RoonCyber.AI

CUSTOMER CASE STUDY ALLURE SECURITY × ROONCYBER.AI

ALLURE SECURITY CLEARS FINANCIAL-SERVICES SECURITY REVIEWS 40% FASTER WITH COMPLETE AI RUNTIME SECURITY

Securing the models, agents, and data pipelines that are an AI-native defense platform's core product.

~40% faster
financial-services vendor security reviews cleared

Minutes
to detect and contain AI-specific threats, not days

100%
of AI assets (models, agents, tools) inventoried and monitored

1.5B+/day
URLs of AI pipeline activity under continuous visibility

Allure Security is an AI-native brand-protection company whose custom vision and NLP models and autonomous agents scan more than 1.5 billion URLs a day to detect brand impersonation, phishing, and deepfake fraud for over 300 enterprise customers, most of them banks and credit unions. When your product is AI, the AI itself has to be beyond reproach, and traditional security tooling was never built to see or defend AI workloads.

THE CHALLENGE

Allure's entire value proposition rests on AI. As the platform scaled, its AI footprint of models, agents, tools, and the sensitive data flowing through them grew faster than it could be tracked or secured by hand. Three pressures converged:

01 **Adversaries target the AI directly.** Data poisoning, model evasion, prompt injection, and agent hijacking are aimed at the very AI that protects customers.

02 **Customers demand demonstrable governance.** Financial-services buyers run deep vendor due diligence and expect provable AI governance before signing.

03 **A security vendor's own AI must be exemplary.** For an AI-native security company, credibility depends on securing its own AI to the same standard it promises customers.

Allure needed security purpose-built for AI, not retrofitted from endpoint or network tools, and it needed more than another stream of alerts. Because RoonCyber.AI operates at runtime, the same execution visibility that surfaces a threat also gives the team the power to act on it.

RoonCyber.AI discovers every AI component, maps how data and attacks move through the pipeline, and detects AI-specific threats in real time. When a workload crosses the line, it does more than raise an alert: it **blocks** the offending process, **contains** the workload to limit the blast radius, and **remediates** the access that enabled it — all from one platform, without adding operational drag to a focused team.

THE ENGAGEMENT

AI DISCOVERY AND INVENTORY

RoonCyber.AI automatically inventoried every model, agent, tool, and LLM across the platform, replacing fragmented, tribal knowledge with a single live map of the AI estate.

AI WORKLOAD ATTACK-PATH VISIBILITY

It visualized how customer brand data, decoy credentials, and threat intelligence move through the model pipeline, surfacing exactly where an adversary could poison data or manipulate an agent.

AI DETECTION AND RESPONSE

It monitored models and agents for poisoning, evasion, prompt injection, and anomalous behavior at runtime, and automated remediation so threats are contained in minutes.

RUNTIME-BACKED COMPLIANCE EVIDENCE

As a complement, the same runtime record produced audit-ready evidence for SOC 2 and customer security reviews, with no separate documentation effort.

- **Complete, continuous visibility** into an AI pipeline processing 1.5B+ URLs a day across 300+ customer brands.
- **Faster financial-services deals,** with vendor security assessments completed roughly 40% faster on the strength of provable runtime security.
- **AI-specific threats contained in minutes, not days,** with automated remediation closing the loop from detection to defense.
- **Audit-ready evidence on demand,** cutting manual preparation for SOC 2 and customer audits.
- **A sharper competitive story:** an AI-native security vendor that secures its own AI with RoonCyber.AI.

> "Runtime is where the truth about AI lives. Seeing exactly what our models and agents do the moment they execute, instead of inferring it from logs after the fact, is what lets us tell more than 300 financial-services customers, with evidence, that the AI defending them is itself secure. RoonCyber.AI gave us that ground truth without slowing us down."
>
> **JOSH SHAUL** — CHIEF EXECUTIVE OFFICER, ALLURE SECURITY

RoonCyber.AI

ABOUT

RoonCyber.AI is a Runtime AI Workload Security platform that reveals what AI is actually doing in production. As AI agents interact with APIs, applications, data, and cloud services, traditional security tools often lose visibility beyond the prompt. RoonCyber.AI provides real-time visibility into AI workload execution, enabling security teams to identify shadow AI, uncover AI-driven attack paths, validate threats, and assess business impact. Built on advanced runtime technology, RoonCyber.AI helps organizations reduce risk, improve security operations, streamline compliance, and confidently adopt AI by exposing what AI is accessing, executing, and impacting across the enterprise.