TrustModel.ai Launches Independent Assessment Platform for the Foundation Models Powering the Global 2000

TrustModel.ai launches independent assessment platform for foundation models powering Global 2000 to evaluate, monitor and certify AI to close assurance gaps

MOUNTAIN VIEW, FL, UNITED STATES, September 25, 2026 /EINPresswire.com/ -- TrustModel.ai, the independent AI assurance company and a Stanford StartX company, today announced its third-party assessment service for foundation models and the Global 2000 enterprises that deploy them. The service delivers rigorous, standardized and continuously verifiable evaluation, monitoring and certification of AI models and agents. Built to work alongside the world's leading consulting firms, it gives enterprises an independent platform next to the advisory partner they already trust.

> Trust is becoming a competitive advantage, and verification is how you earn it. That can't depend on any single lab, consultancy or vendor; it has to be an open, shared standard."
>
> *Karl Mehta, founder and chairman of TrustModel.ai*

The announcement follows a defining moment for the industry: a leading AI lab recently appointed a global consultancy as its first embedded third-party safety evaluator, establishing independent assessment as essential infrastructure for frontier AI. But that model has a structural blind spot. Many of the models enterprises actually run are open-weight models such as Llama, Mistral, DeepSeek, Qwen and Gemma. Once the weights are released, they are fine-tuned, quantized and deployed by thousands of organizations the original developer never sees — and no one commissions an independent assessment of the version running in production. The result is that Global 2000 companies are building on powerful open models with no independent assurance of their safety, fairness, accuracy or compliance.

"When a frontier lab wants to prove its model is safe, someone makes that call and hires an assessor," said Karl Mehta, founder and chairman of TrustModel.ai. "For the open-weight models running inside the Global 2000, there is no such decision-maker, so no one is independently assessing them - even as they're deployed into hiring, lending, healthcare and customer operations. That's the gap we exist to close. We rigorously and independently evaluate the models enterprises actually use, and we make the result verifiable."

TrustModel.ai is not a lab, a model builder or a single-vendor bundle. It operates as an AI Assurance Control Plane with three functions:

1. Evaluate — automated assessment across 10 trust dimensions, rolled into a single 0–100 TrustScore, using multi-judge consensus for reproducibility.
2. Monitor — a live control plane that watches deployed models and agents in production and detects drift.
3. Certify — a certificate authority for AI that issues signed, verifiable and revocable certificates, automatically revoked when a model or agent drifts out of tolerance.



TrustModel.ai Founder and Chairman of TrustModel.ai Karl Mehta

All results map to recognized frameworks — the NIST AI Risk Management Framework, ISO/IEC 42001, the EU AI Act and the OWASP Top 10 for LLM Applications — plus more than 30 sector regulations, so scores are comparable and audit-ready.

Critically, TrustModel.ai is built to the independence standard that governs financial auditing. The company is wholly founder-funded, with no outside investors, so no frontier lab or lab backer has a stake in its findings. It is not owned or governed by any frontier AI developer, has no reseller or channel business with the model developers it evaluates, and its fees are never contingent on a model's score. It builds no competing models and is not tied to any single consultancy: enterprises pair TrustModel.ai with the advisory firm of their choice. The services firm brings people, brand and relationships; TrustModel.ai provides the independent platform and the verifiable credential.

Open assessment for open models

To make its assurance transparent, and to give the community a shared, independent view of the models it relies on, TrustModel.ai is publishing its open-source model assessment methodology and a public TrustScore leaderboard on GitHub: https://trustmodel.ai/open-source. The project provides a transparent, reproducible framework for evaluating open-weight models against the same standards TrustModel.ai applies to frontier systems, so any enterprise, regulator or researcher can see how a model was scored and why.

Convening the Industry at Stanford

The announcement comes ahead of TrustModel.ai's AI Assurance & Governance Summit 2026 at the Stanford Faculty Club on Thursday, Oct. 1, where leaders from AI, enterprise and government will explore how verifiable trust becomes a competitive advantage for foundation models and autonomous agents.

"The standards for trustworthy AI already exist, from NIST's AI Risk Management Framework to ISO 42001 to the EU AI Act," Mehta said. "What's missing is an independent, verifiable way to hold models to them continuously, at the scale enterprises operate, including the open models no one is assessing. Trust is becoming a competitive advantage, and verification is how you earn it. That can't depend on any single lab, consultancy or vendor; it has to be an open, shared standard. That's what we'll work on together at the summit."

Availability

TrustModel.ai's assessment service is available now for foundation-model developers, enterprises and consulting partners. The platform is deployed across multiple clouds and is available on Google Cloud Marketplace. Organizations can request an assessment, and consulting firms can explore partnership, at www.TrustModel.ai.
https://console.cloud.google.com/marketplace/product/zippy-pad-480808-n8/pdx-trustmodel-v1

About TrustModel.ai

TrustModel.ai is the independent AI assurance company and a Stanford StartX company. Its AI Assurance Control Plane evaluates, monitors and certifies AI models and agents across 10 trust dimensions and more than 30 regulatory frameworks, and operates as a certificate authority for AI, issuing verifiable, revocable trust credentials. Founder-funded, structurally independent of the frontier labs and consultancy-agnostic by design, TrustModel.ai gives enterprises, developers and regulators a neutral, verifiable answer to a simple question: Can this AI be trusted? Learn more at www.TrustModel.ai.

Leslie Farrell
Farrell Communications, Inc.
+1 813-493-1805
Leslie@Farrellci.com
Visit us on social media:
LinkedIn
YouTube
X

This press release can be viewed online at: https://www.einpresswire.com/article/944997513

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.