Know Your Agentic AI
Trust should be checkable, not promised.

Rume Dominic
AI accountability engineer
rumedominic.com

Your AI agents act. Can you prove what they did?

Your clients are putting AI agents into real work: spending, writing, deciding, acting on their own. When a board, an auditor, an insurer or a customer asks what an agent actually did, what it was allowed to do, and what it was stopped from doing, most teams cannot answer. That gap is now your risk, and theirs.

WHY THIS LANDS ON YOUR DESK NOW

- **Agents act without a human in the loop.** A logged action is not the same as a provable one.
- **Everyone is starting to ask.** Boards, auditors, insurers and customers want evidence, not assurances.
- **The rules are arriving.** New AI regulation expects a record you can query of what an AI decided. The UK is funding this market toward 18.8 billion pounds by 2035.

WHAT A TELLING-THEM-TO-BEHAVE APPROACH MISSES

Telling an agent in its prompt not to do something is a suggestion, not a control. Under pressure the instruction can be dropped, worked around, or quietly ignored, and the dashboard still shows green. You find out after, not before.

The fix is not a bigger prompt. It is a record and a limit the agent cannot step outside, and that anyone can check.

THE OFFER, AI COVER

Continuous assurance for your clients' AI agents. Not a one-off audit.

We give each agent a tamper-evident record of what it did, what it was allowed to do, and what it was stopped from doing, written so a stranger can verify it in seconds. Real limits are enforced in the system, not the prompt. It runs continuously and renews, so the proof is always current. You carry it to your clients under your name. We do the assurance behind it.

WHY YOU CAN TRUST THIS, IT IS ALREADY CHECKABLE

- **An open receipts standard** for verifiable agent records, published on PyPI.
- **An open auditor** that measures whether an AI built what it was told to, downloaded more than 5,000 times.
- **An independent MSc study** that measured four leading AI coding tools on governance, pre-registered and blinded.
- **A live public-record deployment**, with every record checked and holding.
- **We publish our own limits and mistakes.** If a tool fails its own test, we say so. That is the whole point.

Start with a free check of one agent.

We take one of your clients' agents and show you, in writing, what can and cannot be proven about it today. No cost, no commitment.

15 minutes to begin