

Article

Oct 2026

You can't automate a mess

Automation doesn't solve problems, it amplifies them. And most organisations are not nearly as ready as they think they are.



Andrew Buckels

Head of Advisory Services and Housing
Sector Principal

Email andrew.buckels@waterstons.com

In the [previous piece in this series](#), we looked at what goes wrong when businesses stumble into "doing more with less" rather than choosing it deliberately, and automating a broken process was one of the traps we named. This piece is about why that trap is so easy to fall into, and how to tell whether you're actually ready to automate at all.

It is usually the team holding the pen when the automation request lands, expected to make it work regardless of the state of what's underneath. If the problem is too much work and too little capacity, surely the answer is to remove people from the equation, or at least reduce their involvement. It's a compelling argument. It's also, in a significant number of cases, how organisations end up worse off than when they started. Not because automation is wrong, but because they automated the wrong thing, in the wrong state, at the wrong moment.

What automation actually does

There's a persistent belief that automation is a corrective force, that the discipline of configuring a system will force an organisation to confront its inconsistencies and resolve them. This occasionally happens. More often, it doesn't.

What automation actually does is take whatever process you give it and run it faster, more consistently, and at greater scale than any team could manage. If that process is coherent and consistently followed, the results are genuinely transformative. If it's inconsistent or quietly different depending on who's doing it, automation makes all of that worse, faster, and harder to unpick. The technology isn't the problem. It's doing exactly what it was asked to do. The problem is that what it was asked to do was never properly defined.

AI has raised the stakes

AI has shifted the guardrails we've grown used to, the ones that kept automation fairly predictable and contained. Good, properly governed automation has effectively grown superpowers: it can now reason through exceptions, adapt to variation, and handle judgement calls that used to need a person. Automation built on shaky foundations fares differently. AI will run that mess with more confidence and more speed than ever, making issues harder to spot and harder to unpick. The fundamentals haven't changed. The cost of ignoring them has gone up.

Trap one: automating an inconsistent process

Ask most organisations how they handle a given process, onboarding a supplier, processing an expense claim, and they'll describe it with confidence. There's a process. It's documented. People follow it. Then look at how it's actually done.

In most cases you'll find not one process but several, running in parallel, each the product of a different team's interpretation, or a workaround introduced to solve a specific problem and never removed. Nobody involved is wrong, exactly, but they're not doing the same thing either. Automate this, and you don't get one efficient process, you get several automated inconsistencies running simultaneously.

This isn't a failure of the people involved. It's a predictable consequence of teams growing faster than their processes, solving local problems without visibility of what others are doing. The automation project doesn't fail because the technology was inadequate. It fails because the organisation tried to automate something that was never actually standardised, and the technology simply made that impossible to ignore any longer.

Trap two: automating on top of tool sprawl

The second trap is related but distinct: inconsistent infrastructure rather than inconsistent process. Tool sprawl is the quiet accumulation of overlapping platforms that happens when technology decisions get made locally and reactively, a team adopts its own project tool, a department buys its own data platform, none of it unreasonable in isolation. Collectively, it creates an environment where nobody's sure which system is authoritative, and the honest answer to "where do I find this?" is usually "it depends."

When an organisation tries to automate in that environment, it hits a problem no automation platform can solve: it can't connect systems that were never designed to talk to each other, or establish a single source of truth where none has ever existed. These aren't technical problems. They're organisational ones, and they need organisational solutions first.

Signs you may not be automation-ready

If you're the one being asked to automate it, these are the signs worth checking for before you commit to a timeline.

- The same process is described differently depending on who you ask
- Teams have workarounds that are now load-bearing parts of the operation
- More than one system could reasonably be called the "source of truth" for the same data
- Documentation describes how things were designed to work, not how they actually work
- Nobody can confidently map a core process end to end without caveats

Why organisations skip the hard part anyway

Process standardisation and tool rationalisation are slow, unglamorous, and politically difficult. They mean asking teams to give up how they currently do things, and telling someone their tool is being retired. Automation, by contrast, feels like progress. It has a business case, a timeline, a launch date. The process work that should come first rarely gets the same billing, even though it's what actually determines whether the automation delivers what was promised.

This isn't an argument against automation

Done well, automation is one of the most powerful levers you have for more capacity without burning through your people. The IT leaders who get it right aren't the ones who moved fastest, they're the ones most deliberate about what they were automating and why. Sorting the process first isn't a delay. It's the work. Skip it, and you tend to do the implementation twice, once to automate, and once to fix what the automation broke or exposed.

Real readiness means answering three questions with confidence before a single workflow gets built:

1. Is this process consistent? Not documented, consistent. Does everyone who touches it do it the same way, and if not, have you decided which way is right?
2. Is the data clean and in one place? Not approximately, definitively. Is there a single source of truth, and does everyone know what it is?
3. Do you know what you're trying to achieve? Not "efficiency" in the abstract, specifically. What does success look like, and what will you do if the automation produces something you didn't expect?

If the answer to any of these is uncertain, the automation project hasn't started yet. The process work has.

Automation isn't a shortcut to organisational clarity. It's a multiplier of whatever you already have. If what you have is a mess, the most useful thing you can do isn't find a better tool. It's stop, look honestly at the process, and fix it before asking any technology to run it at scale. AI hasn't changed any of that logic. It's just made the gap between doing it well and doing it badly far wider than it used to be.
