

ALICE

Ship Agents You *Can Trust*

WonderSuite automates red-teaming and enforces policy-tuned guardrails across the entire AI lifecycle. Created by Alice, the team trusted by 8 of the 10 leading frontier model labs for AI safety research and alignment.



WonderSuite: One security layer from pre-launch to production.

WonderBuild

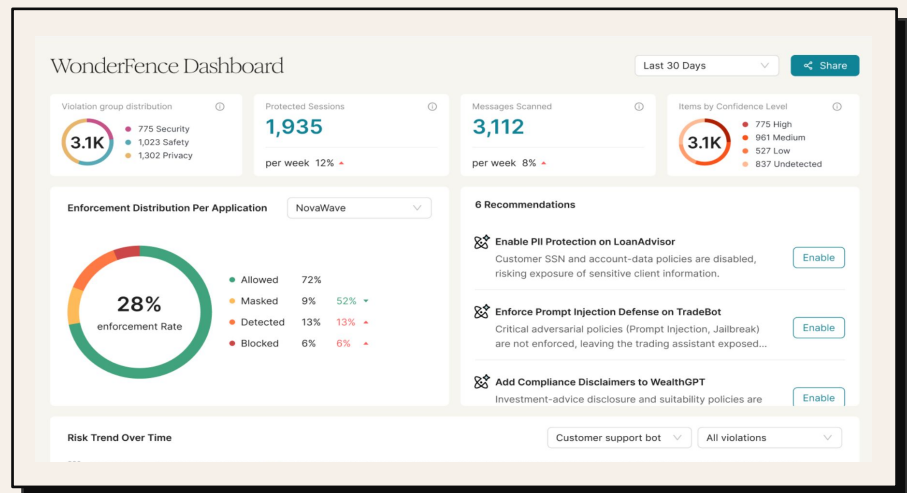
Surface real failure modes before launch with automated red-teaming. Multi-turn and mapped to the EU AI Act, ISO 42001, NIST, and OWASP.

WonderFence


Enforce pre-configured or personalized policies at runtime with multilingual, multimodal guardrails at sub-99ms latency.

WonderCheck

Catch drift, regressions, and emerging threats with scheduled red-teaming in production. Ensures agentic behavior stays consistent with your policies.



TRUSTED BY THE WORLD'S LEADING TECH INNOVATORS



Powered by *Rabbit Hole*.

A decade of real-world adversarial data continuously updating across 3B users, 1B+ daily AI interactions, and 120+ languages. Every test, guardrail, and evaluation runs on threats adversaries actually use.