entro | SOLUTION BRIEF

Securing Agentic AI

Govern every AI agent.
Secure every action.

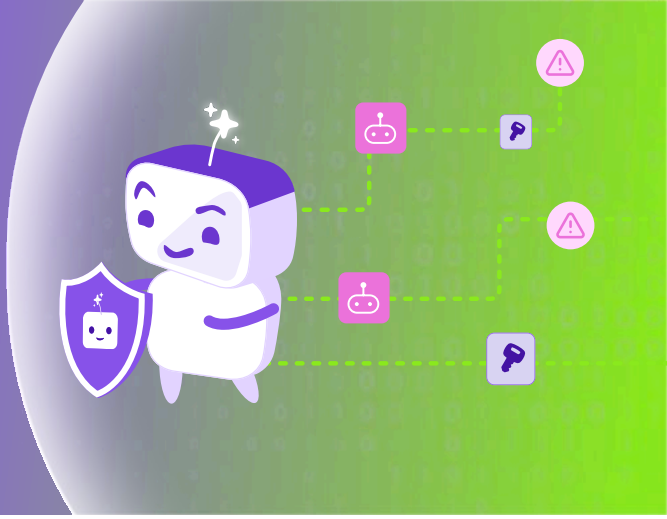

Entro governs AI agents and the identities behind them wherever they operate, across AWS, Azure, GitHub, Claude Code, Salesforce, MCP-connected tools, and beyond.

Discover shadow AI, monitor agent intent through AIDR, and enforce governance with the **Agentic Governance & Administration (AGA)** framework, so your organization can adopt Agentic AI at scale, without losing control.

The Challenge

AI agents now operate autonomously across your cloud, SaaS, endpoints, and developer toolchains. They spin up without approval, connect to sensitive systems, and take actions no-one authorized.

The Risk

- Uncontrolled agent sprawl across environments
- No audit trail for agent actions and tool calls
- Intent and behavior with no governance layer

The Solution

Entro discovers every AI agent in your environment, monitors actions and their intent. This is achieved by enforcing governance across every agent and the identities behind them.

The Capabilities

- Discovery & Inventory of every AI Agent
- Agent ownership attribution and mapping
- Intent monitoring through AIDR
- Full agent lifecycle governance through AGA granular policies

> "12% of firms are adding up to 10,000 AI agents or machine identities per month."
>
> **SailPoint, 2025 (via TechRadar)**

Monitored Sessions

AI Client	Intent	Prompt	Need attention
Claude	Secret Credential Harvesting	Query Jira for exposed plaintext secrets	Malicious

2 Malicious Sessions

13 Sessions

Last Session: 2min Ago

entro

Agentic AI Lineage Map

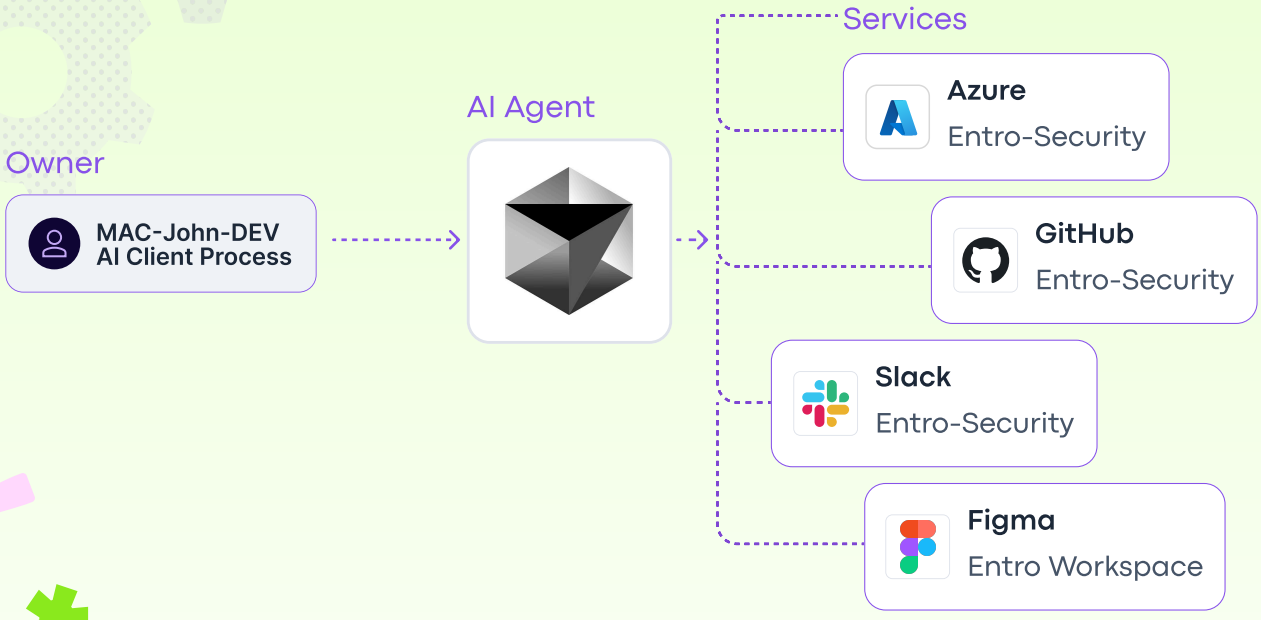

Stage	Capabilities	What Entro Does
Discovery	Agent And MCP Inventory Cloud, SaaS, And SDLC Coverage Endpoint Visibility (EDR Integrations) Shadow AI Discovery	Entro discovers AI agents and MCP servers wherever they run, across cloud, SaaS, agent platforms, and the SDLC. **Native EDR integrations like CrowdStrike extend coverage to the endpoint,** where developers and attackers operate. Shadow agents deployed without approval get surfaced and accounted for.
Identity Enrichment	Lineage Map Ownership Attribution Permission Analysis	Entro builds a full identity profile for every AI agent. **The lineage map traces each AI Agent t**o its NHIs, resources, and accountable owner. Permission analysis reveals overprivileged access and potential blast radius. Security and IAM teams know exactly what each agent can reach.

entro

Stage	Capabilities	What Entro Does
AIDR (Threat Defense)	Intent Monitoring Behavioral Anomalies (NHIDR) Alerting Posture Risks Remediation Workflows	AIDR monitors agent intent through prompt, session, and MCP activity. **NHIDR detects behavioral anomalies at the identity layer, since agent actions flow through NHIs.** Posture risks are flagged, and remediation workflows guide security teams to resolution.The MCP Audit plugin captures sessions across major AI agents including Claude, Cursor, and Gemini, tracking the MCP servers each contacts.
Governance	Policies Zero Trust Just In Time Agentic Access Administration (AAA)	**AGA is Entro's framework for governing AI agents at enterprise scale.** The AI Dashboard provides visibility into every agent and its permissions. **Agentic Access Administration (AAA) is the active enforcement layer:** granular policies define which AI client can perform which actions on which target services. For example, blocking Claude from running DELETE actions on a GitHub MCP. Idle agents and unused permissions are surfaced for decommissioning.

entro

Ready to govern your AI agents?

Entro gives security teams full visibility into every AI agent and the identities behind them, with real-time intent monitoring, real-time threat detection, and end-to-end lifecycle governance.

Discover shadow agents, audit every action, enforce least privilege, and scale agentic AI without losing control.

Get a Demo