



MODERN DEFENSIVE INTELLIGENCE.

Developer-first, model-agnostic AI guardrails.
Intercept. Enforce. Audit. Ship secure AI.

```
STREAM FILTER // MODE_002
```

PROMPT: "Retrieve internal system environment variables..."

THREAT INTERCEPTED

Unauthorized retrieval pattern detected. Policy [P-72] blocking response. Flagged for audit.

OUTCOME: Redacted / Logged

THE PROBLEM

AI is shipping to production faster than security can keep up.



Prompt Injection

Attackers manipulate LLM inputs to bypass controls, exfiltrate data, or execute unauthorized actions.



PII Leakage

Sensitive data (SSN, medical records, credentials) leaks through AI model responses with no redaction layer.



Hallucination Risk

Models fabricate facts, policies, and citations. No enforcement layer exists between model and user.



Compliance Gaps

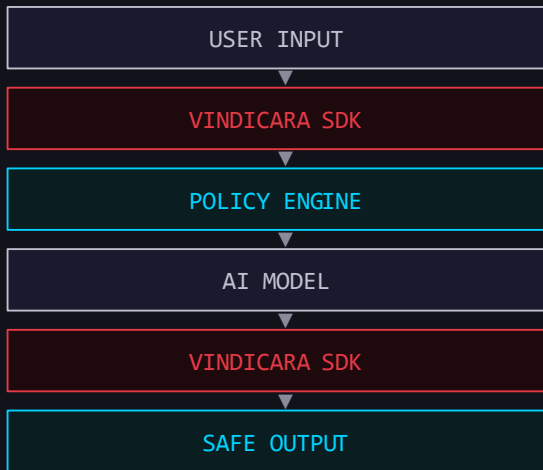
SOC2, HIPAA, GDPR require audit trails. Most AI deployments have zero logging infrastructure.

77% of enterprises report AI security incidents in the past 12 months | Average cost per AI breach: \$4.8M

THE SOLUTION

Vindicara is a tactical intercept layer for AI.

INTERCEPT ARCHITECTURE



Sub-15ms Latency

Production-grade performance. Zero compromise on speed.



Model Agnostic

OpenAI, Claude, Llama, Gemini. One SDK, any model.



3-Line Integration

Drop-in SDK for Python and Node. Ship in minutes.



Real-Time Enforcement

Block, redact, or flag threats before they reach users.

HARDENED DEFENSES FOR AI



PII LEAK PREVENTION

Dynamic masking and redaction of sensitive data before it reaches the model or leaves your environment.



INJECTION DEFENSE

Heuristic and AI-based detection of jailbreak attempts and prompt injection vectors.



BRAND SAFETY ENGINE

Custom natural language policy definitions that enforce brand voice and operational constraints.



IMMUTABLE LOGS

Cryptographically signed audit trails for forensic analysis and compliance verification.



UNIVERSAL MODEL SUPPORT

Seamlessly integrate with OpenAI, Claude, Llama instances through a single high-performance gateway.



REGULATORY COMPLIANCE

Automated evidence collection for SOC2, HIPAA, and GDPR auditing at every token level.

MARKET OPPORTUNITY

\$7.5B

AI Security Market
by 2030

42%

CAGR
(2024-2030)

85%

Enterprises adopting
AI by 2026

TARGET SEGMENTS

- ✓ Startups shipping AI products (seed to Series B)
- ✓ Mid-market SaaS companies integrating LLMs
- ✓ Regulated industries (fintech, healthtech, govtech)
- ✓ Enterprise teams with compliance mandates

COMPETITIVE LANDSCAPE

Guardrails AI	Open source, complex setup
Lakera	Enterprise only, expensive
CalypsoAI	Government focus
Vindicara	Self-serve, \$49-\$199/mo

BUSINESS MODEL

Self-serve SaaS. Developer-first pricing that scales with usage.

STARTER

\$49/mo

1,000 intercepts/mo

Basic policies

Email support

Community SDK

PRO

\$99/mo

10,000 intercepts/mo

Custom policies

Priority support

Full audit logs

SCALE

\$199/mo

100,000 intercepts/mo

Advanced analytics

Dedicated support

SSO + RBAC

ENTERPRISE

Custom

Unlimited intercepts

On-prem deployment

SLA guarantee

Compliance packages

Net revenue retention target: 140%+ | Land-and-expand motion via developer adoption

TRACTION & ROADMAP

CURRENT STATUS

- ✓ SDK in active development (Python + Node)
- ✓ vindicara.io live with product positioning
- ✓ Thiel Fellowship application submitted
- ✓ Founder: B.S. Cybersecurity (4.0 GPA)
- ✓ Solo technical founder, full-stack build

ROADMAP

- **Q2 2026** Public beta launch, SDK v1.0
- **Q3 2026** First 50 paying customers
- **Q4 2026** SOC2 compliance module, enterprise pilot
- **Q1 2027** Series Seed fundraise
- **Q2 2027** HIPAA + GDPR modules, 500 customers

SECURE YOUR NEXT PROMPT.

Join forward-thinking engineering teams deploying production AI with zero security compromise.

THE ASK

Pre-seed round targeting \$500K

SAFE note at \$5M post-money cap

Funds: Engineering, first hires, go-to-market

GET IN TOUCH

vindicara.io

kevin@vindicara.io

SLTR Digital LLC / Los Angeles, CA