RoonCyber.AI

AI BENEFITS BUILT FOR EVERY SEAT AT THE TABLE

VALUE, BENEFITS, AND DIFFERENTIATORS BY PERSONA: CIO · CISO · SOC · AI SECURITY ENGINEERING

RoonCyber.AI turns runtime ground truth into outcomes each stakeholder can act on. Three core capabilities work together so you can adopt AI fast while keeping every workload **discovered, mapped, and protected** end to end, framed below for what each role on your team is accountable for.

THREE CORE CAPABILITIES

AI RUNTIME DISCOVERY

Runtime visibility surfaces every AI workload running in production, including shadow AI, and traces end-to-end downstream vulnerabilities across the systems, APIs, and data each agent actually touches.

Business outcome

Adopt AI with accountability, nothing runs unseen or unowned.

AI VISUALIZATION & REPORTING

RoonCyber.AI creates Security Graphs that turn everything into clear visual dashboards and reports: executive dashboards and summaries, line-item insights, and instant remediation plans. Context and criticality are easy to grasp end-to-end, keeping everyone from practitioners to the board aligned.

Business outcome

Less time building and deciphering reports, with AI insights at your fingertips.

AI PROTECTION

RoonCyber.AI moves from insight to action, extending protection to the kernel level with automated remediation and blocking. Malicious or out-of-policy AI activity is contained the moment it executes, not after a review cycle, closing the loop from detection to defense.

Business outcome

Stop threats at runtime, automatically. Exposure is contained before it becomes an incident.

CAPABILITY × PERSONA AT A GLANCE

The same runtime ground truth, framed for what each role is accountable for.

BENEFIT	CIO	CISO	SOC	AI SEC ENG
Runtime AI Discovery	Live AI inventory	Shadow AI surfaced	Enriched runtime signals	Process-level detail
AI Visualization	One executive view	Defensible evidence	Attack-path maps	Full execution traces
AI Protection	Automated containment	Kernel-level blocking	Pipeline-native response	Runtime enforcement

RUNTIME AI WORKLOAD SECURITY

RoonCyber.AI

HOW IT MAPS TO YOUR TEAM

CIO

SCALE SAFELY, MOVE FAST

"I want to scale AI quickly, but I need to know what's actually running, who owns it, and that it's all under control."

Discover
A live, owned inventory of every AI workload in production. Nothing scales without accountability.

Visualize
A single executive view of the AI estate, so you move fast and stay informed without raw telemetry.

Protect
Automated containment stops risky AI activity before it spreads, so you scale aggressively without becoming the bottleneck.

CISO

DEFENSIBLE CONTROL, NO DRAG

"Agents are doing things I can't see. Shadow AI is spreading. I need to know who's using AI and own all the risks, without slowing the business."

Discover
Surfaces shadow AI and unsanctioned agents the moment they execute: who's using AI, and what it touches.

Visualize
Turns invisible agent activity into evidence you can defend to auditors and the board.

Protect
Kernel-level blocking and automated remediation stop malicious agent behavior at runtime: defensible control without slowing the business.

SOC

AGENTIC-SOC FRIENDLY

"How does this fit the SOC I already run, and the agentic workflows I'm building? I can't afford another silo."

Discover
Runtime detections stream as enriched, agent-readable signals into your existing pipeline, no rip-and-replace.

Visualize
Execution and attack-path maps give analysts (human or agent) instant context to validate and scope incidents.

Protect
Automated response and blocking plug into your existing pipeline, so containment happens inline, not in a separate console.

AI SECURITY ENGINEERING

GROUND TRUTH ON AGENTS

"I need to know exactly what our agents do at runtime: validate behavior, prove guardrails hold, and find what's actually reachable."

Discover
Process-by-process detail. Every tool call, API, and downstream dependency an agent touches, end to end.

Visualize
Trace full execution paths to debug behavior and confirm guardrails hold at runtime, not just in config.

Protect
Kernel-level enforcement blocks agent actions that violate policy at runtime, turning validated guardrails into live controls.

SEE WHAT YOUR AI IS ACTUALLY DOING

Request a technical briefing and live demonstration for your security team. We will show you the complete execution path of your AI workloads, and what your current stack is missing.

contact@RoonCyber.com | RoonCyber.com/demo