FOR SOC ANALYSTS · SECOPS · AI SECURITY ENGINEERS

A Red Team That Updates Itself From This Week's Real AI Attacks

Static playbooks go stale the week you write them. RuntimeAI's Red Team engine (ATSA) ingests real-world AI-security incidents and drafts new attack scenarios on its own — admin-reviewed before anything ships — while a Blue Team of 19 real enforcement services stops what it finds, and a stop-switch you can trust actually works when it matters.

Red **Autonomous Threat Simulation**

100+ scenarios across 9 attack-surface domains, self-updating from real incident feeds — not a static checklist reviewed once a year.

Blue **19 Real Enforcement Services**

Flow Enforcer, WAF, OPA, PII Shield, Fraud Shield, and more — each enforcing autonomously at request time, every time.

THE PROBLEM

Most AI Security Programs Run on Faith, Not Evidence

Playbooks Go Stale

A red-team catalog hand-authored once a quarter is already behind the attacks landing on real AI systems this week.

Containment Takes Hours

Traditional SOAR containment measures in minutes to hours. An AI agent exfiltrating data doesn't wait for an analyst to start a playbook.

Point Tools, No Single View

Red-team findings, blue-team enforcement status, and governance gap data usually live in three different consoles — nobody sees the whole loop.

WHAT RUNTIMEAI DOES

Kill Switch That Actually Stops an Agent

CAPABILITY	WHAT IT DOES
Agent termination	<50ms, globally across all deployment metros — credentials revoked simultaneously so re-authentication can't happen mid-stop
Scoped containment	Kill by agent ID, integration, tag, business unit, risk tier, or platform-wide — precise, not all-or-nothing
Forensic capture	Last 100 agent actions preserved in an encrypted package the instant a kill switch fires — stopping the agent doesn't destroy the evidence
Reversal	One API call or console click, <50ms propagation, every activation and deactivation permanently recorded

WHERE WE'RE HONEST ABOUT WHAT'S NOT DONE YET

Purple Team validation — proving red findings actually trigger blue detection — is a real, named, still-open gap (NS-16). We're closing it, we don't claim it's built.

The **unified Red/Blue/Purple console** is implementation-ready but not yet shipped. Today, ATSA itself, the autonomous scenario pipeline, and all 19 Blue Team services are real and live — the single-pane console is what's coming next.

Agent-to-agent trust verification (ML-DSA-87 signature + session + confused-deputy protection) is built and live-verified against real attack scenarios, but enforcement is still off in production while we complete a shadow-mode observation window.

WHO THIS IS FOR

SOC Analysts
A single triage number instead of correlating six separate alert streams by hand

SecOps
19 enforcement services with real-time status, not a quarterly compliance snapshot

AI Security Engineers
A red-team catalog that stays current without hand-authoring YAML every week

Incident Response
Automated 5-phase response — detect, contain, investigate, remediate, learn — in seconds, not hours

WHY THIS MATTERS TO A SECURITY LEADER

A Stop-Switch That Works

No other commercial vendor has shipped an agent kill switch with this latency, scope precision, and forensic capture combined. Not a claim you can find matched.

Self-Updating, Not Hand-Authored

New AI-security incidents become new test scenarios within days, admin-reviewed before anything ships — no auto-promotion, ever.

One Loop, Not Three Consoles

Red findings, Blue enforcement status, and Purple validation are designed to live in one place — the console is coming, the underlying data already does.

See your real risk score in under an hour

Request a gap assessment against the current attack scenario catalog.

Roshan Shaik
Founder & CEO · RuntimeAI
roshan@runtimeai.io
www.runtimeai.io
calendly.com/roshan-runtimeai/30min