

KEY NuGuard Differentiators

Why NuGuard is the only open-source AI Trust Platform that closes the agentic behavior gap before production.



01

DISCOVER

AI-SBOM Agent Stack Discovery

Maps every tool, permission, and API endpoint reachable by your agent — automatically, in minutes, with no code changes. 53% of organizations have had agents exceed their intended permissions. You cannot govern what you cannot see.



02

DEFINE

Machine-Readable Behavioral Policies

Turns soft guidelines into hard, executable constraints — written in plain English, version-controlled alongside agent code. Not buried in system prompts. Intent without policy is the root cause of the specification gap.



03

VALIDATE

Automated Red Team Testing

Runs boundary probing, injection resistance, and capability enumeration before a single production request is processed — built dynamically from your AI-SBOM and policy set. Validation is a hard gate. Agents that fail never ship.



04

GOVERN

Audit-Ready Evidence, Every Run

Every run produces a coverage score, policy pass/fail results, and remediation guidance. CISOs authorize. Regulators audit under EU AI Act and NIST AI RMF. AI governance becomes a repeatable, automated workflow — not a retroactive scramble.

READY TO SEE IT IN ACTION?

Validate your AI agents before they go live.

Experience our Platform >