

Not Another Taxi

Every AI governance vendor on your shortlist is selling you a fleet. The question is what happens the first time your concern set changes — and it will change this year.

CAUSUM AIGP + MARS THESIS

You are not buying governance. **You are buying a fleet.**

The vendor you are evaluating covers four concern areas well. Fairness, privacy, model risk, and whichever regulation was live when they raised their Series A. Their evidence format is their own. Their evaluation logic is their own. Their experts are on their payroll.

Now name the fifth concern — the one your regulator adds next year, or the one a customer's procurement team asks about in the sales call you have not had yet. Cognitive harm. Supply-chain integrity. A jurisdiction you do not operate in today.

Who do you call? Not them: covering it means they hire an expert and write new logic, on their roadmap, at their pace. So you call another vendor, whose evidence does not compose with the first. And another.

This is the position every enterprise ends up in, and none of them chose it. They just kept calling taxis.

THESIS

Causum is not another taxi

Every generation of enterprise software eventually spawns a “governance solution” that is just another bespoke SaaS product — another taxi company with its own fleet, routes, and medallion system. AI governance is no different: the market is filling with point solutions, each addressing a single concern, for a single framework, in a single vendor’s proprietary format.

Causum is the ride-sharing platform. It is the infrastructure connecting those who need AI governance (**Emitters**) with those who hold specialised expertise in particular concern areas (**Receivers**). It does not compete with governance frameworks. It makes them composable, distributable, and economically viable at scale.

A platform is only as good as the currency that crosses it. Causum’s currency is governance evidence — and evidence is worth trading only if it can be trusted. That trust is not assumed here. It is manufactured, by two protocols working as one.

AIGP

What was governed

Who asked, under what authority, in what jurisdiction, with what consent. The evidence that a request was **permitted**.

MARS

Whether the answer was sound

Reasoning governed against a formal model of the domain, grounded in real evidence, certified with authority. The evidence that a decision was **sound**.

Neither half is the whole story alone. Together they form one unbroken record, and it is that combined record any party can independently re-derive. The evidence is verifiable by construction — the one thing a marketplace of governance cannot function without.

The taxi company owns the cars, the routes and the drivers. The platform owns the matching function — and that is what scales.

Every bespoke governance tool is a taxi company

Consider what a typical AI governance SaaS product actually does.

WHAT THE PRODUCT DOES	WHAT THE MARKET BECOMES
Builds its own everything ▪ Defines its own model of a concern area ▪ Writes its own evaluation logic against it ▪ Stores results in its own evidence format ▪ Employs the experts the logic encodes ▪ Covers the frameworks asked for first ▪ Adds coverage by hiring — on its roadmap	A landscape of closed fleets ▪ Dozens of vendors, each strong in a narrow band ▪ No two evidence formats that compose ▪ Expertise locked inside whoever hired it ▪ Several tools that cannot see each other ▪ Every new concern a new procurement ▪ No basis for comparing organisations

This works for exactly one customer: the one whose needs happen to match the taxi company's routes. For everyone else, you call another taxi. And another.

Captive supply. Limited geography. No interoperability. No network effects.

THE INSIGHT

Uber did not build a better taxi

It built something structurally different — and the difference was not quality of service. It was where the value sat.

TAXI COMPANY	RIDE-SHARING PLATFORM
Owns the vehicles	Connects vehicle owners to riders
Employs the drivers	Enables independent operators
Fixed routes and zones	Matches dynamically by need
Revenue from fares	Revenue from facilitation
Quality by hiring standards	Quality by ratings and evidence
Scales by buying more cars	Scales by attracting more participants

The insight was never that taxis are bad. It was that the matching function is worth more than the fleet.


THE APPLICATION

Causum applies the same structure

BESPOKE GOVERNANCE SAAS	CAUSUM PLATFORM
Builds its own concern model	Hosts any concern model as a Dialect (RFC-038)
Owns the evaluation logic	Enables specialised Receivers to provide evaluation
Limited to its concern areas	Open to any Domain of Concern (RFC-034)
Proprietary evidence format	Open evidence schema — interoperable by design
One vendor’s geography	Jurisdictional governance across any regime
Scales by hiring more engineers	Scales by attracting more concern specialists

THE TWO SIDES

Riders and drivers



Emitters

THE RIDERS

Organisations that deploy, operate or consume AI systems and need those systems governed.

- Enterprise AI teams deploying agents into production
- SaaS vendors with AI features their enterprise customers demand governance over
- Regulated industries — healthcare, finance, defence — with mandatory compliance obligations
- Government agencies accountable for AI decisions affecting citizens

Emitters do not want to become governance experts. They want to emit evidence and have it consumed, evaluated and verdicted by parties who specialise. They want governance the way a rider wants a ride: on demand, by qualified parties, without owning the car.

Receivers

THE DRIVERS

Parties with deep, specialised knowledge in specific governance areas.

- Fairness and bias specialists who measure representational and allocative harm
- Privacy engineers who understand differential privacy, data minimisation and consent models
- Safety researchers who evaluate autonomous system boundary conditions
- Regulatory experts who know the EU AI Act, Japan AI Promotion Act or AU Continental Strategy in depth
- Domain specialists — medical AI evaluators, financial risk assessors, military ethics reviewers
- Cybersecurity analysts evaluating adversarial robustness and model supply-chain integrity
- Cognitive harm researchers assessing dependency patterns and human deskilling

Receivers do not want to build a governance platform. They want to package their expertise as a consumable, subscribable, versioned artifact — a Dialect — and distribute it to the Emitters who need it.

WHY THIS IS NOT A METAPHOR

The structure is the product

The two-sided economy is not a marketing analogy. It is how the thing is built.

The protocol is the matching function. AIGP's core protocol — REGISTER, CHECK, RECORD, TRACE, ANTICIPATE, VERIFY — is the shared substrate connecting Emitters to Receivers. It is not the governance itself; it is the interface that makes governance composable. An Emitter emits evidence in a standard format. A Receiver consumes it through a standard interface. The protocol matches them, not by picking winners, but by defining the interaction contract.

And because the evidence carries the full governed record — AIGP's account of what was governed and Mars's account of whether it was sound — it is not merely interoperable. It is trustworthy across the seam between two organisations that have no reason to trust each other. A Receiver in one country can act on an Emitter's evidence from another without taking anything on faith.

Interoperability is what lets the evidence move. A checkable record is what makes it worth moving.

Dialects are the supply. A Dialect is the distributable, versionable, subscribable expertise package a Receiver publishes — their vehicle, in this analogy.

DIALECT (RFC-038)

Domain of Concern	what class of harm
+ Mediation Observation Model	what to measure
+ Calculation Semantics	how to calculate posture
+ Observer Requirements	who is qualified to verdict
+ Default Thresholds	what levels trigger action

They publish it to the Dialect Registry — the marketplace — and any Emitter can subscribe.

Evidence is the currency. In the taxi model the currency is the fare. Here it is the stream of structured, signed records flowing from Emitter to Receiver: the Emitter produces evidence, the Receiver evaluates it against their Dialect's measurement apparatus, a verdict flows back, and both parties hold an auditable record of the exchange.

But a fare is only worth taking because you can count it. A signature proves a record was not *altered*. It proves nothing about whether the decision behind it was *sound*. Currency that can only be signed, not verified, is a marketplace of assertions.

HALF ONE · AIGP

What was governed

Before an AI system produces an answer, AIGP settles the record: who is asking, under whose authority, in what jurisdiction, with what consent. Consent is a dial, not a checkbox — an Emitter can satisfy it with minimum disclosure, proving what a decision requires without exposing the content behind it.

This is the evidence that a request was **permitted**.

HALF TWO · MARS

Whether the answer was sound

Mars is the governed reasoning system. It gives a domain a formal model — multi-order, multi-perspective, precise — and every answer is reasoned against that model, grounded in gathered evidence, conformance-checked, and certified with declared authority. Where evidence runs out it does not guess: it opens a typed gap.

This is the evidence that a decision was **sound**.

Together they make the currency bankable. The two records join into one chain, and it is that whole chain — not either half — a Receiver can independently re-derive. None of it has to be taken on faith.

That is the difference between a receipt you have to believe and a fare anyone can count.

The platform enables but does not constrain. Causum does not tell Receivers what concern models to build, or Emitters which Dialects to subscribe to. It provides discovery, subscription, distribution to all governed systems, deliberate versioning, and compatibility so different Dialects coexist without conflict. That is the platform function: enabling a market, not owning the supply.

WHAT THIS MAKES POSSIBLE

Four things a point solution cannot do

01

The multinational enterprise

A company deploying AI across the EU, Japan and Australia does not need three governance vendors. They subscribe to three Dialects, all consuming the same evidence stream from the same emitting systems. No duplication, no incompatibility, no bespoke integration.

- `eu_ai_act@3.1.0` --- European regulatory specialists
- `japan_ai_promotion@1.4.0` --- Japanese compliance experts
- `au_continental_strategy@2.0.0` --- AU governance analysts

02

The specialised concern expert

A research group working on cognitive harm — AI-driven deskilling, dependency patterns — packages their measurement apparatus as a complete Dialect with variable definitions, calculation semantics and observer requirements.

■ `cognitive_harm@1.0.0` --- published to the registry

Any organisation worried about cognitive atrophy can subscribe — without the research group building a SaaS platform, and without the enterprise hiring cognitive harm researchers. The expert's reach exceeds their consulting capacity for the first time.

03

The defence organisation

A military AI programme needs governance no commercial vendor offers: autonomous systems under International Humanitarian Law. Wait years for a vendor, build it in-house at expense and without portability, or:

■ `ihl_governance@2.3.0` --- maintained by IHL-specialist Receivers

The Dialect bundles the measurement apparatus. The platform distributes it. The emitting systems produce evidence against it.

04

The framework author

NIST publishes the AI RMF. Today every vendor interprets it independently and incompatibly, so no two organisations' compliance claims mean the same thing.

■ `nist_ai_rmf@2.0.0` --- published by NIST or an authorised party

Every subscriber uses the same measurement apparatus: same variables, same calculation semantics, same thresholds. Cross-organisational comparison becomes meaningful. The framework author becomes a Receiver — packaging their expertise for consumption rather than hoping vendors interpret it correctly.

WHY BESPOKE SAAS CANNOT FOLLOW

The taxi company cannot become a platform by adding cars

These are not execution gaps a well-funded competitor closes. They are structural consequences of owning the fleet.

THE LIMITATION

WHY IT IS STRUCTURAL

It owns the expertise

A SaaS tool embeds one team's understanding of one concern area. Covering a new concern means hiring new experts and writing new logic. That is linear scaling — more cars, not more participants.

It owns the format

If your evidence only works inside one tool, you are locked to that tool. The taxi determines your route, and switching means re-instrumenting everything already governed.

It cannot distribute what it does not own

A vendor cannot resell expertise held by an independent specialist without acquiring or employing them. It is a closed fleet, not an open marketplace, and its coverage is bounded by its payroll.

Its outputs do not compose

Combining two vendors' governance requires custom integration built and maintained by the customer. Two taxi companies cannot share a ride.


It cannot create comparability

Quality is asserted by the vendor's hiring standards rather than demonstrated by evidence anyone can re-derive. Nothing lets one organisation's posture be compared with another's.

The platform model does not compete with taxi companies. **It makes them unnecessary** — because the expertise once trapped inside vendor organisations becomes distributable, composable and available to anyone by subscription.

THE NETWORK EFFECT

Platforms do not add features. They create markets.



More Receivers create better options for Emitters. More Emitters create demand for Receiver expertise. More expertise produces more Dialects. More Dialects expand coverage. Broader coverage drives adoption. Once adoption is widespread, comparison becomes meaningful, regulators gain a basis for endorsement, and participation can move from voluntary to mandatory.

A point solution can fill a slot in an enterprise stack. It cannot create the market dynamics that make universal governance possible — and it cannot join a market it is structurally unable to interoperate with.

THE ECONOMIC MODEL

Subsidise demand, monetise supply

EMITTER SIDE	RECEIVER SIDE
Free to emit Implementing the protocol and emitting governance evidence is free. The specification is open. The SDKs are source-available. The barrier to entry is two JSON files and fifteen minutes. You do not charge the rider to get in the car.	Commercial intelligence ▪ Dialect creation tools ▪ Scoring and posture engines ▪ Certification programmes ▪ Registry hosting

The Uber driver pays a commission, not the rider. Causum earns from facilitation, not from owning the governance logic — which is also why it cannot be accused of grading its own homework.

THE COMMERCIAL MODEL

Open at the point of adoption

Causum’s commercial model is built the way the platform is: open where you start, out of the way as you grow, and priced so that using it is never the reason not to.

Open and unobstructive. The specification is open. The SDKs are source-available. Emitting evidence is free and stays free — there is no gate between an organisation and getting governed. Where Causum is licensed, pricing is volumetric, negotiable, and guaranteed never to rise. No per-seat traps, no lock-in formats, no penalty for growth. The model earns from value crossing the platform, not from standing in front of the door.

Licensing the protocols. AIGP and Mars are licensed together and each stands alone. An organisation can adopt one first, pair them, or slot either beside a component of its own — the boundary between them is a published contract, not a proprietary tie. A licensee receives the conformance suites and reference implementations that turn “are we compliant?” from an argument into a test that is run.

Enablement. Causum does not hand over a specification and walk away. It trains an organisation’s people into certified practitioners and stands beside them through design and deployment, with consulting from the teams who built the protocols. The goal is an operator who owns their own governance, not a dependency on a vendor.

Expertise is transferred, not rented. That is the same principle that governs the platform, applied to the commercial relationship: **enable the customer, do not capture them.**

BOUNDARIES

What Causum is not

NOT THIS	BUT THIS
Another governance SaaS tool	The protocol layer those tools would emit into — and which makes any of them replaceable without re-instrumenting the systems they govern
A competitor to NIST, ISO, the EU AI Act	The substrate that makes those frameworks distributable
A replacement for domain experts	The marketplace connecting domain experts to demand
A single concern model	A registry hosting any concern model
A proprietary lock-in	An open protocol with commercial intelligence on top

THE TITLE, EXPLAINED

Four reasons it is not another taxi

01

The world does not need another point solution

It needs the infrastructure that makes governance expertise distributable, composable and economically viable.

02

Taxis solved the wrong problem

They optimised vehicle ownership; ride-sharing optimised the matching function. Bespoke governance tools optimise their own logic. Causum optimises the connection between governance needs and governance expertise.

03

Platform economics create value products cannot

Network effects, cross-organisational comparability, specialist participation without full-stack investment. These emerge from platforms, not products.

04

The expertise already exists

The world is full of brilliant fairness researchers, privacy engineers, safety specialists and regulatory experts. They are trapped inside consulting firms, research labs and vendor organisations, reaching a handful of clients a year. Causum gives them a distribution channel that scales beyond their headcount.

The matching function is worth more than the fleet.

Every vendor in this market is still buying cars.

EDITORIAL NOTE — RECONSTRUCTED CONTENT

The source document lost several tables in format conversion, and its seven referenced images were not present in the folder. The following were rebuilt from what the surrounding argument requires and should be reviewed against your originals:

- **“Every bespoke governance tool is a taxi company”** — both lists (what a typical SaaS product does; what the market becomes) were empty in the source.
- **“Why bespoke SaaS cannot follow”** — the source table’s left column had been overwritten with labels from the Uber comparison, leaving rows that did not match their explanations. Rewritten as five structural limitations.
- **“What Causum is not”** — row one’s right cell duplicated the SaaS-limitation text rather than stating what Causum is. Replaced.
- **Diagrams** — the two-sided market and network-effect figures are original, built to carry arguments the source made in prose. The five remaining image slots are not reproduced.

All other content is the source text, reordered and edited for register.