



AI Posture

One score. Three vectors.
Bounded by the weakest.



One Number You Can Defend

*Aggregated Intelligence Posture for
governance, risk, and compliance leaders*

A white paper by Sam Rogers

Stewarded by PAICE.work PBC

aiposture.org

June 2026 • Companion to “Aggregated Intelligence”

Executive summary

Every organization adopting AI now runs at least three governance efforts at once. One tracks whether people use AI well. One tracks whether digital systems are ready for agents and partners. One tracks whether the organization meets its AI-specific legal obligations. Each produces its own dashboard. None of them produces the single answer a board, a regulator, a partner, or a customer actually asks for: how ready are you, on the whole, and can you prove it?

AI Posture answers that question with one level, bounded by the weakest of its in-scope vectors. It is an output measure: it scores verified behavior, not policies, intentions, or tool purchases. It is orthogonal to NIST AI RMF, ISO/IEC 42001, and EU AI Act conformance, which govern program design and remediation. An organization can complete any of those programs and still hold a low AI Posture if the program has not yet produced behavior a third party can inspect. The reverse is also true.

This paper sets out the five design choices that make AI Posture defensible: a single readiness number across three independent vectors, an externally verifiable evidence standard, aggregation by the weakest link rather than a flattering average, treatment of the score as a time-stamped signal that decays rather than a certificate that persists, and a clear boundary around what the framework does not claim to do. The recommended action for a GRC leader is concrete: stop reconciling three dashboards by hand, and start reporting one posture whose constraining vector names your next investment for you.

Aggregated Intelligence is the collective output of different intelligences working together toward a clear intent. Today, that means biological and artificial intelligences. Together they can produce far more than they would apart, or far less; we measure the floor each reliably contributes, not the ceiling their collaboration might reach.

The readiness question no dashboard answers

Ask a chief information security officer, a chief AI officer, and a head of compliance the same question on the same day: are we ready for AI? You will get three confident answers that do not combine. The security lead points to agent-readiness scans and integration hardening. The AI lead points to training completion and adoption telemetry. The compliance lead points to an obligation register mapped against the jurisdictions that bind the organization. Each answer is true within its frame. None of them is the answer the board needs, because readiness is not the sum of three frames. It is the weakest of them.

The gap is structural, not a tooling failure. The three efforts measure different actor classes: humans, digital systems, and regulators. They mature on different clocks and they constrain each other in ways a side-by-side dashboard cannot express. A board reading three green

dashboards has no way to see that mature human practice is sitting on top of an unmet legal obligation, or that an agent-ready infrastructure is amplifying poor human judgment. The dashboards are not wrong. They are disconnected, and the disconnection is exactly where risk lives.

What leaders need is a single readiness level they can state out loud, defend under inspection, and act on. That level has to do three things the dashboards cannot. It has to combine the vectors without averaging away a weak one. It has to rest on evidence rather than assertion. And it has to come with an expiry, because readiness is a moving target. The rest of this paper builds that level one design choice at a time.

Why program frameworks leave the question open

The reflexive response to an AI readiness question is to name a framework. We follow NIST AI RMF. We are pursuing ISO/IEC 42001 certification. We have an EU AI Act conformance program. These are serious commitments and they matter. But they answer a different question than the one the board asked. They measure whether the organization has designed a governance program, staffed it, and run its remediation discipline. They measure program inputs.

AI Posture measures the output. It asks whether the work has become visible enough to be checked from outside. This distinction is not academic. An organization can hold an ISO/IEC 42001 certificate and still score Perceiving on AI Posture, because certification confirms a management system exists, not that the system has yet produced externally observable behavioral change across people, infrastructure, and regulation. Conversely, an organization can score Calibrating on AI Posture without formally adopting any named framework, provided the behavioral evidence is there to inspect.

You can pass ISO/IEC 42001 and still be Perceiving. The certificate proves a program exists. AI Posture asks whether the program has changed observable behavior yet.

Because the two kinds of measurement are orthogonal, they do not compete. A governance management system organizes the work; AI Posture reports whether the work has produced defensible behavior. A risk framework identifies what must be managed; AI Posture reports the maturity of the observable management output. A compliance program feeds one vector; AI Posture prevents compliance maturity from being mistaken for whole-organization readiness. Where these frameworks are required, they remain required. AI Posture does not replace them and does not substitute for them. It sits downstream of all of them and reports what they have actually moved.

Intent is signal. Only behavior scores.

If readiness is an output measure, the evidence standard has to be strict, because AI governance claims are already easy to overstate. A framework that accepted internal sentiment, policy aspiration, or vendor adoption as maturity would reward presentation over operating reality. AI Posture takes the opposite position. Every claim under the framework must resolve to an artifact a third party can inspect. Surveys and self-reported confidence do not qualify. Intent is signal. Only behavior scores.

This is why the artifact classes are specific to each vector. People maturity rests on behavioral assessment with verifiable telemetry, measured in a privacy-preserving way, not on a training-completion percentage or an engagement survey. Infrastructure maturity rests on agent-readiness scans, machine-readable declarations, and structured identifiers across the full stack, from internal systems to partner integrations to public surfaces. Regulation maturity rests on an obligation register with verifiable coverage and recorded interpretations tied to authorities, not on a statement that the organization takes compliance seriously.

The evidence standard has a sharp consequence for misrepresentation. A claim that exceeds reality is not a separate offense to be tracked; it is a signal of immaturity inside whichever vector the claim falsifies. An organization that asserts a People maturity it cannot evidence has, by that act, reduced its People score. The framework folds honesty into the measurement rather than policing it on the side. The strongest version of this rule applies to scope itself: an organization that declares a vector out of scope, then is shown to have material activity there, has not earned a low score. It has produced a claim that cannot be trusted, and that invalidates the entire posture assertion for that stamping.

Three vectors, one number

AI Posture v1.0 ships three vectors, each measuring a distinct actor class that can independently constrain the whole. People measures, inside-out, how effectively humans in the organization collaborate with AI. Infrastructure measures, bottom-up through the outer surface, how ready the organization's digital systems are for agent interaction. Regulation measures, top-down, how completely the organization has met its AI-specific obligations across the jurisdictions that bind it. The outer surface, how the organization presents to humans and agents, is not a fourth vector; it is the edge of the same infrastructure.

Vector	Span	What it measures	Artifact class
People	Inside-out	How effectively humans collaborate with AI	Behavioral assessment with verifiable telemetry
Infrastructure	Bottom-up to edge	Readiness of digital systems for agent interaction	Agent-readiness scans, machine-readable declarations
Regulation	Top-down	Completeness of AI-specific obligations met across jurisdictions	Obligation register with verifiable coverage

Table 1: The three v1.0 vectors of Aggregated Intelligence Posture

Each vector is scored on one shared five-level scale, so a single level name carries equivalent weight wherever it appears. The scale runs from awareness to systematized frontier practice, with a distinct level reserved for vectors that genuinely do not apply.

Level	Name	Meaning
0	N/A	The vector does not apply at this time. Defines the scope boundary; excluded from the aggregate. Itself a falsifiable claim.
1	Perceiving	Aware the domain exists but has not acted.
2	Assessing	Has begun inventorying its state but has no deliberate practice.
3	Integrating	Deliberate practice is in place; evidence is starting to accumulate.
4	Calibrating	Practice is measured, tuned, and defensible to outside inspection.
5	Engineering	Practice is systematized; the organization advances the frontier rather than catching up.

Table 2: The five-level maturity model, shared across all vectors

Each vector reaches each level independently. An organization may be Engineering on Infrastructure while Assessing on People. Defensibility at Calibrating is not audience-specific: a Calibrating score must hold up to auditors, regulators, boards, partners, and customers alike, the way a healthy immune system defends against threats it has not catalogued in advance. The question that remains is how three independent vector levels become one number a board can act on. The answer is the most consequential design choice in the framework.

Bounded by the weakest link

AI Posture equals the minimum of its in-scope vector levels. Vectors marked N/A are excluded from the minimum. The rule is normative, not a scoring convenience, and it is deliberately not an average. Domains depend on each other in practice, and the minimum captures the true operating ceiling of a cross-domain claim: the posture is only as strong as the weakest link that supports it.

An average launders an immature vector behind mature ones. Consider an organization that is Engineering on People and Calibrating on Infrastructure but Perceiving on Regulation. An average reads near Integrating and implies a cross-domain posture the Regulation vector cannot support. The minimum reads Perceiving, which is the truth: the organization cannot make a defensible compliance narrative, no matter how strong its human practice is. The same logic runs in every direction. A Calibrating Regulation score does not prove people collaborate with AI responsibly. Agent-ready infrastructure can amplify poor human practice rather than redeem it. Strength in one vector is a capability, never a substitute for aggregate maturity.

The minimum prevents a mature vector from laundering an immature one. A board that sees Perceiving knows exactly which domain bounds the organization, and therefore exactly where the next dollar goes.

Stress-testing the rule against realistic profiles confirms it holds where an average would mislead. A team strong on people but weak on regulation, a compliance function strong on obligations but weak on adoption, a single business unit scoring high while the enterprise scores lower: in each case the minimum names a constraining vector, and the constraining vector is both the honest ceiling and the investment case. This is the rule's quiet advantage. It does not only score the organization. It tells the practitioner what to fix next, because the vector that bounds the aggregate is the vector that, once advanced, raises the whole posture.

A posture decays. A certificate pretends it does not.

A certificate is a statement about a moment, printed as if it were permanent. AI readiness does not behave that way. Regulation can shift in a fortnight when a new jurisdiction begins to apply. Infrastructure changes with every deployment. Human behavior drifts as tools, tasks, policies,

and incentives change. A readiness measure that ignored time would be lying within weeks of being issued.

AI Posture is therefore a time-stamped assertion, not a guarantee of future state. Every report stamps the scope it covers, the date it was instantiated, and a next-review date reflecting the assessor's belief about how long the assertion is likely to hold. Each vector also carries an at-this-level-since date, because duration at a level is itself a trust signal: two organizations at the same level with different tenures are publishing truthfully different signals, and a careful reader weights them accordingly. Level 5 goes further and requires a declared review cadence plus a review artifact produced within the prior window. The organization sets the cadence; the artifact proves adherence.

Freshness, then, is the reader's job, not a gate the framework enforces. A score three months past its next-review date is not invalid. It is a weaker signal, and it should be treated as one. This is also why externally imposed change does not retroactively erase maturity. A new regulation does not reduce the People practice an organization has already evidenced; it reveals where the next assertion must extend its coverage. The posture is progressive. New exposure changes what must be covered next. It does not delete the behavior already proven.

What AI Posture is not

A framework earns trust as much by what it refuses to claim as by what it asserts. AI Posture is not a certification, and it issues no seal. It is not an audit and collects no independent evidence on the organization's behalf; the organization self-assesses and self-asserts against a defined evidence standard, and the framework defines that standard rather than grading the work. It is not legal advice. It does not replace the named program frameworks, and where those are required they stay required.

The framework is also honest about what remains open. The shared level names are intended to carry equal weight across vectors, but whether a reader interprets Calibrating the same way for people, systems, and regulation is a question still under validation. The minimum-vector rule may be correct yet under-explained, since dashboard conventions condition readers to expect averages; the framework treats persistent confusion as a reason to improve the explanation, with a spec issue reserved for the case where the rule misleads even when understood. The current three-vector set is sufficient for early adoption but deliberately open: new vectors are admitted only when they have an observable artifact, a distinct actor class, independent variation, and the ability to constrain the whole. Stating these openly is not a weakness in the framework. It is the same evidence standard turned inward.

The one number you can defend

The case for AI Posture reduces to a single proposition: a board, a regulator, a partner, and a customer all deserve one readiness level that is true, defensible, and current, and three disconnected dashboards cannot produce it. AI Posture produces it by measuring output behavior rather than program inputs, by holding every claim to an externally verifiable evidence standard, by aggregating to the weakest in-scope vector rather than a flattering average, and by stamping the result as a signal that decays rather than a certificate that pretends not to.

For a GRC leader, the next step is concrete. Stop reconciling three readiness stories by hand. Establish the three vectors as named, independently measured inputs, then report a single Aggregated Intelligence Posture whose constraining vector names your next investment without further argument. A fast self-assessment at aiposture.org/assess/ produces an estimated posture as a starting point; verified per-vector measurement follows. The specification is open under CC BY 4.0, the reference code is MIT, and the framework favors no particular implementation. The number you can defend is available now. The only question is whether your organization is reporting it yet.

About AI Posture

AI Posture is the unified governance framework for Aggregated Intelligence readiness, stewarded by PAICE.work PBC, a US public benefit corporation, with a planned transition to an independent steward. The specification, question bank, and rubric are published under CC BY 4.0; reference code is MIT. It is one measure within a wider PAICE portfolio that keeps Aggregated Intelligence legible, human-owned, and governable as it scales.

Read the spec and self-assess: <https://aiposture.org/>

About the author

Sam Rogers is the founder of PAICE.work PBC. He has spent 25 years working across learning ecosystems, data integration, and AI adoption, and has been developing the concept of Aggregated Intelligence since 2022. He writes at sam-rogers.com.

This white paper is provided for informational purposes only and does not constitute legal, compliance, or professional advice. AI Posture is a maturity-assertion framework, not a certification, audit, or compliance guarantee. Organizations self-assess and self-assert; the framework defines an evidence standard, not a grader. Framework terms are in beta and may change before general availability; the specification at <https://aiposture.org/> is authoritative. © 2026 PAICE.work PBC. Specification licensed CC BY 4.0; reference code MIT.

Appendix A: Adjacent-framework crosswalk

AI Posture is positioned as an output maturity assertion. It does not replace governance frameworks, risk frameworks, compliance programs, certifications, audits, or legal advice. The table below summarizes how it relates to each adjacent class without claiming to substitute for it.

Adjacent class	How AI Posture differs	Useful relationship
Governance management systems (e.g. ISO/IEC 42001)	Measures externally inspectable output behavior, not the existence of a program.	The program organizes the work; AI Posture reports whether the work produced defensible behavior.
Risk frameworks (e.g. NIST AI RMF)	Progressive maturity, not exposure scoring. New risk reveals bounded scope; it does not erase evidenced behavior.	Risk frameworks identify what must be managed; AI Posture reports the maturity of observable output.
Compliance and legal conformance (e.g. EU AI Act)	Regulation is one vector, not the whole. Strong compliance does not prove mature People or Infrastructure.	Compliance programs feed the Regulation vector; AI Posture prevents that from being read as whole-org readiness.
Capability maturity models	Applies one shared level shape across actor classes and aggregates by the minimum, not a process average.	Maturity language aids understanding; AI Posture adds cross-vector constraint logic and time-stamps.
Behavior-change / adoption models	People is one vector; sentiment alone is rejected as evidence in favor of inspectable, privacy-preserving behavior.	Adoption work improves the People vector; AI Posture keeps it tied to infrastructure and regulatory constraints.

Table 3: How AI Posture relates to adjacent frameworks without replacing them

Appendix B: A sample posture report

An AI Posture report states the aggregate level, the declared scope, the stamping and next-review dates, each in-scope vector's level and tenure, the constraining vector, and the recommended next action. Vectors marked N/A are listed explicitly so the scope is legible. The illustrative report below shows the format; it is not a real assessment.

```
Aggregated Intelligence Posture: Assessing
Scope: Acme Corp, organizational
Stamped: 2026-06-01      Next review: 2026-10-01

People:      Calibrating      since 2025-09-01
Infrastructure: Integrating    since 2026-02-14
Regulation:  Assessing        since 2026-04-20

Constraining vector: Regulation
Recommended next action: Advance Regulation to Integrating
```

Figure 1: Illustrative AI Posture report (sample data, not a real assessment)