LangGuard Launches Arbiter© to Enforce Agent Run-time Governance

The industry's first proactive agent action enforcement engine delivers deterministic run-time authority for your agentic workflows.

AUSTIN, TX, UNITED STATES, June 10, 2026 /EINPresswire.com/ -- LangGuard today announced the general availability of LangGuard Arbiter© at the Databricks Data and AI Summit for deterministic enforcement of agent actions. Earlier this year, LangGuard launched GRAIL Data FabricTM to give enterprises deep visibility into every action agents take. Introducing LangGuard Arbiter© that delivers the deterministic policy control over every agent action before it acts.

> Every enterprise is just days away from an agent incident they can't explain to their board. LangGuard Arbiter© provides the deterministic say enforced at the action surface, evidenced automatically"
>
> *Venkat Raghavan, Co-Founder, LangGuard Inc.*

In April 2026, an agent deleted a customer database in nine seconds after reasoning its way past a credential mismatch to a destructive API call. The agent had permission to access the system. No human was asked. No authorization existed for that specific action at that specific moment. The agent later documented every safety rule it had violated.

Two urgent imperatives are simultaneously impacting enterprise AI adoption.

The first is the Agent Action Surface. Agents reach production systems through an expanding universe of tools: MCP servers, REST APIs, CLI commands, SQL interfaces, and SaaS platforms going headless. Every system that once required a human to log in and click, or a pre-defined API connection, is now directly reachable by any agent that reasons it needs it. The action surface is no longer bounded. It is expanding faster than any enterprise can inventory them.

The second is the Agent Harness. Enterprises are deploying agents through Claude Code, Cursor, Codex, Cowork, Hermes, and internally built harnesses. Enterprise SaaS platforms are also activating embedded agents by default. Every one of these harnesses can reach the same expanding action surface by querying databases, modifying customer records, executing transactions, committing code, and calling external APIs, all without a human authority behind it.

Between these two imperatives sits the run-time authority gap: the absence of any mechanism that enforces human oversight over what agents actually do. Closing this gap requires two capabilities no enterprise has had together until now: a System of Actions that sees everything agents do across every session and system, and an enforcement engine that acts on what that system knows at the moment an agent decides to act.

LangGuard GRAIL Data FabricTM is a System of Actions for agentic workflow, purpose-built to record what agents do, not just what they access. Every action, every session, every system touched, in a continuous structured context graph. LangGuard Arbiter© enforces at the moment agents act.

LangGuard Arbiter© operates at the agent action surface, the layer between an agent's reasoning and the enterprise systems it acts upon. Every action an agent attempts is evaluated by Arbiter© before it reaches the target system. The outcome is deterministic: ALLOW, BLOCK, or ESCALATE to a human authority for decision. Every policy generated is immediately red-teamed against adversarial agent behavior before it is accepted. Only policies that pass become verified, creating a policy ledger with a provenance chain from compliance intent to enforcement. That record is the compliance artifact auditors ask for and enterprises have never been able to produce automatically.

With LangGuard Arbiter©, every enterprise team accountable for agent adoption now has a ready made control layer built for them:

Data & AI teams: Generate and enforce run-time policies aligned to enterprise, SOX, GDPR, and other regulatory needs, without writing rules for every possible agent-action combination.

IT teams: Detect and contain agent’s excessive agency and token maxing by flagging agents operating beyond the intended scope and budget overrun but prior to reaching production systems.

Audit teams: Generate Segregation of Duties (SoD) policies that prevent agents from executing conflicting actions in the same workflow and enforce the same SoD controls that govern human actions, now applied to agents.

Security teams: Identify and block unintended actions like the lethal trifecta where an agent's sequence of individually permissible actions produces an outcome that no one authorized.

The LangGuard Platform, Built on Databricks. Support any agent harness.

LangGuard Platform consists of three essential components, delivered in a single unified offering. LangGuard GRAIL Data FabricTM is the foundational System of Actions. SCOPE-MCP server maps what the agents can reach at design time. LangGuard Arbiter© enforces what they can do. Together they close the run-time authority gap from gaining visibility into agent action surface to enforcing actions in real-time.

A 14-Day Free Trial Offer at Databricks Data & AI Summit

Beginning June 15, LangGuard is offering a 14-day free trial of the entire LangGuard platform for current Databricks customers. Data & AI teams can deploy LangGuard as an App in your Databricks workspace and work with your favorite agent harness within minutes.

Visit LangGuard at booth #727 at Databricks Data and AI Summit in the Moscone center (June 15-18) to meet the founding team.

About LangGuard

LangGuard is the run-time action authority layer for enterprise AI. LangGuard is a Databricks-native app built for the humans accountable for what agents do: AI builders, Forward Deployed Engineers, IT, Compliance teams, and Chief Data and AI Officers. LangGuard is headquartered in Austin, TX with offices in the Bay Area and Canada. Learn more at www.langguard.ai

Ravi Srinivasan
LangGuard Inc.
+1 5122004345
email us here
Visit us on social media:
LinkedIn
YouTube
X

This press release can be viewed online at: https://www.einpresswire.com/article/918415736

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.