eBook

Secure identity. Secure everything.

The three principles of a modern identity strategy for people, machines, and AI

okta

okta

What does your organization think about when it thinks about identity?

For many, the first answer is security. And for good reason. In a threat landscape dominated by credential theft, phishing, session hijacking, and the rapid rise of autonomous systems, identity has become the primary attack surface. It is how attackers gain access to sensitive systems, escalate privileges, and move laterally across environments.

Today, identity no longer refers only to employees and customers. It includes non-human identities, service accounts, workloads, APIs, and AI agents acting on behalf of users. As organizations adopt AI-driven systems and automation at scale, the number of identities requiring governance and control continues to grow rapidly.

Yet many organizations are still playing catch-up. They remain anchored to outdated perimeter-based security models, relying on network and device controls that cannot adequately protect against identity-centric threats. Fragmented visibility across human, non-human, and AI-driven identities leaves security teams reacting instead of enforcing policy proactively.

AI is accelerating the identity challenge

As organizations deploy autonomous AI agents and automated workflows, systems are increasingly acting on behalf of users. These agents retrieve data, execute tasks, and interact with multiple applications at machine speed. Without a unified identity layer to govern these actions, organizations risk creating a new class of unmanaged identities operating beyond traditional security controls.

Defining a modern identity strategy

A modern identity strategy establishes identity as the central control layer across your technology ecosystem. It delivers unified visibility into human, non-human, and AI-driven activity, identifies vulnerabilities before they are exploited, and enables real-time response to evolving risk.

It also enables organizations to scale securely. As AI agents, services, and applications operate with increasing autonomy, identity must continuously enforce least privilege and policy-based access without slowing innovation.

This resource explores why identity is the foundation of enterprise security in the AI era, including:

- How identity-related threats are evolving and accelerating
- Why traditional identity approaches leave organizations vulnerable
- The three principles that define a modern identity strategy

okta

The fragmentation of the enterprise security stack

The past decade transformed the enterprise tech stack. Cloud services, SaaS applications, APIs, and remote work reshaped how organizations build and operate. Today, AI agents and automated systems are accelerating that change even further. Applications do not just connect. They act. Services do not just store data. They execute tasks. AI systems retrieve information, trigger workflows, and make decisions across systems.

In this environment, the idea of technological unity under a single enterprise license agreement is no longer realistic. Organizations build ecosystems of best-in-breed solutions to stay competitive and move quickly. The result is a powerful and flexible infrastructure. It is also increasingly complex.

Modern tech stacks are highly distributed and deeply interconnected. Every new application, service, API, workload, and AI agent introduces new identities. These identities operate across clouds, on-premises environments, SaaS platforms, custom apps, and infrastructure. Without centralized control, identity becomes fragmented across a tangled web of systems and environments.

Siloed tech creates security blind spots

Fragmentation creates major security liabilities. Human identities, service accounts, machine identities, and AI agents often live in separate systems with inconsistent policies and incomplete oversight. This expands the attack surface and increases the likelihood of unnoticed credential theft, token misuse, and privilege escalation.

Security teams struggle to gain unified visibility across this landscape. Logs are scattered. Policies are inconsistent. Non-human identities multiply faster than they can be governed. As AI systems operate with increasing autonomy, organizations face a new challenge: enforcing control at machine speed.

Many AI agents and automated services operate with persistent credentials and broad access to multiple systems. When these identities are not properly discovered and governed, they can retain privileged permissions long after their intended purpose ends. This creates a growing pool of powerful identities that act across systems with little oversight, expanding both the attack surface and the potential impact of a breach.

Attackers understand this shift. Identity has become the primary attack vector because it is often the least unified control point. According to the 2024 Verizon Data Breach Report, in which Okta participated, 80% of breaches involve some form of compromised identity. The average time to recognize and contain a breach remains nearly 290 days.

It's time to change identity

Every organization understands that identity plays an important role in security. But traditionally, that role has been limited to authentication and access control. Few organizations are using identity in an equally critical role: as the central control layer for enterprise-wide visibility, governance, and enforcement across human and non-human identities.

For example, most organizations rely on identity to enable secure access through authentication. But when identity is more deeply integrated across technology and security ecosystems, it can power continuous risk evaluation, real-time policy enforcement, and automated remediation, even after a session has been established. In environments where AI agents and automated systems act on behalf of users, identity must verify not only who is accessing a system, but what is allowed to happen next.

Centralizing identity in your security ecosystem

To ignore this potential is to overlook what the data already shows: identity has become the primary battleground of enterprise cybersecurity. It is how attackers gain initial access, escalate privileges, and move across environments. It is also how organizations establish accountability, enforce least privilege, and maintain control at scale.

A modern identity strategy acts as a connective layer across your security and technology ecosystem, unifying visibility, governance, and enforcement. It strengthens posture not only at login, but across every interaction, whether initiated by a person, a service, or an AI agent.

To meet the moment, organizations must take a broader and more strategic view of what identity is and what it enables.

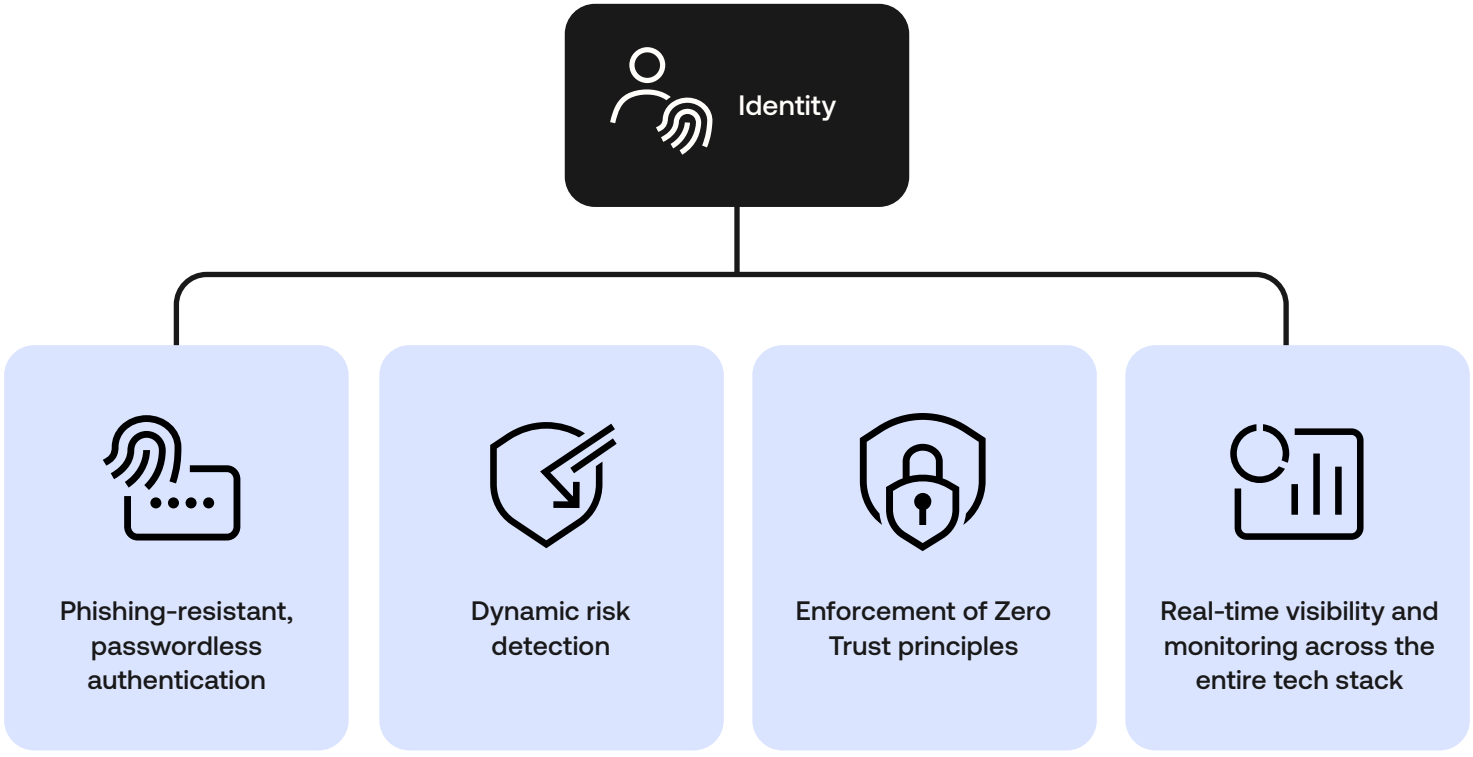

A closer look at the threat

Bad actors are counting on outdated approaches to identity. Fragmented IT and security environments scatter core resources, applications, and identities across disconnected systems and infrastructure, leaving human and non-human identities vulnerable to unnoticed and unaddressed attacks. As AI agents and automated systems gain access to more data and execute more actions, the consequences of weak identity controls grow more severe.

Manual processes

Cumbersome, time-consuming permissioning practices that leave access decisions vulnerable to human error and inconsistent policy enforcement.

Blind spots

Limited visibility into real-time security posture, identity behavior, and permissions across users, services, workloads, and AI agents.

Slow response

Delayed detection and remediation allow attackers or rogue automated processes to exploit identity-related vulnerabilities before controls can intervene.

This problem is not going away. It is accelerating. AI-powered applications, services, non-human identities, and AI agents expand the identity attack surface and increase the speed at which actions occur. Identity-related attacks no longer stop at stolen credentials. Post-authentication threats such as stolen session cookies, token replay, and privilege escalation introduce new challenges for security teams monitoring distributed environments.

Add the risks of well-funded nation-state actors, insider threats, and unmanaged autonomous systems, and the result is a rapidly evolving threat landscape centered on identity.

91%

91% of organizations are already using AI agents, but only 10% have a strategy to manage them

(Okta)

80%

Over 80% of all data breaches stem from Identity-related attacks

(Verizon)

180%

Identity-related attacks are increasing at a rate of 180% YoY

(Verizon)

1.9B

1.9 billion session cookies were stolen from Fortune 1000 employees in 2023

(Fortune)

okta

Identity is security

The risk landscape is full of multiplying threats that converge on identity. As organizations expand their use of cloud services, non-human identities, and AI agents, identity becomes both the primary attack surface and the primary control point.

Identity is not only a source of risk. It is also the greatest opportunity to establish control. By putting identity at the center of your security strategy, organizations can move from reactive defense to continuous enforcement, verifying access and automated actions before risk turns into impact.

When identity is unified across people, services, and AI agents, it becomes the foundation for preventing breaches and maximizing the value of your security and technology investments.

okta

The three principles of a modern identity strategy

Now that we've covered the stakes of taking an identity-first approach to security, let's turn to something more practical: how to move your organization's identity from where it is today to where it needs to be in an AI-driven environment.

Modern, cloud-native identity platforms provide a path to reduce fragmentation and centralize control. They deliver unified, real-time visibility across human and non-human identities, services, and AI agents, enabling IT and security teams to eliminate blind spots, surface risk, and respond faster.

This value can be organized into three core principles:

Comprehensive visibility

Verify that no vulnerabilities across users, non-human identities, or AI agents go unseen or unaddressed.

Powerful orchestration

Enforce real-time remediation and policy-based control across systems, services, and automated workflows.

Broad & deep integrations

Connect identity across your security and technology ecosystem to enable consistent governance and enforcement.

When evaluating the market, organizations should look for a platform capable of delivering across all three dimensions.

okta

Principle #1

Comprehensive visibility

Individually managing access permissions across applications, services, and systems creates exploitable security gaps and inconsistent policy enforcement. As non-human identities and AI agents multiply, these gaps expand beyond user accounts into APIs, workloads, and automated processes.

A modern identity platform must centralize and simplify access lifecycle management while delivering a comprehensive, real-time view of identity activity across your environment. It should provide unified visibility into human identities, service accounts, workloads, and AI agents, helping teams detect vulnerabilities early and enforce policy consistently at scale.

Core capabilities

Access governance and lifecycle management

Tools that centralize provisioning and deprovisioning across systems, automate joiner-mover-leaver processes, and continuously certify access for users and non-human identities.

Posture and threat visibility

Capabilities that monitor identity configurations, detect misconfigurations, and provide real-time insight into identity-related risk across applications, infrastructure, and AI-driven systems.

Privileged and sensitive access controls

Protections for high-impact identities, including administrative accounts, service accounts, and AI agents operating with elevated permissions.

Continuous risk evaluation

Real-time monitoring and signal aggregation that enables rapid detection and response to suspicious activity, whether initiated by a person or an automated system.

Checklist: Can your identity platform...

- ☐ Provide visibility into threats across users, non-human identities, AI agents, and customer accounts?
- ☐ Incorporate third-party signals from across your tech stack for comprehensive, real-time risk visibility?
- ☐ Continuously evaluate identity posture against Zero Trust principles?
- ☐ Identify misconfigurations such as inconsistent MFA enforcement, account sprawl, or over-privileged service accounts?
- ☐ Automate provisioning and deprovisioning when employees change roles or when non-human identities are no longer required?
- ☐ Enforce granular access controls for privileged users, services, and AI agents?
- ☐ Discover non-human identities and monitor their permissions and behavior?
- ☐ Integrate with HR systems and directories for centralized identity lifecycle management?
- ☐ Manage and secure customer identities at scale?

okta

Principle #2

Powerful orchestration

Fragmented security stacks generate vast amounts of risk data. Without a unifying, identity-driven control layer to analyze and act on that data, teams are left correlating logs across systems and reacting after damage has been done. The result is a slow security posture that cannot keep pace with automated workflows, non-human identities, and AI agents operating at machine speed.

A modern identity platform must enable organizations to prevent, detect, and remediate threats in real time. Beyond visibility, it must translate risk signals into automated enforcement actions, confirming that suspicious behavior, whether initiated by a user, a service account, or an AI agent, is evaluated and governed immediately.

Checklist: Can your identity platform...

- ☐ Simplify the configuration of automated remediation and policy enforcement actions?
- ☐ Customize responses based on risk signals, contextual data, and dynamic policy?
- ☐ Trigger protective actions such as step-up authentication, session revocation, or universal logout?
- ☐ Integrate with phishing-resistant authentication methods to strengthen ongoing protection?
- ☐ Evaluate device posture and contextual risk during active sessions?
- ☐ Block malicious IP addresses or anomalous activity in real time?
- ☐ Provide secure, self-service factor recovery without weakening security controls?
- ☐ Continuously evaluate identity activity after authentication, including automated and AI-driven actions?

okta

Principle #3

Broad & deep integrations

Your technology ecosystem is only as strong as the connections between its components. Without seamless integration across applications, infrastructure, security tools, APIs, and AI systems, organizations struggle to enforce consistent identity controls and realize the full value of their investments.

A modern identity platform must connect every part of the environment to enable consistent governance, risk monitoring, and enforcement. Vendor-neutral identity platforms make it possible to unify human and non-human identities across SaaS, cloud, custom applications, and AI-driven systems without placing unnecessary integration burdens on developer and IT teams.

Checklist: Can your identity platform...

- ☐ Integrate with core enterprise SaaS applications such as CRM, productivity, collaboration, ERP, and IT operations tools?
- ☐ Extend identity protection beyond provisioning and single sign-on to enforce control before, during, and after authentication?
- ☐ Integrate with security tools across your stack to strengthen threat detection, risk scoring, and automated remediation?
- ☐ Support APIs, services, and AI agents with consistent identity governance and policy enforcement?
- ☐ Provide no-code or low-code automation capabilities to trigger secure workflows across systems?
- ☐ Offer extensibility to support new applications, services, and AI capabilities as your environment evolves?

Outcomes of a unified security strategy

Unified, identity-first security is more than a concept. In practice, it delivers measurable outcomes that strengthen protection, streamline operations, and enable organizations to scale securely in a more complex environment.

Identity security & breach protection

Strengthen your security posture by protecting every human and non-human identity, enabling organizations to detect, contain, and remediate identity-based threats before they lead to breaches.

Operational efficiency & resilience

Simplify operations and reduce complexity by consolidating fragmented identity systems into a unified control layer that improves agility, lowers operational overhead, and accelerates time to value.

AI visibility & control

Establish governance for AI agents and automated systems by providing centralized visibility, policy enforcement, and secure access controls across emerging AI-driven workflows.

Learn more about the core outcomes of a unified security strategy.

okta

The path to identity-first security

Fully realized, identity-first security supports an open ecosystem that makes it secure and manageable to build, connect, and operate applications, services, and automated systems by default. No more identity silos.

No more costly and time-consuming custom integrations. No more security gaps and blind spots across workforce, customer, and machine-driven environments. Just a unified technology stack with identity at the center, enabling consistent governance and enforcement across every interaction.

Secure identity. Secure everything.

It cannot be overstated: identity is security. To stay ahead of threats and build resilient, long-term protection, security and IT leaders must modernize how identity is governed, enforced, and integrated across their environments.

When identity is unified across people, non-human identities, and automated systems, organizations reduce risk, strengthen operational control, and enable secure innovation at scale. Strong identity foundations protect against increasingly sophisticated attacks while supporting the agility required in today's distributed and AI-driven environments.

To start advancing your identity strategy and receive targeted recommendations from Okta experts, learn more or get in touch with our team.

About Okta
Okta, Inc. is The World's Identity Company™. We secure AI, machine, and human identity so everyone is free to safely use any technology. Our customer and workforce solutions empower businesses and developers to protect their AI agents, users, employees, and partners while driving security, efficiencies, and innovation. Learn why the world's leading brands trust Okta for authentication, authorization, and more at okta.com.

okta

Okta Inc.
100 First Street
San Francisco, CA 94105
info@okta.com
1-888-722-7871