

SHADOW AI IS THE BLIND SPOT RUNTIME CAN SEE

DETECTION, VISIBILITY, AND PROTECTION AT RUNTIME

Shadow AI is any model, agent, pipeline, or tool operating in your environment without your security team's knowledge or approval — the AI equivalent of shadow IT, but faster moving and far more dangerous. Guardrails and asset scanners only see what was registered; runtime is the only way to find, inventory, and stop every one of them.

74%

of employees use AI tools
security hasn't approved

3x

faster AI agent adoption
than traditional shadow IT

0%

of prompt-layer tools
catch agents that skip the
gateway

KEY USE CASES



GUARDRAILS AND GATEWAYS

Only monitor traffic that passes through them — an agent that bypasses the gateway is invisible.



ASSET SCANNERS AND CSPM

Only find what's registered. Unauthorized agents inside approved infrastructure look like normal activity.



LOG REVIEW

Finds shadow AI after the fact — often weeks after it started accessing data.

Runtime is the only approach that catches Shadow AI the moment it starts executing — regardless of whether it was ever registered, approved, or passed through a gateway.

HOW ROONCYBER.AI FINDS IT

AI RUNTIME DISCOVERY

Complete inventory, approved or not

Kernel-level monitoring surfaces every AI workload the moment it starts executing.

AI VISUALIZATION

See what it touched, and who owns it

Traces every system, API, and data store a shadow workload has reached.

AI PROTECTION

Act before exposure grows

Kill the process, quarantine the workload, and remove access — immediately.

SHADOW AI IS THE BLIND SPOT RUNTIME CAN SEE

THE UNAUTHORIZED AI RUNNING IN YOUR ENVIRONMENT RIGHT NOW



UNAPPROVED AI MODELS

Personal or third-party AI models running without security review.



ROGUE AI AGENTS

Autonomous agents deployed with no security involvement or oversight.



UNAUTHORIZED MCP SERVERS

Unapproved servers connecting agents to internal tools and data.



SHADOW AI PIPELINES

Automated workflows processing data outside approved channels.



THIRD-PARTY AI INTEGRATIONS

Vendor AI features embedded beyond what security reviewed.



CLOUD-HOSTED AI AGENTS

Workloads never surfaced in your security inventory or asset register.



AI ACCESSING UNSANCTIONED DATA

Approved agents drifting beyond scope into new data stores



AI WITH EXCESSIVE PERMISSIONS

Workloads that accumulated access beyond original authorization.



EXTERNAL AI CALLING INTERNAL APIS

Outside services calling internal APIs with zero visibility.

 NATIVE CLAUDE INTEGRATION

Ask What's running right now.

What unauthorized AI agents are running in our environment right now?



Three unauthorized workloads detected: an unregistered agent in marketing-vpc with no owner, a shadow pipeline querying the customer database after hours, and an unapproved MCP server connecting to internal Slack and Jira. None passed through your approved gateway.

Reply to Claude...



Sonnet 4.5



SEE WHAT YOUR AI IS ACTUALLY DOING

Request a technical briefing and we'll show you every unauthorized AI workload your current stack has never seen