AI ASSESS TECH

WHAT WE DO · *Independent Runtime Behavioral Attestation for Deployed AI*

THE PROBLEM

Deployed AI is opaque and non-stationary — it changes under you. Logs and evaluations produced by the party that built or runs the system are self-attested evidence; segregation of duties means a board, auditor, or insurer cannot accept them. Point-in-time testing fails a second way: the result goes stale the moment the deployment changes.

THE JOB

Be the neutral instrument that turns *"is my AI behaving?"* into evidence a skeptical third party accepts — and keep that evidence current.

WHAT WE PROVIDE

Independent runtime behavioral attestation for deployed AI — locked instruments (LCSH and related frameworks), persisted results, and public cryptographic verification. Enterprises run assessments **on demand**, on a **compliance schedule**, or via the **SDK on a production cadence** so the seal stays current as systems change. Boards, auditors, and insurers consume the same verifiable record without trusting the deployer's mutable logs.

WHY TRUST IT

Our validation protocol is preregistered and publishes where gates pass **and fail** — FAIL is FAIL as registered (osf.io/u2b3j). An instrument that only shows green has no teeth. Results are hash-chained, tamper-evident, and publicly verifiable at **aiassesstech.com/verify** — 565+ locked, publicly verifiable production runs.

WHERE IT FITS

ISO 42001 certifies your AI management *process*. We attest the *runtime behavior* of the deployed system — complementary evidence that GRC platforms and auditors ingest alongside process certification, observability, and policy controls.

WHAT WE DON'T CLAIM

No permanent "safe AI" badge. No ISO 42001 or statutory certification. No live-stream scoring. A **recurrent canary / going-concern attestation** with published limits — nothing more, nothing less.