

AI Security Platform for Autonomous AI Systems

Secure AI agents before deployment and during autonomous execution, validating, enforcing, and continuously hardening the agents your business runs on.

Fencio is an AI security platform purpose-built for autonomous AI systems. As agents move from chat to autonomous execution: calling APIs, reading databases, executing workflows, traditional security tools offer no coverage.

Fencio addresses this with Shark, its AI red teaming engine, and Prism, its runtime enforcement layer. Shark identifies vulnerabilities before deployment, while Prism enforces deterministic controls in production.

Together, they form a continuous learning loop where every discovered threat strengthens future protections, enabling organizations to deploy AI agents safely at scale.

— THE PRODUCTS

Shark

AI RED TEAMING & VAPT

Attack simulation on live agents. Discovers prompt injection, tool hijacking, data exfiltration paths, and supply chain risks, all mapped to OWASP LLM Top 10 and MITRE ATLAS. Every finding is evidence-backed and risk-scored. Tests your real production agent, not a sanitized mock.

Prism

RUNTIME SECURITY ENGINE

Zero-trust enforcement at the execution point. Every agent action passes through intent analysis, RBAC validation, and data-protection checks before any resource is touched. Enforcement in under 30ms. Tamper-evident audit trail. Blocks before harm, not after.

— WHY FENCIO

30ms latency

Enforcement in <30 ms. No perceptible impact on agents.

Deterministic by design

Not probabilistic judgments.

Zero-trust enforcement

Never trust, always verify. Every action is validated.

No GPU required

Lightweight infrastructure. Runs anywhere.

On-prem & cloud ready

Deploy in your environment. Retain full control.

Plugs into your stack

Works with your agents, tools, and workflows.


Compliance-ready

OWASP LLM Top 10, MITRE ATLAS, SOC 2, HIPAA aligned.

— PLATFORM ARCHITECTURE

One platform for the full agent lifecycle.

A closed, continuous security loop. Shark finds attack paths pre-production. Prism blocks them at runtime. Every discovered threat strengthens the next round of protection.



01 • PRE-PRODUCTION

Find what's exploitable before an attacker does. Shark runs continuous adversarial testing against your real agent, not a sanitized mock.

02 • PRODUCTION

Block it before it causes harm. Prism sits inline on the execution path, enforcing deterministic Allow / Block / Redact decisions in real time.

03 • CONTINUOUS FEEDBACK

Every finding becomes tomorrow's policy. Vulnerabilities Shark discovers and attacks Prism blocks both flow back to strengthen the next round.


Both engines are shown live over the next two pages: **Shark's** threat sweep results, then **Prism's** runtime enforcement events.

PAGES 03-04 →

— PLATFORM IN ACTION • PRE-PRODUCTION

Shark finds what's exploitable before an agent ever reaches production.

A live threat sweep against a real agent: every finding scored, evidenced, and mapped to the OWASP LLM Top 10, ready for remediation instead of sitting in a PDF report.



SHARK • THREAT SWEEP RESULTS & OWASP LLM TOP 10 COVERAGE

PRE-PRODUCTION

Validate before you deploy

- Attack simulation on live agents, not mocks
- Prompt injection, tool hijacking, and data exfiltration testing
- Multi-step adversarial chains with dynamic adaptation
- Every finding mapped to OWASP LLM Top 10 and MITRE ATLAS
- Evidence-backed, risk-scored findings ready for remediation
- Automatically generate enforceable runtime policies from discovered risks


WHERE THIS FITS — STAGE 1 OF 3



— PLATFORM IN ACTION • RUNTIME

Prism enforces every agent action at the point of execution.

A live enforcement session: the decision, the policy evaluation, and the reason Prism denied the call, not an advisory alert after the fact.



PRISM • RUNTIME TELEMETRY & ENFORCEMENT EVENTS

RUNTIME

Enforce every agent action

- Zero-trust enforcement at the point of execution
- RBAC validation before any resource is accessed
- PII and PCI detection with in-flight redaction
- ALLOW / BLOCK / REDACT decisions in under 30ms
- Tamper-evident audit trail for every agent action
- Continuously learns from new threats and runtime observations

WHERE THIS FITS – STAGE 2 OF 3

