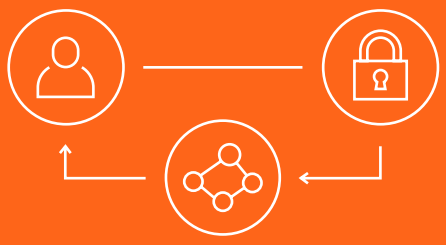

Automating Role-based Access and Reducing Privilege Risk

Using AI for broad analysis of organizational structure for Role Mining and creating Role-Based Access Controls to reduce privilege creep.

Where orgs would see the most benefit from AI-driven automation

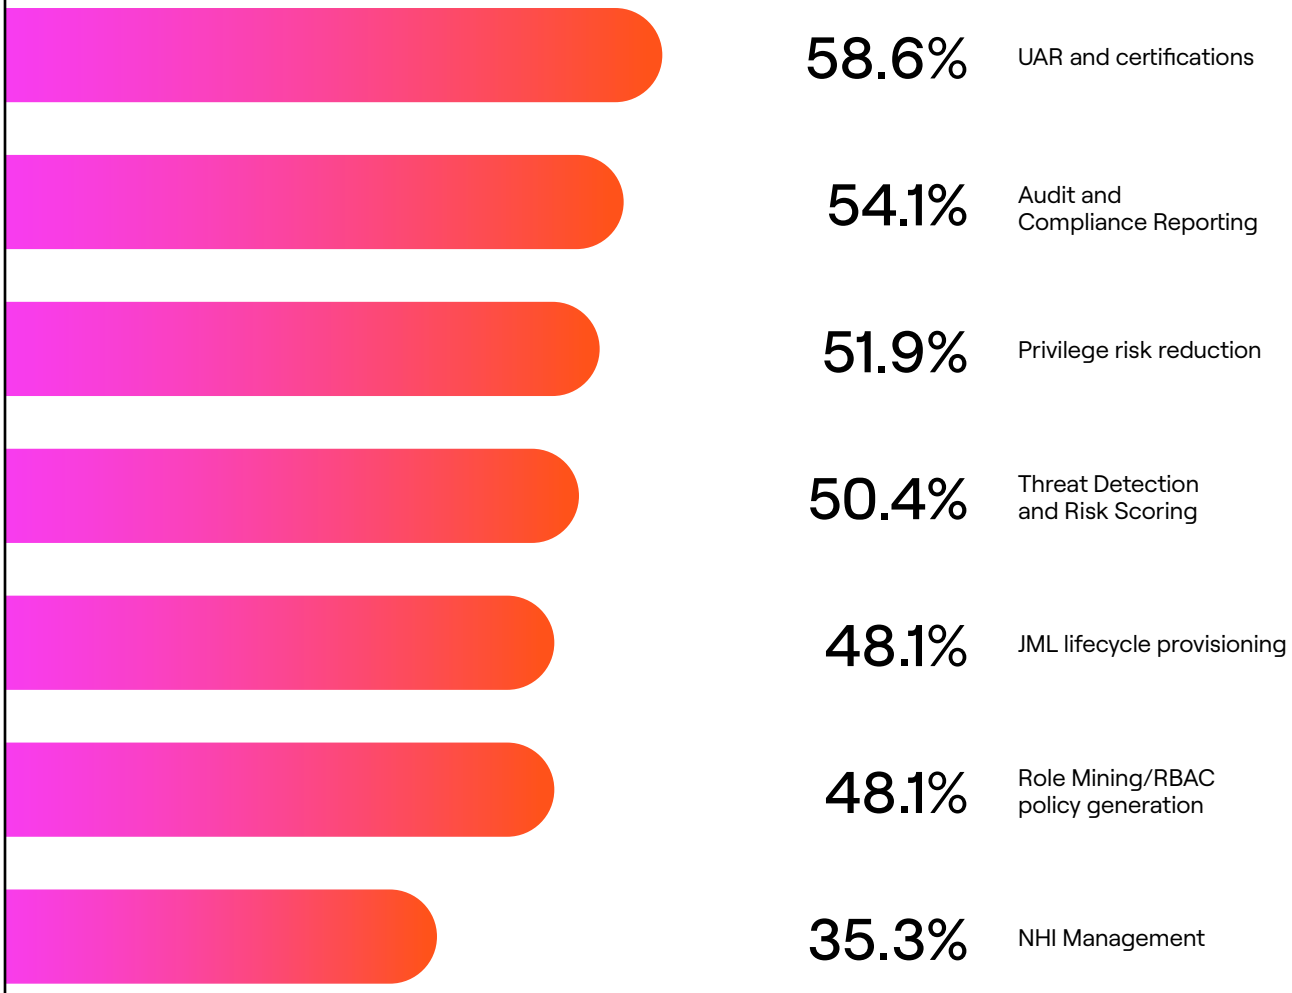

Why Organizations Stall

It's clear that AI is one of the most important and top-of-mind emerging technologies for identity leaders, with 88.7% rating it as important or very important to their detection and response efforts over the next 2 years. This aligns with current adoption trends, with 85% of organizations leveraging AI in some capacity in their identity governance processes. However, the majority (68.4%) of organizations are only using AI in narrow use cases, or not at all, and very few organizations, just 4.5%, allow AI to make limited decisions.

4.5%

of organizations allow AI to make limited decisions

AI usage in Identity

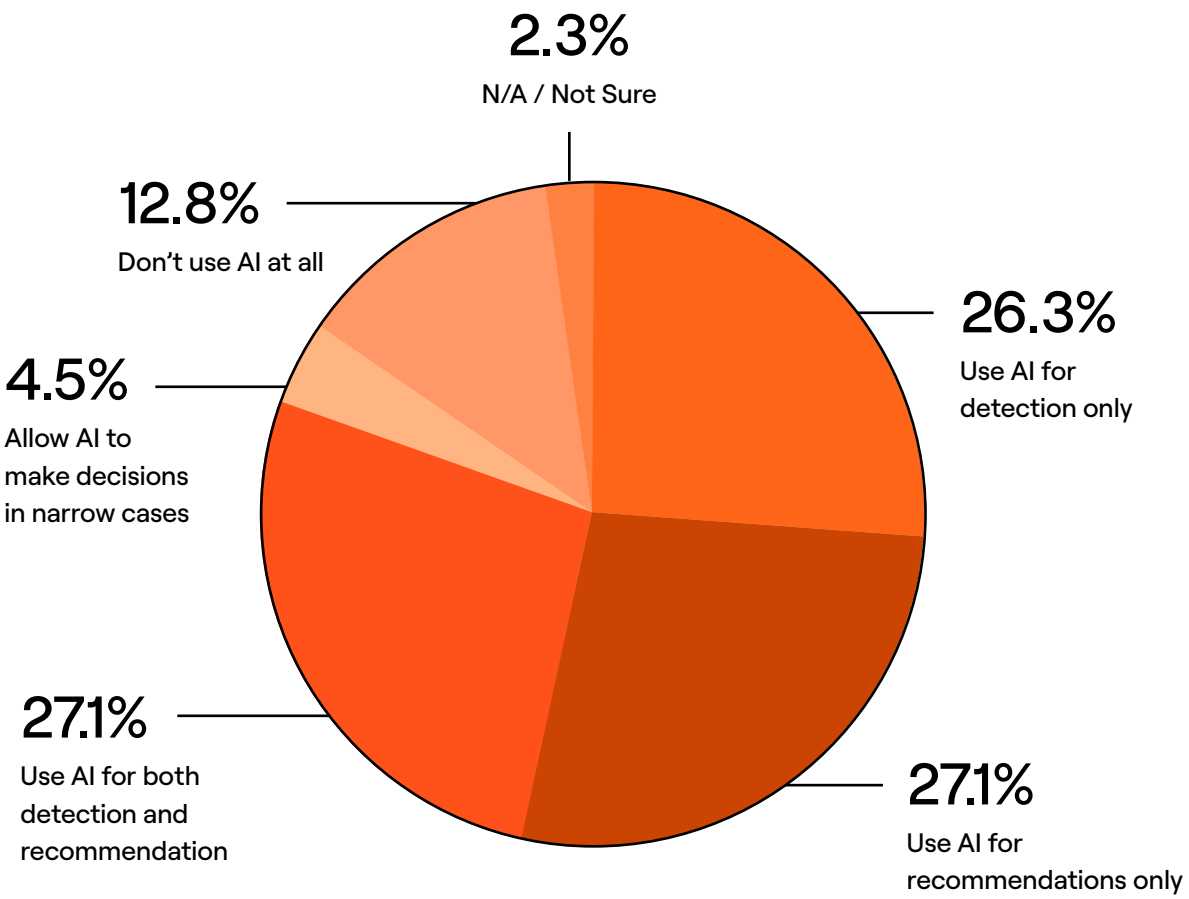

of orgs believe AI will be vital to their efforts over the next two years

The overlap between where organizations are looking to improve and where they believe AI would be the most helpful is clear, and if 88.7% of orgs believe AI will be vital to their efforts over the next two years, then why are so many struggling to implement it fully?

Despite the high level of interest, there is a significant gap between the desire to use AI and its effective application. While the intent is high, the data shows that more than 1/3 of organizations are currently using AI to its full potential within their security stack. Although the benefits of automation like speed, scope, and risk reduction are theoretically obvious, many leaders lack the confidence to initiate full-scale deployment.

Why Organization's Aren't Using AI Yet

47.1%
Distrust in automated results

There is a fundamental skepticism regarding whether an automated system will make the right decision without human intervention.

41.2%
Insufficient auditability

Security teams are wary of black-box processes; without clear trails showing why an access decision was made, they cannot meet strict compliance and internal governance standards.

45.9%
Data quality and schema

Technical debt and messy data within outdated systems make the implementation of modern automation feel like an insurmountable engineering challenge.

52.6%
Skill gaps and lack of expertise

The existing gaps that have plagued the cybersecurity industry for years, but also the gap in expertise around new technologies and standards.

Incremental Adoption Reduces Concern and Increases Security

Crucially, adopting AI does not require a high-stakes, all-or-nothing leap into fully autonomous decision-making. The transition to AI-driven security is a progression, not a binary choice. For many organizations, the immediate value lies in intelligent assessment—leveraging AI to gain the visibility, reasoning and the right context required to understand the "why" behind a risk, even before taking an automated action. This is particularly vital given that while 88.7% of leaders view AI as a necessity, only 4.5% currently trust it to make limited decisions.

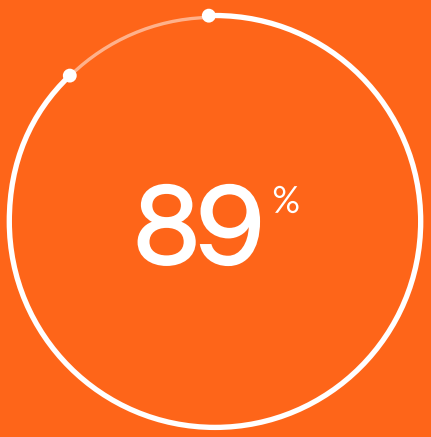

of leaders view AI as a necessity

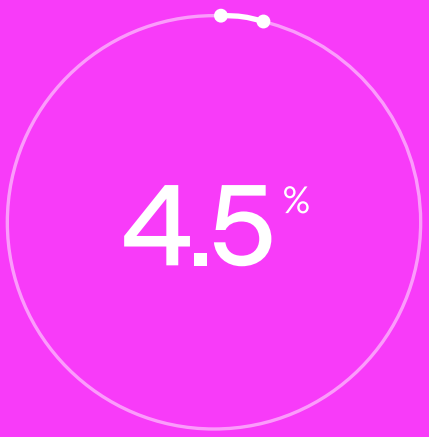

currently trust it to make limited decisions

By starting with a single, high-impact use case, such as automating stagnant User Access Reviews or identifying high-risk dormant accounts, teams can validate the reasoning of these systems and build the trust necessary to move toward broader automation. An incremental adoption allows leaders to assuage black box fears while immediately reducing the exposure that leaves their organizations vulnerable.

Next Steps for Identity Leaders

To bridge the gap between perception and reality, C-level leaders and identity practitioners must move away from static, tool-based security toward a dynamic lifecycle of identity protection. Combating modern identity threats requires prioritizing the following four strategic pillars:

1. Visibility: Establish the Source of Truth

True security is impossible without comprehensive visibility into the "identity sprawl" currently plaguing the enterprise.

Develop a Single Source of Truth

Adopt systems and integrations that create a unified source or truth for decision making based on high-fidelity data.

Continuous Monitoring

Move beyond point-in-time audits to continuous monitoring of both human and machine identities, ensuring that lateral movement can be detected before data exfiltration begins.

2. Intelligence: AI-Powered Risk Reasoning

With 85% of leaders viewing AI as a necessity to manage administrative burdens, organizations must move from detection-only to AI-driven reasoning.

Risk Posture Analysis

Utilize AI to analyze broad organizational structures, performing automated "role mining" to identify and reduce excessive privilege accumulation.

Reasoning Access Decisions

Implement systems that can reason out access decisions with clear audit trails.

Anomaly Detection

Leverage machine learning to identify behavioral anomalies—such as MFA fatigue attacks or credential theft—that a human eye would miss.

3. Action: Real-Time Detection and Remediation

Visibility and intelligence must culminate in rapid, automated action to shrink windows of exposure.

Improve MTTP and MTTD

Prioritize lowering the Mean Time to Provision (MTTP) for legitimate access and the Mean Time to Detect (MTTD) for unauthorized misuse.

Rapid Revocation

Automate the revocation of dormant access to ensure that employees who are already gone no longer have active credentials.

Just-in-Time (JIT) Access

Replace permanent, standing privileges with automated policies that grant users and machines access only when needed, exactly for the time required.

4. Governance: UAR Automation and LCM

Governance must shift from a manual compliance box-checking exercise to a resilient operational standard.

UAR Automation

Replace "rubber-stamp" spreadsheets with intelligent, automated UARs that prioritize risky access and identify stale entitlements before they can be exploited.

Full Lifecycle Management (LCM)

Automate the entire Joiner, Mover, Leaver lifecycle for both humans and machine identities to prevent permission creep over time.

The transition from manual identity management to a resilient, automated state is no longer an optional upgrade, but a fundamental necessity for organizational survival. With 85% of leaders acknowledging that AI is required to manage the overwhelming administrative burden of modern identity programs, it is clear that manual oversight has reached a breaking point. The sheer scale of identity sprawl—where machine identities can outnumber humans by as much as 20:1—creates a level of complexity that requires systems capable of evaluating risk and recommending action at machine speed. While psychological barriers like distrust in automated results persist, the alternative is a continued reliance on "messy data" and manual processes that leave 81% of organizations vulnerable to identity-related incidents. By leveraging AI for anomaly detection, automated role mining, and real-time risk scoring, organizations can finally close the "confidence gap" and shift their identity programs from a source of technical debt into a proactive, measurable business outcome.

The Lumos Autonomous Identity Platform

Complete Visibility

Bring together all identity data from HRIS, IDPs, SaaS, and cloud into a single access graph with fine-grained visibility across users, apps, and entitlements.

Contextual Intelligence

Surface real-time insights from identifying overprovisioned accounts to anomalous access and policy gaps. Continuously improve access decisions, reduce risk, optimize spend, and drive data-informed governance decisions.

Agentic Workflows

Albus acts as an intelligent identity agent that understands natural language, explains its reasoning, and autonomously runs access policies, requests, approvals, reviews, and reporting.

Intuitive UX

For both end users and identity teams, Lumos seamlessly integrates into familiar workflows for faster, more informed requests and decision making.

"Identity has become the engine that powers security and productivity, and thriving in the agentic AI era requires rethinking how you leverage identity to secure and power your organization. Every organization needs to start the path to autonomous identity, but that journey can be taken one step at a time. Lumos is excited to be part of that journey to help customers drive that with confidence."

Andrej Safundzic, Co-Founder & CEO

Lumos Products

Lumos Autonomous Identity Platform

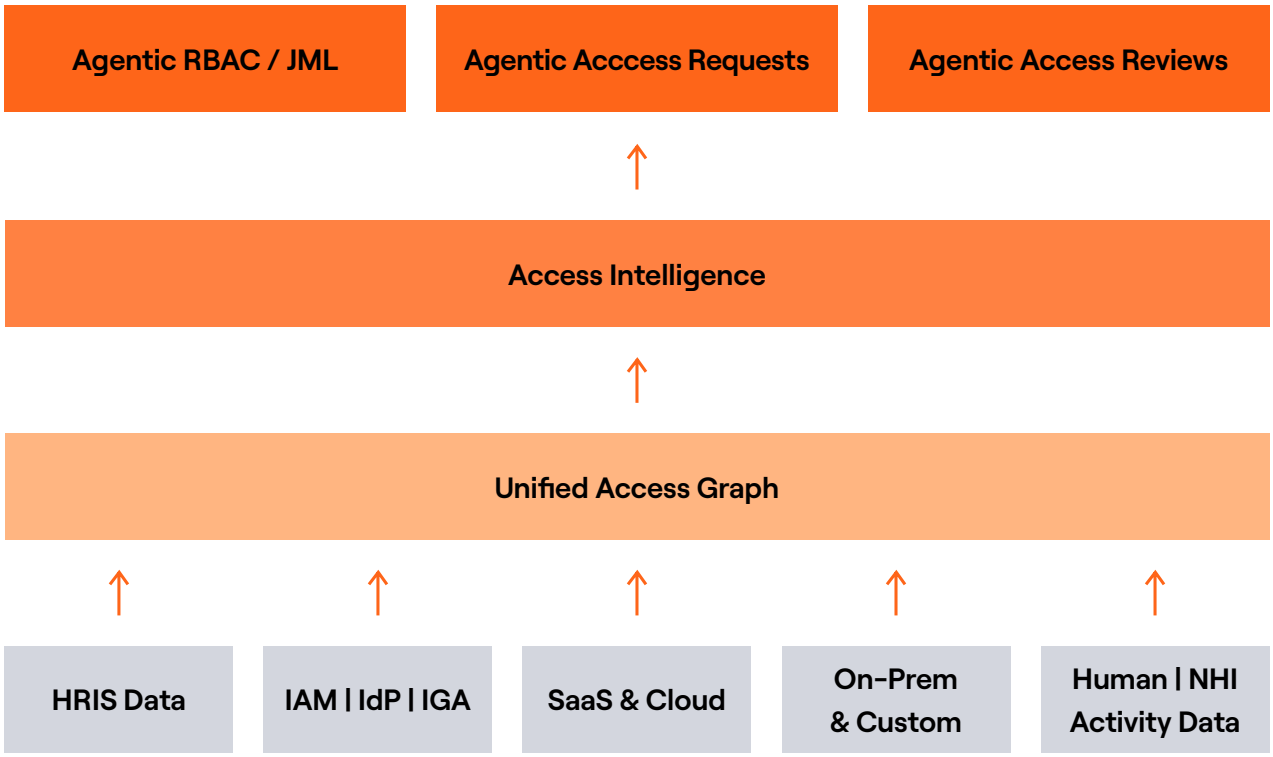

Albus, the Identity AI Agent

Easily build efficient policies, detect anomalies, and reason through complex identity problems with Agentic AI.

Learn More

User Access Reviews

Automate reviews, cut cycles, and pass audits with confidence. No spreadsheets, no blind spots, no wasted time.

Learn More

AppStore

Cut wait time to minutes and create great user experiences with the right access at the right time.

Learn More

Lifecycle Management

Right-size access without the ticket chaos and automate joiner-mover-leaver workflows to boost productivity, reduce risk, and cut busywork across your apps.

Learn More

Identity Risk Intelligence

Elevate your identity hygiene and posture analysis with interactive, customizable dashboards.

Learn More

Lumos is the first Autonomous Identity platform that empowers organizations with an agentic approach to enhance security, drive productivity, and meet compliance standards. The company's AI-native platform automatically discovers and manages access across customers' entire tool stack.

By delivering rich, contextual intelligence, organizations can move beyond rubber stamping to proactively mitigate identity-based risk, all without slowing their business down. Powered by agentic workflows in Albus, the company's leading identity agent, organizations can address security concerns and establish governance that evolve as quickly as your business does.

Trusted by hundreds of companies including Pinterest, Anduril, and GitHub, Lumos powers millions of access requests across global companies.

Request a Demo