Trust3 AI

One Control Plane for Agents & Data in AWS

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

TRUSTED BY

STARBUCKS · INTUIT · AUTODESK · ROBLOX · SAMSUNG · CVS · GENERAL DYNAMICS

THE PROBLEM

AI Agents Deployment is Outpacing Enterprise Governance

Role & Permission Sprawl
Fine-grained controls span S3, Redshift, and beyond. Purpose-level grants become the bottleneck as datasets and agents multiply.

Agent-Level Purpose Enforcement
AgentCore agents, SageMaker endpoints, and Lambda tools multiply fast. Governing by declared purpose is essential for production AI.

Shadow AI in Your AWS Environment
Agents proliferate faster than teams can track across the AWS estate & custom builds. Enterprises undercount deployed agents by 3 to 10x.

MCP & A2A Protocol Exposure
MCP is the new standard for agent-to-tool connections. Every MCP server is a boundary that needs verification and an audit trail.

3,000+
roles in a typical enterprise of 1,000 employees

73%
of enterprises have no complete AI agent inventory*

Dec 2027
2027 EU AI Act full enforcement deadline

The challenge: agents & data are multiplying across AWS faster than teams can govern manually. Purpose-aware access, agent-level enforcement, and audit trails for regulators all need a new layer.

THE SOLUTION

What Trust3 AI Does

01

PBAC + RBAC for AWS Services
Layer Purpose-Based Access Control over data tags, IAM, and Lake Formation. Apply row filters, column masking, and dynamic scoping at query time. Grants auto-expire.

02

Bedrock & AgentCore Scoping
Auto-scope every Bedrock query and AgentCore agent by user identity and declared purpose. A Finance agent never surfaces HR data. One policy replaces dozens of IAM grants.

AWS Agent Governance

03

Discover, score, and govern every AgentCore agent, SageMaker endpoint, and Lambda tool, including shadow AI from CloudWatch. Trust Score gates production.

Runtime Observability

04

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

Audit-Ready Inventory

05

Catalogue every agent & dataset 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.

MCP Threat Protection

06

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

ASK GIA · GOVERNANCE INTELLIGENCE AGENT

Govern agents and data across AWS in plain English

“Grant Risk Analytics agents access to Redshift transaction data for Q4 audit, expiring 01/15.”

HOW IT WORKS

Governance From Inventory to Runtime Control

STAGE 01

Discover

Connect natively to AWS. Inventory every dataset, agent, SageMaker endpoint, RAG pipeline, and Lambda tool, including shadow AI surfaced via CloudWatch.

STAGE 02

Observe

Score every role and agent against 10+ pre-built policy packs (EU AI Act, NIST, HIPAA, GDPR, SOX). Trace every query, prompt, data access, and tool call.

STAGE 03

Secure

Enforce row & column controls natively via Lake Formation. Block unauthorized agents at runtime. Extend governance through MCP & 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.

Trust3 AI

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 · Review

20+
AWS Native Integrations
Across data & AI. Supported by AWS Data & Analytics Competency Partner status.
S3 · Redshift · Lake Formation · Bedrock · AGenCore · SageMaker

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
AWS Data & Analytics Competency Partner with production-grade policy enforcement across AWS, 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.*