movate

IDENTITY-FIRST SECURITY

The Foundation of Trust in the AI Era

Table of Contents

In today's AI-driven enterprises, a single compromised service account can cascade into a full-blown data incident: an attacker hijacks a privileged identity, then leverages an unmonitored AI agent embedded in the environment to quietly exfiltrate customer records, payment logs, or operational data.

This isn't a hypothetical war-game scenario; it mirrors real-world incidents where identity, not the firewall, turned out to be the first and last line of defense. In other words, as systems grow smarter and more interconnected, the weakest link is no longer infrastructure but who or what gets access to it.

Identity-first security treats every entity, including employees, third-party vendors, boots, and AI agents, as an identity that must be continuously verified, strictly privileged, and clearly governed. It replaces the old idea of a "trusted internal network" with a zero-trust reality: every access request is treated as potentially risky, and trust is granted only after continuous validation. In this context, identity-first security is not just a technical control layer; it is the operational foundation of trust for any organization that wants to harness AI without sacrificing security.

01. AI Security Challenges

As more teams use AI tools for email drafting, customer support, and internal reporting, **attackers are shifting from "breaking through the firewall" to "stealing the login."** Instead of attacking the network, they focus on identities, emails, passwords, and service accounts because that's what AI-driven workflows use to move data.

A common scenario: an employee clicks a phishing link in a fake "HR update" email, which installs a credential-stealing script. Within days, the attacker has a valid username and password, logs in as that user, and then uses an approved AI plugin or internal chat-with-data tool to pull sensitive customer lists or financial reports.

This matters because phishing is still one of the most common attack paths.

Google's own security guidance has long shown that account takeovers usually start with stolen credentials, not advanced hacking. In many real-world cases, just one exposed login can lead to access across email, documents, and connected AI tools.

02. Deepfake-style impersonation

Imagine a mid-size company where a project manager gets a video call from the "CTO" asking for urgent approval to change an admin password. The voice, face, and logo all look real, but it is an AI-generated deepfake. The employee almost complies until they notice the background looks slightly off.

03. Shadow AI in daily workflows

"Shadow AI" is also common in small and mid-sized companies.

For example, a marketing team may use an unsanctioned AI writing tool on a personal account and paste in campaign copy or customer emails. A developer may use a personal AI coding assistant and accidentally share API keys or database snippets.

Google-style workflows make this especially relevant because many teams already live in Gmail, Google Drive, Docs, and Meet. If an employee copies sensitive text from a Google Doc into an unapproved AI tool, that data can leave the company's control in seconds.

A simple example: a sales rep asks an AI tool to summarize a Google Sheet of customer contacts, but the sheet includes phone numbers, deal values, and notes. If that tool stores prompts or logs, the company may have no idea where the data went.

In short, AI is making identity the main attack surface; every employee, contractor, and AI-powered tool is now a potential "entry point," and the risk is not just from big-name hackers but from everyday missteps and unnoticed tools.

04. Core Principles of Identity-First Security

Identity-first security means verifying every person and every tool before granting access. A login should not be trusted just because it comes from inside the company network; it should also be checked against device, location, and behavior.

Google's security research shows why these matters:
adding a recovery phone number can block up to

100% of automated bots,

99% of bulk phishing attacks, and

66% of targeted attacks.

A simple example: if an employee signs in from a new country or an unfamiliar device, the system can request additional verification or block access. The same logic applies to AI agents and bots, which should only access approved data, at approved times, through registered service accounts.

Password less login also helps. Google has promoted stronger authentication methods like security keys because stolen passwords are still a major risk. In one well-known case, Google **reported that security keys** stopped phishing attempts across its workforce.

In practice, identity-first security is about making every login, and every AI tool prove it is trusted before it gets access. This continuous validation model transforms security from a one-time checkpoint into an ongoing process.

05. Enterprise Implementation

For most organizations, moving to identity-first security does not have to start as a full-blown redesign; it can begin with a few clear, repeatable steps rolled out gradually across teams and systems.

STEP 1

Inventory and classify all identities

Begin by mapping every "who" that accesses your systems: employees, contractors, service accounts, and any AI tools or bots. Many mid-sized enterprises discover they have far more non-human identities (scripts, bots, APIs, and cloud service principals) than actual employees, often by a factor of 2–5× or more. Recognizing this imbalance is the first step toward managing hidden risk.

*For example, a generic **cloud-based operations team** typically manages dozens of automated workflows, each using its own service account, with many of those accounts holding broad administrator-level permissions that have not been reviewed for several years. Inventorying these identities in an Identity Governance and Administration (IGA) or IAM console makes them visible, classifiable by risk (low, medium, or high), and governable under least-privilege policies instead of remaining "invisible" standing privileges.*

STEP 2

Deploy and centralize IAM
(for example, Okta or similar platforms)

Once identities are mapped, the next step is to move to a centralized identity and access management (IAM) platform. A common scenario is a company using multiple SaaS tools (email, HR software, CRM, and analytics dashboards) with separate logins, each managed ad hoc by different teams.

By integrating these tools with a single **IAM provider** (such as Okta, Microsoft Entra ID, or similar), employees get single sign-on (SSO), and admins can enforce consistent policies such as multi-factor authentication (MFA) for all users and stricter conditional-access rules for finance or admin roles. Centralization, therefore, becomes the backbone of scalable and enforceable security. This fragmentation often leads to inconsistent policies and overlooked vulnerabilities.

STEP 2

Automate access and reviews

Once IAM is in place, automation is essential. Without it, access reviews become manual approvals where reviewers rarely verify actual usage. Automation ensures that security keeps pace with the speed of modern enterprise operations.

Organizations can:

Sync HR systems like Workday or SAP with IAM so that access is automatically removed or updated when employees leave or change roles. For example, an employee's cloud access is revoked within hours of their status changing to "Terminated."

Implement just-in-time (JIT) **access so users only get elevated privileges** for a short window to complete a specific task, reducing standing admin rights. For example, an engineer requests admin rights for 2 hours to fix an issue; after that window, the rights are auto-revoked.

Use AI-driven or **analytics-based tools to flag rarely used**, overly privileged, or anomalous accounts, such as admin accounts suddenly active from unusual locations.

A managed-services provider using automated access reviews found that 15–20% of its admin accounts had not been used in six months and proactively disabled them, shrinking the attack surface for privilege-based attacks. These insights show how intelligence and automation work together to continuously refine security posture.

06. Benefits: Why Identity-First Pays Off

Adopting identity-first security **lowers breach risk by 40–60% through** verified identities, least-privilege access, and continuous checks that block lateral movement without slowing workflows, as seen in a logistics firm's 40% drop in threat events.

It simplifies compliance with automated IAM logs and traceable permissions, enabling a fintech startup to pass audits in half the time. Compared to traditional perimeter-focused approaches, it reduces AI-related risks by scoping agent privileges, cuts compliance overhead via automated reporting, and delivers clear ROI by making attacks harder, responses faster, and AI-driven growth predictable. In essence, security shifts from being a bottleneck to becoming a business enabler.

07. Conclusion

Identity-first security is no longer optional in an AI-driven world; it is the core framework that determines whether innovation can scale safely.

As identities, humans and machine, become the primary gateways to data and systems, organizations must shift from perimeter-based thinking to continuous verification, least-privilege access, and centralized control. By embedding these principles into everyday operations, enterprises can not only reduce breach risks and strengthen compliance but also build a resilient foundation where AI can thrive without compromising trust.

08. About the Author

Mushtaq Ahmad

Mushtaq has almost two decades of IT industry experience and is the global Chief Information Officer at Movate. With expertise in data center technologies, next-generation cybersecurity, cloud, and applications he has assumed various leadership roles and worked across the globe in geographies like the USA, Europe, and APAC.

As the CIO of Movate, he has set the organization's technology strategy and roadmap, and has been driving the organization's efficiency while creating a digitized ecosystem to elevate customer experience and service agility by collaborating with different stakeholders.

About Movate

Movate is an Applied AI services company that helps enterprises translate AI ambition into measurable business outcomes. With over 12,000 employees across 20 global locations, augmented by AI collaborator agents, it brings together human expertise and technology to help organizations rethink operating models for efficiency and growth. Powered by Mova iO, its intelligent outcomes platform, Movate delivers practical, scalable, and industry-contextualized solutions that address real business challenges.

For more details, please mail us at info@movate.com

www.movate.com

movate®