RuntimeAI

RuntimeAI — Budget Displacement vs. New Spend

CONFIDENTIAL

A CISO / CIO Reference Guide

Prepared by: RuntimeAI

Date: June 2026

Audience: Channel Partners, SIs, MSSPs — for use with enterprise clients

The Question

Every CISO and CIO will ask: **"Is this replacing something I already pay for, or is this new budget?"**

The honest answer is **both — and that's the pitch.** RuntimeAI consolidates 4–6 point tools enterprises currently buy separately for AI-era security and governance, while adding a new enforcement capability that no existing tool provides. The net result is typically budget-neutral to budget-positive for the enterprise, with significantly higher coverage and a cleaner stack.

Part 1 — What RuntimeAI Displaces (Budget Recovery)

For the CISO's Budget

Current Tool Category	Representative Vendors	What RuntimeAI Replaces	Displacement %
Shadow IT / CASB	Netskope, Lookout, Microsoft Defender for Cloud Apps	AI Discovery — finds all AI agents, LLM connections, and tool servers across the enterprise. CASB tools were not built for agents; they miss non-browser AI traffic entirely.	50–100% for AI agent scope
API Security Gateway	Salt Security, Noname, Traceable, Imperva	AI Integration Fabric — governs all AI agent API calls, tool server calls, and LLM API access. Traditional API security gateways do not understand agent context, tool invocation chains, or LLM-specific attack patterns.	60–80% for AI API scope
DLP (Outbound)	Symantec/Broadcom, Forcepoint, Digital Guardian, Microsoft Purview	AI Firewall — bidirectional DLP enforcement on all AI agent data flows. Existing DLP tools inspect email and endpoint; they miss LLM prompt and response channels entirely.	40–60% for AI traffic scope
Compliance Automation	Vanta, Drata, Secureframe, AuditBoard	AI Compliance Hub — automated compliance evidence for AI deployments across SOC 2, ISO 27001, HIPAA, EU AI Act, and NIST AI RMF. Existing compliance tools have no AI-specific controls catalog.	30–50% for AI compliance scope

Current Tool Category	Representative Vendors	What RuntimeAI Replaces	Displacement %
WAF / Bot Protection	Cloudflare WAF, Imperva, F5, Akamai	AI Firewall catches prompt injection, model manipulation, and agent impersonation attacks that traditional WAFs have no signatures for.	20–40% for AI traffic scope

For the CIO's Budget

Current Tool Category	Representative Vendors	What RuntimeAI Replaces	Displacement %
LLM Observability / Cost Tracking	LangSmith, Helicone, Langfuse, Arize AI	AI FinOps — per-agent cost attribution, budget enforcement, and LLM spend forecasting. Most enterprises paying for standalone LLM observability tools can consolidate into RuntimeAI.	80–100% for LLM cost tracking
Model Registry / ML Ops	MLflow, Weights & Biases, Neptune.ai	ML Intelligence Hub — model registry, feature store, drift detection, and hybrid scoring engine. RuntimeAI replaces point ML infrastructure tools for teams running models in production.	50–70% for enterprise model governance scope
Shadow IT / AI Inventory	ServiceNow CMDB (AI assets), Axonius, CyCognito	AI Discovery — automated agent and AI asset inventory fed directly into CMDB or accessible via API.	60–80% for AI asset scope

---

Part 2 — What Is Net-New (New Budget Category)

The following RuntimeAI capabilities have **no existing incumbent** — enterprises currently have no budget line for them because the products did not exist.

Capability	Why It Is New	Budget Framing
Agent Identity	Non-human identity management for AI agents is an entirely new problem. Existing IAM tools (Okta, SailPoint, CyberArk) were designed for human login sessions. Agents need sub-second identity verification across distributed tool calls — no IAM vendor has built this.	New line: AI NHI (Non-Human Identity) governance
Agent Kill Switch	The ability to instantly halt a misbehaving AI agent in production is a new control plane concept. No existing EDR, SIEM, or SOAR tool can suspend an autonomous AI agent in real time.	New line: Agentic incident response / containment
Post-Quantum Data Protection	NIST finalized post-quantum cryptography standards in 2024. The migration window is 2025–2028. Most enterprises have not budgeted for PQC yet — this is an emerging mandated spend for regulated industries and federal customers.	New line: PQC migration / quantum-safe data vault
Tool Server Governance	Model Context Protocol (MCP) is the emerging standard for AI agent tool calls. No security vendor currently enforces policy on MCP traffic. RuntimeAI's AI Integration Fabric is the only product that governs tool calls at this layer.	New line: AI Integration security

Capability	Why It Is New	Budget Framing
Autonomous Agent Behavioral Baseline	Per-agent behavioral profiling — what does this agent normally do, and when does it deviate — is a new detection category. Existing UEBA tools model human behavior, not agent behavior.	Extends existing UEBA budget or new AI UEBA line

---

Part 3 — Typical Budget Impact

Scenario A: Enterprise Already Spending on AI Security Point Tools

An enterprise paying for Netskope ($80K/yr) + Salt Security ($60K/yr) + Vanta ($40K/yr) + LangSmith ($24K/yr) = **$204K/yr** across four separate tools — each covering only a slice of the AI security surface.

RuntimeAI Guardian tier covers all four use cases plus adds agent identity, kill switch, behavioral monitoring, and post-quantum protection: **$60K–$120K/yr** depending on agent volume.

Result: Net budget reduction of $84K–$144K/yr with significantly broader coverage. CISO wins the consolidation argument with the CFO.

Scenario B: Enterprise Deploying AI Agents for the First Time

No existing AI security tooling. CISO is being asked to approve AI agent deployments across the organization but has no governance framework.

RuntimeAI provides the entire governance stack — discovery, identity, policy enforcement, audit trail, compliance — as a single platform: **$36K–$120K/yr** (Foundation to Guardian).

Result: New budget, but enables the AI deployment program to proceed. CIO gets the speed; CISO gets the governance. The alternative is a hand-rolled governance framework costing 10–20× in engineering time.

Scenario C: Managed Service / Channel Partner Deployment

A channel partner (SI or MSSP) delivers AI agent implementations for enterprise clients. Adding RuntimeAI as the governance layer inside the managed service: **$36K–$60K/yr** incremental per client, fully managed by the partner.

Result: Partner increases account ARR, client gets AI governance as a managed service, partner earns margin on the RuntimeAI license plus services revenue on implementation and ongoing management.

Part 4 — The CISO Conversation Script

> *"RuntimeAI is primarily a consolidation play for enterprises that have started buying AI-era security point tools — CASB, API security, LLM observability — and are paying for three or four products that still don't talk to each other. RuntimeAI replaces all of those for AI-agent traffic specifically, which is where the risk and the cost is concentrating.*
>
> *The net-new component is agent identity and the kill switch — things that literally don't exist anywhere else. But in most enterprise budgets, that new spend is smaller than what gets consolidated out.*
>
> *For enterprises that haven't started buying AI security tooling yet, RuntimeAI is the single purchase that covers the entire surface area — instead of assembling a stack piece by piece over the next 18 months and still having gaps."*

Part 5 — The CIO Conversation Script

> *"Every AI deployment your teams are running is generating cost that isn't attributed to anyone. Every agent call to GPT-4o or Claude is going on a shared API key, and nobody knows which team or project is responsible for the bill.*
>

> *RuntimeAI's AI FinOps layer attributes every cost to the agent, the project, and the business unit — automatically. CIOs who have deployed this tell me it pays for itself in the first quarter from cost waste they didn't know was happening.*

>

> *The governance piece — knowing which agents are running, what they can access, what they've done — is what allows the CIO to say yes to more AI deployments faster. Without it, every new AI agent deployment is a security review that takes 6–8 weeks. With RuntimeAI, the governance is built in from day one."*

Part 6 — The CFO Conversation Script

> *"The first wave of enterprises that deployed AI agents without governance are now cleaning up data exposure incidents, shadow AI cost overruns, and compliance gaps — each of which costs 10–50× the price of a governance platform deployed upfront.*

>

> *RuntimeAI is the governance platform. The math is straightforward: one incident — a data exposure, a compliance finding, or a runaway AI cost spike — exceeds the annual platform cost by an order of magnitude. This is infrastructure spending, not discretionary."*

Summary — One Sentence Per Executive

- **CISO:** RuntimeAI consolidates your AI security point tools into a single platform — typically reducing stack cost while adding agent identity and kill-switch controls that no point tool provides.
- **CIO:** RuntimeAI attributes every AI agent cost to the right team, enables faster deployment approvals, and ensures every autonomous agent is auditable and compliant from day one.
- **CFO:** One AI governance incident costs 10–50× the annual RuntimeAI platform fee — this is infrastructure spending, not discretionary security.

Platform Tiers

Tier	What's Included	List Price / Year
Foundation	AI Discovery + Agent Identity + Control + Compliance Hub	From $36,000
Guardian	Foundation + AI Firewall + Behavioral Intel + Kill Switch + Ops Center	From $60,000
Fortress	Full platform — all products	From $120,000
Sovereign	Air-gapped / on-premises deployment	Custom

Pricing scales with agent volume. Partner margin: 25%. Volume pricing available.

RuntimeAI | www.runtimeai.io | trial.runtimeai.io

Partner inquiries: roshan@runtimeai.io | https://calendly.com/roshan-runtimeai/30min