

BEYOND COORDINATION

The Programme Manager Is Over: Autonomous Delivery Needs a New Leadership Contract

Why treating AI agents as another deployment will automate transformation's dysfunction.

Giovanni Leonardi

July 2026



The Programme Manager Is Over: Autonomous Delivery Needs a New Leadership Contract

Transformation · No. 440
July 2026

*"When the machine can coordinate the plan,
leadership begins where coordination ends."*

We Are Automating the Wrong Thing

The weekly transformation meeting begins with 42 workstreams, eleven people maintaining the integrated plan, and a 96-page status pack assembled from updates that were already stale when they arrived. Two decisions are waiting for sponsors. Seven dependencies are described differently by the teams on either side. The programme is busy, visible and slow.

Now an AI agent is added. It gathers the updates, reconciles the milestones, drafts the commentary and produces the pack before breakfast. The meeting becomes cheaper. The decisions are still waiting. The dependencies are still disputed. The programme is now busy, visible and slow at machine speed.

This is the danger in front of us in 2026. We are treating AI agents as productivity tools for the programme office when they are becoming participants in the delivery chain: able to interpret a goal, plan a sequence, use permissioned tools, call other agents, act on live information and return with evidence. If we insert that capability into yesterday's operating model, we will not transform delivery. We will automate its dysfunction.

The conventional programme manager — owner of the master plan, collector of status, broker of dependencies and centre of coordination — is over. The title may survive. The monopoly must not.

The Job Was Built for Information Scarcity

Programme management grew around a practical constraint: no one could see the whole system without an organisation dedicated to assembling the view. Information travelled through meetings, spreadsheets, presentation packs and personal networks. Coordination was expensive, so authority accumulated around the people who could perform it.

Agentic systems weaken that constraint. They can monitor delivery signals continuously, compare commitments with evidence, trace dependencies across repositories and workflows, prepare options, and escalate exceptions without waiting for a weekly reporting cycle.

The mechanism matters. Agents change delivery not because their prose is fluent, but because they can close four loops that human coordination leaves open:

- The sensing loop: collecting current evidence from the places where work happens, rather than asking people to restate it.
- The reasoning loop: comparing evidence with outcomes, constraints and dependencies, then identifying the decision that is actually blocked.
- The action loop: carrying out bounded, reversible actions through authorised tools instead of merely recommending them.
- The learning loop: retaining the result, including the exception and the human judgement, so the next cycle starts better informed.

A programme office that uses agents only to write status reports closes none of these loops. It merely improves the typography of delay.

When the machine can coordinate the plan, leadership begins where coordination ends.

End the Coordination Monopoly

The serious claim is not that agents can replace judgement. They cannot. It is that coordination no longer justifies concentrating judgement, information and action in one managerial role.

A familiar pattern exposes the distinction. In a composite programme, an agent was permitted to read delivery systems, test milestone claims against completed work, and flag contradictions. In its first fortnight it processed roughly 1,800 delivery signals and found 63 dependency mismatches. The old weekly process had reported nine.

The important result was not the larger number. It was what happened next. Forty-one mismatches were resolved directly by the teams because the evidence and ownership were clear. Fourteen required a choice between scope and date. Eight exposed disagreements about the intended outcome. The agent accelerated coordination; it could not settle purpose, appetite or consequence.

That is the line we must draw. Machines should carry the coordination load. Humans must carry the accountability load.

Human Work Becomes Harder

There is a comforting story that AI will remove administrative work and leave leaders free for leadership. It is directionally right and practically incomplete. The administrative burden often protected us from the harder work: naming the trade-off, confronting a sponsor, stopping low-value activity, or admitting that two promised outcomes cannot both be achieved.

Autonomous delivery removes some of that shelter.

Old programme contract	Agent-enabled leadership contract
Maintain the integrated plan	Define the outcome and its non-negotiable constraints
Collect and reconcile status	Demand evidence and resolve contested meaning
Broker every dependency	Set rules for local resolution and intervene on true exceptions
Chair governance forums	Make consequential decisions at the point of need
Control change through process	Bound autonomy by reversibility, value and risk
Measure activity and milestones	Hold the system to realised outcomes

This is not less leadership. It is leadership without clerical camouflage.

The strongest objection deserves respect. Complex transformations need a person who holds the whole, understands the politics, protects coherence and remains answerable when automated decisions cause harm. Remove that centre, the argument goes, and accountability fragments just as autonomy expands.

That objection is correct about the need and wrong about the design. Accountability must be whole; coordination need not be centralised. One named human should remain answerable for the mission. But the same person should not also be the routing layer through which every update, dependency and decision must pass. When responsibility depends on informational monopoly, it is not accountability. It is bottleneck management.

Write a New Leadership Contract

Autonomous delivery requires explicit human and machine jurisdictions. Not a vague promise that a person remains “in the loop”, but a contract precise enough to govern action.

We should insist on five principles:

- Outcomes before tasks. Agents receive a measurable purpose, constraints and evidence standards — not an endless queue of disconnected instructions.
- Autonomy follows reversibility. Low-consequence actions that can be undone may proceed within bounds. Irreversible, regulated, financially material or people-affecting decisions require named human authority.

- Evidence replaces narration. Delivery claims must point to observable change. A generated confidence statement is not evidence simply because it is current.
- Exceptions travel to decisions. Escalation should identify the conflict, the options, the consequence of delay and the human who owns the choice.
- Every agent has an accountable steward. Someone owns its permissions, objectives, monitoring, failure modes and retirement. An agent without a steward is unmanaged organisational power.

This contract redistributes the work. The mission owner defines value and accepts trade-offs. Domain leaders own the quality of decisions in their territory. An agent steward controls the digital workforce and its boundaries. An integration lead protects end-to-end coherence. Assurance tests whether actions remain within policy and risk appetite. These may be roles, not new departments. What matters is that the decisions are explicit and no longer hidden inside a generic programme-management mandate.

Do Not Build a Faster Programme Office

Leaders do not need another technology deployment. They need to redesign delivery around a new fact: parts of the organisation can now sense, reason and act without waiting for a person to translate every intention into a task.

The first moves are organisational, not technical:

1. Identify where the programme manager is acting as a human message bus — collecting, translating, chasing and routing information.
2. Give agents access to evidence before asking them to generate commentary. If they cannot see where work happens, they will manufacture certainty from reports.
3. Delegate one bounded decision loop end to end. Define the objective, permissions, stop conditions, evidence and human escalation point.
4. Remove the meeting, report or approval that the loop makes unnecessary. Keeping the old control while adding the new capability creates double bureaucracy.
5. Review failures as operating-model evidence. Ask whether the objective, boundary, data, authority or escalation was wrong before blaming the model alone.

The test is not how many agent pilots exist. It is whether decisions move faster without becoming less accountable, whether teams resolve more dependencies locally, and whether leaders spend less time reconstructing the past and more time choosing the future.

Stop Protecting the Role

Some programme managers will become exceptional leaders of autonomous delivery. They understand systems, sequencing, organisational friction and the consequences of weak governance. But they will succeed only by surrendering the very activities that once made the role indispensable.

Others will defend the pack, the meeting, the master plan and the central office because those artefacts prove that coordination is happening. Agents will then be trained to feed the machinery. Costs may fall. Cycle time will not. Responsibility will become harder to locate because machine-generated activity will create the appearance of control.

We should choose differently.

Let agents track the work. Let them reconcile the evidence, surface the exception and execute the bounded response. Let teams resolve what no longer needs central permission. Reserve human authority for purpose, conflict, ethics, consequence and commitment.

The programme manager as the organisation's coordination engine has reached its end. What replaces it is not leaderless automation. It is a stricter form of leadership: fewer rituals, clearer rights, bounded autonomy and unmistakable accountability.

That is the break with convention now required.