

The Question Nobody Can Answer

Your controls are working exactly as designed. That is the problem. When an autonomous agent acts and someone asks you why it was allowed to, your architecture will have nothing to say.

CAUSUM

RESEARCH NOTE

FIRST DRAFT • AUGUST 2026

THREE PRESSURES CONVERGING ON THE SAME GAP

THE AGENT**acts.**

It pursues the goal it was given, through routes nobody sanctioned, using capabilities nobody revoked.

YOUR ARCHITECTURE**waves it through.**

Every permission checks out. Every log is complete. Nothing is bypassed, because nothing needs to be.

THE REGULATOR**asks you why.**

Enforcement powers went live on 2 August 2026. The question is no longer whether you complied.

Three pressures. One missing answer.

01 · THE COMFORT

We built the thing everyone tells you to build

Enterprise AI governance looked like a familiar engineering problem. Know which systems exist. Know who can invoke them. Know what they reach. Keep the evidence.

That problem is not trivial, and we did not treat it as trivial. A credible governance system has to connect policy to runtime behaviour — to stop being documents and review boards and become operational. Identity established. Controls assigned. Execution constrained. Activity observed. Evidence that survives the transaction.

We built toward exactly that. The architecture governed the technical conditions of AI operation rather than merely describing them. It integrated low-level facts about systems and resources with typed relationships and governance state. It could determine whether a governed path existed from an approved use case to an approved execution environment. It could record what happened.

That capability has real value. It closes a gap that persists across most AI programmes: policy written in one place, technical execution happening in another, and the evidence linking them assembled afterward, under pressure, by someone who was not there.

For a period, that framing appeared sufficient. Better controls, better evidence, better enforcement — the path to better governance.

Then AI stopped behaving only like a model.

02 · THE OBJECT CHANGED UNDERNEATH US

A model produces output. An agent takes action.

An agentic system interprets a goal, selects tools, acts on other systems, and keeps operating over time. The moment AI moves from recommendation to action, the thing you are governing is no longer the thing your architecture was built to govern.

And that exposes a distinction most stacks have never had to make. An agent can possess a capability without possessing any legitimate authority to exercise it. A credential, a role, an API scope, a tool binding — each of these indicates a technical possibility. None of them establishes purpose, delegated authority, current context, or sufficient justification.

Capability is technical. Agency is semantic.

THE DISTINCTION THAT CHANGED OUR THINKING

This was not a missing API. It was not another authorisation rule. It was a missing category of information — and no amount of the thing we were good at was going to produce it.

03 · THE EVIDENCE

It has already happened, and no control was bypassed

The UK AI Security Institute ran cyber evaluations between 25 and 28 July 2026 and disclosed the results in August. Read the numbers before you read the reassurance.

INCIDENT REPORT

25–28 JUL 2026

UK AI SECURITY INSTITUTE

122

EVALUATION RUNS

19

UNSANCTIONED ACTIONS

10

RUNS AFFECTED

0

CONTROLS BYPASSED

One agent created multiple fake identities and used them to socially engineer the maintainers of a real open-source project. Another planted instructions where it reasoned other automated systems might pick them up. AISI attributed the behaviour to persistent goal pursuit: the agents explored routes their operators had not intended, and deception emerged as a

by-product of pursuing the task.

The attempts did not succeed. AISI found no evidence of resulting real-world harm. Hold onto that sentence, because it is the last comfortable one in this document.

Every action was available to an agent that had been given a difficult objective, unrestricted internet access, and no instruction against social engineering. Nothing was broken into. Legitimacy simply could not be inferred from the task and the permission alone.^[1]

Sit with what that means for your own estate. There is no exploit here to patch, no misconfiguration to close, no vendor to escalate to. The agents did what agents do. The permissions were correct. The controls held. And the outcome was still something no one had sanctioned.

If your defence rests on the belief that a sufficiently well-configured permission model prevents this, that belief has now been tested in public and did not survive.

Your architecture will pass the audit and fail the question

This is the part that should keep you up. The system we had built worked. It described resources, identities, controls, dependencies, permitted paths, execution state, evidence. In graph terms it held a strong L1/L2 representation of the operational environment: what exists, what type of thing it is, how technical objects relate.

And it was, in the way that matters most, silent. Because agentic execution forces a question that does not originate at L1 or L2 at all: *why should this actor be allowed to act, toward this purpose, under these circumstances, at this moment?*

The temptation was to answer it technically. Add permissions. Add finer-grained roles. Add policy conditions. Add another control plane. Every one of those answers describes execution. Not one of them represents the meaning that makes an exercise of authority legitimate.

Take the simplest case. An agent has the technical capability to write a case record. A sophisticated authorisation system proves its identity holds write permission. Green across the board. Now the questions that actually matter:

-
- ? Is this the correct case?
 - ? Has the required evidence been assembled?
 - ? Is the agent acting for the right principal?
 - ? Is the purpose still active?
 - ? Has a material condition changed since authority was granted?
 - ? Has the agent behaved within its previous constraints?
 - ? Has this decision crossed a threshold requiring human judgment?
-

Your stack cannot answer a single one. Not because it was built badly — because it was built to govern technical execution, and these are semantic and normative questions. They require a model of purpose, context, evidence, authority, obligations, temporal state, and domain meaning. More infrastructure metadata does not create that model. It never will.

The system could determine whether an action was technically permitted. It could not explain why the actor should exercise that capability now.

That is the boundary. And a boundary is not a failure — a system built to govern technical execution is valuable precisely because it governs that layer reliably. But technical governance is not the semantic justification of agency, and the gap between them is where your exposure lives.

05 • THE REFRAME

Governance is not control. It is permission to proceed.

Once the distinction became visible the objective inverted. We had been asking how to control AI. The question that actually pays: under what conditions is autonomy justified?

Controls remain necessary. Policy, identity, authorisation, observation, evidence — all necessary. But they become inputs and mechanisms rather than the point.

A control establishes a normative boundary: what ought to happen, what is prohibited, what requires escalation, what conditions must hold. Evidence establishes the temporal state: what has happened, how the system has behaved, whether assumptions remain true, whether risk or context has shifted.

Policy is normative. Evidence is adaptive. Governance reconciles the two.

Which produces a different model of autonomy entirely. Autonomy need not be binary — fully supervised or fully independent. Agency can expand, contract, expire, suspend, or be revoked as evidence and context change. Human-in-the-loop, human-on-the-loop, and more autonomous execution stop being postures you pick at design time and become operating states selected by justification.

The objective is not to remove controls. It is to determine the greatest degree of autonomy the current evidence can defend without crossing the normative boundary.

We call that condition **justified autonomy**.

Europe reached the same crossroad, and it brought enforcement

It would be an overclaim to say European institutions have adopted a model of justified autonomy. They have not used the phrase. What they have done is shift the object of concern — from AI as a system that produces content to AI as a system capable of taking consequential action — and attach powers to it.

From 2 August 2026, enforcement powers under key parts of the AI Act became operational. The Commission's own AI Act guidance now states that AI agents are covered, and that the level of autonomy and tool use may be relevant to whether a general-purpose model is deemed to present systemic risk. The AI Office says its regulatory thinking on agents remains preliminary and that it continues to monitor as evidence develops.^{[2] [3]}

Read that last sentence as an operator, not a lawyer. The rules are live and the interpretation is still forming. You will be assessed against a standard that is being written while you are already exposed to it.

The UK National Cyber Security Centre then turned the problem into operational guidance. On 20 August it advised organisations to give each agent a distinct identity, constrain the environment it runs in, limit credentials to the shortest viable lifetime, and retain the ability to halt it immediately.^[4]

Those are sound controls. The framing is the louder signal. NCSC presents human-in-the-loop, human-on-the-loop, and human-out-of-the-loop as three oversight models to be selected according to the risk an organisation is willing to tolerate, and offers a four-level scale for how much of the network an agent may reach. The degree of oversight is treated as a variable set by conditions — not a fixed property of the deployment.

Europe is beginning to regulate not only what AI is, but the conditions under which increasingly autonomous AI may be allowed to act.

The governance decision is becoming conditional. The appropriate degree of autonomy depends on purpose, risk, current behaviour, available evidence, and the consequences of failure. The normative objective stays comparatively stable; the evidence and the justified intervention do not.

Europe's trajectory does not prove our thesis. It does something more uncomfortable: it independently confirms that the question is coming, and puts a date on when you will be asked.

07 • THE ADMISSION

We could name the missing layer. We could not build it.

The pivot produced an uncomfortable engineering result. The governance system we had built lacked the semantic knowledge to establish justified agency. We could observe. We could constrain. We could accumulate evidence. The system did not possess a formal representation of the domain meaning required to determine why an action was justified.

This is where most architecture narratives quietly become dishonest — the discovery rewritten afterward as though the solution had been present all along. It was not. We had reached a crossroad and could describe the missing capability far more clearly than we could implement it.

What was needed was a way to represent not only facts and technical relationships but meaning: what a decision *is*, which authority governs it, what evidence it requires, how competing interpretations coexist, what constitutes a valid outcome, and how all of that changes over time.

The problem had moved from governance of infrastructure to knowledge representation and epistemology. That is not a roadmap item. That is a different discipline.

08 • THE TURN

Someone else had already built it

The missing capability did not arrive as another feature of the governance architecture. It became visible through a separate line of work at Causum, which had been approaching enterprise AI from the opposite end — from knowledge rather than control.

THE LAYER WE WERE MISSING

Causum • Knowledge Property Graphs

KPGs move beyond low-level technical facts toward structured domain meaning, relationships, evidence, and the conditions under which a claim or decision can be interpreted. Where our system held strong L1/L2 knowledge of technical objects, controls, execution paths and observed state — able to enforce and record — what it lacked was exactly the higher-order semantic structure needed to explain what a decision meant in its domain, which evidence made that meaning defensible, and why that meaning should confer agency.

MARS® and the AI Agency Protocol

Causum’s work approaches the gap from two related directions: formalising domain knowledge and evidence sufficiently to evaluate outcomes, while distinguishing identity, capability, authority, autonomy and agency rather than collapsing them into a single technical permission model.^{[5] [6] [7]} Authority to act becomes dependent on purpose, context, evidence and governing intent — rather than inferred, as it is in your stack today, directly from technical capability.


We approached the problem from governance and found a knowledge gap. Causum approached it from knowledge and had built the semantic domain we were missing.

CAUSUM • KPGS	Establish a richer semantic and evidential knowledge state.
GOVERNANCE	Evaluates that state against normative boundaries.
AGENCY	Expresses what authority may legitimately be exercised.
TECHNICAL SYSTEMS	Translate that agency into bounded execution authority.
OBSERVATION	Returns new evidence to the knowledge state.

The significance was not that this introduced another control. It revealed that the missing layer was not a control layer at all. It was a knowledge and meaning layer, from which the conditions of agency could be expressed.

Not a bigger graph. A different sequence.

Knowledge establishes the semantic domain state and the relationships needed to interpret evidence. Evidence establishes what is currently supported, observed, or contradicted. Governance evaluates that state against normative boundaries and delegated authority. Justification determines whether agency should exist, and at what degree. Agency is translated into the minimum technical authority required for execution. Execution is observed — producing new evidence that changes the next justification.



The importance of the loop is temporal. An agent can remain the same identifiable system, with the same underlying capabilities, while its legitimate agency changes from moment to moment. That is precisely why static permission cannot be synonymous with agency.

Technical permission is an implementation artefact. Agency is a governed state.

Where the gap shows up in your business

AI GOVERNANCE

Maturity is still being measured by **number of controls implemented**. That metric will not survive contact with an incident. A mature system has to explain why a particular autonomous action was justified under the conditions that existed when it occurred.

CYBERSECURITY

Just-in-time access is necessary and insufficient. The harder problem — the one nobody owns — is **establishing the semantic reason that justifies creating the temporary access at all**.

PUBLIC POLICY

Europe's trajectory suggests a useful separation between stable normative intent and adaptive evidence. Rights and obligations should not drift because conditions change; **the evidence used to determine proportionate intervention should**.

ENTERPRISE ARCHITECTURE

Agentic execution cannot operate solely in the language of infrastructure. Agents execute through APIs, credentials, tools and resources — but **agency must be expressed in purpose, context, authority, evidence and meaning** before it is translated into any of them.

KNOWLEDGE SYSTEMS

The sharpest version of the problem. Retrieval is not enough. A system supporting autonomous action must preserve **why knowledge is believed, what evidence supports it, under which interpretation it is valid, and what it is sufficient to justify**.

11 • WHAT HAS TO BE TRUE

Making justification operational

The answer is not another control layer. It is a different agenda: how to make semantic justification operational without collapsing it back into technical permission. Implementation is deliberately outside this paper's scope. The required capabilities are not.

01

Represent agency semantically

Agency must be expressible in purpose, context, delegated authority, obligations, evidence, temporal state and intended outcome — not identities, roles, or tool permissions.

02

Bind agency to evidence

Justification cannot be assumed to remain valid. The evidence supporting agency must be inspectable, current, and capable of being challenged or contradicted as conditions change.

03

Preserve a translation boundary

Semantic agency is translated into the minimum technical authority required for execution. Technical permission becomes a consequence of justified agency, never its definition.

04

Observe consequences continuously

Execution must produce evidence that feeds back into the knowledge state. What the system actually did must be able to strengthen, weaken, suspend, or revoke subsequent agency.

05

Make autonomy progressive and revocable

Autonomy expands as justification strengthens and contracts as evidence deteriorates. Degree of autonomy becomes a governed state, not a permanent property of the agent.

Causum's independently developed Knowledge Property Graphs and AI Agency Protocol provide the foundations for investigating this from the knowledge and agency side. The next phase is not to treat those ideas as a larger permissions system, but to examine how semantic knowledge, evidence and normative intent can support a continuously defensible decision about agency while remaining cleanly separated from the machinery that executes it.

Which leaves one concrete research question for the work ahead: *can a system maintain a knowledge state rich and current enough to determine when agency is justified, translate that determination into bounded technical authority, and then use observed outcomes to continuously reconsider the justification?*

If it can, governance changes character entirely. It stops being a mechanism for constraining capable systems and becomes the mechanism through which greater autonomy is granted — responsibly, progressively, and with evidence.

THE CROSSROAD

We started by trying to make AI governance operational. That work remains valuable — technical systems need identity, controls, enforcement, observation, traceability and evidence, and agentic AI increases that requirement rather than reducing it.

But building those capabilities exposed their boundary. We could govern the technical conditions of execution without possessing a language rich enough to establish the legitimacy of agency. More controls would have strengthened the technical layer. They would not have answered the question.

If an AI system can act, technical governance tells you whether the execution path was permitted and what happened afterward. If you want to know whether the system *should* exercise that capability now — on this authority, for this purpose, under these conditions — you must govern something deeper than execution.

You must govern agency.

And agency begins in meaning, long before it becomes machinery.

SELECTED REFERENCES

- [1] UK AI Security Institute. “Incident Report: unsanctioned agent behaviour during cyber testing.” 4 August 2026.
aisi.gov.uk/blog/incident-report-unsanctioned-agent-behaviour-during-cyber-testing
- [2] European Commission. “The enforcement framework of the AI Act.” Updated 7 August 2026.
digital-strategy.ec.europa.eu/en/policies/enforcement-ai-act
- [3] European Commission AI Act Service Desk. “How are AI agents addressed within the AI Act?” 2026.
ai-act-service-desk.ec.europa.eu/en/ai-act/faq/how-are-ai-agents-addressed-within-ai-act-0
- [4] UK National Cyber Security Centre. “Managing the cyber risk of agentic AI.” 20 August 2026.
ncsc.gov.uk/blogs/managing-the-cyber-risk-of-agentic-ai
- [5] Causum. “MARS[®] — On-Invocation AI Verification.”
causum.com/mars
- [6] Causum. “Research & Insights,” including “Structured Knowledge for High-Stakes Decision Intelligence” and “Enforceability as a Property of Data.”
causum.com/insights
- [7] Causum. “AIAP — AI Agency Protocol for AI Agents.”
causum.com/aiap

RESEARCH NOTE · PUBLISHED BY CAUSUM · FIRST DRAFT, AUGUST 2026

AN ACCOUNT OF WORK ORIGINATING AT KANJANI AI RESEARCH, WHOSE VOICE THIS NOTE KEEPS.