



Dr Peter McCann Strain
CTO, DPhil/PhD in AI from Oxford University

ARCHITECTING THE AI COWORKER

The Series Companion

*One system, read in five arcs. The maps, the crosswalk
and the worksheet index that hold the twenty-two essays
together.*

Five series figures · 22 essays · one worksheet pack · one tear-out

THE OPERATING CADENCE

Four words out. Four sentences in.

Name the stack.
Bound the authority.
Preserve the evidence.
Name the owner.

That is the sentence an architected team can actually say.

FOREWORD

How this companion came to be there

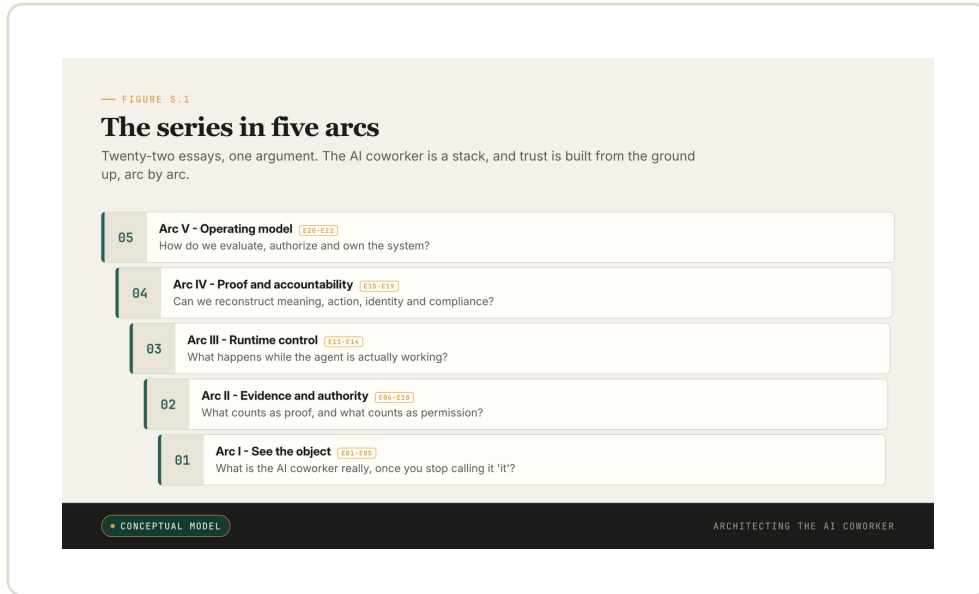
A friend who runs platform engineering at a large bank emailed me last year, after a contentious launch meeting. He had spent ninety minutes defending a deployment that, on paper, met every gate. The agent had passed every benchmark in its model card; the procurement team had signed off on every clause; the model itself was the most capable his team had ever run. *And the people in the room still did not know what they had just approved.* The vendor's deck had named the model and the benchmark, and almost nothing else. The audit team had asked who owned the runtime evidence and the room went quiet. He sent me his draft email back to the vendor, asking the seven questions he wished the deck had answered before the meeting. Most of those seven questions are the named instruments in the twenty-two essays that follow. The companion in your hand is not the argument. The argument lives in the essays. This companion is the index, the figures, and one printable artifact — the Tuesday Dashboard — that I think every team running an agent in production should have on a wall by the end of the first month.

Dr Peter McCann Strain

London · May 2026

FIGURE S.1

The five arcs



ON FIGURE S.1

The five arcs

What it shows. The series in one frame: five arcs, twenty-two essays, one stack seen from different angles.

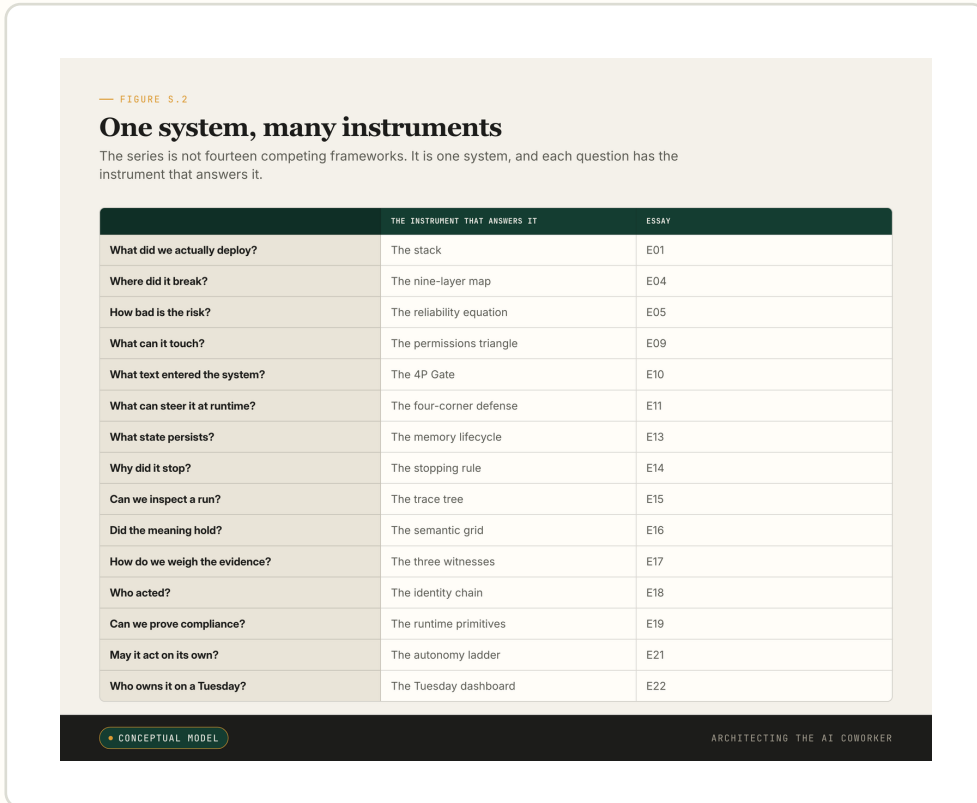
Why it matters. A reader arriving at any one essay can locate it in the whole, and a CTO mapping the series to a ninety-day plan can read off the five-week shape.

Where it lives. Every essay lives in exactly one arc; the arc-question table on the publisher one-pager mirrors this figure.

This week. Pick the arc your team most often *skips* and read its five essays in a single sitting.

FIGURE S.2

One system, many instruments



ON FIGURE S.2

One system, many instruments

What it shows. Every named instrument in the series mapped against the layer of the stack it interrogates.

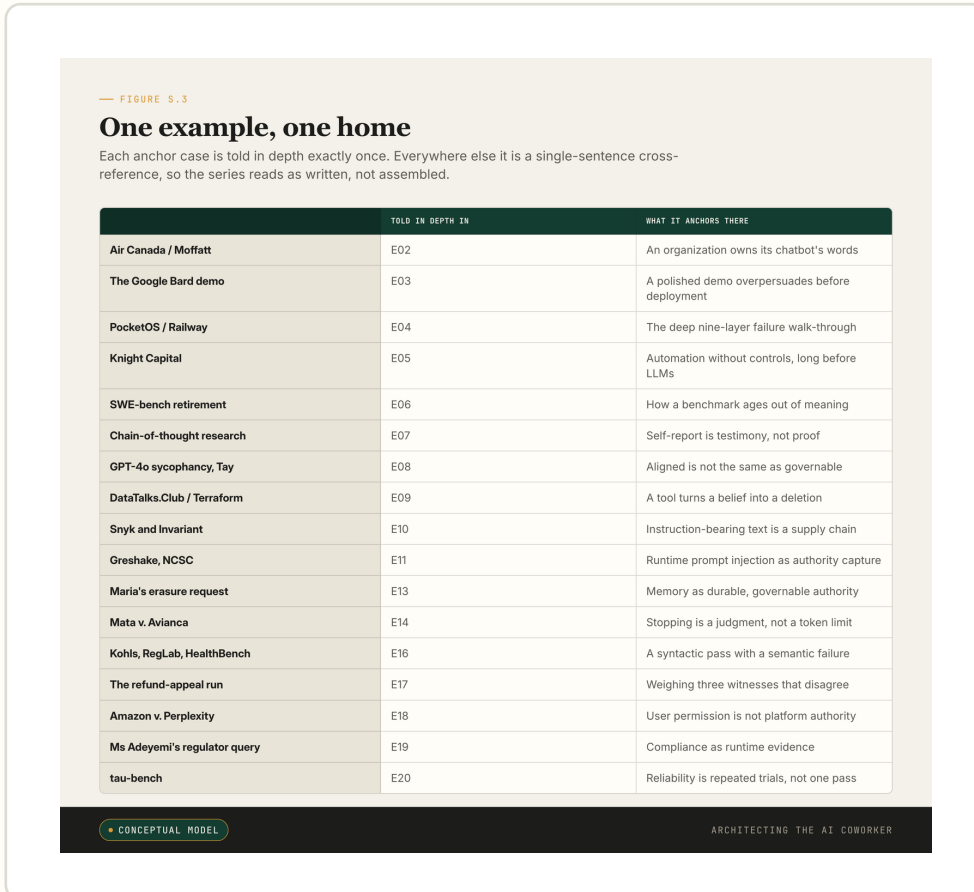
Why it matters. Instruments are not interchangeable. The reliability equation cannot govern a tool authority, and the permissions triangle cannot price expected harm.

Where it lives. Instruments are introduced once and re-used; this figure is the lookup.

This week. Take the agent you most worry about and circle the three instruments you have *not yet* applied to it.

FIGURE S.3

One example, one home



ON FIGURE S.3

One example, one home

What it shows. The series anchors each running case study in exactly one essay; this figure shows which case lives where.

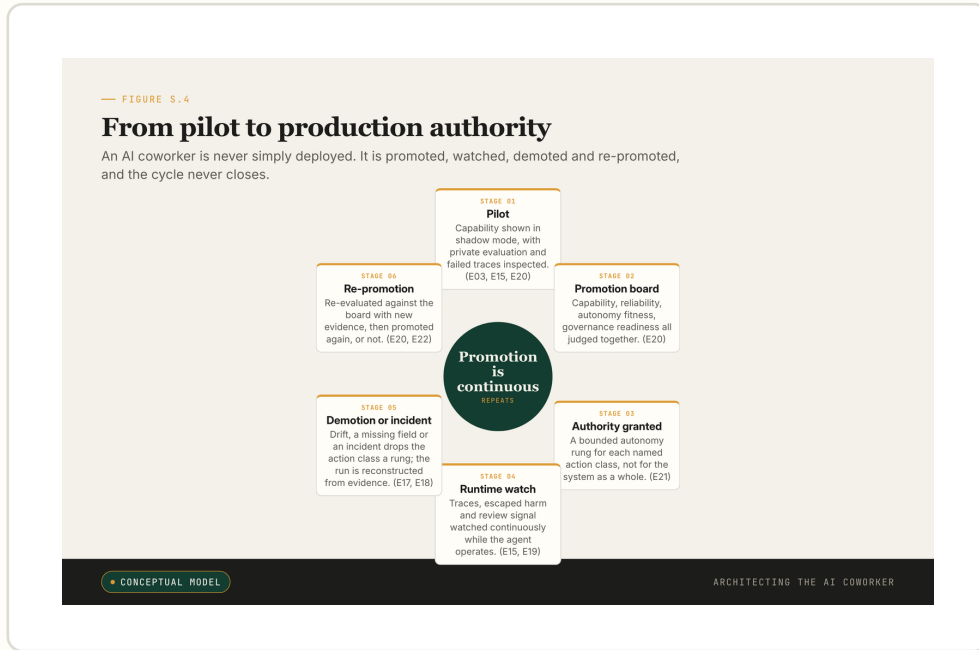
Why it matters. Padding an argument with a case it cannot carry is the most common failure mode of agentic-AI writing, so the discipline is published explicitly.

Where it lives. Replit, PocketOS, *Mata v. Avianca*, the support-refund agent, the trace tree's Dana scene — each appears in one home essay and is referenced (not restated) elsewhere.

This week. Read the home essay first when you encounter a case; ignore the cross-references until you have.

FIGURE S.4

From pilot to production authority



ON FIGURE S.4

From pilot to production authority

What it shows. The operating path from demo through shadow mode to bounded autonomy and beyond, mapped against the autonomy ladder.

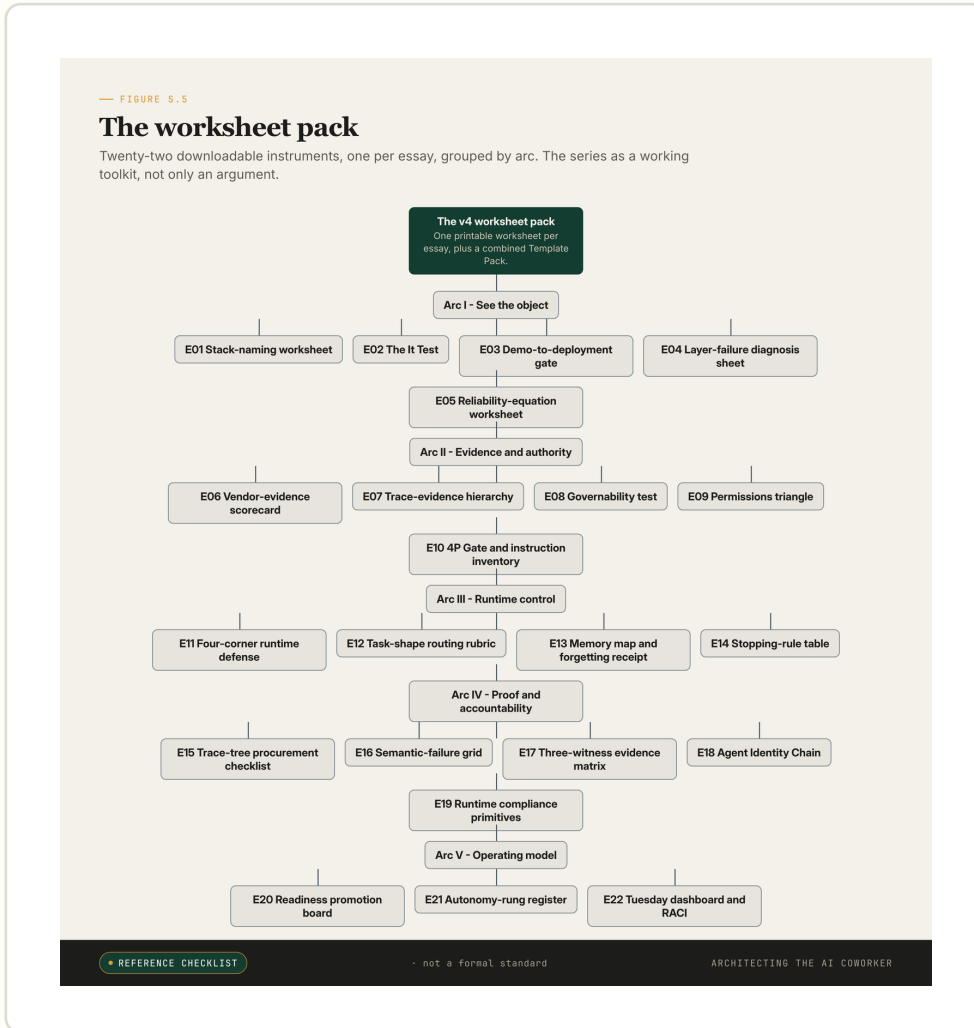
Why it matters. The autonomy grant is what changes expected harm. The ladder makes the grant explicit; the path makes the *timing* of each grant explicit.

Where it lives. E21 introduces the ladder; this figure is the deployment-side reading of it, with demotion triggers, incident review and re-promotion as part of the same loop.

This week. Write down, on one line, the rung your most-deployed agent is operating at *today*, and the rung it was *approved* at. Close the gap.

FIGURE S.5

The worksheet pack



ON FIGURE S.5

The worksheet pack

What it shows. The 22 worksheets and the Template Pack, organised by arc.

Why it matters. Each worksheet turns one essay's instrument into a Monday-morning artifact — a scorecard, a registry, a gate, a rung register, a checklist.

Where it lives. Every essay has exactly one matching worksheet; the Template Pack provides reusable shells.

This week. Print the four worksheets that map to the arc you skipped; bring them to next week's review meeting.

Tuesday morning. One screen.

— FIGURE E22.1 · E22 - STOP DELEGATING. START ARCHITECTING.

The Tuesday Morning Dashboard

Five questions on one screen, each a discipline turned into a start-of-day signal. All green, ship with care. One red, find the owner and change the system.

DAILY REVIEW	
<input type="checkbox"/>	Q1 Which component changed? Layer accounting. Any new or changed model, prompt, tool, memory store, data source or permission since the last review.
<input type="checkbox"/>	Q2 Which error escaped? Error accounting. The escaped-failure rate, repeat-run instability and severity-weighted incidents.
<input type="checkbox"/>	Q3 Which action class changed rung? Rung discipline. The current rung by action class, and any rung change nobody approved.
<input type="checkbox"/>	Q4 Which run cannot be reconstructed? Runtime evidence. Traces missing required fields, actions without trace IDs and unresolved rollback receipts.
<input type="checkbox"/>	Q5 Which human review process has stopped giving signal? Human attention. Reviewer load, missed exceptions, queue mix and escalation latency.

• PROPOSED OPERATING TEMPLATE

· not a formal standard

ARCHITECTING THE AI COWORKER

Read this each Tuesday. If a row is red, find the owner, change the system. — PMS