




AI Agent Governance

Drata discovers every AI agent running in your environment, enforces your policies before an action executes, and produces auditor-grade proof of every decision—so you can answer the questions your board, auditors, and customers are already asking.



Why Drata

Over 8,500 customers trust Drata to continuously monitor their controls from the inside out, and governing agents is a natural extension of how we've built the trust layer between great companies.

Drata's AIQA has processed over two million security questionnaire questions. That data forms the basis of our Trust Graph and uncovered the rise of agent governance as an emerging need. As boards, auditors, and customers begin asking how enterprises govern their AI agents, five questions come up again and again—and today, most companies can't answer them.

Enforce Your Policies Before an Agent Executes

89% of companies can't answer how their AI agents are governed—and only 11% of vendors can prove an audit trail for AI agent decisions. Every CISO now has to answer five questions:

- What agents are running in your environment?
- What permissions do they have?
- Who is running them?
- Are they behaving as expected?
- Can you prove your answers to questions 1-4?

Source: Drata Trust Graph (2.1M security questions over six months)



AI GOVERNANCE OUTCOMES YOU CAN MEASURE

Discover

You know agents are running in your environment, but you don't know how many or what they're doing. The Drata Sensor runs inline and registers every agent at inception, mapping each one to its owner, identity, permissions, and scope. You get a full inventory of all agents in your environment, in minutes.

Govern

Monitoring only tells you what happened after the fact—and notifications aren't governance. Drata Mission Control compares every action to your policies and blocks violations in real time, before they execute.

Monitor

Behavior drifts, and point-in-time monitoring can't keep up with agents that outlive the session that generated them. Drata continuously monitors every command, prompt, and tool call against the policies your team actually sets. The moment an agent operates outside its scope, drift detection flags it to you.

Prove

Agents don't leave an audit trail of their decisions—which leaves you without answers when an auditor asks. Drata logs every decision in a tamper-evident record. In the same platform where thousands of customers run their audits today, Drata ensures agent activity maps to existing frameworks.

FEATURED CAPABILITIES



Drata Sensor

Discover and register every agent at inception, mapping each one to its owner, identity, permissions, and scope.



Mission Control

Define what each agent is allowed to do with policies written as intent, then enforce them inline against every action.



Trust Ladder

Advance each policy from Training to Recommendation to Active on your timeline. Prove it against real traffic before you turn enforcement on.



Inline Enforcement

Block policy violations before they execute, instead of reacting to after-the-fact notifications.



Drift Detection

Use continuous monitoring to flag the moment an agent operates outside its approved scope.



Chain of Custody

See every decision logged in a tamper-evident record, mapped to existing frameworks.



UNLOCK THE POWER OF AUTOMATION

