


BEYOND THE VETO

Built to Say No: The Long Roots of Risk and Compliance's Exile from Transformation

Why the functions everyone now wants as transformation partners were built, over two decades, to do the opposite — and what that inheritance would take to undo.



Giovanni Leonardi

September 2007

Contents

1. Executive Summary	3
2. The Seat at the End of the Table	3
3. Chartered by Catastrophe	4
4. The Independence Trap	5
5. The Case for the Distance	6
6. What Partnership Would Actually Require	7
7. The Partner Who Must Sometimes Refuse	8

Built to Say No: The Long Roots of Risk and Compliance's Exile from Transformation

Transformation · No. 155

Originally written September 2007 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

"A risk function that always agrees is not a partner; it is an accomplice, and a board relying on it is less safe than it believes."

1. Executive Summary

Across two decades of corporate governance reform, organisations built a set of functions — internal audit, compliance, enterprise risk — whose explicit charter was to prevent the last disaster from recurring. Each was created in the aftermath of a shock, and each was designed, deliberately, to stand apart from the business it watched. That distance is what makes their assurance credible. It is also what has left them sitting at the end of the transformation table, consulted late, framed as an obstacle, and quietly routed around by the people actually designing change.

This essay traces the roots of that exile. The argument is that the estrangement between the control functions and the transformation agenda is not a matter of personality, obstinacy, or poor stakeholder management, as it is usually diagnosed. It is structural, and it was built on purpose. The very features that make risk and compliance trustworthy — independence, a mandate to challenge, accountability for what goes wrong rather than for what gets built — are the features that make them poor partners in the ordinary sense of the word.

The fashionable call to turn these functions into "transformation partners" therefore asks for something harder than a change of attitude or a redrawn reporting line. It asks an organisation to hold two things at once: a function close enough to shape design, and independent enough to still refuse. Most organisations resolve that tension by quietly abandoning one side of it — either keeping the function pure and irrelevant, or making it agreeable and useless. This essay argues that the partnership is possible, but only by separating the distinct jobs the second line actually does, moving its involvement upstream from the gate to the drawing board, and changing what it is measured on. And it argues that even done well, the partnership will remain uneasy — because a control function that can never say no has stopped being one.

2. The Seat at the End of the Table

Sit in on the steering committee of almost any large change programme and you can find the pattern within the first ten minutes. The room is organised around the people building the thing: the sponsor, the programme director, the design leads, the delivery partner. Somewhere down the table — often literally at the end, often dialling in — is the representative from risk, or compliance, or both. They are there, as the invitation politely put it, "for completeness."

What is striking is not that this person is marginal. It is how thoroughly everyone has adapted to their marginality. The genuinely consequential design choices have usually been made before this meeting, in rooms to which the control functions were not invited. The compliance representative arrives with a list — and everyone present already knows it will be a list of objections, because that is what arrives from that seat. The programme has learned to schedule the real decisions early and to treat the control function's input as a gate to be passed rather than a source to be consulted. When the objections come, they are absorbed as friction, costed as delay, and, where possible, deferred.

We tend to narrate this as a failure of relationship. The risk team is "not commercial enough"; the compliance officer "doesn't understand the business"; the programme "doesn't engage risk early." All of this may be true in any given case, and none of it explains why the same configuration reappears in organisation after organisation, sector after sector, regardless of the individuals involved. When a pattern is that robust to changes in personnel, the cause is rarely personnel. The seat at the end of the table was built into the furniture long before the current occupants sat down in it.

3. Chartered by Catastrophe

To understand why, it helps to trace where these functions came from — because none of them was designed to be a partner to change. Each was designed to be a check on it.

Internal audit is the oldest of the three and the clearest case. Its purpose, from the outset, was to be the board's eyes inside an organisation the board could no longer see directly — an independent verifier reporting, ultimately, not to management but past it. Everything about the function's construction serves that independence: its reporting line, its access, its right to look where it likes, and its studied distance from the operations it examines. It was never meant to help build; it was meant to tell the board, honestly, what had been built and whether it could be trusted.

The compliance function, in its modern form, is younger and more clearly a creature of scandal. The collapses that opened the decade — Enron, WorldCom, and the wave of restatements and prosecutions that followed — produced, in short order, the Sarbanes-Oxley Act and its now-infamous Section 404, with its demand that management document and attest to the effectiveness of internal controls over financial reporting. In the United Kingdom the lineage runs differently but arrives in a similar place: from Cadbury's early code, through the Turnbull guidance on internal control, to the Combined Code that made "comply or explain" the governing grammar of the boardroom. A whole profession expanded to service these obligations. Its mandate was explicit and backward-looking: ensure that what went wrong last time cannot go wrong here.

Enterprise risk management is the youngest and, in some ways, the most revealing. It arrived with real intellectual ambition — the COSO enterprise risk framework set out to lift risk thinking above the silo of insurance and treasury and make it a whole-organisation discipline, and in banking the new capital accord went further still, putting a price on operational risk and so forcing institutions to measure and hold capital against the possibility of their own failures of process and control. The Chief Risk Officer emerged in these years as a boardroom figure in a way that would have seemed odd a decade earlier. Yet even at its most ambitious, the discipline was framed around loss: risk appetite, key risk indicators, the taxonomy of things that could go wrong. Its vocabulary is the vocabulary of prevention.

Read the founding documents of all three functions and the same sentence is written underneath each: you exist so that the last disaster does not happen again. That is a noble charter. It is also, precisely, a charter to say no.

This is the first root of the exile, and the deepest. These functions were not neglected into their adversarial posture; they were commissioned into it. An organisation that has just been burned does not create a function and ask it to help move faster. It creates a function and asks it to stand in the way of the specific thing that burned it. The orientation toward prevention is not a cultural accident that better leadership could have avoided. It is the job as written.

4. The Independence Trap

If the charter is the first root, independence is the second — and it is the one that turns a difference of orientation into a structural impossibility.

The credibility of every assurance function rests on its distance from what it assures. A compliance sign-off means something only because the person giving it did not design the thing they are signing off. An internal audit opinion is worth having precisely because the auditor has no stake in the verdict. This is not a bureaucratic nicety; it is the entire source of the function's value. The moment the watchdog helps build the kennel, its bark stops being informative.

But distance has a cost, and the cost is paid in exactly the currency that transformation runs on: early, formative involvement. The places where a change programme is genuinely shaped are the early design workshops, the architecture debates, the trade-off conversations where someone decides what the new process will and will not do. These are collaborative, iterative, and messy. To be useful in them, you have to get your hands dirty — to co-own choices, to advocate, to compromise. And every one of those acts spends down the independence that is the control function's reason to exist. The function faces a genuine dilemma, not a false one: it can be in the room where value is designed, or it can retain the standing to judge what was designed, but the structure makes it very hard to do both.

Measurement deepens the trap. Consider how the two sides of a transformation are actually assessed. The delivery organisation is measured on what it ships: scope delivered, dates met, benefits promised. The control functions are measured on the absence of bad things: breaches avoided, findings closed, audits passed clean. This asymmetry is quietly devastating. A programme director who takes a risk that pays off is a hero; a compliance officer who waves through a control weakness that later fails has committed a career-ending error, while one who blocks a perfectly good initiative out of caution pays almost nothing. The incentives do not merely fail to reward partnership; they actively punish the control function for the exposure that partnership requires.

The economics of when the function engages make the pattern concrete. Consider a composite that will be familiar to anyone who has worked near a core-systems replacement. A programme to replace an ageing customer and billing platform runs its design and build largely within the delivery organisation. Compliance is formally "engaged" — a name on a stakeholder map, a standing invitation to a governance forum few of them attend. At the pre-go-live gate, eleven months in, a reviewer reads the data model properly for the first time and realises the new platform retains full customer records indefinitely and cannot honour a request to delete them, in a way the organisation's own

data protection obligations do not permit. The finding is correct. It is also catastrophic in its timing.

Same issue, two moments	At the go-live gate (month 11)	In the design workshop (month 2)
What it is	A blocking finding on data retention	A design constraint among many
Cost to resolve	~£2.4m of rework, a four-month slip	~£180k of design effort, no slip
How it is experienced	Compliance "stopped the programme"	A requirement, like any other
What everyone remembers	The function is an obstacle	Nothing – it simply worked

The point of the composite is not the figures, though figures like these are entirely ordinary. The point is that the same intervention is a partnership in month two and an act of sabotage in month eleven, and nothing about the intervention itself has changed – only its timing. Yet the structure pushes the function relentlessly toward month eleven: excluded from the early rooms by the independence doctrine, pulled in at the gate because that is where the sign-off lives, and then blamed for the expense of a lateness the structure imposed on it. The function is set up to arrive too late and then criticised for arriving too late.

5. The Case for the Distance

It would be easy to end there, with the control functions as sympathetic victims of a structure that wrongs them. But the honest version of this argument has to take seriously the possibility that the exile is not a bug at all – that the distance is right, and the whole partnership agenda is a mistake dressed up as progress.

The case is stronger than its critics allow. Independence, on this view, is not an obstacle to the function's value; it is the value, and it is fragile in a specific way – it erodes through proximity, quietly and without anyone deciding to abandon it. A compliance team that spends its days co-designing the transformation, sharing its authors' hopes for it, invested in its success, will find it almost impossible to be the one that says, at the decisive moment, that the thing they helped build cannot go ahead. Not because they are corrupt, but because they are human. The capture is not a bribe; it is the ordinary loyalty that forms among people who build something together. A function that has been made agreeable enough to be welcome in the design room may have surrendered the one thing that made it worth having there.

There is a harder edge to the argument still. Much of what is sold as "risk partnering" is, on inspection, the business looking for a way to launder its own risk decisions through a function that will bless them. If the control function is close enough to be co-opted, its sign-off becomes theatre: the appearance of challenge with none of the substance, which is worse than no challenge at all, because it manufactures false comfort for a board that is relying on it. On this reading, the awkward officer at the end of the table who insists on their objections and refuses to be managed is not the failure of the system. They are the system working. The friction is not a defect to be engineered away; it is the sound of a real check doing its job.

This case deserves more respect than the partnership literature usually gives it, and any serious answer has to concede its central point: a control function that has lost the capacity to refuse has lost its reason to exist, and closeness genuinely threatens that capacity. The question is whether distance is the only way to protect it — whether the choice really is as binary as both the exiles and the reformers assume.

6. What Partnership Would Actually Require

It is not binary, but the way out is narrower and more demanding than the usual prescriptions — "engage risk early," "build relationships," "give compliance a seat at the table" — suggest. Those slogans are not wrong; they are simply too shallow to survive contact with the structure described above. A genuine answer has to work with the grain of the independence problem rather than wishing it away.

The first move is to separate two jobs that the second line does and that we habitually confuse. One job is to set the rules and judge, independently, whether they have been met. The other is to help the business understand and respond to the risks it faces. These are different activities, and only the first actually requires distance. There is nothing about advising a design team on how to meet a data protection obligation that compromises anyone's ability to later verify, independently, whether it was met — provided the two roles are held by different people, with the judging role kept clean. The error is to treat the whole function as if it were indivisible, and so to keep all of it at arm's length in order to protect the part that genuinely needs to be.

- Split the advisory role from the assurance role, explicitly and by name. The people who help shape the design are not the people who later sign it off. The advisory side can be as embedded as the programme needs; the assurance side stays independent, and is seen to be.
- Move the centre of gravity from the gate to the design. The expensive interventions all happen at the gate because that is the only point where the function has formal power. Give it a defined role in early design, and most of what it would otherwise catch late it shapes cheaply and early — as a requirement rather than a rejection.
- Change what the function is measured on. As long as risk and compliance are rewarded only for the absence of failures and never for the enabling of success, no exhortation to "partner" will survive the first real trade-off. What gets measured on both sides of the table has to include the same thing: change delivered safely, not merely change stopped or change shipped.

The deepest move is the hardest, because it is not organisational but cultural: the business has to take back ownership of its own risk. A great deal of the dysfunction traced here rests on a quiet fiction — that risk belongs to the risk function. It does not. The accountability for a badly designed process, a mishandled obligation, an unfair customer outcome, sits with the people who designed and ran it, and no amount of second-line activity transfers it. When the business believes that risk is compliance's problem, it under-invests in getting things right and then resents the function for the cost of catching them wrong. When the business owns the risk, the control function stops being the department of no and becomes what it should have been all along: expert help with a problem the business already knows is its own.

“The department of no is not a personality. It is what any function becomes when an organisation outsources its conscience and then complains about the invoice.”

None of this dissolves the independence problem. It contains it — by being precise about which parts of the function need distance and which do not, rather than sacrificing the whole function's usefulness to protect a part of it.

7. The Partner Who Must Sometimes Refuse

There is a reason the word "partner" sits awkwardly on these functions, and it is worth ending on rather than smoothing over. A partner, in the ordinary sense, is someone whose interests are aligned with yours and who wants what you want. But the whole point of a control function is that its interests are deliberately not fully aligned — that at some point it must be willing to want something other than the programme's success, and to say so. A risk function that always agrees is not a partner; it is an accomplice, and a board relying on it is less safe than it believes.

So the aspiration cannot be a control function that never obstructs. It has to be one that obstructs rarely, early, and for reasons the business recognises as its own — a function whose default is help, and whose refusals, when they come, carry the weight of someone who has earned the right to them by being useful the rest of the time. That is a harder and more interesting thing than either pole of the current debate imagines. It is not the pure, distant watchdog barking from the end of the table, and it is not the embedded, agreeable adviser who has forgotten how to bark at all. It is the uncomfortable middle that most organisations avoid precisely because it is uncomfortable.

The roots of the exile run deep, and this essay has tried to show that they are structural rather than personal — chartered by catastrophe, entrenched by the genuine logic of independence, and reinforced by measurement that rewards exactly the wrong things. Those roots explain why the problem is so stubborn, and why the easy prescriptions fail. But roots explain a difficulty; they do not excuse a surrender to it. The functions were built to say no because, at the time, saying no was the whole of what was wanted from them. What is wanted now is more than that, and harder: the judgement to know when no is the only honest answer, and the standing — earned, not asserted — to make it heard the rest of the time.