TechnologyAdvice

AUGUST 2026

INDUSTRY INSIGHTS REPORT

Inside the CISO Stack: What's Changing in 2026

SPONSORED BY

FusionAuth

CONTENTS

Inside the CISO Stack: What's Changing in 2026

82%

of detections involved no malware at all

— CROWDSTRIKE

67%

of users accessing AI applications through non-corporate accounts

— VERIZON

88%

of organizations say AI deployment is outpacing their identity controls

— FUSIONAUTH

INTRODUCTION

Attackers aren't breaking in. They're *logging* in.

That single shift, from exploitation to authentication abuse, is rewriting the CISO mandate faster than most security programs can respond. According to CrowdStrike's 2026 *Global Threat Report*[1], 82% of detections involved no malware at all.

Which means that attackers aren't defeating security controls anymore. They're using them.

The perimeter model was already strained. Now AI has collapsed the timeline for replacing it. Organizations are embedding AI agents into production environments, customer-facing applications, and core business workflows faster than their identity and security infrastructure can track. The *2026 Verizon Data Breach Investigations Report (DBIR)*[2] found 67% of users accessing AI applications through non-corporate accounts. FusionAuth's *State of AI and Identity Report*[3] found 88% of organizations say AI deployment is outpacing their identity controls.

The time to prepare for AI has passed. Now the challenge lies in governing the AI that's already deployed.

The core issue is architectural. Most enterprise security stacks were designed around a simple mental model: a human user, a managed device, a corporate network. Verify the person at the edge, grant access, trust the session. That model breaks entirely when the "user" is an AI agent chaining API calls across a dozen systems, authenticating against the same infrastructure as your customers, and holding permissions that were never designed with autonomous actors in mind.

This report examines what CISOs are actually changing in response. Which foundational capabilities are being consolidated, which legacy technologies are being retired, and why identity has moved from a login function to the control plane of the modern security stack.

Organizations that treat identity modernization as a routine technology refresh are misreading the threat. What's actually required is a ground-up rethink of how trust is established across the business.

The decisions CISOs make now will determine whether identity becomes a durable control layer or a persistent source of enterprise risk.

[1]Global Threat Report, CrowdStrike, 2026
[2]Data Breach Investigations Report (DBIR), Verizon, 2026
[3]State of AI and Identity Report, FusionAuth, 2026

HOW THE CISO MANDATE AND STACK ARE CHANGING

> Your job is no longer simply to protect individual systems. It's to make sure every identity, human and non-human alike, earns only the trust it needs, for only as long as it needs it.

The CISO mandate has expanded. The practical response? Don't add another layer of oversight. Instead, redesign how trust is established and enforced across the business.

It's not an easy task. You still have to stop ransomware, reduce cyber risk, respond to incidents, satisfy regulators, and protect customer confidence. But now, you also need to make it possible for your business to deploy autonomous systems without granting them open-ended authority. That escalates identity architecture to a board-level risk decision rather than an implementation detail delegated to an application team.

You should start by focusing on the paths where trust is inherited. Which identities can access production data? Which service accounts hold permissions no one has reviewed in a year? Which customer applications share tenant infrastructure? Which AI agents can call other systems, and who is accountable for the actions that follow? These questions reveal risk that a product inventory cannot.

The next priority? Understand the differences between authentication and authorization. A valid login proves that a credential *worked*, not whether the requested action is appropriate, the permission should still exist, or if the session should remain trusted. You need controls that evaluate access continuously, reduce permissions to the minimum required, and make policy decisions visible enough to audit.

Finally, you need architecture that can change without forcing a disruptive rebuild. AI adoption, data residency rules, tenant-isolation demands, and new identity types will continue to move faster than traditional refresh cycles. This makes standards-based integration, deployment flexibility, and clear data ownership risk controls in their own right.

The takeaway is operational: your job is no longer simply to protect individual systems. It's to make sure every identity, human and non-human alike, earns only the trust it needs, for only as long as it needs it.

AI is changing the risk conversation

AI changes the risk conversation. How? By turning permission design into runtime behavior. An employee can only make one request at a time. An agent can make hundreds,chaining them across systems and acting before a human notices the first decision was wrong. That speed exposes a critical flaw in how most security stacks are built: treating successful authentication as the end of the trust decision.

The infrastructure wasn't built for this. But it's being adapted to the cause. Every AI assistant, autonomous agent, API integration, and machine identity authenticates, receives permissions, and accesses data through infrastructure originally designed for predictable users and workflows. Broad roles, long-lived

85%

of organizations report employees using unauthorized or "shadow AI" tools

— FUSIONAUTH

tokens, and weak revocation processes become much more dangerous when the identity using them can operate continuously.

FusionAuth's findings highlight the mismatch:

85% of organizations report employees using unauthorized or "shadow AI" tools.

65% experienced a confirmed AI identity-related security incident during the past year.

70% already have customers interacting with AI-powered product features.

88% say AI deployment is moving faster than identity and security infrastructure can support.

For CISOs, those numbers point to four immediate tests:

1. Can your organization identify every agent?
2. Can it identify the human or system that delegated its authority?
3. Can it limit the agent to specific data, APIs, actions, and time periods?
4. Can investigators reconstruct the full chain of activity after an incident?

If the answer to any of those questions is no, your organization doesn't have AI *governance*. It has AI *access*.

Mid-market technology leaders report the same operational failures. Over-permissioned access. Incomplete audit logs. Tenant-isolation requirements that can't be consistently enforced by existing platforms. When an autonomous agent chains API calls across a dozen systems, a perimeter check at the first door is meaningless if every interior door opens automatically.

"This is a great illustration of why AI agent authorization is the harder, and more critical, problem than authentication," said Dan Moore, FusionAuth's Senior Director of CIAM Strategy and Standards, boiling the challenge down to its most critical element.

AI already has access to your attack surface. Can your identity architecture govern the actions it's already taking?

FROM SECURITY PRODUCTS TO SECURITY ARCHITECTURE

> Today, attackers rarely begin by defeating your infrastructure.They begin by inheriting trust.

For decades, enterprise security followed a predictable pattern: a new threat emerges, a new product appears to solve it.

Endpoint? Buy EDR.

Privileged access? Buy PAM.

APIs? Add another gateway.

Every problem had its own category, budget, and dashboard. The result was tool sprawl and a security strategy built around protecting individual technologies instead of controlling how identities move through the business.

Today, attackers rarely begin by defeating your infrastructure. They begin by inheriting trust: a stolen credential; an over-permissioned service account; a compromised API token; an AI agent granted access it never should have had. The attack succeeds because the architecture assumes the request is legitimate.

That's why security investment is shifting away from adding another point solution and toward strengthening the systems every request already passes through. Identity, authorization, governance, and detection are becoming the control plane for enterprise security.

That control plane now reaches well beyond your own infrastructure. Modern applications run on customer identities, partner connections, cloud platforms, APIs, and AI services, often across several environments at once. Every one of those requests still passes through the identity layer, so that layer has to hold up under whatever data-residency rules, tenant-isolation requirements, and deployment models each environment demands.

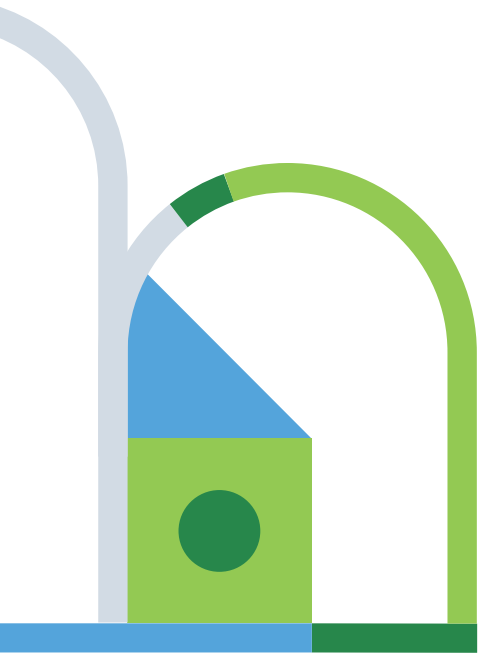

FusionAuth's research found that 57% of organizations identify deployment flexibility as one of the most important criteria when evaluating identity platforms, but only 11% prioritize total cost of ownership. Buyers rank cost low because being locked into the wrong architecture can cost far more than the platform itself.

Deployment flexibility is a control surface that governs three risks that architecture decisions and software and policy can't retrofit:

- **Sovereignty** — where your identity data lives.
- **Blast radius** — whose compromise affects you.
- **Performance** — how quickly authentication and authorization decisions can be made at scale, close to the application.

Specialized security products will still solve specialized problems. The change is that foundational trust decisions can no longer be scattered across disconnected tools. The next-generation stack will be measured by how consistently it authenticates, authorizes, observes, and governs employees, customers, applications, APIs, machine identities, and AI agents.

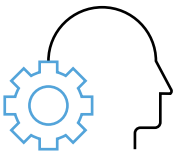

Request resource

- Human user
- AI agent
- Applicaition/API

Authentication
Who or what is making the request?

- Passkey
- Token
- Session
- Credential

Authorization
What is this identity allowed to do?

- Delegated authority
- Specific data
- APIs
- Actions
- Time limits

Continuous evaluation
Should trust continue right now?

- Risk
- Context
- Policy
- Least privilege

Outcome and audit
Allow, step up, restrict, revoke, and log every action

- Audit trail
- Chain of activity
- Reconstruct incidents

WHY IDENTITY AND AUTHORIZATION ARE BECOMING FOUNDATIONAL

Every security decision eventually becomes an identity decision.

Firewalls, endpoint protection, and segmentation still matter. But in a business spread across cloud platforms, SaaS applications, APIs, mobile devices, partner systems, customer portals, and AI services, "where did this request come from" is the wrong question. The right question is, should this identity be allowed to perform this action?

Authentication answers who or what is making the request. Authorization answers the harder part: what that identity may actually do. Every access decision after login depends on those controls being specific, current, and enforceable.

"Written policies don't answer the questions that matter," according to FusionAuth CEO Brian Bell. "Can you scope what each agent can access? Can you see what it's doing? Can you prove what it accessed after the fact? Can you revoke access before a near miss becomes something worse? Architecture answers those questions. Policy alone does not."

That's why identity has become the enforcement layer for Zero Trust. Trust can't be granted once at login and assumed for the rest of the session. Least privilege, adaptive authorization, context-aware policy, and continuous verification decide whether access should continue.

Identity also limits the damage when compromise happens. Passkeys, scoped permissions, adaptive access policies, short-lived credentials, and strong tenant isolation shrink what an attacker can do after they get in.

> Written policies don't answer the questions that matter... Architecture answers those questions.

BRIAN BELL
CEO, FUSIONAUTH

Organizations are starting to prioritize those controls. FusionAuth found that 72% of businesses rank support for machine and non-human identities among their top identity platform evaluation criteria, and 54% prioritize fine-grained authorization.

That tracks with attacker behavior. CrowdStrike found 82% of detections were malware-free, with adversaries using stolen credentials, valid accounts, and session tokens rather than malicious code.

When attackers inherit trust instead of breaking through defenses, identity isn't another layer in the stack. It's the point where the architecture decides whether trust is deserved.

Customer identity has become business risk

Identity isn't just how customers enter an application. Identity is how your organization establishes trust across its digital business.

Customer identity used to be treated as an application feature. Now, it's business infrastructure. Every customer login, partner integration, API request, and AI-powered interaction is a trust decision with consequences for security, compliance, privacy, and revenue.

One compromised identity can trigger regulatory reporting, expose customer data, interrupt digital services, or damage confidence in your products.

As customer identity takes on greater importance across security, compliance, privacy, and revenue, responsibility for it is shifting beyond application development teams. Increasingly it's being evaluated as part of the broader security architecture. CISOs and tech leaders are looking beyond authentication to authorization models, policy enforcement, auditability, tenant isolation, data residency, and the ability to support future business models without rebuilding the foundation.

Stronger security should not create more friction

Security doesn't mean forcing customers to choose between safety and usability. Poor architecture creates that trade-off, but better architecture can remove most of it.

Modern identity platforms can make decisions based on risk rather than applying the same friction to every interaction:

- Low-risk requests move with little interruption.
- Higher-risk activity triggers added verification based on device reputation, impossible travel, unusual behavior, or anomalous login patterns.

Phishing-resistant technologies such as WebAuthn/passkeys, biometrics, and passwordless authentication can improve security and user experience at the same time by replacing weak, reusable secrets with stronger methods that are easier for the user.

The goal is providing more assurance with less unnecessary friction.

85%

of organizations have been asked to demonstrate tenant isolation

— FUSIONAUTH

70%

of organizations already have customers interacting with AI-powered features

— FUSIONAUTH

Architecture has become the competitive advantage

Security is table stakes. Architecture is the advantage.

Organizations no longer evaluate identity platforms by feature lists alone. Now they want to know whether the architecture can support where their business is going next.

Customers, partners, and regulators increasingly expect proof of tenant isolation, data sovereignty controls, auditable access decisions, and secure support for AI-enabled services. These aren't implementation details when they decide whether you win a deal, enter a market, or pass an audit.

FusionAuth's research found that 85% of organizations have been asked to demonstrate tenant isolation, while 56% receive those requests frequently.

About 70% of organizations already have customers interacting with AI-powered features, which raises the stakes on where identity data lives. The research found an architectural split in incident rates. Organizations on multi-tenant SaaS identity platforms reported an 83% rate of confirmed AI-related identity or security incidents, compared with 38% on self-hosted or on-premises deployments.

That gap comes down to blast radius — how far a single failure can spread. On a shared, multi-tenant platform, your identity data sits in the same system as other companies' data, so one stolen credential or misconfigured policy doesn't necessarily stop at the customer where it started. It can reach others in the same environment. Single-tenant deployments give each customer their own isolated instance, so a failure stays contained. No one else's breach reaches you, and yours doesn't reach them.

The deployment model alone doesn't determine the outcome. But it does show that architecture changes the risk profile, which is why deployment flexibility, tenant isolation, authorization, data control, and governance are moving into board-level conversations.

HOW CISOS ARE CONSOLIDATING AND MODERNIZING

Security teams did not set out to build fragmented architectures. They inherited them one urgent purchase at a time. A new threat appeared, a new product category followed, and the organization added another tool.

Repeat the cycle for a decade and you get a predictable result: overlapping policies, duplicate identity stores, inconsistent integrations, dozens of consoles, and analysts spending too much time managing products instead of reducing risk.

Complexity is a security problem. And it can be an expensive one.

Palo Alto Networks' *2026 Identity Security Landscape Report* found that identity fragmentation added an average of 12 hours to incident response for identity-

93%

of organizations are reevaluating identity infrastructure

— FUSIONAUTH

related incidents. It's what happens when you scatter authentication events, permissions, and identity records across disconnected systems. Investigators lose time reconstructing what should have been visible in one place.

CISOs are responding by changing the architecture rather than adding another layer, using a smaller number of foundational systems that consistently establish trust and leaving specialized tools to solve the problems that genuinely require specialization.

From tool sprawl to architectural platforms

Many organizations have accumulated separate technologies: one for workforce authentication, another for customer identity, and others still for single sign-on, multi-factor authentication, privileged access, API authentication, and machine identities. Each tool solves a legitimate problem, but together, they produce duplicated administration and fragmented policy enforcement and lead to inconsistent user experiences and limited visibility across environments.

FusionAuth's research found that 93% of organizations are reevaluating identity infrastructure. About 64% said AI is the primary driver for that reassessment. That's not a normal refresh cycle, it's an architectural redesign prompted by the realization that identity now governs nearly every meaningful interaction in the enterprise.

The time has passed for asking if another identity product is needed. Now it's a question of whether employees, customers, applications, APIs, and AI agents are governed through consistent standards, authorization policies, audit logging, and lifecycle controls.

Retiring legacy authentication

Your organization might have custom authentication, password-based workflows, static API keys, aging directories, and legacy protocols that all solved real problems at the time that you implemented them. But security standards have evolved. Passkeys, phishing-resistant authentication, risk-based MFA, privacy controls, new regulations, and continuous security updates keep adding weight to infrastructure that may have been designed only to validate a username and password. Technical debt eventually becomes security debt.

How do you modernize without increasing the risk of disrupting customers, developers, or business operations? One way is with a phased migration. Legacy and modern identity systems operate together while users move gradually as they authenticate. You can retire inactive accounts, introduce stronger methods, and reduce cutover risk without forcing every user through a single high-stakes migration event.

Modernization becomes an architectural evolution rather than a replacement project with a cliff at the end.

Build versus buy is no longer just an engineering decision

Building authentication once meant building login pages. That description is

> You're not trying to find one platform that pretends to do everything. You want a reliable foundation that handles common trust decisions consistently, with specialized products added where they deliver depth the platform cannot.

now dangerously incomplete. Modern identity systems may need to support passkeys, phishing-resistant authentication, risk-based or "intelligent" MFA, fine-grained authorization, API security, machine identities, audit logging, consent management, evolving standards, data residency, regulatory requirements, and constant security updates.

An internal team may be able to build all of that, but there's a more strategic question that should be asked. *Should* they? Is it worth spending scarce engineering capacity to maintain security-critical infrastructure that customers rarely see when the same people could be building the product customers actually *buy*.

That consideration has a significant impact on whether organizations decide to build or buy their identity solution. Building in-house may offer greater control over architecture, customization, and implementation. But you have to weigh those benefits against the ongoing cost and complexity of maintaining security-critical infrastructure.

Modern CIAM platforms can provide extensibility, deployment flexibility, and operational control while reducing the internal burden of keeping identity systems secure, compliant, and aligned with evolving standards.

Vendor choice has become an architectural decision

Selecting an identity platform is not as simple as comparing features and scoring them against one another. Deployment flexibility, interoperability, open standards, data ownership, portability, tenant isolation, migration paths, and pricing structure all affect how easily the organization can respond to future business, regulatory, and security requirements.

Organizations increasingly want confidence that they control where identity data resides, how it is governed, and how easily the architecture can change, not the vendor. Today's platform decision may determine whether your organization can enter a regulated market, isolate a high-value tenant, move regions, or support a new class of AI agent several years from now.

Even the cheapest identity platform becomes expensive if leaving it later requires rebuilding the application around it.

Consolidate the foundation, specialize where it matters

Consolidation does not eliminate the need for specialized security tools. It clarifies where specialization earns its place. Threat intelligence, offensive security, advanced fraud detection, highly specialized data protection, and complex authorization — these may all still require purpose-built technologies.

You're not trying to find one platform that pretends to do everything. You want a reliable foundation that handles common trust decisions consistently, with specialized products added where they deliver depth the platform cannot.

The emerging model is selective consolidation. Identity, authorization, governance, detection, and policy enforcement form the foundation. Specialist tools solve specialist problems. And the entire stack operates as a coordinated system rather than a pile of unrelated products.

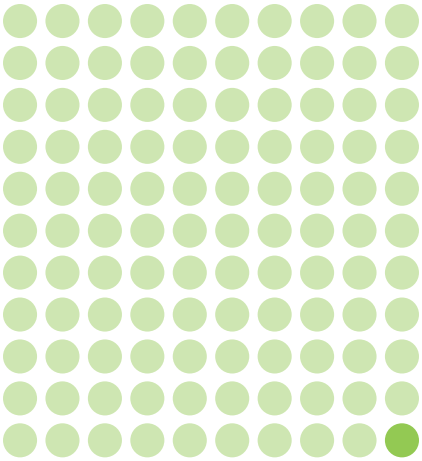

109 to 1

Machine identities already outnumber human identities

— PALO ALTO NETWORKS

HOW AI AGENTS CHANGE IDENTITY GOVERNANCE

AI is not introducing a new identity problem. Rather, it's exposing the fact that most identity systems were designed for a world in which humans were the primary users — and that assumption no longer holds.

Organizations are deploying AI assistants, autonomous workflows, customer-facing agents, and machine-driven processes that authenticate, invoke APIs, retrieve data, and make decisions without a person approving every step. Nobody's asking whether AI belongs inside the enterprise anymore. The more relevant question is whether the identity architecture was built to govern autonomous actors. For many organizations, it was not.

The strength of the next security stack will depend on whether organizations can govern human and non-human identities with consistent visibility, accountability, and control.

AI agents are becoming first-class identities

Traditional identity platforms focused mainly on people. Service accounts, APIs, and machine identities usually operated inside narrow, predefined workflows.

But AI agents are different. They interpret objectives, chain API calls, retrieve information from multiple systems, make decisions, and act on behalf of employees, customers, and other agents. If they can choose the next action, they're not background automation. They're participants in the operation of the business.

From the identity platform's perspective, an agent consumes credentials, receives permissions, accesses sensitive resources, and establishes trusted sessions. In practical terms, it's a first-class identity.

"An AI agent is not a new user to authenticate," said Dave Hayes, FusionAuth's Vice President of Product. "It has no authority of its own; it acts for a human, and that human is where the authority comes from… AI is probabilistic, so your identity layer has to be deterministic."

The scale matters too. Palo Alto Networks reports that machine identities already outnumber human identities 109 to 1.[1]

As AI adoption grows, agents, APIs, service accounts, and autonomous workflows will expand far faster than the workforce. An identity strategy designed around employees alone will govern a shrinking share of the identities actually operating inside the environment.

Authentication gets you in, authorization determines what happens next

Authenticating an AI agent is easy. Deciding what it's allowed to do afterward is where the architecture earns its keep.

Who delegated the agent's authority? Which systems can it access? How long should the permission last? Can every action be traced to an accountable source? Can access be revoked immediately when behavior changes?

[1]2026 Identity Security Landscape Research Report, Palo Alto

> **Authentication establishes trust. Authorization decides whether that trust should continue.**

Those questions move identity beyond login and into continuous authorization and governance.

Long-lived API keys, persistent service accounts, and broad role assignments were already difficult to defend. They become harder to justify when hundreds or thousands of autonomous identities operate across multiple systems at the same time. Organizations are moving toward short-lived credentials, delegated access, least privilege, continuous verification, runtime policy enforcement, and context-aware decisions.

That volume compounds: each of those identities can trigger hundreds of authorization checks in a single workflow, and a remote authorization service costs you a network round-trip on every one. That latency adds up fast.

The authorization decision always remains in the request path, but where it runs determines the latency cost of every check and whether fine-grained authorization is practical. Co-locating the decision point with the workloads it governs enables per-action, least-privilege authorization without relying on broad roles or stale cached permissions. Deployment flexibility is what makes that placement possible.

Authentication establishes trust. Authorization decides whether that trust should continue.

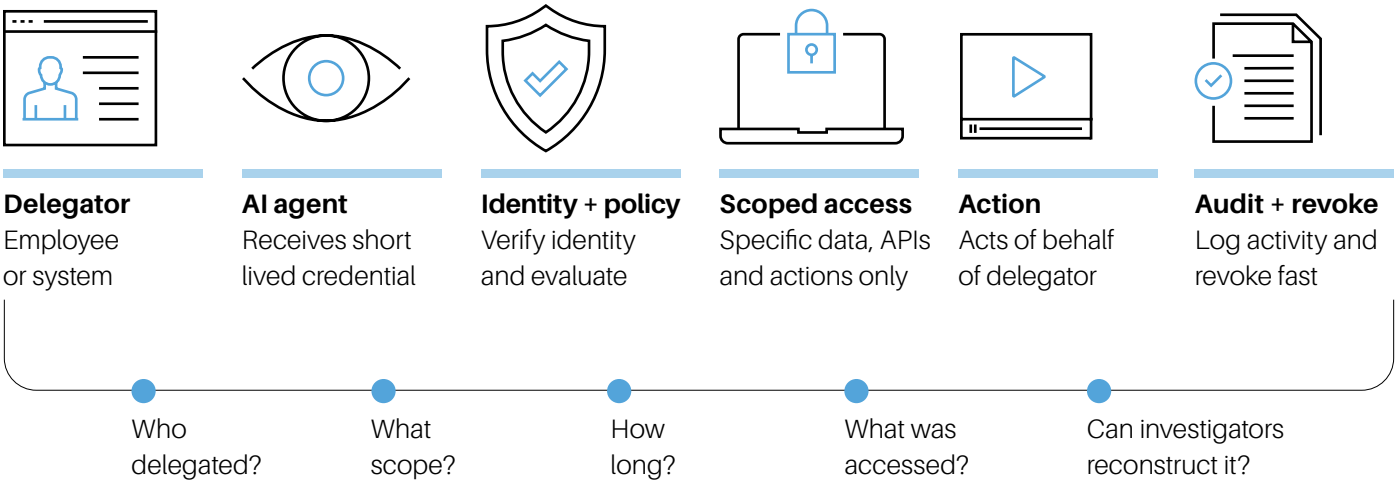

Governance has to extend beyond people

Zero Trust was never really about people. It was about refusing to treat any successful login as permanent proof of safety. The same principle applies to AI identities. Security teams need to know which agent performed an action, what authority it received, what data it accessed, whether the action complied with policy, and whether investigators can reconstruct the full chain later.

FusionAuth's research found that 58% of organizations say AI identity is primarily owned by security teams. Among organizations with AI-powered

> AI will automate more operational work, but it will also require stronger governance than any previous generation of enterprise technology because it can act, adapt, and cross system boundaries at machine speed.

products, 79% experienced a confirmed AI-related identity or security incident.

The same research also identified a maturity gap: auditing, lifecycle management, and access revocation remain among the least developed capabilities for AI identities. These controls become more important as autonomous activity scales, because governance cannot stop at authentication. It has to persist throughout the lifecycle of every human and non-human identity.

Identity becomes the control plane for trusted AI

As autonomous systems multiply, identity increasingly determines whether AI can operate safely inside the enterprise. Future identity platforms will need to authenticate not just employees and customers, but also agents and APIs, services, and machine identities, all through a common foundation. This has to be done while enforcing delegated authority, issuing time-bound credentials, applying context-aware authorization, and preserving verifiable chains of trust across connected systems.

Identity platforms must treat autonomous actors as native participants governed through the same deliberate, integrated controls as human users and other trusted entities.

Standards-based OAuth authorization, delegated access, fine-grained permissions, machine identity management, and policy-driven enforcement become foundational because they apply the same trust discipline regardless of whether the request comes from a person or an agent. You'll have an architecture that makes every identity prove what it is allowed to do.

The future stack gets simpler even as identities multiply

Modern security contains a paradox: the environment is becoming more complex while the architecture is becoming more focused. Specialist technologies continue to handle narrow, high-value problems while organizations consolidate foundational capabilities such as identity, authorization, governance, security operations, and data security into integrated platforms.

AI will automate more operational work, but it will also require stronger governance than any previous generation of enterprise technology because it can act, adapt, and cross system boundaries at machine speed.

The defining challenge will not be managing more users, but different kinds of identities, including employees, customers, agents, APIs, service accounts, autonomous workflows, and technologies that have not been invented yet. Each requires consistent authentication, authorization, governance, and accountability. Organizations that build for that reality can adopt AI without giving every new agent another path around the controls.

CONCLUSION

> Identity doesn't replace endpoint security, network security, application security, security operations, or data protection. It connects them.

Enterprise security was built around a simple assumption: protect the perimeter, verify the user, and trust the systems behind the login. But that assumption no longer describes the enterprise.

Modern businesses run on cloud services, APIs, customer applications, machine identities, and AI agents that authenticate continuously, exchange data across multiple systems, and make decisions without direct human involvement. In that environment, security isn't defined by where a request originates. It's defined by whether that request should be trusted now.

Identity doesn't replace endpoint security, network security, application security, security operations, or data protection. It connects them.

Every one of those controls eventually depends on knowing who or what is requesting access, what actions are permitted, how trust is reevaluated, and whether the decision can be audited and revoked. The organizations that succeed will not necessarily own the largest collection of security products. They will build architectures in which identity, authorization, governance, and policy enforcement operate as a coordinated control plane across employees, customers, applications, APIs, machine identities, and AI agents.

Research suggests many organizations already understand the shift. They are reevaluating identity infrastructure, consolidating fragmented platforms, modernizing legacy authentication, strengthening authorization, and preparing for environments in which autonomous systems operate alongside people. The strategic question is whether the architecture built today will still be capable of governing the identities your business depends on tomorrow.

Attackers aren't breaking in anymore. They're logging in. To respond effectively, you have to stop treating identity as the front door to the business and instead treat it as the decision-making architecture that determines how trust is established, limited, observed, and continuously earned.

Ken Underhill is an award-winning cybersecurity professional, bestselling author, and seasoned IT professional. He holds a graduate degree in cybersecurity and information assurance from Western Governors University and brings years of hands-on experience to the field. Sign up for our Cybersecurity Insider newsletter to get his free daily security briefing covering breaches, threats, and security priorities at **esecurityplanet.com/newsletter**

Sign up