

Stop Sprinkling AI

Why the winners rebuild the work, not the tools

ThoughtFox · AI Transformation · Edel Churchill · July 2026

Before this is about frameworks or research, it is about a single invoice.

A manufacturer we worked with had, some time earlier, deployed an automation to handle supplier invoices. Technically, it worked. It did exactly what it was built to do. And the person it was meant to help carried on processing invoices by hand, opening the system only now and then. The technology worked. The adoption did not.

That gap, between a thing that functions and a thing that gets used, is the whole problem with how most organisations are approaching AI. And it scales all the way up.

Ask executives privately and almost all of them will tell you AI is working. In WRITER's 2026 enterprise survey, 97% said they personally benefit from it. Then ask the harder question. Is it in the numbers? Only 29% see significant enterprise ROI. Over 80% of organisations now use generative AI. Usage is close to universal. Enterprise value is not.

It is the same invoice problem, a thousand times over. The instinct is to treat the gap as a technology problem, to assume a better model or a bigger licence deal will close it. It won't. The organisations stuck in the gap and the ones climbing out of it are often running the same tools. What separates them is what they did to the work around those tools.

This piece is about that difference. Why layering AI onto existing processes produces personal wins but not enterprise returns, why the change is harder than any technology rollout that came before it, and the five-step sequence the small minority who get it right actually follow.

You can't sprinkle AI on a broken process

Here is the pattern we see repeatedly. A team maps an existing workflow, say eight steps. They look for where AI fits. Steps three and four are the painful, obviously manual ones, so a model goes in there to speed them up. Two steps get faster. Six stay exactly as broken as they were. The team reports a productivity win, the pilot looks successful, and six months later nobody can find the impact in the P&L.

McKinsey's 2025 State of AI research puts a number on why this happens. Of all the organisational changes linked to gen-AI success, fundamental workflow redesign is the

change most strongly correlated with EBIT impact, ahead of model choice, tech stack and budget size. And yet only around 21% of firms have redesigned even a single workflow. The lever that matters most is the one almost nobody is pulling.

Contrast that with an optical-software company we worked with. Facing a mature, complex, regulated codebase, the easy move would have been to hand every engineer an AI coding assistant and count the licences. Instead they rebuilt the entire software-development lifecycle around AI. As their own team put it, the step-change came from three structural changes, not a tool rollout.

The insight that drove it is the one most organisations miss. When AI writes the code, the bottleneck moves. Planning, specification and review, the human judgment work, became the critical investment, because that is where the constraint now lived. They did not speed up the old process. They rebuilt it around where the value had shifted to. The result was 95% agent-written code across a 140-engineer team, a twentyfold rise in AI-accepted code in six months, and roughly €6.1M of annualised engineering value in the first year. You do not get numbers like that by decorating the old workflow. You get them by blowing it up.

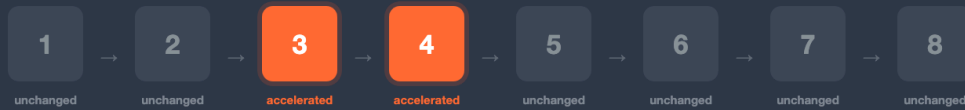
The reason bolt-on AI underdelivers is structural. When you insert a tool into the middle of a process you did not question, you inherit every assumption baked into that process: the handoffs, the approval gates, the "we've always done it this way" steps that existed to compensate for limitations the AI has just removed. You have made a flawed process run faster. Speed on a bad path does not get you to a better destination.

The winners don't accelerate the old steps. They step back and ask the first-principles question. What are we actually trying to accomplish here? Not "how do we do this task faster," but "why does this task exist at all, and what would it look like if we designed it today?" Rebuilt from the goal backwards, an eight-step process often collapses to one or two. At the same food manufacturer, a supplier-onboarding process that had taken days of manual intake was rebuilt to run in minutes, not by speeding up the steps but by removing most of them.

Practical AI, not PowerPoint AI, starts by blowing up the SOP, not decorating it.

THOUGHTFOX · AI TRANSFORMATION

**Sprinkling AI: two steps faster.
Six steps still broken.**



"Speed on a bad path does not get you to a better destination."

Automate, augment, amplify: three different ambitions

It helps to be precise about what "using AI" in a workflow actually means, because the word covers three very different ambitions.

You can **automate**, handing a discrete, well-defined task to a system and taking it off a person's plate. You can **augment**, putting AI alongside a person so they do the same job faster. Or you can **amplify**, redesigning the work so a team achieves something it simply couldn't before, because the constraint that used to shape the whole process is gone.

Most bolt-on AI stalls at augment. Augment is real and worth having, but it is also where personal productivity wins live, and personal wins are exactly the ones that do not automatically scale to the enterprise. Amplification is where the enterprise return sits, and amplification is impossible without redesign. The optical-software company did not augment its engineers. It amplified the whole function, and the €6.1M is what amplification looks like on a P&L. You cannot amplify a workflow you have not rethought. That is the through-line from "AI is working for me" to "AI is working for us," and it runs straight through the work itself.

Why this is genuinely hard, and why your change playbook won't save you

If workflow redesign is the strongest lever, why do so few organisations pull it? Because it collides with something every previous technology rollout managed to avoid: human identity.

You have deployed big systems before. ERP. CRM. The move to mobile. And you survived them all with roughly the same change-management playbook, because those technologies were fundamentally additive. The people stayed the same. The process stayed broadly the same. The technology changed, sometimes significantly, but it sat alongside the existing work rather than replacing the logic of it. You trained people on a new interface, and the core of the job, along with the identity attached to it, survived intact.

AI is different. It rewrites the playbook rather than adding to it. It does not sit beside the job; it absorbs parts of it. And the parts it absorbs are often exactly the parts people built their expertise, their status and their sense of usefulness around. At the optical-software company, engineers moved from writing code to directing and judging AI output. That is a change in what an engineer is for, not a change of tool. The reason the transition worked is that the organisation spent twelve months building AI fluency across the team before the step-change. It treated the human shift as the main event rather than the side effect.

Most organisations don't. And this is where the ROI gap actually opens, along a very human fault line. On their own, people see obvious wins. They pair their expertise with a tool and feel more capable, hence the 97%. But the moment work becomes shared, a different instinct takes over. The natural human response, when the thing you were valued for is being automated, is to justify your position, defend your title, and demonstrate the judgment that makes you indispensable. That instinct is not irrational, and it is not a character flaw. It is what people do when their standing is quietly under threat. It is also, in miniature, the person who kept processing the invoices by hand.

The evidence that this is a people problem rather than a technology problem is hard to ignore. In the same 2026 survey, 54% of C-suite leaders said adopting AI is "tearing their company apart." That is not a statement about model quality. It is a statement about resistance, agency and fear, and it is coming from the people who approved the AI strategy in the first place.

So treat resistance as a signal rather than an obstacle. It tells you precisely where adoption will break, and why. And it is not a signal you can answer with the usual tools. No amount of additional training budget fixes a problem that is really about identity. The standard change-management toolkit, the comms plan, the town hall, the training day, the adoption dashboard, was built for additive technology. It assumes the job stays fundamentally the same and people simply need to learn a new interface. When the job itself is being rewritten, that assumption is the flaw. You cannot communicate your way past a genuine question of "what am I here for now?"

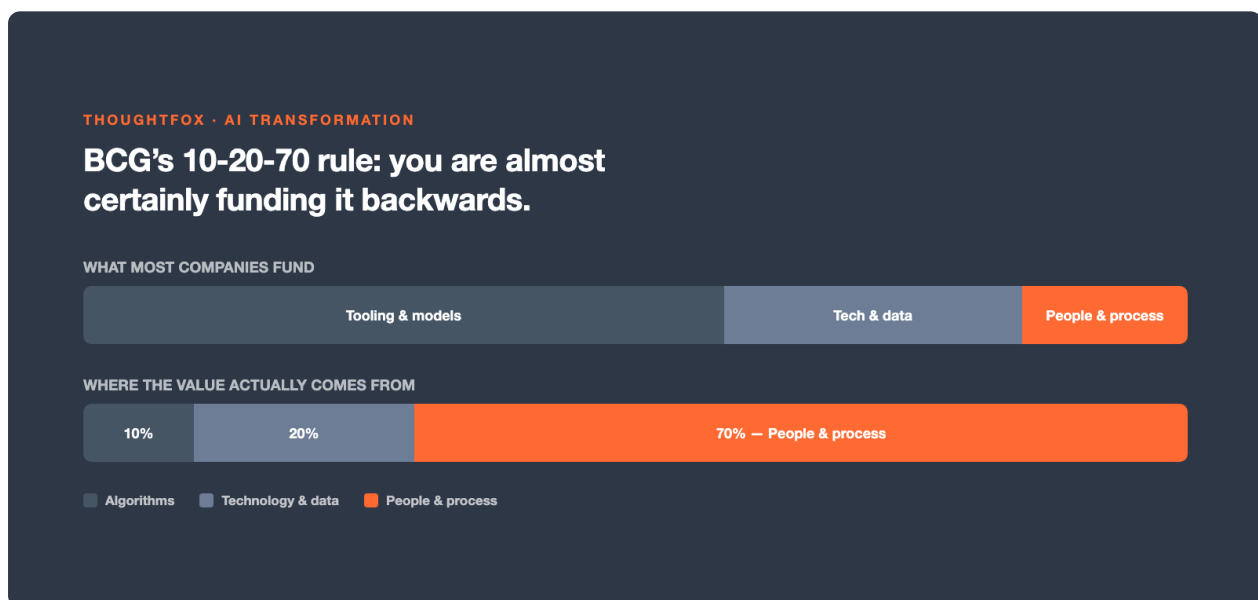
Human-centric AI transformation is not a soft add-on to the technical work. It is the load-bearing part. Get the workflow redesign right and fumble the human transition, and you will have an elegant new process that nobody runs the way it was designed.

You are almost certainly funding it backwards

BCG has been making a related point for years with its 10-20-70 rule. Of the value organisations capture from AI, roughly 10% comes from the algorithms, 20% from the technology and data that support them, and 70% from the people and processes around them. The model is the smallest slice of the value, by a wide margin.

Most companies invert it in practice. They spend the majority of their time, attention and budget on the 10% that matters least (which model, which platform, how many licences) and treat the 70% as an afterthought to be handled by whatever change-management effort gets the time and money left over. The result is predictable: heavy investment in the smallest lever, token investment in the largest, and puzzlement when the returns don't come.

The 10-20-70 split is not a slogan. It is a budgeting instruction. If your AI programme's spend, and more importantly your leadership attention, is not weighted toward people and process, you are funding the transformation upside down.



The CFO test. Look at where the AI budget actually goes this quarter, and where the leadership hours go. If more than a tenth of either is going on tooling and models, you are funding the smallest lever the hardest, and starving the one that drives the return.

How the successful few operate: the five Rs of AI-native change

Roughly 5% of organisations are getting AI transformation right at enterprise scale. These are the ones attributing real, measurable EBIT impact to AI rather than a collection of personal productivity anecdotes. They aren't luckier, better funded, or working with better

models. They run a different sequence, one built for a technology that changes the job rather than adding to it. We call it the Five Rs of AI-native change.

1. Resource. Put the money, and the leadership attention, where the value actually is: people and process. AI literacy, change management, and workflow redesign, not another round of tooling. This is BCG's 70% made operational. The optical-software company spent a full year building AI fluency before it touched the SDLC, and that year was the reason the redesign held. Most programmes are under-resourced on exactly the dimension that determines whether they succeed.

2. Redesign. Don't automate the existing workflow. Decompose it and rebuild it around human-AI collaboration from first principles. Start from the goal, not the current steps. Find where the bottleneck moves once AI is doing the routine work, and put your human judgment there. Build deliberate human checkpoints at the decisions that matter. This is where the eight-step process becomes two, and where the difference between marginal speed and genuine amplification is decided.

3. Relearn. Unlearn the old job and relearn the new one. This is the step most transformations skip, and the reason they stall. People cannot run the new playbook while still being measured and rewarded for the old one. Relearning has to be led explicitly: naming what the role used to be, acknowledging honestly what people are being asked to let go of, and giving them a credible, valued version of what the role becomes. "Engineer" becoming "director of AI output" is a promotion if it is led as one, and a threat if it is not. The goal is not to lose people. It is to help them grow into a different and often more valuable job.

4. Ritualise. Adoption that depends on enthusiasm decays within a quarter. Adoption built into rituals sticks. The optical-software company ran Communities of Practice and a standing AI-enablement function. Make working AI-native a repeated, visible team habit rather than a one-off training event. Rituals make the new way of working normal, social and continuous, and they surface resistance early, while it is still cheap to address.

5. Reinforce. What gets graded gets done. Recognise the new behaviours, measure them, and make them part of how performance is actually judged. When code review becomes the primary quality gate, review is what you value and reward, not lines written. Without reinforcement, the organisation quietly drifts back to the old way the moment leadership attention moves on, because the old way is still what the incentives point to.

Notice the balance. Only one of the five steps is primarily about technology. Four are about people and process. That is not a stylistic choice. It is the 70% expressed as a sequence of actions.

The five Rs of AI-native change



The compounding payoff

The reason this matters is not any single project's ROI. It is that the Five Rs compound.

A workflow rebuilt properly makes the next rebuild easier, because the organisation has learned how to decompose and redesign rather than just accelerate. A team that has genuinely relearned one role takes on the next change faster, because it has done the hard identity work once and knows it survives the process. Behaviours that are ritualised and reinforced stop being an initiative and become the culture, and culture is what turns a handful of successful pilots into organisation-wide capability.

We saw this vividly at Kepak, the Irish food manufacturer. We worked directly with their internal IT and business analyst teams. After a four-session activation programme, team members went home and, without being asked, started rebuilding their own problems. A legacy migration that had sat on the backlog for months was done in a single unassigned day. A physical, multi-desk approval process was automated and running before the programme even finished. Nobody scheduled that work. The capability had simply become theirs.

That is the mechanism behind the high-performer gap. It is not that the top 5% found a better model. It is that each cycle of redesign and relearn makes the next one cheaper and faster, while the organisations still sprinkling AI on top pay full price for every disappointing pilot, over and over. The gap compounds because the capability compounds, and that is why it is widening rather than closing.

The mid-market angle: your size is an advantage, if you use it

Most of the headline research here is drawn from large enterprises. But the organisations with the most to gain, and a genuinely different set of conditions, sit in the mid-market, roughly 200 to 5,000 employees. If that is you, the Five Rs still apply, but the balance of advantage and risk is not the same as it is for a global enterprise.

Start with the advantage, because it is real and under-used. Mid-market firms have structural agility that large enterprises would pay dearly for. Fewer layers. Shorter approval chains. Less accumulated process debt. Leadership close enough to the work to understand what redesign would mean. Kepak is the case in point. A traditional, multi-site manufacturer took twenty people, most of them not AI specialists, from a maturity baseline to two working, owned tools in a single build day, with a governance framework applied by the third session. A global enterprise would still have been aligning stakeholders.

And here is the part that should change how mid-market leaders think about this. As Kepak's own team concluded, the gap is not about technology access. Anyone can access the same models. The gap is about organisational capability: the frameworks, governance, skills and culture that decide whether those models deliver value or stay interesting experiments. The tooling has commoditised. The 10% and the 20% are available to a 300-person firm on the same terms as a 30,000-person one. Which means the differentiator is almost entirely the 70%, precisely the ground on which a nimble mid-market firm can out-manoeuvre a slower, larger competitor.

Now the traps, because the same size that helps you also concentrates the risk.

The first is thin resourcing. Mid-market firms rarely have a dedicated transformation office or a change-management bench. This makes resourcing the 70% harder, and far more important. Under-resourcing the people side is the most common way mid-market AI programmes quietly fail, because there is no spare capacity to catch the transition when it wobbles.

The second is key-person concentration. In a mid-market business, critical expertise often lives in a handful of heads, and those people frequently are the process. That raises the stakes on Relearn sharply. Alienate one senior specialist whose role is being redesigned and the whole initiative can stall. Handled well, though, it cuts the other way. Bringing five or six key people genuinely through the transition is far more tractable than doing it across thousands.

The third is that you can't afford to waste cycles. A large enterprise can quietly absorb a graveyard of failed pilots. A mid-market firm feels every one. The proportional cost of

sprinkling AI on top is higher, which raises, rather than lowers, the return on getting the redesign right the first time.

The net is this. Your size is an edge in an AI-native transition, but only if you use it deliberately. You can move faster than your larger competitors on redesign, and lead the human change more personally than they ever could. What you cannot do is paper over a fumbled people transition with scale. For the mid-market, the 70% is not just where the value is. It is where the whole bet is.

Frequently asked questions

Why do most companies see personal AI wins but no enterprise ROI?

Because most AI use stops at "augment": a person pairs their own expertise with a tool and feels more capable. That is real, but it is a personal win, not an enterprise one. Enterprise return only shows up when the surrounding workflow is redesigned, not just accelerated. WRITER's 2026 survey found 97% of executives report a personal benefit from AI while only 29% see significant enterprise ROI, which is the augment-versus-redesign gap showing up in the numbers.

What is workflow redesign, and why does it matter more than the AI model you choose?

Workflow redesign means decomposing a process from first principles, asking what it is actually trying to achieve, and rebuilding it around human-AI collaboration rather than inserting AI into the existing steps. McKinsey's 2025 State of AI research found it is the organisational change most strongly correlated with EBIT impact, ahead of model choice, tech stack, or budget size, yet only around 21% of firms have redesigned even a single workflow.

What is BCG's 10-20-70 rule?

It is BCG's finding that roughly 10% of the value organisations capture from AI comes from the algorithms, 20% from the supporting technology and data, and 70% from the people and processes around them. Most companies spend their budget and leadership attention in the opposite proportion, over-investing in the smallest lever and under-investing in the one that actually drives returns.

What are the "Five Rs" of AI-native change?

Resource, Redesign, Relearn, Ritualise, and Reinforce. Resource means funding people and process ahead of tooling. Redesign means rebuilding the workflow from the goal backwards. Relearn means helping people let go of the old version of their role and grow into the new

one. Ritualise means embedding the new way of working into repeated team habits rather than a one-off training event. Reinforce means measuring and rewarding the new behaviours so the organisation doesn't drift back to the old ones.

Why doesn't the standard change-management playbook work for AI adoption?

Because it was built for additive technology, ERP, CRM, mobile, where the job stayed the same and people just learned a new interface. AI absorbs parts of the job itself, often the parts people built their expertise and identity around. A comms plan, a town hall, or a training day cannot answer the genuine question an employee is asking, which is "what am I here for now?" That is why 54% of C-suite leaders in the same 2026 survey said AI adoption is "tearing their company apart": it is a signal about identity and resistance, not about model quality.

Is a smaller, mid-market company at a disadvantage in AI transformation?

No, if it uses its size deliberately. Mid-market firms (roughly 200 to 5,000 employees) have fewer layers, shorter approval chains, and leadership close enough to the work to lead redesign and relearning personally. The risk is thin resourcing and key-person concentration, since critical expertise often sits with a handful of people. Handled well, bringing five or six key people through the transition is far more tractable than doing it across thousands, which is exactly the advantage a large enterprise doesn't have.

Where to start on Monday

You don't close this gap with a bigger AI budget. You close it by changing what you do with the budget you have.

Pick one workflow, ideally one that is painful, high-value, and visible enough that success will be noticed. Resist the urge to find where AI "fits" in it. Instead, decompose it from first principles and ask what it would look like if you designed it today. Then treat the human transition, the relearning, as the main event rather than the clean-up. Resource it properly, build it into rituals, and reinforce the new behaviours in how you measure and reward people.

The companies stuck in pilot purgatory are still sprinkling AI on top and wondering why the numbers don't move. The ones closing the ROI gap made a different decision. They rebuilt the work, and brought their people through the change rather than around it. One of them started with a single invoice that nobody was using, and ended with a team that builds its own tools.

If you are serious about enterprise return, stop asking which tool to buy. Start asking which workflow to blow up first, and whether your people are set up to rebuild it.

Sources: McKinsey, The State of AI 2025 (QuantumBlack) · BCG, 10-20-70 principle · WRITER, 2026 AI Adoption in the Enterprise survey (vendor survey, credible but not independent research) · ThoughtFox and Kepak, Tech Team AI Activation case study (2026) · ThoughtFox client engagement, optical-software sector (2026), client anonymised.

© ThoughtFox · thoughtfox.ai