Trust3 AI Integrates with Databricks Agent Bricks to Give Enterprises Full Control Over AI Agents

SAN FRANCISCO, CA, UNITED STATES, June 17, 2026 /EINPresswire.com/ -- Trust3 AI today announced its integration with Databricks Agent Bricks, delivering an independent trust layer that gives enterprises unified visibility, governance, and control over AI agents, regardless of where they run. The integration equips organizations with the cross-platform trust enforcement, runtime observability, and token usage monitoring needed to deploy autonomous agents responsibly at production scale.

> “Trust3 AI provides the independent trust layer for agent ecosystems discovering shadow agents, validating behavior, enforcing governance, and delivering unified visibility across platforms.”
>
> *Neeraj Sabharwal, Co-Founder of Trust3 AI*

Databricks Agent Bricks is a governed enterprise platform for building and scaling AI agents, with support for model choice, MCP-based integrations, unified governance, and end-to-end monitoring. Trust3 AI extends that control plane into the real-world environments where agents rarely stay within a single system, running instead across multiple models, tools, data systems, agent platforms, and orchestration frameworks.

Trust3 AI also plans to integrate with Databricks’ new open-source project, OmniAgent, by providing a custom policy engine plugin. This extension allows organizations to implement tailored governance controls, ensuring that contextual policy enforcement remains robust and consistent across all custom agent workflows.

As enterprises move from pilots to production, agent sprawl is emerging as a significant operational and security challenge. Agents deployed outside approved processes create blind spots that traditional governance tools were never built to catch. According to a customer-backed account, one large oil and gas company identified shadow agents running outside approved processes within less than five minutes of deploying the Trust3 AI collector, and quickly shut down agents that were not supposed to be active.

The integration also strengthens token monitoring, treating token usage as more than a cost metric. Token activity functions as a trust signal: sudden spikes can indicate runaway agents, prompt abuse, misconfigured tools, or agents accessing data in unintended ways. By surfacing

these patterns in real time, Trust3 AI helps teams catch issues before they escalate into operational failures or security incidents.

Trust3 AI delivers an independent, platform-agnostic trust layer for enterprise AI systems, built specifically because agents do not stay within a single stack. As organizations run agents across multiple models, tools, and orchestration frameworks, they need trust enforcement, runtime observability, and governance controls that follow the agent wherever it operates, not controls locked inside any one platform.

A core part of that governance framework is the Trust3 AI Trust Score. Every governed agent carries a numeric Trust Score on a 1.0 to 10.0 scale that summarizes how well it meets the organization's governance requirements. The score gives IT, compliance, and risk teams a shared, real-time signal for prioritization without requiring them to manually review individual policy results for every agent in the environment. Published score bands map directly to trust and risk levels: agents scoring 7.75 and above are classified as High trust with Limited risk, while scores below 4.15 indicate Critical trust gaps and High operational risk. Scores are not static.

They drop when policy violations go unresolved, when agents lack assigned ownership, when documentation is incomplete, or when unapproved models are detected. They recover as teams remediate violations, assign owners, and complete required fields. This continuous scoring creates a clear, measurable feedback loop between governance action and agent risk posture.

Combined with the agent-building and optimization capabilities of Databricks Agent Bricks, the integration allows enterprises to innovate quickly while maintaining the governance, security, and compliance their environments demand.

About Trust3 AI

Trust3 AI is the agent control plane for the enterprise. Founded by Balaji Ganesan, Don Bosco Durai, and Neeraj Sabharwal, the company delivers AI-powered governance for data, AI, and access intelligence through its One Control Plane architecture, powered by a Unified Trust Layer built on AI-native metadata and context. The platform enables enterprises to discover every agent, observe every decision, and secure every action across any framework, cloud, and data source. Trust3 AI is a member of the NVIDIA Inception Program and the Snowflake Startup Accelerator, and plans to join the Databricks Built On program.

Book a demo on https://trust3.ai/demo/

Ibby Rahamani
Trust3 AI
+1 510-413-7300
email us here
Visit us on social media:

LinkedIn