

Sui Sentinel

The Crowdsourced AI Red Teaming platform on
Sui

suisentinel.xyz

The Problem

AI Systems Are Powerful. And Dangerously Unprotected.

AI Systems derive their power from

- Model Intelligence
- Access to sensitive Data
- Autonomous decision making

That same combination makes them a critical attack surface

Problem

Unlike traditional software, AI is attacked through language by anyone, not just expert hackers.

Existing cybersecurity tools were built for code vulnerabilities. They don't work here.

The Gap

- Traditional Red Teaming Works. But It Doesn't Scale.
- Current solutions (Lakera, Mindguard, Prompt Security) offer red teaming as a managed service. expensive, slow, and limited to who they can hire.
- There is no marketplace. No continuous stress-testing. No way for a company to know in real time how resilient their AI system actually is.

Our Solution

Sui Sentinel: Crowdsourced AI Red Teaming

- Companies deploy their AI systems as **Sentinels**, setting a bounty pool and message fee. A global community of red teamers competes to break them.
- Every attack is verified inside a **Trusted Execution Environment** using cryptographic attestations. Successful attacks trigger **instant, trustless on-chain payouts** via Sui blockchain.

How It Works

Defenders

1. Set up their AI system, define public/private instructions.
2. fund a reward pool.
3. receive a full attack dataset + resilience score as an audit report.

Attackers

1. Pay a small per-message fee to attempt jailbreaks, data extraction, and adversarial attacks.
2. Break the Sentinel, earn the bounty.

Protocol

A DSPY-powered jury model evaluates each attack. Results are cryptographically attested and settled on-chain automatically..

The Economics

Each message fee splits three ways



Failed attacks grow the bounty, attracting more attackers

The Economics

The longer a Sentinel holds, the bigger the prize. This creates a self-reinforcing flywheel. more attackers, more data, better defenses.

Defenders also earn **continuous Sentinel token rewards** proportional to their pool size, incentivizing early liquidity.

Why Sui ?

01

Verified authenticity

cryptographic proof that an attack genuinely happened, neither party can dispute

02

Trustless Settlement

No intermediary needed to release funds

03

Transparent Incentives

Token rewards, pool mechanics, and fee splits are all on-chain and auditable.

A centralized platform cannot offer this credibly.

Available Market

The AI security market is projected to reach **\$60B+ by 2028**, driven by enterprise AI adoption and regulation.

Our Web2 competitors are service businesses with linear scaling. We are a platform, it scales with participants, not headcount.

Target segments: AI-native startups, enterprise GenAI deployments, LLM API providers, and any company using AI agents with access to sensitive data.

OUR JOURNEY

June 2025 Won Sui Overflow Cryptography track prize

Oct 2025 Presented at Sui Fest Singapore

Jan 2026 Contract audit completed by Ottersec team

Feb 2026 Live On Sui Mainnet

Building in stealth for months. prioritizing product quality over early noise

We are looking for investors and advisors who see AI security as the next critical infrastructure layer. and want in early.

OUR MISSION

*To make AI security a public good
accessible, verifiable, and powered
by the crowd.*

OUR VISION

*To crowdsource AI red teaming at
scale and help build safer systems.*

Thank you. We would love to connect with you. Any feedback is appreciated.



satyam@cyphronixsoftware.xyz



<https://suisentinel.xyz>

Twitter/Telegram: @satyambnsal