WORKFORCE · A WORKING FRAMEWORK BY JAMES TUTTIETT

The AI Fluency Framework

Four tiers that describe how someone works with AI, not what they know about it.

Most AI training programmes measure awareness: who has done the course, who has a licence. None of that tells you whether work is being done differently. This framework replaces the question "have they been trained?" with "how do they work now?", and gives every role a target tier.

It is built around three behaviours (confident, curious, secure) and applied inside each practice (Build, Scale, Run), so a tier means something specific for an engineer, a change lead or an operations analyst rather than one generic standard for everyone.

Tier 1: AI Aware

Understands what the tools can and cannot do. Uses them occasionally. Knows the safe-use rules and follows them. This is the floor for everyone in the organisation, and it is reached in days, not months.

Tier 2: AI Productive

Uses AI every day on their own tasks: drafting, analysis, research, preparation. Can show time saved and quality maintained. Most knowledge workers should be here within a quarter of starting, and most organisations stall at this tier because nothing has changed except the individual.

Tier 3: AI Operational

Has redesigned at least one workflow or process around AI, not just their own habits. Builds reusable prompts, skills or templates that colleagues pick up. Owns the quality of AI-assisted output in their area and can explain where it should not be used.

Tier 4: AI Orchestrating

Directs agents and pre-built skills across multi-step work. Designs how a team, a function or a client engagement operates with AI in the loop. Accountable for outcomes rather than tasks. This is the tier that changes the economics of a business, and it is where senior practitioners need to be.

How to use it

- Set a target tier for every role, not every person. A Head of Practice and a first-year consultant need different tiers.
- Move people one tier at a time. The jump from Productive to Operational is the hard one, because it needs permission to change how work is done, not just a tool.
- Measure by what has changed in practice: workflows redesigned, skills reused, outcomes owned. Licence activation and course completion are not measures.
- Bake safe and responsible use into the tools rather than a policy document. Secure is a behaviour, not a sign-off.

James Tuttiett · Global Director of Transformation & AI, FDM Group

OPERATING MODEL · A WORKING FRAMEWORK BY JAMES TUTTIETT

Build, Scale, Run

Three practices organised around what the client needs done, not what the firm can do.

Practice structures usually describe capability: technology, data, transformation, operations. Clients do not buy capability. They have something to build, something that works and needs to grow, or something that has to keep running better and cheaper. Three verbs let the client name what they need in one sentence, and make the firm choose whether it is a talent business or a consultancy.

Build

Something new gets made: a product, a platform, a capability. Senior technologists at the front, small teams, AI in the delivery stack by default. Priced on what ships, not who is in the room.

Scale

Something that works needs to work across the organisation: adoption, change, capability transfer, operating model redesign. This is where most AI programmes fail, because the pilot proved the technology and nobody owned the change.

Run

Things have to keep working: operations, service, resilience, continuous improvement. A Run engagement that is not using automation to take cost and risk out is not a credible one.

Where AI sits

Not in a fourth box. AI is inside each practice by design, and the Head of each practice is accountable for it. A separate AI practice signals that it is optional. It is not.

How to use it

- Test the structure with one question: can a client say which practice they need in a single sentence? If not, the structure is for you, not them.
- Give each Head of Practice ownership of content as well as delivery. The practice has to own what it teaches.
- Expect the move to surface an identity question the firm has been avoiding. That is the point of doing it.
- Keep the specialists as doers, in squads under the practice, not as a management layer.

James Tuttiett · Global Director of Transformation & AI, FDM Group

COMMERCIAL · A WORKING FRAMEWORK BY JAMES TUTTIETT

The BUILD Qualification Lens

Five questions that tell you whether an AI opportunity is real before you spend a week on it.

Most AI opportunities fail qualification for reasons that were visible in the first conversation: no real buyer, a use case chosen for novelty, an incumbent who will absorb the work, or a deal shape the client cannot sign. BUILD is a five-lens check designed for account managers who come from a recruitment or staffing background and need a repeatable way to qualify consulting work.

B: Buyer

Who owns the problem and the budget, and are they the same person? An AI opportunity sponsored by an innovation team without a line-of-business owner will not survive the next budget round. Look for a buyer whose own numbers move if this works.

U: Use case

Is it specific, measurable and painful today? "Explore AI in operations" is not a use case. "Cut the time to produce a client onboarding pack from four days to one" is. If the client cannot name the current cost, there is no business case yet.

I: Incumbent landscape

Who is already in the account and what will they do when this surfaces? The big systems integrators will claim AI work as an extension of what they already run. You need a reason the incumbent cannot or will not do this, or a sponsor who does not want them to.

L: Left-behind appetite

What does the client want to be left with when you leave? Some want a working system and nothing else. Others want capability: a build-operate-transfer model, consultants who convert to permanent, their own teams upskilled. The answer changes the delivery model and the commercial shape, so ask it early.

D: Deal shape

Can this be bought the way it needs to be delivered? Outcome-priced work needs a client who can sign for an outcome. Fixed-scope work needs a scope. If the only contract vehicle in the account is day rates, the proposition has to be shaped to fit that or the deal will stall in procurement.

How to use it

- Score each lens red, amber or green after the first meeting. Two reds means do not invest further until they turn.
- Use the lenses as discovery questions, not a checklist filled in afterwards. They are written so a non-technical account manager can ask them.

- Left-behind appetite is the lens most teams skip and the one that most often decides whether the engagement is a one-off or a relationship.

DELIVERY · A WORKING FRAMEWORK BY JAMES TUTTIETT

The 2+1 AI-Native Pod

The smallest team that can own an outcome, with AI agents as the capacity multiplier.

The traditional delivery model answers "how many people do you need?" with a bench. The AI-native model answers "what outcome do you need?" with the smallest senior-led team that can own it, and uses AI agents to expand execution capacity horizontally rather than adding headcount. The pod is the unit of delivery, the unit of pricing and the unit of accountability.

The shape

A senior technologist who owns the outcome and the architecture. A junior delivery lead who runs the day-to-day, learns on the job and grows into the senior role. Plus AI agents, directed by both, doing the execution work that used to need three or four more people: documentation, testing, analysis, first-draft code, reporting.

Why senior-led

AI makes execution cheap and judgement expensive. The constraint in complex programmes is no longer tooling, it is credible technical leadership at the front. Put the senior person where the judgement is needed and let the agents do the volume.

Why the junior still matters

Because the senior bench of the future has to come from somewhere. The junior role is redesigned around orchestration and judgement rather than manual execution, which is how early-career capability gets built when the entry-level tasks have been automated.

How it is priced

On the outcome the pod owns, not the days it works. The efficiency AI creates is passed back to the client as value rather than hidden in a rate card. That is what makes the model defensible against a cheaper bench.

How to use it

- Start with one pod on one outcome the client can measure. Do not try to convert a whole programme at once.
- Use spec-driven development as the working method: the specification comes first, the agents build to it, the humans review against it.
- Expect the client to ask where the rest of the team is. The answer is the outcome, delivered.

James Tuttiett · Global Director of Transformation & AI, FDM Group https://jamestuttiett101.github.io/Project1/frameworks.html

GO TO MARKET · A WORKING FRAMEWORK BY JAMES TUTTIETT

The Hackathon as the Front Door

One day, the client's engineers, their data, their problem. Working software and a sponsor by the end of it.

The most effective AI sales motion I have found is not a pitch. It is a day in a room with the client's own engineers, their real data and a problem they actually have, building something that works before everyone goes home. A slide deck produces a follow-up meeting. A hackathon produces working software and a sponsor who has seen it run.

Before: the three things that decide it

A named sponsor with a budget, who will be in the room at the end. A real problem chosen by the client, not by you, that is small enough to make progress on in a day and painful enough that someone cares. And data access and security sign-off sorted in advance, because the day is lost if the first four hours are spent on permissions.

During: small teams, spec first, ship

Mixed teams of client engineers and your consultants, three or four people each. Every team writes a one-page specification before touching code: what it does, what good looks like, what is out of scope. Agent-assisted tooling does the heavy lifting. The target is something that runs on the client's data by the afternoon, not a prototype on synthetic data.

After: the sponsor conversation

The day ends with each team demonstrating to the sponsor. The conversation that follows is not "would you like a proposal?" but "which of these do you want in production, and what would it take?". Scope the engagement from the thing that worked.

What it is not

Not a training day, not a showcase of your tools, not a hackathon on your problem. If the client's people are watching rather than building, it is a demo, and demos do not create sponsors.

How to use it

- Run it inside your own business first. Every rough edge is found by colleagues rather than clients, and you get an internal proof point.
- Keep the format identical across clients. Repeatability is what makes it a motion rather than an event.
- Capture the specification and the working output as the engagement's starting point. The hackathon is day one of delivery, not a pre-sales cost.