REPORT · AI GOVERNANCE · TRUSTED AI MATURITY

Measured for Trust.

Why AI governance needs a maturity score, not a policy binder.

AUTHORED BY

Taiye Lambo & Dr. Tuboise Floyd

FEATURING THE TAIMSCORE™ FRAMEWORK
IN AFFILIATION WITH HISPI
EDITION 01 · 2026 · ATLANTA, GA

/ EXECUTIVE SUMMARY

A maturity score for institutional AI governance.

AI is reshaping how institutions operate. A board charter and a binder of policies do not prove governance holds when a model drifts, a vendor pushes an update, or an operator overrides a recommendation under pressure.

This report advances one argument. Governance that has never been measured is governance that nobody can defend. A maturity score converts the question every board now asks, how good is our AI governance, from an exchange of adjectives into a baseline that survives personnel turnover, vendor churn, and the next audit cycle.

The framework referenced throughout, TAIMScore™, the Trusted AI Model Score, was developed by Taiye Lambo and the Holistic Information Security Practitioner Institute (HISPI). It scores AI governance maturity across 72 controls organized into four domains — Govern, Map, Measure, and Manage — and aligns to NIST AI RMF, ISO/IEC 42001, SOC 2, and the EU AI Act at once. Human Signal serves as an authorized affiliate partner.

CONTENTS

01 / WHY THIS MATTERS

Failure rarely starts with a bad model.

Most governance breakdowns begin somewhere quieter. An institution writes the charter, names the committee, and inventories its use cases. Then it never builds the instrumentation to verify the system behaves as authorized, or the capacity to intervene when the system drifts. The paperwork exists. The control does not.

That distinction carries weight in regulated and high-stakes settings. A misrouted customer chatbot embarrasses a brand. A misgoverned model inside a clinical workflow, a utility's operational stack, or a public-sector eligibility decision degrades outcomes people depend on, and does so quietly until an incident makes the gap visible.

> *A documented policy is not the same as an executable control. One describes intent. The other can act at the moment a decision goes wrong.*

Frameworks have multiplied to meet this moment. The NIST AI RMF supplies vocabulary. Executive guidance, sector directives, and emerging law fill in the rest. Yet a gap persists between holding a framework and knowing where an institution actually stands inside it. Possession is not measurement.

The binder.

Ask most institutions where their AI governance stands today and the honest answer points at a document, not a number. A maturity score replaces the document with a position you can defend, compare across time, and improve.

02 / THE STANDARDS LANDSCAPE

Trust sits at the center of a standards crosswalk.

TAIMScore™ does not replace the international standards an institution already answers to. It organizes them. The framework reads as a concentric crosswalk: trust at the core, wrapped by the lifecycle processes that produce it, the risk management that bounds it, and the governance that authorizes it. Each layer maps to recognized normative references, which is what makes a resulting score defensible rather than self-asserted.

LAYER	NORMATIVE REFERENCES	WHAT THE LAYER COVERS
Governance	ISO/IEC 38507	Governance implications of the use of AI by organizations, and the board's accountability for it.
Risk Management	ISO/IEC 23894 NIST AI RMF 1.0 BS ISO/IEC 23053	Risk management guidance and the reference architecture for AI systems using machine learning.
Trust	ISO/IEC 42001 ISO/IEC 24028	The AI management system standard and the foundations of trustworthiness in AI.
Life Cycle	ISO/IEC DIS 5338 ISO/IEC 22989	AI system life cycle processes and the concepts and terminology that anchor them.

Read from the outside in, the message holds: governance authorizes, risk management bounds, lifecycle processes operate, and trust results. A program strong on the outer rings but hollow at the core has authorized and bounded a system it cannot yet trust. The crosswalk shows leaders which ring carries their weakness before an incident does.

03 / THE FOUR DOMAINS

Seventy-two controls. Four domains.

A TAIMScore™ assessment decomposes maturity across the four functions of the NIST AI RMF rather than reporting one undifferentiated grade. Weakness in one domain propagates into the others, so a complete assessment scores all four. A recurring pattern surfaces across documented failures: institutions cluster their strength in Govern and Map, then fall away in Measure and Manage, where verification and intervention have to happen in production.

DOMAIN	CONTROLS	SCOPE
GOVERN	19	Accountability structures, AI policy documentation, executive ownership, escalation authorities. Who owns the decision, what the escalation path is, who holds deactivation authority.
MAP	20	AI system inventory, risk categorization, model provenance, stakeholder impact mapping. You cannot govern what you have not mapped.
MEASURE	18	Performance monitoring, bias and fairness evaluation, outcome verification, continuous scoring. Governance without measurement is assumption.
MANAGE	15	Incident response, model retirement, vendor oversight, continuous governance operations. The domain where policy becomes practice.

A high Govern score beside a low Manage score names a precise, addressable condition. The institution has decided who is accountable but has not built the means to honor that accountability when a system deviates in production.

04 / HOW THE ASSESSMENT WORKS

Six steps to a score.

The methodology runs the same way for a single model or an enterprise portfolio. Each domain is scored against its controls, then aggregated into a single rating that prices the institution's exposure and benchmarks maturity over time.

01 Inventory the AI systems. Identify every system in the environment, including third-party and vendor-supplied models, classified by use case, data access, and deployment scope.

02 Score GOVERN. Apply the 19 controls to accountability structures, executive ownership, policy documentation, and escalation paths.

03 Score MAP. Apply the 20 controls to risk categorization, model provenance, and stakeholder impact.

04 Score MEASURE. Apply the 18 controls to performance monitoring, bias tracking, and outcome verification.

05 Score MANAGE. Apply the 15 controls to incident response, model retirement, and continuous operations.

06 Generate the score. Aggregate the domain scores into a TAIMScore™ rating used to prioritize remediation, document audit readiness, and benchmark maturity over time.

The output is not a certificate. The score states a position, names the gaps, and gives leadership a defensible artifact to act on and to show.

05 / A WORKED EXAMPLE

Air Canada, scored.

An airline's customer chatbot told a grieving passenger he could claim a bereavement fare retroactively. No such policy existed. When the passenger acted on the answer and later sought the refund, the airline denied it and argued the chatbot was a separate entity. A tribunal disagreed and held the airline liable. The model behaved as a model does. The governance around it did not exist where it needed to.

Scored against TAIMScore™, the failure implicates specific controls rather than a vague shortfall.

CONTROL	SUBJECT	FAILURE
GOVERN 1.1	**Accountability structure**	No defined structure governed what the chatbot was authorized to represent. The system carried no boundary separating informational content from binding commitments, so the institution absorbed liability it had no mechanism to prevent.
GOVERN 1.7	**Human review escalation**	No mechanism required human confirmation before the system issued definitive answers on refunds and fare policy. A bereavement query is exactly the high-stakes, emotionally charged context this control exists to gate.
MANAGE 1.1	**Incident escalation path**	When the refund request arrived, no process recognized a chatbot-generated policy claim as a systemic risk event. The denial that followed compounded the exposure rather than containing it.

None of these controls is exotic. Each is a baseline expectation under any serious framework. Their absence in a legally consequential deployment is a structural failure, not a technology failure. This case is one of twelve in the Human Signal Failure Files™, where real incidents are scored against the same four domains.

06 / WHAT AN ASSESSMENT GIVES YOU

Six deliverables.

1 **A quantified maturity score.** A shared, repeatable baseline that ends the argument over whether governance is "robust" or "developing." The number matters less than the discipline it imposes.

2 **A domain-level map of where governance breaks.** Maturity broken out across Govern, Map, Measure, and Manage, so leaders see which domain carries the risk before an incident shows them instead.

3 **A gap register tied to controls.** Specific deficiencies bound to specific controls, ranked by exposure. This converts "improve AI oversight" into a work queue someone can own, fund, and verify.

4 **Board-ready and acquisition-ready evidence.** An artifact that answers the audit committee, the contracting officer's diligence, and the insurance renewal with evidence rather than assertion.

5 **A prioritized remediation roadmap.** Gaps sequenced by where exposure runs highest and intervention costs least, rather than spread thin across every deficiency at once.

6 **A re-assessment cadence.** A fixed rhythm that keeps pace with model retraining, vendor updates, and staff rotation. Governance verified once and never re-verified decays in silence.

The gap register carries the most operational weight. Three questions decide whether real structure surrounds any AI system, and a credible assessment forces an answer to each.

Q1 Who owns the decision?

Q2 What does the escalation path look like?

Q3 What accountability exists without the vendor in the room?

07 / WHAT MEASUREMENT CANNOT DO

Honest about the limits.

A maturity score measures structure. It cannot supply judgment. An institution can post a strong score and still field a system that should never have entered the workflow, because no instrument substitutes for the human decision about whether an AI use case belongs at all.

Assessment also depends on candor during evidence collection. An organization that performs for the assessor rather than discloses to the assessor buys a flattering number and forfeits the value. The score reflects the honesty of its inputs.

> *A score does not eliminate risk. It locates risk, prices it, and assigns it an owner. The owning remains a leadership act.*

Who should measure now.

The case applies with the most force to institutions large enough to draw board and regulatory attention yet too small to seat a Chief AI Officer. For a CISO, CIO, or chief risk officer in that position, a current third-party assessment often marks the difference between a credible posture and a hopeful one.

The defense and enterprise community learned the parallel lesson in cybersecurity the expensive way, through years of self-attested postures that dissolved on contact with auditors and adversaries alike. Maturity measurement arrived because assertion had failed. AI governance now stands where security stood then.

08 / START HERE

Measure the distance.

You defend what you can see, and you can only see what you measure. A TAIMScore™ assessment establishes the baseline, names the gaps against 72 controls, and sets the cadence to keep them closed. The first step costs a conversation.

ENGAGE

Establish your baseline.

Start with a readiness conversation, establish the baseline with a TAIMScore™ assessment across Govern, Map, Measure, and Manage, and set the cadence that turns governance from a one-time project into an operational discipline.

REQUEST A READINESS CONVERSATION →

humansignal.io/measured-for-trust#readiness

A HUMAN SIGNAL REPORT
TAIMScore™ created by Taiye Lambo / HISPI
Edition 01 · 2026 · Atlanta, GA

Independence is not a feature. It is the product.