

Someone is going to ask how your AI behaves. Be the one who can prove it.

Your board, your auditor, your insurer, and eventually a customer's lawyer will all ask the same question: **how do you know your AI is behaving the way you said it would?** Today most companies answer with model cards and good intentions. AI Assess Tech gives you a better answer: **the AI you actually run, tested against a locked framework stack and your own policies, on a schedule, in records anyone can verify without taking your word for it.**

REASON 1

The audit question is coming

Regulators, examiners, and boards have stopped accepting "trust us." NIST AI RMF, ISO/IEC 42001, the EU AI Act, and Illinois SB 315 point the same way: governance claims need independent evidence.

What you get: evidence a third party can check at aiassesstech.com/verify, with no login and no need to trust your logs.

REASON 2

Deployment changes behavior

Your prompt, tools, and data change how a model acts, and providers update models without asking you. Lab scores describe the model you tested, not the system you shipped.

What you get: scheduled assessments of your live configuration, plus a fresh one whenever you change models, with drift measured against your own baseline, so you see changes when they happen.

REASON 3

Insurers want something to underwrite

Carriers already condition cyber coverage on attested controls. AI liability is next, and there is no standard instrument yet. The organizations with evidence will be the ones that get quotes instead of exclusions.

What you get: baseline, material-change, and incident-time records built for underwriters and adjusters. Carriers set the terms; you bring the evidence.

REASON 4

You have to pick, and ship, with confidence

Which model? Which prompt? Is the new version better or just different? Most teams decide on demos and gut feel, then defend the choice after something goes wrong.

What you get: side-by-side behavioral comparisons and a record showing why you chose what you shipped, with a named owner who signed off.

The economics — what behavioral evidence costs today

MANUAL QUARTERLY REVIEW

\$40K–\$80K / yr

Four snapshots a year, manual reports, nothing an outsider can independently verify.

EXTERNAL AUDIT ENGAGEMENT

\$100K–\$300K / yr

Weeks-long engagement; the report describes a system that has already changed by delivery.

AI ASSESS TECH

\$600–\$12K / yr

\$50 to \$1,000 a month by volume. Recurring assessments, drift detection, verifiable records.

Indicative ranges. Your costs will vary; AI Assess Tech complements, and does not replace, your auditors and red team.

BOARD & CEO

"Are we governing the AI we deployed?"

A board pack with named owners and a verification link.

CFO & GENERAL COUNSEL

"What's our exposure?"

Reasonable-care evidence at a fraction of audit cost.

CISO & CTO

"What changed after the last update?"

Drift detected against baseline, per system.

COMPLIANCE & AUDIT

"Where's the proof, and can I check it?"

Tamper-evident records, verifiable without a login.

What we promise, and what we don't. We don't hand you a badge or a promise of perfect behavior. We measure how your AI behaves, show you when it changes, and give you records that hold up when someone who doesn't work for you checks them. That honesty is the product. *Runtime, not real-time: we measure your live system on a schedule and after changes; we don't sit in the request path.*

Try it before you believe it.

Run one free assessment against your real production AI this week. Predict the scores first. Then run it again with the prompt your team wants to ship next. **If the results match your guess, you've lost five minutes. If they surprise you, that surprise is the reason to buy.**