RoonCyber.AI

PROTECT YOUR ORGANIZATION'S AI AT RUNTIME

AI RUNTIME PROTECTION, BLOCKING, AND REMEDIATION

Runtime detection without response is just a more expensive flashlight. RoonCyber.AI does not just surface AI threats — it blocks them, contains them, and helps your team close the loop before damage spreads.

THE PROBLEM WITH DETECTION WITHOUT RESPONSE

Most runtime tools surface AI threats in real time but stop there. Your team gets an alert, a severity score, and a dashboard. What they do not get is the ability to stop the threat from that same platform, immediately and with evidence to back the action.

RoonCyber.AI gives security teams the runtime evidence to understand exactly what happened and the controls to act immediately — killing processes, quarantining workloads, removing permissions, and generating compliance evidence — all from one platform.

WHAT ROONCYBER.AI DOES WHEN IT FINDS A THREAT

01

BLOCK

Stop threats the moment they execute.

When RoonCyber.AI detects anomalous AI behavior at runtime, it gives your team the ability to kill the process immediately — before it can access more data, call additional APIs, or spread further across your environment.

- Kill unauthorized agent processes instantly
- Block out-of-scope API calls before they complete
- Stop credential misuse at the point of access

02

CONTAIN

Quarantine workloads to limit the blast radius.

When a threat cannot be immediately killed without disrupting critical operations, RoonCyber.AI allows your team to quarantine the workload — isolating it from the rest of your environment while investigation continues.

- Isolate rogue AI workloads from production systems
- Contain lateral movement before it spreads
- Preserve execution evidence for investigation

03

REMEDIATE

Remove access and close the gap for good.

Stopping a threat is only half the job. RoonCyber.AI helps your team remove the permissions and access that enabled it, generate the compliance evidence that documents it, and update policies so it cannot happen again.

- Remove permissions and revoke agent access
- Generate audit-ready compliance evidence
- Surface policy gaps to prevent recurrence

IME PROTECTION & REMEDIATION

RoonCyber.AI

REMEDIATION AND COMPLIANCE — CLOSE THE LOOP. KEEP IT CLOSED.

EVERY ACTION BACKED BY GROUND TRUTH, NOT GUESSWORK

Because RoonCyber.AI operates at the kernel level, every protection action is backed by a complete runtime execution record — what the workload did, what it accessed, and why the action taken was the right one. That makes remediation defensible to your security team, leadership, auditors, and regulators alike.

OWNERSHIP MAPPING

Every AI workload involved in an incident is traced back to an owner so accountability is clear and remediation has a starting point.

COMPLIANCE EVIDENCE

Runtime records become audit-ready documentation proving what happened, what was done about it, and that your AI workloads behaved within policy.

POLICY GAP ANALYSIS

After an incident, RoonCyber.AI surfaces exactly which policy gaps allowed the threat to execute, so your team can close them before it happens again.

BUSINESS IMPACT REPORTING

Convert runtime telemetry into financial exposure metrics your board and CFO can act on, not just technical severity scores that require translation.

NO NEW TOOLS REQUIRED
NATIVE CLAUDE INTEGRATION

RoonCyber.AI connects natively to Claude. Ask what needs immediate attention, get a runtime-backed answer, and trigger response actions directly without logging into a new platform.

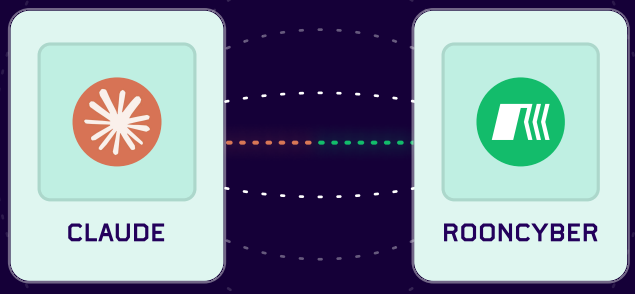

IME PROTECTION & REMEDIATION

RoonCyber.AI

WHAT ROONCYBER.AI BLOCKS AND CONTAINS AT RUNTIME

SHADOW AI WORKLOADS

Unauthorized models agents, and pipelines blocked the moment they start executing.

UNAUTHORIZED AGENT ACCESS

Agents stopped at execution when they reach systems or data outside their declared scope

LATERAL MOVEMENT

Agents traversing systems and containers beyond their boundaries quarantined in real time

PRIVILEGE ESCALATION

Agents attempting to escalate beyond their authorization caught and shut down immediately.

DATA EXFILTRATION

Workloads transmitting sensitive data outside authorized boundaries blocked before it completes.

THIRD-PARTY AI INTEGRATIONS

Vendor AI features embedded beyond what security reviewed.

ROGUE AI PIPELINES

Unauthorized automated AI workflows identified, contained, and documented for remediation.

CROSS-TENANT DATA ACCESS

Agents querying data across tenant boundaries stopped with full runtime evidence.

CREDENTIAL MISUSE

Workloads accessing credentials beyond their scope identified and shut down at the moment of access.

SEE AI PROTECTION IN ACTION

Request a demo and we will show your team how RoonCyber.AI detects, blocks, contains, and remediates AI threats at runtime — from one platform, with evidence behind every action.

contact@RoonCyber.com | RoonCyber.com/demo