

The independent third eye for enterprise AI

As AI moves from **assisting** the enterprise to **running** it, we are the **only** independent, deterministic auditor built to referee that AI. Independence here is **structural, not a claim of superiority** — the vendor that sold it, and an AI grading an AI, cannot fill this seat. **Deterministic. Independent. Beholden to no vendor.**

THE SHIFT — AI IS NOW DOING THE WORK, DECIDING AND ACTING

AI has stopped being a copilot. It migrates the system, writes the code, closes the books, and runs agents inside the business. On **29 September 2026** the White House and every major AI company — Anthropic, OpenAI, Google, Meta, xAI, Nvidia — signed an accord making an **independent external auditor** a pillar of AI safety: the makers themselves conceding AI cannot be trusted on its own word. Meanwhile the vendors keep shipping agents that *act* — **Microsoft** gave Dynamics 365 agents **write access** to finance records; **Oracle's Fusion Claw** issues an auditable "Outcome Receipt." The category is real, and it is arriving faster than anyone can govern it.

THE GAP — THE FOX GUARDS THE HENHOUSE, AGAIN

The accord names the seat — and leaves it empty. It does not stipulate *who* carries out the audits, and the companies engage their own. A supplier grading its own agents, or picking its own auditor, is not independent assurance — and no board, auditor or regulator can take the doer's word as proof. The referee of AI cannot itself be an AI, and cannot be the party that built or deployed it. That is the one thing the accord left open. We took it.

THE FOUR QUESTIONS THE AI VENDORS CANNOT ANSWER ABOUT THEMSELVES

Which AI?

Vendor-neutral fit — which AI platform, agent or model, assessed on your own estate, not a pitch.

Did it do it right?

Independent reconciliation of what the AI actually did against your own pre-state — says-done vs is-done.

Can you trust the autonomy?

Agent authority, spend / blast-radius, drift and segregation of duties — governed, from the trail it leaves.

Can you evidence it?

A deterministic, traceable record for the board, the auditor and the regulator — every claim to source.

WHY ONLY US — THE MOAT

- **Deterministic** — you catch a probabilistic AI with rules and recomputation, not another AI (an AI checking an AI can hallucinate the check).
- **Independent** — we don't build or deploy the AI, so we can flag it; a vendor grading its own agent structurally cannot.
- **Zero-egress** — verified against your own source, on your machine; nothing leaves, nothing trusted on faith.
- **No-fabrication is our DNA** — PRAIS was born refusing to assert anything that doesn't trace to a captured input. Turned outward, that *is* hallucination detection.

FROM THE WEDGE TO THE CATEGORY

Today we referee AI-delivered ERP transformation and the agents running your ERP — where our deterministic engine, our buyers and our channel already are. The same engine, re-aimed, is the third eye for **all** of enterprise AI: any high-stakes work where the doer can't be the proof.

THE TAILWIND — VERIFICATION IS GOING FROM OPTIONAL TO REQUIRED

The White House accord itself floats codifying independent audit into law; as the EU AI Act, DORA and sector rules land, independent verification of AI stops being a nicety and becomes a duty. We are the independent party that can produce it — **evidence, not assurance theatre.**

THE THIRD EYE — ONE ENGINE, TURNED OUTWARD AT THE AI

Capture your own data, read-only and zero-egress → deterministic rules → an evidence-traced verdict, with nothing fabricated. The discipline that made PRAIS trustworthy on its own reports is the discipline that catches what an AI over-claims. Prove what's provable, gate what isn't, overclaim nothing — the independent record the board and the regulator will demand. This is **engineering, not a questionnaire.**
Roltrader · the third eye for enterprise AI.