Tautai

Agency at Work

What agency is made of, and what a team can do about it

Version 0.1 · 17 November 2026

Krishan Mathis · tautai · krishan@tautai.net

What this paper is for

This paper explains a claim, not a product. It sets out what agency is made of, reads the same model at the person, the team and the organization, and says what a diagnosis may claim, what the practice does not promise and where a team should start. Its limits sit in the same place as its claims.

The occasion is a pattern that repeats in conversations about capability. A team is given new tools, a licences budget and a course catalogue. The question asked is whether people are equipped, and the answer is counted in completions. Six months later the tools are in place, the certificates are on the wall, and the work has not changed. Something was measured. It was not the thing that decides whether any of it converts.

Most teams are offered the dimension they already hold. The model below asks a different first question: which dimension binds this unit, and is the binding constraint even inside the unit at all.

In five sentences. Agency is the socially produced capacity to direct your own course, and it is limited by its weakest dimension. Three dimensions gate it — capacity, room to act and will — and two aim them, direction and identity. The same five elements read at three resolutions: the person, the unit, the organization. Most capability programs invest where an organization is already strong, which is why the training completes and the work does not move. The question that decides the outcome is which dimension binds — and a block found in a team whose cause sits one level up is an organizational-design finding, not a training need.

What in this paper is which

The claim is old enough to have parts in different states, and the difference matters to anyone who might quote it.

Element	Status
The five dimensions and the weakest-dimension rule, at the person	the manuscript's claim, Agency Is All You Need, draft v0.1 — not accepted canon
The nine team-layer concepts	accepted working canon, 2026-10-12 (D-004) — a working commitment, not proof of truth
The organization resolution	proposed; the organization as a navigating unit is the working frame (D-015), and two questions stay open
The instruments and the consulting package	draft v0.1 (imported 2026-10-12); no measured results are recorded in this field
The conditions report	a draft template of 2026-10-20, not yet used in an engagement

Nothing in the paper is a research finding. It is a model, a set of concepts and a practice, and the practice has not been measured against its own claims.

1. The problem: three answers, each one dimension

Talk to a leadership team about capability and you will hear one of three answers. Each answers one dimension of a five-dimension problem, and each is offered as if it answered all five.

More access. Licences, laptops, the new platform, an intranet full of courses. Access is the cheapest part of agency to buy and the easiest to count. It is also, on its own, exposure: people are handed tools whose results they may not be able to check, in a unit that has not said who may override them. The field's accepted canon states the rule plainly — agency is built from conditions, not from a trait, and access without those conditions is exposure.

More training. Courses, certifications, a reskilling catalogue. Training builds capacity, and capacity is one of the three doors. It does nothing for the other two. A skilled person in a unit that will not let them decide is not more agentic for the course; they are better equipped to notice that they cannot act. That is the failure the model calls frustration, and it is the most common misreading of a capability gap.

More will. Ownership, mindset, empowerment, an away-day with a motivational speaker. Where capacity and room exist, this can be the missing piece. Where they do not, it produces motion: plans made with adrenaline, decided without a read of the room, and quietly reversed. The model calls it recklessness, and the base rate of the survivors is not encouraging.

The AI twist makes the mistake expensive. Automating the junior work is the most defensible efficiency argument a team has, and it removes the occasions on which the next expert was made. A unit can raise its output this quarter and lower its capacity to evaluate that output in three years. The canon names it the pipeline paradox, and it is accepted working canon, not a hypothesis the field is still testing.

2. The idea: five dimensions, three that gate

Agency has five dimensions. Three are operational and gate each other. Two frame the configuration and aim it.

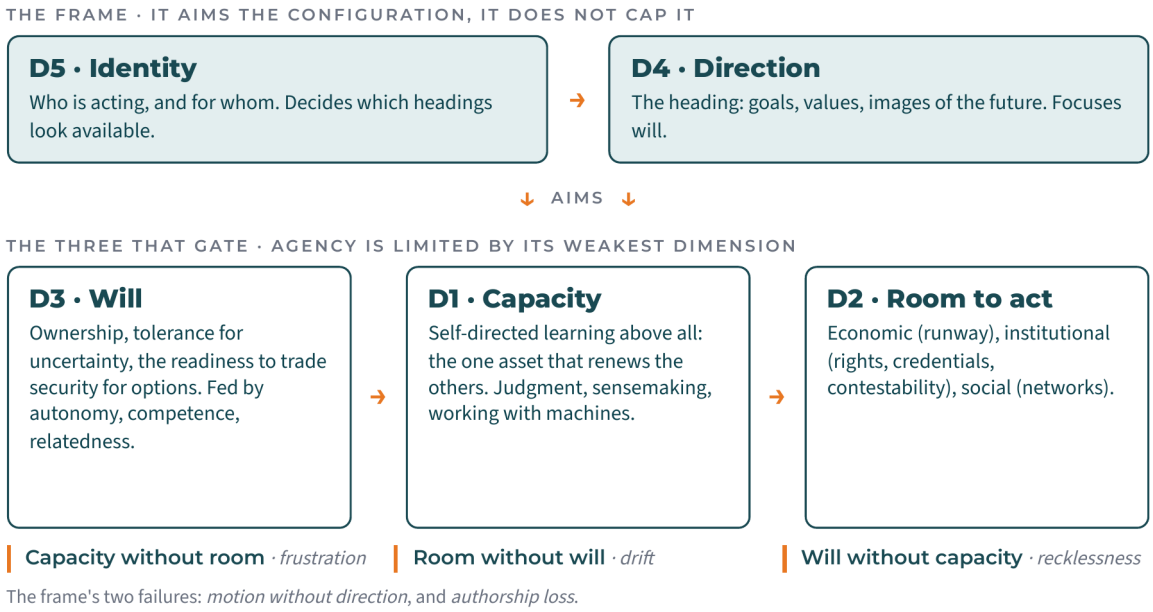


Figure 1 — The five dimensions. Capacity, room and will gate each other; direction and identity aim them. The rule is the sentence that makes the five a model: agency is limited by its weakest dimension.

Capacity (D1) is the ability to function and to keep functioning as conditions change. Its master capability is self-directed learning — working out what you need to know, acquiring it and applying it without an institution assigning it — because it is the only asset in the set that renews the others. Working with intelligent machines, judgment under uncertainty and sensemaking sit beside it. The stock of domain knowledge a person holds is not capacity; it is inventory, and inventory expires.

Room to act (D2) is the objective space of action, in three arenas. Economic room is income, buffers and obligations, condensed into runway: the time your resources buy at your current burn rate. Institutional room is rights, status, credentials and the gatekeepers who administer them, and, just as decisive, contestability — whether you can appeal a decision made about you. Social room is networks, family systems and communities, which are resource and constraint in the same breath.

Will (D3) is the disposition to act: ownership of your own course, tolerance for uncertainty and the readiness to trade the illusion of security for real options. It is fed by autonomy, competence and relatedness. Starve any of the three and motivation does not simply drain; more often it curdles into the driven, self-punishing kind, where people work harder and the work costs them.

Direction (D4) is the heading — goals, values and images of the future. **Identity (D5)** is the answer to who is acting and for whom. Direction and identity do not cap the configuration the way a weak operational dimension does. They misaim it. The two failures they own are motion without direction, which is the most flattering of all the failures because from outside it looks like success, and authorship loss, where a role-story does the steering.

The interaction rule dictates the strategy: work on the weakest dimension first. This sounds obvious and is systematically violated, because people prefer to invest where they are already strong. The capable take another course, the resourced buy another buffer, the willful make another leap. The rule also explains an inequality the field's companion volume documented: when opportunity is distributed without regard to all three dimensions, it flows to those who can already convert it.

Three borders belong to the word. Agency is not hustle culture, which performs will and burns the other two. It is not positive thinking; confidence is downstream of evidence, and the strongest evidence is a mastery experience you gave yourself. And it is not a license to blame. Agency is socially produced and gated by room, so when a person whose room was constrained from the start fails to flourish, the model's verdict is that the weakest dimension was chosen for them.

3. The same model at three resolutions

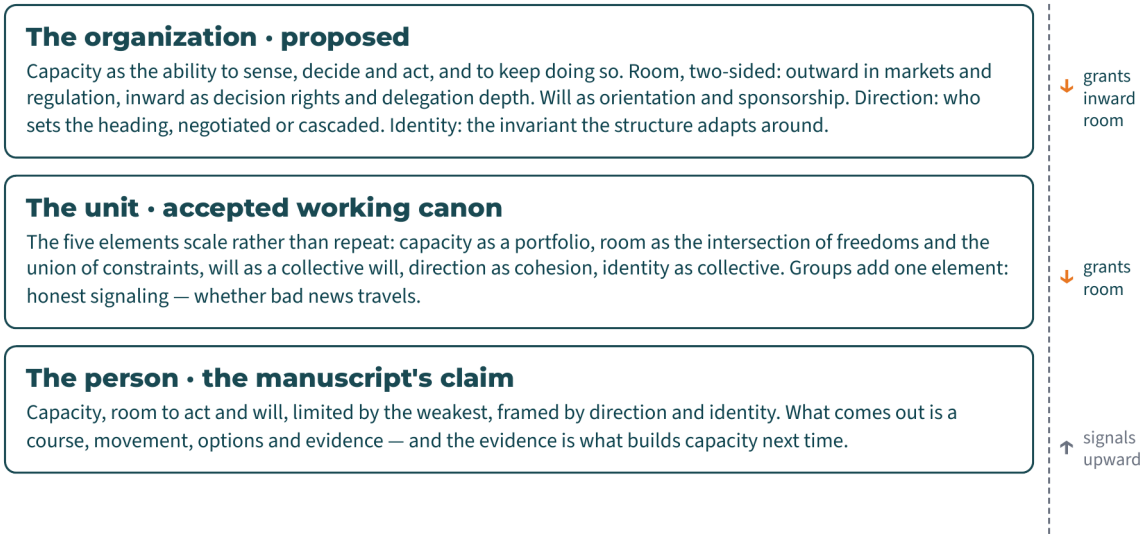


Figure 2 — One subject at three resolutions. Downward each level grants the room the level below acts in; upward each supplies capacity and signals what it cannot absorb. The statuses differ, and the figure says so.

At the **person**, the model is the manuscript's claim: the five dimensions, the gate, the two frame failures. It is drawn from a literature rather than from a study of teams, and it is draft.

At the **unit** — a team, a crew, a department with its own customers — the five elements scale rather than repeat, and groups add one element a person does not have: honest signaling, whether bad news travels. A unit that filters bad news navigates on a false map. The team layer is the field's only accepted canon, nine concepts agreed as a working commitment on 2026-10-12:

Concept	One line
The team is the unit of adaptation	The four capabilities live in teams, not in org charts
Agency is infrastructure	Agency is built from conditions, not a trait; access without it is exposure
Team capability is emergent	Trust, judgment and coordination arise from interaction, not role definitions
Hybrid teams: humans above the loop	The boundary principles that fit human teams fit agent fleets
The human remainder	Meaning, motivation, mentorship, relationship and judgment stay scarce
The new signaling stack	Degrees, then micro-credentials, then evidence, then behavior
The pipeline paradox	Automating junior work consumes the future senior talent
Orchestration replaces coordination	What remains of management is designing human+AI work and holding direction
The dissolved coaching layer	The work of the agile-role layer stayed; its host role went

At the **organization**, the model is proposed. Room becomes two-sided: the room the organization has outward, in markets and regulation, and the room it grants inward, as decision rights, delegation depth and the resource bargain. Capacity becomes the ability to sense, decide and act, and to keep doing so. Identity is read from the sister field that owns it rather than restated here. Two constraints travel with this level, and both are refusals: there is no company grade, and viability is indifferent to purpose, so a well-allocated organization can be a hospital or a cartel. The level answers whether an organization can direct its course, not where it should point it.

The cross-level rule is the most useful sentence in the model for anyone who works in a team: **a dimension at one level is a condition at the level below**. The team's latitude is the organization's delegation depth. The team's information access is the organization's reporting design. That is why the gate is always measured one level up, and why a diagnosis that stops at the team will keep finding the same block.

4. What a reading may claim

Every element of a finding carries one of three marks: **measured** (a value on a stated scale, from a stated instrument, with n), **read** (a categorical judgment by a named reader from named evidence), or **asserted** (a statement

without instrument or evidence). An element moves up only when evidence or an instrument is attached. It never slips down silently.

The practice reads team conditions with seven instruments, each usable in a quick mode of a 60 to 90 minute workshop or survey, or combined into one profile:

Instrument	What it reads
Team agency audit	Decision latitude, information access, mentorship, safety, AI usage patterns — the conditions for agency
Cognitive load assessment	Whether team structure fits cognitive capacity: domains, handoffs, boundaries, supervision load
Hybrid readiness assessment	Task inventory, supervision capability, interaction-design maturity for work with agents
Learning infrastructure audit	Mentorship density, junior work inventory, course-versus-capability ratio
Pipeline risk assessment	Tenure structure and tacit-knowledge concentration — the seniority cliff
Psychological safety baseline	The safety foundation the other instruments stand on
AI fluency assessment	Practiced skill per audience tier

Two of them form a deliberate pair, and the pair carries most of the method: the audit reads the **conditions** and the fluency assessment reads the **skill**. Run together they produce a two-by-two that decides what to do next. Conditions without skill is a unit that will hand its judgment to a tool. Skill without conditions is a unit that could check the tool and is not allowed to.

No item asks about individual performance, results are reported in aggregate, and no finding can be used in a performance conversation. The instruments are client-safe by construction: the unit sees its own profile, and the profile is written as questions the unit can act on.

5. Where the practice begins: the binding dimension

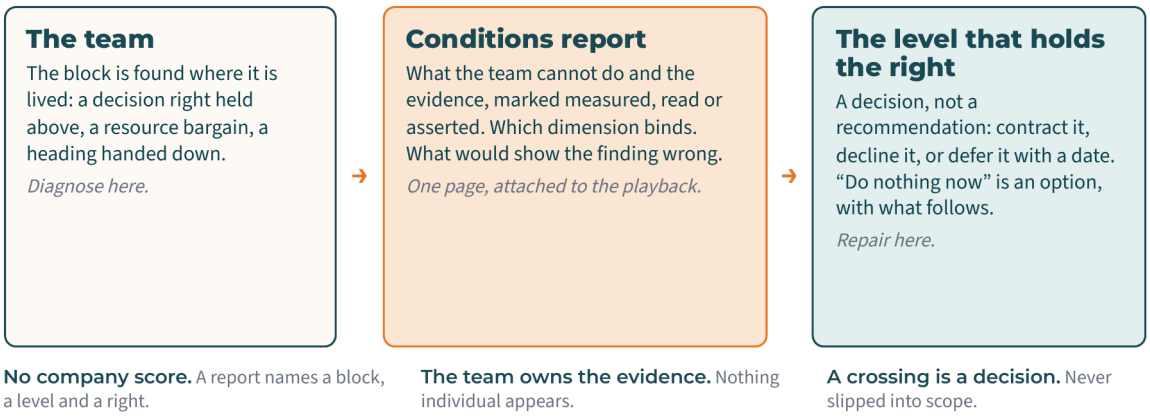


Figure 3 — A block is diagnosed where it is found and repaired where the right sits. The crossing is a decision for the client, not a recommendation slipped into scope.

The engagement begins with the conditions, not with the tools, in five practice streams: team structure, people flow, continuous learning, roles in transition and AI fluency. They are delivered through six phases, of which the first is qualification and the last is an exit; the package is explicit that it leaves.

Two readings decide the route. First, the conditions-versus-skill grid from the pair above. Second, whether the binding dimension is inside the unit or above it.

If the block is inside — low mentorship density, a handoff that exceeds the team's cognitive capacity, no agreement on which decisions the team owns — the work is team work, and the instruments say where to start.

If the block is above — a decision right held two levels up, a resource bargain the unit cannot change, an identity question only leadership can answer — the finding leaves as a **conditions report**. One page per finding, attached to the playback: what the team cannot do and the evidence, which dimension binds, which level holds the right, the scope line, the crossing as an option with a decision, and the observation that would show the finding wrong. Three rules travel with it. No company score: a report names a block, a level and a right, never a grade. The team

owns the evidence: only what the team agreed appears, and nothing individual. And a crossing is a decision, not a recommendation.

This is the part of the method that most engagements get wrong, and it is a design choice rather than a technique. A team development program that absorbs an organizational block converts it into a team deficit. The unit tries harder against a constraint it cannot move, and the report that should have gone up becomes a workshop instead.

6. What the practice does not promise

No reading delivers: a score, a benchmark against other teams, a maturity level, a certificate, a promise that a team becomes agile or high-performing, or the assurance that agents will make the team more productive.

Two results are explicitly admissible. **The binding dimension sits above this team, and the finding goes up.** And **this team should not adopt agents yet** — with the reason and with what would change that judgment. A practice that cannot produce those two results is a sales instrument.

7. How you would know it is wrong

Every finding in the conditions report carries the observation that would refute it. That discipline is at the level of the finding. The model itself is another matter, and the field has stated no falsification targets at this level. What follows is the set of tests the model implies. None of them has been run.

- **The weakest-dimension rule.** If teams whose binding dimension was addressed did no better at re-measure than teams that invested in their strongest dimension, the gate would add nothing to a list.
- **The unit layer.** If a unit's conditions — latitude, information access, mentorship, safety — predicted nothing about how well it corrected its own course, the team layer would be decoration.
- **The direction of the gate.** If units that lack room were rescued by more capacity, or units that lack capacity were rescued by more room, the cross-level rule would be wrong about where to look.
- **The pipeline paradox.** If units that automated their junior work showed no change in the supply of people who can evaluate what comes back, the argument that links delegation depth to evaluation capacity would fail.

Two things would weaken the model without refuting it, and the field should say so: if the five dimensions proved impossible to read independently in practice — if capacity and will could not be told apart by a reader using the instruments — the model would lose its diagnostic use before it lost its truth.

8. What this means for a team without a mandate

A team can read most of its own conditions without permission. Four questions do most of the work, and all four are answerable in a room.

May we override? Name the last three decisions the team made about its own work, and who could have reversed them. If the answer is nobody in the room, the block is room.

Could we tell if the output were wrong? Take the last piece of model output the team used. Ask who checked it, against what, and what they would have done if it had been wrong. If nobody can answer, the block is capacity, whatever the training record says.

Where do the next experts come from? Name the work that used to make them and where it went. If a tool absorbed it and nothing replaced it, the block is the pipeline, and it is already running.

Does bad news travel? Ask what happened to the last person who reported something unwelcome. The answer is the unit's signaling condition, and a unit that filters bad news is navigating on a false map.

What the team cannot do itself is grant itself room. That is the upward half of the method, and it is a decision for whoever holds the right.

9. Lineage

The model has a long lineage and a short evidence base, and both should be stated. The three operational dimensions come from three traditions that developed independently: Bandura on agency and self-efficacy, where confidence is earned through mastery experiences rather than installed by talk; Sen on real freedom against nominal ability, and on the person as agent rather than patient; Deci and Ryan on why people with skill and opportunity still do not move. The split between capacity and room is built on Sen's conversion problem. Direction

and identity draw on the agency literature and on the sibling field that owns identity as a subject. Room draws on Hirschman for exit and voice, on Granovetter for weak ties, and on Beer for the warning that the same toolkit can concentrate agency or distribute it.

The AI-side arguments — delegation depth bounded by evaluation capacity, the pipeline paradox, humans above the loop — are this field's own, and they are the least tested part of it. The team-layer canon that carries them was accepted as a working commitment, which is a governance state and not proof. The paper is written so that a reader can see which sentence rests on a literature, which on an argument and which on an institution's decision to commit.