Nuggets Labs Publishes Enterprise AI Governance Framework as Industry Converges on AI Action Control

Vendor-neutral framework defines Action Governance - the missing control layer between AI access and AI execution - as regulatory and enterprise risk accelerate

LONDON, UNITED KINGDOM, April 13, 2026 /EINPresswire.com/ -- Nuggets Labs, the research arm of Nuggets, the trust infrastructure company for AI actions, today announced the Enterprise AI Governance Framework, a freely available, vendor-neutral model enabling enterprises to govern AI systems at the point of execution, not just access, and prove that every AI action was authorized.

The release comes as enterprises converge on the need to control AI actions at runtime. The EU AI Act's high-risk AI obligations take effect in August 2026, the Colorado AI Act becomes enforceable in June, and enterprise AI is moving rapidly from analytical use cases into autonomous execution - agents that initiate transactions, modify infrastructure and access sensitive records without persistent human oversight.

The framework was first covered by Biometric Update, which noted that Nuggets' model addresses a governance gap that "most [companies] cannot prove whether an AI-initiated action was authorized, by whom, or under what constraints."

The Shift from Access to Execution
Traditional identity and access management governs one question: can this actor reach this system? As enterprises move to autonomous AI, that question is no longer sufficient. The new question is: should this specific action - by this actor, under this authority, within these

constraints - be permitted to execute right now?

The Enterprise AI Governance Framework introduces Action Governance as the answer: a distinct control layer that sits between access and execution, evaluating identity, delegated authority, intent and policy constraints before an AI action takes effect.

"IAM tells you who got in. Action Governance determines whether what happened next was authorised." said Seema Khinda Johnson, Co-founder and CCO of Nuggets. "Most enterprises deploying autonomous AI today cannot produce a verifiable answer to that question. That is an untenable position as regulators and boards begin demanding one."

What Makes This Framework Different
While the developer community has rightly focused on runtime security and policy enforcement for agent systems, these approaches answer how AI actions are controlled, not whether they are authorised in the first place. The framework addresses a distinct challenge: enterprise-grade governance for regulated environments, built around accountability to regulators, auditors and boards.

Several characteristics set it apart:
Consent as a first-class governance primitive. The framework defines consent as an enforceable, signed, auditable artifact specifying what actions are permitted and on whose behalf - essential where the basis for an action must be provable, not just logged.

Cryptographic, portable audit evidence. The framework requires tamper-resistant proof that every action was authorized and executed within defined constraints. This is not a log. It is a cryptographic artifact that can be produced to a regulator, an auditor or a board.

Cross-cloud, cross-system neutrality. Action Governance must be consistent whether AI systems operate on AWS, Azure, GCP, SaaS platforms or on-premise. Governance scoped to a single cloud or framework is not enterprise governance.

Vendor-neutral and freely available. Built to complement existing investments in IAM, security and compliance infrastructure - not replace them - and designed for CISOs, CIOs, Chief Risk Officers and Heads of AI, not only the teams building agents.

The Governance Gap Is Now a Regulatory Gap
Most enterprise AI deployments are now entering High Risk and Critical Risk governance tiers without the infrastructure those tiers require. Under the EU AI Act, non-compliance for high-risk AI carries fines of up to €30M or 6% of global annual turnover.

"The question facing every enterprise deploying autonomous AI is not whether they need this control layer. The question is whether they establish it before something goes wrong, or after."

The Trust Stack
The framework is anchored by a four-primitive trust stack - Identity Authority Intent Action - and eight infrastructure primitives: identity, authority, intent, consent, policy, enforcement, verification and audit. It maps to NIST AI RMF, the EU AI Act and ISO AI governance standards, extending each with the execution control layer none currently address.

Availability
The Enterprise AI Governance Framework is freely available now at nuggetslabs.com.

About Nuggets
The trust layer for autonomous AI. Nuggets extends existing IAM and cloud infrastructure to make AI actions provable, auditable and compliant at the point of execution. Built for enterprises operating in regulated environments.

Dom Gilmore
Nuggets
contact@nuggets.life
Visit us on social media:
LinkedIn
