



RuntimeAI

KILL SWITCH IN THE LOOP.

RUNTIMEAI · LIVE DEMO WEBINAR

Ten Attack Surfaces. One Kill Switch.

RuntimeAI Kill Switch — sub-50ms containment, live

A live demo with Roshan Shaik, Founder & CEO, RuntimeAI

Wednesday, October 14 · 3:00 PM ET / 12:00 PM PT

30-minute live demo · 10 minutes of Q&A

**CONTAINMENT
AT MACHINE SPEED**

An autonomous agent decides and acts in milliseconds. By the time a human notices, the action is already taken. This session is about the control that operates on the same clock as the agent: RuntimeAI's Kill Switch. Roshan Shaik walks through the ten attack surfaces an agent actually touches, then demonstrates live how one kill switch governs all ten — and what evidence remains after containment.

10

Attack surfaces covered by one kill switch

<50ms

Containment, out-of-scope sessions terminated

1

Platform, instead of 28+ point solutions

The Ten Attack Surfaces

- | | |
|------------------------------|-------------------------------|
| 01 Infrastructure & Egress | 06 Memory & Context |
| 02 Identity & Authentication | 07 Agent-to-Agent (A2A) |
| 03 Authorization & Policy | 08 Behavioral Intelligence |
| 04 Prompt & Injection | 09 Supply Chain & Ecosystem |
| 05 Tool Call & Action | 10 Output, Audit & Compliance |

One kill switch governs all ten. We cover them in the first ten minutes, then run the switch live.

The Kill Switch — The Main Event

Contain on the agent's clock

Out-of-scope sessions terminated in under 50ms; no human in the loop, no waiting on an approval queue.

One switch, ten surfaces

The same control applies across infrastructure, identity, tool calls, memory and A2A; no per-vendor patchwork.

Regulation-ready by design

Containment is moving from best practice to regulatory expectation, and every stop is logged as evidence.

Featured Speaker



Roshan Shaik

Founder & CEO, RuntimeAI

Hosted by Nancy Faith Anello, RuntimeAI

Roshan Shaik is a technology executive and founder dedicated to building security, control, and governance infrastructure for autonomous AI. With deep expertise in AI middleware, cryptographic identity systems, and enterprise compliance, Roshan leads RuntimeAI in enabling organizations to deploy agentic workflows safely at scale.

www.runtimeai.io



REGISTER

Scan to register for the live demo

Wednesday, October 14 · 3:00 PM ET / 12:00 PM PT

[LinkedIn event](#)



START YOUR FREE 30-DAY TRIAL

No credit card required.

Scan to start your free 30-day trial at trial.runtimeai.io

What It Takes To Cover The Same Ten Surfaces

All ten attack surfaces — RuntimeAI versus the point-solution stack.

WITH RUNTIMEAI

= live capability today

WITHOUT RUNTIMEAI

= vendor/tool needed instead, plus the capability gap

ATTACK SURFACE	WITH RUNTIMEAI	WITHOUT RUNTIMEAI
01 Infrastructure & Egress	eBPF kernel-layer inline dropBlocked-destination BFP mapAudit event on drop	Vendors: Palo Alto NGFW · Cisco Secure Firewall / Umbrella · Zscaler. Gap: no agent-aware egress enforcement across any of them.
02 Identity & Authentication	Vendor attestation APIIdentity/DNS DoL resolverSPIRE SVID rotation	Vendors: HashiCorp Vault · SPIRE (self-managed) · Okta AI Agents · Microsoft Entra Agent ID · SGNL + Pangea. Gap: no per-agent identity at this granularity — most are service-account level or single-vendor-locked; nothing for Entra, no broader coverage for SGNL.
03 Authorization & Policy	OPA inline multi-layer chainKill switch L1/L2/L3 + breakerEntitlement + quota enforcementHuman approval gates	Vendors: OPA (self-managed) · AWS IAM · Azure RBAC · Cedar. Gap: no behavioral scope enforcement; JetStream Security and IBM watsonx.governance have no published performance data, and wait-time governance is a documented weakness rather than real-time runtime enforcement.
04 Prompt & Injection	Direct injection defenseIndirect injection defense (MCP Firewall)A2A message injectionJailbreak detection (rules-based)	Vendors: Lakera Guard · Prompt Security · Rebuff (open source). Gap: no coverage for indirect or multi-modal injection across any of them.
05 Tool Call & Action	Multi-layer enforcement chainUnauthorized tool blockKill switch L1/L2/L3Kills shadowing at execution	Vendors: Zenity (posture management). Gap: no live enforcement; the closest analog but policy/posture-focused, not live enforcement. No commercial solution matches RuntimeAI's multi-layer chain — alternative is custom middleware or static OWASP rules.
06 Memory & Context	Memory enc/decr (ML-KEM-1024)Write-time injection classificationRAG adversarial chunk filterMemory provenance chain	Vendors: Pinecone / Weaviate · Custom encryption layer · IBM Guardium. Gap: no injection detection on writes, no RAG poisoning defense, no provenance chain in any of them.
07 Agent-to-Agent (A2A)	A2A detection (SVID header)Payload injection inspectionCross-tenant mTLS (Bot-CA)Identity/DNS revocation	Vendors: Istio / Envoy (mTLS transport only). Gap: no A2A payload inspection, no cross-tenant agent identity standard, no authorization decisions.
08 Behavioral Intelligence	Risk score 0–100 (pattern-based)MCP anomaly detectionFraud shield velocity rulesIdentity fraud detection	Vendors: Datadog LLM · Splunk UBA · Exabeam · IBM QRadar. Gap: all async, none live; none are agent-native at tool-call granularity — Noma Security stops at detection, no enforcement or kill switch.
09 Supply Chain & Ecosystem	Vendor attestation APICommerce Rails (Stripe cards)MCP marketplace vetting	Vendors: Snyk / Gripe (SBOM) · Socket.dev (supply chain). Gap: no agent-specific framework vetting, no virtual card per agent, no MCP server regress model in any of them.
10 Output, Audit & Compliance	Audit Black Box (SHA-256 hash chain)Legal holds + exportCost hard-stop enforcementSOC 2 / ISO 27001 / EU AI Act / NIST evidence	Vendors: Microsoft Purview · AWS Config · Splunk / Elastic (audit) · Vanta / Drata. Gap: compliance workflow, not real-time blocking; none offer an agent-native audit chain or per-agent cost enforcement.

28+ vendors and platform competitors required to approximate RuntimeAI's coverage — RuntimeAI = 1 platform.

One control plane.
Ten surfaces.
Live today.

...and stitching 28+ point solutions together still leaves you without agent-native tool-call governance, Memory Vault-grade context protection, or a published, benchmarked kill switch.