wallarm

DATA SHEET

PRODUCT DATA SHEET · PLATFORM

Wallarm AI Control Platform

Discover what AI is running. Observe what it does. Enforce when it crosses a line. Prove it is governed.

SUMMARY

Your AI agents are deciding, calling, and accessing data right now.

Most security teams cannot tell you which agent made which call, whose request triggered it, or whether anything it touched violated a policy. That is not a tooling gap. It is a governance gap. The tools built to secure API traffic in AWS were not built to govern agentic AI, and the tools built to monitor infrastructure were not built to close the loop between what is deployed and what is actually running.

Wallarm AI Control Platform closes that loop with the Wallarm AI Control Loop. The same platform that has protected your APIs for years now governs the AI those APIs serve. One platform, one control plane, for AI security and API security together.

3	4	1
PRODUCTS IN ONE CLOSED LOOP	JOBS: DISCOVER OBSERVE ENFORCE GOVERN	CONTROL PLANE FOR AI & API SECURITY

THE WALLARM AI CONTROL LOOP

AI Governance is four jobs that must work together.

Discover what is running. Observe what it does. Enforce when it crosses a line. Govern with evidence that holds up in an audit. Close any one of the four and you have left the others open.

Discover 01

Know what AI is running in your AWS estate. Every account, every cluster, every workload, every API the AI talks to — including the AI nobody told security about.

Observe 02

See what AI is doing at runtime. Data accessed, decisions made, what crosses the boundary, who triggered each call. AWS findings land on the same graph with full behavioral context.

Enforce 03

Stop AI behavior that violates policy. Block at the connection level. Revoke compromised sessions by user or trace ID. No pod restart, no deploy cycle. Stop attacks against the APIs your AI exposes.

Govern 04

Produce continuous, audit-ready evidence. Coverage heatmap, AI-SBOM, session audit logs, sensitive data flow records, Security Hub triage logs — built before the audit, not because of it.

An Amazon GuardDuty finding in a random account out of 1,200 means nothing without context for what AI workloads run there, what agents made external calls, and what data flowed through them. AWS surfaces the finding. Wallarm provides the context that makes it actionable — and enforces before it becomes a breach.

THE THREE PRODUCTS

Three products. Three layers of the stack. One closed loop.

From the AWS estate at the base, up through AI workloads on Kubernetes, out to the API edge where external traffic hits. Each product owns its layer. Each carries its share of the loop.

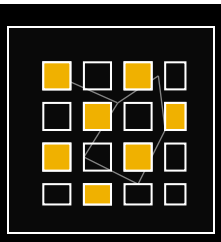

01

Infrastructure Discovery

SAAS AGENT · THE AWS ESTATE

DISCOVER OBSERVE

Reads your AWS estate across every account and region, diffs each scan, and ingests Security Hub findings against the assets they describe. CEL-based finding policies let security teams extend and triage without forking the platform. Attack-path analysis traces the route from any internet-facing entry point to a given asset.

Inventory + topology
EC2, VPC, EKS, Lambda, API GW, ALB/NLB, IAM, Bedrock — typed graph.

Findings on a graph
Amazon GuardDuty, Inspector, Macie, Access Analyzer, Config — in plain language.

Attack-path analysis
Traces exposure path from internet gateway to any asset through security groups and load balancers.

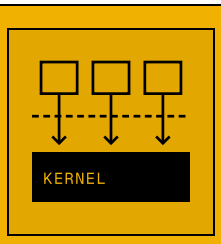

02

AI Hypervisor

KUBERNETES DAEMONSET · THE RUNTIME

DISCOVER OBSERVE ENFORCE GOVERN

One DaemonSet on EKS. Zero application code changes. Within minutes of labeling a deployment, see every AI workload, every outbound connection, every piece of sensitive data — and stop it on policy.

Auto-discovery
11+ model providers, MCP servers, agent frameworks. Shadow AI surfaced.

Per-call attribution
Kernel runtime-task-ID correlation. No trace headers required.

Kernel enforcement
RST injection + conntrack flush on in-flight flows. BPF-optional.

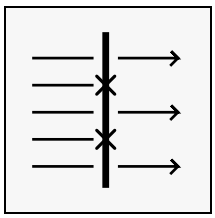

03

API Security

INLINE WAAP · THE API EDGE

DISCOVER ENFORCE

Every AI system is an API consumer, an API producer, or both. The attack surface does not change because the consumer is an agent. Inline blocking across REST, GraphQL, gRPC, WebSocket, and SOAP.

Inline blocking
OWASP Top 10, injection, BFLA/BOLA, volumetric — in real time.

Shadow API discovery
Passive discovery of undocumented endpoints from production traffic.

Business logic
Excessive fetching, parameter manipulation, scraping patterns WAFs miss.

THE ECONOMIC ARGUMENT

Security features do not get bought. Business risk does.

Close any one of the four and you have left the others open. The three products are integrated by design, not by configuration.

Benefits.

FOR EVERY ROLE THE AI CONVERSATION TOUCHES

01 CIO / EXEC

Reportable posture.

- AI governance across every workload, running inside your existing AWS environment.
- One platform for AI security and API security — one procurement decision.
- Walk into an EU AI Act or SOC 2 review with evidence built continuously, not assembled the week before.

02 CISO

Actionable findings.

- Know what AI is running, what it is doing, and who triggered each action.
- Amazon GuardDuty findings land on the same graph as runtime behavioral context — alerts you can act on.
- Compromised sessions revoked on in-flight connections. No pod restarts, no downtime.

03 CTO / PLATFORM

Zero code changes.

- Deploy, label, instrumented in minutes. No application changes, no sidecars.
- Language-agnostic across Python, Go, Node, Java, Ruby, and Rust — no per-language SDKs required.
- Enforcement runs inside your EKS environment. Nothing leaves your boundary.

04 SECURITY ENGINEERING

End-to-end attribution.

- Per-call attribution from user request across every service hop — no trace headers required.
- PII detected inline at the SSL boundary. TLS stays intact end to end.
- CEL-based finding policies with every triage action audit-logged.

05 COMPLIANCE

Audit-ready, always.

- Coverage heatmap, AI-SBOM, session audit logs, and sensitive data flow records generated continuously.
- EU AI Act and SOC 2 audits answered with live evidence. No retroactive assembly.
- Read-only AWS posture with verifiable scope — survives review of the security tool itself.

● NEXT STEP

See the closed loop, end to end.

A 30-minute walkthrough across Discovery, Hypervisor, and API Security on a sample environment — no code changes required to follow along.

WALLARM.COM / REQUEST-DEMO — 30 MIN · LIVE

Schedule a Demo →