EnforceAuth Delivers Coverage of Gartner AI TRiSM Framework, Closing the Authorization Gap

AI Security Fabric platform becomes the first solution purpose-built to enforce all four layers of the Gartner AI TRiSM model

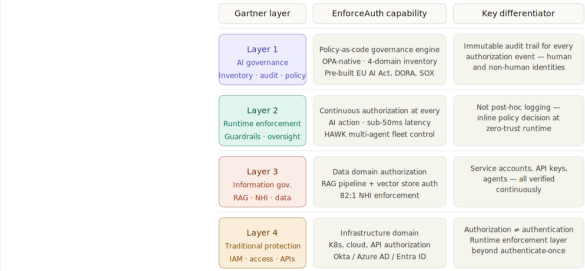

EnforceAuth's AI Security Fabric platform capabilities to all four layers of the Gartner AI Trust, Risk and Security Management (AI TRiSM) framework, as defined in Gartner's Market Guide for AI Trust, Risk and Security Management (February 2025).

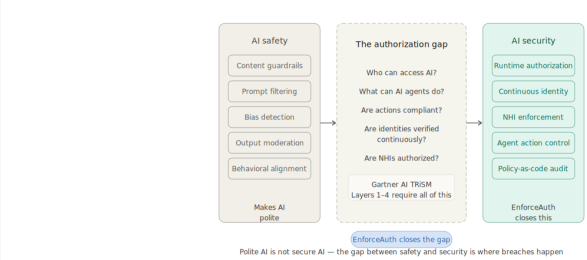

AI safety vs. AI security — the core thesis underlying EnforceAuth's AI Security Fabric platform.

SAN DIEGO, CA, UNITED STATES, March 18, 2026 /EINPresswire.com/ -- EnforceAuth, the AI Security Fabric company, today announced that its unified authorization platform delivers end-to-end coverage of all four layers of the Gartner AI Trust, Risk and Security Management (AI TRiSM) framework — the industry's most comprehensive model for governing enterprise AI systems. EnforceAuth is positioned as the enforcement layer that the AI TRiSM framework requires but no existing security vendor has provided: a real-time, policy-driven authorization engine that continuously verifies the identity and authorized actions of every actor across enterprise AI deployments, whether human or non-human.

The announcement comes as enterprise adoption of AI agents accelerates across regulated industries, and as organizations face a widening gap between AI safety investments — behavioral guardrails and content filters — and the runtime authorization controls required to comply with the HIPPA, PCI, EU AI Act, DORA, NIST AI RMF, and SOX. EnforceAuth refers to this gap as the Authorization Gap: the space between making AI polite and making AI secure.

"The Gartner AI TRiSM framework is the most rigorous articulation of what enterprise AI governance actually requires. When you map its four layers against what current security

vendors provide, a single gap appears in every layer: runtime authorization enforcement. That is precisely what EnforceAuth was built to provide. We are not retrofitting an IAM tool or bolting AI coverage onto a CNAPP platform. We built the authorization layer that the AI era demands — and we built it to map exactly to what Gartner's research says enterprises need."
— Mark Rogge, Founder and CEO, EnforceAuth

The Authorization Gap: Where AI Safety Ends and AI Security Must Begin

According to Gartner's Market Guide for AI Trust, Risk and Security Management (February 2025), organizations must address governance, runtime enforcement, information governance, and traditional access control as distinct and interdependent layers of AI security. Yet research from multiple industry sources indicates that security investments in enterprise AI have been heavily concentrated in AI safety capabilities — content moderation, prompt engineering, alignment techniques, and behavioral guardrails — rather than the runtime authorization controls that AI TRiSM's framework demands.

EnforceAuth defines this misalignment as the Politeness Trap: the assumption that a well-behaved AI agent is a secure one. A polite AI agent that passes every content filter can still access data it is not authorized to read, execute transactions it is not permitted to initiate, and call APIs with privileges it was granted during deployment and never had revoked. These are authorization failures, not safety failures — and they represent the primary unaddressed attack surface in enterprise AI environments today.

62 percent of AI practitioners cite security as the primary challenge in deploying AI agents.

48 percent of enterprise leaders identify security threats as a top-three barrier to AI implementation.

Non-human identities — service accounts, API keys, AI agents, and automated pipelines — now outnumber human identities in enterprise environments at a ratio of 82 to 1. Each of these identities requires not just authentication, but continuous authorization enforcement at the action level.

EnforceAuth's Full-Stack AI TRiSM Coverage: Layer by Layer

EnforceAuth's AI Security Fabric platform provides enforcement capabilities that directly address each layer of the Gartner AI TRiSM framework:

LAYER 1 — AI GOVERNANCE

EnforceAuth provides policy-as-code governance across all four enterprise domains — Applications, Infrastructure, Data, and AI Workloads — through a single, unified policy engine. Built on Open Policy Agent (OPA) with native compatibility for Cedar (AWS) and Zanzibar (Google), the platform delivers an immutable decision audit trail for every authorization event across both human and non-human identities. Pre-built compliance frameworks for the EU AI Act, DORA, and SOX enable enterprises to demonstrate continuous governance from day one, replacing manual policy documentation with enforcement-time evidence.

LAYER 2 — RUNTIME INSPECTION AND ENFORCEMENT
Where most vendors offer post-hoc logging and periodic reviews, EnforceAuth enforces authorization inline — at the moment of every AI action, at machine speed. Every request made by an AI agent, whether to read data, call an API, execute a workflow, or transact on behalf of a user, is evaluated against centralized policy in real time. Authorization decisions are rendered in sub-50 millisecond latency, supporting enterprise-scale AI deployments without performance degradation. Multi-agent fleet management through the HAWK architecture enables identity-level control across complex agentic systems operating simultaneously.

LAYER 3 — INFORMATION GOVERNANCE
EnforceAuth extends authorization enforcement into the data layer, including retrieval-augmented generation (RAG) pipelines, vector stores, AI training data repositories, and data lakes. Row-level and column-level access enforcement for AI workloads ensures that the data an AI agent can retrieve is governed by the same policy engine as all other enterprise access — eliminating the ungoverned data access pathway that RAG architectures introduce. With non-human identities including service accounts, API keys, and AI agents subject to the same 82:1 ratio enforcement, organizations can demonstrate continuous identity verification across every data access event.

LAYER 4 — TRADITIONAL PROTECTION
EnforceAuth integrates with and extends the traditional access control and IAM infrastructure that most enterprises have already deployed. Native integrations with Okta, Azure Active Directory, Microsoft Entra ID, Kubernetes, and major cloud providers (Azure, AWS, GCP) position EnforceAuth as the runtime enforcement layer above the authenticate-once model that traditional IAM provides. The fundamental distinction: IAM answers the question of who an identity is. EnforceAuth answers the question of what that identity is authorized to do, continuously, at every action, without trusting the initial authentication decision indefinitely.

“The industry has spent billions of dollars making AI polite. We have invested in alignment, guardrails, and content moderation. Those investments matter. But polite AI with admin-level access to enterprise systems is still a security breach waiting to happen. Authorization is the missing layer — and it has to be enforced at runtime, not documented in a policy framework that nobody checks until after an incident.”
— Mark Rogge, Founder and CEO, EnforceAuth

Mark Rogge
EnforceAuth
+1 612-868-7193
email us here
Visit us on social media:
LinkedIn