

THE POLITICAL ECONOMY OF INTELLIGENCE

# From *AI Intent* to *AI Reality*

Why the economics of intelligence requires an  
independent decision clearinghouse



---

*Causum is the independent clearinghouse through which  
sophisticated machine and human decision capability  
becomes adjudicable.*

CAUSUM  
POSITION PAPER

# An argument in six parts

## I

---

### Beyond the Model

- 01 The Market Has Moved Beyond the Model
- 02 From Consuming Intelligence to Delegating It

## II

---

### Delegated Cognition

- 03 Delegated Cognition Changes the Governance Problem
- 04 The Intent–Reality Gap
- 05 Evidence Is Not Judgment

## III

---

### The Adjudication Layer

- 06 The Missing Layer: Adjudication
- 07 Making Machine Activity Adjudicable
- 08 Conform, Diverge, Refer
- 09 From Adjudication to Assurance

## IV

---

### Outcomes, Markets and Autonomy

- 10 When an Outcome Becomes an Economic Claim
- 11 Where Causum Sits
- 12 A Clearing Layer for Delegated Intelligence
- 13 Justified Autonomy

## V

---

### The Decision Clearinghouse

- 14 The Independent Decision Clearinghouse
- 15 Two Routes to the Economics of Intelligence
- 16 The KPG: A Common Decision Substrate
- 17 The Clearinghouse Network Effect

## VI

---

### Evidence and Market Entry

- 18 Market Evidence: United States and Europe
- 19 Capability Inequality and the First Market
- 20 From Expert Systems to Cleared Judgment

---

Conclusion · Position Summary · Editorial Note on Evidence

## THE POSITION

Artificial intelligence has stopped merely producing information. It now performs consequential work. Models have become agents; agents are wired into enterprise systems, tools, data, workflows, and one another. Around them, organizations are building *enterprise AI harnesses* — the runtime, context, permissions, controls, and observability that turn model capability into reliable machine work.

The harness answers how to operationalize machine intelligence. It does not answer the question that matters most once cognition and authority are delegated to machines:


---

*Did what happened in AI Reality remain justifiable against what was legitimately intended?*

---

Causum exists for that boundary.

### CAUSUM IS NOT



Causum is the protocol, evidence, and knowledge infrastructure that reconciles consequential machine activity with legitimate intent — and presents it for determination by an authorized human or institutional authority.

### IN SHORT

Causum makes machine activity  
*adjudicable.*


# Beyond the Model

01

## The Market Has Moved Beyond the Model

The first phase of generative AI was a race for model capability: parameters, benchmarks, context windows, latency, inference cost. That race continues, but the enterprise problem has outgrown it. A capable model is now one component of a system expected to do useful work.

The real operational unit looks like this:



~~What can the model do?~~



*What can the enterprise safely and reliably allow machine intelligence to do?*

Fig. 1 — The operational unit of enterprise AI, and the question it changes.

Call it an AI harness, agent harness, control plane, or agent runtime — the terminology varies, the direction does not. Enterprise value depends less on what a model can infer than on whether machine intelligence can be given the right context, constrained authority, reliable tools, observable execution, and clear operational boundaries.

*Causum accepts that evolution as necessary. Its position begins where that question ends.*

## From Consuming Intelligence to Delegating It

The economics of AI are shifting just as fast. Tokens matter because they measure computational demand. But no enterprise buys tokens for their own sake. It buys intelligence because it expects work completed, decisions improved, processes changed, outcomes delivered. The value chain runs:




Fig. 2 — The economic progression of intelligence, from computational demand to realized value.

As agents take on longer-running work, the question changes from *how much inference did we buy?* to *what work did that intelligence complete, what did it cost, and what value resulted?* That is the beginning of an economics of intelligence.

Causum goes one step further. In agentic systems, intelligence is not merely consumed — it is *delegated*. Machines receive objectives, interpret context, select tools, make intermediate judgments, initiate actions, and deal with other machines and people.

The moment intelligence is delegated, economics collides with authority, responsibility, legitimacy, and institutional trust. The question becomes as political-economic as it is computational:

- |     |                                                                                                              |    |                                                    |
|-----|--------------------------------------------------------------------------------------------------------------|----|----------------------------------------------------|
| i   | <i>Who may delegate cognition?</i>                                                                           | ii | <i>What authority accompanies that delegation?</i> |
| iii | <i>What limits apply?</i>                                                                                    | iv | <i>Who bears the consequences?</i>                 |
| v   | <i>Who determines whether an outcome legitimately satisfies the purpose for which authority was granted?</i> | vi | <i>Who can rely upon that determination?</i>       |

### DEFINITION

The **Political Economy of Intelligence**: not simply the price of machine cognition, but the production, allocation, delegation, control, consequences, and value of machine intelligence within institutions and markets.

# II

## PART TWO

# Delegated Cognition

03

## Delegated Cognition Changes the Governance Problem

Traditional software executes instructions. Agentic AI interprets objectives.

That is a new kind of delegation. The human or institution no longer specifies every step. It supplies an objective, context, authority, tools, constraints, and an expected outcome — and the machine decides parts of the path in between. We call this *delegated cognition*.

Delegation does not dissolve responsibility. Nor does technical authorization confer substantive legitimacy. An agent can be properly authenticated, permitted to call a tool, compliant with every runtime control, fully observable — and still produce a consequence that departs from the intent under which it was authorized. Three questions must be kept apart:

### AUTHORIZATION

Was the machine *permitted* to act?

### EXECUTION EVIDENCE

What did the machine *actually do*?

### ADJUDICATION

Was what it did *justifiable* against legitimate intent and governing knowledge?

Fig. 3 — Three related questions. They are not interchangeable.

Today's stack answers the first two well. The third is the one that determines whether a result can be relied upon.

## The Intent–Reality Gap

Public policy has always distinguished policy as declared from policy as realized. Intent passes through institutions, interpretation, implementation, discretion, and local conditions before it becomes anyone's lived outcome.

Agentic AI recreates that implementation problem — at machine speed.




Fig. 4 — Policy, controls, authorization and telemetry each govern a step. None independently proves the consequence stayed consistent with the intent that originated it.

An enterprise can define policy, set controls, authorize an agent, monitor execution, and keep exhaustive telemetry. All of it is valuable. None of it, on its own, proves that what actually happened stayed consistent with the legitimate intent that set it in motion. This is the **Intent–Reality Gap**.

The gap is not a defect. Interpretation is inherent in complex systems — especially when objectives are qualitative, context-dependent, or governed by several sources of authority. The defect is the absence of an independent mechanism to reconcile what was intended with what became real. And it grows more costly with every step machines take from recommendation toward execution.

---

*Governance at the beginning of a process does not automatically prove governance at its consequence.*

---

## Evidence Is Not Judgment

Mature harnesses make machine activity observable. That is essential: there is no reliable adjudication without reliable evidence. But it is not sufficient.

#### AN EXECUTION RECORD MAY ESTABLISH

- which principal or machine initiated activity
- what authority was presented
- which policies and controls were applied
- what context or evidence was available
- which tools or services were invoked
- what actions occurred, and when
- what resulting state or outcome was observed

Answers: *what happened*.

#### IT DOES NOT ADJUDICATE

A **flight recorder** does not adjudicate an aviation incident.

A **ledger** does not adjudicate a commercial dispute.

**Telemetry** does not determine whether an outcome satisfies a contractual, regulatory, professional, or institutional obligation.

Asks: *ought it count?*

Fig. 5 — Evidence makes a matter inspectable; it does not make it adjudicable.

Evidence tells you *what happened*. It cannot tell you whether what happened ought to count as a legitimate realization of intent. That distinction is the foundation of Causum's position.

---

Evidence makes a matter  
*inspectable*.

Causum is designed to make  
the matter *adjudicable*.

# III

## PART THREE

# The Adjudication Layer

06

## The Missing Layer: Adjudication

A layered market structure is forming around enterprise AI. Every layer has an owner — except one: the reconciliation of legitimate intent with evidenced machine reality. That layer is adjudication, and it is where Causum sits.

|                                |                                                                                                                               |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <b>Models</b>                  | provide machine reasoning capability.                                                                                         |
| <b>Agents</b>                  | organize capability around objectives and action.                                                                             |
| <b>Enterprise AI harnesses</b> | provide context, tools, runtime control, identity, policy enforcement, observability, and operational reliability.            |
| <b>Governance systems</b>      | establish policies, responsibilities, risk boundaries, and permitted behavior.                                                |
| <b>Evidence systems</b>        | establish what occurred.                                                                                                      |
| <b>Causum</b>                  | <i>The reconciliation of legitimate intent with evidenced machine reality.</i>                                                |
| <b>Legitimate authority</b>    | enterprise authority, domain expert, assurance provider, regulator, contractual authority, or another authorized institution. |

Fig. 6 — The emerging enterprise AI stack, and the layer it is missing.

Causum does not presume that one machine should be the final judge of another. The legitimate adjudicator may be an enterprise authority, a domain expert, an assurance provider, a regulator, or a contractual authority.

Causum supplies the common substrate on which a matter is constituted, evidenced, evaluated, and — when necessary — referred to that authority. The underlying relationship:



Fig. 7 — The relationship Causum makes explicit and machine-scale.

## 07 Making Machine Activity Adjudicable

Causum is a protocol and knowledge architecture for representing every relationship needed to examine consequential machine activity. It rests on four complementary capabilities.

### AIAP

#### AUTHORITY & AGENCY

Who may act, on whose authority, for what purpose, within what scope, and under what temporal or contextual limits.

*Agency is earned, bounded, and revocable — not assumed.*

### AIGP

#### ADMISSIBILITY

Whether a proposed invocation is admissible within authority, policy, consent, posture, and governing conditions — before consequential execution proceeds.

*Technically capable is not legitimately entitled.*

### KPG

#### GOVERNING KNOWLEDGE

A Knowledge Property Graph of domain-specific assertions, properties, provenance, and relationships relevant to determination.

*Explicit domain knowledge, not implicit model interpretation.*

### MARS

#### EVIDENCE & CONFORMANCE

Evaluates outputs or actions against domain specifications, producing evidence-bearing conformance or recomputation witnesses.

*Not universal truth — an inspectable basis for determination.*

Fig. 8 — The four capabilities of the Causum architecture.

**AIAP — Authority and Agency.** AIAP represents the authority under which a machine acts: who may act, on whose behalf, for what purpose, within what scope, and under what temporal or contextual limits. Its principle is simple — agency is earned, bounded, and revocable, never assumed.

**AIGP — Admissibility.** AIGP decides whether a proposed invocation is admissible under the applicable authority, policy, consent, posture, and conditions — before consequential execution begins. A system can be technically capable of an action without being legitimately entitled to take it.

**KPG — Governing Knowledge.** A Knowledge Property Graph holds the structured, domain-specific assertions against which machine activity is interpreted — not a loose graph of context, but governing knowledge with properties, provenance, and relationships. Adjudication then rests on explicit domain knowledge, not on a model's implicit interpretation.

**MARS — Evidence and Conformance.** MARS evaluates machine outputs and actions against defined domain specifications, producing evidence-bearing conformance or recomputation witnesses. It does not claim universal truth. It makes the basis for a determination inspectable, reproducible where appropriate, and available to an authorized adjudicator.

Together they make possible a bounded form of *justified autonomy*: autonomy whose authority is explicit, whose exercise is observable, whose evidence persists, and whose legitimacy can always be challenged.

08

## Conform, Diverge, Refer

Not every machine action can — or should — be forced into a binary automated verdict. Causum recognizes three determination states, and treats referral as a first-class outcome:

### CONFORM

The evidenced reality is sufficiently consistent with applicable intent, authority, and governing knowledge.

### DIVERGE

Material evidence indicates that reality departs from the applicable intent, authority, or governing knowledge.

### REFER

The matter cannot legitimately be resolved at the current machine or institutional layer and requires an appropriate authority.

Fig. 9 — Three determination states. Referral is a first-class outcome, not a failure mode.

**REFER is the critical one.** A system that forces every ambiguity into an automated yes or no has not solved the authority problem — it has hidden it inside another algorithm. Responsible autonomy begins with knowing when machine determination is not enough.

Causum does not try to eliminate human or institutional judgment. It creates the conditions for applying that judgment efficiently, consistently, and on the strength of evidence.

09

# From Adjudication to Assurance

Adjudication is also the bridge to assurance. Once AI does economically consequential work, internal confidence is no longer enough. Customers, counterparties, regulators, insurers, boards, and investors need a basis for relying on claims about machine activity and its outcomes. The chain that results:




Fig. 10 — From AI Intent to reliance. Separation of functions is deliberate.

No single organization needs to perform every link — and by design, none should. The result is a neutral clearing function that never requires Causum to become the universal authority.

# IV Outcomes, Markets and Autonomy

10

## When an Outcome Becomes an Economic Claim

As AI pricing moves toward completed work and outcomes, the line between execution and adjudication stops being philosophical. It becomes commercial.

If a provider is paid for an outcome, that outcome is no longer a technical result. It is an *economic claim*. And when an enterprise relies on an agent's work to authorize a transaction, change a customer's state, satisfy an obligation, or trigger another machine process, the result carries consequences far beyond the runtime that produced it.

### SIX QUESTIONS EVERY CONSEQUENTIAL OUTCOME MUST ANSWER

|                                                          |                                                          |                                                     |
|----------------------------------------------------------|----------------------------------------------------------|-----------------------------------------------------|
| 01<br>Did the claimed outcome actually occur?            | 02<br>What machine and human activity contributed to it? | 03<br>Under what authority did that activity occur? |
| 04<br>Against which definition is success being claimed? | 05<br>Did the result conform to legitimate intent?       | 06<br>Who is authorized to make that determination? |

### THE ANSWERS DECIDE

payment

reliance

assurance

escalation

accountability

further delegation

Fig. 11 — These are not observability questions. They determine what an outcome can be used for.


None of these is an observability question. Together they decide whether an outcome can support payment, reliance, assurance, escalation, accountability — or further delegation.

*As machine cognition becomes cheap and abundant, trusted judgment about machine-produced reality becomes scarce. That scarcity creates a new institutional — and economic — function.*

# Where Causum Sits

Causum is complementary to the AI systems enterprises are already building. It replaces none of them.

CAUSUM DOES NOT REPLACE



Instead, Causum connects these capabilities around a problem that only becomes visible once AI exercises consequential, delegated cognition: *the relationship between AI Intent and AI Reality*.

THE MARKET POSITION IN SIX LINES

|                        |                                                           |
|------------------------|-----------------------------------------------------------|
| Governance             | <i>defines intent.</i>                                    |
| Authority              | <i>bounds agency.</i>                                     |
| The harness            | <i>operationalizes intelligence.</i>                      |
| Evidence               | <i>establishes machine reality.</i>                       |
| Causum                 | <i>makes the Intent–Reality relationship adjudicable.</i> |
| Legitimate authorities | <i>determine whether reliance is justified.</i>           |

Fig. 12 — Each function stays where it belongs. Causum occupies the one no one else holds.

## A Clearing Layer for Delegated Intelligence

The long-term implication is larger than any single deployment. Machine intelligence will be produced by heterogeneous models, vendors, clouds, enterprises, consultancies, internal teams, and autonomous systems. No single provider can credibly serve as the final authority over every consequential outcome that ecosystem produces.

A scalable AI economy needs a way for heterogeneous machine activity to be understood across institutional boundaries. Causum provides the neutral protocol and evidentiary substrate between those who *produce* consequential machine activity and those legitimately authorized to *adjudicate or assure* it.



Fig. 13 — The clearing function: Causum stands between those who act and those who judge.

This is a clearing function — not because Causum declares itself the judge, but because it makes heterogeneous machine activity capable of being presented to the right one.

## Justified Autonomy

The destination is neither unrestricted autonomy nor permanent human approval of every machine action. Both extremes squander what machine intelligence can do.

Causum advances **justified autonomy**: authority that expands when evidence supports it, contracts when conditions change, expires when its basis no longer holds, and refers matters when machine determination is insufficient.




Fig. 14 — Autonomy as an evidence-bearing institutional state, not a binary feature.

Autonomy stops being a binary feature or a system designer's declaration. It becomes an *evidence-bearing institutional state*.

That changes the relationship between governance and innovation. Governance is no longer a brake applied to machine capability. Properly designed, it is the mechanism by which greater autonomy is responsibly earned.

---

Speed is made possible *by governance*.

---

## The Independent Decision Clearinghouse

Causum's strategic role extends beyond adjudicating the output of enterprise AI. It can operate as independent market infrastructure: the place where evidence-bearing results and complex decision matters are constituted, routed, adjudicated, and returned as determinations that authorized parties can use.

*The unit of analysis changes. It is not the model, the agent, or the harness. It is the matter requiring determination.*

One organization arrives with an AI-produced result and its supporting evidence. Another arrives with no meaningful AI capability at all — only a sophisticated decision that exceeds the expertise it can economically maintain. Both enter the same clearing process.



Fig. 15 — Three economic actors, one common infrastructure.

Causum need not own the underlying AI, nor make every judgment itself. Its institutional role is to let heterogeneous machine and human work enter a common adjudication process.

15

## Two Routes to the Economics of Intelligence

There will be no single operating model for extracting value from machine intelligence.

Large, AI-intensive enterprises will internalize the capability. An enterprise AI harness connects models, agents, context, knowledge, tools, controls, runtime evidence, and enterprise systems into a repeatable production environment. At scale, these organizations move beyond token economics toward *AI Value Economics* — measuring intelligence consumed against machine work performed, human intervention, risk, realized outcomes, and business value.

But not every organization needs — or can justify — owning that infrastructure. Causum opens a second route: consume sophisticated decision capability through the clearinghouse instead of internalizing the entire intelligence-production stack.

### ROUTE ONE

## Enterprise Harness

**Internalizes** sophisticated machine decision capability. Fits large, AI-intensive enterprises that can own the production stack.

### ROUTE TWO

## Causum Clearinghouse

**Externalizes access** to sophisticated, evidence-bearing, adjudicated decision capability — for any organization.

### THREE WAYS INTO THE CLEARINGHOUSE




Fig. 16 — Complementary market structures, not substitutes.

The routes are complementary, not rivals. The same clearing infrastructure serves the enterprise with a mature harness, the AI-enabled firm working with frontier models, and the organization with no AI capability at all.

---

*Organizations need not own the means of producing sophisticated machine intelligence to participate in the value it creates.*

---

16

## The KPG: A Common Decision Substrate

The Knowledge Property Graph is what lets Causum operate above any particular model or AI stack.

A KPG represents the assertions relevant to a matter — with their properties, relationships, provenance, governing context, and state. Those assertions may come from machine output, enterprise data, policy, regulation, contracts, domain experts, external evidence, or prior adjudication.

### SOURCES OF ASSERTION



Fig. 17 — Neither machine nor human is treated as inherently authoritative. The KPG relates and evaluates both.

So the clearinghouse never needs to ask, “*Did you use our AI?*” It asks a more important question: “*Can this matter be represented as evidence-bearing assertions capable of adjudication?*”

For a smaller firm, that means bringing an AI-enabled result into Causum without deploying a full enterprise harness. For a company without sophisticated AI, it means bringing a difficult decision and letting the platform assemble the knowledge, evidence, machine assistance, and expertise around it.

17

## The Clearinghouse Network Effect

Causum's durable value lies less in software features than in the network that forms around the clearinghouse.




Fig. 18 — A two-sided network. Each matter makes the next one cheaper to clear.

As participation grows, more matters can be constituted, more domain KPGs mature, more adjudicative expertise becomes reachable — and more organizations obtain sophisticated judgment without rebuilding that capacity themselves.

This reframes the relationship with frontier-model and cloud vendors. Their improvements in identity, policy enforcement, execution traces, and observability make the clearinghouse's inputs better. Causum need not compete to own the execution substrate. It is not merely an AI governance product; it is potential market infrastructure for the exchange and adjudication of sophisticated decision capability.

---

*Causum's defensible role begins where a producer's authority to describe its own execution ends.*

---

18


## Market Evidence: United States and Europe

The clearinghouse is a Causum market proposition, not an established industry category. But current evidence on both sides of the Atlantic supports the structural conditions from which such a market could emerge.

### AI USE, BY FIRM SIZE

share of firms, %

#### UNITED STATES · CENSUS BTOS, NOV 2025 – JAN 2026



#### EUROPEAN UNION · EUROSTAT, 2025




Fig. 19 — Capability is diffusing, but not evenly. Large firms adopt AI at roughly three times the rate of small ones. Sources: U.S. Census Bureau (2026); Eurostat (2026).

US

**Uneven internalization.** Census research on the 2026 Business Trends and Outlook Survey found 18% of firms used AI in a business function — 32% on an employment-weighted basis. Very large firms in Information, Professional Services and Finance reached roughly 50–60%. AI capability is diffusing; it is not being internalized evenly. [\[U.S. Census Bureau, 2026\]](#)

---

EU

**A pronounced size divide.** Eurostat reports 20.0% of EU enterprises used AI in 2025: 55.0% of large enterprises, but only 18.9% of SMEs. Many smaller organizations could benefit from sophisticated machine-assisted judgment without reproducing a large enterprise's infrastructure and specialist capacity. [\[Eurostat, 2026\]](#)

---

US

**Evidence is becoming part of execution.** Amazon Bedrock AgentCore enforces policy centrally, outside agent code, and records policy evaluations through observability. In July 2026 AWS added unified per-agent observability across traces, prompts, structured logs, and output. Richer execution evidence makes a better input to independent clearing — complementary, not competitive.

[\[AWS AgentCore Policy, 2026\]](#) [\[AWS AgentCore Observability, 2026\]](#)

---

US

**Accountability across the lifecycle.** NIST's AI Risk Management Framework treats accountability and transparency as core to trustworthy AI, emphasizing monitoring, testing, evaluation, verification and validation. Its 2026 critical-infrastructure work extends trustworthiness requirements across lifecycles and supply chains. This does not mandate a clearinghouse — but it supports evidence that travels across organizational boundaries. [\[NIST AI RMF\]](#) [\[NIST Critical Infrastructure Profile, 2026\]](#)

---

EU

**Traceability and human oversight.** The EU AI Act requires high-risk systems to support automatic event logging, effective oversight by natural persons, and enough transparency for deployers to interpret outputs. It does not prescribe Causum's architecture; it reinforces the distinction between machine execution, persistent evidence, human oversight, and institutional responsibility. [\[EU AI Act\]](#)

---

EU

**Relief for smaller firms.** The 2026 AI Omnibus extends simplification to small mid-caps, widens access to regulatory sandboxes, and reduces administrative burden while retaining safeguards. The underlying problem is the same: smaller organizations need sophisticated AI and governance capability without the fixed institutional cost of the largest enterprises. [\[European Commission, 2026\]](#)

---

#### FOUR STRUCTURAL CLAIMS THE EVIDENCE SUPPORTS

- |   |                                                                            |   |                                                                                                             |
|---|----------------------------------------------------------------------------|---|-------------------------------------------------------------------------------------------------------------|
| 1 | AI adoption and internal capability are unevenly distributed by firm size. | 2 | AI systems are moving toward more autonomous, consequential work.                                           |
| 3 | Execution platforms increasingly produce policy and runtime evidence.      | 4 | U.S. and EU frameworks emphasize accountability, traceability, oversight, and demonstrable trustworthiness. |

The Causum-specific proposition remains a thesis to be tested. It follows directly from these conditions: *if sophisticated intelligence can be produced anywhere, but the capacity to establish justified reliance remains institutionally uneven, a neutral clearinghouse creates value* — by converting heterogeneous evidence and complex decisions into a common adjudicable form, and connecting them with qualified judgment.

## Capability Inequality and the First Market

Access to a frontier model does not give every organization equal capacity to turn machine intelligence into trusted economic action. The surrounding capabilities — specialist skills, data, compute, integration, governance, evidence, domain knowledge, assurance — remain costly and unevenly distributed.

The evidence is consistent. Eurostat attributes the size gap to large firms' greater ability to manage implementation complexity, achieve economies of scale, absorb investment, set strategic priorities, and supply expertise. [\[U.S. Census, 2026\]](#) [\[Eurostat, 2026\]](#) OECD research is sharper still. Its 2025 G7 paper names connectivity, data, algorithms and compute, skills, and finance as prerequisites for SME adoption. Its 2026 D4SME survey finds that time, maintenance costs, and skills gaps still hinder implementation: most AI-using SMEs remain on simple off-the-shelf tools, and few run tailored solutions or broad deployments. [\[OECD, 2025\]](#) [\[OECD D4SME, 2026\]](#)

---

*The real divide is not access to AI. It is access to the institutional machinery that turns intelligence into trusted decisions.*

---

Large organizations internalize that machinery through harnesses, specialist teams, governance systems, and assurance processes. Capability-constrained organizations cannot justify those fixed costs — even when their decisions are just as consequential.

Causum's proposition is to **rightsize that asymmetry**. Not to make every organization technologically equal, and not to claim it has already reduced inequality — but to reduce how much access to sophisticated decision capability depends on owning enterprise-scale AI infrastructure.

That defines a practical first market: capability-constrained organizations facing decisions whose required intelligence, expertise, evidence, or assurance exceeds what they can economically maintain. The customer need not buy an AI transformation or understand the protocol stack. *It brings the matter requiring determination.*



Fig. 20 — The first-market path: bring the decision, not the AI stack.

This follows an established pattern: smaller firms externalize specialist capacity. OECD's 2026 survey shows SMEs relying substantially on external channels — including consultants — to source digital skills; earlier OECD work describes external knowledge markets and cloud services filling the same role. Causum extends that logic upward, from externally sourcing software or technical skill to externally accessing *structured, adjudicated decision capability*. [\[OECD External Skills, 2026\]](#) [\[OECD Digital Transformation of SMEs\]](#)

The thesis is deliberately narrower than “democratizing AI.” Causum seeks to lower the institutional cost of participating in the economics of intelligence — so that a smaller organization need not reproduce the capital, workforce, governance stack, and assurance machinery of a large enterprise to obtain a decision of comparable institutional quality.

It is also the bootstrap path. Constrained organizations generate real decision matters. Those matters create demand for domain KPGs, evidentiary conventions, and qualified adjudicators. As those assets accumulate, the same infrastructure

accepts AI-enabled results from sophisticated firms and evidence from enterprise harnesses. **The first market builds the network the broader clearinghouse requires.**

# From Expert Systems to Cleared Judgment

Where have organizations historically tried to encode or distribute scarce expert judgment? Expert systems took root in bounded domains where decisions were consequential, specialist knowledge was expensive, rules and evidence could be represented, and an expert could not economically sit beside every decision. Those same traits are a useful — not definitive — signal for where Causum should enter.




Fig. 21 — Two models for distributing scarce judgment.

Causum does not need to reproduce the expert. It uses structured knowledge and machine intelligence to constitute the matter, reduce the search and analytical burden, and expose uncertainty and conflict — then invokes scarce human expertise only where authority or professional judgment actually matters. That changes the economics of distributing sophisticated judgment.

---

## THE DOMAIN-ENTRY TEST

- |                               |                                      |                                                |                                                           |                                        |
|-------------------------------|--------------------------------------|------------------------------------------------|-----------------------------------------------------------|----------------------------------------|
| 1                             | 2                                    | 3                                              | 4                                                         | 5                                      |
| Scarce or expensive expertise | Repeatable but non-trivial decisions | Representable evidence and governing knowledge | Meaningful consequences from an unsupported determination | Full in-house capability is uneconomic |

| DOMAIN                                           | REPRESENTATIVE DETERMINATIONS                                                                                                                         | STATUS          |
|--------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|
| <b>Cybersecurity, GRC &amp; compliance</b>       | Accept a risk, approve a supplier, close a finding, escalate                                                                                          | Early candidate |
| <b>Specialty insurance</b>                       | Constitute unusual risks; refer material ambiguity to an authorized underwriter                                                                       | Early candidate |
| <b>Manufacturing &amp; industrial operations</b> | Return equipment to service, disposition a defect, justify a supplier substitution                                                                    | Early candidate |
| <b>Accounting, tax, audit &amp; advisory</b>     | Unusual or cross-jurisdiction matters for regional and specialist firms                                                                               | Candidate       |
| <b>Healthcare decision support</b>               | Strong lineage; liability, safety, privacy and device regulation raise the bar                                                                        | Later           |
| <b>Others</b>                                    | Procurement and supply-chain risk, specialty lending, engineering and environmental compliance, energy, agriculture, construction, public procurement | Potential       |

Fig. 22 — Candidate domains. A market hypothesis, not a demonstrated ranking.

**Cybersecurity, GRC and compliance.** Small and mid-sized organizations routinely face control, supplier-risk, architecture, regulatory, and remediation decisions without dedicated specialists. NIST's small-business guidance acknowledges the need for outside providers; workforce research shows persistent gaps in risk and governance skills. The matter is never just whether a control fired — it is whether the evidence, requirements, and business context justify accepting risk, approving a supplier, closing a finding, or escalating. [\[NIST Small Business Cybersecurity\]](#) [\[ISC2, 2025\]](#)

**Specialty insurance.** Underwriting and claims share every trait that drew expert-system research: scarce expertise, repeatable but exception-rich decisions, structured evidence, explicit authority limits, and referral paths. Causum would not replace the underwriter. It would constitute unusual risk matters, organize evidence and underwriting knowledge, prepare assertions, and refer material ambiguity to an authorized underwriter. [\[Insurance Expert Systems Research\]](#)

**Manufacturing and industrial operations.** Fault diagnosis, configuration, maintenance, and process control were classic expert-system problems. OECD highlights predictive maintenance, quality assurance, and supply-chain optimization in European manufacturing — alongside skills, data, and investment barriers for

smaller firms. Should equipment return to service? Is a defect disposition acceptable? Is a supplier substitution justified? Each determination maps directly onto downtime, scrap, quality, and operational risk. [\[OECD AI in Manufacturing\]](#)

**Accounting, tax, audit and advisory.** Expert systems were applied to taxation, auditing, financial planning, and business-risk assessment. Regional and specialist firms sell judgment but cannot maintain expertise across every jurisdiction or unusual matter. Causum could give them structured governing knowledge, machine-assisted preparation, and external specialist adjudication — while preserving the line between analytical assistance and professional authority. [\[Accounting Expert Systems Research\]](#)

**Healthcare** has the most famous expert-system lineage, but clinical liability, patient safety, privacy, and medical-device regulation raise the threshold. It demonstrates the architecture's reach rather than a first commercial domain.

On this basis, cybersecurity/GRC, specialty insurance, and industrial decision support warrant particular investigation as early domains. Cybersecurity already shows substantial outsourcing and structured control evidence. Insurance has a native language of authority, evidence, referral, and exception. Manufacturing ties expert determinations to directly observable value.

---

*Expert systems began with a problem: “We cannot put an expert next to every decision.”*

Expert systems tried to distribute expertise by reproducing the expert. Causum distributes judgment by clearing access to knowledge, machine intelligence, evidence — *and the expert, only where necessary.*

---

## CONCLUSION

The AI market is rapidly solving machine capability. It is solving the operationalization of that capability through agents, harnesses, context systems, controls, and execution infrastructure. It is beginning to solve the economics of consuming and allocating machine intelligence.

### One institutional problem remains.

As cognition and authority are delegated to machines, enterprises need a way to establish whether evidenced machine reality remains justifiable against legitimate human and institutional intent.

|                                  |                      |
|----------------------------------|----------------------|
| Model capability alone           | cannot solve it.     |
| Authorization alone              | cannot solve it.     |
| Observability alone              | cannot solve it.     |
| Another AI as unquestioned judge | should not solve it. |

It requires evidence, governing knowledge, explicit authority, and legitimate adjudication. That is the position Causum is built to fill.

**Causum makes delegated machine activity adjudicable** — *the bridge from AI Intent to AI Reality, and from machine execution to justified reliance.*

# Causum, in seven lines

|                                |                                                                                                                                                                                |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CATEGORY                       | Adjudication infrastructure for delegated machine intelligence                                                                                                                 |
| PROBLEM                        | The Intent–Reality Gap                                                                                                                                                         |
| CORE FUNCTION                  | Make consequential machine activity adjudicable                                                                                                                                |
| ECONOMIC CONTEXT               | The Political Economy of Intelligence                                                                                                                                          |
| OPERATING PRINCIPLE            | Justified autonomy                                                                                                                                                             |
| DETERMINATION STATES           | <div>CONFORM</div> <div>DIVERGE</div> <div>REFER</div>                                                                                                                         |
| RELATIONSHIP TO THE AI HARNESS | Complementary. The harness operationalizes and evidences machine work; Causum makes the relationship between legitimate intent and evidenced reality capable of determination. |

EDITORIAL NOTE ON EVIDENCE

This paper separates two things deliberately. The broad market developments — agentic execution, enterprise AI harnesses, delegated machine work, token and outcome economics, AI governance, and execution evidence — are observable industry and research trends. The Intent–Reality Gap, the framing of the Intent–Reality relationship as an adjudicable object, the clearinghouse proposition, and the Causum architecture are Causum's positioning and thesis.