

One Control Plane for Agents & Data in Databricks

Trust3 AI brings complete visibility, protection, and observability to every Databricks agent, endpoint & dataset.

TRUSTED BY

STARBUCKS

INTUIT

AUTODESK

ROBLOX

SAMSUNG

CVS

GENERAL DYNAMICS

THE PROBLEM

AI Agents Are Outpacing Enterprise Governance

Role & Permission Sprawl

Access control doesn't scale for AI. Managing permissions across workspaces, serving endpoints & Mosaic AI agents at scale becomes unworkable.

Agent-Level Enforcement Gap

Mosaic AI agents multiply fast. Governing access by declared purpose, not just identity or role, is the critical next layer for production-ready AI.

Shadow AI Blind Spots

Agents proliferate faster than any team can track. Enterprises undercount their deployed agents by 3 to 10x.

Unprotected Agent Protocols

MCP is becoming the standard for agent-to-tool connections. Every MCP server is an emerging attack surface.

73%

of enterprises have no complete AI agent inventory*

100+

AI agents in a typical Fortune 1000 enterprise*

Dec 2

2027 EU AI Act full enforcement deadline

The challenge: AI agents are accessing sensitive data faster than any team can govern manually. Visibility, enforcement at the agent level, and audit trails for regulators all need a new layer.

THE SOLUTION

What Trust3 AI Does

PBAC + RBAC Access Control

Layer Purpose-Based Access Control over Unity Catalog permissions and data tags. Apply row filters, column masking, and dynamic scoping at query time. Grants auto-expire.

01

Genie Space Scoping

Auto-scope every natural language query by user identity and declared purpose. A Finance user never surfaces HR data. One policy replaces dozens of workspace roles.

02

Databricks Agent Governance

Discover, score, and govern every Agent Bricks workflow, serving endpoint, and RAG pipeline, including shadow AI. Trust Score gates production.

03

Runtime Observability

Trace every agent action, prompt call, and tool invocation. Baseline in-scope behavior, flag out-of-scope automatically. Catch model drift before it becomes an incident.

04

Audit-Ready Inventory

Catalogue every agent with identity, purpose, risk, and approvals. One-click evidence packs for EU AI Act, HIPAA, GDPR, and NIST AI RMF. Query it all in plain English with GIA.

05

MCP Threat Protection

Treat every MCP connection as untrusted. Strip injections, enforce scope minimization, maintain tamper-evident audit logs per agent.

06



ASK GIA · GOVERNANCE INTELLIGENCE AGENT

Govern your AI agents across Databricks in plain English

"Grant Risk Analytics agents access to transaction data for Q4 audit, expiring January 15."

HOW IT WORKS

Governance From Inventory to Runtime Control

STAGE 01

Discover

Connect natively to Databricks. Automatically inventory every agent, serving endpoint, RAG pipeline, and MLflow model, including shadow AI surfaced via logs.

STAGE 02

Observe

Score every agent against 10+ pre-built policy packs (EU AI Act, NIST, HIPAA, GDPR, NERC CIP). Trace prompts, tool calls, and data access.

STAGE 03

Secure

Enforce row and column controls natively via Unity Catalog. Block unauthorized agents at runtime. Extend governance through MCP and A2A chains.

TRUST SCORE SYSTEM

A Composite Rating, Kept Current

Every agent gets a Trust Score across four governance dimensions, updated automatically

Policy Compliance

35%

Adherence to assigned governance policies.

Scope Adherence

25%

Operating within defined boundaries and permissions.

Security Posture

25%

Credential management and access patterns.

Behavioral Baseline

15%

Consistency with learned behavioral patterns.

WHO IT'S FOR

Built for IT, Security, and Compliance



IT / CIO Office

Unified inventory across every platform instead of fragmented admin consoles. Prove governance posture to the board with a monthly Governance Posture Report.



CISO / Security

Runtime enforcement blocks violations before they occur. Trust Score surfaces at-risk agents instantly. Preventive controls auditors prefer over detective alerts.



Compliance & GRC

Continuous policy monitoring with reproducible, timestamped evidence for NIST AI RMF, EU AI Act, GDPR, HIPAA, and SOX audits. Audit-ready at any moment.



Operations & Platform

Register agents, connect platforms, and manage approval workflows. Trust3 auto-discovers new agents so governance stays current without manual tracking.

WHY TRUST3

Key Differentiators

10+

Governance Agents

A2A workforce orchestrated by GIA. Author policies, classify data, score roles, and surface violations automatically.

[Author](#)[Classify](#)[Score](#)

10+

Data & AI Compliance Packs

Shipped with the product and kept current as standards Evolve.

[NIST](#)[EU AI ACT](#)[GDPR](#)[HIPAA](#)[SOX](#)[+ MORE](#)

100%

Audit Trail Coverage

Every prompt, tool call, and data query recorded with its policy evaluation.

[Prompts](#)[Tool Calls](#)[Data Access](#)[Evidence](#)

"By 2027, *60% of firms may fail to realize AI value* due to poor governance."

GARTNER · 2025

**Trust3 AI**

Built on Privacera Heritage

The only AI governance vendor with production-grade policy enforcement lineage across Databricks, Snowflake, and cloud data platforms.

Request a Demo

See Trust3 AI governing agents in your stack.

demo@trust3.ai
www.trust3.ai

Request a Demo →

* Based on an internal survey of 437 technology experts.