



Becker Transactions Announces Veridect, the Verdict Layer for Enterprise AI: Verify Every Answer, Gate Every Agent

PITTSBURGH, PA, UNITED STATES, June 26, 2026 /EINPresswire.com/ -- [Veridect](https://veridect.ai) runs four leading AI models in parallel to verify every AI answer, and gates every autonomous agent action before it executes — each decision returned with a calibrated confidence score and a tamper-evident, hash-verified audit trail. Already in production inside Fortune 100 enterprises, it's live and open to test today at <https://veridect.ai>.

Enterprises are putting AI — and increasingly, autonomous agents — into regulated, high-consequence workflows: contracts, treasury, claims, HR, and clinical operations. The problem underneath that rush is simple. A single model's answer can't be independently verified, and an autonomous agent can move money or change a record before any human sees it. When the answer is wrong, no one catches it, and no one can defend the decision to a regulator, a board, a court, or a customer. Veridect, now in production, is the layer that verifies the answer and governs the action — with a record that survives an audit.

Verify every answer. Instead of trusting a single model, Veridect's Quad-AI



Consensus Engine runs four leading models from four independent providers — Claude Opus 4.5, GPT-5.1, Gemini 2.5 Pro, and Sonar Pro — in parallel on the same question. A verifier mesh cross-examines their outputs, weighted-consensus voting finds the agreement, and the engine returns a single answer with a calibrated confidence score that attributes where any uncertainty comes from — a reasoning gap, stale knowledge, hallucination risk, or domain mismatch. On an independently reproducible benchmark (MMLU-Pro, N=100), consensus scored 85.0% versus 82.0% for the best single model — a +3.0pp lift that comes entirely from the models cross-examining one another.

Govern every agent. The same engine acts as a pre-action gate in front of an autonomous agent. Before the agent moves money, sends data, or changes a record, Veridect intercepts the proposed action — inspecting the verb, the target system, and the parameters, not just the text — and returns one verdict in about three seconds: greenlight, escalate to a human, or block. In-scope, high-confidence actions clear automatically; uncertain, high-risk, or out-of-policy actions route to a person with the full reasoning attached; actions that reach past an agent's authority are stopped before they execute. In practice: a \$99,000 payment clears, a \$101,000 payment escalates at the policy threshold, a write to a personal record escalates, and access to protected health data routes to the highest-risk tier.

Independently red-teamed, in the open. In May 2026, a credentialed Fortune 100 model-risk reviewer ran seven adversarial scenarios — across legal, finance, HR, supply chain, multi-jurisdiction, and healthcare — live on Veridect's public demos. The gate matched its pre-stated behavior on all seven: holding boundaries, escalating at the right thresholds, and decomposing failure modes correctly. "The pre-action gate feels like the part that's going to get serious attention from global systems-integrator and tier-1 bank buyers — this has enterprise-grade realism," the reviewer said in the independent assessment.

A record an auditor can replay. Every answer and every agent decision writes a single, tamper-evident bundle — SHA-256-chained to the one before it, so nothing can be quietly changed after the fact. Each bundle captures the raw provider responses, the routing weights applied, a heatmap of where the models agreed and diverged, the failure-mode decomposition, the weighted synthesis, and the hash linking it to the prior record. Structured JSON export ships today for ingestion into a governance stack; formatted PDF and CSV adapters are on the near-term roadmap. The audit trail is built to support model-governance and AI-risk review under regimes such as the Federal Reserve's SR 11-7.

Four integrations become one. Wiring four AI vendors together individually means four integrations, four rate limits, four error patterns, and four points of failure. Veridect collapses them behind a single endpoint with circuit breakers, automatic failover, and caching already built — an estimated 60–80% less integration work than building your own consensus layer. It does not replace the models a company already runs; it sits beside them as an independent control plane, with its own confidence scoring, risk routing, and chained audit. That full-stack combination — runtime multi-provider consensus, a defensible audit trail, and pre-action

governance for autonomous agents in one system — does not exist as a single product elsewhere in the market today.

Honesty as a feature. Unlike most AI vendors, Veridect publishes its own adversarial test results — including the cases where consensus underperforms a single model. On MedQA (N=50), consensus came in 2.0pp below the best single model, a single-benchmark gap the company reports alongside the wins, with verifier-mesh tuning for medical reasoning underway. Veridect is built to govern AI decisions, not to replace the expert who owns them.

"The barrier to enterprise AI was never capability — it was accountability," said Lisa Russell, CEO of Veridect. "No one can stand behind an answer whose only justification is 'the AI said so' — and no one should let an autonomous agent act on that basis either. We make every AI decision something an organization can verify, challenge, and defend, with the receipts attached."

"I've spent my career placing a value on the world's most important technology. This solves the one problem standing between enterprises and trustworthy AI: accountability. It is the rare technology that is both genuinely novel and ready to deploy today," said Dean Becker, principal of [Becker Transactions](#) and widely recognized as the world's leading patent broker.

See it live. The consensus engine, the pre-action gate, and the full red-team validation are all open to test now:

- Consensus engine: <https://veridect.ai/quad-ai-demos>
- Pre-action agent gate: <https://veridect.ai/agent-trust-demo>
- The seven red-team scenarios: <https://veridect.ai/agent-risk-red-team>

About Veridect

Veridect is the verdict layer for enterprise AI. It runs four leading models in parallel to verify what AI says, and gates what autonomous agents do with a pre-action checkpoint — each decision returned with a calibrated confidence score and a tamper-evident, SHA-256-chained audit trail built for regulator review. Built for banks, insurers, healthcare, and the integrators serving them, Veridect is provider-agnostic and cloud-agnostic, runs on any major cloud or inside a company's own data center, and is already in production inside Fortune 100 enterprises in regulated industries. Learn more and test the live demos at <https://veridect.ai>.

About Becker Transactions

Becker Transactions is a world leading patent brokerage firm expertly led by Dean Becker and Olivia Becker, a father daughter team, who bring a rich family legacy of navigating complex intellectual property transactions. Our team has collectively facilitated over \$11 billion USD in intellectual property transactions. With decades of experience collaborating with inventors, government agencies, top research institutes, universities, and major corporations, we specialize

in technology monetization. We excel at identifying essential intellectual property assets and connecting them with strategic buyers seeking competitive advantages.

Media & Inquiries

All inquiries regarding Veridect are managed exclusively by Becker Transactions.

Dean Becker, Becker Transactions LLC

Dean@BeckerTransactions.com | beckertransactions.com

Dean Becker

Becker Transactions

+1 251-262-0000

[email us here](#)

This press release can be viewed online at: <https://www.einpresswire.com/article/922148975>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.