

PRACTITIONER FIELD NOTE

The hours were never the product

Why the consulting selloff is a repricing, not an ending, and what the first governed step looks like for an organisation that can neither stop nor wave everything through.

Jason Cruz Falzon

Creator of the Cognitive Delivery Framework

Version 1.0 · cognitivedelivery.co.uk

What the market is actually pricing

A consultancy was never bought for the hours. It was bought to carry risk a client could not carry alone, and to buy pattern knowledge built across many clients. AI is collapsing the way that got priced. It is not touching the thing being bought.

Over a year, Accenture has lost more than half its market value. Capgemini is down around 40 per cent. The reading attached to those numbers is that AI automates the work consultancies do, so clients no longer need them. Revenue at the same firms has in several cases still grown, which is the first sign the simple story does not fit.

The simple story mistakes the billable hour for the product. The hour was a unit of pricing. It was a proxy for something a client could not assemble on its own: experience drawn across many engagements, and a counterparty willing to stand behind a recommendation when it carried consequences.

The proxy is what is collapsing

Every prior shift in this industry repriced the unit while leaving the job intact. Offshore moved where the hours were done. Fixed-price moved who carried delivery risk. Managed services moved the unit from time to outcome. In each case the hour changed and the underlying job, carrying risk and lending cross-client experience, stayed constant.

AI is doing it again, faster. It compresses the production work that used to fill the hours. What it does not do is carry accountability for what gets shipped, or supply the judgement that comes from having seen the same failure across thirty organisations.

WHERE THE READING GOES WRONG

Watching billable hours as the health metric is like watching a thermometer to diagnose the engine. The hours are falling because the production work compressed. That tells you nothing about whether the risk-bearing and the cross-client judgement, the things actually being bought, are intact or not.

Capability raises the floor, not the ceiling

There is a reason the thing being bought gets scarcer, not more abundant, as the models improve. A more capable model raises the floor of plausible-looking output. Code reads as competent. A recommendation reads as considered. An analysis reads as sound.

Plausible-looking output is precisely what an organisation without governance cannot challenge. When the floor of plausibility rises, an ungoverned decision becomes harder to question, not easier to catch. Capability is the most measured variable in the room and the least decisive one.

"Capability without governance is speed without control."

This is why the competitive question has moved. It is no longer how many hours a firm can bill, or even how capable its tools are. It is whether the firm can govern AI-assisted delivery at the point where decisions actually get made, rather than document that it intends to.

The distance that now matters

There is a gap between a control that is written down and a control that demonstrably changes what ships. A policy nobody consults at the moment of action is documentation. A review step that is signed but not engaged with is a signature. An escalation route that has never been used is a diagram.

That distance is the whole game. When model vendors will embed their own engineers and govern their own output, "we have convened a committee and are drafting a policy" is not a competitive position. It is a tell that the controls have not reached the point of action.

THE STOP-EVERYTHING REFLEX

Halting delivery while a committee writes a policy is itself the symptom. It signals that the controls are being designed away from the work, in a room, rather than instrumented around the work as it happens. By the time the policy lands, the teams have been shipping with AI for months without it.

The first governed step

The reframe most organisations need is this. A governance programme is not a precondition you complete before developers begin. It is something you instrument around work already happening. None of the steps below requires stopping, and none requires buying anything.

Each step is a move you can make this quarter. Read on its own, each is method-neutral. Underneath, each maps to one dimension of how AI-assisted delivery is actually governed, named in small caps so you can see the spine.

01 **Make one decision reconstructable**

Take a single change that shipped with AI assistance last week. Try to answer, from records that already exist: what was asked of the model, what it produced, and who decided it was fit to ship. If you cannot, you have found your first gap. Fix it for one team before you write it into a policy for forty.

PROVENANCE INTEGRITY

02 **Name who is accountable, not who is capable**

For that same change, write down who is answerable for what shipped. Not who is most skilled with the tool. The right axis is exposure to the output, not peak capability. If the answer is "the model" or "the team, collectively," you have located a supervisory gap worth closing first.

SUPERVISORY CLARITY

03 **Put the scope where the work happens**

Check whether the constraints a change must respect are visible at the point the work is done, or whether they live in a document nobody opens at that moment. A scope that sits outside the runtime context is a scope that does not constrain. Move one constraint to where the decision is made.

SCOPE DEFINITION

Three more, then the test

04 Ask whether review still happens

Look at how much now merges without a human meaningfully engaging, as opposed to a human signing. Volume has grown faster than review capacity almost everywhere AI is in use. A signature is not engagement. Pick the highest-volume pipeline and ask honestly what review still means there.

REVIEW SUBSTANCE

05 Test one control at the point of action

Choose a single control you believe is in force. Try to ship something that should violate it. If it ships anyway, the control is documented, not operating. This is the fastest way to find the distance between your level four on paper and your level five in practice.

CONSTRAINT FIDELITY

06 Calibrate to exposure, not enthusiasm

Map where the most capable AI use sits against where the output carries the most consequence. The two are rarely the same place. Govern first where exposure is highest, not where capability is most exciting. That is where an ungoverned decision costs the most.

CAPABILITY CALIBRATION

A CONTROL THAT GOVERNS

- It changes what ships at the moment of action.
- Someone has been stopped by it in the last month.
- It is referenced in the work, not just in the policy.

A CONTROL THAT PERFORMS

- It is documented, signed, and never consulted.
- No one can recall it ever blocking anything.
- It lives in a folder, away from where work happens.

What this means for the firms

Return to where this started. The selloff reads as an ending only if you believe the hours were the product. They were not. The work that endures is the work AI makes scarcer: bearing risk on behalf of a client, and lending judgement built across many engagements that a single organisation cannot build alone.

The model is changing the way it has changed before. The unit moves. The job holds. A firm that can show, not assert, that it governs AI-assisted delivery at the point of action is selling exactly the thing the market is short of. A firm still counting hours is measuring the variable that is supposed to fall.

"The unit of pricing is collapsing. The thing being priced is not."

None of this requires a stop. The first governed step is small, runs on records that already exist, and starts with one team and one change. The organisations that come through this will not be the ones that paused longest to write the best policy. They will be the ones that instrumented the work they were already doing.



Jason Cruz Falzon

Creator of the Cognitive Delivery Framework

jason@cognitivedelivery.co.uk · cognitivedelivery.co.uk

cognitivedelivery.substack.com