

BEYOND HUMAN COORDINATION

Stop Managing Work That Agents Can Execute: A Manifesto for Autonomous Delivery

Why transformation leaders must redesign governance before agentic work redesigns the profession

Giovanni Leonardi

March 2025



Stop Managing Work That Agents Can Execute: A Manifesto for Autonomous Delivery

Transformation · No. 343
March 2025

"The profession's future is not to remain in every loop; it is to decide which loops deserve a human."

1. The Profession Must Choose

In a controlled five-day pilot, an AI agent is given a workstream objective, forty-six source documents, access to a planning repository and a defined set of permitted actions. It decomposes the objective into 312 activities, finds nineteen dependency conflicts, drafts the weekly report, prepares six stakeholder messages, checks evidence against acceptance criteria and updates the risk log. Human specialists intervene nine times: twice to correct ambiguous scope, three times to approve external communication, twice to resolve conflicting policy, and twice because the agent's confidence is unjustified.

The experiment is imperfect. It is also decisive.

We are no longer discussing software that merely helps a programme manager write faster. We are beginning to see systems that can pursue an objective, choose a sequence, use tools, inspect results and continue without waiting for an instruction at every step. Their reliability is uneven, their authority must be bounded, and their judgement is not human judgement. None of that makes the direction less consequential.

The transformation profession now faces a choice. We can treat agents as improved assistants and preserve our existing roles until the work is redesigned around us. Or we can accept that autonomous execution changes the unit of delivery, then redesign governance, leadership and professional value before adoption outruns discipline.

We should choose the second path.

2. Stop Calling It Assistance

The language of assistance is already too small.

An assistant waits for a request. An agent is assigned an outcome. An assistant produces an artefact. An agent takes actions across a sequence. An assistant leaves coordination to a person. An agent can maintain state, revisit incomplete work, call another capability, test a result and escalate an exception.

That difference matters because transformation has built an enormous professional layer around coordination:

- converting objectives into tasks;
- chasing owners and dates;
- reconciling plans;
- preparing status reports;
- comparing evidence with gates;
- routing decisions;
- maintaining risks, actions and dependencies;
- repeating instructions across teams.

This work is necessary. Much of it is also structured, legible and vulnerable to autonomous execution.

We should not defend it merely because it has accumulated senior titles, established methods and comfortable career paths. A profession earns its future by solving the next problem, not by protecting the labour required by the previous toolset.

Coordination is a cost of complexity, not proof of leadership.

3. We Must Stop Confusing Presence with Control

Our programmes often equate human presence with governance. A person attends the meeting, signs the report, owns the action and sits inside the escalation chain. Yet the work may still be poorly controlled. Status is late, evidence is selective, dependencies are missed and decisions drift because no one can hold the full system in view.

Agents expose this weakness. They can be persistent where people are episodic. They can inspect every action rather than the sample presented at a meeting. They can compare a decision with hundreds of downstream tasks, maintain a complete history and surface an exception when a threshold is crossed.

That does not make them accountable. It makes them potentially useful instruments of accountability.

The central design question is therefore not, “Can the agent replace the programme manager?” It is, “Which delivery decisions may be delegated, under what authority, with what evidence, and to whom does the exception return?”

This is a governance question. If the profession leaves it to technology suppliers, procurement teams or isolated experiments, we will deserve the operating models imposed upon us.

4. The Objections Are Real — and Still Not a Strategy

The strongest case against autonomous delivery deserves respect.

Agents can fabricate a plausible answer. They can fail silently, repeat an error at scale, act on incomplete context, expose sensitive information or optimise a measure while damaging the outcome. Their reasoning may be difficult to inspect. Permissions that are harmless for a person taking one action may be dangerous for a system taking ten thousand. When several agents interact, failure can travel through a chain before a human recognises the first mistake.

There is also the human consequence. Delivery roles will change. Some will disappear. Entry-level work that taught people how programmes function may be automated before we create a new route to experience. Authority may concentrate among those who control systems and data. A transformation that ignores these effects is not responsible transformation.

But resistance is not an answer to any of this. Organisations will adopt agents because the economic pressure is obvious: more work pursued continuously, shorter coordination cycles, lower administrative cost and the possibility of scaling expertise without scaling every layer around it. A profession that offers only caution will be bypassed by leaders who hear only the opportunity.

The risks are not reasons to preserve the current model. They are requirements for the next one.

“Autonomy without boundaries is negligence. Boundaries without autonomy are theatre.”

5. Replace Job Descriptions with Authority Envelopes

Autonomous delivery requires a clearer operating contract than most human programmes currently possess.

Every agent should operate inside an authority envelope: the explicit boundary of outcomes, data, tools, expenditure, communications and decisions it may touch. The envelope must define what the agent may execute, what requires human approval, what must be escalated, and what is prohibited.

Decision class	Agent role	Human role
Reversible, low-impact action	Execute and record	Review by exception
Action within agreed tolerance	Execute, test and continue	Set tolerance and monitor
Material commitment or external promise	Prepare recommendation	Approve before action

Decision class	Agent role	Human role
Ambiguous policy or conflicting objective	Stop and present options	Exercise judgement
Safety, legal, ethical or employment consequence	No autonomous decision	Retain accountable authority
Unexpected behaviour or lost evidence	Suspend sequence	Investigate and re-authorise

This is more rigorous than placing a named person beside every action. It forces the organisation to specify decision rights before execution begins. It also reveals how much present governance depends on ambiguity: a human uses relationships, memory and tacit permission to navigate work the formal model never resolved.

Agents will not tolerate that ambiguity gracefully. They will expose it through stoppage or, worse, through confident action. That exposure is an opportunity to make transformation governance real.

6. Build Evidence into the Work

The promise of autonomous delivery is not speed alone. It is the possibility that every meaningful action carries its own evidence.

An agent can record:

- the objective it was pursuing;
- the sources and version of information used;
- the action selected and alternatives rejected;
- the authority under which it acted;
- tests performed and results observed;
- confidence, exception and escalation;
- the next action triggered.

This record should not become another ocean of logs that nobody reads. It should support three levels of control.

1. Transaction control: Was this action permitted and correctly executed?
2. Sequence control: Is the workstream progressing within outcome, cost and risk tolerances?
3. System control: Are repeated local decisions moving the transformation toward the intended enterprise result?

Humans should govern the second and third levels, with the first increasingly automated. Where patterns indicate drift, humans intervene in the rule, authority envelope or objective rather than manually correcting each transaction.

The shift is profound. Today, assurance often arrives after work has been performed and summarised. In an agentic model, assurance can become continuous and embedded—provided the evidence is designed at the beginning.

7. Redefine the Human Contribution

The profession must be candid about what remains uniquely valuable.

It is not typing the update, moving the date or convening another meeting. It is not being the human message bus between teams. It is not translating ambiguity into heroic personal effort so the organisation can avoid resolving it.

Human contribution sits where consequences, values and contested interests meet:

- choosing outcomes when stakeholders want incompatible things;
- deciding what risk is acceptable and for whom;
- reading power, trust and readiness that are not present in formal data;
- negotiating commitments no workflow can compel;
- recognising when the stated objective has become the wrong objective;
- accepting accountability for decisions with human consequence;
- creating meaning when people experience change as loss.

These are not decorative “soft skills”. They are the work that remains when execution becomes abundant.

The profession’s future is not to remain in every loop; it is to decide which loops deserve a human.

8. Our New Professional Commitments

We need a manifesto because incremental role adjustment will not be enough. The transformation profession should commit now to seven principles.

1. We will design for autonomous execution, not bolt agents onto human bureaucracy. A faster status report does not justify the technology. Workstreams must be reconsidered from outcome to evidence.
2. We will make decision rights machine-clear and human-accountable. Every delegated action will have an authority envelope and a named accountable executive.
3. We will preserve human judgement where consequences demand it. Employment, ethics, safety, material commitments and unresolved policy will not be handed to convenience.
4. We will require traceable evidence, not persuasive fluency. An agent’s confidence is not proof. Actions must connect to sources, tests and permissions.
5. We will measure outcomes rather than agent activity. Thousands of completed tasks can still produce the wrong transformation.
6. We will redesign professional development. If agents perform the administrative work through which people once learned, we will create deliberate apprenticeships in judgement, diagnosis and governance.
7. We will share the productivity gain honestly. Lower coordination cost should fund better capability, resilience and service—not merely conceal new overload.

These principles are not a complete control model. They are the position from which one can be built.

9. A Concrete Starting Point

Organisations should not begin by giving an agent an entire critical programme. They should select one bounded workstream with high coordination effort, reliable source material, reversible actions and measurable outcomes.

A practical first design might allow an agent to:

- decompose approved scope;
- maintain the integrated plan;
- gather evidence from named repositories;
- prepare internal communications;
- test completion against defined criteria;
- update risks and dependencies;
- recommend, but not approve, changes to cost, scope or external commitments.

Set a thirty-day authority envelope. Record every intervention. At the end, do not ask only how many hours were saved. Ask:

- Which decisions were safely delegated?
- Where did ambiguity cause stoppage or error?
- Which human interventions added judgement rather than clerical correction?
- Did evidence improve?
- Did the workstream outcome move faster without increasing risk?
- What authority can be expanded, and what must remain constrained?

In the opening composite pilot, nine interventions are not evidence of failure. They are the map of where the operating model is still unclear. The next cycle should reduce avoidable intervention by improving scope, policy and evidence—not by lowering the threshold for human concern.

10. The Window Is Closing

The profession has survived previous waves of methods, platforms and automation by absorbing them into familiar roles. Agents are different because they do not merely change the tool used to perform work. They can change who—or what—performs the sequence.

We can resist and become custodians of meetings, reports and controls that machines increasingly execute. We can embrace without discipline and become sponsors of fast, opaque failure. Or we can lead: define the boundaries, retain accountable judgement, redesign work and make evidence inseparable from action.

That is the next frontier. It is not the disappearance of transformation professionals. It is the end of transformation professionalism built on compulsory human coordination.

We should not wait to be asked whether agents belong in delivery. They are arriving. The question is whether the profession will shape autonomous delivery—or merely be reshaped by it.