

BEYOND THE VELOCITY TRAP

The Acceleration Paradox: Why AI Makes Change Faster and Trust Slower

The production bottleneck has broken. A second clock — the pace at which people come to trust what changes around them — has not. A reflection on the gap that now defines transformation.

Giovanni Leonardi

August 2026



Contents

1. Executive Summary	3
2. The strange quiet after the fast win	3
3. Two clocks that were always different	4
4. Why trust keeps its own time	5
5. The strongest objection: trust always lags, and always catches up	6
6. The mistake practitioners keep making	7
7. Trust debt, and how to pay it down	7
8. The longer view: a different race	9

The Acceleration Paradox: Why AI Makes Change Faster and Trust Slower

Transformation · No. 385
August 2026

“The winners of this era will not be measured by how fast they can ship. They will be measured by how fast they can be believed — and those are not the same race.”

1. Executive Summary

For most of the last thirty years, the binding constraint on transformation was production. Change was slow because making it was slow: writing the code, drafting the target operating model, migrating the data, building the business case, working the change through committee after committee. The craft of the transformation professional was, in large part, the craft of moving heavy things through that slow machine.

That constraint has broken. By this year, an organisation with competent tooling can generate in an afternoon what once took a quarter — a migration plan, a first-cut policy rewrite, a working prototype, an options analysis with the numbers already in it. The production clock has been compressed almost to nothing, and the reflex of the profession has been to celebrate it as pure gain.

The argument of this essay is that the celebration is premature, because acceleration has done something subtler than make us faster. It has exposed a second clock that we never previously had to manage on its own terms — the clock of trust — and that clock has not sped up at all. Trust is built from track record, from locatable accountability, and from comprehension, and each of those accrues at a human pace that no model can compress. The faster we manufacture change, the wider the gap opens between what an organisation can now do and what its people can absorb, verify, and believe.

This is the acceleration paradox: the same capability that removes the production bottleneck actively widens the trust deficit, because it strips out precisely the slow, visible, human deliberation from which trust used to be a by-product. The practitioners who mistake this for a communications problem — something to be closed by explaining harder — will keep losing ground. The ones who treat trust as a structural quantity to be designed for, and paced deliberately, will define what good transformation looks like in this era. The constraint has moved. It is no longer can we build it. It is can we be believed.

2. The strange quiet after the fast win

The pattern announces itself most clearly in the anticlimax. A team ships something that, by every measure of the old world, is a triumph: a decisioning capability rebuilt in six weeks rather than nine months, a reporting layer regenerated over a weekend, a policy library rewritten and internally consistent for the first time in a decade. The velocity is real and the demonstration lands. And then, instead of the adoption curve bending upward, there is a strange quiet.

The quiet has a texture, and anyone who has stood in the room will recognise it. The frontline keeps a parallel spreadsheet "just to check." The senior reviewer waves the output through in the meeting and quietly reworks it afterwards. Usage telemetry shows people opening the tool, glancing at what it produced, and then doing the task the old way anyway. Nothing has failed. The thing works. It is simply not believed, and being unbelievably turns out to be almost indistinguishable, in its operational effect, from being broken.

We are fluent, as a profession, in reading the first situation — the slow build — and we have almost no vocabulary for the second. When a programme was late, we knew what to do: find the bottleneck, add capacity, re-sequence the plan. When a programme is fast and inert, the old instincts misfire. Adding more capability to a system that people already do not trust does not raise adoption; it lowers it further, because each new increment of change resets whatever fragile confidence had begun to form. The velocity we are proud of is the very thing suppressing the return.

3. Two clocks that were always different

It helps to be precise about what has actually changed, because the trust problem is not new. It is simply newly visible.

Production and trust have always run on different clocks. It is just that, in the slow world, the two were close enough that we could treat them as one. A change that took nine months to build also took nine months during which people watched it being built — saw the pilots, met the team, absorbed the reasoning, formed a view about whether the thing was sound. The slowness of production was, without anyone designing it that way, a trust-manufacturing process. Confidence accreted in the gaps. By the time the change landed, much of the belief it needed had already been laid down alongside it.

Compress production to near-zero and you do not compress that second process; you sever it. The change now arrives fully formed, with none of the accompanying period in which understanding could keep pace. What used to be built into the timeline now has to be built after it, against a system that is already moving on to the next increment. This is the mechanism beneath the paradox, and it is worth stating plainly.

Slow production did not merely delay change. It quietly manufactured the trust that change required, by giving comprehension and confidence time to accrue alongside the build. Acceleration keeps the change and discards the by-product.

Once the two clocks are visibly decoupled, the differences between them stop being academic. Production time is a function of tooling and capacity — the things AI has just transformed. Trust time is a function of something else entirely, and that something does not answer to compute.

Dimension	Production time	Trust time
Governed by	Tooling, capacity, model capability	Track record, accountability, comprehension
Direction	Forward — output is generated ahead	Backward — belief is earned from what has already held
Response to more capability	Gets faster	Often gets slower, as each change resets the record
Can be bought	Largely yes	No — it can only be accrued or designed for
Failure mode	Delay	Quiet non-adoption, shadow workarounds

The table looks tidy, but the lived consequence is not. Organisations have spent two years buying down production time and assuming trust time would follow. It has not, and the assumption was never sound.

4. Why trust keeps its own time

If trust simply lagged capability by a fixed interval, this would be a scheduling problem and we could wait it out. The harder truth is that trust is governed by mechanisms that acceleration does not slow but actively disrupts. Three of them matter most.

The first is that trust is retrospective. We come to trust a person, a process, or a system by watching it behave predictably over repeated exposure. Trust is inferred from a track record, and a track record takes calendar time to accumulate — there is no forward-dating it. This is why a newly deployed agentic capability, however impressive its evaluations, starts from zero in the eyes of the people asked to rely on it. Worse, acceleration keeps resetting the very record it needs. Each rapid change alters the thing being judged, so the track record never stabilises long enough to become the basis of confidence. People are, quite rationally, reluctant to trust a system that is different this month from what it was last month.

The second is that trust requires a locatable owner. We extend confidence to processes whose failures can be attributed and answered for — where, when something goes wrong, there is someone who understood the decision and can be held to it. Much of the current wave of automation quietly erodes this. When an autonomous agent composes the analysis, another summarises it, and a human signs a recommendation they did not construct and cannot fully reconstruct, accountability becomes diffuse. The reviewer who "approves" an output they could not have produced and cannot interrogate is not really accountable for it; they are laundering it. People sense this, and they withhold trust from systems where they cannot find the person behind the judgement.

The third is that trust requires comprehension. We trust, in the end, what we understand — or what is vouched for by someone we trust who understands it. Speed attacks comprehension directly. When change arrives faster than anyone can build a working mental model of it, the rational response is not confidence but caution. And the opacity of the underlying systems compounds the problem: a tool that cannot show its reasoning asks to be trusted on faith, at precisely the moment its behaviour is changing too fast for faith to form.

- Trust is retrospective — it is inferred from a record that takes calendar time to build, and acceleration keeps resetting the record.
- Trust needs a locatable owner — and layered automation diffuses accountability until no one can be found behind the judgement.
- Trust needs comprehension — and change that outpaces understanding produces caution, not belief.

None of these is a soft or attitudinal barrier that will yield to better messaging. They are the load-bearing structure of how confidence forms, and each has a natural pace. That pace is the second clock, and it is the one we have failed to manage.

5. The strongest objection: trust always lags, and always catches up

The serious counter-argument deserves its strongest form, not a caricature, because a great many capable people hold it.

It runs like this. Trust has lagged every significant technology in history, and it has always caught up. The first users of the automated ledger, the electronic trading system, the online banking channel all faced exactly this suspicion, and within a few years the suspicion looked quaint. Trust lagged, reliability accumulated, and the lag closed — often faster than the sceptics predicted, because machines, once proven, tend to earn confidence more quickly than people do. On this view the current unease is simply the adjustment phase of a familiar cycle. Wait, let the systems prove themselves, and the trust clock will catch the production clock as it always has. To treat this moment as a new and permanent condition is to mistake a transient for a structural fact.

There is real force in this, and any honest treatment has to concede the pattern is genuine. Trust did lag those technologies, and it did catch up. But the analogy conceals a disanalogy that is decisive, and it sits in the mechanism of catching up rather than the fact of it.

Trust caught up with earlier tools because those tools were legible and their failures were bounded and attributable. The automated ledger did the same arithmetic every time; when it erred, the error was traceable, repeatable, and correctable, and the record it built was stable enough to be trusted precisely because the tool was not changing underneath you. The trust-building mechanism — a stable system accumulating an attributable track record — was intact, and time did the rest. What is different now is not the length of the lag but the health of that mechanism. Systems that change every few weeks never accumulate a stable record; systems whose failures are correlated, opaque, and hard to attribute deny people the very thing from which trust is inferred. The earlier technologies waited for trust; this one, deployed carelessly, corrodes the machinery by which trust is

made.

So the objection is right that trust can catch up — and wrong that it will do so automatically here. It caught up before because the conditions for catching up were present. Remove those conditions — stability, attribution, legibility — and the lag does not merely persist; it can widen indefinitely. The catch-up is not a law of nature. It is a thing that has to be engineered, and this era, left to its own momentum, engineers the opposite.

6. The mistake practitioners keep making

Faced with a visible trust deficit, the reflex of the profession is to reach for the communications playbook. Run the town halls. Sharpen the narrative. Appoint the champions, publish the FAQs, commission the adoption campaign. Treat the deficit, in other words, as a problem of perception — a gap between how good the thing is and how good people believe it to be, to be closed by explaining the thing harder.

This is a category error, and an expensive one. The trust deficit is not, in the main, a perception gap. It is an accurate perception of a structural gap. People do not trust the accelerated system because the conditions under which trust forms are genuinely absent: there is no stable record, no locatable owner, no time to comprehend. Explaining harder does nothing to supply any of those. It often makes matters worse, because a confident communications push against a well-founded caution reads as pressure, and pressure applied to a trust decision produces compliance, not belief — the parallel spreadsheet simply moves off the shared drive and onto someone's laptop.

“A confident campaign aimed at a well-founded caution does not create trust. It creates the appearance of adoption laid over the reality of quiet refusal.”

The tell is when adoption metrics and confidence diverge — logins up, reliance flat. That divergence is not a messaging failure to be closed with more messaging. It is the organisation telling you, precisely and truthfully, that the structural conditions for trust are not yet in place. The work is not to argue people out of their caution. It is to build the conditions that make the caution unnecessary.

7. Trust debt, and how to pay it down

The most useful way I have found to hold this is by analogy to a concept the technology side of the house already understands in its bones: technical debt. When you ship fast by cutting corners on the underlying structure, you take on a debt that accrues interest and eventually has to be paid, usually at a worse exchange rate than if you had done the work up front. Acceleration in transformation creates an exactly parallel liability. Every

increment of change shipped faster than trust can form leaves a balance outstanding — a trust debt — and that debt, too, accrues interest in the form of workarounds, shadow processes, second-guessing, and the slow corrosion of the organisation's willingness to rely on anything you deliver.

Consider a composite that is true to what recurs in this era. A lending operation rebuilds its credit-decisioning workflow around an agentic capability. On the production clock the result is spectacular: median time-to-decision falls from around six days to roughly four hours, and the model's decisions, measured against subsequent outcomes, are at least as sound as the human baseline. By every metric the old world used, it is a triumph — and it is booked as one.

Then the trust debt begins to service itself. Underwriters, unable to reconstruct why a given decision was reached, start overriding cases they would once have accepted; the override rate climbs from a historic three per cent toward the high teens. Each override drags a fast decision back onto the slow clock, and the effective cycle time — the one the customer actually experiences — quietly rises again. A second-line review queue forms to adjudicate the overrides, staffed by the very people the automation was meant to free. Within two quarters the operation is spending more human effort per decision than before it accelerated, and the celebrated four-hour figure survives only in the slide that first reported it. Nothing broke. The debt simply came due.

What makes the illustration worth dwelling on is that the debt was avoidable, and cheaply, by paying trust its own time up front rather than borrowing against it. Paying it down is not mysterious; it is a matter of designing for the second clock as deliberately as we now design for the first.

1. Restore a locatable owner. For any consequential decision, ensure there is a named human who genuinely understands the basis of it and can answer for it — not a signatory laundering an output, but an owner who could reconstruct and defend the judgement. Accountability that can be found is the precondition of trust, not a compliance afterthought.
2. Show the reasoning, not just the result. A system that can expose why it reached a conclusion lets comprehension keep pace with output. Legibility is not a nicety here; it is the mechanism by which the trust clock is allowed to run at all.
3. Pace rollout to comprehension, not to capability. The fact that you can ship the next increment this week does not mean the organisation can absorb it this week. Deliberate, visible stability — holding a capability constant long enough for a track record to form — is now a design decision, not a delay.
4. Preserve reversibility. People extend trust far more readily to changes they know can be undone. A credible path back lowers the stakes of belief and, paradoxically, accelerates it.
5. Measure reliance, not adoption. Track whether people actually depend on the output when it matters — whether the parallel spreadsheet has been retired — rather than whether they have logged in. Reliance is the only honest measure of trust.

There is a deliberate friction in several of these, and that will offend the instinct of the moment, which is to remove all friction as waste. But not all friction is waste. Some of it is the trust-manufacturing process that slow production used to supply for free, now reintroduced on purpose because the thing it made is still needed. The art is to add friction precisely where trust is formed and nowhere else — a very different discipline

from the reflexive elimination of every delay.

8. The longer view: a different race

Step back far enough and the shape of the era becomes clear. For a generation, competitive advantage in transformation went to those who could overcome the production constraint — who could build faster, migrate faster, ship faster than their rivals. That race is ending, not because speed stopped mattering but because speed stopped being scarce. When everyone can generate the change in an afternoon, the ability to generate it confers no advantage. It becomes table stakes, and the contest moves elsewhere.

Where it moves is to the second clock. The organisations that will pull ahead in the years just in front of us are not the ones that deploy the fastest; almost everyone will deploy fast. They are the ones that become trustworthy the fastest — that build the stable records, the locatable accountability, and the legible reasoning from which reliance can form, and that have the discipline to pace change to what people can actually absorb. This is a genuinely different capability from the one the profession spent thirty years building, and most organisations have not yet noticed that the thing they are good at is no longer the thing that is scarce.

There is a version of the future in which we get this wrong at scale — where organisations keep winning the production race, keep accruing trust debt, and slowly discover that they have built enormous capacity to do things their own people will not rely on and their own customers will not believe. And there is a version in which the discipline reinvents itself around the constraint that actually binds, treating trust as the engineered quantity it has become. Which of those we get is not determined by the technology. It is determined by whether we can unlearn the reflex to treat speed as the goal.

The paradox, in the end, is also the instruction. AI has made change faster than trust can keep up — and so the work of transformation, for the first time, is less about accelerating the change than about doing the patient, structural, unglamorous work of earning the right to be believed at the speed we now move. The winners of this era will not be measured by how fast they can ship. They will be measured by how fast they can be believed — and those are not the same race.