

You are not delegating to a model. You are architecting a team.



Dr Peter McCann Strain

CTO, DPhil/PhD in AI from Oxford University

Swipe >>

— THE OPERATING RITUAL

An autonomy ladder still needs owners; a rung with a name but no real authority is decoration. The work is to turn the stack, the evidence, the rungs and the attention into a team that can run the system on Tuesday morning.

— THE REFRAME

A team, not a hire.

THE OLD QUESTION

We shipped Claude. We use GPT. We built an agent.



THE QUESTION THAT HOLDS UP

Which humans hold authority, evidence, attention and responsibility when the system acts? Read the series as a team org chart, not a tool catalogue.

~30 deployed agents

The MIT AI Agent Index (2026 arXiv) documents around thirty deployed agents and uneven public transparency around design, safety, evaluation and third-party testing. The operational read: low documentation is the norm, not the edge case. So naming your stack (the components and the owner of each) is what lets you inspect failures instead of improving by superstition.

SOURCE

Endsley (Human Factors, 2017) on the automation conundrum and Lee & See (Human Factors, 2004) on trust calibration are the peer-reviewed anchors for human-attention discipline; the MIT AI Agent Index (2026 arXiv, accepted at ACM FAccT 2026) is the anchor for the capability-and-transparency picture; the September 2025 FTC companion-chatbot inquiry and Garcia / Raine litigation are procedural records, never stacked as findings of liability.

— CADENCE

Five operating disciplines, walked as three paired roles.

- 01 Run **layer accounting** (name each stack layer) and **error accounting** (measure escaped failure).
- 02 Hold **rung discipline** (an action class is what work an agent may do; a rung is how much autonomy it gets) and capture **runtime evidence** (the trace, the per-run record of decisions and tool calls, treated as a work product that the verifier, the independent checker, judges).
- 03 Treat **human attention** as a scarce production resource, not a magic phrase.

— THE ARTIFACT

Five questions on one morning screen.

Which layer changed?

find the moving part

Which error escaped?

read the harm path

Which rung changed?

demote or justify

Which run is missing?

rebuild evidence

Which review stopped giving signal?

name the owner

Five questions on one screen. All green: ship with care. One red: find the owner, change the system.

— PUT IT ON THE CALENDAR

Pick your most ambitious agentic deployment this week. Assign the five disciplines to named people, one responsible and one accountable per row. Flag any row with no one who can stop, demote or redesign the system within the week.

— RITUAL DRIFT

Folding all five disciplines into one AI lead. A single owner cannot run layer accounting, rung discipline and human attention in the same week. Name one accountable per row before the next launch, or expect the post-mortem sentence: everyone thought someone else owned it.

— RUN THE OPERATING LOOP

Stop Delegating. Start Architecting.

Read the full essay – the argument, the sources, the figures and a reader-ready working artifact.

Substack petermccannstrain.substack.com · Medium @peter.mccann.strain ·

LinkedIn peter-strain-dphil-15a607128

New essays twice weekly, 2 June – 21 July 2026.

— THE STACK SO FAR

E22 · Essay 22 of 22 complete · Arc V: Operating model

YOU JUST ADDED

The Tuesday Dashboard and framework crosswalk

STACK LAYER LIT UP

Whole stack

YOU CAN NOW ASK

operate the whole AI-coworker team through named disciplines.

NEXT

This closes the series. The next move is yours.



Dr Peter McCann Strain

CTO, DPhil/PhD in AI from Oxford University

I build production AI systems and write about making agentic AI useful, inspectable, governable and safe enough for real work.

Follow on Substack for the full 22-essay series
petermccannstrain.substack.com