



ACON – THE REMEDIATION PAPER

Delivering remediation after the review.

*A practitioner's guide to the programme that decides
whether the Section 166 ever really ends.*

An ACON white paper by Caitlin Taylor • August 2026 • v2 public edition



A practitioner's guide to the programme that decides whether the Section 166 ever really ends

This paper is for firms whose review is already over. The skilled person has reported, the findings are agreed, the remediation plan is committed – and the response programme that produced it is winding down. It is a companion to Running the bank's side of a Section 166 (ACON, 2026), and it stands alone: nothing here assumes ACON ran, or even saw, the response. Part I explains why remediation programmes fail so predictably. Part II sets out what good looks like – usable whether or not you ever work with me. Part III describes how ACON delivers remediation as a standalone engagement, from Stage 2 alone.

Executive summary

When the skilled person's report lands, most firms breathe out. It is exactly the wrong moment. The review was scoped by the regulator, run by the reviewer, and paid for whatever happened. The remediation is the firm's alone – and it is the part the regulator will actually judge. The FCA's commissioning practice includes appointing skilled persons specifically to test whether earlier remediation worked; its enforcement record includes firms fined years after weaknesses were first identified, because the fix stalled, drifted, or quietly never happened. A remediation that fails re-review is not a delay. It is a second review, a second fee, and a supervisory relationship that starts every subsequent conversation from behind.

Almost every remediation programme is born the same way: in a hurry, from a plan written to reassure rather than to deliver, staffed by a response team that is exhausted and already leaving. This paper works through the two reflexes that shape that moment and the seven missteps that follow from them (Part I); the decisions, structures and closure disciplines that a well-run remediation actually uses – including the closure pack and the standalone Verification & QA function (Part II); and how ACON packages those disciplines into a gated Stage 2 engagement that mobilises in three weeks and ends on signed acceptance (Part III).

PART I – WHY REMEDIATION PROGRAMMES FAIL

I. Remediation is not the victory lap

A remediation programme inherits everything the response programme was and adds three properties of its own – and each one defeats a conventionally run programme:

The commitments are already made. The remediation plan was submitted to the regulator under the pressure of the report's arrival. Its dates are, in effect, contractual; under Principle 11, slippage against them is itself a reportable event. Most plans were written in weeks, by people who would not have to deliver them, to a standard of "credible in a meeting" rather than "deliverable in a business unit." The programme starts life measured against a posture.

The audience has changed sides. During the review, the skilled person was a consumer of evidence the firm produced reactively. During remediation, the firm must prove a negative – that the weakness is gone – proactively, to an examiner who may not arrive for two years. Everything must be built to be tested later, by someone with no context and no goodwill: dated, attributable, contemporaneous evidence, or it did not happen.




FIG 1 – THE INDEPENDENCE BOUNDARY – UNCHANGED AFTER THE REPORT

The energy is spent. The response consumed the firm's best people for one to two years. The natural instinct at report delivery is demobilisation – just as the longest, least glamorous, most consequential phase begins. Remediation is a marathon run by people who think they have just finished one.

2. The two reflexes

Where response programmes fail on two beliefs ("structure takes too long" and "governance kills agility"), remediation programmes fail on two reflexes – both triggered in the fortnight after the report lands:

Reflex one: announce fast, plan later. The board wants the regulator reassured; the regulator wants dates. So dates are produced – from the report's recommendations, not from root causes; from the top down, not from the business units that must deliver them. The plan is a press release with milestones. Institutions rarely approach remediation in a way suited to its regulatory character: they structure it like an ordinary change programme, with the added handicap of a knee-jerk, reactive start. Teams overpromise under pressure from senior management and set themselves up to under-deliver – and under-delivery on a committed remediation is not a project variance, it is a supervisory event.

Reflex two: fix everything at once. The review exposed more than it cited. The findings become the pretext for a broader transformation – a new platform, a target operating model, a culture programme – and the regulatory commitments are welded to a discretionary portfolio three times their size. The ambition is usually right. The coupling is always wrong: when the transformation slips (and at industry norms of 30–45% overrun, it will), the regulator's milestones slip with it.

Both reflexes feel like decisiveness. Both are the opposite: decisions made before the programme exists, that the programme then spends years absorbing.

3. The seven missteps

The two reflexes cash out in seven recognisable failure modes. Any practitioner who has lived a large remediation will recognise most of them; the unlucky have met all seven.

1. Fixing the finding, not the cause. Actions are mapped to the *wording* of each finding rather than the failure underneath it. The symptom clears; the underlying control still does not operate; the re-review finds the same weakness wearing new documentation. Root-cause discipline is the foundation of every corrective action that follows – skip it and everything built on top is decoration.

2. **"We fixed it" isn't "we can prove we fixed it."** The most common failure at validation, on both sides of the Atlantic: work genuinely done, evidence reconstructed afterwards, credit refused. Examiners credit board-approved policies with dates, training records with names, samples showing controls operating without ad-hoc manual intervention. Firms that remediate well but document poorly are marked non-compliant at validation — and the clock restarts.

3. **The "while we're at it" programme.** Reflex two, matured. Regulatory milestones held hostage by a transformation timetable; closure evidence citing future internal deliverables ("addressed by the new platform in Q3"), which a regulator reads as "not done." Ambition is not the error — coupling is.

4. **Remediation run as a project sprint.** Treated as ordinary change with a deadline, the programme gets a heroic plan, an exhausted team, and controls that only operate with programme support. Regulators increasingly test sustainability, not completion: whether the fix survived contact with business as usual. A sprint cannot prove durability; only an operating discipline can. The human cost compounds it — high programme intensity and repeated resets drive exactly the staff turnover that destroys institutional memory mid-remediation.

5. **Nobody in the business owns anything.** "The first line will own remediation" is how silent risk transfer starts. Actions close into a void: job titles instead of names, no operator for the new control, nobody who signed for it. Six months later the control has no owner and the finding is back — and the FCA's enforcement record shows deficiencies sitting on risk registers for years precisely this way.

6. **Three versions of the truth.** The board pack, the regulator update and the audit trail drift apart; the weekend reconciliation tax begins; credibility — still the programme's only real currency — drains with every discrepancy. Remediation multiplies the audiences of the response (the supervisor, the board, internal audit, and possibly a monitor or returning skilled person), so divergence is found faster and forgiven less.

7. **The invisible slip.** Without a live critical path, a moved date is invisible until it detonates downstream — by which time it is weeks old and reportable. On remediation the stakes are higher than on the response: the dates were commitments in the plan the regulator approved, and the supervisory tolerance that existed during the review has been spent.

4. The numbers

The re-review is real. The FCA commissioned 83 skilled person reviews in 2023/24 and 48 in 2024/25; testing the effectiveness of earlier remediation is an established commissioning purpose. A firm's remediation is not finished when the firm says so — it is finished when it survives that test.

The enforcement record. The pattern of "stalled remediation" runs through FCA final notices: weaknesses identified, logged, and fined years later, unfixed. Commerzbank London was fined £37.8m in 2020 after the FCA had repeatedly flagged AML weaknesses that remediation failed to resolve; Starling Bank was fined £29m in 2024 after failing to implement a formal programme to comply with a financial-crime requirement it had itself agreed. In the US, Citigroup's 2020 \$400m OCC penalty was explicitly for failing to remediate longstanding deficiencies identified in earlier orders. The supervisory grace period for "work in progress" has, in the FCA's own assertive-supervision era, evaporated.

The overrun tax, re-priced. Projects overrun budgets and schedules by 30–45% on average; roughly 70% of large transformations miss their objectives; weak governance is consistently among the leading causes. On a remediation running at £400–500k a month, a "normal" overrun is seven additional months of run-rate — plus, uniquely to this context, a Principle 11 conversation, a longer shadow over every supervisory interaction, and a materially higher chance of the re-review that starts the cycle again.

The sponsorship factor still governs. Prosci has ranked active, visible sponsorship the top contributor to change success in every benchmarking study since 1998. On remediation the sponsor is typically an SMF holder whose personal attestation will close the programme. Sponsorship is not a soft factor here. It is a control — and the signature at the end concentrates the mind wonderfully.

PART II — WHAT GOOD LOOKS LIKE

Usable whether or not you ever engage ACON. It reflects the audiences a remediation must satisfy, verified against published supervisory expectations and hardened on a live multi-year banking S166 that ran to signed acceptance.

5. The first 30 days: inherit, separate, baseline

By day 30 the remediation's fate, like the response's before it, is largely set. The calls that matter:

Days 1-5 — Audit the inheritance. Findings, commitments, correspondence, and whatever tooling and half-built trackers the response left behind — catalogued against a single register before anything is re-planned. What was promised, to whom, by when, in which document. Most programmes discover their real commitments in month six, from the regulator. Do it in week one, from the files.

Days 1-10 — Separate the lanes. Every candidate action classified: regulatory commitment or internal goal. The regulatory lane is bounded, baselined and protected by change freezes. The internal lane — the "while we're at it" work the findings genuinely justify — is governed in its own lane, with its own gates and its own business case, from the same golden source. This is the severability rule: no regulatory closure may ever depend on an internal deliverable. It is the single decision that makes holistic ambition safe.

Days 5-15 — Name the accountable humans. The SMF owner who will sign the final attestation; a sponsor who understands the role; a programme director empowered to hold gates; and a named BAU owner — a person, not a job title — against every finding. Ownership named at day 15 is cheap. Ownership discovered missing at closure is a follow-up review.

Days 10-20 — Re-baseline honestly, once. The committed plan was written under reflex one. If its dates are undeliverable, the worst strategy is discovering that publicly, one slipped milestone at a time. Build the real plan bottom-up with the business units, then have the honest conversation with the supervisor early, once, with a credible replacement — early self-reporting is a control, not a confession. Every date kept after that conversation must then actually be kept.

Days 15-30 — Define closure before delivering anything. The closure standard — what a finding must prove to close — is agreed now, with internal audit and, where possible, socialised with the supervisor: root cause closed, design effectiveness, operating effectiveness on sampled evidence, sustainability with named ownership. Programmes that define closure at the end negotiate it; programmes that define it at the start measure it.

6. The closure pack: the unit of proof

A remediation is delivered finding by finding, and each finding closes on a pack — assembled as the work completes, never reconstructed after it. The ACON standard pack runs to twelve tabs: closure summary; finding and root cause; actions delivered with authority trail; design effectiveness; operating effectiveness testing (methodology, population, sample, results, defect disposition); the Verification & QA certificate;

evidence index and manifest; BAU integration statement with named, trained ownership; sustainability and monitoring plan with thresholds; the sign-off matrix; residual risk and dependencies; and the regulator correspondence extract mapping every commitment to where it was met.

Two disciplines matter more than the tab count. First, evidence is generated as a by-product of doing the work — if a pack cannot be assembled from the golden source in a day, the programme has an evidence problem it has not yet noticed. Second, the signatures run in the right order: the business owner first, because the business is accepting the risk; the SMF last, on a pack already independently certified — which is what makes the attestation survivable.




FIG 2 – ONE SOURCE, SIX VIEWS – THE CLOSURE EVIDENCE BASE

7. The Verification & QA function: closure as measurement, not negotiation

Delivery teams mark their own homework optimistically everywhere, always — not from bad faith, but proximity. The answer, proven at scale on the reference programme, is a standalone Verification & QA function through which every closure passes: **tested** (design walkthrough, then attribute-based operating-effectiveness testing on a documented sample), **QA'd** (a second reviewer checks the testing itself — method, sample integrity, evidence provenance), **validated** (pack completeness against the closure standard), **retested** (defects returned with severity-graded notes; critical defects invalidate the closure basis, majors trigger targeted retest on a fresh sample), and only then **passed** — a certificate with names and dates at every step, without which the close gate cannot convene.

The independence rules are absolute even when the team is small: the builder never tests, the tester never QAs their own work, and the function cannot be overruled by a delivery deadline. Its registers are read-only visible to internal audit from day one; its MI — first-time-pass rate by workstream, defect density trended, closure funnel aging — is the most honest picture of the programme's real state that a board will ever see.

8. Ten questions your board should ask when the report lands

1. What exactly did we commit to, in which documents, and who has read them all?

2. Which of these actions close findings, and which are things we want to do anyway — and are those separable today?
3. Are the plan's dates the business units' dates, or the meeting's dates?
4. Who owns each finding's remediation in the business — by name, in writing?
5. What is our closure standard, and would internal audit co-sign it?
6. Who independently verifies closure, and can a deadline overrule them?
7. If a milestone moved yesterday, when would we know — and when would the FCA?
8. Could we evidence, today, everything we will one day claim was done?
9. What happens to each new control on the day the programme stops paying for it?
10. Who signs the final attestation, and what will they be shown before they do?

A programme that cannot answer 5, 6 and 8 quickly does not have a delivery problem yet. It has a proof problem — and on remediation, proof is the product.

PART III — THE ACON ANSWER

9. Stage 2 stands alone

ACON's delivery system was built in two stages — Response, then Remediation — and the second was designed from the start to be engaged on its own. If the response is already done — by you, by your advisers, by anyone — ACON mobilises directly into Stage 2: **Handover & Mobilisation → Remediation Delivery → BAU Integration & Close**. Three phases, three hard gates, one golden source. You do not need to have run the review with me; you need a findings list, and the willingness to name owners.

Mobilisation (weeks 1–3) executes Part II's first 30 days as a pre-built system rather than a design exercise: the inheritance audit, the two-lane scope with the severability rule enforced in the workbook itself, named owners in both lanes, the honest baseline, and the closure standard agreed before delivery starts. The golden-source workbook — SLA-timed workflows for actions, evidence and closure packs, one refresh updating every report — is generated, tested and live by week three, on the Microsoft estate the bank already licenses, trusts and audits.

Delivery (week 4 onward) runs the regulatory lane behind change freezes and the internal lane behind its own gates, with same-day breach visibility on a live critical path and the six-audience reporting model of the response system intact — supervisor, board, audit, workstreams, legal and, when they return, the skilled person, all reading one dataset.

Close (gated, finding by finding) runs every pack through the Verification & QA function to a signed business-unit acceptance. No signature, no closure; no certificate, no gate. At programme end you keep the populated toolkit, the SOPs, a trained team — and a closure file built for the day the regulator asks whether the remediation worked. Because someone will ask. That is rather the point.




FIG 3 – TWO STAGES, SIX GATES – STAGE TWO STANDS ALONE

10. The economics, briefly

Remediation is where the response's cost logic inverts: the fee is gone, and everything that remains is run-rate – which means everything that remains responds to how the programme is run. The levers are the same as the response system's (no build phase; manual reporting cut by up to 80% on the reference programme's measurement; a small senior cell instead of a pyramid; slippage attacked at the root) plus one unique to Stage 2: the re-review you never have. A second skilled person engagement – fee, response costs, and two more years of supervisory shadow – is the single most expensive line item a remediation can generate, and the entire system is designed so that it never appears.

11. Standing, people, engagement

ACON is my bank-side delivery practice – permanently bank-side: not a skilled person, not seeking to be one, and structurally incapable of the independence conflict that position would create. The system was created by remediation-experienced practitioners; I was design and delivery lead on a c.£120m multi-year banking S166 that ran, after the model's adoption, to 100% on-time regulatory submissions with every phase passing skilled person review first time – including the closure model this paper describes. Engagement is a fixed-term contract with me – from Stage 2 alone, for the full Section 166, or for remediation and target operating model programmes in regulated environments more broadly. The best moment to talk to me is before the remediation plan is baselined. The second-best moment is now.

A note on what this paper does not say

Part II describes what a well-run remediation does; it does not describe how the ACON suite achieves it. The workflow models, data architecture, control logic, closure-pack generation and V&QA tooling are proprietary, protected as trade secrets and by copyright, and demonstrated live under NDA.

© Agile Controls (ACON), 2026. All rights reserved. This document is free to read and share unmodified – no registration required. It may not be used to develop, or assist in developing, competing products or services. ACON™ and Agile Controls™ are trade marks of the proprietor.

Appendix — research base (published sources)

FCA skilled person commissioning data (83 reviews 2023/24; 48 in 2024/25) and re-review commissioning practice; FCA Handbook SUP 5, EG 3.3, SUP 15, Principle 11; FCA final notices — Commerzbank London (2020, £37.8m), Starling Bank (2024, £29m), Canara Bank (2018); IIA analysis of 97 FCA final notices ("stalled remediation"); OCC consent-order validation practice — Citigroup (2020, \$400m failure-to-remediate), TD Bank (2024); consent-order sustainability and independent-validation expectations (Crowe; Protiviti; practitioner walkthroughs); remediation programme failure research — BeyondFS "Fine Prevention" white paper (knee-jerk initiation, overpromising, Definition of Done, closure frameworks, staff turnover); root-cause and evidence-standard guidance (Bridgeforce; audit issue-lifecycle standards); McKinsey overrun and McKinsey-Oxford project-value studies; transformation failure research (~70%); Prosci sponsorship benchmarking; controls-testing sampling conventions. Full citations available on request.