



BEYOND THE CLOSING GATE

The Compliance Programme That Never Ends

Why regulated firms keep paying twice for the knowledge they throw away at the closing gate

Giovanni Leonardi

November 2017

The Compliance Programme That Never Ends

Programme · No. 175

Originally written November 2017 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“There is no slide that shows the fines you did not receive.”

1. The programme that closed on a Friday

Every regulatory programme has its closing ceremony. I have sat through more than one: the final gate cleared, the benefits statement signed, a slide with a single large green tick, and a real sense of relief around the table. One such programme — two years, the better part of £12m, a peak team of forty — was declared closed on a Friday afternoon. The specialists went back to their business units or on to the next engagement. The regulatory-change analyst who actually understood how the new reporting obligations mapped onto the firm's product taxonomy rotated onto a different account inside a month.

The following spring a consultation paper landed that reopened the very obligations the programme had just embedded. There was no team to receive it. There was a service-desk ticket, an anxious email chain, and eventually a hastily convened group reading the old programme's documentation, trying to reconstruct why a particular control had been built the way it had. Nobody could say. The knowledge had walked out at the closing ceremony.

I want to argue with this pattern, because it is not the story of one badly run programme. It is the predictable result of a category error we make repeatedly: we treat regulatory divergence as a project, when it has become a permanent condition of operating.

2. A project has an end date; a condition does not

The textbook framing is comfortable and wrong. A regulation is published; a programme is stood up to implement it; the programme delivers and closes; business as usual absorbs the result. That sequence assumes a world in which the rules eventually settle — in which compliance is a destination you can arrive at and then stop.

That world, if it ever existed, is gone. For most of the years after the financial crisis the regulatory current at least ran one way: towards harmonisation, common standards, a single rulebook. What we are living through now is the opposite. The obligations multiply and, more importantly, they diverge. A firm operating across the UK and the Continent is watching MiFID II come into force in a matter of weeks, preparing for the

General Data Protection Regulation the following May, standing up ring-fencing ahead of the 2019 deadline, and absorbing the Senior Managers regime — while, over all of it, the Article 50 negotiations make it impossible to say which of today's shared rules will still be shared in two years. The same transaction, booked in two jurisdictions, is beginning to attract two different and drifting sets of requirements.

Divergence is not a bigger version of compliance. It is a different kind of problem: not implement the rule but maintain a firm that can keep implementing rules that no longer agree with each other.

Once you see it that way, the closing ceremony looks less like success and more like the deliberate dismantling of the one capability the firm most needs.

3. Why the fiction survives

If the project framing is so plainly mismatched to the reality, why does it persist? Not because practitioners are naïve. It survives because it is the only language our governance and funding machinery understands.

Money inside a large organisation is released against a business case, and a business case wants a defined scope, a fixed cost, an end date, and a benefits line that can be booked and closed. A programme is the vehicle purpose-built to carry those four things. A permanent condition has none of them: its scope is open-ended, its cost recurs, it never ends, and its benefit is the avoidance of a harm that, done well, never materialises and so is never noticed. There is no slide that shows the fines you did not receive.

So the organisation does the only thing its own accounting will permit. It repackages a standing obligation as a temporary project, funds it as change, delivers it, and closes it — because closing it is what releases the next tranche of capital for the next thing. The regulatory reality is continuous; the funding model is discrete; and the gap between them is filled, every time, by the quiet loss of institutional memory.

- The scope is written to be finite, so anything genuinely open-ended is defined out of it.
- The team is contract-shaped, so it disperses precisely when its knowledge becomes most valuable.
- The benefits are booked at closure, so there is a positive incentive to declare the work done rather than ongoing.

None of these is a failure of individual competence. They are what the system rewards.

4. "But you cannot fund a programme forever"

The strongest objection to everything I have said is not weak, and it deserves stating at full strength. It runs like this: programmes have end dates for good reason.

Open-ended change work is how you get zombie programmes — the ones that consume budget for years, never quite deliver, and cannot be killed because no one can define what finished would look like. Discipline, closure, and the hard question "are we done yet?" are not bureaucratic obstacles; they are what stops good money following bad. Declare compliance a permanent programme, the argument goes, and you have licensed exactly the unaccountable spending that governance exists to prevent.

This is correct — but it is an argument against the permanent project, not against the permanent capability, and the two are not the same thing. The mistake is to assume the only alternative to a project that ends is a project that never does. There is a third option, and every organisation already uses it for other standing needs: you fund it as a capability, on the run side of the ledger, with the accountability that run-cost carries — not the accountability of a project.

"We do not run a "security programme" that closes when the firm is finally secure. We staff a security function, because the threat is permanent. Divergence deserves the same honesty."

The distinction matters because it changes what accountable means. A project is accountable for delivering a defined scope on time and to budget. A capability is accountable for a service level — for keeping the firm's rule-to-control mapping current, for turning a new consultation paper into an impact assessment inside a fixed number of days, for never again convening a group to reconstruct why a control exists. Those are measurable, and the capability is killable if it underperforms against them. Discipline is not lost; it is relocated to where the work actually lives.

5. The standing capability

What does the alternative look like in practice? Not a forty-person army kept permanently on the payroll — that would be the zombie the objection rightly fears. It is a small, senior, permanent core with a defined surge model around it.

1. A permanent core that owns the maps, not the projects. The most valuable artefact the closed programme produced was not the delivered control; it was the mapping between obligations and controls — which rule drives which check in which system. Treat that map as a maintained asset with a named owner, not as programme documentation that is filed and forgotten.
2. A horizon-scanning function whose only job is to see the next divergence coming while it is still a consultation paper, not a deadline. The spring surprise in my opening was only a surprise because no one was paid to be watching.
3. A pre-agreed surge model — framework contracts, retained specialists, a standing call on internal capacity — so that when a genuinely large change lands, the firm scales up around a core that remembers, rather than assembling forty strangers

from scratch.

4. A funding line that names the truth: run-cost for the standing capability, change-cost for the genuine one-off build. The category error is, at root, an accounting error, and it is fixed in the accounting.

The economics are not subtle. The rediscovery I described — the group reconstructing a control's rationale from old documents — is not free. On the programme I have in mind it consumed something close to three months of senior specialist time to answer a question the departed analyst could have settled in an afternoon. Multiply that across every gap between every programme and the saving from standing the team down reveals itself for what it is: a deferral, repaid later with interest.

The test of whether a firm has made this shift is brutally simple. When the next unexpected consultation paper lands — and in this climate it always does — is there a name against it by the end of the week, or is there a service-desk ticket?

6. What the textbooks leave out

The transformation literature is full of guidance on running a programme well: how to gate it, resource it, govern it, and close it cleanly. All of that advice is sound, and none of it saves you here, because the literature quietly assumes the thing it is helping you build is temporary. Its silence on what happens after closure — on how a firm holds regulatory knowledge across the gaps between programmes — is the omission that matters most in an age of divergence.

The lesson I would carry into any regulated organisation now is not a better closing gate. It is the recognition that some of what we habitually call programmes were never projects at all. They were standing conditions we financed as if they were events, and paid for twice — once to build the capability, and again, six months later, to rediscover what we had just thrown away.