

The truth about AI and Transformation

Jon Ward

Executive summary

I recently wrote that “there is no such thing as an AI transformation”. I still believe it. What exists is business transformation that uses AI as one mechanism among others based on a clear definition of organisational purpose and strategy. This paper is the other half of the argument.

AI earns its place in a transformation when it is tied to a specific strategic goal, defined as a solution with a business case, assessed for the harm it could do, bounded by operational controls, and fitted with a way to stop it should it start to misbehave. Omit any one of those elements, and you have an experiment, not a transformational capability.

From AI transformation to AI for business transformation

Organisations do not transform because they bought an AI platform or model. They transform when they change how they create and deliver customer value. I saw the same mistake with Agile. Many Agile transformations stalled because they promised business benefits yet treated the transformational element as merely changing how software teams work. Localised, parochial adaptation, whether involving AI or Agile, does not deliver Enterprise Agility; Enterprise Agility is the rapid adaptation of an enterprise's value creation model. It is a broad business goal involving business operations facilitated by technology, not a purely IT goal.

The evidence is starting to show. S&P Global Market Intelligence found that 42 per cent of companies abandoned most of their AI initiatives in 2025, up from 17 per cent the year before. The cause is rarely the algorithm. It is architectural, cultural, and, importantly, governance-related. The pilot was never connected to a business outcome. Many had AI solutions seeking business problems rather than solving them.

AI for business transformation puts business operations and strategy at the centre and asks where AI can change what is possible or solve a specific business problem. AI helps decide *what* to change through continuous market sensing and metrics derived from a live Value Creation Model, a Digital Twin. It also becomes part of *how* change is delivered, by automating a process, informing a decision or acting inside a value stream. Most deployments are incremental, and that is fine. AI becomes transformational only when it changes the value creation model itself. Leaders need to know which of the two: experiment or strategic capability they are funding. And I am not saying everyone must change at speed or perish. Your market and how fast your customers shift set the pace. The discipline below works at either pace.

Start with the strategy map, not the use-case list

Most organisations plan for AI backwards. They list what AI can do in their sector and hope relevance follows. It rarely does. An AI opportunity is contextual: only ever right or wrong relative

to a specific strategy. A low-cost insurer might welcome fully automated claims handling. A premium insurer that wins on the quality of its claims experience might find the same automation quietly undermines why customers chose it.

So, I start with Strategy and create a Strategy Map. Strategy Mapping results from applying nine well-known techniques to create a picture so leaders can decide where change is needed, where AI can be applied, and how to prioritise investments.

The Strategy Choice Cascade sets "*Where to Play and How to Win*" as the bet the further techniques test. The Business Model Canvas baselines the value proposition and resources those techniques then build on. The Wardley Map takes that value chain and decides what to build versus buy, which shapes how the value streams are resourced later. The Value Proposition Canvas grounds a candidate segment in real jobs, pains and gains, and its named pains and gains become the most credible source for the Reason-to-Buy Map's claims when a segment-specific canvas already exists. Together, these techniques show how the organisation wins. AI opportunities then fall out of it, each traced to a strategic claim it strengthens.

What leaders hold, and what they adapt

Most AI initiatives do not fail in the data centre. They fail on a Tuesday afternoon in a steering committee where nobody is sure who owns the decision the model is now making.

Leaders hold strategic intent, values and purpose constant. They adapt decision processes, team structures and governance to fulfil that intent. An AI may not have that compass or the ability to adjust outputs as circumstances change. Consequently, an organisation must name an owner for every AI-informed decision and choose deliberately whether a system augments judgement or automates it.

AI also exposes structural weaknesses. At one client's IT function, we calculated the ratio of testers and software engineers to managers, project managers, delivery managers and architects, which we called enablers, and senior leaders agreed the number of enablers compared to engineers was far higher than expected. That ratio, supported by a productivity analysis, triggered an organisational redesign. AI created this moment.

If an agent now performs work that previously passed through three approval layers, the question is: why were those layers in place? Does the operational risk being managed justify that level of scrutiny, or were three layers of approval overkill? How does the use of AI further the organisational strategy or improve operational performance?

Defining the AI solution

Strategic analysis tells you where in the business operation an AI solution is needed. OKRs then set the improvement goals it must deliver and how to measure the impact. Only then, when the business outcomes are specified, is it time to adopt AI governance mechanisms.

I use an AI Authority and Governance Toolkit, which runs from design through deployment to ongoing oversight. It defines the solution through four artefacts, built in order.

Artefact	The question it answers	What it hands to the next step
Business case	Should we do this, and is AI the right answer?	A named owner, a value to measure and a decision inventory
Risk assessment	What is the worst plausible outcome, and who bears it?	A risk tier and decision boundary conditions
Controls	Who, or what, may decide what, and at which thresholds?	An authority model, Threshold Matrix, an autonomy state and when human intervention is required
Stop-use mechanisms	How do we halt it and return to a safe state?	Tested “kill” criteria and a reversion runbook

Figure 1. Four linked artefacts that turn an AI opportunity into a governed solution.

Business case

The first and obvious question is whether this solution should be AI at all. With rising concerns about token cost, a human, a rules engine, or a process fix will sometimes do the job at a fraction of the cost and risk. A good business case names the strategic claim it strengthens, sets measurable success criteria and a path from pilot to production, is candid about data readiness, and prices the cost of governing the system, not just building it. It also starts the decision inventory: the decisions the system will influence or make. If the team cannot write that list, they do not yet understand what they are automating.

Risk assessment

Asking what the worst plausible outcome is if the AI misbehaves, and which leader is responsible, sets a risk category. The Risk Tier is set by the level of organisational harm it could cause; it is not the department or the model's sophistication. For example, a CV screening tool has a high impact even if it is only a keyword filter. The EU AI Act gets this right by tiering risk on impact, not algorithm. Each material risk then gets a boundary condition, which becomes both a quality-assurance test and a monitoring criterion in live operation, with a named owner.

Controls

A Threshold Matrix sets out in advance what the system may decide alone, what needs human approval, and what must be escalated. Behind it sits **organisational drag**: the delay caused by each review and approval a decision passes through. Drag is not the enemy. The failure is the wrong amount in the wrong place. Autonomy should be a setting, not a switch. A system moves from shadow mode to assisted to autonomous only on evidence, and a named person approves each step.

Stop-use mechanisms

Agentic AI reduces opportunities for human intervention; it reduces executive intuition. Therefore, you must design stop mechanisms before go-live. A kill switch checklist defines what halts or downgrades the system, such as boundary breaches, confidence drops or volume spikes, and who owns the decision to kill. A Reversion Model defines how to return to a known-good state, including what happens to half-finished actions. Stop-use also means retirement. A framework that can approve an AI system but cannot retire one is only half built.

A kill switch nobody has pressed is a hypothesis, not a control.

Keeping it honest after go-live

Models drift, business data changes, and scope creeps one reasonable request at a time. So governance is a lifecycle discipline. Measure how long approvals and interventions actually take. Review for drift before an incident reveals it. Log every time authority is exercised, because that log is the evidence the governance is live. With Provision 29 of the UK Corporate Governance Code asking boards to declare that material controls work, that evidence will matter. What you learn then feeds back into the strategy map. That loop is what makes this Enterprise Agility rather than a project plan.

Where to start

You probably have at least one AI initiative that cannot be linked to a strategic outcome, or cannot yet be switched off. Start there. I am working with a small number of organisations in financial services, insurance, healthcare and telecoms in three ways. I can take one live or proposed use case with your team from strategy map to stop mechanisms, producing a board-ready definition of the solution. I am looking for a few organisations to pilot the Strategy Mapping Toolkit and the AI Authority and Governance Toolkit, and help shape them. And my five-day AI Leadership and Strategy for Business Transformation programme with Informa builds these disciplines in your leaders.

Let's talk about one AI initiative you are not yet sure you could stop.

jonathan.ward@beneficialconsulting.co.uk | beneficialconsulting.co.uk

About the author

Jon Ward is Director of Beneficial Consulting, working at the intersection of Enterprise Agility and AI governance. He brings more than 30 years of transformation experience across organisations including LSEG, JP Morgan and British Telecom. He is the author of the soon-to-be-published *The Supercharged Enterprise* (Business Expert Press) and *Developing Agile Coaches*, and an acknowledged contributor to PMI's Manifesto for Enterprise Agility.

Further reading

Ward, J. (2026). [Why Your AI Roadmap Should Start With a Strategy Map, Not a Use-Case List.](#)

Ward, J. (2026). [Strategy Mapping: Using AI to Do the Legwork.](#)

CIO Dive (2025). [AI project failure rates are on the rise: report](#) (S&P Global Market Intelligence data).

Jonathan Ward



Beneficial Consulting Ltd
36-38 Wigmore Street,
London
W1U 2BP

Web: <http://www.beneficialconsulting.co.uk>
Mobile Tel: +44 (0) 7802 884598