veza

Report

State of Identity & Access

HUMAN_IDENTITY

NON_HUMAN_IDENTITY

AI_AGENT

2026

Table of Contents

Foreword

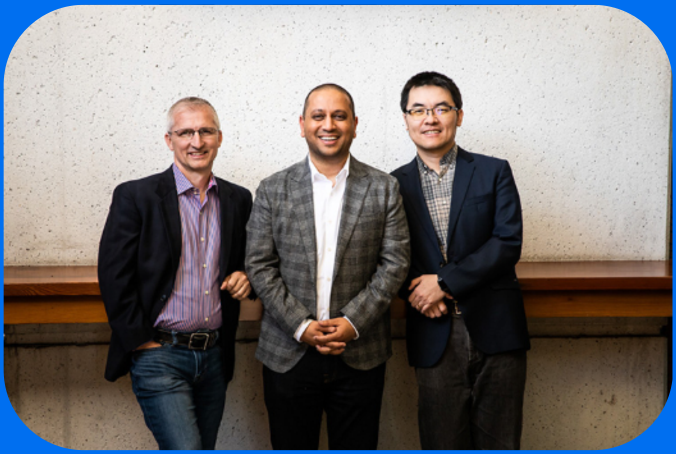

Rob Whitcher | Co-founder & Chief Architect
Tarun Thakur | Co-founder & CEO
Maohua Lu | Co-founder & CTO

You've heard the adage: everything in moderation. Too much of a good thing becomes a bad thing, and access is no exception. The right access accelerates productivity, fuels automation, and unlocks the value of AI. But the wrong access? That's where things unravel. Because access itself isn't the villain. Neither are non-human identities. Nor the rapidly growing fleet of AI agents powering today's enterprises.

The real villain is bad permissions: the excessive entitlements, over-privileged access, forgotten SaaS API tokens, dormant accounts, and invisible privilege creep quietly growing behind the scenes.

One would think that by now, we would have learned from our past mistakes. While the current wave of AI is undeniably transformative — and disruptive — it's not without precedent. We've seen similar seismic shifts before.

Take the move to distributed computing, which emerged out of a need for speed and agility. That transition brought its own chaos: shadow IT, unsecured systems, and a lack of oversight. Back then, it wasn't uncommon to find customer records and credit card numbers stored in plain text files under an IT admin's desk — in an open terminal, in an unlocked office.

Fortunately, we've come a long way since then. But fast forward twenty years, and we may look back on today's unchecked permission sprawl and AI-fueled automation with the same disbelief — wondering how we ever let it get this far.

If we don't move forward with intention — guided by visibility, accountability, and security by design — we are not just inviting risk; we are engineering disasters at scale. Now is the time to act, before hindsight becomes regret.

In an era where AI-driven systems can take action at machine speed, bad permissions don't just create inefficiency; they create systemic risks. Worse yet, this exponential permission sprawl bars enterprises from answering cybersecurity's most pressing question: **"Who can and should take what action on what data and what resources?"**

What follows is a look into how permission sprawl, the modern remote workforce, non-human identities (NHIs), and AI agents are shaping identity security. These insights are based on Veza's visibility into more than 230 billion permissions across 160 enterprises–all of which has also been analyzed and distilled into Veza's 2026 State of Identity and Access Report.

Introduction

Identity security has always carried its share of complexity, but recent developments in AI and the rise of NHIs have pushed that complexity into uncharted territory. Organizations now span hybrid environments that stitch together legacy directories, cloud IAMs, SaaS platforms, data systems, internal apps, and an exploding population of machine identities. In these environments, permissions multiply at an unprecedented speed.

And as they do, they surface a simple truth: **identity is the new perimeter, and the only way to secure it is to understand and control the permissions behind every identity.**

The urgency intensified last year as new technologies, regulations, and threat vectors converged. Suddenly, permissions data—often overlooked, rarely normalized—became one of the most critical foundations of enterprise security.. Not because access is dangerous, but because **bad permissions** create dangerous gaps that attackers can exploit.

The last decade of identity tooling (IGA, IAM, PAM) focused on workflows, provisioning, and role structures. But 2025 made one thing painfully clear:

The battleground isn't workflow, it's contextual authorization.

Every permission—whether in Okta, Azure AD, AWS IAM, GitHub, Snowflake, or a custom internal system—determines what an identity can do. When those permissions are misconfigured, overly broad, ungoverned, or simply forgotten, they become the villain of the story: a shadow layer of risk hiding in plain sight.

As companies embraced zero trust, hybrid work, AI-driven automation, and accelerated cloud adoption, failing to understand permissions became a strategic liability. Identity teams fought on multiple fronts: access requests, audits, privilege violations, as well as incident response, and yet blind spots persisted, and access debt continued to build.

Why? **Bad permissions grow faster than any manual process can keep up with.**

Today, permissions often number in the hundreds of millions across a single enterprise, making manual oversight impossible. The need for continuous visibility, normalization, and automation is no longer optional. So, we measured the problem.

Veza now monitors more than 230 billion permissions, providing continuous visibility and least privilege enforcement across hybrid ecosystems. Studying this data—hundreds of millions of entitlements across 160 enterprises—we uncovered the trends that defined identity and access last year and will shape 2026 ahead.

Let's dig in.

Executive Summary

The 2026 State of Identity & Access (SOIA) Report reveals a stark truth: identity security is no longer a technical discipline—it's a board-level risk. Across industries, uncontrolled growth in human and machine accounts, inconsistent MFA adoption, and unchecked permission sprawl have turned identity into the primary breach vector.

Data from analysis of millions of identities and billions of entitlements across several large enterprises show that nearly 3.8 million dormant accounts (38% of all IdP users) remain active, machine identities now outnumber humans by 17:1, and the average worker holds 96,000 entitlements.

Recent attacks underscore this reality. The Jaguar Land Rover (JLR) cyberattack, one of the most economically damaging incidents with an estimated financial impact exceeding £1.5 billion, started with a single credential compromise. Further, breaches at Change Healthcare and Colonial Pipeline didn't require advanced exploits; they leveraged inactive or weakly protected accounts. Meanwhile, threat actors such as Volt Typhoon demonstrate how credential misuse can be weaponized for long-term persistence in critical infrastructure. The United States recorded nearly half of all global ransomware incidents, at 47%.

Identity sprawl and debt—the accumulation of dormant accounts, orphaned credentials, and excessive entitlements—is now a systemic business problem. The rapid increase in Agentic AI is poised to exacerbate this with a massive influx of autonomous non-human identities that will be permissioned to execute complex tasks. Adversaries can weaponize compromised AI agent identities to execute high-velocity ransomware and data exfiltration attacks. Boards, regulators, and insurers are increasingly demanding proof of control over "who can do what, where, and when?" This report quantifies that gap while showing how far most organizations currently lie from real-time authorization maturity and least privilege at scale.

The Jaguar Land Rover (JLR) cyberattack started with **a single credential compromise.**

We are on a collision course. Identity is not only the most vulnerable— and actively targeted— entry point in the enterprise, we must now also reckon with an explosion of 'non-human identities', including AI agents to sprawling cloud apps. The volume of permissions security teams are expected to manage has jumped from millions to billions almost overnight. Veza's 2026 SOIA report lays bare the identity and access insights you simply can't afford to ignore. The trends in this data are clear and accelerating: as organizations scale and attack surfaces diversify, securing identities and non-human identities isn't just a 'best practice' anymore— it's table stakes for survival.

Nicole Perlroth

Cybersecurity Author & Journalist

About Nicole Perlroth

Nicole Perlroth spent over a decade as The New York Times' lead cybersecurity reporter, where her groundbreaking work on Chinese cyberespionage helped lead to the first U.S. hacking charges against members of the Chinese military. Her reporting on commercial spyware was nominated for the Pulitzer Prize.

Her bestselling book, This Is How They Tell Me the World Ends, an exposé on the global cyber arms race, won the FT-McKinsey Business Book of the Year Award and the Arthur Ross Foreign Policy Book of the Year Prize. It was also inducted into the Cybersecurity Canon Hall of Fame and optioned for both scripted TV and documentary film.

Since leaving The New York Times in 2021, Perlroth has served on the Department of Homeland Security's Cybersecurity and Infrastructure Security Advisory Committee (CISAC), launched the cyber moonshot fund Silver Buckshot Ventures, and is a Venture Partner at Ballistic Ventures.

Methodology

The SOIA 2026 Report combines proprietary analysis on millions of identities from hundreds of large enterprises across various industries, such as financial services, healthcare, technology, retail, and the public sector.

Sources and Data Scope

Enterprise Dataset: Aggregated analysis of millions of identities and billions of entitlements across identity providers, cloud platforms, SaaS applications, and data systems.

Timeframe: Data observed through 2025.

Threat Intelligence Validation: Findings aligned with external reports, including *Verizon's 2025 Data Breach Investigations Report, CrowdStrike's 2025 Global Threat Report, and Expel's 2025 Incident Response Data.*

Government and Advisory Context: Key insights referenced from CISA, NSA, and FBI advisories on credential-based attacks, with emphasis on nation-state campaigns like Volt Typhoon.

Data Normalization

Metrics were normalized across tenants using median and interquartile range (IQR) to reduce outlier bias. Definitions were standardized as follows:

Dormant accounts: Accounts that are inactive ≥ 90 days.

Orphaned accounts: Identities that remain active with no legitimate owners (for example, when an employee leaves the organization).

Machine identities: Non-human identities (NHIs) like service accounts, API keys, tokens, secrets, certificates, or other automation credentials.

Effective permissions: Entitlements that result in active create, read, write, or delete capability on a resource, these *Create, Read, Write and Delete* capabilities are effective permissions.

All figures represent aggregated anonymized data.

Introduction

Identity security is the #1 enterprise risk. According to the Verizon 2025 Data Breach Investigations Report, identity is now the primary vector in most ransomware breaches. Veza's 2026 State of Identity & Access (SOIA) Report provides a comprehensive view of how uncontrolled identities are fueling today's most damaging breaches.

What we found is clear:

Dormant accounts represent a large fraction (38%) of all accounts, providing potential backdoors for adversaries.

Machine identities (NHIs) continue to grow faster than human identities, multiplying the potential attack surface at a 17:1 ratio over humans.

Permission sprawl has reached historic levels, with the average worker holding tens of thousands of entitlements, most of which are unmonitored.

NHIs often lack MFA, and have no defined owner - making them easy for adversaries to steal and hard for enterprises to govern.

In both the Change Healthcare ransomware incident and the Colonial Pipeline shutdown, adversaries gained initial access by compromising credentials on systems that lacked strong MFA or other identity controls. Once inside, they moved laterally, escalated privileges, exfiltrated data, and deployed ransomware, forcing operational shutdowns and widespread disruption.

This report reframes identity debt (the buildup of dormant accounts, unchecked non-human identities, and unused entitlements) as a systemic business threat on par with financial debt, supply chain failure, or regulatory non-compliance. The data underscores an urgent need to prioritize identity-led controls, continuous validation, and remediation before exposure compounds into enterprise-wide risk.

Major threat reports (below) in the last six months flag identity as the top attack vector. Our findings show why: dormant accounts, entitlement sprawl, weak MFA, and help-desk abuse create simple, repeatable paths that adversaries exploit.

Major threat reports

CrowdStrike: In its 2025 Global Threat Report, CrowdStrike reported that 79% of attacks were malware-free and emphasized that "identity is the new battleground."

Identity Defined Security Alliance (IDSA): According to their 2025 Trends in Identity Security report, two-thirds of respondents have designated Non-Human Identity (NHI) management as a top 5 priority in their security plans, recognizing that machine identities now significantly outnumber human identities.

Expel: In its 2025 annual report, Expel found that 68% of all security incidents investigated were identity-based, with compromised credentials and misused access as top vectors.

MITRE: Based on real-world adversary behaviours, MITRE ATT&CK data shows that over 50% of observed attack techniques target identity, including privilege escalation, credential access, and lateral movement tactics.

Cisco Talos: In its 2024 Year in Review, Cisco Talos reported that identity-based attacks accounted for 60% of all incident response cases. These attacks frequently involved the misuse of valid credentials and targeted systems like Active Directory and cloud APIs. Additionally, ransomware actors leveraged valid accounts for initial access in nearly 70% of cases.

The Identity Security Reality Check

Category	2025 Snapshot	Year-Over-Year Change	Why It Matters
Dormant Accounts	~3.8 million accounts (38% of all IdP users) active after 90+ days of inactivity resulting from ineffective lifecycle management and timely deprovisioning	↑ from 730K (20% of all IdP Users) (2024)	Every dormant identity is an open login for adversaries.
Orphaned IdP Identities	824,000 accounts (8% of all IdP users) due to improper or incomplete deprovisioning of these accounts that are no longer associated with a human owner	↑ 40% YoY	No owner, no accountability, becoming easy targets for misuse by adversaries.
Machine Identities (NHIs)	17:1 ratio vs. human users	↑ from ~16:1 (2024)	Non-human accounts continue to dominate the identity landscape.
Permissions Sprawl	230 billion total permissions	NA	Vast entitlement noise obscures real access risk.
Average Human Entitlements	96,000 permissions per worker	NA	Least privilege remains aspirational, not operational.
MFA Gaps	13% of enterprise users lack MFA	No significant change YoY	MFA deployments continue to have significant gaps and deployment has plateaued.
Ex-Employees with Active Access	78,000 users still retain credentials	↓ from 92K (2024)	Deprovisioning lags continue to fuel insider risk.
Over-Permissioned Users	16.5% of total permissions belong to inactive users	↑ from 12% (2024)	Permission creep compounds dormant account exposure.
Application Concentration	A small number of large service providers control a majority of all enterprise permissions	Stable	A few large service providers such as Microsoft, Okta, AWS, define the enterprise risk surface.

Demystifying Bad Permissions

When we examine bad permissions, we categorize them into four types:

Over-privileged

1

Where an identity is granted more access than necessary to do a job.

Residual

2

Where access should've been revoked after a termination, job change, or completion of a temporary task.

Ungoverned

3

Where local users are created outside the purview of identity tools like SSO and IGA.

Policy-violating

4

Where security rules are not being followed due to lack of visibility (e.g. separation of duties, requirement for MFA)

Bad permissions are pervasive. How many do you have? It depends on your policies and level of regulatory oversight, but it's not uncommon to see thousands of troubling permissions. Turning on Veza is a bit like opening the hood on a car that's never had its oil changed.

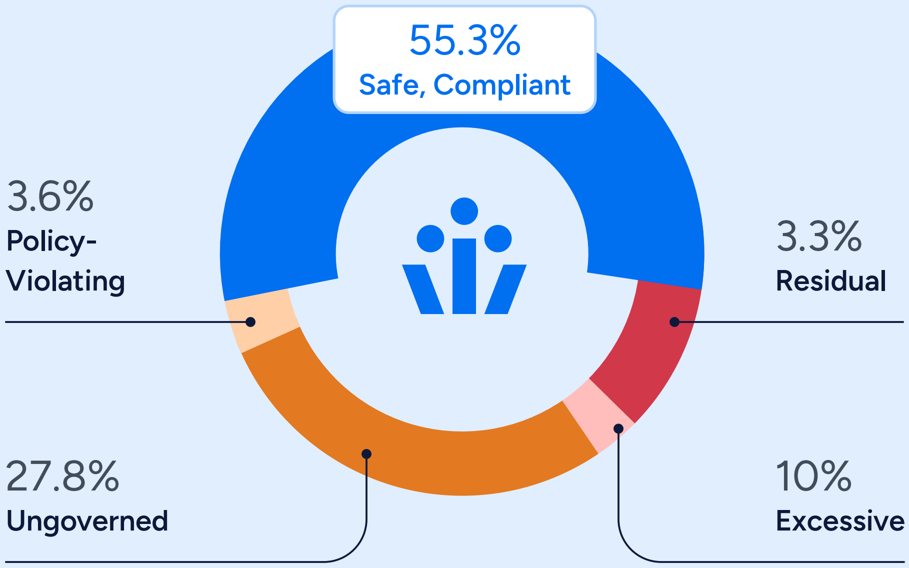

Enterprises have a lot of bad permissions because new ones are created throughout the organization every day. Over-privileged permissions are created every time an IT staff grants access on the basis of a mislabeled group name, making their best guess about access. Residual permissions are created every time an employee leaves or changes jobs. Ungoverned permissions are created when access is provisioned locally to a cloud-based system or SaaS app. Because identity and security teams cannot see these bad permissions, they cannot fix them. Thus the permissions accumulate as “access debt”. With more debt comes increased risk of insider threats, breaches, and compliance headaches.

Compared to 2024, the overall security posture has deteriorated. “Safe, Compliant” permissions dropped from 70% in 2024 to 55.3% in 2025, driven primarily by a significant increase in “Ungoverned” permissions (5% → 27.8%). A major contributing factor is the rising use of “local” accounts that bypass the required process for linking back to IdP users, underscoring the persistent tension between strong security controls and day-to-day business needs.

With billions of permissions to manage, security and identity teams are struggling to maintain and enforce the principle of least privilege across their organizations. Excessive privileges, dormant accounts and over permissions are running rampant all across. The latest State of Identity and Access Report by Veza illustrates these threats and underscores a key tenet: identity risk is everywhere, and it's growing faster than most teams realize. Every security leader should study this report and use it to inform their roadmap of understanding and countering these threats before they become impossible to address effectively.

Phil Venables

Cybersecurity leader, Partner at Ballistic Ventures,
and former CISO, Google Cloud

Rise of the Machines

Non-human identities (NHIs) scale value, but without governance, they scale the blast radius too.

Enterprises used to think of identity as an HR problem: focusing on the joiner, mover and leaver (JML) of human users. That view is dangerously outdated, however. Today, NHIs — service accounts, API keys, workloads, RPA bots, and now AI agents and unsanctioned MCP servers — outnumber human users by a factor of 17:1. The access associated with every one of those is a potential entry point. The numbers tell the story:

2,188 NHIs

(0.01% of all NHIs)

control 80% of cloud resources across enterprises. A small handful of accounts can bring an entire business to a halt.

This concentration of privileges in a relatively small number of NHIs indicate the potential for huge improvements in reducing risk- assuming an organization has the right level of visibility to know which accounts are the highest priority to secure.

Expel's 2025 report confirms adversaries are increasingly targeting API tokens and machine credentials because they offer persistence and scale without drawing attention. Adversaries don't need to chase privileged employees anymore; crown jewels are already exposed through uncontrolled machine identities.

Non-human identities (NHIs) outnumber human identities by a factor of

17:1

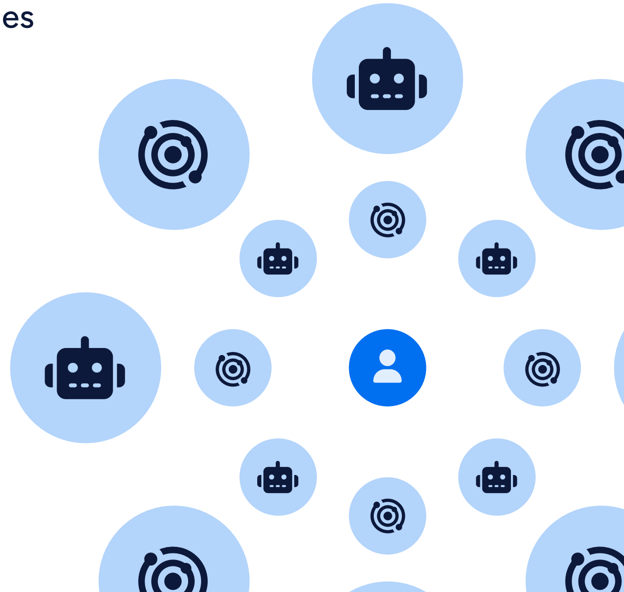

Why it Matters

1

AI agents and automation are multiplying faster than oversight can keep up. Every new workflow, model, or API spins up fresh entitlements, most of which are never reviewed.

2

Persistent attack surface: Unlike employees, machine identities don’t leave when they “quit.” They persist indefinitely unless explicitly revoked, quietly expanding the enterprise’s risk footprint.

3

Business continuity risk: A single overlooked bot, token, or API key can unlock access to critical infrastructure or data.

4

Audit blind spots: Traditional IAM and IGA systems still treat NHIs as static records, not as active entities with behavioral risk profiles.

5

Invisible growth: Every CI/CD pipeline, automation, and AI integration spawns new credentials without security review or lifecycle governance.

6

Persistence by design: NHIs are not subject to HR offboarding or compliance processes. Once created, they live forever unless deactivated.

7

Amplified blast radius: A single compromised NHI often has broader entitlements than any human user, making every compromise exponentially more damaging.

8

Uncontrolled dominance: Non-human identities (NHIs) now hold the majority of infrastructure-level permissions yet remain largely unmonitored and unmanaged.

Ransomware crews increasingly avoid noisy malware, leaning on valid credentials, admin tooling, and gaps in backups and MFA. That tradecraft amplifies the risk of powerful NHIs.

In effect, enterprises are drowning in identities but starving for control. And adversaries are the ones taking advantage.

Executive Takeaways

- NHIs now dominate the identity attack surface. They outnumber human users 17:1. Unlike employees, NHIs don’t leave the company; they accumulate, silently expanding exposure.

- A tiny number of NHIs hold disproportionate power. Just 2,188 (0.01% of all NHIs) machine identities control 80% of cloud resources.

- Plan for NHI adaptation: Contain the high-impact access set for NHIs now. Drive the NHI-related risks down quarter over quarter.

Human Risk is the Biggest Risk

Identity risk isn't just about hackers or rogue insiders. It's about the uncontrolled permission sprawl that HR and IT leave behind. Every new hire adds accounts, entitlements, and access to critical applications. But when people leave, that access often doesn't leave with them.

Our analysis shows the scale of the problem:

The average entity today carries 96,000 permissions.

Across enterprises, we found 78,000 ex-employees (3% of all employees) still retain active credentials.

Even when HR marks identities as inactive, 38% still have live entitlements in critical systems like Salesforce, SAP, and Workday. That gap is not an IT housekeeping issue but a board-level exposure with direct business consequences.

38%

of all accounts are dormant, with live entitlements.

Why it Matters

1

Compliance Risk

Regulators expect clean, provable user access reviews. Lingering permissions and ex-employee accounts are audit failures waiting to happen.

2

Operational Risk

A single forgotten account can enable ransomware to spread across finance, HR, or operations systems.

2

Reputational Risk

Few headlines resonate more than “former employee still had access to payroll” or “contractor stole data after termination.”

78,000 ex-employees (3% of all employees) still retain active credentials.

The most dangerous employee is the one that’s already gone.

Executive Takeaways

Every employee is over-entitled. The average worker now has 96,000 permissions.

Departures don’t mean deprovisioning. 78,000 (3% of all employees) ex-employees still have active access in enterprise systems.

HR + IT gaps = exposure. 38% of “inactive” users in identity systems still retain entitlements in core apps.

Insider risk is often accidental. Most incidents aren’t malicious employees; they’re forgotten accounts and permission sprawl.

Incomplete offboarding and non-immutable backups provide low-hanging fruit for adversaries. Groups like BianLian weaponize these identity lifecycle gaps, turning simple hygiene failures into catastrophic breaches.

Orphaned Accounts & Weak MFA = Access Time Bomb

The fastest-growing source of enterprise breaches isn't malware, it's forgotten identities and weak authentication. Dormant accounts, orphaned logins, and incomplete MFA deployments are the low-hanging fruit adversaries count on. The numbers are staggering:

8%

We found 824,000 (8% of all IdP users) orphaned IdP accounts with no HR match across enterprises

13%

13% of enterprise users still lack MFA, leaving them completely exposed. Even where MFA exists, many rely on weak verification methods: over 61,000 Okta users (6% of all Okta users) still use SMS or email MFA, which adversaries can easily bypass.

Multiple assessments place identity misuse at the center of most breaches. The early-2025 spike mapped to crews that weaponize weak hygiene. They succeed where MFA coverage is inconsistent and dormant identities linger.

Edge devices and VPNs remain common initial access points. Treat them as identity controls, not just network gear.

In the Change Healthcare breach, MFA was removed for certain staff while traveling, and adversaries exploited that exact gap to walk in. CrowdStrike 2025 reports weak or absent MFA was exploited in nearly 60% of credential-based breaches. The result? Adversaries don't need to invent new exploits. They weaponize what's already there: a stale account, a missing MFA prompt, a forgotten service credential. One gap is all it takes to compromise an entire business, disrupt critical operations, or, as in the case of Colonial Pipeline and Change Healthcare, destabilize entire industries.

Why it Matters

1

Business continuity risk

Healthcare payments, fuel pipelines, and financial operations have all ground to a halt because of single missed MFA or forgotten accounts.

2

Technical debt = identity debt

Every orphaned account is a potential ransomware launchpad.

2

MFA gaps break the chain

MFA is the single strongest line of defense, yet inconsistent deployment leaves predictable holes.

> *When you see the complexity of relationships across billions of permissions at scale, the magnitude of the problem of achieving Least Privilege comes into focus. It's really no surprise that enterprises are struggling with it- it's an immensely challenging and important problem.*

Rich Dandliker

Head of Strategy

Executive Takeaways

- Orphaned accounts are everywhere. Over 824,000 active identities (8% of all identities) have no HR match.

- MFA adoption is inconsistent. 13% of users lack MFA entirely; many rely on weak SMS/email factors.

- One lapse can trigger systemic failure. Both Change Healthcare and Colonial Pipeline were breached via dormant accounts without MFA.

- Identity controls are business-critical. Strong MFA and continuous identity hygiene are now as essential as firewalls once were.

- Service-desk manipulation to add new MFA devices or force resets remains a primary on-ramp. Treat help-desk identity flows as high-risk controls.

Adversaries don't need zero-days.

They need one forgotten account.

Identity Security Next Steps for C-Level Leaders

The 2026 State of Identity & Access report makes one thing clear: identity security is now a business-critical discipline. Adversaries are focused on exploiting credentials, processes, and organizational blind spots related to identity and access. The path forward for executive leaders is to treat identity as an enterprise control surface, with the same rigor applied to finance, supply chain, and compliance.

To regain control and reduce identity debt, C-level executives should act on five priorities:

1 Make Identity a Board Metric

- Require regular reporting on the number of active vs. dormant accounts, orphaned identities, and non-human identity growth.
- Track identity risk as a measurable form of operational debt — one that compounds unless remediated.

2 Establish Continuous Visibility Across All Accounts

- Demand a unified view of "who can do what, where, and when" across human, machine, and AI identities.
- Replace snapshot audits with ongoing validation and real-time access analytics.

3 Mandate Zero-Tolerance for Dormant and Orphaned Access

- Set an executive-level goal: eliminate dormant accounts (>90 days) and orphaned identities within one quarter.
- Tie access hygiene metrics to performance indicators for CIO, CISO, and business unit leaders.

4 Operationalize Governance for AI and Machine Identities

- Require lifecycle management and least-privilege enforcement for all non-human identities, APIs, bots, service accounts, and AI agents.
- Treat these identities as living assets with ownership, expiration, and continuous review.

5 Extend Identity Governance beyond the compliance minimum

- Require lifecycle management and least-privilege enforcement for all high-risk access: privileged access, access to restricted data.
- Use Just-in-Time (JIT) access for elevated roles with time-bound, approver-gated elevation instead of permanent privileges.
- Conduct regular access reviews to revoke outdated, unused, or unnecessary entitlements for all High-Risk access.

Identity is now the control plane of the modern enterprise. Organizations that quantify, normalize, and continuously validate access will not only shrink their attack surface but also gain the operational assurance boards, regulators, and insurers demand. The next generation of resilient enterprises will be those that treat identity security as a measurable business outcome.

The average entity today carries 96,000 permissions.

The Veza Platform

Access Graph Learn more	Best-in-class graph database that unravels complex entitlements to reveal the true effective permissions for every identity, at scale.
Access AI Learn more	Generative AI-powered capabilities to accelerate threat hunting and remediation, enabling SOC teams to quickly discover and fix identity security risks.
Access Hub Learn more	Removes IT bottlenecks, maximizes productivity and ensures least privileges through an easy-to-use self-service access hub for employees.

> *People cannot manage what they cannot see. Veza's Access Graph brings that visibility to life, and what we're uncovering today is only the tip of the iceberg. With the rapid proliferation of AI agentic identities, the access problem is getting exponentially harder. As an industry, we must take decisive action to enforce and realize true least privilege before the gap becomes unmanageable.*

Maohua Lu

Co-Founder & CTO

veza

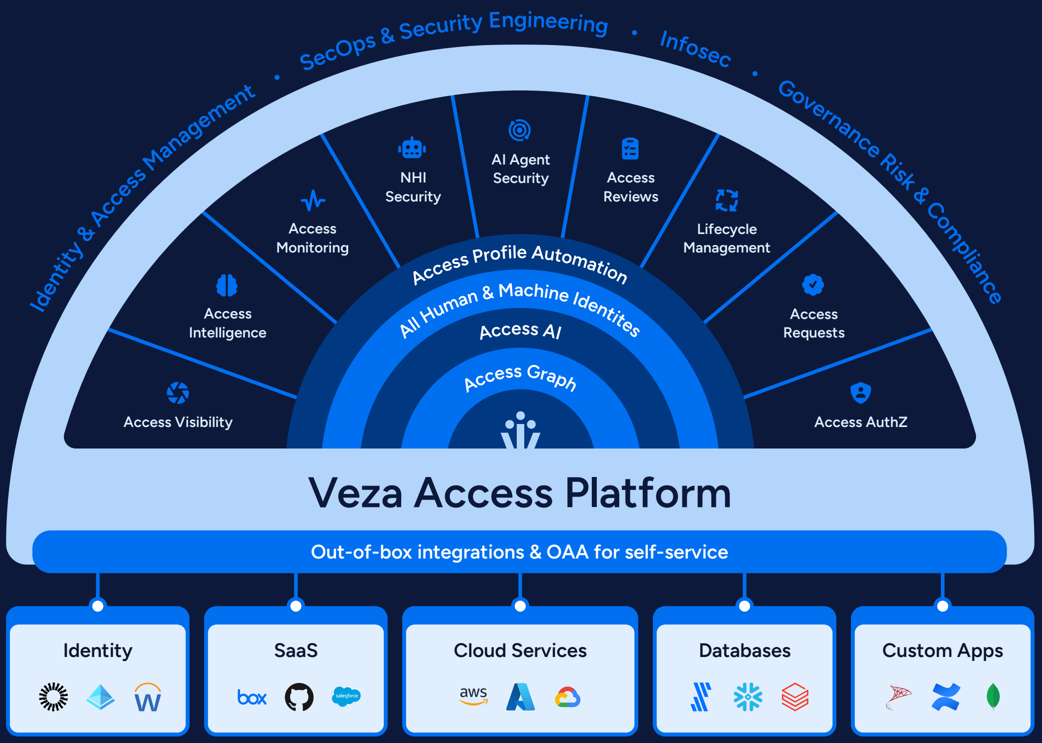

Veza Products

Access Visibility

Learn more

Eliminates visibility gaps by instantly mapping all human and non-human identities to effective permissions across all systems, including apps, on-premise, cloud services, and data systems.

Access Intelligence

Learn more

Provides actionable insights to reduce identity risks with visibility into privileged users, dormant accounts, policy violations and misconfigurations with Veza's 2,000+ pre-built queries.

Access Monitoring

Learn more

Monitors activity by identities and roles on key resources to identify over-privileged permissions, right-size roles, and trim unneeded access and entitlements to sensitive resources.

Access Reviews

Learn more

Automates user access certifications through comprehensive campaigns, prioritizing risky access first and giving reviewers the context they need to approve or reject.

Lifecycle Management

Learn more

Automates the joiner, mover, leaver processes by provisioning and deprovisioning access throughout the users' lifecycle.

Access Requests

Learn more

Streamlines how access is granted from the start and ensures that every permission is provisioned/ deprovisioned correctly and securely by adhering to the principle of least privilege.

Access AuthZ

Learn more

Provides visibility, control, and automated execution needed to achieve true least privilege at scale by automating access grants and revocations across cloud, SaaS, on-prem, and custom apps.

Separation of Duties

Learn more

Prevents internal fraud and compliance failure by discovering and mitigating toxic combinations and separation of duties violations within and across platforms.

NHI Security

Learn more

Provides full visibility and control over Non-Human Identities (NHIs) with a complete inventory of service accounts, keys, and secrets, eliminating ungoverned access and 'shadow' service accounts.

Veza is the identity security company. Identity and security teams use Veza to secure identity access across SaaS apps, on-prem apps, data systems, and cloud infrastructure. Veza solves the blind spots of traditional identity tools with its unique ability to ingest and organize permissions metadata in the Veza Access Graph. Global enterprises like Wynn Resorts, and Expedia trust Veza to visualize access permissions, monitor permissions activity, automate access reviews, and remediate privilege violations. Founded in 2020, Veza is headquartered in Los Gatos, California, and is funded by Accel, Bain Capital, Ballistic Ventures, GV, Norwest Venture Partners, and True Ventures. Visit us at veza.com and follow us on LinkedIn, Twitter, and YouTube.

Request a demo