USE CASE

Answer the question every buyer asks about your agent

Aribot shows what your agent can touch, keeps a verifiable record of every action, and threat-models the agent itself.

THE PROBLEM

Your product is an AI agent. Every customer security team asks the same thing: what can it reach, who approved that, and what did it do last Tuesday. The honest answer today is a diagram from launch week and a promise.

HOW ARIBOT HANDLES IT

1. **Threat-model the agent**
   Model the agent's architecture layer by layer, including its tools and data access.
2. **Govern its access**
   Give each agent identity the scopes it uses and refuse the rest at the gateway.
3. **Record every action**
   Each decision lands in a hash-chained ledger that a customer's auditor can verify.
4. **Call it from your stack**
   Your own agents and coding assistants call Aribot over MCP with the same controls.

WHAT YOU SEE IN THE DEMO

- Agent Assurance: granted versus used, with verdicts
- Ledger Integrity: the chain verified on demand
- MCP Platform: governed calls from assistants

WHO IT IS FOR

Head of AI Engineering

Founding Security Engineer · CTO

WHAT YOU GET

- A threat model of the agent, not just the app around it
- Granted versus used access for every agent identity
- A verifiable record of agent actions for customer reviews
- Evidence for ISO/IEC 42001 and EU AI Act questions

#1

Prompt injection ranks first in the OWASP Top 10 for LLM Applications 2025; excessive agency is among the most expanded entries

Source: OWASP GenAI 2025

EU AI Act · general-purpose AI · 2 Aug 2025

GPAI obligations and AI Office powers apply; customers ask how your agent is controlled.