

The Guild Service Ladder

Independent, vendor-neutral engagements — from first look to continuous assurance.

AI speed amplifies hidden weaknesses. Code that would have taken a quarter to write now lands in a week — and so do its flaws. Hallucinated APIs, shallow tests, over-permissioned agents, structural debt, and silent security regressions all ship at the same velocity as the features that impress in a demo. The faster a codebase grows under AI assistance, the more it needs an independent party whose only job is to see what the speed conceals. **We sell no tooling, so every recommendation is the one we'd give a friend.**

THE LADDER · FOUR RUNGS, STEP ON AT ANY ONE

1

ENTRY

Diagnostic

A short, focused conversation and surface assessment of your AI-assisted practice — how AI is used, where the obvious risks sit, whether deeper review is warranted. The cheapest way to find out whether you have a problem worth solving. It ends with an honest answer, not a sales pitch.

2

INSPECTION

Audit

Independent, candid review of an AI-assisted codebase: hallucination detection, security vulnerabilities, test integrity, dependency & license risk, architecture coherence, data & error handling, agent configuration. The audit tells you what your codebase *actually contains* — not what the tooling reports it contains.

3

REMIEDIATION

Optimization

Once the audit has surfaced what is broken, bloated, or wasteful, optimization removes it — reducing token waste, eliminating dead weight, tightening structure, and lowering the ongoing cost of running an AI-assisted codebase. The audit finds the problems; optimization eliminates the conditions that create them.

4

ASSURANCE

Continuous Security Assurance

A standing, vendor-neutral relationship that verifies your AI-assisted code and agent behaviour keep meeting the open security canon — OWASP Agentic Top 10, OWASP LLM Top 10, SAFE-MCP, NIST AI Agent Standards, and applicable regulation — on the cadence your work actually moves at. Point-in-time audits expire; continuous assurance does not.

GUILD AMBITION

The impartial layer that watches the watchers.

The market is filling with instruments — security fabrics, governance platforms, evaluation harnesses, agent monitors. Each is a watcher sold by a vendor with a roadmap to protect and a story to tell about its own coverage. Useful tools. But a tool cannot grade itself, and a vendor cannot grade its competitors. The Guild is the party with nothing to sell except the accuracy of its assessment: we verify whether your tooling enforces the standards it claims to, whether your audits cover the surface that matters, and when a dashboard is reassuring and wrong.

Independence is not a marketing posture. It is a structural commitment: no platform, no fabric, no proprietary taxonomy, no referral fees — just craft-grounded human judgment, against open standards the whole industry can read.

WHERE TO STEP ON

Not sure you have a problem?

Start with the **Diagnostic** — lowest cost, lowest commitment.

Shipping AI code to production?

Start with an **Audit**. Scrutiny before launch beats scrutiny after.

Already audited — now what?

Move to **Optimization**. Findings only pay off once remediated.

Running agents at scale?

Engage **Continuous Assurance**. Point-in-time can't keep up.