



CAPABILITY BRIEF · 2026

# Govern the AI. *Prove it.*

The runtime control plane for AI governance and non-human identity security.

# Every enterprise has lost track of who is acting on their behalf.

Non-human identities — service accounts, API keys, agentic workloads — now drive the majority of cloud-native activity. AI agents make decisions, call tools, and move data with no governance layer between them and the systems they touch.

**100:1**

Non-human identities now outnumber humans in cloud-native organizations.

Palo Alto Networks, 2026

**70%**

Of identity incidents in 2025 originated from AI-related activity.

Microsoft Secure Access Report, 2026

**68%**

Of organizations cannot distinguish AI agent activity from human.

CSA + Aembit Study, 2026

Existing identity tools were built for people. Existing AI guards were built for models. Nobody owns the intersection — until now.

# One platform. Two capabilities. Six agents.

ModelCop *is* the integration. Built on the principle that every prompt has an identity, and every identity has a chain of accountability — from request to model to data to the human who owns it.

## 01 · RUNTIME AI GOVERNANCE

## Every prompt, governed at the wire.

*Sub-50ms p99 added latency. Transparent to applications. Hard enforcement, not advisory.*

- 11 LLM providers governed inline — Anthropic, OpenAI, Bedrock, Azure, Vertex, and more
- 8-class data classifier: PII, PHI, PCI, ITAR, CUI, secret, internal, public
- Hard block, redact, warn, or log — with HTTP 451 and a forensic record on every decision
- Policy engine surfaces explainable decisions to investigators, auditors, and the LLM caller

## 02 · NHI IDENTITY SECURITY

## Every identity, with a chain of accountability.

*Every prompt grounded in the identity that requested it. Every identity grounded in a human owner.*

- 41 connectors auto-discover NHIs across Okta, Entra, GitHub, Vault, AWS, GCP, Azure
- Tier 3 enforcement: rotate, disable, or revoke against the source system — not just alert
- Behavioral baselines, just-in-time access, ephemeral broker tokens for AI agents
- Single-query forensic chain: prompt → NHI → human owner, in three clicks

● 03 · BUILT TODAY

# Demonstrable today. Production-ready in six months.

ModelCop is not a slide deck. It is a working platform with the integrations, dashboards, and compliance mapping in place. Active design partners onboarding now.

158

Backend API endpoints

16

Persona-tuned dashboards

4

Industry tenants seeded with data

41

Identity connector kinds

11

LLM providers governed inline

23

Compliance frameworks mapped

● 04 · COMPLIANCE COVERAGE

**NIST AI Risk Management Framework**

Govern, Map, Measure, Manage

**EU AI Act**

Articles 9, 10, 12, 13, 15

**NYDFS Part 500 · DORA**

Financial services cybersecurity

# Built by a CISO *for* CISOs.

The category leaders in AI runtime security and non-human identity were acquired in 2025 and 2026. The buyers are bundling those capabilities into broader platforms. Mid-market and enterprise buyers want a focused vendor with deep CISO empathy — not a feature on a sprawling stack.

## David Stanton

*Founder & CEO · ModelCop LLC*

ModelCop was built in response to the question David's clients kept asking through 25+ years of consulting: *"Are there any proven AI security and governance solutions for non-human identities?"*

- 25+ years in technology, compliance, and cybersecurity professional services
- 9+ years in interim and CISO leadership across regulated industries
- Big 4 / top consulting: PwC · Accenture · Protiviti
- Direct relationships with Fortune 1000 CISOs across banking, healthcare, manufacturing, and defense

# Tell us what you're trying to govern.

We take a small number of design partners each quarter. Reach out and we will set up a 30-minute call to walk you through the platform and understand your environment.

**EMAIL**

hello@modelcop.ai

We reply within one business day.

**WEB**

modelcop.ai

Capability overview · contact form.

**ENTITY** ModelCop LLC · Dallas, Texas  
**SOC 2 TYPE II** In progress · Q4 2026

**TRADEMARK** USPTO Serial 99841440  
**ISO 42001** Controls implemented