EXECUTIVE BRIEFING

The Control Plane for the Autonomous Economy

AI agents are running in your infrastructure today — buying services, executing workflows, accessing sensitive data — with no identity, no governance, and no kill switch. RuntimeAI closes that gap.

THE PROBLEM

Three Gaps No Existing Stack Closes

🔍

Shadow AI Proliferation

Most enterprises don't know what AI agents are running inside their own infrastructure. Agents operate without declared identity, registry, or governance.

1M+ AI Endpoints

Exposed on the internet · scan · No auth · No audit trail

⚡

No Runtime Governance

Traditional security stacks — SIEM, EDR, Zero Trust — weren't built for agents. They can't baseline agent behavior, issue agent identities, or enforce fail-closed policy at agent runtime.

50ms Window

Time for an agent to exfiltrate data, corrupt a DB, or escalate access

💀

Breach Cost Is Accelerating

Canvas: 275M records, $12M ransom, 6-day dwell. ADT: 5.5M customers breached. Comcast: $117.5M settlement. All had best-in-class security stacks. None had agent runtime governance.

$117.5M+

In direct damages — ADT + Comcast alone — no kill switch in place

RECENT INCIDENTS

The Breach Record

COMPANY	INCIDENT	GOVERNANCE GAP	RUNTIMEAI RESPONSE
Palo Alto Networks	PAN-OS zero-day RCE (CVSS 10.0) actively exploited — attackers pivoting through enterprise perimeters before patch	Perimeter bypass into enterprise networks; agents escalate after perimeter breach	Policy enforcement above network layer — agents blocked from escalating even when perimeter is compromised
Canvas / Instructure	ShinyHunters: 275M students, 3.65TB exfiltrated, $12M ransom, 6-day dwell time. No AI behavioral alerting fired.	No AI agent behavioral monitoring — data movement anomaly undetected for 6 days	Behavioral baselining flags anomalous data movement in minutes; kill switch terminates offending agent instantly
ADT + Comcast	Best-in-class stacks; both breached; $117.5M+ in	No runtime governance on AI workflows —	Fail-closed enforcement: no agent acts outside declared

PLATFORM ARCHITECTURE

One Platform. Five Control Pillars.

01 · DISCOVER	02 · GOVERN	03 · PROTECT	04 · OPERATE	05 · COMPLY
Know	**Control**	**Shield**	**Automate**	**Certify**
Know Your Agent (KYA)	AI Integration Fabric	PII Shield · Fraud Shield · Audit Black Box	Autonomous Operations	PQ Data Platform
Finds every AI agent, shadow or declared, scores identity and risk, establishes behavioral baseline	Secure connectivity between AI agents and enterprise tools — behavioral analysis, agent identity verification, fail-closed policy enforcement	Inline data protection, fraud detection, quantum-safe tamper-proof audit trail for every agent action	Agent lifecycle, FinOps, health monitoring, automated alerting, SIEM export, incident response	NIST-certified quantum-safe encryption, crypto asset inventory, automated SOC 2 / FedRAMP / HIPAA evidence

PRODUCT CATALOG

9 Categories · 50+ Capabilities Across the Full AI Governance Stack

⚡ AI SECURITY

- Know Your Agent (KYA)
- Shadow AI Discovery
- AI Integration Fabric
- <50ms Kill Switch
- Real-Time Policy Engine (<5ms)
- Agent Observability
- Agent Audit Black Box
- Agent Behavioral Intel

⚙️ AUTONOMOUS OPERATIONS

- Agent Lifecycle Management
- Agent Cost Control & FinOps
- Security Alerting & SIEM Export
- Agent Memory Vault
- Sovereign LLM Platform

🔑 NHI SECURITY

- Non-Human Identity Governance
- Agent Identity Verification
- Service Account Lifecycle
- API Key Inventory & Rotation
- Credential Posture Management
- Machine-to-Machine Auth

☁️ CLOUD SECURITY

- Cloud AI Workload Protection
- Multi-Cloud Agent Governance
- Cloud-Native Policy Enforcement
- K8s AI Security
- Cloud Compliance Posture
- Blast Radius Analysis

🛡️ DATA SECURITY

- Agent PII Shield
- PQ TokenVault
- PQ Secure DataShare
- PQ Transit Shield
- QuantumVault

🔐 POST-QUANTUM SECURITY

- PQ CryptoGuard
- PQ Sign
- PQ Migrate
- PQ Comply
- PQ Policy Engine
- PQ Developer Platform

🏢 AGENTIC ENTERPRISE

- QuantoSign
- Agent Commerce Rails
- Agentic Contract Manager
- Agentic Procurement Hub
- Agent Marketplace
- Agent Developer Hub

🚨 FRAUD DETECTION

- Agent Fraud Shield
- Behavioral Anomaly Detection
- Identity Abuse Detection
- Real-Time Risk Scoring
- Identity Drift Detection

🌐 KINETIC AI

- Edge Agent Governance
- Network Function AI Governance
- AI Hub Infrastructure Governance
- Sovereign Deployment Controls

COMPLIANCE & CERTIFICATIONS

Compliance by architecture — not bolt-on. Automated evidence generation, continuous monitoring, and audit-ready reporting from day one.

✓ SOC 2 · ✓ FedRAMP · ✓ ISO 27001 · ✓ HIPAA · ✓ FINRA · ✓ GDPR · ✓ CCPA · ✓ PCI-DSS · ✓ NIST PQC · ✓ NIST AI RMF · ✓ EU AI Act · ✓ ISO 42001 · ✓ MITRE ATLAS · ✓ OWASP LLM Top 10 · ✓ NIST CSF 2.0 · ✓ EO 14110 · ✓ DORA

DEPLOYMENT & INTEGRATION

- Fully SaaS
- On-Premises
- Air-Gapped
- BYOC — AWS · Azure · GCP · Oracle · Equinix

AI / LLM

OpenAI · Anthropic · Gemini · Llama · Mistral · Cohere · DeepSeek · Grok · Sovereign LLM

SECURITY STACK

Splunk · Sentinel · CrowdStrike · Palo Alto · Cisco · XSOAR · ServiceNow · IBM QRadar · Wiz · Qualys · Tenable

IDENTITY & GRC

Okta · Azure Entra · CyberArk · Vanta · OneTrust

COLLABORATION & PRODUCTIVITY

Slack · Microsoft Teams · Jira · Confluence · PagerDuty

VERTICAL INDUSTRIES

Financial Services
Trading, underwriting, fraud — runtime governance for AI in regulated finance

Healthcare & Life Sciences
Clinical AI, claims processing — HIPAA-grade audit and identity for every agent

Government & Defense
FedRAMP & ITAR-compliant air-gapped deployments for national security workloads

Telecom & Networks
AI Hub, edge agents, and network function governance — Equinix's core use case

Technology & SaaS
Multi-tenant AI governance, shadow AI in dev pipelines, agent-to-agent controls

E-Commerce & Retail
Autonomous purchasing agents and supply chain AI with real-time risk scoring

Energy & Utilities
Fail-closed policy prevents catastrophic AI actions on critical infrastructure

Manufacturing
IIoT agents and supply chain automation with behavioral monitoring and audit

Legal & Compliance
AI-generated contracts and eSign with quantum-safe signing and attestation

Education & Research
AI tutoring and research agents with student data protection under FERPA/COPPA

WHY CHOOSE RUNTIMEAI?

Vendor-Neutral Control Plane

Not locked to any cloud, security platform, or AI vendor. Works across your existing stack — Okta, CrowdStrike, ServiceNow, Azure, AWS — without rip-and-replace. Every AI agent, every environment, one control plane.

<50ms Kill Switch · <5ms Policy

Every agent action evaluated in under 5ms. Rogue agents terminated in under 50ms — mid-execution, before damage is done. No pass-through mode. Fail-closed by default.

Quantum-Safe & Air-Gap Ready

NIST-certified post-quantum cryptography across all products. Full platform deployable on-premises or completely air-gapped — no data egress, no cloud dependency, no compromise on sovereignty.

Ready to see it live?

30-minute demo tailored to your AI agent environment and security requirements

Roshan Shaik
Founder & CEO · RuntimeAI
roshan@runtimeai.io