


BEYOND LAYER COUNTING

What Sits Between Strategy and Agents?

The most compelling restructuring cases in a decade are solving for the right cost and the wrong function

Giovanni Leonardi

April 2026



What Sits Between Strategy and Agents?

Leadership · No. 115

April 2026

“This is talent work of the most consequential kind, and it is invisible to any system that measures management by its outputs rather than its judgements.”

The redesign deck lands on a Monday. Two layers of management removed, three hundred roles reclassified, and a projected annual saving that makes the business case self-evident. The agentic platform handles task routing, status aggregation, and the weekly reporting cadence that once consumed a third of every team lead's calendar. The chief operating officer calls it the most compelling restructuring she has seen in a decade.

By Thursday, the first questions arrive — not from the people whose roles are marked for removal, but from the engineers and analysts two levels below them. Who approves my development plan? Who do I talk to when two priorities collide and neither product owner will blink? Who notices when I am struggling before I have missed a deadline?

The questions are not about reporting lines. They are about the human work that lived inside them — and that no one had mapped before deciding to delete the layer.

1. The Flattening Thesis — And Its Blind Spot

The argument for removing middle management has never been stronger, and it has never been more dangerous.

The strong version is straightforward: a significant proportion of what middle managers do is information relay. They translate strategy downward into tasks. They aggregate status upward into dashboards. They attend meetings on behalf of people who cannot attend all the meetings. In any organisation of moderate complexity, a remarkable share of managerial time is consumed not by judgement or development but by the mechanical work of moving information between layers that cannot see each other directly.

Agentic systems do this work cheaply, reliably, and without the latency that turns a three-day decision into a three-week one. When a platform can decompose an objective into tasks, assign them based on capacity and skill, track completion, surface exceptions, and compile the reporting that three layers of management once produced by hand — the business case for those layers becomes difficult to defend. The flattening thesis is not wrong about the diagnosis. A great deal of middle management was overhead, and the honesty to say so is overdue.

But the diagnosis is not the remedy. And the remedy being adopted across the industry — compress the layers, redistribute the surviving work upward and downward, let the

platform handle the connective tissue — is built on a category error so fundamental that it will take two years of organisational damage to surface.

The flattening thesis treats the reporting function and the human function as though they are the same role. They were never the same role. They happened to live in the same person because organisations never separated them — and because the reporting function was visible while the human function was not.

2. Two Roles, One Title

Consider what a competent middle manager — not a bureaucratic one, a competent one — actually does across a working week.

Some of it is relay: cascading objectives, compiling updates, chairing alignment meetings, translating between the vocabulary of the executive corridor and the vocabulary of the team. This is the work that agents handle well, and the work that made the role look like pure overhead to anyone measuring it from above.

But alongside the relay work, a different kind of labour was happening — labour that rarely appeared in any job description and almost never surfaced in a restructuring review.

A programme director in a large financial services transformation — the kind of role now routinely marked for compression — might spend a Tuesday morning reviewing a workstream lead's approach not because the status was red, but because something in the last two stand-ups suggested the lead was solving the wrong problem and had not yet realised it. That intervention — a thirty-minute conversation that redirects three weeks of effort — does not appear in any dashboard. No agent would have flagged it. It required pattern recognition built across years of watching people work, and the relational trust to deliver a course correction without it landing as a rebuke.

The same director might spend Thursday afternoon in a conversation that has no project code at all: a senior analyst is being considered for a leadership role elsewhere in the organisation, and the director is the person who knows both the analyst's capability and the receiving team's culture well enough to judge whether the move will develop or destroy her. This is talent work of the most consequential kind, and it is invisible to any system that measures management by its outputs rather than its judgements.

What agentic platforms cannot replicate is not information processing. It is the holding work — the absorption of ambiguity, the development of people, the maintenance of context that spans organisational boundaries, and the exercise of judgement in situations where the data does not tell you what to do because the situation has not fully declared itself.

3. The Inconvenient Selection Problem

Here is where the even-handed account becomes uncomfortable for both sides.

The flattening evangelists are right that many managers were never doing the human work at all. They were hired for their technical competence, promoted into management because no other career path existed, given no development in the skills the role actually demanded, and evaluated on the one thing that was easy to measure: whether information moved on time. These managers were the relay function, and replacing them with agents is not a loss — it is a correction of a staffing error that persisted for decades because the cost of the error was distributed and invisible.

But the managerial nostalgists are equally wrong to defend the layer as a category. The argument that middle management is "essential" slides past the question of whether any particular organisation's middle management was ever performing the essential function. In many cases it was not — and the managers themselves knew it, which is why the flattening thesis landed with such force. It described a reality that practitioners at every level recognised but could not say aloud.

The honest position sits between the two, and it is harder to act on than either: the relay function is overhead and should be automated; the human function is load-bearing and should be preserved — but only where it was actually being performed, and only in people who were selected and developed for it. This means the redesign question is not how many layers to keep, but what judgement and care work must remain human at each level, and whether the people currently occupying those levels were ever chosen for their capacity to do it.

We are not, in other words, facing a structural question about layers. We are facing a talent question about what was inside them — and it is a question most organisations have never asked, because the answer requires an honesty about managerial capability that few leadership teams are prepared to exercise.

4. What Compression Actually Demands

The organisations I see handling this well — and they are fewer than the conference circuit suggests — are doing something more difficult than deleting layers. They are redesigning the managerial role around the work that agents cannot do, and then selecting for it with a rigour that would have been unthinkable five years ago.

In practice, this means three things.

1. The separation of the relay function from the human function becomes explicit. Status aggregation, task decomposition, resource scheduling, and routine escalation move to the platform. What remains in the human role is context-holding, people development, cross-boundary judgement, and the kind of ambiguity absorption that requires someone to sit in a situation long enough to understand what is actually happening before deciding what to do about it.
2. The evaluation framework changes. Managers are assessed not on throughput or reporting cadence but on the development trajectories of their people, the quality of the judgements they made under ambiguity, and their ability to hold context across organisational seams that the platform cannot see. This is harder to measure, which

is precisely why it was never measured — and why so many organisations ended up with relay managers in roles that demanded development managers.

3. The hardest step: acknowledging that some current managers cannot make the transition. The skills that made someone an effective relay manager — diligence, process discipline, reliable execution of defined procedures — are not the skills that make someone effective at the residual human role. Some managers will grow into it. Others were never hired for it, never developed for it, and cannot be expected to perform it simply because the rest of their role has been automated away.

This is not a comfortable message. It asks leadership teams to make judgements about individual capability that most organisational cultures are designed to avoid. But the alternative — compressing the layers, hoping the human work redistributes itself, and discovering eighteen months later that it simply vanished — is not more humane. It is merely more deniable.

5. The Question That Remains

The middle manager debate as currently framed — do we need them or don't we? — is the wrong question asked at the wrong altitude. Of course the relay function is being automated; it should have been automated years ago, and the only reason it persisted was that the technology to replace it had not yet arrived and the cost of the human workaround was tolerable.

But "middle management" was never one thing. It was two things wearing the same badge, and the flattening moment is forcing a separation that was always overdue. The organisations that treat this as a cost exercise — how many layers can we remove? — will find that they have cut away capability they did not know they had. The organisations that treat it as a design exercise — what human judgement must sit at each level, and who is capable of exercising it? — will build something more resilient than anything the old hierarchy produced.

The layer that merely relayed was always overhead. The layer that developed people, held context, and absorbed ambiguity was always load-bearing. The only thing that has changed is that we can no longer afford to pretend they were the same thing.