


BEYOND THE EIGHTH WAVE

What Twenty-Five Years of Transformation Taught Me About the Next Five

Eight technology waves, five constants, and a credo for leading through the agent era

Giovanni Leonardi

July 2026



Contents

1. Executive Summary

2. The Announcement We Have Heard Eight Times

3. What Every Wave Was Actually Testing

3.1 Operating model before technology

3.2 Decision rights before org charts

3.3 Data honesty before dashboards

3.4 Absorption capacity before ambition

3.5 Named human accountability before every other control

4. The Objection That Deserves an Answer

5. The Arithmetic of Scar Tissue

6. What We Commit To

7. The Next Five Years

3

3

4

4

4

5

5

5

6

7

7

8

What Twenty-Five Years of Transformation Taught Me About the Next Five

Leadership · No. 124
July 2026

"The agent does not remove the need for an operating model. It removes the margin of error that let us survive without one."

1. Executive Summary

Eight waves in twenty-five years, and each one arrived with the same announcement: the rules of delivery have changed. ERP said it. Offshoring said it. Agile, cloud, digital, machine learning, generative AI — each said it. Agents are saying it now, louder than any of their predecessors. This manifesto is written against that announcement. It argues that every wave changed what we build and none of them changed what makes building succeed, and that the five constants which decided the outcome in 2001 will decide it again in 2031: operating model before technology; decision rights before org charts; data honesty before dashboards; absorption capacity before ambition; and named human accountability before every other control. It takes the strongest version of the "this time is different" case seriously — that agents act rather than merely assist — and answers it: an actor without an operating model is not a transformation, it is an incident waiting to be named. It closes with the only leadership stance the next five years reward: not the leader who predicts the technology best, but the one who refuses, one more time, to let its novelty excuse the fundamentals.

2. The Announcement We Have Heard Eight Times

We have been told the rules have changed before.

We were told it when the first enterprise resource planning programmes promised a single version of the truth, and we discovered that a single system cannot manufacture agreement about what the truth is. We were told it when work crossed oceans on the strength of a rate card, and we learned that a process nobody had written down cannot be handed to anyone, at any price. We were told it when agile arrived and the plan was declared dead, and we watched organisations discover that a backlog is a plan with the dates removed. We were told it when the data centre became someone else's building, when every company became a technology company, when the model could predict, when the model could write, and now — in the summer of 2026 — when the model can act.

Eight waves. Eight announcements. And beneath every one of them, the same five things decided whether the money turned into anything.

We are not cynics about the technology. The waves were real; each of them changed what was possible, and the organisations that sat them out paid for it. But we have become, across a quarter of a century, very hard to impress by a claim that the fundamentals have been suspended — because we have heard the claim made with total conviction eight times, and we have stood in the room afterwards, eight times, when the post-mortem reached the same page.

This is that page, written down once, so that we do not have to rediscover it in 2029.

The waves changed what we build. They never changed what makes building succeed. A technology that can do more does not reduce the need for the fundamentals; it raises the cost of neglecting them, because the failure arrives faster and at greater scale.

3. What Every Wave Was Actually Testing

Strip the vocabulary from each wave and the same examination is being set. The technology is the question paper. The organisation's answer is always written in the same five subjects — and it is always the same subjects that are failed.

3.1 Operating model before technology

The first constant is the oldest and the most stubbornly ignored. A technology is a capability; an operating model is the arrangement by which an organisation turns capability into outcome — who does what, in what sequence, with which hand-offs, under which rules. Every wave has been sold as a capability and bought as though the operating model would follow.

It does not follow. It has to be built, and it has to be built first, because the technology will be shaped to whatever operating model it finds. Configure a system around undesigned processes and the system faithfully automates the confusion. Offshore an undesigned process and you export the confusion to a time zone that cannot ask the person in the next seat what it means. Deploy an agent into an undesigned process and you have given the confusion the ability to execute without waiting.

We have watched this pattern in its 2026 form: a function that could not state, on one page, who owned each step of its intake process — but that had already licensed a platform to "agentify" it. The platform was not the problem. The platform was an extremely efficient way of discovering that the problem had never been solved.

3.2 Decision rights before org charts

The second constant is the one leaders most readily agree with and least readily act on. Every wave is accompanied by a restructure: the ERP centre of excellence, the offshore captive, the agile tribe, the cloud platform team, the digital factory, the AI centre. Boxes

are redrawn; reporting lines move; the org chart is published with some ceremony. And the programme stalls in exactly the same place as before, because nobody redrew the thing that actually governs behaviour — who may decide what, with whose money, and who can say no.

Org charts describe accountability at rest. Decision rights describe it in motion. It is the motion that delivery needs. We have seen a reorganisation that created a new, well-funded transformation function and left every spending decision above a modest threshold with the same three finance directors who had blocked the previous programme. The chart said the function was empowered. The decision rights said it was not. The decision rights were right.

3.3 Data honesty before dashboards

The third constant becomes more dangerous with each wave, because each wave makes the dashboard more beautiful. A status report in 2001 was a spreadsheet that could be argued with. A portfolio dashboard in 2026 renders in real time, draws on forty sources, and is summarised in confident natural language by a model that has been asked to be helpful. The honesty of the underlying data has not improved at the same rate as the confidence of its presentation — and that gap is where programmes die.

Data honesty is not a technology question; it is a cultural one. It means that the person who knows the milestone will be missed says so in the week they know it, not the week it becomes undeniable. Every layer of abstraction between the person who knows and the person who decides is a place where green can be applied to red. We have seen a programme whose twelve workstreams reported green for nine consecutive months while its integration test plan quietly lost a quarter of its scope. The dashboard was accurate. It reported what it had been told.

3.4 Absorption capacity before ambition

The fourth constant is the one we are worst at measuring, and so the one we most often violate. An organisation can only absorb so much change in a given period — so many new ways of working, so many retrained roles, so many altered controls — before it stops absorbing and starts merely enduring. Ambition that exceeds absorption does not produce faster transformation. It produces a backlog of half-adopted changes, each one undermining the next.

Every wave has tempted leadership to set ambition by what the technology could do rather than by what the organisation could take in. The test is simple and rarely applied: for the population that has to change, how many material changes have landed on them in the last twelve months, and how many of those are actually embedded rather than merely announced? We have seen a branch network asked to absorb its fourth customer-facing system in three years while the second was still being worked around with paper. The fourth was the best of the four. It made no difference.

3.5 Named human accountability before every other control

The fifth constant is the one the agent era puts most directly under pressure, and the one that matters most. Every control we have ever designed — stage gates, steering committees, risk registers, assurance reviews, model governance, audit trails — is a

mechanism for making a named human being answerable for a consequence. Controls that do not end in a name are not controls. They are documentation.

Each wave has offered a new place to hide the name. The vendor's methodology. The offshore partner's SLA. The self-organising team. The platform. The model. Now the agent. And each time the hiding place has been discovered only when something has gone wrong and the room has turned to ask who owned it — and nobody did.

*"A control that does not end in a named human being
is not a control. It is a description of the accident we
have agreed to have."*

4. The Objection That Deserves an Answer

There is a serious version of "this time is different", and it would be dishonest to caricature it.

It goes like this. Every prior wave gave people better tools; this one gives organisations additional actors. An ERP system waited to be used. An agent does not wait. It plans, executes, and chains actions across systems at a pace and scale no human workforce can match. The argument is that a doctrine written for a world where humans did the work cannot simply be carried forward into a world where software does — that operating models, decision rights, and accountability are concepts built around human agents, and that applying them to synthetic ones is a category error of the same kind as applying factory discipline to knowledge work.

We take the objection seriously because its premise is correct. Agents do act. That is exactly what makes the constants more binding, not less.

An actor without a defined operating model is not a transformation; it is an incident waiting to be named. Consider what it means to give a system the ability to act inside a process nobody has written down: it will act on the implicit version, the one that lives in the habits of the three people who have done the job longest — and it will act on it at scale, at speed, and without the judgement those three people apply when the implicit rule does not fit the case. The agent does not remove the need for an operating model. It removes the margin of error that let us survive without one.

The same is true of decision rights. A human who lacks clear authority hesitates, asks, escalates. An agent that lacks a clear boundary on its authority does not hesitate; it proceeds to the limit of its permissions, which are usually far wider than anyone intended because they were set for convenience rather than governance. The question "who may decide what" does not become obsolete when a synthetic actor enters the room. It becomes the single most important thing to have answered before it does.

And accountability. The objection's hidden hope is that agents might carry some of it — that the model, the platform, or the vendor becomes the place the name lives. That hope

is the ninth hiding place, and it is no more real than the eight before it. When a chained sequence of automated actions produces a consequence a regulator, a customer, or a board wants explained, the explanation will be demanded from a person. Either that person knew they were accountable before the sequence ran, and designed for it – or they find out afterwards, in the room.

So yes: this time the actors are different. The examination is the same. The paper is harder.

5. The Arithmetic of Scar Tissue

We are not writing this from foresight. We have no better view of what agents will do by 2031 than anyone else, and we are suspicious of anyone who claims one. We are writing it from scar tissue – from the accumulated, repeated, slightly embarrassing experience of watching the same five things decide the outcome, wave after wave, while the conversation in the room was about something else.

The arithmetic is worth stating plainly. Across the programmes we have seen turn from red to something that lasted, the recovery almost never came from a change of technology. It came from one of five moves, and usually from two of them together: the operating model was finally drawn; a decision right was finally moved; a true number was finally put in front of the people who could act on it; the ambition was finally cut to what the organisation could absorb; or a name was finally attached to a consequence. None of those moves cost much. All of them were available at the start. Each of them had been deferred because the technology was new, and novelty felt like a reason.

Novelty is not a reason. It is a feeling. Every wave has persuaded intelligent leaders that the fundamentals could wait until the technology had proved itself – and every wave has proved itself by exposing, expensively, which fundamentals were missing.

The leaders who came through the earlier waves with their credibility intact were not the ones who predicted the technology best. Some of them predicted it badly. They were the ones who treated each wave as a new examination in the same five subjects, and refused to let the excitement in the room become a reason to skip the revision.

6. What We Commit To

A manifesto is a set of refusals. These are ours, for the next five years, and we offer them as a standard to be held to rather than a summary to be admired.

1. We will not buy a capability before we have drawn the operating model it serves. If the process cannot be put on a page – steps, owners, hand-offs, exceptions – the technology is not ready to be deployed into it, however ready the technology is.
2. We will not publish an org chart in place of a decision-rights map. For every material decision the transformation depends on, we will be able to name who decides, with

whose budget, and who holds the veto — and we will move those rights before we move the boxes.

3. We will not let the dashboard be more confident than the data. We will design reporting so that the person who first knows a thing is the person who first reports it, and we will treat nine consecutive months of green as a risk in its own right.
4. We will set ambition by absorption, not by capability. Before we commit the organisation to a change, we will count the changes already on the population that must absorb it, and we will sequence rather than stack.
5. We will name a human being before we name any other control. No agent, model, platform, partner, or team will be allowed to stand where a name should stand. If we cannot say who is accountable for a consequence, we will not permit the action that could produce it.

And one commitment that sits underneath all five:

6. We will not let novelty excuse the fundamentals — one more time. We have been offered the excuse eight times. We will be offered it again, more persuasively than ever, by a technology that can genuinely do remarkable things. We will admire the technology and decline the excuse.

7. The Next Five Years

The agent era will bring real change. Some of the work we spent careers doing will be done differently, faster, and by something other than us. That is not a threat to the doctrine; it is the doctrine's next test. Organisations will be sorted, as they always have been, not by who adopted the technology first but by who had the five constants in place when it arrived.

We know how this goes, because we have seen it go this way eight times. The organisation that has drawn its operating model will deploy agents into processes that can hold them. The organisation that has mapped its decision rights will know exactly what an agent may and may not do, and why. The organisation that tells itself the truth will notice when the automation drifts, in the week it drifts. The organisation that has respected its own absorption capacity will embed what it adopts. And the organisation that puts a name before every control will be the one standing in the room, when the question is asked, with an answer.

The rest will discover, one more time, that the rules had not changed. They never do. Only the cost of forgetting them does.