

THE FRONTIER

What does it take for an organisation to keep becoming capable?

THE FRONTIER

What is your organisation's success quietly using up?

Your AI programme may be working. Drafts arrive sooner, backlogs fall, experts spend less time on routine. This paper asks what else that work used to produce: the practice through which people became able to judge, the routes through which expertise passed on, and the conditions under which an organisation stayed capable when its best people were absent.

The argument in one sentence

Organisations are redesigning work faster than they can see what that work was forming. Responding well needs capability to be treated as something the organisation deliberately forms, evidences, governs and renews, not something it assumes. I call an organisation that does this a Frontier Capability Organisation (FCO).

What this paper contributes

- 1. Capability as a configuration, not a stock.** The ability to achieve an outcome depends on people, practice, tools, relationships, capacity and legitimate authority together. Decisions made on a skills view and a capability view can differ, and that difference is testable.
- 2. Productive versus renewable capability.** A test for any AI or operating-model decision: does it improve today's output while weakening the organisation's ability to reproduce the judgement it depends on?
- 3. Evidence, judgement and authority kept deliberately apart.** People and agents increasingly act together. The paper proposes a chain in which evidence informs, people decide, trust is bounded, and nobody's authority grows silently out of a dashboard.
- 4. Renewal as operating design, not goodwill.** Practice, succession, stewardship and capability debt become explicit decisions with owners, costs and review points.

None of these is proven. Each is stated so it can be tested, and Appendix F says what would weaken or break it.

Who this is for, and what you take away

If you are...	You will leave with...
An executive or board member	A way to ask whether you are becoming more capable or only more productive
A technology or AI leader	A clearer answer to who does, decides, reviews, overrides and remains accountable
A people, learning or operations leader	A view of where capability is forming, concentrating or quietly draining away
A manager or practitioner	A practical route to widen the path through critical work
A partner, consultant or freelancer	A way to create value without engineering dependency
A researcher or sceptic	Clear claims, their status, and the evidence that would count against them

How to read it

Read straight through for the argument. The coloured boxes are optional depth:

Box	What it does
QUESTION	The question that carries you into what comes next
PAUSE + CONNECT	A moment to link an idea to your own work
QUICK EXPLAINER	A term or distinction in a few lines
SIDE JOURNEY	An analogy or neighbouring idea
DEEPER DIVE	More depth for readers who want it
IN PRACTICE	A scenario: red where it goes wrong, green where it goes better
RESEARCH CHALLENGE / HYPOTHESIS	Where the argument is tested or still open
FROM THE NOTEBOOK	Original drawings, preserved as provenance
TAKE IT FURTHER	Where to go next if you want to act on an idea with Empower Your World

CONTENTS

Front door

The Frontier: what is your organisation's success quietly using up?

MOVEMENT 1 · DISCOVER

- I · What work was secretly doing
- II · Inside the excavation
- III · What capability actually is

MOVEMENT 2 · CONSTITUTE

- IV · The questions stopped staying in their lanes
- V · The moment we try to see capability, we change the problem
- VI · The right answer is sitting on the table

MOVEMENT 3 · OPERATE + TEST

- VII · The organisation is already paying for capability
- VIII · Do not start an FCO programme

MOVEMENT 4 · ATTACK + REOPEN

- IX · This could go badly wrong
- X · I do not want to end by pretending we have arrived

APPENDICES

USE IT

- A · Where to begin: entry-route prompts
- B · Experiments to run together
- C · Harbourside, two years on
- D · Tools from the paper
- E · A capability evidence record

CHALLENGE IT

- F · Claim status, propositions and challenge register
- G · Intellectual neighbourhood

GO DEEPER

- H · Where to go deeper next
- I · Frontier Research Library roadmap

REFERENCE

- J · Glossary of terms
- K · Research notes, sources and figure accounts

MOVEMENT 1 · PARTS I-III

DISCOVER

What was the work secretly forming?

What work was secretly doing

Begin with success. Then ask what else the work was producing.

A very successful Monday morning

Imagine a very successful Monday morning. First drafts arrive sooner. The backlog falls. Routine corrections no longer consume the senior reviewer's day. The intervention is doing what the business case promised. Then a new colleague asks a question the dashboard cannot answer: if the system now does the first assessment, where will I learn enough about assessment to recognise when it is wrong?

I want to stay with that question for a moment. Not with a failed system, an alarming prediction or a defence of the way we used to work. With success. The output is better, yet something else may have changed underneath it. The report still appears. The practice that once helped somebody become able to write it may not. Perhaps that practice was mostly wasted effort. Perhaps the new work will teach more. Either way, I cannot tell from the report alone.

That's what drew me into the question. Work produces things, but it can also produce people who become able to do things. We notice the first output because it has a name, a deadline and a customer. The second is less tidy. It accumulates through an awkward question, an explanation that doesn't survive scrutiny, a difficult correction, or somebody seeing why an apparently routine case isn't routine at all.

STATUS · Proposed · C1

A slow task does not become a valuable apprenticeship simply because generations had to endure it. Recopying information and navigating broken systems may teach endurance without teaching much judgement. I am not interested in saving that work. I am interested in finding out what was bundled inside it. If useful practice and pointless friction travelled together, we need to separate them before deciding what to remove, what to replace and what deserves to disappear.

WORK PRODUCES THINGS. IT CAN ALSO PRODUCE PEOPLE WHO BECOME ABLE TO DO THINGS.

What else did the work produce?

I used to think the drawing below was about digital capability. It belongs to earlier work in which we were helping people use low-code tools to improve real problems. At GSK, I helped shape a Power Platform strategy with colleagues, including Darlene Schwartz. We wanted business users to own more of their digital transformation, with technology colleagues as partners. The aim was not to send every local problem into a central delivery queue, nor to leave people building alone. It was to bring knowledge of the work together with the capabilities needed to change it responsibly. [1]

That intention becomes clearer in the practical advice. Find an opportunity in your own role. Start with something small and achievable. Try it. Share the attempt and the block. Ask for help, and help somebody else. A report might improve or a manual process might be automated. Those were the visible results. Alongside them, people could learn to frame a messy problem, ask better questions about data, explain themselves across disciplines, handle risk and decide when what they had built was good enough to use.

The app was one output. People becoming able to do the next piece of work was another. We were making room for people to act, encounter consequences and learn with one another. Courses mattered, but they couldn't supply every real opportunity or every colleague with time to explain a difficult correction. The historical 70:20:10 drawing puts those routes beside one another. I keep it for that relationship, not as a rule for dividing anybody's learning time.




Figure 1. Learning was happening through the work and through other people, not only in the classroom.

Preserved as drawn.

FROM THE NOTEBOOK
Preserved as drawn

The percentages are historical labels. The value here is the relationship between learning, other people and real work.

Hidden did not mean good

The phrase hidden curriculum is useful, but it can make the past sound kinder than it was. Some people learned through excellent supervision. Others learned through avoidable

stress, unpaid effort, exclusion or luck. One person received worthwhile cases and patient correction; another was kept at the edge of the work. The fact that a system produced experts tells us very little about the people it did not help, or the cost it asked them to carry.

The same route can pass on good judgement and bad judgement. An experienced colleague may recognise a subtle risk, or dismiss an account that does not fit the local convention. “This is how we do things here” can contain accumulated wisdom. It can also contain a workaround whose original purpose disappeared years ago. Once I noticed that, preserving the hidden curriculum stopped being a sensible objective. Understanding it became the objective: what should survive, what needs redesign, what should be unlearned, and what should be retired?

To make that less abstract, I use Amira and Graham from my earlier White Paper. Harbourside Group and its people are fictional illustrations, not a client account. Amira learns customer-resolution work by reading case histories, drafting assessments and having her reasoning corrected. Graham can notice when an unusual complaint is not really about the apparent problem. Graham is four years from retirement. Nobody has yet asked who will do what he does. If an AI agent supplies Amira’s draft, Graham may still be essential for the exceptions. Better throughput would not tell us whether anybody is becoming able to do Graham’s work. [2]

There are three plausible readings of that change. The removed drafting may have taught little. Reviewing may teach more. Or the output may survive while exposure, correction and progressive responsibility diminish. The third possibility has to compete with the other two, not win because I’ve told a worrying story. The decision is what arrangement of work to accept, and who can learn, act and carry its consequences.

Simon Wardley’s practitioner challenge, preserved in an endnote to my earlier White Paper, is that technology creates new work and engineering practices as well as removing old ones. That’s an important alternative, not a finding verified here. An AI-supported learner might encounter more varied cases, get faster explanations and reach useful judgement sooner. New work somewhere in the economy doesn’t guarantee that this route is forming here. Somebody still has to look. [3]

Review is not automatically that route. If Amira can inspect the case evidence, explain why she accepts or rejects a draft, encounter convincing mistakes and receive correction, reviewing might teach a great deal. If she is asked only to approve what appears plausible, she may become fluent in approval. Graham’s time matters too. Calling his additional work mentoring does not create capacity. Faster drafts can leave every difficult exception with the same practitioner, while the dashboard makes the redistribution look like a saving.

THE HIDDEN CURRICULUM WAS NEVER AUTOMATICALLY GOOD. IT WAS SIMPLY HIDDEN.

PAUSE + CONNECT

Do not preserve the old work

The question is not whether the removed task was “good learning”. Separate useful practice from pointless friction. Preserve the exposure, correction, judgement or relationship that matters, then redesign where that learning can happen.

Borrowing from the future

At first, I was asking whether useful practice might be lost. That was too small a question. An organisation can retain the practice and still make itself dependent on one person who knows how to interpret it. It can postpone development, obscure a dependence in a partner arrangement, or keep reproducing competence that no longer fits the need. The common difficulty is not simply a missing skill. It is tomorrow's work depending on capability that today's arrangement is making harder to access, renew or recover.

I use capability debt to name that exposure: the future constraint created when a system consumes, concentrates, postpones, obscures or fails to renew capability on which it will later depend. It's a decision lens, not a balance-sheet number. What are we relying on? How is that reliance changing? What would rebuilding involve if the arrangement stopped working? An honest description of a dependence we hadn't seen is more useful than a persuasive total. [4]

QUICK EXPLAINER

Three debts, four postures

Capability debt is the broad exposure. Two forms come under the most direct pressure from AI. Judgement debt builds up where agent activity expands faster than people's capacity to interpret, challenge and govern it: outputs approved faster than anybody can meaningfully review them. Practice debt builds up where removing work also removes the route through which future experts form: people who have completed the training but never done the thing under real conditions, with real correction.

Posture matters as much as type. Intentional debt is a visible, bounded, governed decision to tolerate temporary fragility. Accidental debt arrives without recognition or owner. Structural debt is built into incentives, role design, funding or governance. Compounding debt erodes the capacity to repay itself. Graham is compounding debt: the more difficult work reaches him, the less time he has to form a successor, so even more difficult work reaches him.

One more distinction keeps the conversation honest. Productive capability density asks how much valuable capability a team can enact for its size and cost. Renewable capability density asks whether judgement, stewardship and the ability to develop successors are spread widely enough for continuity. A team can score highly on the first and poorly on the second, becoming exceptionally productive while growing more dependent on people it cannot replace. The question I keep returning to is simple: how much of today's performance depends on capability we cannot currently reproduce? [15]

STATUS · Proposed · C2

Borrowing can be rational. During a crisis, concentrated review may be the safest temporary answer. A rare specialism need not be distributed everywhere. A small organisation may be better served by an external expert than by constructing an internal pathway it cannot sustain. I do not want the word debt to condemn those choices. I want it to make their conditions visible. How long can we carry this exposure? Who owns the decision? How concentrated may it become? What happens if the expert, partner or service is unavailable?

Those questions change the business case without cancelling the gain. Protected practice and a second reviewer cost time. Better evidence access may cost money. Their value might appear in continuity during absence, fewer reopened cases or reduced dependence on emergency help, not in cash saved this quarter. I would not invent those benefits to make the arithmetic work. Compare alternatives, state the time horizon and keep hoped-for value separate from realised value. The temporary exception becomes dangerous when it quietly turns into the permanent operating model.

The wider employment debate can't settle this workplace question. The Department for Science, Innovation and Technology's policy brief Entry-level hiring in the UK: a snapshot, published with LinkedIn on 8 June 2026, considers economic conditions and skills mismatch while leaving AI causality open. That makes the argument less comfortable, usefully so. A labour-market pattern isn't a measurement of how Graham's judgement forms or how Amira acquires it. Neither is a task model. [5]

STATUS · Supported · adjacent phenomenon only; AI causality not established

DEEPER DIVE

What the entry-level evidence does and does not show

CLOSEST TO THE CONCERN

The World Economic Forum, with PwC, describes a risk to be managed.

Its 2026 report on AI and entry-level work observes that AI is changing how early-career workers learn, develop and exercise judgement, and that organisations may have to balance short-term efficiency against capability-building. The language is careful: a risk, not a settled effect.

HIRING SIGNALS

Weakness is visible; the cause is contested.

Adzuna reported UK graduate adverts in July 2026 down 45.6 per cent on a year earlier, a count of adverts rather than hires, with no attributed cause. The government snapshot with LinkedIn found the largest declines in roles where AI has improved, while stating plainly that this is not yet causal evidence. LinkedIn's own analysis attributes current weakness mainly to macroeconomic conditions.

A MORE SPECIFIC MECHANISM

Practice can disappear before skill visibly declines.

Verschuere and Cameron separate substituted from complemented tasks and report that US job starts for 22 to 25 year olds in the most substitutable occupations fell by around a third, while declining to claim definitive attribution. PwC's 2026 analysis of US entry-level adverts found the most AI-exposed roles were around seven times more likely to ask for traditionally senior skills such as leadership and creativity. If that holds, organisations are asking entrants to arrive with judgement that used to be formed on the job.

TAKEAWAY

Consistent with the concern. Not proof of it.

Several of these sources say explicitly that they do not establish AI causality, and that caution travels with the claim wherever it appears in this paper. [5] [15]

I would begin with one consequential workflow. Before changing it, ask what people encounter, what they must explain, which errors are corrected and who supplies the correction. Record exception quality, supervisor time, dependency concentration, time to independent judgement and the people able to provide cover. Then agree a small redesign and decide in advance what would make us revise or stop it. Training satisfaction and throughput can contribute, but neither answers whether somebody can exercise judgement in context.

If the AI-enabled route forms judgement just as well without deliberate reconstruction, I'd have to rethink the concern. If observing development requires surveillance, ranking or the removal of challenge, we've failed. The question I started with, "What can the system do for us?", has become "What will changing this work make possible, or make harder, for the people who must act next?"

PART II

Inside the excavation

Preserve the archaeology. Let questions connect without pretending the destination was inevitable.

I DIDN'T START WITH A MODEL. I STARTED WITH QUESTIONS THAT WOULD NOT STAY APART.

Provenance: The drawings show the enquiry changing shape. They do not prove the destination was inevitable.

Inside the excavation

When I put the Power Rangers diagrams, Simon's Theory of Community sketches and the later Empower Your World (EYW) Capability Ecosystem and AI vs the Brain Muscle notebook pages beside one another, I can see why that question didn't arrive fully formed. Some try to explain learning, others ownership, participation, value or the limits of a technological change. They don't look like chapters of the same argument. A neat origin story would make it too easy to mistake what I now think for what I understood then.

PowerPoint was originally where I did the thinking. Rebuilding the slides every time an idea changed was slow. Drawing on the iPad made it easier to change, connect, abandon and revisit ideas. Then a SEACON deadline arrived and the rough drawings went straight into the slides. I apologised for them. The audience suggested keeping them. That changed how I saw the drawings: the rough artefact helped me think and invite a conversation; a formal model had a different job. Making one more polished wouldn't necessarily make it better at the other. [6]

PAUSE + CONNECT

What I mean by excavation

Follow a question until the next dependency changes the question itself. Then decide whether to dig, act, test or park the rabbit hole. Excavation does not mean FCO was the treasure waiting at the bottom.

I thought these were different problems

At GSK, we could start with any of several questions. Could business teams own useful solutions with technology colleagues as partners? Could governance vary with consequence rather than treat every experiment alike? Could people find real opportunities to develop?

Could a community make help easier to reach? Could a useful local solution travel without being rebuilt? I didn't start with a plan to join all of them.

The partnership question is particularly important. Empowerment did not mean permission to ignore security, privacy, architecture or support. We were trying to bring business knowledge and technical capabilities together, with shared ownership and controls appropriate to the work. A person needs room to attempt something, but the attempt has consequences somebody must understand and own. That tension survived. What changed was my need to distinguish much more precisely between being able to see a problem, being able to judge it and having legitimate authority to act. [1]

Measurement opened another route. Platform reporting could show what existed, what people used and what might need support. Useful visibility, certainly. A usage count still didn't tell me what value had reached whom, or whether somebody had become more capable. The earlier work distinguished hypothesised from actual value. I now see a discipline I'd want in any business case: don't let what we hoped would happen become what we say happened.

I'd thought connection made help easier to reach and good practice easier to spread. Then I noticed the other side. Every question could land on the same generous person's desk. Confidence, proximity and existing relationships could decide who got help. A popular shortcut could travel faster than a careful objection. I hadn't lost faith in community. I'd stopped treating connection as automatically good, or willingness to help as a free resource.

Some drawings didn't survive as answers. Crawl, Walk, Run, Fly can help explain progression; it isn't a mandatory maturity path. Network-growth sketches made a relationship interesting without giving me a mathematical warrant for ever-larger communities. The AI Transformation Flywheel came before I changed the question. I keep those abandoned or parked directions in view so the past doesn't become a trail of clues arranged to make the present look inevitable. [6]

A safe tool didn't make the work safe

Safe AI adoption made it harder to keep these conversations apart. A tool could be approved while nobody knew who could stop its response. A reviewer could be named without the evidence, time or right to intervene. A team could deliver more while leaving nobody ready to take over. These weren't different names for one problem. They were different decisions inside the same work. I began asking who was looking at them together.

Return to the fictional Harbourside complaint. An agent recommends an ordinary response to a case that ought to be escalated. A person approves it. It is reassuring to say that a human was in the loop, yet that tells us remarkably little. Did they understand the context? Could they inspect what the agent had used? Could they stop the response? Were they authorised to change it? A human can touch the process without having meaningful control over its consequence. The comforting phrase can hide the very questions it seems to answer. [2]

QUICK EXPLAINER

Human in the loop does not tell us enough

A human can be present without having the context, time, authority or ability to change the outcome. At this point the important question is simply whether meaningful human judgement exists. Part VI returns to the governance detail.

I have found it helpful to resist joining these distinctions too quickly. Capability to recognise a problem is not authority to act on it. A record is not a judgement. A judgement is not a permission. Good intentions do not supply a working tool, and a working tool does not supply time for correction. Connection is useful precisely because the things being connected remain different. If I collapse them into one reassuring label, I have made the diagram simpler and the decision worse.

Partnership raises another practical question. If a partner supplies the capability we need, what happens when they can't? Working together can help, and it can create dependence and coordination costs. Before celebrating the connection, I'd want to know who can keep the work going and what recovery would take. [7]

Then the tunnels started connecting

Can people learn through changed work? Can the capability reach where it's needed? What warrants confidence? Who can legitimately decide and act? What needs to change when yesterday's competence no longer serves today's need? I use Frontier Capability Organisation to name the attempt to connect these questions. A course catalogue, a productivity dashboard or an approval record can help answer one without settling the others.

The serious sceptic asks: "Does any of this actually require FCO?" Organisational learning, human factors, sociotechnical systems, workforce planning and governance already address much of this ground. FCO doesn't earn its place by being coherent. If existing learning, operating-model, workforce or governance practices reveal the exposure and support the same legitimate decision more simply, they should win. Connecting the questions has to change something we'd otherwise miss. If FCO adds terminology but changes no consequential decision, the sceptic is right.

The earlier work and drawings show how I reached these questions, not that FCO works at scale. Market resonance, my own practice and the first EYW journey aren't independent validation. The professional perspectives used during development were structured challenges, not a panel of real external reviewers. I still need people outside the construction to challenge it. [8]

That's why I want to preserve more than the latest diagram. What was I asking before I changed my mind? What survived? What did I abandon? What haven't I tested? Keeping those distinctions makes a correction less like betraying the original idea and more like learning. A polished version can make an early idea look more certain than it was.

DEEPER DIVE

Why safe AI adoption kept turning into a capability question

START WITH THE ORDINARY FAILURE

A system can be approved and the work around it can still be unsafe.

Return to the Harbourside complaint. The agent is approved. Privacy, security and model risk have been considered. Yet a consequential case is still mishandled because the person reviewing the response cannot see enough context, has too little time, or does not have the authority to stop what happens next. The tool passed governance. The working arrangement did not.

THE CONNECTION

Experimentation exposes governance; governance exposes operating-model choices; changed work exposes capability.

Experimentation asks what AI can do here. Governance asks what should be permitted, owned and controlled. The operating model has to answer who decides, reviews, overrides and remains accountable. Once those roles and workflows change, a further question appears: where will the judgement needed to perform them come from, and how will it renew as the technology changes again?

WHY IT KEEPS RESTARTING

The chain is not a one-off journey.

Each new class of agent, each new capability and each new regulatory expectation restarts it. An organisation that treats AI adoption as a programme with an end date may find itself back at the beginning repeatedly, each time with a slightly thinner base of experienced people to draw on. [16]

TAKEAWAY

Capability is not a training line at the end of an AI programme.

It becomes part of the design of the work. The chain is not a maturity model and may not apply equally everywhere. If established governance and operating-model disciplines already surface the capability consequence and support the same decision more simply, use them.

I started looking back to ask how I'd got here. Now I'm asking what must stay distinct when I connect the questions. Capability, evidence, judgement and authority are close enough to confuse. A manager who mistakes one for another can make a very different decision. Before asking an organisation to become frontier-capable, I need to understand what capability lets us see, and how it can be wrong.

**A POLISHED VERSION CAN MAKE AN EARLY IDEA
LOOK MORE CERTAIN THAN IT WAS.**

What capability actually is

Move from person-centred skill to contextual capability, evidence, judgement and renewal.

What capability actually is

A person can know how to perform an action and still be unable to achieve the required outcome here. They may lack the case history, a working tool, a colleague's help, time, permission or a way to challenge the instruction. Calling that a skills gap sends us towards the person when the obstruction may be somewhere else. I keep coming back to capability because it makes me look at the situation too.

Capability doesn't sit inside one person

Suppose Graham moves to another organisation. Much of his skill travels with him. This customer's history doesn't arrive automatically. Nor do the relationships that let him ask a quick question, the permissions to inspect a case, or the local understanding of when to escalate. Skill remains essential. Capability becomes visible in what Graham, in this situation, can achieve. Asking whether he's skilled is useful; asking what he can reliably achieve here, for whom and under which constraints changes where we look. [2, 9]

Capability keeps that context in view. A single good outcome can show that something happened. It need not establish dependable performance across situations. Equally, a rare specialist need cannot be judged by demanding frequent repetition its context never provides. The evidence should fit the actual demand and consequence. "Valued outcome" also needs a question attached: valued by whom? Efficiently delivering the wrong response does not become legitimate simply because it is repeatable. [9]

STATUS · Architectural · C3

QUICK EXPLAINER

Driving: skill, qualification and capability

"Can drive" is an unhelpfully broad statement. A licence records categories and restrictions. Real capability also depends on vehicle, conditions, experience, practice, consequence, confidence, rules and legitimate permission. Some differences are qualifications; others are judgement or human factors.




Figure 2. "Can drive" hides very different capabilities.

Illustrative synthetic asset · derived for explanation · no readable personal data.

SIDE JOURNEY

The useful question is situated

A regular-car entitlement says nothing about safely driving a lorry. Racing experience does not automatically make somebody the right person to tow in dense fog. Teaching, assessing and maintaining vehicles introduce different capabilities again. What matters is what somebody is trying to achieve, where and under which conditions.

An earlier notebook tried to work through abundance and scarcity. It set data, information, knowledge, expertise, experience and judgement alongside an ecosystem of people, assets and delivery. Knowledge could become abundant while capability remained scarce. I can still recognise the question. What I no longer want is a ladder that makes capability look like the final, largest quantity of knowledge. That would leave out much of what makes useful action possible.




Figure 3. An earlier question: if knowledge becomes abundant, why can useful action still be difficult?

From the Notebook. The original drawing records the enquiry; this readable descendant does not promote the sketch into canonical architecture.

FROM THE NOTEBOOK
Preserved as drawn

An early attempt to understand why abundant knowledge does not automatically make capability abundant.




Figure 4. Look around the outcome, not only at the person: skill is one ingredient in a working configuration.

Formal explanatory descendant, not canonical FCO architecture.

Look at Graham through that picture and one ingredient stands out.

QUICK EXPLAINER

Capacity is a filter

Graham has capability. His capacity is the problem. Capacity is whether capability can actually be brought to bear in the time, place and quantity needed, and it narrows in stages. Potential capacity may exist but not be available. Available capacity may not be accessible to the team that needs it. Accessible capacity may not be agency-enabled, because the person lacks permission or support to act. Even then, it may not be mobilisable in the moment.

Authority sits beside this chain, not inside it. It is granted, never possessed, and capability, whether human or agent, never creates authority by itself. [9]

PAUSE + CONNECT

Potential is not availability

Capability moves through states: latent, available, enacted and renewable. Potential is not availability, availability is not enactment, and enactment is not renewal. Many skills systems measure the first and quietly imply the other three.

Participation is not just willingness

Head, Heart and Hands gives me a simple way into another common misreading. Understanding is not the same as wanting to participate, and wanting to participate is not the same as being able to. Head asks whether the purpose, evidence, alternatives and path

are intelligible. Heart asks whether participation feels meaningful and plausible: whether identity, belonging, recognition and confidence in a possible future have a place. Hands asks whether the opportunity, agency, time, tools, support and manageable burden actually exist. It is a design sequence, not a psychological test. [10]

I find the last distinction especially easy to lose. Somebody may understand the initiative, care deeply about the work and still have a week with no room to attempt it. Another may have time but no accessible tool. A third may have the tool and no safe way to say that the task should not be done. If we jump from missing activity to missing motivation, we prescribe encouragement for a practical obstruction. Another course may help understanding; it cannot repair an inaccessible system or make unpaid development time fair.

There is also the question of working with others. Who helps somebody interpret the situation? Who makes an opportunity reachable? Who provides correction without leaving the learner ashamed to ask again? Who can share responsibility rather than becoming the person who must rescue everything? Relationships affect every part of participation, including whether the people outside the familiar network get a genuine route in. [10]

PAUSE + CONNECT

Head, heart, hands, and with others

Understanding something, caring enough to move, being able to participate in practice and being able to do so with other people are different conditions. Capability can fail when any one is missing.

Those answers need to change what a manager does. Protect time for the person supplying correction. Give newcomers a route that doesn't depend on already knowing somebody. Make it safe to question a shared practice. A cheerful community channel won't do those jobs on its own, and recognition isn't payment in hours.

What can one piece of evidence really tell us?

The moment I try to observe capability, I can create a new problem. Evidence makes some work visible and leaves other work in shadow. If practice becomes a score, people can optimise for the score. If a record becomes a leaderboard, learning becomes competition. If missing evidence becomes evidence of absence, people who were never offered the opportunity can be made to look less capable by the system itself. The purpose of observation has to stay attached to a legitimate decision, not an appetite to make people comparable. [11]

Take one of Amira's reviewed complaints. We might preserve the case context, the evidence she considered, her reasoning, the correction she received and the scope in which she was allowed to act. That tells us about this event. An interpretation may suggest improving reliability at a defined kind of judgement. A later human decision may widen her permitted work within bounds. The record, interpretation and decision are three different things. None should silently do the job of the next. Part V returns to why that separation matters so much once an organisation starts watching.

The omissions matter too. A person may do valuable work that leaves little digital trace, or have good reason not to disclose sensitive material. They may not have been offered a relevant case. They may have exercised judgement by not proceeding, which a completion count could miss. An account of capability that ignores opportunity and visibility can turn an organisational failure into an apparent personal deficit. The person needs a way to correct the interpretation and challenge the use being made of it.

This is not a case for recording every correction. The strongest objection is that a context-rich method becomes a second administrative job. I would start with the minimum account needed for one consequential decision, keep disagreement and stop collecting what does not change that decision. If people perform for the record instead of learning, or paperwork increases without better judgement, the method should be revised or abandoned. Reflection prompts are not validated person measures, and developmental information must not quietly become employment scoring. [11]

You can't learn judgement just by recognising the answer

Recognising an answer is not the same as constructing one. A practitioner has to decide what kind of situation this is, which evidence matters, what is uncertain, what consequence is acceptable and where their own scope ends. Exposure, correction, consequence, reflection and relationships seem important to that development. I treat them as conditions to investigate, not a sufficient recipe. Repeating cases can reinforce a bad convention. Feedback can reward fluency. A good explanation can hide an inability to act when the context changes.

In the Impact of AI essays I described a working heuristic that helps explain why this matters. For perhaps eighty per cent of situations, the pattern is familiar enough that a good procedure gets to a good outcome. The other twenty per cent does not behave like that: the case that passes every check and still feels wrong, or the classification two reasonable reviewers would score differently. The proportions are illustrative rather than measured. The shape is the point. Organisations rightly spent years converting judgement into procedure, and the cost stayed hidden because enough people had come up through the harder version of the work to handle the twenty per cent. AI is now automating much of the eighty per cent that used to be the training ground for it. [1]

Years after the Power Rangers work, I drew a note headed "AI versus the brain muscle". If technology did more of the mental work, were there kinds of practice we'd cease to exercise, and would some be hard to recover? I had a picture and questions. Making the worry visible meant I could interrogate it instead of merely feeling it. [12]

The muscle analogy can suggest a simple decline that the evidence here doesn't establish. I'd rather ask something we can observe: can people still construct a judgement, recognise an exception and act within legitimate bounds when the convenient first answer is absent? That moves me from a curve on a page to a comparison that might tell me I'm wrong.

Before changing the workflow, capture supervisor load, exception quality, dependency concentration, time to independent judgement and successor readiness. Afterwards, examine the same questions in comparable contexts. Selected cases without an AI first draft might help distinguish recognition of a suggested answer from the ability to construct one. Prior experience, case selection and changing difficulty could all distort the result. This is a proposed test, not a comparison already run. The method has to make rival explanations visible rather than arrange a demonstration of its own answer.

Supervisor load is a particularly revealing part of that account. Amira may improve because Graham works additional hours, selects easier cases or rescues each failure before it becomes visible. The pathway could be effective and still unsustainable. Record denied opportunities, adverse effects and the time correction requires. Agree what would make the owner stop or revise the change before looking at results. This is not an objection to supervision. It is a refusal to treat experienced attention as an unlimited input.

My earlier paper drew on Olle ten Cate's work on entrustable professional activities in medical education. Bounded trust and supervision are useful neighbours for this question. Domain specialists still need to ask what travels into AI-enabled service work and what

depends on the professional setting in which those ideas developed. The inherited citations and their verification limits are in the notes. [13]

Better independent judgement, less concentrated dependence and acceptable supervisor load would strengthen the redesign in its tested context. The same improvement without deliberately connecting the decisions would weaken the case for doing so. If the gains disappear when Graham leaves, we haven't established persistence or transfer. Those results must be allowed to change what we do next.

Blog Post - AI vs the Brain Muscle

30 March 2026 15:36




Figure 5. A worry I could now question, rather than just feel.

From the Notebook. The curves record questions; they are not evidence of irreversible decline.

FROM THE NOTEBOOK
30 March 2026

A worry made visible so it could be questioned, tested and changed.

Capability can be wrong

There is a final difficulty I do not want to tidy away. An organisation may become very good at doing something it should stop doing. It may deliver consistently while excluding a group, preserve a once-useful habit after its conditions disappear, or reward a workaround that bypasses a legitimate control. More capability is not necessarily better capability. Propagation magnifies what travels, including error. That is why renewal has to include challenge, unlearning and retirement, not only faster formation. [14]

Unlearning is not an instruction to forget. A practice can be bound to incentives, identity, tool limitations or an old need that has not actually gone away. Removing the habit without changing those conditions can leave people unable to deliver. I would ask why it existed, what outcome it served, whose experience challenges it and what replacement makes safe

action possible. What still works deserves care. Continuity does not, by itself, establish legitimacy.

The people affected need a route to challenge how the organisation sees them. Who defines good capability? Who sees the evidence? Can someone correct an interpretation? Will mentoring and community contribution receive time, or only another expectation? The fictional employee-forum questions in my earlier paper remain uncomfortable in the right way. A system that makes people more observable and less able to challenge has failed its developmental purpose. Permission to collect evidence isn't permission to use it for every later convenience. [14]

A small practical decision is now possible without assembling an entire architecture. Choose work whose outcome matters. Examine the configuration, the practice that forms judgement, the capacity and authority needed for review, and the habits that should not survive. Agree a bounded change with the people affected. Protect correction and challenge. Compare what happens with what we expected, including rival explanations. Completing a course should not automatically widen authority. An absent record should not become a verdict on a person.

I'd change my mind if deliberate formation reproduces unequal access, or useful evidence can't be gathered without excessive burden or coercive visibility. I'd narrow the approach if it works only where supervision is abundant or outcomes are easy to observe. These aren't small implementation details. They're reasons to stop, change the design or limit where we use it.

The cycle underneath all of this

Pull the threads of Parts I to III together and a cycle appears. I call it the Capability Renewal Cycle. It is a public explanatory model, not a specification, and it describes how capability forms and renews for a person or a team.

It starts from a real need, not a course catalogue. The need creates a legitimate opportunity for practice: consequential work with supervision and boundaries. Practice produces evidence of what was attempted and what judgement was required. Evidence can support bounded trust within a declared scope. Trust opens opportunity, and where opportunity carries decision rights, authority is granted separately. Opportunity enables contribution. Contribution generates learning, if anybody captures it. Learning feeds renewal, changing pathways and roles so that capability continues beyond the people who hold it today. The strongest renewal evidence is often a successor who can now do what their predecessor did, with decreasing support. Renewal then reveals the next need. [17]

STATUS · Architectural

A cycle can run and still be unhealthy. Capability is morally neutral: an exploitative system can learn, plan succession and regenerate very effectively. So I read the cycle alongside three flows.

Public explanatory model




Figure 6. The Capability Renewal Cycle: eight stages, with capability, agency and benefit circulating at the centre.

Adapted from the White Paper v1.0 (Figure 9): the note on renewal now sits beside its stage. A public explanatory model, not canonical FCO architecture.

Capability distribution asks who becomes more able, and whether judgement spreads widely enough to reduce fragile concentration. Agency distribution asks who gains meaningful power to originate, challenge, decide and steward. Benefit distribution asks who receives the economic, developmental and reputational value created, and how much of it is reinvested. A system is empowering only when it progressively distributes all three while reinvesting enough for the people who come next. It can be productive, and even self-renewing, while still concentrating authority, burden or benefit. Part IX returns to that risk.

I began by asking what AI could do, on a successful Monday morning. Now I'm asking what the organisation must remain or become capable of doing: for whose outcomes, with what opportunity to learn and act, and what needs renewing or retiring? That's the question Part IV will take into the organisation itself. Before any architecture earns our attention, it has to help people make a better decision. The next useful answer is a small encounter with reality that can answer back.

So this is where Movement 1 lands. Capability does not live wholly inside a person. If that is true, learning, work, evidence, opportunity and authority can't stay separate conversations.

QUESTION

Pull the camera backwards

If capability depends on the whole working situation rather than skill alone, the organisation can no longer leave learning, evidence, opportunity and authority in separate lanes. What happens when those questions collide?

MOVEMENT 2 · PARTS IV-VI

CONSTITUTE

What changes when the excavation becomes an organisation?

The questions stopped staying in their lanes

At what point does a collection of sensible functional answers become one organisational capability problem?

I had been trying to solve a much smaller problem

I didn't set out to invent an organisation. I was trying to answer an ordinary question: how do organisations adopt AI safely?

It behaved badly. Safe adoption needed governance. Governance needed an operating model. The operating model depended on work, and changed work changed where people practised, where judgement formed, who could recognise an exception and who could act when people, technology and circumstances moved.

So I started drawing the journey as an excavation, not because the answer was buried at the bottom but because each shaft changed the question. Some went nowhere. Some exposed another domain. Some, FCC among them, had to be parked. Good exploration isn't digging forever. Judgement includes knowing when to surface and let reality have a turn.

PAUSE + CONNECT

Your own excavation

Think of a question at work that started small and kept exposing the next dependency. Where did you stop digging? Was that a decision, or simply where the time ran out?

Then imagine the successful Monday morning again

The intervention worked. Drafts arrived faster. The backlog shrank. Experienced people stopped spending so much time on routine first passes. This time, instead of staying with the new joiner, pull the camera backwards.

Operations sees productivity. Finance sees a saving. Technology sees a scalable pattern. Risk asks for controls. Learning and development may see somebody who needs training. A manager sees a dependency on Graham. Amira sees fewer chances to make the first judgement herself. None of those views is wrong.

The problem is that they can all be right separately while the organisation becomes wrong collectively.

THE ORGANISATION CAN OPTIMISE EVERY PIECE AND STILL FAIL TO NOTICE WHAT ITS NEW WAY OF WORKING IS MAKING POSSIBLE, OR IMPOSSIBLE.

Why now? Organisations have always muddled through

A fair challenge: none of this is new. Experts have always retired, juniors have always learned on the dull work, and organisations have generally muddled through.

My answer is that muddling through relied on time. The work shifted over a few years and people grew into its new shape, or someone noticed the old routes had gone. What feels different now isn't any single pressure. It's how many are arriving together, and how fast.

DEEPER DIVE

Seven pressures arriving at once

ENTRY ROUTES ARE NARROWING FIRST

AI is absorbing routine, junior and first-draft work fastest: often the very work through which people used to learn. The external evidence is consistent with that concern, though it does not establish AI as the cause (Part I).

CAPABILITY IS MOVING FASTER THAN ROLES

What an agent can be trusted to do changes in months, not years. Knowledge becomes abundant while the judgement to use it, the long climb from data to information, knowledge and wisdom, still has to be formed in people.

THE BUSINESS CASE POINTS AT THE TRAINING GROUND

Under pressure to show a return on rising AI costs, organisations reasonably start with high-volume, repeatable work. That is frequently the practice ground for the next generation of experts.

NOBODY QUITE AGREES WHAT "AI" MEANS

Within one vendor estate alone, a leader may be weighing assistants such as Copilot, agent builders such as Copilot Studio, development platforms such as Foundry, first-party app agents and multi-agent orchestration across platforms. Each carries different governance, operating-model and capability consequences.

GOVERNANCE AND OPERATING MODELS ARE MOVING TARGETS

Controls, decision rights and roles are being redesigned while the technology underneath them keeps changing.

EXPERIENCE IS APPROACHING RETIREMENT

Many organisations rely on experienced people nearing the end of their careers, at the same moment the routes that formed their successors are narrowing.

PEOPLE MOVE TO KEEP UP

Cost-of-living pressure encourages people to change organisation for a pay rise, taking context, relationships and half-formed judgement with them.

Any one of these could probably be absorbed. Together they compress the time muddling through relied on. The question isn't whether organisations have faced capability problems before. It's whether the informal ways they solved them will still be there when they're needed. [25]

MUDDLING THROUGH ALWAYS RELIED ON TIME. AI IS TAKING THE TIME AWAY.

An allotment, not a factory

Before I define anything, here is the picture I reach for: an allotment. Shared land, divided into plots with different purposes. Shared rules protect the common ground. Experiments happen in bounded beds. Someone watches the weather from the top of the hill, and seeds, tools and advice change hands over the fence. Leadership is distributed, and the place is always preparing for another season.

That's much closer to what I mean than factories and pipelines. The observation tower is the Observatory. The experimental beds are bounded experiments. The soil is trust and culture, the thing everything else grows from. The swap table is exchange. The allotment society is governance and stewardship. Nobody owns the whole allotment, yet it renews itself season after season.




Figure 7. The Frontier Capability Allotment: cultivating capability, connecting people, exploring the frontier.

Illustrative explanatory visual. Some labels, such as Culture Spine, Capability Marketplace and Intelligence Engine, come from earlier EYW vocabulary. Appendix D.10 shows the crosswalk from allotment to organisation.

SIDE JOURNEY

Where the allotment stops

Capability does not grow simply because conditions are pleasant. Organisations have formal authority, legal obligations and unequal power. Not every contribution is voluntary, and nobody should be expected to garden for free. Much of the rest of this paper exists to make those boundaries explicit. [17]

Hang on. Haven't we seen all of this before?

Yes, much of it. That point matters.

Organisational learning, dynamic capabilities, communities of practice, human factors, knowledge management, workforce planning, enterprise architecture and governance all illuminate important pieces of this territory.




Figure 8. Capability sits at the centre of an interconnected system of organisational and societal domains.

Simplified orientation map for the main argument. Appendix G sets out the detailed knowledge domains, depth of engagement, perspectives and review considerations.

That picture explains something uncomfortable. In a typical organisation, almost every one of those domains has an owner somewhere: HR owns talent, technology owns platforms, risk owns controls, learning owns development, finance owns value and communities mostly look after themselves. Each can do its part well. In my experience, the capability question sits in the gaps between them, which may be why so many organisations have not yet asked it at all. Nobody's job description includes the joins. [24]

So the research challenge is not “Can I put all of those ideas in one diagram?” I can. The harder question is whether connecting them changes a consequential decision.

What's new, and how would we know?

Here is the honest version of the neighbourhood. Each discipline below already does serious work on part of this territory. The third column is what I think FCO adds. The fourth is how you could find out that it adds nothing. Appendix G gives the longer account of each. [26]

Existing discipline	What it already gives us	What FCO adds	How to test that it adds anything
Dynamic capabilities (Teece, Pisano and Shuen)	How firms integrate, build and reconfigure competences as conditions change	How formation, opportunity, evidence, judgement and authority connect to those moves	Does a configuration view change a decision a dynamic-capabilities analysis would leave alone? (P1, B1)
Communities of practice (Wenger)	Learning as participation in shared practice, meaning and identity	Governance of burden, opportunity, extraction and the route into legitimate action	Does resourced stewardship renew a community without killing its energy? (B9)
Entrustable professional activities (ten Cate)	Bounded trust and progressive entrustment in medical education	Bounded trust in AI-enabled service work, kept separate from authority	Are chain-based entrustment decisions easier to explain and challenge? (P2)
AI risk management (NIST AI RMF 1.0)	Govern, Map, Measure and Manage across an iterative lifecycle	The capability around governance: recognising new situations, bounded authority, renewal	Can named roles make oversight real without a new bottleneck? (B3)
7 Powers (Helmer)	Seven sources of durable advantage	How renewal strengthens or undermines them; no eighth power	Does renewal raise optionality without automatically creating power? (P6)
Wardley mapping	How components evolve, commoditise and retire	Which context, judgement and trust must stay renewable as execution commoditises	Does a renewal view change a build, buy or retire decision? (P1)

RESEARCH CHALLENGE

The novelty test

If an established discipline or operating model already connects formation, work, evidence, judgement, legitimate authority and renewal more simply, use it. FCO gains nothing by renaming territory that is already adequately mapped.

Tests P1 · Break condition B1

External research, practitioner signals and market interest each make the problem more or less plausible; none of them shows that FCO works. Appendix F.2 sets out how I weigh each kind of evidence, and what external and market signals can and cannot tell us.

My current judgement is that there is still a gap worth testing. The domains often meet in the organisation only after a problem has crossed their boundaries. FCO makes the connection itself the object of design. Not because the pieces are new, but because consequences can appear in the joins.

STATUS · Proposed · C4

FCO EARNS ITS PLACE ONLY IF CONNECTING THE QUESTIONS CHANGES WHAT SOMEBODY CAN SEE, DECIDE OR PROTECT.

So what, exactly, is a Frontier Capability Organisation?

A Frontier Capability Organisation is not an organisation that has completed a transformation. It is one that deliberately improves its ability to sense consequential change, form and mobilise capability, connect people and intelligence, create opportunity, act with trustworthy judgement and renew itself as the conditions around it change.

That definition describes organisational abilities rather than a structure. Direction translates strategy into capability questions. Capability visibility reveals where ability is forming, available, concentrated or degrading. Frontier exposure creates supported encounters with consequential work before the highest-stakes moment arrives. Ecosystem connection recognises that no organisation should build everything itself. Judgement and decision keep evidence, interpretation, authority and accountability distinct. Renewal forms successors, adapts practice and retires what no longer serves.

Two foundations sit beneath those abilities. The first is trust designed into how evidence, authority, challenge, correction and remedy work. The second is an information foundation that connects capability to context, enactment and outcome without converting people into scores. Neither can be supplied by a learning programme, a technology platform or a policy document alone, which is why the proposition is organisational.

The definition has two siblings: a public promise to participants and an operating responsibility for the people who run the system. Appendix D.1 sets the three side by side.

RESEARCH HYPOTHESIS

Why connecting the joins should matter

If an organisation connects three things that usually sit in separate functions (how capability forms in real work, what evidence it can legitimately see, and who has authority to act on that evidence), then four kinds of decision should change: how work is redesigned, who is entrusted with what, where investment goes, and when practices are retired. That should happen because fragility that sits between functions becomes visible to someone with the authority to act on it. You would expect to see less dependence on single people, successors forming with decreasing support, and stable outcomes on unusual cases, without person scores or hidden evaluation. If those effects don't appear, or appear just as well without FCO, the claim weakens. (Tests P1–P6; break conditions B1–B7, Appendix F.)

DEEPER DIVE

Projects, products and capability ask different questions

ONE DECISION, THREE LENSES

Should Harbourside automate the first assessment of customer complaints?

A project lens asks whether the change can be delivered safely to an agreed scope. A product lens asks whether the service improves persistently for customers. A capability lens asks what the new arrangement will still be able to do when Graham is unavailable, where Amira will practise judgement, what human-agent supervision the work needs, and how those capabilities will renew.

THE POINT

Capability-led operation extends the other two. It does not replace them.

Bounded change still needs projects. Persistent customer value still needs product ownership. The additional question is what those systems depend on, form, consume or concentrate over time. A project consumes capability and then disbands, often dispersing the learning. A product team builds deep capability that may concentrate in a few people. Both depend on capability they do not themselves renew.

WHAT FCO IS NOT

Not a platform, maturity score, compulsory restructure or replacement profession.

Technology may support FCO characteristics; it cannot manufacture trust, authentic practice, legitimate authority, stewardship or successor commitment. If good learning, workforce, governance, product or architecture practice already connects those decisions adequately, FCO should not add another layer of terminology.

RESEARCH CHALLENGE

Does the capability lens produce a materially better decision?

That is the novelty test. Coherence is not enough. If the same decision emerges with less conceptual machinery, the simpler account should win.

The architecture begins to crystallise

At this point architecture starts being useful, because it can hold distinctions steady while the situation changes. For the organisation as a whole, I use nine movements. They are the organisational counterpart to the Capability Renewal Cycle in Part III: the cycle describes how capability forms for a person or team; the movements describe how an organisation notices a capability question and responds to it.

Movement	The question the organisation is answering
Sense	What is changing, and where might capability be forming, concentrating or degrading?
Frame	Which question is worth investigating, for whom, and what would change if we knew the answer?
Interpret	What does the evidence say, with what confidence, and what else could explain it?
Direct	Who has the legitimate authority to decide, and what bounded response will they sponsor?
Form	Where will people practise and develop the capability we need?
Mobilise	How do we bring existing capability, capacity and partners to the need now?
Propagate	How does what we learned travel to others without being rebuilt each time?
Create value	What outcomes and value followed, for whom, and at what cost?
Renew	What continues, changes, transfers or retires so we can respond again?

It isn't a flywheel promising perpetual acceleration, and it isn't a compulsory workflow. It can slow down, go backwards or be told by reality to stop. Part VII traces one Harbourside challenge through all nine.

QUESTION

The proposition has a consequence

If an organisation intends to renew capability deliberately, it needs to see something about what is forming, concentrating, disappearing or becoming fragile. What can it legitimately see, and what happens when that visibility starts changing people's opportunities?

The moment we try to see capability, we change the problem

How can the organisation learn enough to renew itself without creating a surveillance system or mistaking a dashboard for truth?

“If only we had more data about our people...”

That sentence is seductive. I have spent years wanting better organisational visibility. If we could see the experts, the emerging capability, the dependencies, the people almost ready for more, surely we could make better decisions.

Then put Amira inside the sentence.

Suppose a dashboard says: NOT YET READY. What, exactly, does the organisation think it knows? That Amira has tried the work and cannot yet perform it reliably? That Amira has not completed the expected course? That nobody has observed enough consequential practice? That Graham has not had time to let Amira try? That the team is too busy to tolerate a learning curve? Those are very different claims hiding behind the same label.

PAUSE + CONNECT

A missing record may describe opportunity, not capability

“We have not seen it” is not the same statement as “it is not there”. Before turning missing evidence into a judgement about a person, ask whether the organisation ever created a fair opportunity to produce the evidence.

A familiar people-planning side journey

I have seen organisations use people-planning conversations where performance and future capability are separated, often through something like a 9-box grid. The exercise is usually more human than the grid makes it look: some data, some examples, somebody’s memory of a difficult assignment, somebody else’s confidence, and a fair amount of judgement.

Used well, that kind of conversation can spot a “diamond in the rough”: somebody whose present evidence is incomplete but who may deserve a consequential opportunity to demonstrate more. Used badly, the same conversation can turn thin evidence and reputation into a durable label.

QUICK EXPLAINER

Performance, evidence, potential and capability are not synonyms

Performance asks what happened. Evidence records what can legitimately support a claim. Potential is an interpretation about what might become possible. Capability is contextual ability to achieve an outcome. An opportunity is one way reality can answer back.

That is precisely why capability intelligence cannot simply mean collecting more data. It has to include the quality of the question, the lineage of the evidence, alternative interpretations, the conditions in which the evidence was produced, and a route to challenge the resulting claim.

Then the dashboard starts teaching people what counts

IN PRACTICE

Bad scenario · the score escapes its purpose

A capability indicator becomes visible to managers. Opportunities start flowing towards people with stronger scores. People notice which activities create visible evidence. Quiet work becomes less attractive. Mentoring that is hard to count loses status. Missing evidence starts to look like low capability. Nothing in the original design brief said “change behaviour”. The measure does it anyway.

IN PRACTICE

Better scenario · the question controls the evidence

The organisation begins with one bounded question: are difficult complaints becoming dependent on too few people? It gathers minimum-sufficient evidence about cases, correction, opportunity and concentration. People carrying the work can challenge the interpretation. The evidence is not silently reused for unrelated employment decisions. At review, leaders must decide what the evidence changes.

The difference is not necessarily a more advanced analytics stack. It is a better question, purpose, boundary and consequence. Measurement has entered the system, so the measurement regime has to be governed as part of the system too.

**THE MEASURE IS NOT OUTSIDE THE SYSTEM.
ONCE IT CHANGES ATTENTION, OPPORTUNITY OR
BEHAVIOUR, IT IS PARTICIPATING IN THE
OUTCOME.**

STATUS · Supported · C5

So what is the Observatory?

Chloe, who manages Harbourside's customer resolution team, first tried to understand her team's capability with a spreadsheet: each person, the case types they could handle, their skill ratings. It was accurate on the day she finished it and wrong a fortnight later, when Graham took leave, two analysts moved team and the agent was updated to handle a new category of complaint. The spreadsheet was an inventory. What Chloe needed was a way of seeing movement.

Capability is not a stock that sits still. It moves in seven ways. Formation is capability coming into being through practice and supervision. Availability is capability becoming accessible where and when it is needed. Enactment is capability being exercised on consequential work and producing evidence. Concentration is capability gathering in a few people, usually because it is efficient to route hard work to them. Degradation is capability weakening through disuse, departure, stale evidence or overload. Adaptation is capability changing to fit new conditions. Renewal is capability continuing beyond the people who hold it now. Risk and opportunity live in those movements, not in the inventory.

The Observatory is my name for the capability that sees this movement. It is how a Frontier Capability Organisation pays attention: the distributed capability through which it notices consequential change, frames questions worth investigating, brings evidence and competing interpretations together, supports accountable judgement, follows what happened next and feeds the learning back into renewal. Stewards, practitioners, managers and communities all contribute to it. It is closer to an eye at the centre of a loop than a report read once a quarter. It informs accountable judgement; it does not hold decision authority simply because it sees more.

Its purpose is easy to misread, so I hold it to an explicit hierarchy. The primary outcome is better judgement, not more information. The organisational effect is better decisions. The long-term consequence is renewal. The long-term aspiration is organisational wisdom: learning that survives repeated application, challenge and changing context. An Observatory is not valuable because it gathers more data. It is valuable if people decide better because of it. [18]

For Helen, at the housing association Harbourside serves, the value would be practical. If the Observatory does its job, the first sign will not be a report. It will be the fortnight Graham is on leave and her team's complex cases are still handled on time, with replies that still explain their reasoning. She will know nothing about the Observatory. She will notice that the service has stopped depending on one person's diary.

No shortcut from observation to truth

This is where language that initially looks bureaucratic starts protecting people. Observation is not interpretation. Interpretation is not a claim. A claim is not judgement. Judgement is not an authorised decision.

That chain can be performed quickly. It can even involve the same person. Keeping the distinctions visible gives us somewhere to ask: what did we actually observe? What meaning did we add? What else could explain it? What evidence would change our mind? Who is allowed to attach a consequence?




Figure 9. From evidence to legitimate action: four stages connected by governed boundaries.

Explanatory illustration. Evidence informs, judgement decides and authority enables; none should silently do the job of the next.

QUICK EXPLAINER

From signal to decision

I keep a sequence of things separate, because merging them is how organisations mistake a hunch for a finding. An observation records what was noticed. A signal is an observation judged worth attention. A pattern relates observations. An interpretation offers provisional meaning. A hypothesis is testable. A claim is a bounded assertion with evidence, scope, confidence and a review point. A recommendation is advice. Judgement integrates evidence, values and consequence. A decision is an authorised choice. Action is execution.

The Observatory starts with a question frame: the purpose, scope, affected people, assumptions, evidence need and review point of a question. Before gathering evidence, be clear what it is meant to answer, for whom, and what would change as a result.

At Harbourside, Chloe's first question frame was deliberately modest: how can complex-case judgement keep forming while the agent handles routine drafting? It named the people affected, the outcome that mattered, the evidence that might help and the date the question would come back for review. What came back was not a rating. [18]

PAUSE + CONNECT

Confidence without scores

A useful claim reads like this: strong evidence of enactment in context A, moderate evidence of transfer to context B, weak evidence of successor readiness. That tells a manager what is known, how well, and where to look next, which a single rating cannot. The vocabulary runs from exploratory through plausible, contextually supported and operationally assured to renewal assured. These are confidence statements about claims, never grades for people.

Seeing well is itself a capability, and like any capability it fails in recognisable ways.

RESEARCH CHALLENGE

Seven ways an Observatory fails

Dashboard substitution: the dashboard replaces the questions it was meant to inform. Consensus compression: disagreement is smoothed away. Source laundering: one claim repeated through several channels is mistaken for independent corroboration. Narrative capture: a compelling story outruns its evidence. Decision theatre: the form of evidence-based decision-making without its substance. Memory ossification: past learning treated as permanently true. Signal overload: so much is observed that nothing is attended to.

One warning matters especially for AI. A hundred agent-generated summaries of the same source are not a hundred sources. Synthesis produced at scale is one transformed evidence lineage unless genuinely independent origins are shown. [18]

Relates to B3

Then the darker possibility appears

Have I spent several parts arguing that capability is contextual and human only to build a more sophisticated way of watching everybody? That would be quite an achievement.

It is possible. Visibility distributes power. What the organisation can see can affect who gets opportunity, who is trusted, who gets challenged, who is invested in and who becomes dependent upon whom. A capability-intelligence system that ignores this is not neutral. It is simply ungoverned.

My boundary is therefore stronger than “be ethical”. Developmental evidence should not quietly become a universal employment score. Missing evidence should not default to deficiency. People should be able to challenge consequential interpretations. Observation should remain attached to purpose and context. Sometimes the correct answer to “could we infer this?” is “we should not”.

So here is where the second movement turns. An organisation can't renew capability it can't see moving. The moment it starts to see, it changes who gets opportunity, who holds power and how people behave.

DEEPER DIVE

What information foundation would an FCO actually need?

START WITH THE DECISION

The purpose is not to build a universal profile of everybody.

Begin with the consequential question: where is complex-case judgement becoming concentrated, what evidence would justify a wider route, or which capability will a strategic change require? Collect the minimum needed to improve that decision and keep it challengeable.

TAXONOMY

Name the capability areas without pretending the labels are the capability.

A taxonomy gives people a navigable vocabulary. It helps route questions and find pathways. It does not prove that a person or organisation can act.

ONTOLOGY

Keep different kinds of things different.

The ontology separates capability from methods, artefacts, outcomes, assessments and pathways, and relates evidence, context, opportunity, authority and value. That shared grammar stops a course, activity or job title silently becoming a capability claim.

PLATFORM ROLE

Name, observe, interpret, route and learn.

A capability intelligence platform can support that pattern. EYW is the first AI-capability domain through which it can be tested; the platform is not the FCO itself.

BOUNDARY

Shared meaning is not universal visibility.

Domain relevance does not create permission to profile a person. No covert observation, universal score, automated adverse decision, inference of deficiency from absent evidence, ungoverned repurposing or unchallengeable record.

PAUSE + CONNECT

Where Part V leaves us

Responsible evidence makes better judgement possible. It does not make the decision, and it does not manufacture the authority to act.

PART VI

The right answer is sitting on the table

What turns evidence and judgement into legitimate action when people and agents can both act at speed?

Eleven seconds

A draft appears on your screen. The customer is disputing a charge. The agent recommends a standard goodwill payment and a polite letter closing the matter. It reads well. It is probably right.

You can't see the case history, only the draft. There are forty more in the queue, and the team's target is to clear them by lunchtime.

You have about eleven seconds.

Do you approve it?

PAUSE + CONNECT

Before you read on

Be honest about what you would do, and why. Then ask a harder question: if it turned out to be wrong, who would answer for it?

Who decided?

Six months after the drafting agent arrived, that is exactly what happened at Harbourside.

The complaint wasn't really about a charge. It was a potential regulatory breach that should have been escalated. An analyst approved the draft. The customer received the payment and a letter that implicitly closed the matter. Weeks later, the issue surfaced and the review asked the obvious question.

Who decided?

The analyst said the agent had recommended it. The agent's owner said the analyst had approved it. The analyst's manager said the AI team had designed the approval process.

Everyone had touched the decision. Nobody owned it.

A HUMAN WAS IN THE LOOP. NOBODY WAS IN CONTROL.

This is the uncomfortable truth behind one of the most reassuring phrases in AI governance. “Human in the loop” tells you that a person was present. It tells you nothing about whether they could see the evidence, had time to think, had the authority to refuse, or would be supported if they did.

QUICK EXPLAINER

Which human? Which loop? Doing what?

A human may review a recommendation, authorise an action, supervise delegated execution, remain available for escalation, or simply click approve. Those are radically different control relationships hiding behind one reassuring phrase.

“Human on the loop” can be useful language when a system acts within delegated bounds while a person supervises and can intervene. The phrase only helps if the delegation is intelligible: who granted it, for what purpose, with which limits, what evidence justified it, what requires escalation, and how quickly it can be revoked.

Five questions nobody had answered

Harbourside’s incident wasn’t a technology failure. The agent did what it was allowed to do. The failure was that five questions had never been answered:

- Who does the work?
- Who decides?
- Who reviews, with what visibility and time?
- Who can override, and how quickly?
- Who remains accountable?

Harbourside had answered the first. The other four had been left to assumption, and assumption is where accountability goes to disappear.

The principle underneath is simple, and I think it will matter more every year: agents may exercise delegated authority. They may not originate accountability. Every consequential action has to trace back to a named person or legal entity who answers for it.

Five roles answer those questions: the accountable party, who answers for the outcome; the human authority holder, whose authority the agent exercises; the agent principal, on whose behalf it acts; the agent operator, who runs it; and the agent itself. One person may hold several roles, but they stay distinct, and the absence of a named person doesn’t erase the accountability. [20]

	Acts	Decides	Reviews	Can override	Accountable
Agent	Drafts, classifies	No	No	No	Never
Agent operator	Runs and maintains	No	Monitors	Can stop	For operation
Reviewing practitioner	Reviews the case	Within scope	Yes	Yes	For review
Human authority holder	Delegates	Yes	Samples	Yes	For delegation
Accountable party	Owens workflow	Sets policy	Assures	Yes	For outcome

Agents may exercise delegated authority. They may not originate accountability. Every consequential action traces to a named human or legal accountable party.

Figure 10. A human-agent responsibility map for one workflow.

Reproduced unchanged from the White Paper v1.0 (Figure 15). Illustrative allocation for one Harbourside workflow.

IN PRACTICE

Failure scenario · Who decided?

An approval process that everyone touched and nobody owned. The goodwill payment closed a case that should have been escalated, and the review found accountability spread so thinly that it had effectively vanished.

IN PRACTICE

What Harbourside changed

Harbourside named an accountable owner for the workflow. It separated goodwill recommendations from escalation recommendations, and required a human decision, not just an approval, for anything with a possible regulatory dimension. Reviewers got the case history, not just the draft. It also recognised that reviewing agent output well is itself a capability that has to be formed and renewed.

A few weeks later, Helen asked her account manager who was accountable when AI helped decide something about one of her tenants. This time Harbourside could name a person.

Go back to your eleven seconds. Would anything on that list have changed what you did?

Then the right answer arrives

Now run the story the other way. This time everything goes right.

Graham spots the pattern. Amira can explain why the ordinary response is wrong. The agent surfaces the relevant evidence. The control catches the exception.

Excellent. We have fixed the knowing problem.

The right answer is sitting on the table. Who is actually allowed to do something about it?

A GOOD ANSWER IS NOT PERMISSION. CAPABILITY IS NOT AUTHORITY. EVIDENCE IS NOT A DECISION.

That distinction sounds obvious until work speeds up. Expertise creates influence. Good evidence creates confidence. A reliable agent creates pressure to delegate more. Slowly, being able to do something and being allowed to do it collapse into the same shortcut. FCO keeps them apart on purpose. Capability is credible ability in context. Judgement integrates evidence, uncertainty and values. Authority is the legitimate right to make or approve a consequential decision. Good design connects them; good governance makes the connection inspectable.

STATUS · Architectural · C6

Amira's first real decision

Chloe believes Amira is ready to handle some complex cases on her own. Graham mostly agrees.

In most organisations, that decision would be made informally, on reputation, and would be hard to explain or challenge later. Graham is four years from retirement, and nobody has yet asked who will do what he does. So this decision matters more than it looks.

An FCO makes it through a chain. Each step matters because it stops one thing from quietly turning into the next: a definition of the capability, evidence of attempts, a claim no broader than the evidence, a human decision, bounded trust, and a review.




Figure 11. From evidence to bounded entrustment: evidence informs, people decide, and trust does not automatically create authority.

Reproduced unchanged from the White Paper v1.0 (Figure 14).

Read the chain slowly once. The capability definition says what the capability is, in which contexts, for what valued outcome. An evidence event records an attempt: what was tried, what happened, what judgement it needed. A claim is never broader than that evidence. An authorised person then decides; the evidence doesn't. Bounded trust follows, within a declared scope and never as a number. Where an opportunity carries the right to decide or act, authority is granted explicitly and separately. Review closes the loop, so trust can widen, narrow, lapse or be challenged. [19]

IN PRACTICE

Amira's bounded entrustment

Chloe extends trust for Amira to handle complex billing complaints on her own, below a stated threshold, and not cases involving bereavement or vulnerability. Graham reviews a weekly sample. There's a review date in three months. Amira can see the evidence and challenge it. If she does well, the scope widens. If something goes wrong, it narrows, without becoming a permanent mark against her. It isn't a promotion and doesn't touch her pay; those are separate decisions with their own criteria.

One of her first cases comes from Helen's housing association. Amira spots that the dispute is really about a transfer that was never completed, and says so in the reply. Helen has no idea it was handled by someone in their first year. She notices the answer makes sense.

That's what a successor looks like as it forms. It's small. It's bounded. It's also the first sign that Graham's knowledge might outlast Graham.

QUICK EXPLAINER

What counts as evidence, and how it is kept

Different evidence does different jobs. Declared evidence is what people report about themselves: valuable, and easily biased. Demonstrated evidence is observed practice in context: strong, specific and costly to gather. Outcome evidence matters, though outcomes have several causes. Relational evidence about help, review and blockage is revealing and especially open to surveillance misuse. Governance evidence records the authority and accountability someone has been entrusted with; recognition is not proof of current competence. Renewal evidence, continuity through handover or departure, is the strongest and the slowest to emerge. Stress and counterfactual evidence shows what happens under absence, overload or changed assumptions, and has to be designed and read with care.

The record rules matter as much as the classes. Evidence is appended and superseded, never silently overwritten. Every record carries a review date, a confidence statement in words and a right of reply. Purpose limits use: evidence gathered for development does not quietly become evidence for evaluation. No combination of evidence is ever collapsed into a single number about a person. Appendix E sets out a minimum capability evidence record. [19]

Not every decision deserves the same weight

There is an equal and opposite failure. If every action needs the same senior person's approval, the control system simply builds a new bottleneck, and Graham becomes it.

DEEPER DIVE

How much control is proportionate?

SEVEN FACTORS

Consequence, vulnerability, scale, automation, reversibility, novelty and boundary crossing.

Where these are low, light controls and light evidence may be enough. Where they are high, stronger controls, fuller evidence and independent review are needed.

CLASSES OF CONSEQUENCE

From informational to irreversible.

An illustrative ladder runs from informational, through advisory, consequential and reversible, consequential and hard to reverse, to irreversible or rights-affecting. An agent summarising a public document for internal use sits at the light end. An agent issuing a payment to a vulnerable customer sits at the heavy end.

A FLOOR PROPORTIONALITY CANNOT GO BELOW

Challenge, exit and remedy remain.

Approving a system is not authorising a specific action in a specific business state. Authority before action; delegation before automation; override before autonomy; remedy before scale. An agent's own logs are evidence, never acceptance.

Agents turn out to be useful stress tests here. They force authority that was always implicit into the open. Humans have always worked through delegated authority. Agents just make the boundaries impossible to ignore.

Then a case refuses to behave

All of this works when the organisation knows what kind of situation it is facing.

Then, on a Tuesday afternoon, one arrives that doesn't fit.

Every check passes. The categories almost fit. The agent follows the rules correctly. Amira reads it twice and something is wrong. She can't say what. There's no field for it, no escalation route that matches it, no procedure that describes it. The rule isn't missing. The description of the situation itself has started to wobble.

RESEARCH HYPOTHESIS

The Recognition Boundary

There are two different gaps. In the first, the situation is known but the response is uncertain: improve the procedure, the practice or the escalation. In the second, the situation itself hasn't yet been recognised or named. A procedure can't help yet. People have to compare, question, reframe and learn before anyone can decide what should become repeatable.

Proposed · relates to P5




Figure 12. The Recognition Boundary, presented as a research hypothesis rather than a proven fact.

Explanatory illustration of a research hypothesis, not an established finding.

This may be too neat. “We didn’t see that coming” isn’t proof of a capability gap. Missing data, poor interfaces and ordinary procedural failure can all produce surprise. The idea earns its place only if it helps diagnose different problems and leads to better responses.

Still, it explains something I keep seeing. Organisations write runbooks for the situations they understand, then automate the runbooks. That is exactly right for the 80 per cent. At the frontier, writing the runbook too early can harden the wrong frame around a situation nobody yet understands. The people who can recognise that a situation is new are exactly the people whose formative practice we’re automating away.

RULES CANNOT TELL YOU WHEN THE RULES NO LONGER APPLY.

An organisation needs both disciplines: formalise what has earned stability, and keep enough human and collective capability to recognise when reality no longer fits the formalisation.

What I’m claiming, and what I’m not

I think the case is now strong that capability, evidence, judgement and authority can’t be managed as separate concerns, and that AI makes connecting them urgent because it changes who does the work and how fast consequences travel.

I haven't shown that FCO is the only way to connect them. I haven't proven the Recognition Boundary. I haven't shown that capability can be made visible enough to be useful under strong protections against surveillance. Nor have I shown that organisations will fund the people who can see past the rules when delivery pressure rewards the people who follow them.

None of that is a reason to stop. It's the reason to make the next problem small enough for reality to answer back.

How the pieces fit together

Every model in this paper answers a different question. None of them is the FCO on its own; together they describe one system.

Question	Model	Where
What must the organisation be able to do?	Six abilities: direction, capability visibility, frontier exposure, ecosystem connection, judgement and decision, renewal; resting on two foundations: designed trust and an information foundation that never turns people into scores	Part IV
How does capability form and continue in people and teams?	The Capability Renewal Cycle: need → practice → evidence → trust → opportunity → contribution → learning → renewal, with capability, agency and benefit circulating	Part III
How does the organisation respond to a capability question?	Nine organisational movements: sense, frame, interpret, direct, form, mobilise, propagate, create value, renew	Part IV, traced in Part VII
How does it notice and learn?	The Observatory: question frames, confidence-bearing claims and distributed attention that informs but never holds authority	Part V
How does evidence become action without becoming power?	The evidence-to-entrustment chain: definition → evidence → claim → human decision → bounded trust → opportunity or authority → review	Part VI
Who answers when people and agents act together?	Five questions, five roles: agents may exercise delegated authority but never originate accountability	Part VI
Where does all this sit?	The cross-section: visible work, capability movement and the enabling infrastructure beneath	Part VII

AN FCO USES SIX ABILITIES TO RUN NINE MOVEMENTS, KEEPS THE RENEWAL CYCLE TURNING IN PEOPLE AND TEAMS, AND STAYS LEGITIMATE THROUGH THE OBSERVATORY, THE EVIDENCE CHAIN AND CLEAR ACCOUNTABILITY.

Back at Harbourside, it's ten past four.

Amira is still looking at the case that doesn't fit. The agent has drafted a perfectly reasonable response. The queue is growing. Graham is in a meeting.

She knows something is wrong. She can't name it.

The rulebook says approve.

QUESTION

Part VII widens the lens

What Amira does next depends on far more than Amira: on whether anyone protected the time she needs, on whether Graham can be reached, and on whether the organisation is paying for the capability to notice this or quietly spending it. What happens when you look at the whole organisation?

MOVEMENT 3 · PARTS VII-VIII

OPERATE + TEST

What happens when the model has to survive contact with reality?

The organisation is already paying for capability

What changes when formation, mobilisation, community, value and renewal have to work together every day?

What Amira did next

Amira doesn't approve it. She holds the case, writes two lines on why it feels wrong and adds them to the file. Holding is the easy part. Harbourside's rule since the incident is that when an agent's recommendation conflicts with a person's judgement, the agent doesn't get to escalate its own authority: the case waits for human review. [20]

The hard part is that the only person who can tell her whether she's right is Graham, and Graham is in a meeting. He comes out at twenty to five, reads the case once and agrees with her. Then he mentions, almost in passing, that he's off for a fortnight from Monday.

Amira did the right thing. The organisation got lucky that Graham's meeting finished before the queue did.

Scene one • Graham takes a fortnight off

The complaint is resolved. Graham catches the difficult cases. Amira is learning. The agent handles routine work. Leaders can see the evidence and the delegation is bounded. If I stopped the story there, the organisation would look reassuringly capable.

Then Graham takes two weeks off.

Cases queue. Amira can handle some but not all. The agent continues because its approved scope has not changed. A manager quietly borrows somebody from another team. The dashboard reports a temporary throughput issue. None of those responses is irrational. Together they reveal something the successful day concealed: capability existed, but too much of it was available through one route.

By Thursday, Chloe has a whiteboard covered in names, and most of the arrows point at the same empty desk.

**A SYSTEM CAN SUCCEED REPEATEDLY WHILE
TEACHING ITSELF A DEPENDENCY.**

Formation is not mobilisation

QUICK EXPLAINER

Being capable is not the same as being able to contribute

Capability may exist and still remain stranded. Somebody can have relevant experience but no time, permission, relationship, tool, team or route into the work. Formation asks how capability develops or becomes accessible. Mobilisation asks how it is assembled and brought into legitimate action now.

This distinction annoys simple workforce planning. If the organisation has trained twenty people, it is tempting to say the capability exists. Perhaps it does. If every important decision still lands with Graham, the operational system is telling a different story.

The reverse can happen too. An organisation can mobilise the same trusted people so efficiently that success concentrates future opportunity around them. Being reliable creates more work. More work creates more evidence. More evidence creates more trust. The people outside the loop remain “not yet ready” because the system keeps sending reality to those who are already proven.

Trusted people receive consequential work. Consequential work produces evidence. Evidence produces more trust. Unless opportunity is designed deliberately, yesterday's reliability becomes tomorrow's concentration.




Figure 13. The reinforcing opportunity loop and the deliberate wider route.

Explanatory illustration.

The person everybody calls

IN PRACTICE

Reliable today · fragile tomorrow

A difficult issue appears. The same expert is called because the expert is fast, trusted and usually right. The outcome is good. The next issue goes to the same expert. Then the next. The organisation is not failing. It is teaching itself that only one route is safe.

IN PRACTICE

Use some work to widen the route

Some cases become deliberate opportunities for supported participation, correction and successor development. The expert remains valuable but stops being the only path through the problem. This is not automatically the right answer. It makes the decision about concentration, burden and renewal explicit.

The expert may reasonably object. “I am already doing the difficult work. Now I am expected to create my successor as well?” That objection matters. Calling the additional burden mentoring does not create capacity. A renewable system has to fund correction, review and handover rather than treating experienced attention as an infinitely reusable public good.

There's a second answer to Graham's objection, and it's about Graham rather than the organisation. He does not have to become a manager to be recognised as more valuable. People can advance through deeper expertise, broader application, wider decision scope, facilitation, stewardship and successor development, community leadership, human-agent orchestration, or standards that others adopt. If developing successors counts as part of Graham's contribution rather than a tax on

it, his expertise becomes a route through the work instead of a bottleneck in it. Capability evidence can inform those decisions. It never automatically decides pay, promotion, discipline or redundancy. Appendix D.2 sets out the fuller career model. [21]

That's what scene one resolves. Formation is not mobilisation: capability can exist and still be stranded behind one person's diary, and success will keep narrowing the route unless somebody deliberately widens it.

Scene two • Sam's inbox

Sam opens the laptop at quarter past seven. Forty-three unread messages in the AI community channel. Twelve begin with "quick question". Three are from people Sam has never met, who were told "ask Sam, Sam knows". Sam does know. That is the problem.

I believe in community. Much of my own learning and opportunity has come through people sharing, challenging, connecting and making room for others. That is exactly why I do not want to romanticise it.

A community can spread capability across boundaries that a formal hierarchy struggles to cross. It can also route every awkward question towards the same generous people. It can create belonging and opportunity. It can also turn enthusiasm into unpaid moderation, mentoring, translation and emotional labour.

Communities become renewable through a recognisable progression: participation, contribution, stewardship, renewal. People observe, learn and ask. Some contribute by answering questions, sharing assets and running sessions. Some of those become stewards, taking ongoing care of purpose, trust, boundaries and succession. A community renews when stewardship passes to new people without the community losing its capability. The most important line in the architecture on this may be that a steward's highest achievement is to create conditions in which the community no longer depends on them.

Three risks recur. Burnout, because stewardship is work and it concentrates on the most committed people until they stop; open-source maintainers and Wikipedia administrators both show the pattern. Gatekeeping, because informal status can decide who is welcomed and who progresses. Extraction, because an organisation that benefits from contribution is tempted to treat it as free. Repeated reliance on voluntary contribution should trigger a decision: fund the capacity, narrow the scope, transfer the responsibility, pause or stop. [22]

SIDE JOURNEY

When "thank you" becomes the funding model

Recognition matters. It is not the same as time, authority, workload protection or money. A community that produces organisational value while relying on invisible stewardship is not costless. It has moved the cost somewhere the budget cannot see.

The FCO question is therefore not “Do we have a community?” It is “What capability is forming and moving through this community, who is carrying the burden, who receives opportunity, and what keeps the whole thing healthy enough to continue?”

IN PRACTICE

Sam’s community, before

Sam coordinates Harbourside’s internal AI community. It began as five people swapping prompts at lunchtime and now has several hundred members and a monthly showcase. Most questions still come to Sam. The same dozen people run every session. New members ask a few questions and drift away. Need and practice are strong; evidence is weak; trust and opportunity are concentrated; learning lives mostly in Sam’s head. There is no route from member to contributor to steward to successor.

IN PRACTICE

Sam’s community, resourced

Three decisions changed it. Sam’s stewardship became a recognised part of the role, with a time allocation. Two co-stewards were developed and given the same recognition. A rule was agreed: when a community activity becomes something the organisation depends on operationally, such as supporting production AI systems, it moves into a funded service with an owner rather than remaining voluntary. Voluntary participation stayed voluntary. Work the organisation relied on was resourced.

Sam still answers questions. The difference is that answering them is now part of the job, with time set aside for it, and two other people can answer them too. Community is infrastructure. When the organisation depends on it, the organisation pays for it.

Scene three • Janet pulls the camera back

Janet, Harbourside’s chief operating officer, books a room for an afternoon with no slides. On the wall go six sheets of paper: customer resolution, cyber incident response, a new product bet, the AI community, a specialist in claims who retires next spring, and a partner contract up for renewal. “Every one of these,” she says, “is a capability question we have been answering separately, in different meetings, with different budgets.”

Different worlds need different maps

Capability also refuses to live inside one universal hierarchy. The language that makes sense in medicine may not fit software, social care, sales, education or community stewardship. Even inside one organisation, different domains carry different evidence, consequences and traditions of practice.

QUICK EXPLAINER

A Domain Mesh is not one giant ladder

The aim is not to force every capability into one maturity scale. A mesh lets different domains keep legitimate local meaning while making important relationships visible: where capability transfers, where extra evidence is needed, where authority changes and where people can move or collaborate.

That sounds less tidy than a universal taxonomy because it is. The gain is not aesthetic. It is the possibility of connecting fields without pretending that identical labels mean identical capability.

Before the afternoon ends, Janet asks everyone in the room what they would expect to notice if Harbourside really were renewing capability. The new joiner, the expert, the manager, the partner and the customer all answer differently, and their answers pull against one another. That tension is the point: the proposition is only serious if it makes those conflicts discussable and gives someone the authority to resolve them. Appendix D.3 sets the eight perspectives side by side.

Now look at the whole organisation


The Amira and Graham thread has taken us far enough to see the mechanics. A Frontier Capability Organisation has to do more than repair one workflow. It responds to several capability needs at once: today's customer service, tomorrow's cyber risk, a strategic product bet, an experienced practitioner approaching retirement, a community carrying knowledge across boundaries, a partner supplying scarce expertise and agents changing the division of work.

Most descriptions of an organisation show its visible structure: the organisation chart, the process map, the product list. They show where work happens. They do not show how capability moves, or what makes that movement possible. I find a cross-section more useful. The top layer is visible work: projects, products, services and customer interactions. The middle layer is capability movement: formation, mobilisation, propagation, entrustment and renewal. The bottom layer is the enabling infrastructure: the Observatory, the Domain Mesh, trust and its protections, shared meaning, the operating model's roles, rhythms and capacity, community and stewardship, and the technology platform.

I call that bottom layer capability infrastructure: the connected human, organisational, semantic, trust, intelligence, community and technical conditions through which capability can be formed, found, applied, shared, assured and renewed. Like roads and water, it is noticed mainly when it fails. Time for mentoring, routes for practice, trusted relationships, clear authority and safe experimentation usually stay invisible until a crisis shows they were missing. Each element has to earn its place by serving formation, mobilisation, evidence or renewal, and each needs an owner.

[20]

Most of what makes an organisation capable sits below the surface



A simplification of layers, not an organisation chart.

Figure 14. The FCO cross-section: visible work, capability movement and the enabling infrastructure beneath them.

Reproduced unchanged from the White Paper v1.0 (Figure 16).

The cross-section is a simplification of layers, not a literal organisation chart. Visible work also forms capability, and the layers interact constantly.

The nine organisational movements introduced in Part IV become concrete when one real challenge is traced through them. This is how Harbourside responded to the drafting agent. It is an illustration of how the mechanisms connect, not a compulsory workflow.

Movement	What happens at Harbourside	Mechanism involved
Sense	The Observatory notices rising dependency on Graham and falling entry-level practice after the agent is introduced.	Observatory; signals
Frame	A question is agreed: how can complex-case judgement keep forming while the agent handles routine drafting?	Question frame
Interpret	Evidence shows strong productive capability, weak renewable capability and an accidental capability debt.	Evidence; confidence statements
Direct	Janet, the COO, sponsors a bounded response and names an accountable owner.	Operating model; decision authority
Form	A supervised practice route: selected cases reach analysts without an agent draft, with reasoning reviewed.	Practice pathways
Mobilise	Graham's review time is protected and a second reviewer developed; a delivery partner supports the redesign with a transfer plan.	Capacity; mobilisation; partner
Propagate	The redesign is shared through the AI community and adapted by the claims team.	Community; Domain Mesh
Create value	Unusual-case quality improves; dependency falls; customer outcomes are monitored.	Outcome, impact and value
Renew	Amira is entrusted with a bounded scope; Graham has two developing successors; the practice route is revised.	Evidence chain; renewal

From one capability question to an organisational portfolio

An organisation never has only one capability question. It is forming some capabilities, borrowing others, retiring several and discovering new needs at the same time. The strategic question is not “How do we become capable?” It is “Which capability must exist, where, in which configuration, by when, and through which route should we obtain and renew it?”

The route changes the system. Building creates practice and possible differentiation, but demands time and successors. Buying can be efficient, but needs retained judgement and an exit route. Borrowing buys speed, and risks dependency. Exchange and partnership distribute capability and value across boundaries, and need reciprocity and transfer. Automation changes the practice through which people learn and the supervision through which action stays legitimate. Redesign may remove the need. Retirement frees attention by ending what no longer serves.

Route	When it may fit	The FCO question
Build	Strategic, judgement-rich or central to future direction.	What practice, evidence, capacity and succession will make it renewable?
Buy	Mature service; little need to differentiate internally.	What retained judgement, governance and exit capability remain essential?
Borrow	Temporary or exploratory demand.	How will insight transfer, and how will dependency stay bounded?
Exchange	Reciprocal sharing across teams, communities or organisations.	What is given and received, and is the exchange fair over time?
Partner	Value and capability created across organisational boundaries.	How are responsibility, learning, benefit and renewal shared?
Automate	Stable, governable work where human formation is preserved elsewhere.	What practice disappears, what supervision appears, and who stays accountable?
Redesign	The work can change so the capability is needed differently or less.	What does the redesigned work still require people to learn?
Retire	The capability or practice no longer serves direction.	What memory, transition, protection and replacement are required?

The portfolio decision is economic, but it is not reducible to price. Cost, speed and quality matter. So do option value, reversibility, concentration, supervision, customer dependency, transfer and recovery. FCO does not prescribe the answer. It makes the consequences of the route harder to hide. Appendix D.4 explores strategic advantage and how anyone would know an FCO is healthy.

SIDE JOURNEY

Total Football: connected capability without erasing specialism

THE USEFUL PART

The whole system can reconfigure around the situation.

The Dutch style of the early 1970s is remembered for players moving fluidly between positions and reading the game collectively. Specialist positions still existed. The team's capability was the ability to reconfigure while keeping its shape.

WHERE THE ANALOGY STOPS

Not everybody does everything.

Fluid mobilisation is not permission to remove specialist depth, ignore capacity or blur accountability. A goalkeeper does not become a striker because the system is adaptable. The challenge is flexible composition with clear roles, support and decision rights.




Figure 15. The Frontier Capability Organisation: the allotment translated into an organisation, with each part mapped to an FCO mechanism.

Illustrative explanatory visual; some labels use earlier EYW vocabulary, such as Culture Spine, Capability Marketplace and Intelligence Engine.

By the end of the afternoon, Janet's six sheets of paper have become one portfolio. Each question has a route, an owner and a review date, and somebody can explain why that route and not another. That's what scene three resolves: an FCO isn't one workflow repaired well. It's every capability question in the organisation answered on purpose.

Scene four • The budget meeting

The finance director has one slide and one question. "You want protected review time, a second reviewer and paid stewardship for the community. That is real money this year. Show me where it comes back." Nobody in the room thinks the question is unfair.

So far it is easy to nod. Of course we should invest in practice, successors, stewardship, memory and renewal. Then delivery pressure arrives and the argument meets the budget.

An FCO is not always cheaper. Sometimes the deliberate formation of capability costs more now. Mentoring takes time. A second reviewer looks inefficient. Preserving reserve capability can resemble paying for inactivity. Learning capture delays the next task. A partner exit route costs money before anybody wants to leave.

Part I called this capability debt, and borrowing can be rational. The budget question is narrower: what future constraint have we accepted, for how long, and how would we recover?

This reframes economics without pretending everything can be reduced to a neat return-on-investment calculation. Today's value matters. So does whether today's arrangement leaves the organisation more or less able to create value tomorrow.

THE ORGANISATION IS ALREADY PAYING FOR CAPABILITY. THE CHOICE IS WHETHER IT PAYS DELIBERATELY THROUGH RENEWAL OR LATER THROUGH FRAGILITY, DEPENDENCY AND REWORK.

Value, impact and the commercial question

Janet tries first. "Capability is value we haven't used yet. Value is capability we have used. If we keep spending the second without topping up the first, sooner or later we run out of both."

"That's a metaphor," says the finance director. "Where does the money actually move?"

It's a fair challenge, so Janet walks it through. Capability becomes available. Something connects it to a need. Action creates outcomes, impact and value. Some return is recognised and reinvested. Future capability is preserved or grown. The engine can fail at every join: capability left stranded, activity that produces nothing anyone values, good outcomes credited to the wrong cause, one party capturing the value while another carries the burden, reinvestment that quietly vanishes under delivery pressure.

"And I suppose everything counts as value now."

"No. That's the trap." Naming several kinds of value is meant to show who gives and who receives, not to count everything as a benefit. Meaning doesn't replace fair pay. Community contribution doesn't make infrastructure free. Capability growth doesn't justify an intervention whose cost can't be sustained. [22]

Value surface	What may improve	Evidence worth examining
Revenue + growth	New offerings, markets and faster mobilisation.	Adoption, pipeline, time to market and realised revenue.
Productivity + cost	Less reinvention, delay, rework and bottleneck dependence.	Cycle time, cost to serve, reopen rates, throughput and hidden review time.
Customer value	Better outcomes, continuity, confidence and less engineered dependency.	Outcome quality, remedy, repeat demand and customer independence.
Capability resilience	Faster formation, transfer, recovery and successor readiness.	Coverage, concentration, recovery time and performance during absence.
Workforce + opportunity	Stronger routes into practice, mobility, contribution and retention.	Opportunity distribution, proficiency, moves, succession and avoidable attrition.
Risk + trust	Earlier challenge, clearer authority and safer human-agent action.	Incidents, overrides, unresolved dissent, audit findings and remediation.
Intangible + strategic	Reusable assets, organisational memory, relationships and option value.	Reuse, learning transfer, response speed and future choices made possible.

The seven value streams behind those surfaces, each seen at Harbourside, are in Appendix D.5.

Capability ROI: a decision discipline, not an attribution machine

"Fine," says the finance director. "If I fund it, how will I know it worked?"

"Partly," Janet says. "Late. And without being sure what caused it."

That is the honest answer. A supervised practice route may improve Amira's judgement, and so may her own reading, a good manager and a quiet year for complaints. Claiming the route caused a specific financial return would be false precision.

So Capability ROI is a sequence of questions rather than a formula. What was invested, including the output forgone during practice? Which capability changed, in which configuration and context? What evidence supports that, with what confidence? Which outcomes and value followed? Which dependencies, risks or future options changed? What else contributed? Compared with what alternative? What should now be continued, narrowed, stopped, outsourced, automated, partnered or redesigned?

Value doesn't arrive in one form, so the answer is read through six lenses, from financial and operating through customer, capability, human and ecosystem to risk and social value. Appendix D.6 sets them out.

Three disciplines keep it honest. Separate association from cause. Expect delayed effects: a review that looks only at the first quarter will systematically undervalue renewal and overvalue automation. State the counterfactual. Before investing, state what evidence would show the investment is working and what would make you stop. Harbourside's practice route carried both: within two quarters, analysts on the route should handle a defined category of complex cases with review, with no rise in escalations and Graham's direct caseload reduced; if customer outcomes on routed cases deteriorated materially, the route would pause for redesign. [22]

RESEARCH CHALLENGE

The counterfactual is part of the business case

Doing nothing is not a zero-cost baseline. It may preserve output today while concentration, practice debt or recovery cost grows.

Equally, an FCO-shaped intervention is not automatically the right alternative. Ordinary management improvement, better product ownership, workforce planning, improved tooling or a simpler training route may achieve the same outcome with less burden.

A credible investment decision compares plausible routes over an explicit time horizon, keeps hoped-for value separate from observed value, and states what evidence would stop the investment.

Tests P3 and P6

For Harbourside, the investment includes selected cases completed without an agent draft, Graham's protected review time, a second reviewer, changes to case selection and the information needed for review. The visible benefit might be stable or improved complex-case outcomes. The less immediate benefit might be another credible route through the work. The costs include slower throughput, reviewer burden and governance overhead. The decision is not whether renewal is morally good. It is whether this configuration creates enough present value and future option value to justify those costs compared with the alternatives.

What would you do? You're the finance director. The cost is real and lands this year. The benefit is uncertain and lands later. Do you fund it, and on what condition?

IN PRACTICE

A supplier rewarded for reducing dependency

Helen's housing association agreed to share two simple measures: how often a complex case had to be reopened, and how often her own team could resolve a similar query without calling Harbourside at all. The second is an unusual thing for a supplier to track. It rewards Harbourside for making its customer less dependent, which is exactly the point.

The finance director looks at the slide for a long moment. Then comes the question that changes the meeting: "What would doing nothing cost us in three years?"

Nobody can answer it precisely. Everybody in the room can see that the answer isn't zero. That's what scene four resolves: doing nothing is a decision with a price, and the business case has to show it.

What actually changes

Stories make this feel real. They can also hide how little has been specified. So here is Harbourside before and after, in the plain terms an operating model uses: roles, rhythms, decision rights, working artefacts and money. None of it is a new department. Most of it is a change in who is named, when people meet and what gets written down. [27]

	Before	After
Who owns a capability question	Nobody. Questions surface in different meetings with different budgets, usually when something breaks.	A named decision owner for each question: Janet for customer resolution, the security lead for incident response, the product sponsor for the market bet.
Who answers for an agent workflow	Everybody touched it; nobody owned it.	A named accountable owner; a human authority holder whose bounded delegation the agent exercises; an operator who runs it.
Who defines capability	Implicitly, in job descriptions and in Graham's head.	Domain stewards hold versioned definitions, with affected teams taking part and a route to challenge.
Who reviews	Graham, around his caseload.	Reviewers with protected time and the case history, a second reviewer in development, and reviewing agent output treated as a capability in its own right.
Who carries the community	Sam, voluntarily.	A community steward with a time allocation, two co-stewards, and a rule that work the organisation relies on becomes a funded service with an owner.
Rhythms	Capability comes up when someone resigns, an incident lands or Graham goes on holiday.	Capability-question framing when a change is proposed; evidence reviews on the dates every claim and trust decision carries; capability-debt reviews against named owners and dates; portfolio decisions that include renewal; experiment reviews at each loop; renewal decisions to continue, narrow, deepen, scale or stop.
Decision rights over people's scope	Entrustment by reputation, informal and hard to challenge.	Bounded trust decisions made by a named person, scoped, dated, open to challenge, and kept separate from pay and promotion.
Decision rights over agents	Whoever clicks approve.	A human decision, not an approval, wherever a regulatory dimension is possible; hold and override routes; authority before action and delegation before automation.

	Before	After
Working artefacts	A skills spreadsheet that was accurate for a fortnight; a throughput dashboard.	A question frame; an evidence record; a bounded trust decision; a dependency map; a capability debt profile; a decision record; a health profile for teams and capabilities, never for people.
How money flows	Productivity gains are booked. Review, mentoring and stewardship are unfunded and invisible.	Renewal is funded alongside productivity change. Every capability investment carries a proof obligation and a stopping condition. Reinvestment is an explicit decision, and stewardship and review time are budget lines.

Read the right-hand column again and notice what isn't there: no maturity score, no capability index, no platform at the centre. The change is mostly in what the organisation decides on purpose that it used to decide by default.

What the four scenes add up to

Taken one at a time, each scene looks like an ordinary management problem. A key person goes on holiday. A generous colleague is overloaded. A chief operating officer notices that her budgets do not talk to each other. A finance director wants to see a return. Organisations have handled all of these for years. Put side by side, they show the same pattern from four angles.

In every scene, the organisation was productive and fragile at the same time. Graham's team delivered while depending on one route through the work. Sam's community spread knowledge while consuming the person who held it together. Janet's six capability questions were each well managed and collectively invisible. The budget meeting could see the cost of renewal but not the cost of its absence. This is the gap between productive and renewable capability from Part I, now at the scale of a whole organisation.

It is also where the distinctions of Movement 2 stop being academic. Formation without mobilisation leaves capability stranded. Mobilisation without formation concentrates it. Community without resourcing turns generosity into a hidden cost. A portfolio without a capability view chooses routes on price alone. An investment case without a counterfactual treats doing nothing as free. Each mistake is reasonable on its own. Together they produce an organisation that looks capable right up to the moment it needs to be.

So the claim of Movement 3 is narrower and stronger than "invest in people". Every organisation is already making capability decisions, mostly by default, through how it designs work, routes difficult cases, funds communities, chooses suppliers and judges business cases. FCO asks whether those decisions should be made on purpose, with the evidence, authority and time horizon they deserve.

The finance director agrees to fund a hundred days, not a programme. One condition comes attached, and it's the right one: "Write down now what would make you stop."

QUESTION

Architecture now owes us evidence

A system this interconnected cannot be validated by admiring the diagram. What is the smallest experiment that would let reality disagree with it?

Do not start an FCO programme

How do we make an enormous organisational proposition small enough for reality to answer back?

My first implementation recommendation is: do less

A large proposition invites a large transformation programme. New workstreams. New governance. A maturity model. A steering committee. A platform. Possibly a logo.

I think that would be exactly the wrong first response.

FCO has not earned the right to reorganise anybody. It has earned the right to ask sharper questions about one consequential piece of work and see whether those questions improve a real decision.

MAKE THE PROBLEM SMALL ENOUGH FOR REALITY TO ANSWER BACK.

Scene five • A hundred days, from the inside

This scene goes back in time. Amira's entrustment in Part VI didn't begin with Chloe's decision. It began months earlier, with an experiment that very nearly didn't start.

Return to the AI-assisted first pass. Instead of debating whether automation destroys or improves capability in the abstract, choose one real workflow. Understand what happens now. Capture the visible output and the less visible conditions: difficult cases, corrections, reviewer time, opportunity, dependency, escalation and independent judgement.

Then change one bounded thing. Perhaps Amira completes selected cases before seeing the agent's draft. Perhaps Graham reviews reasoning rather than only output. Perhaps the agent must expose uncertainty or ask for escalation in a defined class of case. Preserve a way to stop. Decide in advance what result would make the team narrow or abandon the change.




Figure 16. A bounded experiment: real work, one change, evidence and a clear decision.

Explanatory illustration of the route Part VIII follows, from consequential work to a decision to continue, narrow or stop.

“Start small” is not enough implementation guidance. A small experiment can still carry hidden employment consequences, customer risk, unpaid supervision or evidence that escapes its purpose. The practical route needs a sequence that protects legitimacy while producing a decision.

After Day 0, the work moves through five loops over roughly the first hundred days. Frame chooses the consequential question and valued outcome. Observe establishes what is known, unknown, claimed and evidenced. Intervene changes one bounded condition. Learn examines outcomes, burden, distribution, rival explanations and unintended effects. Renew decides what transfers, changes, stops or receives further investment. The loops overlap and can repeat; later evidence can send the team back to its frame. Appendix D.8 shows what each loop asked and produced at Harbourside.

Day 0 · The gate that nearly failed

The tempting experiment is simple: give selected complaints to Amira without the agent's draft and compare the results. It sounds small and sensible. It still changes who carries risk, creates developmental evidence about a person, consumes Graham's attention and may affect customers.

So before Day 1, Harbourside has to answer six things. Who is the accountable decision owner, and who is affected? What is the experiment for, and what will its evidence never be used to decide? Whose time does it consume, and is that time protected? Who sees the evidence, and for how long? How can anyone challenge it, stop it or recover if it goes wrong? What benefit, burden and risk does everyone expect, written down before the result arrives?

Five of the six have answers quickly. The third doesn't. Graham has no protected time. The plan assumes he will review Amira's reasoning around his caseload, which in practice means on top of everything else.

What would you do? The experiment is technically ready. The sponsor is keen. Do you start anyway?

Harbourside didn't. Day 0 took three weeks. The employee forum's questions were answered before anybody was asked to take part, and two hours a week came out of Graham's queue. The gate would also have failed if Amira's developmental evidence could slip into performance management, if customers carried a risk nobody had acknowledged, or if nobody had practical authority to pause the change. Day 0 is complete only when the people carrying the consequences can understand and challenge the arrangement. It isn't bureaucracy before learning. It's the minimum boundary that keeps learning from becoming extraction, hidden evaluation or unmanaged consequence. [28]

Day 1 · Amira's first Monday without the draft

Monday, nine o'clock. Twelve cases arrive in Amira's queue with no agent draft attached. Graham has blocked out two hours on Wednesday to go through her reasoning, not just her answers. Chloe has written down, before anyone sees a result, what would make her stop. It feels oddly small for something with so many words behind it. That is deliberate.

Day 30 · Early signs

A month in, the early signs look good. Outcomes on the routed cases are stable. Amira's written reasoning is getting sharper: she has started to say why a case is unusual, not just that it is. Throughput has dipped more than anyone expected, which the team records rather than explains away.

Day 60 · A result we hoped not to get

Then the review-time figures arrive.

IN PRACTICE

A result we hoped not to get

Output remains fast. Amira's explanations improve. Graham's review time doubles. Difficult cases still concentrate around Graham, who is now both expert and training infrastructure. The experiment has not failed. It has revealed that the apparent gain depends on a hidden input the business case did not price.

There's a second, quieter finding. Two analysts are on the route. The evidence for the first is moderate; for the second it is weak, because fewer suitable cases reached her. The analyst with more suitable cases is also the one Graham finds easier to work with. Nobody designed that. It is exactly the kind of unequal opportunity that Part IX will worry about.

Day 90 · The honest decision record

On Day 90, Chloe writes the decision down, including the parts that are uncomfortable to sign.

IN PRACTICE

What an honest decision record could say

Continue the practice route for one more cycle in customer resolution: difficult-case reasoning improved and customer outcomes stayed stable.

Redesign the supervision model: review burden doubled and is not sustainable.

Do not scale to the claims team yet: the evidence and authority boundaries are different.

Commission a simpler comparator: improved case review without the wider FCO mechanics. If it achieves the same outcome with less burden, use it.

Three things change at once. Stewardship of the practice route moves from Graham alone to a shared arrangement with Chloe. A second reviewer begins development. Case selection is made fairer between the two analysts. For the first time, Graham is developing successors rather than simply carrying the work. [28]

The experiment does not validate the whole FCO. It can test whether the integrated questions improve this decision. A positive result in customer resolution does not establish transfer to cybersecurity, social care or another organisation. Scaling needs a new legitimacy check, because consequence, participants, evidence and authority may all have changed.

PAUSE + CONNECT

A useful experiment does not have to produce a happy ending

It can strengthen the proposition, narrow it, reveal another dependency, show that a feared capability loss does not occur, or demonstrate that the intervention creates more burden than value. The result is useful if it changes what we do next.

Where organisations begin

Harbourside came in through AI. That's common, not compulsory. There are at least thirteen legitimate starting points, from risk and governance to succession, community and resilience, and Appendix A offers a first prompt for each. Harbourside's own route changed twice: it became succession-led when Graham's position surfaced, then community-led when Sam's community was drawn in. The route matters less than the direction.

Empower Your World and Journey 01 are where I let the idea become inconvenient

Empower Your World and Journey 01 matter because they make parts of this argument tangible. A participant has to reason through AI adoption, risk, bounded decisions, evidence and what happens when the rules run out. Assets can be opened, tested, misunderstood, improved or rejected. That is far more useful than keeping FCO safely coherent on paper.

QUICK EXPLAINER

What is Journey 01, and what can it prove?

Journey 01, Make AI Adoption Safe, is the first Empower Your World Journey: a practical route through six connected questions between an AI idea and governed use. It moves through a Front Door, Risk Size, Control Floor, Living Loop, Evidence Chain and Response Net. Here I use it to make FCO questions tangible and testable.

It is a reference implementation, not proof. Empower Your World and Journey 01 are one experimental pathway through the FCO proposition. If they work well, they provide evidence about that implementation in that context. They do not prove FCO universally. If they expose confusion, burden or weak transfer, the architecture should learn from it. The public Journey 01 materials are open to explore at empoweryour.world.

STATUS · Experimental · C9

This is also where commercial reality can become intellectually useful. A paid engagement does not have to be disguised validation. It can fund a consequential experiment with explicit boundaries, evidence, alternatives and stop conditions. The ethical test is whether the engagement is designed to learn, including results that weaken the offer, rather than to manufacture a case study.

AI is the first laboratory, not the whole world

I keep using AI because it is where this enquiry became urgent, where authority and judgement are being redesigned quickly, and where EYW already provides a medium for bounded experiments. That consistency helps the reader follow the argument.

It also creates a risk. A sufficiently AI-shaped paper can claim to be domain-agnostic while never leaving the domain that created it.

SIDE JOURNEY

Remove AI from the room

Take a non-AI capability problem: a customer service depending on one experienced practitioner, a community whose stewardship rests on two volunteers, or a partner relationship with no credible exit route. If the FCO questions still change the decision, domain independence becomes more plausible. If they do not, the proposition should become narrower.

Change the work, not only the tool

The practical move is wider than deciding what AI should do. Ask what the new arrangement does to practice, exposure, relationships, correction, judgement, authority and accountability. The tool may be excellent. The surrounding work may still be incapable of using it responsibly or renewing what it consumes.

That is the real purpose of the bounded experiment: not to prove the architecture, but to make the work answer questions the architecture cannot settle alone.

What success would look like

It would be easy to end here with a maturity model. I deliberately do not. Success is described through observable behaviours, each tested against a simple question: how would the people the organisation serves experience it? The customer test guards against an FCO becoming an internal

exercise, and against a subtler failure, an organisation becoming more capable while making its customers less so. Appendix D.9 lists twelve characteristics, each with what you would observe and how a customer would experience it.

After two years, Harbourside could point to seven of the twelve in at least one part of the organisation, with evidence of mixed strength. It could not yet show fair distribution of opportunity, where monitoring had revealed a pattern it was still addressing, or customers becoming more capable, which it had not seriously attempted. That honest partial answer is what success looks like early on. Anyone claiming all twelve after two years should be asked for the evidence.

So the practical advice is simple to say and hard to follow. Do not begin with an FCO programme. Begin with one consequential problem small enough for reality to answer back.

One line in the Day 90 record got less attention than it deserved. The claims team, hearing good things, had already started using the practice route informally: no Day 0, no protected review time, nobody deciding what the evidence could be used for. The idea had begun to travel. Its safeguards hadn't.

QUESTION

Reality gets a vote

If the proposition can be tested, it must also be allowed to lose. What are the ways FCO-shaped ideas could make an organisation worse?

MOVEMENT 4 · PARTS IX-X

ATTACK + REOPEN

Where does the model fail, and what remains unresolved?

This could go badly wrong

What happens when attractive language about capability, community, evidence and renewal becomes a cover for surveillance, extraction or bureaucracy?

The architecture has accumulated enough virtue words to become dangerous

Capability. Trust. Community. Evidence. Opportunity. Renewal. Put enough of those words on a page and almost any intervention can begin to sound benevolent.

That is not a theoretical communications problem. Systems built in the name of development can rank people. Communities can hide labour. Governance can protect itself rather than people. “Temporary” capability debt can become a permanent excuse. Experiments can be designed so that only the preferred answer counts as learning.

Part IX therefore has one job: turn the proposition against itself before anybody else has to.

A story from two years on

Picture a Harbourside that adopted every word of this paper and none of its discipline.

The launch went well. There was a capability dashboard, a community of practice with its own logo, and a register of capability debt with owners and review dates. Eighteen months later, a restructuring needed criteria. The capability profiles were sitting right there, so they were used, just this once. Nobody announced a ranking system. People simply noticed which kinds of work produced visible evidence and stopped volunteering for the quiet kind.

Sam was thanked by name in the annual report and left three months later. The capability debt for Graham’s team had its review date moved six times. Every agent recommendation now carried a human approval, recorded in an audit trail, given in an average of eleven seconds. And a consulting case study described the whole thing as a successful FCO transformation, with a testimonial from a sponsor who had moved on before the results came in.

Nobody did anything wicked. Every step was reasonable on the day it was taken. That is what makes this kind of failure dangerous: it is assembled out of convenience, one sensible decision at a time. The rest of Part IX takes that story apart, failure by failure, and asks what would have to be true for it not to happen.

Failure one • The developmental record becomes an employment score

What happened. A restructuring needed criteria, and the capability profiles were sitting right there. They were used just this once. People noticed which kinds of work produced visible evidence and stopped volunteering for the quiet kind. The loop from Part VII did the rest: clearer evidence attracted better work, and better work produced clearer evidence.

The moment it could have been caught. The meeting where somebody said “just this once”. That was a proposal to change what the evidence was for, and it was handled as an administrative convenience.

What would you do? It's the restructuring meeting. The profiles are better than anything else you have, the deadline is Friday and nobody is proposing anything sinister. Do you use them?

The safeguard. Purpose limits that actually bite. Any proposal to use developmental evidence in an employment decision is made explicitly, consulted on and open to challenge, which is exactly what the Day 0 gate promised the employee forum.

The defence will be that the data is useful. It may be. The failure is not usefulness; it is silent migration of purpose and consequence. Observation collected for development should not acquire employment authority simply because the fields already exist.

RESEARCH CHALLENGE

The institutional contradiction test

THE CONTRADICTION

The organisation says the capability record exists for development. Then promotion pressure arrives.

Managers understandably prefer people with clearer evidence. Those people receive better work, produce more evidence and become safer choices. Nobody needs to introduce a ranking system formally. The institution can contradict its own developmental promise through ordinary incentives and convenience.

THE TEST

Promises count only when the surrounding conditions make them credible.

An organisation cannot credibly claim to value capability renewal while rewarding only immediate output, removing mentoring time, relying on unpaid stewardship, automating developmental work without replacement, or withholding authority to act on what learning reveals. A manager asked to develop successors while measured solely on throughput has been handed a contradiction, not a responsibility.

TAKEAWAY

FCO fails its own test if it makes people more legible to the organisation while leaving the organisation less accountable to them.

The safeguard is institutional, not cosmetic: incentives, capacity, decision rights, challenge, correction, purpose limits and routes to remedy have to back the promise.

Relates to B6

Failure two • Capability debt becomes a story we tell ourselves

What happened. The capability debt for Graham's team had an owner and a review date. The review date moved six times.

The moment it could have been caught. The second time it moved. A moved date should force a decision, not produce a new date.

The safeguard. A debt is governed only when it has an owner, a boundary, a review that changes something and a route back to rebuilding the capability. Accepting a bounded period of debt can be legitimate. Not deciding is not.

IN PRACTICE

Failure scenario · Temporary since the year nobody remembers

A team relies on one specialist during a difficult transition. Leaders record the exposure as capability debt and agree to revisit it. Delivery pressure continues. The review date moves. The specialist keeps rescuing the system. “Capability debt” becomes evidence that the risk is understood rather than evidence that anything is changing.

A named debt without an owner, boundary, review or reconstitution route is not governance. It is vocabulary applied to fragility.

Failure three · Community becomes extraction with better branding

What happened. Sam was thanked by name in the annual report and left three months later.

The moment it could have been caught. Long before the annual report: when the same dozen people were still running every session and every question still landed with Sam. Part VII's rule says repeated reliance on voluntary contribution should trigger a decision.

The safeguard. Fund the capacity, narrow the scope, transfer the responsibility, pause or stop. When the organisation depends on community work operationally, it moves into a funded service with an owner.

IN PRACTICE

Failure scenario · The helpful people become infrastructure

A community grows. Questions are answered quickly. Stories of generosity are celebrated. Three people quietly moderate conflict, translate difficult material, welcome newcomers and rescue events. Their contribution is praised and remains outside workload, progression and funding decisions. The organisation calls the community self-sustaining because the cost is being paid by people rather than a budget.

Failure four · Governance becomes the work

What happened. The launch went well. There was a dashboard, a logo and a register with owners and review dates. All three kept running. None of them changed a decision.

The moment it could have been caught. The first review where nobody could say which decision it had changed.

Controls can breed because controls are visible. A new review answers a real failure. Later, nobody can explain which decision the review changes, but removing it feels unsafe. Evidence is collected because policy requests it. Meetings exist to confirm that other meetings occurred. The architecture created to preserve judgement slowly replaces judgement with proof of procedural obedience.

The safeguard. Every control has to keep earning its place.

QUICK EXPLAINER

A safeguard can become the new problem

A control is not justified forever because it was once useful. Ask what harm it prevents, what decision it changes, what burden it creates, and what evidence would allow it to narrow or retire.

Failure five · Human oversight becomes theatre

What happened. Every agent recommendation carried a human approval, recorded in an audit trail and given in an average of eleven seconds.

The moment it could have been caught. The audit trail itself. An eleven-second average is a measurement of theatre, recorded and filed as evidence of control.

The safeguard. Part VI already set it out: answer the five questions, give reviewers the case history and the time, and require a decision rather than an approval wherever the consequence demands it.

The human is present. Control is absent.

IN PRACTICE

Failure scenario · Human oversight becomes theatre

A constructed scenario used to expose a control failure: human presence is not treated here as evidence of meaningful human authority, context or ability to intervene.

Capability systems distribute power

Any system that makes capability visible also distributes ability, work, visibility, authority, opportunity, risk, recognition and benefit. Capability-led must not come to mean people-led extraction. Eight questions about power matter most, and each needs a safeguard that exists in practice, not only in policy.

Power question	Required safeguard
Who defines capability?	Named stewardship, affected-domain participation, versioning and challenge.
Who creates and sees evidence?	Purpose limitation, proportional visibility, participant protection and correction.
Who interprets it?	Separation between observation, claim, decision, reward and authority.
Who gets opportunity?	Visible criteria, access routes, support and monitoring of unequal distribution.
Who grants trust and authority?	Named human decision owners, reasons, scope, review and right of reply.
Who carries the work?	Capacity protection, fair recognition and reward, no assumption of free community labour.
Who receives the benefit?	An explicit value and burden view, reinvestment, and customer or participant outcomes.
Who can pause, leave or refuse?	Non-retaliation, portability, exit and recovery routes proportionate to context.

Participation is not consent. People may take part because they want to, because their role requires it, because it is the only route to opportunity, or because declining would be noticed. Formal choice with a practical penalty is not choice, and stewardship, mentoring and community labour are where that risk is highest. Exit should never be impossible, indefinitely deferred or punitive. Defining opportunity, allocating it and evaluating those who take it should not sit with the same people, or incumbent experts end up defining capability, controlling progression and receiving the benefit of the system they govern.

IN PRACTICE

The employee forum's four questions, answered

Before Harbourside's plan reached the executive team, the employee forum's representative asked four questions. None was hostile. All were exactly right.

Who decides what counts as capability? Named domain stewards, with affected teams participating, versioned definitions and a route to challenge. Who sees the evidence about me? You do, and those with a stated purpose you can see; it cannot be repurposed without governance. Will this be used in redundancy selection? Capability evidence does not silently become redundancy or disciplinary scoring; any proposal to use it that way must be explicit, consulted on and challengeable. Will you give us the time? Mentoring, review, stewardship and community work are resourced, rotated and recognised, not assumed.

Answers on paper are not enough. They are tested only in practice, over time, by people who are free to say when the answers are not being honoured.

In the public sector, and in ecosystems spanning public and private bodies, one further condition applies. When capability is shared, borrowed or federated across organisations, it must still be possible to say who is accountable to citizens for a decision. Distributed capability must not become distributed unaccountability.

Perhaps FCO simply does not travel

The proposition may turn out to fit AI-enabled knowledge work unusually well and other settings poorly. It may require governance capacity that smaller organisations do not possess. It may be too broad to operate without specialists translating it back into the disciplines it tried to connect. It may add useful language without adding enough decision value to justify the complexity.

Those are not hostile interpretations. They are conditions that can make the claim more accurate.

Validation becomes marketing

This is the failure I need to watch most closely. I have built the proposition. I am building the reference implementation. I would like people to engage commercially. Those facts create incentives even when my intentions are good.

An experiment becomes theatre when an unwelcome result is reframed as "valuable learning" while the offer continues unchanged; when testimonials are promoted beyond their context; when market interest is treated as evidence of efficacy; or when the people designing the intervention also control what counts as success and who gets to challenge it.

RESEARCH CHALLENGE

Can the work disagree with its author?

The strongest protection is not a disclaimer. It is predeclared revision criteria, evidence that can travel independently, affected-party challenge, transparent boundaries and visible changes when the result is inconvenient.

Relates to B1–B7

PAUSE + CONNECT

Your pre-mortem

It's two years from now. Your FCO effort failed. Write down why.

Be specific. Which promise wasn't backed by time or money? Which evidence went somewhere it shouldn't? Whose generosity became infrastructure? Which control could nobody explain any more? Then look at your list and ask which item you could prevent this quarter. A pre-mortem costs an hour. The story at the start of this Part cost two years. [29]

**THE CENTRAL PROBLEM CAN BE REAL WHILE THE
PROPOSED RESPONSE BECOMES HARMFUL,
HEAVY OR SIMPLY UNNECESSARY.**

QUESTION

After the attack, what still stands?

Separate the problem we have evidence for, the architecture we have authored, the experiments we have begun, the opinion I hold, and the future I think is worth testing.

I do not want to end by pretending we have arrived

What do I believe, what could become possible, what should we test next, and where does the frontier widen beyond the organisation?

The tempting ending

There is an ending available in which the pieces connect, the architecture looks substantial and the paper announces a new organisational model. It would be tidy. It would also claim more than the evidence earns.

There is an opposite ending in which I present the open questions, invite the reader to make up their own mind and quietly avoid saying what I think.

I do not want that ending either.

A CHALLENGEABLE PROPOSITION STILL NEEDS A POINT OF VIEW.

Horizon 1 • What I believe now

I do not know whether every part of FCO will survive contact with reality. I do not expect it to.

I do believe the central problem is real enough that ignoring it is becoming harder to justify. Organisations are changing work faster than they can always see what that work was forming. Capability remains fragmented across learning, workforce, technology, governance, community and operations. Evidence is becoming easier to gather while legitimate interpretation and authority remain difficult. Humans and agents increasingly contribute to the same outcomes without automatically sharing accountability.

My current judgement is that these are not separate problems. They are different faces of the organisation's ability to remain capable.

I should also be plain about what I have not established. I have not shown that FCO is the only way to connect these questions, that capability can be made visible without that visibility being misused, that the Recognition Boundary is real, or that organisations will fund renewal when delivery pressure rewards the present. Those gaps are why this Part ends with experiments rather than a conclusion.

PAUSE + CONNECT

Five statuses, not one conclusion

Supported or observed: evidence exists in defined contexts. Architectural: a relationship or boundary FCO deliberately specifies. Experimental: learning from a bounded implementation. Proposed: an interpretation still requiring tests. Open: a question whose answer could materially change the thesis.

DEEPER DIVE

What would make me narrow or abandon parts of FCO?

WHY THIS MATTERS

A challengeable proposition needs findings that are allowed to count against it.

Open questions are not enough if every result can be absorbed as “learning”. The stronger discipline is to say what a result would change.

FORMATION CLAIM

Changed work can hollow out future capability even while output improves.

Weakened if AI-enabled work forms equal or better judgement without deliberate reconstruction or unsustainable hidden supervision. Revision: narrow the claims about lost developmental work and study the new formation mechanisms instead.

INTEGRATION CLAIM

Connecting work, evidence, judgement, authority and renewal changes consequential decisions.

Weakened if established disciplines connect those decisions as well with less conceptual overhead. Revision: reduce FCO to a translation layer, or retire the redundant construct.

OBSERVABILITY CLAIM

Capability can be made visible enough to support renewal without becoming surveillance.

Weakened if useful evidence repeatedly needs coercive visibility, creates gameable person scores or imposes unacceptable burden. Revision: narrow observability, return more judgement to local practice, or abandon the mechanism.

RENEWAL CLAIM

Deliberate renewal can reduce fragility and improve future options.

Weakened if benefits disappear when experts leave, do not transfer across contexts, or the renewal labour is itself unsustainable. Revision: change the intervention, economics or scope rather than declaring the architecture successful.

TAKEAWAY

The central problem may stay real while the proposed response becomes smaller, different or unnecessary.

That would not be a failure of the research. It would be the research doing its job.

Heart, head, hands, and with others

Looking back, the journey itself now seems to echo another pattern. I encountered something that mattered before I understood it. I tried to make intellectual sense of it. I began building and testing responses. Then the work became something that had to be challenged, extended and ultimately carried with others.

That is a reflection, not another architecture. It may simply be the familiar shape of serious enquiry. I still find the parallel useful: heart notices, head frames, hands test, and work this large eventually has to become with others.

Horizon 2 • What could become possible if this is valid

The future I find compelling is not an organisation that worships capability. It is an organisation that creates present value without becoming blind to the conditions of future value.

For executives, that could mean making AI, investment, workforce and operating-model decisions with a clearer view of the capability they create, consume, concentrate or renew. For teams, work could become something designed not only to produce output but to form judgement, successors and resilience. For individuals and freelancers, capability could be recognised through situated evidence and opportunity rather than reduced to job title or credentials.

For technology and consulting partners, success could include what remains possible after the engagement: what judgement transferred, which dependency is appropriate, what must remain external, and whether the client has a credible route to govern, challenge, renew or exit. For customers, the question becomes sharper still: did this make us more capable, or merely more dependent?

SIDE JOURNEY

Competitive advantage and human advantage

Competitive advantage asks whether an organisation can adapt and create value better than alternatives. Human advantage asks whether expanding machine capability is accompanied by meaningful human judgement, agency, participation, learning and the ability to shape what happens next. The proposition is that these advantages need not be enemies.

At the edge of the map • Is the next frontier FCC?

The enquiry began with safe AI adoption. It moved through work, capability, evidence, judgement and authority into the organisation. At the edge sits a larger question I am not ready to answer.

What happens when the capability challenge crosses the boundary of the organisation?

If entry routes, developmental work and meaningful participation change across an economy, the issue may not be only how one organisation remains capable. It may concern how people, communities, professions, education providers, organisations and institutions help capability form, travel, become legitimate and renew across society.




Figure 17. Four horizons: grounded proposition, plausible future, invitation to test and an open FCC research frontier.

Explanatory illustration. The fourth horizon is an open research question, not a proposal.




Figure 18. Four horizons, sketched: understood at granular scale, understood at organisational scale, tested in multiple organisations, and perhaps a way of looking at capability across a lifetime.

From the Notebook. Preserved as drawn; the rising curve is a question about scale, not a forecast.

FROM THE NOTEBOOK
Preserved as drawn

The horizons are deliberately hand-drawn. The further out they go, the less certain they are.

The sketch calls this a fourth horizon. I've left it unnumbered in the text because it isn't part of what I'm asking anyone to test yet.

RESEARCH HYPOTHESIS

Frontier Capability Civilisation

FCC is not a proposal in this paper. It is the possible horizon beyond FCO: multiple capability-forming systems connecting while preserving autonomy, meaning, challenge, correction and exit. The reader should see the door. PS03 should not pretend we have already explored the territory beyond it.

STATUS · Open · C12

Monday morning, one year on

A year after the drafting agent arrived, it is another very successful Monday morning at Harbourside. First drafts still arrive sooner. The backlog is still down. Graham is somewhere on a walking holiday with his phone switched off. He retires in three years. For the first time, somebody can answer the question nobody had asked: two people already do parts of what he does, and both are still learning. Sam's community runs its Monday session without Sam; it's someone else's turn on the stewardship rota.

A new joiner asks the question that opened this paper: if the system does the first assessment, where will I learn enough to recognise when it is wrong? This time somebody has an answer. Amira, who once asked a version of the same question, shows her the cases that arrive without a draft, the Wednesday review slot and the record she will be able to read and challenge.

Mid-morning, a complaint arrives that passes every check and still feels wrong. Amira works out why, writes it down and takes it to the person with authority to decide. Nobody calls Graham. Nothing on the dashboard records any of it.

That is the point. The output looks exactly as it did a year ago. What has changed sits underneath it: another route through the work, a newcomer with somewhere to practise, and an organisation that can lose its most experienced person for a fortnight without anybody noticing.

Your Monday morning

If this paper has changed how the problem looks, the first useful move is not to announce an FCO programme. It is to take one decision already in front of you and ask a sharper question.

Reader	A different Monday-morning question
Executive / board	What capability does this transformation create, consume, concentrate or make harder to renew?
CIO / CAIO / technology leader	Who does, decides, reviews, overrides and remains accountable, and how are people becoming capable of those responsibilities?
COO / operational leader	Where does successful delivery depend on a route that disappears when one person disappears?
Finance leader	What would doing nothing cost over three years, and what evidence would make us stop this investment?
People / learning / talent	Is “not yet ready” describing capability, or the absence of a fair opportunity to produce evidence?
Manager / team	Which real work could widen the safe route through a capability without turning an expert into unpaid infrastructure?
Individual / freelancer	What can I demonstrably achieve, under which conditions, and what next opportunity would let reality test the boundary?
Consultant / partner	What capability remains with the client after I leave, and which continuing dependency is deliberate and fair?
Researcher / sceptic	What simpler explanation could produce the same result, and which evidence would distinguish it?

Horizon 3 • Come and test it

The next proposition is not “adopt FCO”. It is “bring a consequential capability problem”.

Choose work that is already changing. Examine what it produces besides output. Identify where capability forms, concentrates or disappears. Make evidence and authority explicit. Design one bounded intervention with the people affected. Agree what would make the team continue, narrow, stop or rethink it. Then let reality answer back.

That invitation can take different forms. An executive may bring a transformation decision. A practitioner may bring a workflow. A freelancer may bring the gap between experience and visible opportunity. A partner may bring a dependency or transfer challenge. A technology organisation may bring a human-agent operating model. A researcher may bring the simpler explanation that makes part of FCO unnecessary.

Here is what I would like to test, and with whom. Each experiment is bounded, starts with the Day 0 legitimacy gate, and is designed so that an unwelcome result counts. Each also leaves the people taking part with something useful, whatever the result.

If you are...	An experiment we could run together	What would count against FCO	Appendix B
An enterprise leading AI transformation	Redesign one AI-assisted workflow so judgement keeps forming, over a bounded hundred days.	Judgement forms just as well without redesign, or the review load is unsustainable.	B1
An enterprise worried about key-person risk	Map who can really enact three to five critical capabilities, and name the debts and owners.	Existing skills data shows the same picture.	B2
A leader deploying agents	Review one agent workflow against the five questions and five roles, and make oversight real.	Review stays nominal, or approvals slow the work without fewer incidents.	B3
A smaller or mid-sized organisation	Use Journey 01, Make AI Adoption Safe, to see whether developing judgement transfers into real decisions.	People complete the journey without applying it.	B5
A partner or fractional CAIO	Test whether an engagement can create value while measurably increasing the client's own judgement.	Client independence makes the work commercially unviable for both; or everything still waits for the CAIO.	B7, B8
An individual or freelancer	Build a learner-owned account of capability and use it to open one credible next opportunity.	The record adds burden and changes nothing.	B6
A community steward	Resource stewardship and contributor routes, and see whether the community renews without losing its energy.	Participation falls once it is formalised.	B9
A sceptic or researcher	Take FCO questions to a problem with no AI in it, and see whether they change the decision.	They make no difference, so the claim should narrow.	B10

This is, unapologetically, also where the next stage of my own work sits. The intellectual proposition becomes valuable only if it helps real people see, decide or protect something consequential. Paid work can fund that experiment. It does not get to predetermine the result.

QUESTION

Want to run one of these?

Appendix B describes eleven experiments in detail: who each is for, how long it might take, what we would aim to learn, what would count against the proposition and what the co-experimenter gains. If one of them describes a problem you already have, I would genuinely like to hear from you through Empower Your World at empoweryour.world.

**BRING REALITY. CHALLENGE THE PROPOSITION.
BUILD ONLY WHAT EARNS THE RIGHT TO
CONTINUE.**

My current answer

I may be wrong about parts of the architecture. I no longer think we can afford to ignore the question.

If FCO is useful, it will not be because this paper got everything right. It will be because organisations, practitioners, partners and researchers can use the proposition to see consequential things earlier, make better choices, test those choices honestly and leave themselves more capable of meeting whatever comes next.

That is the frontier I think is worth exploring.

**THE GOAL IS NOT TO PROTECT FCO FROM
REALITY. IT IS TO MAKE FCO PRECISE ENOUGH
THAT REALITY CAN CHANGE IT.**

OPTIONAL DEPTH · A-K

APPENDICES

Use it · Challenge it · Go deeper · Reference

Appendix A • Where to begin: entry-route prompts

Every route starts from a problem somebody already recognises. What makes it an FCO route is what happens next: the Day 0 legitimacy gate and the five loops in Part VIII.

Nobody wakes up wanting to buy a Frontier Capability Organisation. People sponsor a credible response to a problem they already recognise, and FCO characteristics emerge as those responses become connected and renewable. The route matters less than the direction. A risk-led start that fixes a control without asking what capability the control depends on has not moved towards renewal. Harbourside entered through an AI-led route, which became succession-led when Graham's position surfaced, then community-led when Sam's community was drawn in.

Route	Typical starting point	A first prompt
AI-led	Moving beyond pilots or governing agents.	What capability does this deployment consume, require or remove?
Risk-led	An incident, failed control or unchallenged automation.	What capability did the control depend on, and who holds it?
Governance-led	Making oversight real in human-agent systems.	Can reviewers actually see, understand and stop what they approve?
Workforce-led	Departure, restructure, succession gap or reduced entry hiring.	Why did the gap form, and where will the next entrant practise?
Learning-led	Courses are not producing capability.	Where is the consequential practice that should follow the course?
Problem-led	A consequential problem existing capability cannot solve.	Which configuration would solve it, and what would we learn?
Customer-led	Customer dependency, poor continuity or declining trust.	Are customers becoming more capable or more dependent?
Product-led	A product team needs capability it cannot recruit.	Can we build, borrow or exchange it, and how will it renew?
Community-led	An internal community needs to become sustainable.	Who carries the work, and is there a route to succession?
Domain-led	A domain such as cybersecurity needs living pathways and standards.	Who stewards the definitions, and how are they challenged?
Ecosystem-led	Partners, suppliers or a sector building capability together.	Is the exchange fair, and is dependency bounded?
Succession-led	Critical expertise concentrated in people who will leave.	Who could do this in three years, and what would they need to practise?
Resilience-led	High performance depends on a fragile core.	How much of today's performance could we not currently reproduce?

TAKE IT FURTHER

Start where your problem already is

If your starting point is AI adoption, Journey 01, Make AI Adoption Safe, is the first Empower Your World route through it. Its public materials are open to explore at empoweryour.world; registering adds the self-guided route, community and research updates.

If your starting point is one of the other twelve routes, bring the problem itself. That conversation is how the next experiments in Appendix B begin.

Appendix B • Experiments to run together

These experiments are proposals, not completed studies. Together they are designed to test the opportunities this paper describes with the audiences most likely to recognise them: larger organisations running platform and partner estates, smaller organisations without a dedicated governance function, and the partners, fractional leaders and individuals who work across many organisations. Every experiment begins with the Day 0 legitimacy gate, states in advance what would count against the proposition, and keeps developmental evidence separate from any assessment of people.

Experiment	For, and EYW route	Shape and time	What we would learn, and what would count against	What the co-experimenter gains
B1 • Practice-route redesign	Enterprise. EYW Coached, with Assurance where an independent view is needed.	One AI-assisted workflow; Day 0 gate, then five loops over about 100 days.	Does deliberate practice keep judgement forming while the agent handles routine work, at acceptable supervisor cost? Counts against: judgement forms just as well without redesign, or review load is unsustainable.	A redesigned workflow, a successor route for a critical expert and an honest decision record.
B2 • Concentration and succession scan	Enterprise or mid-market. Workshop, then Coached.	Three to five critical capabilities; six to eight weeks.	Does a capability-and-configuration view reveal fragility that a skills matrix misses? Counts against: existing skills data shows the same picture.	A dependency map, and named capability debts with owners and review dates.
B3 • Human-agent accountability review	Enterprise or fractional CAIO. Workshop or Assurance, using Journey 01's Control Floor and Evidence Chain.	One agent workflow; four to six weeks.	Can named roles and proportionate controls make oversight meaningful without a new bottleneck? Counts against: review stays nominal, or approvals slow the work without fewer incidents.	A responsibility map, a revised control floor and a reviewer capability plan.
B4 • One-question Observatory pilot	Enterprise. Capability tier.	One question frame; around twelve weeks.	Can minimum evidence improve one capability decision without person scores or surveillance? Counts against: useful evidence needs intrusive data, or the decision does not change.	A working question frame, confidence-bearing claims and a governed evidence pattern.
B5 • Safe adoption to real decisions	SME and mid-market. Journey 01 Public, Registered, SELF, Workshop and Coached.	A cohort working through Journey 01; eight to twelve weeks, plus a follow-up checkpoint.	Does a journey that develops judgement transfer into real decisions without a large governance function? Counts against: completion without application.	A working AI front door, people who can size risk, and a record of value created.

Experiment	For, and EYW route	Shape and time	What we would learn, and what would count against	What the co-experimenter gains
B6 · Learner-owned evidence	Individuals and freelancers. Journey 01 SELF, using its starting point, reflections, application and looking-back records.	About eight weeks, with a later checkpoint.	Can learner-owned evidence support a credible next opportunity without becoming a score? Counts against: the record adds burden and changes nothing.	A portable account of their capability, in their own words.
B7 · Partner capability transfer	Delivery and technology partners. Partner tier.	One client engagement.	Can a partner create value while measurably increasing the client's judgement rather than its dependency? Counts against: client independence makes the work commercially unviable for both.	A differentiated offer, evidence of transfer and deeper client trust.
B8 · Fractional CAIO baseline	Fractional CAIOs. Fractional CAIO tier.	Two or three client organisations; about six weeks each.	Can one part-time leader set up governed AI capability that keeps working between visits? Counts against: everything still waits for the CAIO.	A reusable method and continuity for clients.
B9 · Community stewardship renewal	Any organisation with an internal community.	About three months.	Does resourcing stewardship and contributor routes renew a community without killing its voluntary energy? Counts against: participation falls once it is formalised.	Named stewards, a succession route and a clear rule for when community work becomes a service.
B10 · Remove AI from the room	Any organisation; researchers welcome.	A non-AI capability problem, such as an expert handover; eight to twelve weeks.	Do FCO questions change a decision outside AI? Counts against: they make no difference, so the claim should narrow.	A better decision on a real problem, and a test of transferability.
B11 · Capability ROI retrospective	Finance-led, usually enterprise.	An AI programme already under way; about six weeks.	Does a six-lens review change a continue, narrow, stop or reinvest decision compared with the existing business-case review? Counts against: it reaches the same decision at greater cost.	A sharper investment decision and a counterfactual the board can see.

Each opportunity in the paper is tested by more than one experiment, so no single result has to carry the whole proposition.

Opportunity to test	Experiments
Work forms capability as well as output	B1, B5, B6, B10
Concentration, succession and renewal	B1, B2, B9
Evidence without surveillance or scoring	B4, B6
Meaningful human-agent accountability	B3, B8
Value, ROI and the counterfactual	B11, B7
Partner and ecosystem transfer	B7, B8
Transferability beyond AI	B10, B9

The experiment charter, at a glance

Every experiment in this appendix starts from the same one-page charter. The questions are open for anyone to use; the working charter, Day 0 checklist and evidence templates are shared with co-experimenters.

Charter section	The question it answers
Question frame	What consequential question are we testing, for whom, and what would change if we knew the answer?
Decision owner	Who has legitimate authority to start, pause, change or stop this, and who is affected?
Purpose and limits	What is the evidence for, and what will it never be used to decide?
Protected time	Whose time does this consume, and how is that time protected?
Evidence access	Who sees the evidence, for how long, and how can it be challenged or corrected?
Stop and recovery	What would make us stop, and how would we recover if it goes wrong?
Expected value	What benefit, burden and risk do we expect, written down before the result?
Counts against	What result would weaken the proposition or this intervention?
Decision point	When do we decide to continue, narrow, deepen, scale or stop, and who records it?
What you gain	What does the co-experimenter keep, whatever the result?

TAKE IT FURTHER

Run one of these with me

If one of these experiments describes a problem you already have, start a conversation through Empower Your World at empoweryour.world. Experiments can run through the Journey 01 routes, a one-day Workshop, Coached support, a Fractional CAIO arrangement or a Partner engagement, depending on who you are and what you need.

Each experiment is designed so an unwelcome result counts. That is part of the offer, not a risk to it.

Appendix C • Harbourside, two years on: one organisation, five capability questions

A fictional, analytical story. It holds several capability needs in view at once so the mechanisms can be seen working together, and it is honest about what each route costs.

Two years after the drafting agent arrived, Janet keeps a single page on her office wall. It is not a dashboard. It lists five capability questions Harbourside is answering at the same time, each with a different route, a different owner and a different kind of evidence.

The first is where the story started. Complex-case judgement had concentrated in Graham while the agent changed how entrants practised. Harbourside chose to redesign the work and build: a protected practice route, a second reviewer and bounded agent delegation. The evidence it watches is unglamorous: stable customer outcomes, Graham's review burden, how ready his two successors are, and whether dependency is falling. In the second year Graham took a month's leave, and the queue barely noticed.

The second question has a very different shape. Cyber incident response is scarce, specialist and episodic; Harbourside will never need enough of it to build a full team. So it partners: an external response firm, with internal staff shadowing exercises and an agreed exit route. The security lead owns that decision. The evidence is response quality, how much judgement stays inside the building, and how dependent Harbourside has become. Here, the right answer was deliberately not to build.

The third is a new market proposition that needs capabilities nobody at Harbourside yet has together. The product sponsor borrows time-bounded experts, assembles a cross-domain team and insists on reusable decision assets, so the knowledge stays when the experts leave. Success is measured in time to a credible decision, what transferred, and the options it created, set against what else that team could have done.

The fourth is Sam's community, which was moving knowledge across the organisation beautifully while running on three volunteers. Stewardship is now funded and rotated, there is a visible route from member to contributor to steward, and support for production AI systems has moved into a service with an owner. Harbourside watches how the burden is shared, how many new stewards emerge, and whether the community keeps going when its founders are away.

The fifth is the one people forget. A manual validation step had consumed hours every week for years and no longer changed any decision. Harbourside redesigned the process, retired the control, recorded why it had existed, and watched for incidents in case the old check had been doing something nobody had noticed. Retiring things is part of renewal too.

None of these five sits in the same domain, follows the same route or reports to the same person. What connects them is that each is framed as a capability question, each has a named decision owner, each has evidence chosen for the decision rather than for the dashboard, and each comes back for review. From the inside, that is what an FCO looks like: not one transformation, but a portfolio of honest decisions that keep the organisation able to respond again.

The Observatory does not own this portfolio. It helps the organisation notice and frame the questions, preserve evidence and competing interpretations, and follow consequences. Decision authority remains distributed: the COO sponsors customer-resolution change; the security

authority owns incident-response decisions; the product sponsor owns the market bet; community stewardship is negotiated with participants; control retirement remains accountable to the relevant risk owner.

The Domain Mesh does not merge these domains into one maturity ladder. Customer resolution, cybersecurity, product strategy and community stewardship keep different evidence standards and authorities. Shared interfaces allow the organisation to state need, context, evidence, confidence, dependency, value and review point consistently enough for portfolio judgement.

The ontology and taxonomy therefore perform different jobs. The taxonomy makes capabilities navigable. The ontology keeps capability, evidence, opportunity, authority, outcome, value and pathway distinct and expresses their relationships. Neither proves capability. Both reduce the semantic mistakes that otherwise allow a course, activity, job title or popular contribution to become a capability claim.

The commercial decision is portfolio-wide. Harbourside may accept temporary dependency in cyber response while investing in customer-resolution succession and retiring a low-value control. The aim is not internal ownership of everything. It is explicit choice about which judgement must remain, which dependencies are bounded, where capability should propagate and what value justifies the burden.

Appendix D • Tools from the paper

These are the working tools behind Movements 2 and 3, moved here so the story could keep moving. Each can be used on its own in a leadership conversation or a workshop.

D.1 • Three definitions, three jobs

DEEPER DIVE

Three definitions, three jobs

THE PUBLIC PROMISE

Help people, organisations, communities and human-agent systems continually form, mobilise, evidence and renew the capability needed to respond to accelerating change.

THE ORGANISATIONAL ABILITY

Sense change, form and mobilise capability, connect people and intelligence, create opportunity, act with trustworthy judgement and renew faster than the environment changes.

THE OPERATING RESPONSIBILITY

Deliberately form, acquire, mobilise, govern and renew human, organisational, community, ecosystem and agent capability while preserving agency, trust and future capacity.

WHY THREE

They are not competing slogans. One tells participants what the organisation promises; one tells leaders what abilities have to exist; one tells operators what must be governed.

D.2 • Careers that do not only go up

DEEPER DIVE

What's in it for me? Careers that do not only go up

THE QUESTION PEOPLE ASK PRIVATELY

Pull has to be designed for.

What made the GSK Power Rangers community spread was not a mandate. It put the individual at the centre and answered what was in it for the person being asked to change. An FCO has to answer the same question: help people build credible futures through visible capability, meaningful practice, proportionate evidence, opportunity, fair recognition and more than one route to contribution.

ADVANCEMENT IN SEVERAL DIRECTIONS

Graham does not have to become a manager to be recognised as more valuable.

People may advance through deeper specialist expertise, broader contextual application, greater decision scope, leadership and facilitation, stewardship and successor development, community leadership, human-agent orchestration, or reusable assets and standards. That matters for the many experts who are excellent at their work and have no wish to manage people.

THE COMMITMENT

Capability evidence does not automatically determine pay, promotion, discipline or redundancy.

It may inform workforce decisions; it is not the decision. Missing evidence does not prove missing capability. Quiet contributors who work privately, intermittently or through care must still receive fair opportunity. Opportunity must be monitored, because if opportunity depends on evidence and evidence depends on opportunity, existing inequality can be reproduced in a more legitimate-looking form.

OUTCOME

Future Confidence.

The reasoned belief, supported by evidence, opportunity and a visible route, that effort can open a credible future. Not a promise that everyone will be promoted. [21]



Figure 19. Eight directions in which a career can advance, without everyone having to become a manager.

Reproduced from the White Paper v1.0 (Figure 13). Recognition does not automatically create authority, pay or promotion.

D.3 • Eight seats, eight acceptance tests

A whole-organisation proposition should be recognisable from the lives and decisions inside it. The perspectives below are not predicted benefits. They are acceptance tests: if the organisation claims to be renewing capability, what would different participants reasonably expect to notice?

Perspective	Core question	What credible progress would look like
New joiner	Where can I practise consequential work before I am expected to possess judgement?	A visible route through supported practice, correction and widening scope.
Experienced practitioner	Am I valued only for output, or also for developing the capability around me?	Protected review time, shared responsibility and successors who can act.
Manager	Who can achieve which outcome, under what conditions and with what support?	A contextual capability view rather than only courses, roles or ratings.
Executive	Are we becoming more capable, or only more productive?	Investment decisions changed by concentration, renewal and customer consequence.
AI / transformation leader	How do pilots become governed work without creating fragility?	Every consequential deployment names practice, supervision, authority and renewal.
Community steward	Will this continue beyond the people currently carrying it?	Funded stewardship, contributor routes and credible succession.
Partner	How do I create value while transferring ownership rather than engineering dependency?	Success includes retained client judgement and an intentional dependency boundary.
Customer / beneficiary	Am I receiving better value and becoming more capable rather than more dependent?	Continuity, explainable judgement, remedy and growing independence where appropriate.

These perspectives create tensions rather than a harmonious scorecard. The throughput gain that helps the executive may increase review burden for the practitioner. Wider evidence may help the manager while increasing exposure for the individual. A partner may lose recurring revenue when a client becomes more independent. The proposition is serious only if it makes those conflicts discussable and attaches authority to the people who must resolve them.




Figure 20. One organisation seen from eight seats, each with its own question.

Reproduced from the White Paper v1.0 (Figure 12). Harbourside and its people are fictional composites.

D.4 • Strategic advantage and organisational health

DEEPER DIVE

Strategic capability advantage

THE PROPOSITION

The deepest advantage may not be one rare capability today.

It may be the repeatable ability to sense, form, combine and renew the next capability before it becomes urgent.

ESTABLISHED LENSES

FCO does not claim an eighth power.

Hamilton Helmer's 7 Powers asks about scale economies, network effects, switching costs, brand, cornered resources, process power and counter-positioning. The more modest question is how capability renewal strengthens, preserves or undermines those sources of power: a process advantage that depends on a few experts who are not being replaced is weaker than it looks. Simon Wardley's evolution mapping suggests AI may commoditise some knowledge and execution while raising the strategic importance of context, judgement, trust and orchestration, and reminds us that evolution includes retirement.

THE BOARD QUESTION

Optionality, not guaranteed defensibility.

How effectively can we create, combine or access the capability we will need tomorrow, while keeping enough judgement, agency and accountability to use it well? [21]

DEEPER DIVE

How would anybody know an FCO is healthy?

Every framework that describes a desirable state is eventually turned into a maturity score, and the score becomes a target. Goodhart's warning applies. There is no FCO maturity score. Health is a profile, discussed by the people responsible for acting on it, at the level of teams, capabilities and domains, never individuals.

TWELVE DIMENSIONS

Direction, formation, availability, evidence, mobilisation, judgement, opportunity, reciprocity, stewardship, renewal, agency and sustainability. Each carries a status rather than a number: strong evidence, mixed evidence, insufficient evidence, active concern, or not yet assessed. Insufficient evidence is a legitimate answer, not a low rating.

NINE WARNING SIGNATURES

Claimed capability with weak evidence. High delivery dependent on a small, exhausted core. Courses without authentic practice. Evidence without opportunity, or opportunity without support. Agent output without adequate human judgement. Community activity without progression. Recognition based on visibility rather than contribution. No correction, appeal or successor route. Productivity rising while renewable capability declines. Before its journey, Harbourside showed at least five. None appeared in its dashboards. [21]

Dimension	Diagnostic question
Direction	Can change be translated into meaningful capability questions?
Formation	Can knowledge, rehearsal, practice, feedback and reflection connect?
Availability	Can capability be accessed with capacity, timing, permission and context?
Evidence	Are claims proportionate, current, contextual and challengeable?
Mobilisation	Can capability combine and move towards consequential needs?
Judgement	Can uncertainty and human-agent output be interpreted and challenged?
Opportunity	Can emerging and established capability find legitimate work?
Reciprocity	Are burden, contribution, value and benefit sufficiently visible and fair?
Stewardship	Are review, maintenance, mentoring, memory and care supported?
Renewal	Can the system form successors, adapt and retire weak practice?
Agency	Can participants choose, question, pause, change route and seek remedy?
Sustainability	Are human, economic, social and institutional burdens supportable?

Each dimension takes a status, never a number: strong evidence, mixed evidence, insufficient evidence, active concern, or not yet assessed.

TAKE IT FURTHER

Use the profile on a real team

The questions above are yours to use. A facilitated health-profile conversation, with the evidence prompts behind each dimension, is part of the Empower Your World Workshop and Coached routes at empoweryour.world.

D.5 • Seven value streams, seen at Harbourside

DEEPER DIVE

Seven value streams, seen at Harbourside

CAPABILITY FORMATION

Need, practice, evidence, entrustment. Amira moving from supervised practice to bounded independence. Value: new, evidenced capability for consequential work.

CAPABILITY MOBILISATION

Available capability reaching a need. Graham's judgement made available to the claims team for a quarter while their reviewer is on leave. Value: capability reaching the need without recruitment.

CAPABILITY PROPAGATION

Experience, representation, reuse, local adaptation. The practice route adapted by the claims team. Value: learning that does not have to be reinvented.

JUDGEMENT AND DECISION

Signal, question, evidence, interpretation, decision, review. The Observatory supporting Janet's decision. Value: better decisions, earlier.

HUMAN-AGENT OPERATION

Delegation, execution, supervision, learning. The drafting agent inside a clear delegation. Value: automation's productivity without unsupervised fragility.

RENEWAL

Successors forming, practice revised after an agent update, an obsolete manual check retired. Value: continuity and adaptability.

ECOSYSTEM VALUE EXCHANGE

Contribution recognised and partly reinvested. A partner transferring ownership; community contributors recognised. Value: sustained participation beyond the organisation's boundary. [22]

D.6 • Capability ROI: the six lenses

Value does not arrive in one form, so Capability ROI reads every investment through six lenses. Financial and operating ROI covers the conventional measures: cost, revenue, productivity and cycle time. Customer or beneficiary ROI asks whether the people the organisation serves receive better judgement, continuity and value, and whether they become more capable rather than more dependent. Capability and renewal ROI asks whether renewable capability has grown: fewer single points of failure, more successors, faster formation. Human ROI asks whether participants have gained agency, opportunity, future confidence and progression. Ecosystem ROI asks whether partners, communities and the wider ecosystem have gained capability and sustained participation. Risk, resilience, environmental and social value is considered where relevant. At Harbourside, the customer lens was the one that changed the conversation. [22]

Canvas section	The question to answer
Investment	What are we putting in, including time, attention and output forgone during practice?
Capability change	Which capability should change, in which configuration and context?
Evidence	What would show that change, with what confidence?
Value, through six lenses	Financial and operating; customer; capability and renewal; human; ecosystem; risk, resilience and social.
Dependencies and options	Which dependencies, risks or future options does this change?
Rival causes	What else could produce the same result?

Canvas section	The question to answer
Counterfactual	What would doing nothing cost over the time horizon, and what simpler alternative could we try?
Proof obligation	What evidence, by when, would show it is working?
Stopping condition	What evidence would make us pause, narrow or stop?
Decision	Continue, narrow, stop, outsource, automate, partner or redesign?

TAKE IT FURTHER

Take a real investment through the canvas

Experiment B11 in Appendix B applies the canvas to an AI programme already under way. If you would like to run it with your finance and transformation leads, start at [empoweryour.world](https://www.empoweryour.world).

D.7 • The four flywheels

SIDE JOURNEY

Four flywheels, and the anti-pyramid test

Under the right conditions value compounds. The individual flywheel is one person's renewal cycle: scope widening on evidence, moving from receiving review to giving it. The organisational flywheel: developed people develop others and captured learning reduces reinvention. The ecosystem flywheel: partners transfer ownership and communities develop contributors who give back. A fourth, ecosystem learning, encloses the others: methods revised on evidence and failures shared rather than hidden.

A flywheel exists only where its transitions can be observed: specific people whose scope widened, specific dependencies that reduced, specific practices that travelled. Otherwise it is a picture. One test matters above the rest. Later participants must not primarily subsidise earlier ones. A flywheel that turns because newcomers' unpaid contribution funds incumbents' benefit is extraction with good branding.

D.8 • The five loops at Harbourside

Loop	Harbourside question / action	What must be produced
Frame	How can complex-case judgement keep forming while routine drafting is assisted by an agent?	Question, affected contexts, valued outcome, alternatives and decision owner.
Observe	What happens now, including supervisor load, opportunity and concentration?	A baseline separating known, unknown, claimed and evidenced.
Intervene	Selected cases reach Amira before the agent draft; her reasoning receives protected review.	One changed condition, bounded authority, support and stop rules.
Learn	Output stays fast; explanations improve; Graham's review time doubles.	Outcomes, burden, distribution, rival causes and unintended consequences.
Renew	Change case selection, develop a second reviewer, narrow or deepen the route.	An explicit stop, revise, deepen or scale decision with an owner.




Figure 21. The first hundred days: a Day 0 legitimacy gate, then five overlapping loops ending in a decision.

Reproduced from the White Paper v1.0 (Figure 17). Indicative for one challenge; loops can repeat.

D.9 • Twelve characteristics and the customer test

FCO characteristic	What you would observe	How the customer would experience it
Capability is seen as flow	Leaders discuss where capability is forming, concentrating and degrading, not only headcount.	Fewer surprises when a key person is unavailable.
Automation decisions are capability decisions	Every significant automation proposal states what practice it removes and how that learning is replaced.	Quality on unusual cases holds as routine work is automated.
Practice is designed into work	Entrants have supervised, consequential practice routes, including in AI-enabled work.	The people handling their case can explain their reasoning.
Evidence informs, people decide	Entrustment decisions are explicit, bounded, recorded and challengeable.	Clear accountability when something goes wrong.
Trust is bounded and reviewable	Scope widens and narrows on evidence without becoming a permanent label.	Consistent quality from a wider range of people.
Human-agent accountability is clear	Consequential agent workflows name who acts, decides, reviews, overrides and is accountable.	A person who can explain and correct an AI-assisted decision.
Stewardship is resourced	Mentoring, review and community work have allocated time and recognition.	Knowledge does not leave when one person leaves.
Opportunity is fairly distributed	Criteria and access routes are visible, and distribution is monitored.	Service that reflects a broader range of perspectives.
Renewal is funded	Successor development is budgeted alongside productivity change.	Services that improve rather than degrade.
Participants can challenge and exit	People can see, correct and dispute evidence about them, and decline without penalty.	Staff who are confident rather than defensive.
Customers become more capable	Where appropriate, capability transfers to customers rather than deepening dependency.	Growing confidence and independence.
Learning is shared and methods change	Failures are reviewed openly and methods revised on evidence.	Problems fixed, not repeated.

D.10 • From allotment to organisation



Figure 22. FCO crosswalk: the same principles in the allotment and in the organisation.

Illustrative explanatory visual; labels include earlier EYW vocabulary.

Appendix E • A capability evidence record

A minimum record behind the evidence chain in Part VI. The system logs contextual evidence events that may support a claim; it never logs a capability as a fixed possession. Every record is appended and superseded, never silently overwritten.

Record element	What it captures
Who and what	The person or configuration, the capability definition and version, and any agent contribution, recorded separately from human contribution.
Context	Where, when, under what conditions, consequence and level of supervision.
Observation	What was attempted, what happened, and what judgement was required.
Evidence class	Declared, demonstrated, outcome, relational, governance, renewal, or stress/counterfactual.
Quality	Confidence in words, currency, review date, and known limits or rival explanations.
Governance	Purpose, who may see it, retention, and how it can be challenged or corrected.

A worked example: progressive entrustment to facilitate an AI governance decision workshop. Entrustment at step 5 is bounded by audience, risk and format. Step 6 counts as positive evidence. The strongest evidence arrives at step 7, when somebody else can now do the work.

Event	What happens	Evidence contribution
1. Observe	Attends a facilitated AI governance decision workshop and completes structured reflection.	Declared understanding and observed participation.
2. Practise	Facilitates one section under direct supervision.	Demonstrated practice within narrow scope.
3. Repeat	Co-facilitates two workshops in different contexts.	Repeatability and early transfer.
4. Outcomes	Participants can explain the decision, evidence, uncertainty and review point.	Associated outcome, with a causality caveat.
5. Entrust	A named human authorises independent facilitation for a defined audience, risk and format.	Bounded trust with explicit scope.
6. Correct	Review reveals weak representation of affected users; the practitioner adapts and re-demonstrates.	Challenge responsiveness and adaptation.
7. Renew	Supervises another facilitator, who performs safely with decreasing support.	Successor and renewal evidence.

An illustrative record for Amira, using the same structure:

Record element	Amira's record
Who and what	Amira; handling complex billing complaints independently.
Context	Complaints below a stated financial threshold, excluding bereavement and vulnerability.
Observation	Six months of reviewed cases, with her reasoning and the corrections she received.
Evidence class	Demonstrated.
Quality	Confidence: contextually supported. Review date in three months.
Governance	Decision owner Chloe; weekly sample review by Graham; visible to Amira with a right of reply; not an input to pay or promotion.

TAKE IT FURTHER

Keep your own record

The record structure above is open to use. Journey 01's self-guided route builds learner-owned records on the same principles: what you attempted, what you learned and what you would do next, in your own words. Explore it at empoweryour.world.

Appendix F • Claim status, propositions and challenge register

This register distinguishes what is supported, architectural, experimental, proposed or open. It is a navigation aid for challenge, not a scorecard of certainty.

ID	Claim family	Status	Basis	Hard test
C1	Work produces capability as well as output	Proposed	Supported by adjacent learning literature and situated examples	Could new work form equal or better capability?
C2	Productivity can improve while renewable capability weakens	Proposed	Research proposition	Do replacement pathways preserve judgement, succession and manageable supervisor load?
C3	Capability exceeds skill	Architectural	Authored configuration with intellectual neighbours	Does contextual configuration alter consequential decisions beyond a skill inventory?
C4	FCO contributes through connection	Proposed	Authored architecture / novelty claim	Can an established discipline connect the decisions more simply?
C5	Measurement changes opportunity and behaviour	Supported	Supported principle; implementation effects remain contextual	Can minimum evidence improve decisions without distortion or ranking?
C6	Evidence does not manufacture authority	Architectural	Architectural boundary	Where may policy legitimately automate permission?
C7	Recognition Boundary	Proposed	Research hypothesis	Does the distinction predict different useful interventions?
C8	Capability debt and renewal economics	Architectural	Authored decision lens	Do debt profiles improve investment, resilience or reconstitution decisions?
C9	Empower Your World and Journey 01 as reference implementation	Experimental	Experimental pathway, not proof	Does learning transfer beyond completion or facilitated settings?
C10	FCO can backfire	Proposed	Illustrative scenarios + architecture safeguards	Can misuse controls reduce harm without destroying usefulness?
C11	Competitive and human advantage can coexist	Proposed	Authored future proposition	What outcomes demonstrate both adaptability/value and meaningful human agency?
C12	FCC	Open	Open research frontier	Does organisational logic travel meaningfully beyond the organisation?

F.1 • Status key

Every claim in the register, and the small number of status labels in the main text, uses one of five statuses. **Supported:** evidence exists in defined contexts. **Architectural:** a relationship or boundary FCO deliberately specifies. **Experimental:** learning from a bounded implementation. **Proposed:** an

interpretation still requiring tests. **Open:** a question whose answer could materially change the thesis. The main text carries no more than twelve labels, attached to the central claims only, so that it never reads like a ledger. A status describes the state of the evidence for a claim. It is never a confidence grade for a person.

F.2 • External validation and market signals

External validation evidence is used to strengthen, qualify, contradict or contextualise a proposition. Market evidence is used to understand comprehension, demand, adoption, buyer value and possible packaging. Neither is proof that FCO works. Signals remain attached to their source, limitations and claim ceiling; repeated publication of one origin is not independent corroboration.

Signal family	What it contributes	What it does not establish
Changing early-career work and hiring	Makes developmental routes and changing work worth investigating; introduces rival macroeconomic and skills explanations.	That AI caused the labour-market change or that workplace judgement formation has weakened.
Work and operating-model redesign	Shows current organisational interest in rearchitecting work, governance and human-agent operation.	That FCO is the necessary model or that its mechanisms improve outcomes.
Connected ecosystems and partnership	Makes distributed capability, exchange and accumulated ecosystem advantage commercially legible.	A durable moat, fair value exchange, low coordination cost or bounded dependency.
Interest in the FCO/EYW proposition	Tests whether audiences recognise the problem and whether organisations may fund experiments.	Research validity, efficacy or transferability.

DEEPER DIVE

What external and market signals can, and cannot, tell us

EXTERNAL VALIDATION

Independent signals can strengthen, qualify or contradict a claim without proving FCO.

Microsoft’s 2026 Work Trend Index argues that organisational factors account for roughly twice the reported impact of AI that individual effort does; it is vendor research based on self-reported surveys. McKinsey’s 2026 State of AI survey finds most respondents report personal productivity gains while far fewer attribute earnings impact, with high performers more likely to have redesigned workflows. Both are associations, not proof that any single practice causes performance. [16]

MARKET SIGNALS

Demand and resonance answer a different question.

The External Validation and Market Signals register tracks signals such as practical AI baseline-control frameworks that operationalise govern, comply, register, access and track controls. They show the market building control architecture. They also expose the gap this paper is interested in: what happens when predefined controls do not determine an adequate response. Signals like these show the problem is legible and commercially relevant. They do not establish that the FCO mechanism works. [23]

THE DISCIPLINE

Every signal keeps its role.

External evidence informs plausibility and challenge. Market evidence informs comprehension, demand and package design. Experiments inform the bounded implementation tested. None is promoted into universal proof because it is convenient.

How this paper weighs evidence.

The strongest version of this paper must neither hide behind cautious footnotes nor decorate the argument with sources that happen to agree. External evidence has three jobs here. It can establish an adjacent phenomenon, such as weak entry-level hiring or the importance of workflow redesign. It can make a proposed mechanism plausible, such as routine work acting as developmental practice. It can also introduce a rival explanation that makes the claim narrower and more honest.

Market signals answer a different question again. Audience resonance, organisational enquiries, willingness to pay and adoption of an experiment indicate whether the problem is legible and whether a response may be commercially viable. They do not validate efficacy. A paid engagement can be rigorous evidence about one bounded implementation, but only if the success criteria, alternatives and unwelcome results remain visible.

Evidence type	Question it can answer	Legitimate use	Claim ceiling
External research / data	Does an independently observed phenomenon or relationship exist?	Plausibility, qualification, contradiction and context.	Not proof that FCO works.
Practitioner signal	Are people encountering a related operational problem?	Language, mechanisms to investigate and contextual experience.	Not independent or representative evidence.
Market signal	Do audiences recognise the problem and support a response?	Comprehension, demand, buyer value and package design.	Not efficacy, transferability or causal validation.
Bounded experiment	Did this intervention change this decision or outcome in this context?	Implementation evidence, burden, alternatives and revision.	Not universal validity or organisation-wide proof.

F.3 • Testable propositions and break conditions

Six propositions can be tested, and seven conditions would weaken or break the argument. They are stated as invitations to test, not defences to maintain.

ID	Testable proposition	How it could be tested
P1	Capability definitions that include context and configuration produce different decisions from skill inventories alone.	Matched groups make the same workforce decisions with and without configuration-rich definitions; compare decisions and later outcomes.
P2	Evidence-bearing practice improves the defensibility of entrustment decisions.	Compare how well chain-based and conventional entrustment decisions can be explained, challenged and justified afterwards.
P3	High productive but low renewable capability shows identifiable fragility.	Assess both densities across teams and observe responses to disruption such as a key departure.
P4	Capability and opportunity must propagate together for sustainable participation.	Compare pathways offering evidence without opportunity, opportunity without support, and both together.
P5	AI-enabled operating-model change alters the pathways through which judgement develops.	Task-level analysis of the learning value of work before and after AI introduction.
P6	Renewal strengthens strategic optionality but does not automatically create defensible power.	Comparative case studies over several years.

ID	Condition that would weaken or break FCO	Consequence
B1	An existing model already integrates the complete system more coherently.	The novelty claim fails.
B2	Context-rich capability definitions do not improve decisions beyond simpler approaches.	The Part III reframing is not worth its cost.
B3	Meaningful observability requires unacceptable surveillance or distortion.	The proposition fails on its own ethical terms.
B4	The architecture cannot distinguish healthy from failing systems in ways leaders act on.	The Health Profile is decoration.
B5	Renewal measures are too slow, ambiguous or gameable to govern.	Renewal cannot be managed.
B6	Institutional incentives prevent renewal whatever the evidence.	The architecture describes an organisation that cannot exist.
B7	Customer, participant or societal outcomes do not improve.	Nothing else matters.

F.4 • The irreducible core: five things I cannot remove

By this point the excavation had produced far too many concepts to carry around in my head at once. I needed a smaller test. Strip away the terminology, architecture and diagrams. What would I be unwilling to lose?

1. Productivity is not capability. More delivered today can coexist with greater dependence tomorrow.
2. Capability is more than skill. Useful action depends on the person, tools, environment, relationships, conditions, opportunity and legitimate permission around the work.
3. Judgement needs bounds. Being able to reach a sound conclusion does not automatically give somebody authority to make the consequential decision.
4. Capability has to renew. Formation, correction, succession, adaptation and retirement continue after the first successful change.
5. Reality has to answer back. A coherent model is not evidence that using it will improve an organisation.

Appendix G • Intellectual neighbourhood

This appendix is the longer prose behind the "What's new, and how would we know?" table in Part IV. Each neighbour below already does serious work on part of the territory this paper covers. For each, I say what it gives us, where FCO overlaps and what FCO tries to add. None of these sources endorses FCO, and none should be cited as evidence that FCO works. If any of them, used well, connects the decisions as adequately and more simply, it should win.

G.1 • Adjacent disciplines and what FCO tries to add

Dynamic capabilities

Teece, Pisano and Shuen define dynamic capability as the firm's ability to integrate, build and reconfigure internal and external competences in rapidly changing environments. Later formulations emphasise the capacity to purposefully create, extend or modify the resource base. FCO overlaps strongly with that territory; its additional research question is how capability formation, opportunity, evidence, judgement and legitimate authority connect to those adaptive moves.

Communities of practice

Wenger's communities-of-practice theory treats learning as social participation and explores relationships among practice, meaning, community and identity. That is highly compatible with FCO's interest in capability forming through consequential participation. FCO adds governance questions about burden, opportunity allocation, extraction, evidence and how capability moves into legitimate organisational action.

Entrustable professional activities

Olle ten Cate's work on entrustable professional activities in medical education makes explicit the supervisor's judgement about when a trainee can be trusted to carry out a professional task with less supervision, and more recent commentary frames workplace training as a form of succession. Bounded trust and supervision are close neighbours of the evidence-to-entrustment chain in Part VI. What FCO tries to add is the separation of trust from authority, and the application of bounded trust to AI-enabled service work and human-agent supervision. Domain specialists still need to ask what travels from the professional setting in which those ideas developed. [13]

AI risk management

NIST's AI Risk Management Framework 1.0 organises AI risk work around Govern, Map, Measure and Manage, with governance cross-cutting the other functions and risk management treated as iterative across the lifecycle. FCO does not replace this kind of AI governance. It asks what capability the organisation needs around governance so people can recognise emerging situations, interpret evidence, exercise bounded authority and renew the system over time.

7 Powers

Hamilton Helmer's 7 Powers asks about scale economies, network effects, switching costs, brand, cornered resources, process power and counter-positioning. FCO does not claim an eighth power. The more modest question is how capability renewal strengthens, preserves or undermines those sources of power: a process advantage that depends on a few experts who are not being replaced is weaker than it looks. [21]

Wardley mapping

Simon Wardley's evolution mapping suggests AI may commoditise some knowledge and execution while raising the strategic importance of context, judgement, trust and orchestration, and it reminds us that evolution includes retirement. FCO's interest is in which of those components have to stay renewable, and who keeps them so. [21]

Research note: these adjacent frameworks and theories illuminate component questions. They should not be cited as evidence that FCO itself is validated.

G.2 • Intellectual foundations and disciplinary landscape

This section makes visible the breadth and depth of inquiry behind FCO, EYW and the emerging FCC horizon. It is both a record of the architecture journey and a practical guide for assembling the people needed to challenge it. The depth scores describe how far the architecture has had to engage with each domain in order to frame, connect or test the capability challenge. They do not claim academic or professional mastery.

Reading the depth assessment

Depth	Meaning	Interpretation
10	Foundational	The architecture cannot coherently exist without sustained engagement with this domain.
9	Very deep	The domain shapes multiple papers, mechanisms and boundary conditions.
8	Deep	The domain is critical to design, implementation, legitimacy or success.
7	Significant	The domain strongly influences outcomes and must be represented in review.
6	Moderate	The domain provides important supporting knowledge and boundary tests.
5	Contextual foundation	The architecture requires awareness and informed challenge but has not yet gone deeply.

Tier 1 • Core foundational domains

Knowledge domain	Depth	Why engagement was required
Organisational Design	10	The FCO is a proposed organisational form and system of work.
Capability Management	10	Capability formation, mobilisation, propagation and renewal are the central subject.
Learning & Development	10	Practice, pathways, reflection, transfer and lifelong renewal underpin Capability.
Enterprise Architecture	10	The suite integrates principles, semantics, trust, information, platforms, controls and operating models.
Operating Model Design	10	Roles, authority, rhythms, Capacity, decisions and accountabilities must work together.
Change & Transformation	10	The architecture addresses repeated transition rather than one-off implementation.
Knowledge Management	10	Memory, representation, reuse, provenance and knowledge flow are core infrastructure.
Systems Thinking	10	Feedback, dependencies, delays, emergence and unintended consequences shape the model.
Value Theory	10	Legitimate value, burden, benefit, exchange and reinvestment are first-class concerns.
Community Building	9	Trust, contribution, propagation, belonging and stewardship often operate socially.

Tier 2 • Strategic and executive domains

Knowledge domain	Depth	Why engagement was required
Strategy	9	Purpose, choice, fit, investment and renewal connect Capability to direction.
Business Architecture	9	Capabilities, value streams, organisation and information must align.
Portfolio Management	8	Investment and attention move across experiments, horizons and domains.
Product Operating Models	8	Persistent ownership and outcomes inform the move beyond projects.
Service Management	8	Value co-creation, transition, support, reliability and continual improvement matter.
Governance	9	Decision rights, legitimacy, challenge, remedy and stewardship shape every layer.
Leadership Theory	8	Distributed leadership, sponsorship, succession and stewardship matter.
Decision Science	8	Evidence, uncertainty, bias, judgement and consequences require explicit treatment.
Risk Management	8	Early risk participation, proportionate control and assurance enable safe action.
Organisational Psychology	8	Motivation, identity, behaviour, power, trust and groups affect Capability.

Tier 3 • Human development domains

Knowledge domain	Depth	Why engagement was required
Adult Learning	9	Capability forms through context, consequential practice, reflection and feedback.
Behavioural Change	9	Adoption depends on conditions, confidence, incentives, habits and social proof.
Coaching & Mentoring	8	Guided practice and developmental relationships accelerate judgement.
Talent Development	8	Pathways, mobility, succession and Opportunity connect people and need.
Career Development	8	Capability identity and movement matter beyond fixed roles.
Human Motivation	8	Participation must create meaningful and legitimate value.
Human Potential	9	The architecture seeks to reveal and enable Capability rather than classify labour.
Human Flourishing	8	Wellbeing, agency, belonging, purpose and sustainable contribution matter.
Identity & Purpose	7	Work, contribution, community and change shape personal meaning.

Knowledge domain	Depth	Why engagement was required
Lifelong Learning	9	Capability requires repeated renewal across changing lives and contexts.

Tier 4 • Economics and market domains

Knowledge domain	Depth	Why engagement was required
Economics	8	Scarcity, allocation, incentives, cost, return and externalities shape feasibility.
Network Economics	8	Participation and connection may increase value and concentration risk.
Ecosystem Design	9	Multiple actors create, exchange, receive and reinvest different forms of value.
Incentive Design	8	Rewards must support healthy behaviour without displacing purpose or trust.
Exchange Theory	9	Reciprocity, contribution, access, Opportunity and circulation are central.
Social Capital	8	Relationships, trust and reputation affect access and mobilisation.
Reputation Systems	7	Recognition can support trust but may create gaming and exclusion.
Platform Economics	7	Shared infrastructure, multi-sided participation and capture risk affect sustainability.
Value Creation Theory	9	Capability must connect to outcomes without collapsing value into money.
Sustainability Models	8	The system must return enough value into future Capability and stewardship.

Tier 5 • Technology and AI domains

Knowledge domain	Depth	Why engagement was required
AI Adoption	10	AI disruption and adoption were the original problem context.
Human-Agent Collaboration	10	Work, judgement, accountability and learning span humans and Agents.
Data Architecture	8	Capability Intelligence needs governed data, lineage and context.
Information Architecture	9	Meaning, navigation, representation and findability are central.
Knowledge Graphs	7	Typed relationships can connect Capability, Evidence, people and Domains.
Intelligence Systems	9	The Observatory converts signals and questions into interpretation and learning.
Observability	8	System health, drift, burden and renewal require visibility.

Knowledge domain	Depth	Why engagement was required
Platform Architecture	8	Secure, composable services must support Journeys, Communities and Domains.
Trust & Safety	8	Participation requires safeguards, contestability and legitimate boundaries.
Security & Governance	8	Identity, access, privacy, jurisdiction and accountability are foundational.

Tier 6 • Measurement and science domains

Knowledge domain	Depth	Why engagement was required
Measurement Theory	9	Capability and value cannot be reduced to simplistic scores.
Assurance	9	Claims, Evidence, controls and review must remain proportional and challengeable.
Evaluation Methodology	8	Interventions need credible outcome and impact evaluation.
Evidence-Based Practice	8	Decisions should use explicit Evidence while retaining uncertainty.
Research Design	8	Hypotheses, tests, counterexamples and learning guide development.
Scientific Method	7	Falsifiability and disciplined observation counter confirmation bias.
Systems Measurement	8	Measures must account for interactions, delays and indirect effects.
Maturity Models	8	Progressions may guide development but can create false certainty.
Metrics & Indicators	8	Leading, lagging and balancing indicators reveal health and harm.
Falsification Methods	7	The architecture must invite Evidence that could show it is wrong.

Tier 7 • Social and community domains

Knowledge domain	Depth	Why engagement was required
Sociology	7	Institutions, norms, power, networks and roles shape participation.
Community Psychology	7	Belonging, empowerment and collective efficacy affect Capability.
Inclusion & Accessibility	7	Capability Infrastructure must reduce barriers and avoid encoding exclusion.
Neurodiversity	6	Different cognitive styles challenge standardised assumptions about talent.
Education	7	Curriculum, assessment, transition and practice connect learning and participation.

Knowledge domain	Depth	Why engagement was required
Social Mobility	6	Access to relationships, Evidence and Opportunity affects progression.
Civic Participation	6	Capability can support contribution beyond employment.
Volunteerism	6	Purpose-led contribution reveals non-market value and stewardship.
Collective Intelligence	8	Groups and networks can sense, interpret and solve beyond individuals.
Social Networks	7	Connection patterns influence knowledge, Opportunity and resilience.

Tier 8 • Future society domains

Knowledge domain	Depth	Why engagement was required
Future of Work	9	Automation, role change and human-Agent work are driving conditions.
Retirement & Later Life	6	Capability, purpose and contribution continue beyond employment.
Human Purpose	8	If work changes, meaning and contribution need wider foundations.
UBI & Post-Work Economics	6	Reduced dependence on employment raises participation and purpose questions.
Civilisational Systems	7	FCC extends Capability thinking to institutions and society.
Workforce Evolution	8	Roles, professions, transitions and pipelines continue changing.
Human Development Across a Lifetime	8	Capability accumulates, changes, decays and renews across life stages.
Capability-Based Society Models	8	FCC asks what changes when Capability and contribution become organising principles.

Perspectives assessed

Perspective	Representative lenses	Question the architecture must answer
Individual human	Learner; practitioner; contributor; mentor; explorer; career changer; retiree	Can a person understand, build, Evidence, apply and renew Capability with agency, dignity and fair access?
Team	Capability team; delivery team; product team; Community team	Can people coordinate, learn and deliver without concentrating knowledge or exhausting a few experts?
Organisation	Executive; manager; employee; architect; HR; transformation lead; CIO/COO	Can the organisation align strategy, work, authority, technology, trust, learning and value?
Ecosystem	Customer; supplier; sponsor;	Can actors exchange value

Perspective	Representative lenses	Question the architecture must answer
	partner; steward; investor; regulator	legitimately while preserving the commons and future Capability?
Society	Education; workforce; Community; inclusion; wellbeing; future generations	Who gains agency and benefit, who bears burden, and what is preserved for those who follow?
Human + AI	Human; Agent; human-Agent team; human-Agent organisation	How are judgement, accountability, learning, authority and value distributed?

Reviewer and stress-test framework

The architecture needs a council of disciplines rather than one “perfect reviewer”. Each reviewer should test the work within their expertise and at the boundary where their discipline meets others.

Review cohort	Perspectives to include	Primary challenge
Essential	Enterprise Architecture; Organisational Design; Organisational Psychology; HR/Talent; L&D; Change; Product Operating Model; Strategy	Does the organisational proposition cohere and translate into credible human and strategic practice?
Highly valuable	Economics; Systems Thinking; Community Building; Knowledge Management; Service Management; Data/AI Governance	Are feedback loops, incentives, services, knowledge, trust and ecosystem dynamics credible?
High-value diversity	Education; Neuroscience; Neurodiversity; Behavioural Science; Social Enterprise; Public Sector Transformation	Which assumptions fail outside conventional enterprise contexts, and who could be excluded?
Implementation reality	Product; Engineering; Security; Data; Finance; Operations; Legal; Procurement	Can the proposition be built, funded, governed, secured, procured and operated?
Lived experience	Learners; practitioners; managers; Community stewards; career changers; customers; partners	Does the architecture describe real needs, motivations, burdens and journeys?

FINAL OBSERVATION

The breadth of disciplines is not a decorative claim about ambition. It records what the architecture has had to connect to describe the capability challenge coherently. FCO does not replace those disciplines. It creates a shared architecture through which their insights, responsibilities and tensions can interact.

TAKE IT FURTHER

Appendix I maps candidate Research Library routes into the questions that deserve deeper treatment. Specialist review, challenge and co-experimentation can continue through Empower Your World without turning this public paper into a complete implementation method.

Appendix H • Where to go deeper next

These are optional routes for readers who want more than this paper offers. Nothing below is presented as published unless it already is. “Planned” describes a direction, not a commitment; Appendix I shows the fuller Research Library roadmap.

Route	Where it is covered now	Planned, not yet published
Value and economics	White Paper v1.0, chapters 20 to 22 (with targeted reviewers now).	Capability Ecosystem, Value Exchange and Economics architecture paper, in the FCO Architecture Suite being prepared for public release; Research Library volume RL10.
Work redesign and Journey 01	Journey 01, Make AI Adoption Safe, public materials.	Research Library volume RL06, Human-Agent Capability Systems; a work-design guide covering work with and without AI.
Evidence, entrustment and the Observatory	Parts V and VI of this paper; Appendix E.	Observatory, Intelligence and Judgement architecture paper; Research Library volumes RL04 and RL05.
Failure, capture and safeguards	Part IX of this paper; White Paper v1.0, chapter 19.	Trust, Security, Jurisdiction and Data Governance, and Requirements, Controls, Acceptance and Compliance architecture papers; Research Library volumes RL08 and RL11.
Transferability beyond AI	Experiment B10 in Appendix B.	Domain-specific packs and comparative studies.
Act on it with Empower Your World	Journey 01, Make AI Adoption Safe, public materials at empoweryour.world .	Further Journeys, Workshop, Coached, Fractional CAIO, Partner and Assurance routes, and the experiments in Appendix B.
FCC	Part X of this paper.	Research Library volume RL13, exploratory only.

Appendix I • Frontier Research Library roadmap

The earlier deep research is not discarded. It is being separated from PS03 so readers can choose depth. This map is a research backlog, not a release commitment. Individual volumes can be published when demand, evidence or experiments make a topic worth prioritising; PS03 can then receive a surgical “go deeper” route rather than a substantive rewrite.

Candidate Research Library volume	Central question	PS03 route	Current position
RL01 • Work as Capability Formation	What does changed work form besides output, and what has to replace developmental work when it disappears?	I–III	Strong source base
RL02 • Capability, Skill and Situated Action	What is organisational capability when skill, context, capacity, relationships and authority remain distinct?	III	Strong source base
RL03 • Capability Debt and Organisational Fragility	When is concentration or postponement rational, when does it compound, and how can recovery be governed?	I / VII	Strong source base
RL04 • Capability Intelligence and the Observatory	How can organisations see formation, concentration and degradation without building surveillance systems?	V	Strong source base
RL05 • Evidence, Judgement and Entrustment	When does evidence justify a bounded next opportunity, trust or authority?	V–VI	Strong source base
RL06 • Human-Agent Capability Systems	How should doing, deciding, reviewing, overriding, delegation and accountability work when agents act?	VI	Strong source base
RL07 • Capability Renewal	How do practice, evidence, opportunity, contribution and succession create another credible route through the work?	VII	Strong source base
RL08 • Power, Agency and Capability	Who defines capability, receives opportunity, carries burden and can challenge consequential interpretations?	V / IX	Strong source base
RL09 • Community as Capability Infrastructure	Under what conditions do communities form and renew capability without gatekeeping, burnout or extraction?	VII / IX	Strong source base
RL10 • The Economics of Capability	How should value, renewal cost, dependency, option value and counterfactuals inform investment?	VII–VIII	Substantial source base
RL11 • Measuring Capability Without Destroying It	How do measurement, targets, confidence and missing evidence alter the system being measured?	V / IX	Strong source base
RL12 • Recognition, Judgement and the Edge of the Map	When is a problem a missing response, and when has the organisation not yet recognised the situation well enough?	VI / X	Emerging research
RL13 • FCO to FCC	Which organisational capability principles, if any, travel meaningfully to ecosystems and society?	X	Exploratory only

The remaining research frontier includes transferability beyond AI-enabled knowledge work, interoperability across capability domains, the sustainability of supervision and community stewardship, anti-surveillance capability intelligence, bounded human-agent delegation, evidence of renewal, and whether any FCO logic travels meaningfully to the societal horizon suggested by FCC.

Appendix J · Glossary of terms

FCO · Frontier Capability Organisation: the authored proposition connecting capability formation, work, evidence, judgement, legitimate authority and renewal.

EYW · Empower Your World: the developing reference and experimental ecosystem through which parts of the proposition can be made tangible.

J01 · Journey 01, Make AI Adoption Safe: the first EYW Journey and a controlled reference implementation, not proof of FCO.

Observatory · The distributed capability through which an FCO pays attention: it notices change, frames questions, weighs evidence and competing interpretations, and supports accountable judgement. It informs authority; it never holds it.

Domain Mesh · Relationships across distinct capability domains without forcing them into one universal hierarchy.

Capability debt · The future constraint created when a system consumes, concentrates, postpones, obscures or fails to renew capability on which it will later depend. It can be intentional and governed, or accidental.

Recognition Boundary · A research hypothesis distinguishing recognised situations with uncertain response from situations not yet adequately recognised or named.

FCC · Frontier Capability Civilisation: an open research horizon beyond FCO, not a proposal established by this paper.

Capability Renewal Cycle · A public explanatory model of how capability forms and renews for a person or team: need, practice, evidence, trust, opportunity, contribution, learning, renewal.

Bounded trust · A human, contextual, reviewable willingness to rely on someone within a declared scope. Never a universal or numeric score, and never authority by itself.

Capacity · Whether capability can actually be brought to bear in the time, place and quantity needed. Capacity narrows from potential to available, accessible, agency-enabled and mobilisable.

Capability infrastructure · The connected human, organisational, semantic, trust, intelligence, community and technical conditions through which capability is formed, found, applied, shared, assured and renewed.

Capability ROI · A decision discipline for capability investment: what was invested, what changed, on what evidence, compared with what alternative, and what should now continue, change or stop.

Accountable party · The identifiable person or legal entity who answers for the outcome of a consequential action. An agent can never be the accountable party.

Decision owner · The named person with legitimate authority to make a specific decision, such as sponsoring an experiment or extending bounded trust, and who must give reasons and accept review.

Question frame · A short statement of a question's purpose, scope, affected people, assumptions, evidence need, consequence and review point, written before evidence is gathered.

Evidence-to-entrustment chain · Capability definition, evidence event, claim, human decision, bounded trust, opportunity or grant of authority, and review. Each step stops one thing from silently becoming the next.

Productive and renewable capability density · Productive density is how much valuable capability a team can enact for its size and cost. Renewable density is whether judgement, stewardship and successor development are spread widely enough for continuity.

Day 0 legitimacy gate · Six conditions an experiment must meet before it starts: a named decision owner and affected people; a stated purpose and what the evidence will not decide; protected time; visible evidence access and retention; challenge, stop and recovery routes; expected benefit, burden and risk stated in advance.

Decision record · A written account, made at the end of an experiment loop, of what will continue, change, stop or scale, with the reasons, including results that weaken the case.

Health profile · An unscored, twelve-dimension profile of capability health for a team, capability, domain or organisation. Never a profile of a person.

Theory of change · The stated chain from what FCO connects, through the decisions that should change, to the effects you would expect to observe, and what would count against it (Part IV).

Pre-mortem · A reader exercise: imagine your FCO effort has failed two years from now and write down why, then decide which cause you could prevent now.

Appendix K • Research notes, sources and figure accounts

Sources and boundaries: keep the scholarship available without turning the reading experience into an evidence ledger.

These notes distinguish the historical material, illustrations and proposed tests behind the narrative. AI-assisted drafting and synthesis were used. The argument is returned for author review, not represented as accepted by readers or independently validated at scale. A numbered note points to a reader source or explains a boundary; it is not evidence that a source endorses FCO.

[1] Simon Owen and Darlene Schwartz, Power Platform Strategy, GSK, March 2021, title slide and slides 2, 6, 19, 24, 27–28; the supplied file is titled Power Ranger Strategy – Update Mar2021. The strategy concerns business-led digital transformation with technology partnership, proportionate governance, shared ownership, blended capabilities and learning through real opportunities. Slides 10 and 22 distinguish visibility and hypothesised from actual value. Only high-level context is used here; confidential categories, contacts and photographs are not reproduced. The retrospective interpretation does not claim that FCO existed in 2021. Figure 1 is a separate notebook drawing, not a slide capture.

[2] Simon Owen, Frontier Capability Organisation: The Next Organisational Evolution Beyond Projects, Products and AI Adoption, Empower Your World, September 2026, White Paper v1.0, Targeted Review Edition, chapters 1–5 and 16–17, publication-details fictional-composite notice; Figure 7 in chapter 4. Harbourside, Amira and Graham are fictional composites used to expose decisions, not reported events or client evidence. The successful Monday is an imagined situation, not a measured intervention. Neither illustration establishes that automation harms development.

[3] Owen, White Paper v1.0, endnote 2, preserves Simon Wardley’s practitioner challenge about new work and engineering practices. That inherited reference is an intellectual counterargument, not an independently verified finding here. The concern would weaken if new work provides a better developmental route without the proposed intervention. Workplace-learning and communities-of-practice propositions in this manuscript require specialist challenge; the supplied synthesis is not a substitute for an independently verified primary-literature review.

[4] Owen, White Paper v1.0, chapters 20–22; EmPOWER Your World, Adoption, Transition and Capability Propagation Architecture, supplied edition v2.0, and Capability Ecosystem, Value Exchange and Economics Architecture, supplied edition v1.0. These are authored architecture sources, not independent efficacy studies. Capability Debt is used as a proposed decision lens. No formula, balance-sheet liability or realised return is claimed. Costs, time horizons, concentration and recovery need local evidence.

[5] Department for Science, Innovation and Technology and LinkedIn, Entry-level hiring in the UK: a snapshot, policy brief, 8 June 2026, <https://www.gov.uk/government/publications/entry-level-hiring-in-the-uk-a-snapshot>. Economic conditions and skills mismatch qualify the hiring picture; AI causality remains open. The separate linked EU/UK article at <https://news.linkedin.com/2026/eu-hiring-and-growth-2026> has no established named author/team or publication date in the available check. It is retained as an uncertain research lead, not used in the main prose. Ben Verschuere and Angus Cameron, Hiding in Front of Us: AI Impact on Jobs, Liminal Capital, August 2026, <https://www.liminal-capital.com/documents/liminal-ai-jobs-impact.pdf>. Task-based modelling separates substitution from complementarity; causal strength differs. This August note does not establish workplace judgement formation. Reference check: 7 October 2026.

[6] Simon Owen's own account, 7 October 2026, is the source for PowerPoint as a thinking medium, the move to iPad drawing, the SEACON deadline, apology and audience suggestion to keep the drawings. No quotation, measured effect or detailed scene is added. The existing documentary record retains a SEACON waypoint dated 1 November 2020; that dated record and the author's recollection are distinct sources, not a complete chronology. The named Power Rangers diagrams, Simon's Theory of Community sketches, EYW Capability Ecosystem and AI vs the Brain Muscle notebook pages are selected artefacts, not an asserted continuous sequence. Crawl, Walk, Run, Fly remains a teaching aid; network-law sketches remain challenged; the AI Transformation Flywheel is an intermediate account. Drawings make questions discussable, not true.

[7] Broad operating-model and ecosystem-advantage signals motivate questions, not specific control claims or research conclusions. No claim of a durable capability moat is supported here. Material with unresolved attribution is not used as a citation.

[8] Owen, White Paper v1.0, Appendix I, The Journey Behind the Proposition, records development and structured stakeholder perspectives. These were not a formal external reviewer panel. EYW and its first safe-AI-adoption journey are possible bounded learning settings, not independent validation. Integration with organisational learning, human factors, sociotechnical systems, workforce planning and governance remains a proposition to test. Comparative decision value, transfer and efficacy at scale are open.

[9] EmPOWER Your World, Introducing the Frontier Capability Organisation, supplied edition v2.1, section 11; Universal Information and Relationship Architecture, supplied edition v2.0; Methods, Measurement and Assurance Architecture, supplied edition v2.0, section 2. The public prose explains the uncontested contextual distinction, not an authorised exact definition. Figure 3 is the unchanged explanatory Figure 7 from Owen's September 2026 White Paper, not a newly authorised formal diagram.

[10] EmPOWER Your World, Personas, Participation and End-to-End Journey Architecture, supplied edition v2.0, section 6.1 and section 14. Head, Heart and Hands concern understanding, meaningful and plausible participation, and the practical conditions for participation. They are not a psychometric test or scoring funnel. The relational question in the narrative, working with others, is an enquiry rather than a fourth formal dimension.

[11] EmPOWER Your World, Universal Information and Relationship Architecture; Observatory, Intelligence and Judgement Architecture; Methods, Measurement and Assurance Architecture, supplied editions v2.0. The distinction between an enactment record, its interpretation and a bounded authority decision is architectural, not proof that an implemented method achieves it. Reflection prompts remain prompts, not person measures. Missing records do not establish missing capability; collection and reuse require legitimate purposes, context and challenge.

[12] Simon Owen, AI vs the Brain Muscle, notebook page, 30 March 2026, 15:36. The curves and questions document a worry, not neurological evidence, a diagnosis of AI users or proof of irreversible mental decline. The proposed comparison is a prose development of the question, not a test already performed. Case selection, prior experience, changing difficulty and hidden supervisor labour remain rival explanations.

[13] Olle ten Cate, "Entrustability of professional activities and competency-based training", *Medical Education* 39(12), 2005, pp. 1176–1177, <https://doi.org/10.1111/j.1365-2929.2005.02341.x>; and "Medical education from a succession perspective: what can the apprenticeship model bring to competency-based education?", *Academic Medicine* 101(8), 2026, pp. 917–923,

<https://doi.org/10.1093/acamed/wvag071>. These references are inherited from Owen, White Paper v1.0, endnotes 27–28, and not independently reverified here. The 2026 article is scholarly commentary, not empirical research. Bounded trust and supervision are an intellectual neighbour; transfer into AI-enabled service work requires domain-specialist examination.

[14] Owen, White Paper v1.0, chapter 19; EmPOWER Your World, Adoption, Transition and Capability Propagation Architecture and Methods, Measurement and Assurance Architecture, supplied editions v2.0; Trust, Security, Jurisdiction and Data Governance, supplied edition v2.0, and Requirements, Controls, Acceptance and Compliance Architecture, supplied edition v1.0. These sources support the design requirement for challenge, renewal and legitimate evidence use, not a validated unlearning intervention. Participant remedy, supervisor sustainability and fair access need implementation evidence.

[15] Owen, White Paper v1.0, chapter 3, drawing on the Universal Information and Relationship Architecture and the Frontier Capability Operating Model for Capability, Judgement and Practice Debt and the productive/renewable density distinction; posture taxonomy from the EYW Capability Renewal research synthesis. Entry-level sources inherited from White Paper endnotes 7-12: World Economic Forum with PwC, Artificial Intelligence and the Future of Entry-Level Work, 22 June 2026; Adzuna, UK Job Market Report, 31 July 2026 (adverts, not hires; no attributed cause); DSIT with LinkedIn, 8 June 2026 (not causal evidence); LinkedIn, The EU and UK Labour Markets, 8 June 2026; Verschuere and Cameron, Liminal Capital, August 2026 (US data); PwC 2026 Global AI Jobs Barometer, 15 June 2026. None establishes AI causality.

[16] Owen, White Paper v1.0, chapter 2. Microsoft, 2026 Work Trend Index Annual Report, 5 May 2026 (vendor research; self-reported survey; correlational); Tinkoff et al., The state of AI in 2026, McKinsey & Company, 25 August 2026 (self-reported survey; associations only).

[17] Owen, White Paper v1.0, chapter 6. The Capability Renewal Cycle is a public explanatory model mapped to the six-process Capability Renewal Loop and the architecture value chain in the White Paper's Appendix A. Figure 6 is adapted from White Paper Figure 9: the note on renewal now sits beside its stage; the stages and flows are unchanged.

[18] Owen, White Paper v1.0, chapters 14-15, drawing on the Observatory, Intelligence and Judgement Architecture and the Methods, Measurement and Assurance Architecture. The outcome hierarchy, question frames, confidence vocabulary and seven failure modes are architectural positions, not evidence that an implemented Observatory achieves them. Chloe and Helen are fictional composites.

[19] Owen, White Paper v1.0, chapter 16 and Appendix D. Evidence classes, append-and-supersede, review dates and anti-collapse rules are architecture positions. Figure 11 is reproduced unchanged from White Paper Figure 14.

[20] Owen, White Paper v1.0, chapters 17-18, drawing on the Capability Infrastructure Architecture, the Frontier Capability Operating Model, and the Requirements, Controls, Acceptance and Compliance Architecture. Figures 10 and 14 are reproduced unchanged from White Paper Figures 15 and 16. The Harbourside incident is a fictional composite.

[21] Owen, White Paper v1.0, chapters 9, 11 and 13. Helmer, H. (2016), 7 Powers: The Foundations of Business Strategy, Deep Strategy LLC; Wardley, S. (2016-2018), Wardley Maps, CC BY-SA 4.0; Goodhart, C.A.E. (1975), with the popular formulation from Strathern, M. (1997), European Review 5(3). FCO does not claim an eighth power.

[22] Owen, White Paper v1.0, chapters 20-25, drawing on the Capability Ecosystem, Value Exchange and Economics Architecture and the Adoption, Transition and Capability Propagation Architecture. Heath, M. (2025), Report on Burnout in Open Source Software (literature review and qualitative analysis); Asikin-Garmager et al. (2024), Wikipedia Administrator Recruitment, Retention, and Attrition, Wikimedia Foundation Research. Capability ROI is a decision discipline; no realised return is claimed.

[23] EYW External Validation and Market Signals Register, October 2026 entries on practical AI baseline-control frameworks. Market signals indicate legibility and demand; they are not evidence of efficacy.

[24] Simon Owen, The Interdisciplinary Nature of the Capability Challenge, explanatory map, 2026 (Figure 8); the domain list and depth ratings are the author's assessment of the enquiry's breadth, not a measured result. Figures 7, 15 and 22 are explanatory illustrations created for EYW, not canonical FCO architecture. Figure 18 is Simon Owen's notebook sketch of the four horizons, October 2026, preserved as drawn.

[25] The seven pressures in Part IV are the author's synthesis. The entry-level evidence and its limits are set out in note [15]. The points on experienced staff approaching retirement and on cost-of-living job moves are observations, not quantified here. Product names are examples from one vendor estate, not an endorsement or a complete list.

[26] The "What's new, and how would we know?" table in Part IV uses only neighbours already cited in this paper: Teece, Pisano and Shuen, Wenger and NIST AI RMF 1.0 (Appendix G); ten Cate [13]; Helmer and Wardley [21]. The "What FCO adds" column is the author's judgement, restated from Appendix G and Appendix D.4. The test column restates propositions P1, P2 and P6, break condition B1 and Experiments B3 and B9. No new literature or empirical claim is introduced.

[27] "What actually changes" draws on Owen, White Paper v1.0: chapter 3 (the operating model funds review capacity, alternative practice, succession and renewal alongside productivity change); chapter 9 (portfolio routes; accepting bounded debt with a named owner and review date); chapter 13 (health profile, never of a person); chapter 15 (stewarding domains and review points); chapter 16 (bounded trust declares subject, scope, context, evidence basis, owner and review point); chapter 17 (five accountability roles; hold, review and override when an agent conflicts with human judgement; authority before action); chapter 19 (named domain stewards; promises and the conditions that make them credible); chapter 21 (proof obligations and stopping conditions); chapter 23 (community stewardship); chapter 25 (decision owner; the five loops); and Appendix C of this paper. The grouping into roles, rhythms, decision rights, artefacts and money, the naming of the six rhythms, and the "before" column are the author's synthesis for a fictional composite, not a reported implementation.

[28] The hundred-day diary in Part VIII draws on Owen, White Paper v1.0, chapter 25: the six Day 0 conditions; Day 0 at Harbourside taking three weeks and involving the employee forum; selected cases reaching two analysts without an agent draft; a Learn loop that found stable outcomes, moderate and weak evidence confidence for the two analysts, a larger-than-expected throughput dip in the first month, a distribution concern and informal adoption by the claims team; and a Renew loop that shared stewardship with Chloe, began developing a second reviewer and made case selection fairer. The near-failure on Graham's protected time and the doubling of his review time are narrative illustrations of the Day 0 and Learn-loop conditions. Day 30, 60 and 90 are narrative markers; the White Paper's loop timings are indicative. Harbourside is a fictional composite.

[29] Klein, G. (2007). "Performing a Project Premortem". Harvard Business Review, September 2007. The pre-mortem is used here as a reader exercise, not as evidence.

A note on visual language. The figures are meant to become more formal through Parts IV to VI, where relationships are stable enough to specify, and to loosen again at the frontier and in Part X, where the map is not the territory. Formalise what has earned stability; keep the pencil for what hasn't.

ACCESSIBLE FIGURE ACCOUNTS

What the figures show

Figure 1 shows three connected learning modes: 10% formal training, 20% learning from others and 70% continual learning whilst doing. People are shown in a taught group, a network and individual uncertainty or discovery. The small classroom-board strokes add no research finding. The percentages are historical labels, not validated allocation targets. The purpose is to make the relationship between doing, help and formal learning visible without asking the reader to decipher miniature handwriting.

Figure 2 is an illustrative synthetic image of the back of a UK-style driving licence, with no real personal data. A grid lists vehicle categories from mopeds and motorcycles through cars, vans, lorries and buses to towing combinations, each with a symbol, start and expiry dates and a column for restriction codes. Some categories are held and others are blank. The point is that "can drive" hides very different, separately permitted capabilities.

Figure 3 distinguishes organisations, freelancers and individuals within an ecosystem bounded by trust. It names network acceleration, connection, sharing, collaboration and events. Assets include frameworks, accelerators, modules, templates, case studies and best practices. Capability labels include vision, strategy, security, governance, operations, delivery, development and industry experience. Delivery separates Simon's scalable delivery, network-member delivery and freelancer delivery. Arrows suggest increased capability density, problems solved, reduced resourcing costs and network growth. Those are aspirations, not demonstrated effects. Below, data, information, knowledge, expertise, experience and judgement form an exploratory progression. The notes contrast abundant knowledge with scarce capability and place EYW between them. Neither the progression nor the commercial promise is validated.

Figure 4 places "a difficult complaint resolved well" at the centre of eight dependencies: customer and case history; people and their competences; agents and tools; data and context; relationships; capacity; permissions and authority; governance. The connecting lines express dependencies, not a temporal sequence or an established causal mechanism. The original small footer says a skill is one ingredient within people and their competences, and that removing an element may make the outcome fail. "May" matters: this is a possibility, not a universal prediction.

Figure 5 is headed "Blog Post – AI vs the Brain Muscle" and is dated 30 March 2026 at 15:36. A rising curve is labelled rate of technology change. A descending curve is associated with loss of mental muscle, with media, social media and AI alongside it. Handwritten questions ask when that muscle is lost and cannot be recovered, and whether there is a scale of criticality or loss of self. The account preserves what the sketch asks; it does not adopt those questions as findings. No curve establishes irreversibility. The narrative moves from the analogy towards observable tests of practice, judgement and dependence.

Figure 6 arranges eight stages clockwise: need, practice, evidence, trust, opportunity, contribution, learning and renewal, with renewal leading back to need. At the centre, three flows labelled capability, agency and benefit must circulate. Notes state that trust is bounded and does not by itself create authority, and that renewal is selective.

Figure 7 is an illustrated allotment. Callouts map parts of the scene to FCO ideas: the wider environment to external context, a watchtower to the Observatory, conversations over the fence to communities, the soil to a culture of trust, curiosity and stewardship, a produce table to a capability marketplace, experimental beds and a greenhouse to experimentation, a shed to strategic architecture, a gardener with a tablet to the human-agent model, and the allotment society to leadership, stewardship and governance. A strip along the bottom shows a renewal and learning cycle: observe, question, experiment, learn, adapt, share, embed, renew.

Figure 8 groups more than forty knowledge domains into eight tiers, from core foundational domains such as organisational design and capability management, through strategy, human development, economics, technology and AI, measurement, and social and community domains, to future society. A central wheel connects organisations, people, technology, strategy, economy, community, measurement and society around capability. Lower panels list the perspectives to assess from and the reviewers who should stress-test the architecture.

Figure 9 shows four stages in sequence: observation (what is happening?), interpretation (what might it mean?), judgement (what should we do?) and permission (who can authorise it?). Three governed boundaries separate them: from evidence to meaning, from meaning to judgement, and from judgement to authority. Notes beneath read: evidence does not interpret itself; plausible does not mean proven; good judgement does not grant permission. A closing principle reads: separate the stages, strengthen the system.

Figure 10 is a matrix of five roles (agent, agent operator, reviewing practitioner, human authority holder, accountable party) against five functions (acts, decides, reviews, can override, accountable). The agent drafts and classifies but never decides or holds accountability. A note reads: agents may exercise delegated authority; they may not originate accountability.

Figure 11 shows a seven-step chain in two rows: capability definition, evidence event, claim and human decision, then bounded trust, opportunity or authority, and review. Gates read “evidence informs, people decide” and “trust does not automatically create authority”.

Figure 12 shows two territories divided by a porous boundary, drawn as a river. On one side, a recognised situation with an uncertain response, served by signposts for procedure, escalation and practice. On the other, a situation not yet recognised, where signposts point to comparison, questioning, reframing and learning. Arrows show new patterns and anomalies crossing between them. A label marks it as a research hypothesis to test and refine, not a proven fact.

Figure 13 shows a loop in which trusted people receive consequential work, the work produces evidence, and the evidence produces more trust, concentrating opportunity. Around it, deliberate participation, access, support, safe practice, learning in the flow of work and visible contribution widen the route. A side panel shows the fragile alternative: when opportunity flows through only a few, dependency, bottlenecks, lower diversity and higher risk follow.

Figure 14 shows three stacked bands: visible work on top of capability movement, which sits on enabling infrastructure. It is a simplification of layers, not an organisation chart.

Figure 15 is an illustrated organisation laid out like an allotment. Callouts map each area to an FCO idea: the environment to external context, a lookout to the Observatory, a tree whose roots carry trust, curiosity and respect to culture, a market stall to capability exchange, an experimentation zone and experimental beds to bounded experiments, a tools shed to strategic architecture, a person with a tablet to the human-agent model, and a table of commons decisions to leadership, stewardship and governance. A strip along the bottom shows a continuous renewal and learning cycle: observe, question, experiment, learn, adapt, share, embed, renew.

Figure 16 shows a bounded experiment in seven steps: consequential work, baseline, one bounded change, protected stop or revise conditions, evidence and rival explanations, reality answering back, and a decision to continue, narrow or stop. A legitimacy gate sits before the start, with a checklist covering objective, people involved, proportionate approach, clear authority and a commitment to stop or revise.

Figure 17 shows a landscape with four numbered horizons: what we believe now (a grounded proposition), what could become possible (plausible expansion), an invitation to test (real-world experiments) and FCC (an open research horizon beyond the organisation), shown as a door opening onto open questions. A bridge labelled FCO rests on five pillars: work, capability, evidence, judgement and authority.

Figure 18 is a hand-drawn rising curve across four horizons. Horizon 1: the problem and opportunity understood at granular scale. Horizon 2: the opportunity understood at organisational scale. Horizon 3: experiments and tests carried out in multiple organisational scenarios. Horizon 4, with a question mark: is this a way of looking at capability across a lifetime, a Frontier Capability Civilisation?

Figure 19 places “you” at the centre of eight directions of advancement: deeper specialist expertise, broader contextual application, greater consequence or decision scope, leadership and facilitation, stewardship and successor development, community and cross-domain translation, human-agent orchestration, and innovation, assets and standards.

Figure 20 places one Frontier Capability Organisation at the centre of eight perspectives, each with its own question: customer or beneficiary, early-career entrant, experienced practitioner, manager, AI or transformation leader, executive sponsor, community steward, and partner or fractional AI leader.

Figure 21 shows a Day 0 legitimacy gate (owner, purpose, capacity, access, challenge, stop, benefit, burden) followed by five overlapping loops: frame, observe, intervene, learn and renew, ending in a decision to stop, revise, deepen or scale.

Figure 22 places the allotment and the organisation side by side, joined by a central column of shared principles, from “context shapes everything” to “learn, adapt and renew”. A side panel lists principles that hold in both worlds, including people at the centre, learning is continuous and stewardship over control.