FOR CISOS · CIOS · CDOS · RISK & THE BOARD

RuntimeAI

One platform to govern, secure, and prove every AI agent in your enterprise

Your teams are adopting AI — coding assistants, agents, and chatbots — faster than security can govern it. RuntimeAI gives you a single control plane to see it all, keep it in policy, stop it instantly, and prove it to your board and regulators.

THE STAKES FOR THE BUSINESS

Ungoverned AI Is Now a Board-Level Risk

📊

A reportable incident waiting to happen

AI agents touch customer data and make decisions with no identity, policy, audit, or off-switch. Under most breach-notification regimes, that's a disclosure obligation and a clock already running.

🔎

You can't govern what you can't see

Most enterprises run several times more AI activity than security knows about. You cannot put it in policy, report on it, or stand behind it at an audit until you can see it.

🌐

Data residency is now a board question

Regulators increasingly expect data residency and provable governance. AI that quietly sends customer data to outside models puts the company — and the board — on the hook.

Recent public incidents show the cost of the blind spot: **Comcast** *($117.5M settlement),* **ADT** *(5.5M customers breached), both with best-in-class security stacks and no agent runtime governance.*

ONE PLATFORM

Five Outcomes, Not Five Tools

Know	Control	Stop	Protect	Prove
See every AI agent, model & tool running across the enterprise	Your policy enforced automatically on every AI action	Shut down a rogue agent in under 50 milliseconds	Customer data masked at the source — never leaked to outside models	A tamper-proof record — compliance evidence on demand

WHERE TO START — SAME PLATFORM, ONE URGENT ENTRY POINT

AI Coding Agents	Govern Copilot, Cursor & Claude Code across your engineering org	**Shadow AI**	Discover the AI already running before your next audit finds it
AI Middleware	One governed layer between your apps and every LLM	**Agent Fleet Governance**	Identity, policy, and kill switch for every deployed agent

WHAT YOU GET

Outcomes Your Board Will Recognise

✅ **Complete AI visibility**
One inventory of every AI agent, tool, and model — across cloud, developers, and SaaS.

🛡️ **Control without slowing teams**
Your rules enforced automatically — teams keep moving fast, security keeps the guardrails.

⚡ **Contain any incident instantly**
Stop one agent, a class, or all AI activity in under 50ms — before damage spreads.

🔒 **Data stays where you need it**
PII masked before it leaves; deployable air-gapped, on-prem, or fully within your own infrastructure for data residency.

📋 **Audit-ready on demand**
A tamper-proof trail turns the next compliance audit from a scramble into a download.

🚀 **Live in weeks, not quarters**
Overlays your existing AI — start with the most urgent area, expand across the estate.

ONE PLATFORM — THE CAPABILITIES INSIDE

AI Discovery · Agent Identity (KYA) · AI Firewall · MCP Gateway · LLM Broker · PII Shield · Kill Switch · QuantumVault · Memory Vault · Audit Black Box · Cost Intelligence · Compliance Hub

Full architecture & capability detail in the RuntimeAI technical briefs (for architect reviews).

WHY NOW

AI governance is now a **board agenda item**, not just a security concern — regulators are actively enforcing breach-notification and data-residency requirements, and the enterprises that get ahead of it turn governance into a competitive advantage and a clean audit.

👤 **Roshan Shaik**
Founder & CEO, RuntimeAI

25+ years in cybersecurity, with leadership roles at **Okta**, **Cisco**, and **Equinix**; **BITS Pilani** alumnus. Roshan now leads RuntimeAI, the AI Governance & Control Plane — bringing identity, policy, kill-switch, and audit to enterprise AI.

linkedin.com/in/roshanshaik

Bring your AI under control — in 30 minutes.

A strategic briefing tailored to your enterprise — the risks, the outcomes, and the fastest first win. No deep tech required.

Roshan Shaik
Founder & CEO · RuntimeAI

roshan@runtimeai.io
calendly.com/roshan-runtimeai/30min