

First principles

If most transformations fail, why not stop and rebuild from first principles?

Most transformation plans begin after the decision is already made. The pace-layer trap, and the questions that get you out of it.

The argument

Most transformation plans begin too late, the real decision is made before the executive team ever sees the roadmap.

By then the platform has a sponsor, the agent has a use case, the partner has a proposal, and the business case has enough confidence to keep the room moving. What gets missed is the quieter question, whether this work should still exist in its current form before anyone automates it.

AI is not landing into clean organisations, it is landing into old workflows, inherited controls, unclear ownership, cultural debt and operating models built for a slower world. The technology layer moves quickly, the organisation around it is expected to keep up, and that is where many programmes start to tear.

Jeroen Tas and Rob Theunissen describe it in Changing Fast and Slow, the operational layer moves fast, the strategic slower, and the existential layer, culture, identity, risk appetite and belief, slowest of all. Most AI programmes change the operational layer while assuming the other two will catch up, and they rarely do.

What most programmes change	What they assume keeps up
The operational layer, fast	Strategy and culture, slower and slowest

From my office, I have been writing down the questions I wish more rooms asked before the spend begins, which workflow deserves to change, what decision is being moved, who owns the outcome when the agent is right and who owns it when it is wrong. That is the work behind the AI BVF and the Workflow Coroner, slowing the first decision down enough that the rest of the programme has a chance.

On the radar

Three signals from the week.

Each one pushes the first-principles question to the front.

Europe · 23 June 2026

High-risk

AI Act classification consultation closes

The European Commission's consultation on high-risk AI classification closed this week, and classification is becoming a design question rather than a legal clean-up after the build. Before automating a workflow, an organisation has to ask what kind of decision is being made, who it affects, what harm could follow, and whether the system belongs near that work at all.

India · 23 June 2026

96%

of govt leaders advancing sovereign AI

An IDC-Dell study reported that 96 percent of Indian government leaders are advancing sovereign AI strategies, with digital public infrastructure as both governance framework and innovation platform. The move from AI as a tool decision to AI as national operating architecture forces the basic questions back into view, whose data, whose infrastructure, whose trust, whose resilience, whose control.

Five Eyes · 22 June 2026

Months

the window to a cyber capability shift

The Five Eyes cyber agencies warned that frontier AI could change offensive and defensive cyber capability within months, and cyber can no longer be treated as a technical issue. When AI systems gain access, autonomy and speed, the questions become unavoidable, what can it touch, who can stop it, what happens when it is wrong, and who owns the decision.

Against the grain

Stopping feels expensive when the programme already has momentum, but the expensive part is usually continuing with a poor question because the calendar has moved on. A week spent rebuilding the problem from first principles can save a quarter spent making the wrong work faster.

The reframing question

Which layer are you really changing, operational, strategic or existential, and which layers are you quietly assuming will keep up?

What's next

Issue 13 lands Friday 3 July 2026, the workforce question behind agentic transformation, who owns the work when the organisation becomes part human, part agent, and the old org chart no longer explains how decisions really move.