EnforceAuth Announces Integration NVIDIA OpenShell to Deliver Continuous Authorization Governance AI Agent Deployments

Somewhere in your enterprise, An AI agent is running right now that shouldn't be....

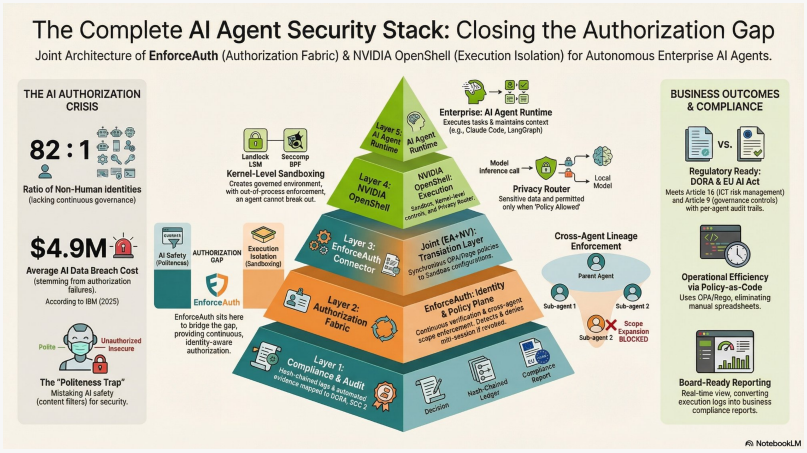

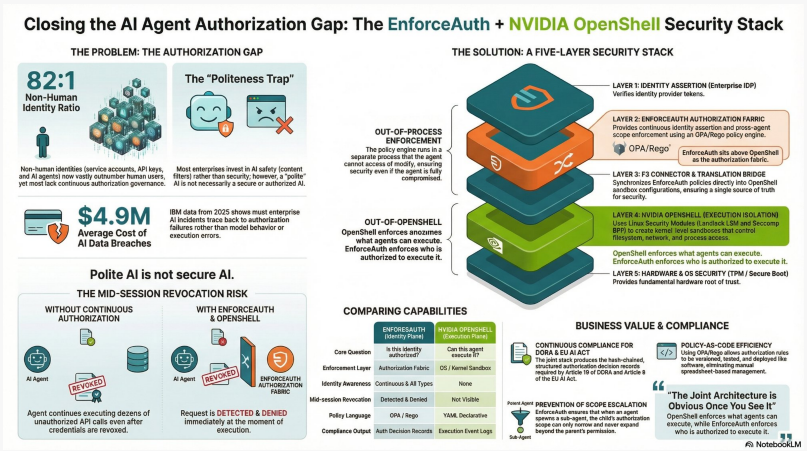

SAN DIEGO, CA, UNITED STATES, March 31, 2026 /EINPresswire.com/ -- EnforceAuth, Inc. today announced support for NVIDIA OpenShell, combining kernel-level execution isolation with continuous identity and authorization governance for enterprise AI agent environments. The integration — arriving in the coming weeks — addresses a structural gap that has remained unresolved in enterprise security architecture as AI agent adoption accelerates.

The announcement coincides with remarks from NVIDIA CEO Jensen Huang at GTC 2026, where he described the current period as the beginning of an inflection point. Industry analyst research has identified authorization governance across non-human and AI identities as an emerging priority for enterprise security teams.

Market Context

Enterprise AI adoption has created a class of non-human identities operating at machine speed, with permissions granted at provisioning and rarely revisited. Research indicates that 82 non-human AI identities exist for every human user in enterprise environments today, and the average cost of an AI-related data breach reached $4.9 million in 2025 (IBM) — with

authorization failure cited as the underlying cause rather than model behavior.

The security industry has historically addressed AI risk through endpoint protection and behavioral controls. Endpoint and network security platforms provide threat detection and execution monitoring. They are not designed to govern whether an AI agent remains authorized to act after it has been provisioned — a distinct and unaddressed problem. Behavioral guardrails and content filters address model outputs. They do not revoke credentials mid-session or enforce scope boundaries across agent-to-agent interactions. The result is an Authorization Gap: no existing security category continuously governs AI agent identity at the moment of action.

The Integration: OpenShell and EnforceAuth

NVIDIA OpenShell — an open source runtime that enforces policy-based security, network and privacy guardrails that make autonomous agents safer to deploy. — NVIDIA Official Press Release, March 16, 2026, GTC 2026, San Jose

NVIDIA OpenShell enforces what agents can execute — providing kernel-level sandboxing, filesystem and network control, and process isolation via Landlock LSM and Seccomp BPF. It answers the question: can this agent perform this operation?

EnforceAuth addresses the prior question: is this agent authorized to operate at all? The two systems answer different questions. Both are required for a complete security posture. Neither is sufficient on its own.

Five-Layer Architecture

The joint architecture operates across five layers:

- Layer 5 — Agent Runtime: Claude Code, LangGraph, OpenClaw, and other agent frameworks execute within a governed perimeter.
- Layer 4 — OpenShell: Kernel isolation via Landlock LSM and Seccomp BPF governs execution.
- Layer 3 — Connector: OPA/Rego policies are translated directly into OpenShell sandbox configurations.
- Layer 2 — Authorization Fabric: Continuous identity verification, cross-agent scope enforcement, and real-time mid-session revocation are applied to all human and non-human identities.
- Layer 1 — Compliance Record: Hash-chained audit logs are generated automatically, mapped to DORA, SOC 2, and EU AI Act Articles 9 and 16.

A key capability of the joint stack is scope escalation prevention. When a parent agent spawns a sub-agent, that child agent's authorization scope can only narrow, never expand. Execution-layer tools do not have visibility into what the parent agent was authorized to do. EnforceAuth enforces this constraint at every hop — closing a vulnerability class that runtime sandboxing alone cannot address.

QUESTION NVIDIA OPENSHELL ENFORCEAUTH
Core Question Can this agent execute this? Is this agent authorized to act at all?
-Identity Awareness -None -Continuous — all identity types
-Mid-Session Revocation -Not visible -Detected & denied in real time
-Sub-Agent Scope -Not enforced -Child scope cannot exceed parent
-Policy Language -YAML declarative -OPA / Rego — policy-as-code
-Compliance Output -Execution event logs -Hash-chained auth decision records

The two systems answer different questions. Both are necessary. The joint stack is the complete answer.

Statement from EnforceAuth
"Jensen Huang identified policy engines as a required component of enterprise AI at GTC — not firewalls, not content filters. Policy engines. That is the category EnforceAuth was built for. NVIDIA built the most powerful AI execution infrastructure in the world. EnforceAuth provides the authorization layer that makes it safe to run inside an enterprise. OpenShell governs what agents can execute. EnforceAuth governs who is authorized to execute it. Any enterprise running one without the other has a gap. That gap carries a $4.9 million average cost." — Mark Rogge, CEO and Founder, EnforceAuth, Inc., San Diego, CA, March 2026

Deployment and Availability
Any enterprise currently running NIM microservices, DGX Cloud, OpenClaw, LangGraph, or Claude Code is operating inside the Authorization Gap. EnforceAuth is in active deployment conversations with organizations across financial services, healthcare, and critical infrastructure — sectors where authorization failure carries regulatory penalties.
The NVIDIA OpenShell integration arrives in the coming weeks. A live technical briefing is scheduled for enterprise security leaders and AI infrastructure architects.

Key Statistics
- 82:1 Non-human AI identities to every human user in enterprise environments today
- $4.9M Average cost of an AI data breach in 2025 (IBM); root cause: authorization failure, not model behavior
- 0 Existing security vendors providing continuous authorization governance across the AI agent identity plane prior to this announcement

Mark Rogge
EnforceAuth
+1 612-868-7193

email us here
Visit us on social media:
LinkedIn
