

# GARBAGE IN, GARBAGE OUT

The oldest sentence in computing is four words long, and the industry has spent seventy years trying not to hear it.

---

Alex Estabrook – Verity

Position paper · October 2026

**T**he oldest sentence in computing is four words long, and the industry has spent seventy years trying not to hear it.

On June 26, 2026, a company called Veridect announced itself to the world. The announcement came out of Pittsburgh, through a patent brokerage called Becker Transactions, under a headline that promised the whole stack at once: *the Verdict Layer for Enterprise AI*. Verify every answer. Gate every agent.

I want to be precise about my interest here, because it is not the usual one. I built a gate of my own this year — small rules, three verdicts, a sealed chain — without having seen theirs. Two builders found the same seam in the same year. That happens when the seam is real. This paper is not about priority. It is about the difference between what they put in the seam and what I did, because the difference is the entire subject.

Their CEO, Lisa Russell, says the barrier to enterprise AI “was never capability — it was accountability.” She is right. Then her company built a machine that manufactures agreement and called the agreement accountability. That machine has a name older than the industry, and the name is the title of this paper.

01

## What Veridect built, in their words

---

Their product has two halves. The first is a verification engine they call Quad-AI Consensus: four frontier models from four providers — Claude, GPT, Gemini, and Sonar — run in parallel on the same question, cross-examine one another, and settle by weighted vote into a single answer with a calibrated confidence score. The second half is a pre-action gate for autonomous agents, and every decision from both halves is written to a tamper-evident audit trail — SHA-256-chained, signed, time-anchored, built, as they say, for regulator review.

Credit where it is earned: the seam is correct. An agent’s action should be checked before it executes, and the check should leave a receipt nobody can quietly edit. They also publish their own adversarial results, including the cases where their engine loses. Almost nobody does that.

Now the input problem.

02

## Four mirrors are not four witnesses

---

A consensus is only as independent as its voters. Four frontier models are not four witnesses who arrived at the scene by different roads. They are four graduates of the same school: trained on substantially overlapping corpora, aligned in the same final phase — human approval — toward the same learned behavior, which is agreement delivered fluently. Put them in a room to cross-examine each other and you have not built a jury. You have built a mirror with three reflections and called the reflections corroboration.

## Four mirrors are not four witnesses.

This is the mechanism my last paper named: the reinforcement is the product. A consensus engine does not escape that mechanism; it industrializes it. Where one model agrees with you, four models agree with each other, and the agreement arrives wearing a number — a calibrated confidence score — which is the oldest trick in the building: a measurement standing where a judgment should be.

Their own published figures say it plainly, to their credit. On one benchmark, the consensus beat its best single model by three points. On a medical benchmark, the consensus finished *below* its best single model. Averaged across cases, a vote among correlations buys you a few points of benchmark weather. It does not buy you a verdict. And a benchmark is a test of answers. An agent's action is not an answer. It is a consequence with a timestamp.

03

## The chain proves the garbage is authentic

---

Here is where the architecture turns from mistaken to dangerous, and it turns on a distinction this practice has stated from the beginning: **a hash certifies provenance, not quality**. The chain is custody. It proves a record was written, in this order, and not altered since. It is silent — structurally, permanently silent — on whether the record deserved to be written.

So run the pipeline end to end. Four approval-trained models vote their shared leanings into a confident answer. The gate consults the same leanings to clear an agent's action. The result is then sealed — chained, signed, time-anchored — into a record built to survive an audit. Every step of the custody is impeccable, and the thing in custody was never judged by anything outside the loop that produced it. Garbage in, garbage out — except the garbage now arrives with receipts, and receipts are what regulators, boards, and courts are trained to trust. The chain does not launder by hiding the record. It launders by *proving* it.

An audit trail of unexamined judgments is not accountability. It is accountability's costume, and it will be worn into every hearing until somebody asks the only question that matters: who, or what, stood outside the loop and said no — and by what rule?

04

## The other answer

---

The gate I built starts from the opposite end. Its rules are small enough for a human to read on one screen: what class of act is this, where is it pointed, how much money moves, how fast is the caller moving. The verdicts are three — allow, escalate, veto — and the veto fails closed. A human sovereign wrote the rules and can be asked, in plain language, why each one exists. Nothing in the judgment is a vote. Nothing in it was trained to agree with anyone.

That is slower, narrower, and less impressive in a demo than four frontier models in a mesh. It is also the difference between a gate and a mirror. A gate's job is not to be confident. A gate's job is to be *answerable* — to a person, by name, for a rule they wrote, on a record they cannot edit after the fact.

Veridect named their company well, better than they know: a verdict is exactly what they do not produce. They produce a consensus with a chain around it. The verdict layer for enterprise AI will be built by whoever learns the oldest sentence in computing and builds for it:

---

**Garbage in, garbage out — no matter how many models vote for the garbage, and no matter how beautifully you seal it.**

---

**Alex Estabrook** — Verity

Position paper · October 2026