

BEYOND THE SIGN-OFF

The Death of the Approval Chain: What Agentic Workflows Reveal About How Organisations Actually Change

Why removing the human checkpoints tells us more about our organisations than about our machines

Giovanni Leonardi

July 2025



Contents

1. Executive Summary	3
2. The afternoon the wall came down	3
3. What the chain was actually for	4
4. What agents dissolve, and what they only relocate	5
5. The friction did not leave. It moved.	6
6. What this actually reveals about how organisations change	7
7. Living without the chain	8
8. Coda	9

The Death of the Approval Chain: What Agentic Workflows Reveal About How Organisations Actually Change

Transformation · No. 419

July 2025

“Organisations change when the distribution of accountability changes — and only then. Everything else is adoption, not change.”

1. Executive Summary

Something quiet has happened to the way work gets authorised. The chain of human sign-offs that used to govern how a change moved through an organisation — the request, the review, the second review, the board, the release gate — is dissolving under a class of software we have started calling agentic: systems that do not merely suggest an action but plan it, execute it, and close it out end to end. The instinctive reading is a story about speed, and speed is certainly the headline. But the more interesting story is what the dissolution exposes. The approval chain was never primarily a mechanism for making better decisions; most of its links changed nothing. It was an architecture for distributing accountability, manufacturing legibility, and pacing change to the speed of human comprehension. When agents collapse the execution, those older functions do not vanish — they are forced into the open, and we discover we had been leaning on them far more than we admitted. This essay argues that agentic workflows are less a productivity event than a diagnostic one. They reveal that organisations change not when they adopt new tools but when the distribution of accountability shifts, and that this is precisely the place organisations resist hardest. The death of the approval chain is not the death of approval. It is the end of pretending the chain was ever really about the decision.

2. The afternoon the wall came down

Consider a change that, until recently, would have been unremarkable in its slowness. A minor adjustment to how a customer record is validated: a day of development, then the queue. It needed the team lead's sign-off, then a peer reviewer, then the change-advisory board's weekly slot, then a security acknowledgement, then a release window, then a post-implementation review. Eleven names touched it before it reached production, and the elapsed time from idea to live was a little under three weeks. Nobody thought this strange. It was simply how change of any consequence moved: through people, in sequence, each adding a signature to a growing ledger of assent.

Now the same change is proposed by an agent that has read the ticket, written the code, generated the tests, run them, checked the result against a policy file, and opened the path to production — in an afternoon. The eleven names are still notionally in the process, but most of them are now looking at a summary the agent produced, of work the agent did, validated by checks the agent ran. The wall that used to stand between intention and execution has not been climbed. It has come down.

The reaction inside organisations to this moment is rarely straightforward celebration. There is relief at the speed, and beneath the relief, an unease that is harder to name. The unease is the subject of this essay, because it is pointing at something true. We are not merely uncomfortable that the work is faster. We are uncomfortable because the speed has stripped the paint off a structure we had stopped looking at, and we do not entirely like what is underneath.

3. What the chain was actually for

To understand what agents are dissolving, we have to be honest about what the approval chain was doing, which is not the same as what it claimed to do. Its stated purpose was quality: many eyes catch what few eyes miss. Its actual purposes were several, and only the first was about the decision itself.

- It distributed accountability. A change that eleven people approved was a change no single person had to own. The chain was, among other things, a device for spreading risk across enough shoulders that no one shoulder would break under it. This is why approval processes grow after incidents and almost never shrink: each new signature is someone else who can be pointed to.
- It manufactured legibility. The ledger of sign-offs was an audit trail — a story the organisation could tell afterwards about who knew what and when. In regulated environments this was not optional, and even outside them the trail was how the institution reassured itself that it was in control.
- It paced change to human comprehension. The sequence was slow partly because slowness was doing work: it gave the organisation time to notice, object, and absorb. The gaps between the gates were where quiet vetoes happened — the senior engineer who saw the pull request and had a word before the board ever met.
- It performed assurance. A great deal of the ritual existed to be seen existing. The board that met weekly to approve changes was, in no small part, a signal to auditors, regulators, and executives that changes were being approved by a board that met weekly.

Now the awkward figure. In one environment I looked at closely, a change-advisory board reviewed roughly forty changes a week. When the records were examined honestly, more than nine in ten were approved without a single modification, and the median time spent on each was under ninety seconds. The board was not, in any meaningful sense, deciding. It was ratifying. The genuine engineering judgement — the part that actually altered outcomes — had happened earlier and elsewhere, in the review between two people who understood the code. The board was the assurance layer, the legibility layer, and the accountability-spreading layer, wearing the costume of the decision layer.

The uncomfortable truth that agents force into view is that most of the approval chain was never adjudicating anything. It was distributing blame, generating paperwork, and slowing the world down to a speed at which humans felt in control. Take those functions seriously and you can automate the chain. Pretend they were about decision quality and you will automate the wrong thing.

This is worth sitting with, because it explains the unease. When an agent replaces the ninety-second ratification, it removes almost no decision-making, because almost none was happening there. What it removes is the feeling of control that the ritual supplied. And a feeling of control, once you have depended on it for years, is not nothing to lose.

4. What agents dissolve, and what they only relocate

It is tempting to conclude that if the chain was mostly theatre, agents have simply cleared away a stage set nobody needed. That conclusion is too quick, and the strongest version of the opposing case deserves a hearing rather than a caricature.

The serious counter-argument runs like this: the approval chain was indeed inefficient, but its legitimate content — the rules, the risk thresholds, the compliance requirements — can be encoded. Turn the policy into a file the agent reads. Express the guardrails as code. Let the machine enforce, deterministically and at scale, what a bored human enforced erratically and slowly. On this view, agentic workflows do not weaken governance; they finally make it real, replacing a ritual that pretended to check everything with a system that actually checks the things that matter, every time, without fatigue. The board that rubber-stamped forty changes was worse governance than an agent that genuinely evaluates four hundred against an explicit policy.

There is a great deal of force in this, and where it is right we should say so plainly: for the codifiable part of governance, the machine is not merely faster but better. A policy that is executed identically every time is a real improvement on a policy that lived in the varying moods of a Thursday-afternoon committee.

But the argument contains a quiet assumption that does not survive contact with practice — that the content of the approval chain was fully codifiable in the first place. It was not. The chain absorbed two things that resist encoding. The first is tacit judgement: the reviewer who cannot articulate why a change smells wrong but is usually right, whose objection is pattern recognition trained on scars the policy file has never heard of. The second is contextual novelty: the situation the rules did not anticipate, where the correct action is to stop and ask a human precisely because no rule applies. The approval chain, for all its theatre, was a container for the unarticulated and the unprecedented. Policy-as-code encodes the rules we can state. It is silent on the rules we enforce without being able to write down, and blind to the cases that no rule yet covers.

“We automated the part of governance we could describe, and discovered how much of it we had never

learned to describe at all.”

So agents do not so much dissolve governance as sort it. The explicit, repetitive, codifiable portion is genuinely absorbed and improved. The tacit and the novel are not absorbed; they are displaced — pushed out of the chain and left, for a while, with nowhere to live. This is the mechanism beneath the unease. It is not that oversight has been abolished. It is that a specific, hard-to-replace species of human oversight has been evicted from the place it used to sit, and has not yet been rehoused.

5. The friction did not leave. It moved.

Follow the displaced oversight and you find the second revelation. Governance, when you automate its visible form, does not disappear. It relocates — upstream and downstream of the work — and in relocating it often becomes harder to see and more dangerously concentrated than it was in its slow, distributed, theatrical form.

Upstream, the real decisions migrate into the design of the agent itself: the policy file, the system prompt, the choice of which actions require a human and which do not, the definition of what counts as an exception. These are now the governing acts. But they are made once, early, often by a small technical group, and then they operate silently at scale. The weekly board that visibly ratified forty changes has been replaced by a configuration decision that invisibly governs four hundred. The governance did not shrink; it moved to a place with far less scrutiny than the board ever had.

Downstream, oversight reappears as monitoring and exception handling — the human who watches the dashboards, investigates the anomalies, and answers when the agent escalates. And here the relocation produces its sharpest pathology, the one I would put at the centre of any honest account of this moment. Call it the last-human problem.

In the old chain, accountability was spread across eleven people, each owning a small, comprehensible slice. In the new arrangement, a single release manager or platform owner is frequently named as accountable for the entire automated flow — for three hundred agent-executed changes a week that they could not possibly inspect individually, and that they are nonetheless expected to answer for when one of them causes harm. We have taken accountability that was diffuse and survivable and concentrated it onto one person who has responsibility without the practical capacity to discharge it. This is not a lightening of the governance burden. It is a transfer of an impossible burden onto a single named owner, dressed up as efficiency.

Dimension	The approval chain	The agentic workflow
Where decisions live	Distributed across many sign-offs, in sequence	Concentrated upstream in policy and configuration
Who is accountable	Diffused across everyone who signed	Named on one owner of the whole flow

Dimension	The approval chain	The agentic workflow
What is visible	The ritual — meetings, tickets, signatures	The exceptions — dashboards, escalations, logs
Speed	Paced to human comprehension	Paced to machine execution
The failure mode	Slowness and rubber-stamping	Invisible policy error at scale, and the overwhelmed last human

This is the gap between transformation intent and transformation reality that recurs in every era of automation, wearing a new costume each time. The intent was to remove friction. The reality is that friction is close to conserved: it was not removed but moved, and sometimes moved somewhere worse. Leaders who bought agentic workflows to eliminate the drag of approvals frequently find, a year on, that they have not eliminated the drag so much as relocated it into a configuration nobody reviews and a job nobody can actually do. The organisation that does not understand this ends up faster and more fragile at the same time — which, for anyone who lived through earlier waves of well-intentioned automation, is a deeply familiar place to arrive.

6. What this actually reveals about how organisations change

Step back from agents specifically, and the deeper pattern comes into view — the one that makes this more than a technology essay.

We tend to describe organisational change as a matter of adopting things: new tools, new processes, new platforms. On that account, an organisation changes when it swaps an old capability for a better one, and the difficulty of change is essentially the difficulty of the swap — training, migration, the learning curve. Agentic workflows expose how shallow this account is. The technical swap here is, if anything, easy; the tools are astonishingly capable and the migration is often trivial. Yet the change is experienced as wrenching, and it is resisted with an intensity the mere swap could never explain.

The resistance makes sense only when you see what is really changing. What agents alter is not primarily the tooling but the accountability architecture — who owns what, whose judgement counts, where blame can come to rest, how the institution reassures itself that it is in control. That architecture is the thing organisations are actually made of. Processes are its visible surface; the approval chain was one such surface. And organisations defend their accountability architecture far more fiercely than they defend any particular tool, because it is the structure through which people hold their standing, their safety, and their sense of the ground being solid.

This is why the same organisations that adopt an AI coding assistant in a week will spend a year in anguish over letting an agent merge without a human gate. The assistant leaves the accountability architecture untouched — a human still signs. The autonomous merge rewrites it. The difference in adoption difficulty is enormous, and it has almost nothing to do with the technology and almost everything to do with which of the two touches the structure of who answers for what.

- Organisations change when the distribution of accountability changes — and only then. Everything else is adoption, not change.
- They resist precisely at the points where accountability is redistributed, which is why the hardest transformations are never the technically hardest ones.
- A tool that leaves the accountability architecture intact will be adopted quickly and change little. A tool that rewrites it will be resisted fiercely and change everything.
- This is why so many transformations that look complete on the surface — the platform is live, the tool is rolled out — turn out to have changed nothing that matters. The accountability architecture was left standing, and it quietly reasserted the old behaviour through the new system.

Agentic workflows are, in this light, a forcing function. They do not let the accountability architecture stay hidden, because they cannot function without touching it. In demanding an answer to the question the approval chain let us avoid — who, actually, is accountable for this, and how do they discharge that accountability in practice — they hold up a mirror to the structure we would rather not examine. The technology is the lesser story. The mirror is the greater one.

7. Living without the chain

If the chain is going, and much of it should go, then the task is not to mourn it or to reconstruct it out of nervousness. It is to keep deliberately the functions that were real and discard the ones that were only ritual. That is a matter of judgement rather than method, but a few principles have earned their place.

1. Separate the reversible from the irreversible, and govern them differently. The great error is to apply one approval posture to everything. Most changes are reversible and their cost of error is low; these should flow through agents with monitoring and no human gate, because a human gate there buys nothing but delay. A small minority are irreversible or high-consequence; these deserve a genuine human decision — not a ninety-second ratification but an actual one. The chain's failure was applying the same ceremony to both. The opportunity is to stop.
2. Make the upstream governance as visible as the old board was. If the real decisions now live in policy files and configuration, then those artefacts must be reviewed with at least the seriousness the weekly board pretended to bring — versioned, debated, owned, and open to challenge. The worst outcome is governance that moved upstream and went dark. Bring light with it.
3. Refuse to name a last human who cannot do the job. If one person is accountable for a flow they cannot inspect, either the accountability is a fiction — in which case say so — or the flow must be designed down to a scope a human can actually stand behind. Concentrated accountability without capacity is not governance; it is a search for someone to blame, arranged in advance.
4. Preserve a home for tacit judgement and novelty. The reviewer's unarticulated unease and the case no rule anticipated need somewhere to go. Build the escalation path that lets an agent stop and ask, and protect the human expertise that answers, because that expertise is exactly the part automation cannot absorb and exactly the

part that atrophies fastest when the chain that exercised it disappears.

None of this is a methodology, and it is not meant to be. It is the shape of a stance: keep the judgement, drop the theatre, and above all be honest about which was which. The organisations that will do well are not the ones that automate fastest, nor the ones that cling longest to the human gate. They are the ones willing to look at the exposed structure without flinching and decide, deliberately, what it was really for.

8. Coda

The approval chain is dying, and the temptation is to read its death as a story about machines becoming capable enough to do without us. That reading misses the point. What is dying is not our involvement in change but a particular arrangement for managing our anxiety about it — a long sequence of signatures that let responsibility stay comfortably diffuse and let the institution feel in control by moving slowly. Agents have taken that arrangement apart, and in doing so they have asked us a question we spent a generation building processes to avoid.

The question is not whether the machine can be trusted to act. It is whether we can be honest about what our approvals were ever for. Answer that honestly and the death of the chain is a liberation: the end of a ritual that consumed enormous effort to adjudicate almost nothing. Answer it dishonestly — rebuild the theatre in software, name a scapegoat, let the real decisions slip upstream into the dark — and we will have automated the least valuable thing we owned while losing the most valuable, and called it progress. Which of those we do will not be decided by the technology. It never is. It will be decided by whether we are willing to look at the structure the falling wall has exposed, and finally tell the truth about it.