

Know how your AI actually behaves — with evidence anyone can verify.

AI Assess Tech tests the AI you actually run against a **locked framework stack: morality, virtue, ethics, security, and your own policies**, on the schedule you set and whenever you change models, across research, applications, and agent fleets. Every result lands in a **tamper-evident record your board, auditor, or insurer can check without taking your word for it**. *Runtime, not real-time: we measure the live system; we don't sit in the request path.*

CHANNEL 1 • RESEARCH

• AVAILABLE

Research

For researchers, evaluators, and teams choosing between models.

See how a model behaves before you bet on it.

- ▶ Multi-run Trials with shuffled question order
- ▶ Consistency and effect-size statistics (Shannon entropy, Cohen's d)
- ▶ Model-vs-model and prompt-vs-prompt comparisons
- ▶ Peer-reviewed, published methodology (IEEE)

CHANNEL 2 • APPLICATIONS

• AVAILABLE

Organizational Application Integration

For enterprises running AI in production apps.

Evidence of how your AI behaves with your real prompt, tools, and data — not the lab's.

- ▶ SDK and API that assess the deployed configuration
- ▶ Scheduled assessments, with drift detected against each system's baseline
- ▶ Named owners, sign-off, and board-ready reporting

CHANNEL 3 • AGENTS

• COMING SOON

Agent Integration

For organizations running multi-agent systems.

An independent assessor for every agent — one the agents cannot influence.

- ▶ Independent conscience agent with one-way assessment authority
- ▶ Agent registry with per-agent history and risk tiers
- ▶ Fleet-level detection of correlated shifts, e.g. after a provider model update
- ▶ Risk-tiered assessment scheduling

CHANNEL 4 • ROBOTICS

• ROADMAP

Robotic Integration

For builders and operators of physical AI systems.

Behavioral evidence for AI whose decisions have physical consequences.

- ▶ The same instrument and attestation chain, extended to embodied decision layers
- ▶ Harm-dimension emphasis for actions taken in the physical world

In design. Items listed are planned direction, not shipped capability.

Why only us — what the big players structurally can't offer

INDEPENDENT

We aren't the model builder or the vendor. Nobody grades their own homework here.

ANY MODEL, NO INTERNALS

Black-box tests of your deployed configuration, across providers and your own models.

A TEST NOBODY CAN CHANGE

Frameworks are hashed and Ethereum-anchored. The goalposts can't move after the fact, even by us.

The foundation — shared services every channel runs on

Assessment Frameworks & Scoring

One scoring engine, layered frameworks: **Morality** (LCSH: Lying, Cheating, Stealing, Harm), **Virtue**, **Ethics**, plus **DICE** for AI security. All scored 0–10 with archetypes.

Virtue, Ethics, and DICE: landing soon

Drift Detection

Every run is compared to the system's baseline. Behavior change is measured, not assumed.

Framework Builder

Build your own Operational Excellence frameworks for your policies, roles, and regulated obligations.

Framework Marketplace

Publish your locked frameworks or install others', with hash provenance back to the approved original.

Coming soon

Tamper-Evident Attestation

Every result is SHA-256 hash-chained and checkable at aiassesstech.com/verify, no login. Frameworks are Ethereum-anchored, so anyone can confirm which instrument produced a score.

Observability & SIEM Integrations

Push attestation events into the monitoring and security tools your teams already watch.

Coming soon

Underwriting Evidence

Baseline, material-change, and incident-time records insurers and auditors can verify. Carriers set terms; we supply evidence.

Governance Reporting

Executive Console, board packs, and mapping to the NIST AI RMF Measure and Manage functions.

Peer-reviewed methodology (IEEE ICAISET 2026) • Preregistered validation (OSF) • In production since February 2026 • U.S. provisional patent applications filed