Aggregated Intelligence

The case for measuring human-machine collaboration independently while it can still be seen

An executive brief

Sam Rogers, Founder of PAICE.work PBC

PAICE.foundation

June 2026 • Companion to “One Number You Can Defend”

Executive summary

The intelligence that determines an organization's outcomes is not its people alone, and not its models alone. It is the two working together, well or badly. 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. Artificial intelligence is one substrate inside that system, not the system itself. The organizations that will compete on it are the ones learning to measure it now, while those intelligences are still legible enough to each other to be measured at all. That legibility is finite. The asymmetry between a bounded human mind and an effectively unbounded machine only grows, and the window in which we can still follow the collaboration narrows as it does.

This brief makes a claim aimed at the leaders and investors deciding where durable advantage will sit. Aggregated Intelligence can be measured, but only a measure independent of the tools it scores can be trusted, and independent measurement of a new category earns an authority that vendor-coupled products cannot. That is the bet. PAICE is building the measurement layer behind it now, across people, infrastructure, and regulation, with AI Posture as one combined score that sits on top. The companion brief, “One Number You Can Defend,” details that instrument; this one makes the case for the approach, and for moving before the window narrows.

The unit of value is aggregated, not artificial

Most AI strategy measures one half of the thing that produces results. Adoption dashboards and training-completion rates measure the human half. Model evaluations, latency budgets, and pipeline telemetry measure the machine half. Both are real, and neither captures the unit that actually generates outcomes: the collaboration between them. A skilled team on top of a capable model can still produce nonsense if the handoffs are wrong, and an unremarkable model in disciplined hands can outperform a stronger one used carelessly. The output is a property of the pairing, not of either side on its own.

Aggregated Intelligence is the name for that pairing treated as a single measurable system. It reframes a familiar picture. We tend to call AI a jagged intelligence, capable in one place and surprisingly weak in another, on the quiet assumption that we are the smooth ones. We are not, and AI is the first instrument that gives us an outside view of our own cognition: the overconfidence, the pattern-matching that skips verification, the comfort with an authoritative answer nobody actually checked. Measuring Aggregated Intelligence means assessing that combined system honestly, rather than grading the human and the machine on separate report cards that never reconcile.

The window is open now, and it is closing

Human intelligence is bounded by the neurons in a skull. The systems built out of the sand that thinks carry no such bound. Reasonable people disagree about how fast and by how much; almost no one disagrees about the direction. For now the asymmetry is still manageable, because we sit close enough to our machines to audit them, to notice when an answer is confidently wrong, and to share enough language to collaborate on purpose rather than by accident.

That proximity is the asset, and it is wasting. As the gap widens, an organization that never built the habit of measuring its human-machine collaboration will find itself directing systems it can no longer follow, with no instinct for the work. The capability cannot be bought at the moment it is needed; it is accumulated through practice, on evidence that is still inspectable, against errors that can still be caught. The cheapest time to baseline is now, and the value of having done so rises precisely as visibility falls. That gap between low cost and rising payoff is the whole reason to move early.

> *We have worked with non-deterministic systems for as long as we have had colleagues. We just called it other people. The discomfort with AI is not its difference from us. It is the recognition.*

What cannot be measured cannot be governed

The recurring failure of the pre-AI era was getting what we asked for instead of what we wanted, and absorbing the gap at human speed. Entire industries were built on the premise that stated demand is a fair proxy for real need, and the premise was wrong often enough that the wreckage is visible from a distance. A fast, tireless system does not remove that gap. It scales it. The competence that matters now is the ability to see the distance between intention and output, and to close it before it compounds.

That competence is measurable, but only against a strict standard. A measure that accepts sentiment, policy aspiration, or tool purchases as evidence rewards presentation over operating reality, which is exactly the gap it was meant to expose. Useful measurement of Aggregated Intelligence has to rest on behavior a third party can inspect: how people actually work with AI, how systems actually respond to agents, how obligations are actually met. Intent is signal. Only behavior counts. It is the discipline that makes a financial statement worth reading, applied to the collaboration itself.

Why the measure has to be independent

A measure owned by the vendor whose tools it scores is not a measure; it is marketing with a number attached. The history of every rating that came to matter, from credit to crash safety to

nutrition labeling, runs the same way: the score earns trust in proportion to the distance between the scorer and the scored. Aggregated Intelligence is no different. For a readiness number to survive a board, a regulator, a partner, or an acquirer, it has to come from a measure with no stake in flattering the result.

This is also where durable advantage sits, which is the part that should interest an investor. Independence is not only an ethic; it is a moat. A neutral measure that arrives first defines the vocabulary the category later argues in, and vocabulary, once it is adopted by regulators, standards bodies, and the people writing the next procurement template, compounds in a way a feature cannot. The measurement is deliberately decoupled from any single implementation and is moving toward an independent steward, so the authority belongs to the measure rather than to a product. And because it spans several surfaces, the calibration data it gathers across them combines into a research asset no single-product competitor can assemble. The independence that makes the number trustworthy is the same property that makes the position defensible.

The approach, already shipping

None of this is a thought experiment. The measurement layer exists today across three actor classes that can each independently constrain the whole. People measures how reliably humans collaborate with AI, drawn from inspectable behavior rather than self-report. Infrastructure measures how ready an organization's digital systems are for agent interaction. Regulation measures how completely its AI-specific obligations are met across the jurisdictions that bind it. Each is a distinct measure with its own reference implementation, and each produces evidence a third party can check.

Actor class	What independent measurement looks at	Reference implementation
People	How reliably humans collaborate with AI, from inspectable behavior	PAICE.work
Infrastructure	How ready digital systems are for agent interaction	Siteline.to
Regulation	How completely AI-specific obligations are met across jurisdictions	EveryAILaw.com
Combined	One readiness score, bounded by the weakest of the three	AIPosture.org

Table 1: The measurement layer beneath an aggregated-intelligence posture

AI Posture sits on top of that layer as one combined score, bounded by the weakest of the three, because an organization is only as ready as its least-ready dimension and an average would launder the gap that actually limits it. It is the first worked example of the approach, and

the companion brief, "*One Number You Can Defend*," sets out exactly how that instrument is built and defended.

Measurement is the visible top of a wider approach. The same thesis, keeping Aggregated Intelligence legible, human-owned, and governable as it accelerates, runs through the rest of the portfolio. Each component holds one property of the human-machine system true at agentic speed and scale. They are not features of the score; they are the surrounding infrastructure that makes a defensible score possible in the first place.

What it keeps true	PAICE component
Coordination between agents stays human-legible and human-owned	Turnfile
Trust boundaries between humans and agents are declared and enforced	GuideCheck
Agents refuse and disengage gracefully at the limits of their remit	Graceful Boundaries
Agent skills carry verifiable version and integrity across hands and machines	Skill Provenance
Organizational knowledge stays machine-readable without ceasing to be human-authored	Knowledge-as-Code
What agents build stays usable by every human, not only the median one	A11y Audit
Identifiers machines emit stay safe for humans to read and transcribe	HardGuard25
Fast-moving capability information stays tied to the human questions that need it	AI Tool Watch
The legal landscape stays navigable at agentic speed and scale	Obligation First
The public record of AI harms and rulings stays citable and current	AI Incident Law
Civic records stay open and verifiable as recordkeeping is automated	PubLedge

Table 2: The wider PAICE portfolio, framed by what each keeps true as Aggregated Intelligence scales. The set is extensible; new components are admitted as gaps appear.

This breadth is not incidental, and it is the second half of the investor case. Each surface contributes calibration data the others cannot, which is what turns a set of open standards into a

research position a single-product competitor cannot copy without first shipping the same breadth. The discipline this brief argues for is real, it is shipping, and the measurement at its center is already producing numbers organizations can put their name to.

Use this time

The argument reduces to a sequence. Aggregated Intelligence, not Artificial Intelligence alone, is the unit that determines outcomes. It can be measured, but only against inspectable behavior, and only a measure independent of the tools it scores will be trusted or will last. The window in which the collaboration is still legible enough to baseline is open now and narrowing. The organizations and the backers that establish an independent measure while it is open will hold the vocabulary, the data, and the muscle memory when the rest can no longer see well enough to follow.

The specifications are open, the reference code is permissively licensed, and the measurement favors no particular vendor, because charging rent on the standard would waste time the situation does not allow. Read the approach, measure once against it, and watch which dimension bounds you. If you can build a better measure, build it and say so. But do not wait for the window to make the decision for you. Use this time. It is the part of the work that does not come back.

About this work

PAICE builds independent measurement for Aggregated Intelligence: the collective output of different intelligences working together, today biological and artificial. The measurement layer spans people, infrastructure, and regulation. AI Posture combines it into one score bounded by the weakest. Sam Rogers has used the term Aggregated Intelligence in his writing since 2022, and founded PAICE.work PBC (in 2025) as its first formal steward. Specifications are published under CC BY 4.0 and reference code under MIT, maintained by PAICE.work PBC, a US public benefit corporation, with a planned transition to an independent steward.

Start here: https://aiposture.org/