

EnforceAuth Launches out of Stealth

SAN DIEGO, CA, CA, UNITED STATES, February 6, 2026 /EINPresswire.com/ -- EnforceAuth today announced the launch of its AI-native Security Fabric, a decision-centric security platform purpose-built to govern autonomous software, AI agents, and machine identities at enterprise scale. As enterprises accelerate AI adoption without corresponding governance, the platform delivers continuous oversight, real-time policy enforcement, and full auditability to ensure every automated action is secure, compliant, and accountable.

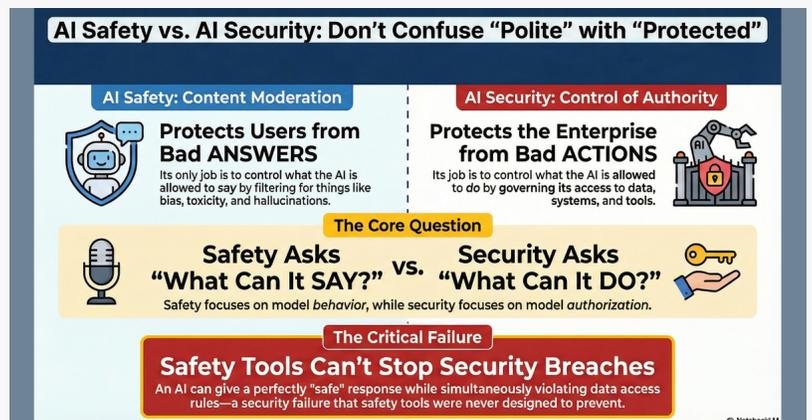
The \$3.4 Trillion Blind Spot Enterprise AI is no longer a pilot program. It's the operating system of modern business. By 2028, Gartner projects that agentic AI will be embedded in 33% of enterprise software applications. At least 15% of daily business decisions will be made autonomously — no human in the loop, no approval workflow, no audit trail.

And yet, the security industry is still solving yesterday's problem.

Today's identity and access management (IAM) platforms were built for a simple question: Should this person be allowed in? But the question has fundamentally changed. In an autonomous enterprise, the question is: Should this action be allowed to happen? That's a different problem entirely. And no one is solving it — until now.



ENFORCEAUTH



AI Safety vs. AI Security: Don't Confuse "Polite" with "Protected"

AI Safety: Content Moderation	AI Security: Control of Authority
 Protects Users from Bad ANSWERS Its only job is to control what the AI is allowed to say by filtering for things like bias, toxicity, and hallucinations.	 Protects the Enterprise from Bad ACTIONS Its job is to control what the AI is allowed to do by governing its access to data, systems, and tools.
The Core Question	
 Safety Asks "What Can It SAY?" Safety focuses on model behavior, while security focuses on model authorization.	Security Asks "What Can It DO?" 
The Critical Failure	
Safety Tools Can't Stop Security Breaches An AI can give a perfectly "safe" response while simultaneously violating data access rules—a security failure that safety tools were never designed to prevent.	

© NotebookLM

EnforceAuth AI Safety vs. AI Security: Don't Confuse "Polite" with "Protected"

Introducing the AI Security Fabric

Unlike traditional IAM systems that gate access at the perimeter, EnforceAuth's Security Fabric operates at the decision layer — evaluating every action initiated by an AI agent, automated workflow, or machine identity against policy, context, and risk tolerance in real time.

"Enterprise security has been playing defense at the front door while AI agents are already inside, making decisions, moving data, and triggering workflows without anyone watching," said Mark Rogge, Founder and CEO of EnforceAuth. "Our Security Fabric doesn't just authenticate who's in the building. It governs what every actor — human, agent, or machine — is allowed to do, in real time, with full accountability. This is how you secure the autonomous enterprise."

Why Now: The Governance Gap Is Already a Crisis

The numbers tell a stark story. AI adoption is nearly universal — 95% of U.S. companies are using generative AI and 88% report regular use. But governance has not kept pace: only 18% of enterprises have fully implemented AI governance frameworks. The result is what EnforceAuth calls the Authorization Gap — the dangerous delta between what AI systems can do and what organizations can verify they should do.

This gap is widening fast. The OWASP Top 10 for Agentic Applications, released in late 2025, codified an entirely new class of risks tied to autonomous action, identity sprawl, and runtime behavior. The EU AI Act's enforcement deadlines for high-risk systems begin in August 2026. And the average enterprise now faces an 82:1 machine-to-human identity ratio — most of those identities operating with excessive privileges and minimal oversight.

Responsible AI programs are beginning to respond: 69% of strategic-stage organizations now invest in real-time testing and observability for AI agents. But without a decision-centric security layer, these investments remain incomplete. Autonomous agents can still trigger unauthorized workflows, cross system boundaries, and execute high-stakes actions without proper scrutiny.

From Access Control to Decision Control: A New Category

EnforceAuth's Security Fabric represents a category shift in enterprise security — from identity-centric to decision-centric architecture. The platform delivers four core capabilities:

- Decision-Centric Architecture. Policies govern every action, not just access. The platform evaluates context, scope, and risk before authorizing execution — whether the actor is a human, an AI agent, or an automated workflow.
- Native Agentic AI Support. Purpose-built for AI systems and machine identities, the Fabric supports delegated authority, time-bounded permissions, and contextual decision evaluation — capabilities that static RBAC systems were never designed to deliver.
- Continuous Observability and Audit. Every decision is recorded and auditable in real time. Anomaly detection and compliance monitoring are built into the core, not bolted on.
- Operational Control and Accountability. Security teams define who — or what — can do what, under which conditions, creating clear authority chains and unbroken audit trails that satisfy regulators and boards alike.

Built by the Team That Wrote the Playbook

EnforceAuth was founded by Mark Rogge, a veteran enterprise security executive who served as CRO at Styra — the company behind Open Policy Agent (OPA), the open-source standard for cloud-native authorization that was acquired by Apple in 2025. Before Styra, Rogge held leadership roles at GitLab and Weights & Biases, where he helped scale the company to unicorn status.

"At Styra, we solved policy-as-code for the cloud-native era," Rogge said. "But the world has moved on. AI agents don't fit in a Kubernetes pod with a sidecar proxy. They reason, delegate, and act across system boundaries. Securing them requires a fundamentally different architecture — one that treats every decision as a governable event. That's what we've built."

Availability and Early Access

EnforceAuth is offering early access through a free enterprise waitlist. Select organizations will partner with the company to refine features, validate deployment patterns, and influence the future of AI governance. Fortune 500 enterprises in regulated industries — including financial services, healthcare, and critical infrastructure — are encouraged to apply.

For details or to join the waitlist, visit enforceauth.com.

About EnforceAuth

EnforceAuth is an AI-native security company headquartered in San Diego, California. Founded by enterprise security veterans from Styra (acquired by Apple), GitLab, and Weights & Biases, the company builds technologies to govern decisions made by autonomous software, AI agents, and machine identities. Its mission is to help enterprises embrace AI safely by ensuring every autonomous action is authorized, accountable, and auditable.

Contact

Mark Rogge, Founder & CEO EnforceAuth, Inc.

Email: mark@enforceauth.com

Web: www.enforceauth.com

Mark Rogge

EnforceAuth

+1 612-868-7193

[email us here](#)

Visit us on social media:

[LinkedIn](#)

[YouTube](#)

This press release can be viewed online at: <https://www.einpresswire.com/article/889776160>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.