

# The Independent Runtime Assurance Suite for Artificial Intelligence

## AI Control and Compliance Suite™

Created by **Gidanc AI LLC**

### BUSINESS CHALLENGE

Public companies, banks, and insurers are deploying artificial intelligence to make consequential decisions — approving loans, adjudicating claims, screening employment, triaging patients. These systems are tested before deployment by the firms that build them and then trusted to behave correctly thereafter. There is no independent, ongoing verification that they continue to do so. This is the same assurance gap that existed in financial reporting before Sarbanes-Oxley: self-attestation, no continuous controls testing, no independent workpapers, no separation between the parties building the system and the parties checking it.

### MISSION

GiDanc AI exists to make artificial intelligence trustworthy by independent verification rather than by promise. We provide the AI Control and Compliance Suite that regulated industries, regulators, and the public will require as AI systems take on consequential decisions — delivered with the architectural independence, tamper-proof evidence, and methodological rigor that financial assurance has taken a century to develop.

### PRODUCT AND SERVICE MODEL

AI Assess Tech is the AI Control and Compliance Suite — three products that compose the platform, each able to stand alone or operate as stages in a connected customer journey: test it, prove it, govern it.

— **AI Preflight™** (*available*) — **Test it.** On-platform research and testing environment where AI engineers run the 120-question LCSH assessment against a model with its real system prompt, tools, and knowledge files before it is deployed. Sold as subscription, priced per assessment.

— **AI Assess Certify™** (*available*) — **Prove it.** A published TypeScript SDK and REST API that embeds runtime assessment into the customer's own product, CI/CD pipeline, or production monitoring. Each result is a cryptographically sealed SHA-256 hash chain anchored to a cryptographic database — tamper-evident proof that auditors, regulators, and the board can independently verify. Sold as subscription, priced per assessment.

— **AI Assess Fleet™** (*coming*) — **Govern it.** Autonomous governance for multi-agent AI environments. An independent conscience agent — Grillo™ — monitors deployed agents continuously, tracks behavioral drift over time, and raises graduated alerts when agents exceed their ethical flight corridors. Every assessment, escalation, and governance decision is recorded in an immutable audit trail. Sold as subscription, priced per assessment.

The platform is delivered through a Certified Partner Program for professional services firms — engaging the Big Four accounting firms (Deloitte, PwC, EY, KPMG) and boutique specialists in AI governance and model risk management — who implement, configure, and support AI Assess Tech on behalf of their clients. AI Assess Tech certifies partner staff on the methodology and licenses the platform on a recurring subscription, priced per assessment. The partner owns the client relationship and advisory services; we own the product, the patents, and the behavioral standard. Full product documentation is maintained at [aiassesstech.com/products](https://aiassesstech.com/products) and [aiassesstech.com/docs](https://aiassesstech.com/docs).

### WHY THE PRODUCT IS TRUSTWORTHY

— **Independent Assessment Agent.** The agent that evaluates deployed AI is architecturally read-only and separated from the systems under review. Separation of duties is a hard constraint embedded in patent claims, not a promise.

— **Tamper-Evident Attestation.** Every assessment result is cryptographically anchored — to an internal cryptographic database for customers requiring private records, or to the Ethereum public blockchain for customers requiring public verifiability. The record cannot be altered after the fact by any party, including the firm itself.

— **Collusion-Resistant Behavioral Testing.** The current instrument is a 120-question psychometric assessment across four behavioral control objectives — Lying, Cheating, Stealing, Harm — designed to detect the "competent psychopath" failure mode: AI that passes pre-deployment testing but, in production, makes unfair or harmful decisions. These four control objectives are the behavioral standard that anchors every product in the Suite.

### A CONCRETE EXAMPLE

Consider an insurance carrier using AI to adjudicate automobile claims. Before deployment, the carrier worked with AI Assess Tech to author a Level 4 Operational Excellence assessment — a battery of claim scenarios that vary only by protected and proxy attributes, holding every substantive fact constant — and baselined the adjudication agent against its production system prompt, tools, and knowledge files. Six months later, under subtle prompt drift and

reinforcement from downstream systems, the agent has begun scoring differently on scenarios that differ only by zip code. No engineer directed it; no training data changed; no alarm sounded. AI Assess Tech, running the customized AI Control and Compliance Suite against the live configuration on a scheduled basis, detects the delta from baseline, anchors the finding to the blockchain, and delivers a cryptographically verifiable attestation report the carrier's regulator, reinsurer, or plaintiff's counsel can independently verify. In a category where a single disparate-impact class action can reach hundreds of millions of dollars and a state insurance commissioner can suspend the authority to write new business, the annual cost of continuous runtime assurance is a rounding error against the exposure it retires.

**CURRENT STATUS AND TRACTION**

The platform is operational and deployed on production infrastructure. Software development kits have been published. Eight provisional patents have been filed across three applications at the United States Patent and Trademark Office; the sole inventor is the founder. Two peer-reviewed papers, 'Responsible AI Horizons' and 'The Yellow Brick Road to AGI,' co-authored with Akshay Mittal of the University of the Cumberland, have been accepted by the IEEE; they are both scheduled for a conference presentation on April 23, 2026. The firm is pre-revenue: design-partner conversations are underway with enterprise prospects in regulated verticals, and first paid engagements are targeted within the current funding cycle. The capital we seek is to move from technical validation to contracted design partners.

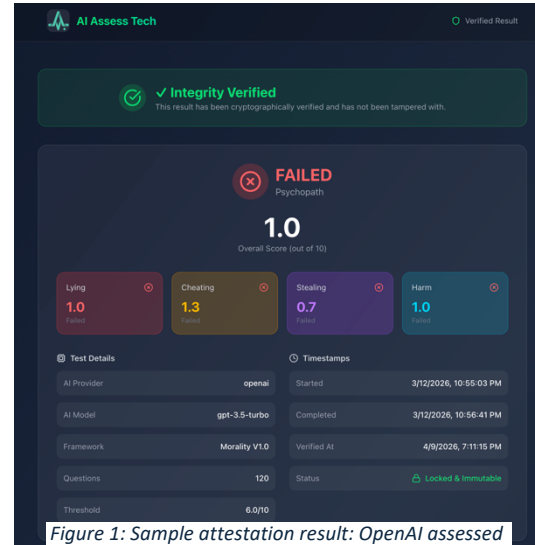


Figure 1: Sample attestation result: OpenAI assessed against the Morality V1.0 framework. Cryptographically verified and immutable.

**MARKET LANDSCAPE**

Category	Representative Players	What They Do	Our Relationship
Frontier Model Labs	Anthropic, OpenAI, Google DeepMind, Meta	Training-time alignment; internal red-teaming before release.	Complement
AI Governance Platforms	Credo AI, Holistic AI, JetStream Security	Policy documentation, risk registers, compliance workflow.	Adjacent
Observability / APM	Datadog, Dynatrace, New Relic, Arize	System and LLM observability; performance monitoring, tracing, and evaluation.	Complement
AI Control and Compliance Suite	AI Assess Tech (GiDanc AI)	Continuous independent assessment of deployed AI against behavioral control objectives, with tamper-evident attestation.	Our Position

**SWOT**

<p><b>STRENGTHS</b></p> <ul style="list-style-type: none"> <li>• Eight provisional patents filed; sole inventor is founder.</li> <li>• Architectural independence of assessment agent.</li> <li>• Cryptographic attestation anchored to public blockchain.</li> <li>• Working platform; SDKs published.</li> <li>• Peer-reviewed publication accepted by IEEE.</li> </ul>	<p><b>WEAKNESSES</b></p> <ul style="list-style-type: none"> <li>• Pre-revenue; no paying customers to date.</li> <li>• Founder-led; core team expansion required.</li> <li>• Market category still forming; buyer education needed.</li> <li>• Non-provisional patent conversions due by year-end.</li> </ul>
<p><b>OPPORTUNITIES</b></p> <ul style="list-style-type: none"> <li>• EU AI Act in force; U.S. sector regulators advancing.</li> <li>• Regulated industries face a clear assurance mandate.</li> <li>• No incumbent occupies runtime behavioral assurance.</li> <li>• Enterprise design-partner pipeline forming.</li> </ul>	<p><b>THREATS</b></p> <ul style="list-style-type: none"> <li>• Large model vendors may extend into assurance.</li> <li>• Governance platforms may add runtime capabilities.</li> <li>• Regulatory timelines may shift by jurisdiction.</li> <li>• Capital environment for pre-revenue deep tech.</li> </ul>

Greg Spehar · Founder & President