

BEYOND CONSENSUAL FICTIONS

Benefits Realisation Died — Value Assurance Is What Replaces It

On building the machinery that holds portfolio investment to account

Giovanni Leonardi

April 2026



Benefits Realisation Died — Value Assurance Is What Replaces It

Portfolio · No. 117
April 2026

"Doctrine without machinery is aspiration."

1. The Quarterly Fiction

There is a meeting that takes place in most large organisations once a quarter. It goes by different names — benefits review, value tracking, realisation checkpoint — but its choreography is invariant. A programme manager presents a spreadsheet of promised benefits against a baseline established two years earlier. The figures are green, because the programme manager wrote both the forecast and the actuals. The portfolio board receives the spreadsheet, notes its greenness, and moves to the next agenda item. Nobody in the room believes the numbers. Nobody says so.

This is benefits realisation management as it is actually practised, and it is finished — not because it has been formally retired, but because its institutional credibility is exhausted. The question that matters now is what replaces it. The answer is emerging, unevenly, under the name value assurance, and it is not a rebrand. It is a different discipline with different machinery, different ownership, and a different relationship to the truth.

2. What Actually Failed

The autopsy of benefits realisation is brief, because the cause of death is structural rather than accidental. Three design flaws made it unrecoverable.

First, the benefits case was an advocacy document. It existed to win funding. The numbers in it were chosen to clear a hurdle rate, not to describe reality. Everyone involved understood this — the programme sponsor who signed it off, the finance partner who reviewed it, the portfolio board who approved it. The fiction was consensual and embedded at inception.

Second, the programme that promised the benefits was the programme that reported on them. This is the moral hazard at the centre of the entire practice. No institutional mechanism existed to separate the claim from the audit. Self-reporting on your own promises, in a culture where programme cancellation is career damage, produces exactly the result you would expect: persistent green status until the money runs out, followed by quiet closure and a benefits shortfall that nobody formally records.

Third, harvesting was positioned after delivery. Benefits were supposed to be "realised" once the programme handed over its outputs. But the gap between an output delivered

and a benefit materialised is where all the difficulty lives — and it is precisely the gap that benefits realisation management had no operational machinery to occupy. The programme team disbanded. The benefits spreadsheet moved to a shared drive. Nobody harvested anything.

These are not implementation failures that better training or tooling might fix. They are architectural defects. The practice was designed around documents rather than data, around self-attestation rather than independent measurement, and around retrospective accounting rather than forward-looking decision-making. It cannot be repaired. It must be replaced.

3. The Four Inversions

Value assurance, as it is beginning to be practised in the organisations that have moved furthest from the old model, rests on four principles that directly invert the failures above.

Benefits Realisation	Value Assurance
Benefits case — asserts a fixed return to win funding	Value hypothesis — proposes a falsifiable claim, testable against operational metrics
Self-reported milestones — programme tracks its own progress against its own baseline	Streamed measurement — operational data drawn from systems of record, independent of the programme
Programme-owned — the team that promised the benefits reports on their delivery	Portfolio-owned — an independent function assures value claims the way finance assures the accounts
Post-hoc harvesting — benefits assessed after programme closure	Tranche-gate testing — value hypotheses re-tested at every funding decision

The shift in language from benefits case to value hypothesis is deliberate and consequential. A benefits case asserts: this programme will deliver twelve million pounds in cost avoidance. A value hypothesis proposes: if we deploy this capability and the operating units adopt it within these parameters, we expect to observe a reduction in processing cost of this order, measurable through these operational metrics, within this timeframe. The hypothesis is falsifiable. The benefits case was not — that is the entire difference.

The shift from self-reported milestones to streamed measurement is equally structural. Value assurance requires that the metrics which would confirm or refute the hypothesis are identified before funding is approved, and that they are drawn from operational systems — processing volumes, cycle times, error rates, unit costs, customer contacts per transaction. These exist in systems of record and can be fed into a measurement layer without anyone manually compiling a spreadsheet. The operating metric does not care whether the programme team thinks the benefit has been realised. It simply reports what is happening.

4. What the Machinery Actually Looks Like

Doctrine without machinery is aspiration. The organisations making this work have built a specific operational stack, and it is worth being concrete about what it contains.

The foundation is a measurement layer that connects to operational systems and holds the baseline and live metrics for every funded initiative. This is not a dashboard project — it is plumbing. It holds agreed metric definitions, automates data extraction, and provides the single version of truth against which hypotheses are tested. In one portfolio I have observed, this layer draws from eleven source systems across four business units and refreshes daily. The investment to build it was significant. The alternative — continuing to accept self-reported benefits figures that nobody believed — was more expensive in misdirected capital.

Above the measurement layer sits a hypothesis register: a structured record of what each initiative claims it will change, by when, measured how. It replaces the benefits map. It is living — hypotheses are re-stated at each tranche gate in light of what the metrics show. A hypothesis that is not confirmed is not a failure to be hidden; it is a signal that triggers a funding conversation.

The decision layer is the tranche gate protocol: a standard set of questions asked at each funding release, grounded in the measurement layer's data. Has the preceding tranche's value hypothesis been confirmed, partially confirmed, or refuted? Have the conditions changed? Is the portfolio-level value envelope still intact, or has this initiative's revised position reduced the total expected return below threshold? These are not rhetorical questions. They are answerable, because the measurement layer provides the data to answer them.

The whole stack is held together by the value assurance function itself — typically two to five people in a portfolio of thirty to fifty initiatives, reporting to the investment board with the same independence as internal audit. They do not manage programmes. They do not write business cases. They maintain the measurement infrastructure, prepare the hypothesis assessments for each tranche gate, and publish a quarterly portfolio value position that the board can trust because it is built from operational data rather than programme self-reports.

5. The Cultural Precondition

Every element described above is implementable. The measurement technology exists. The tranche-funding model is well established. The hypothesis register is a structured document. None of it is conceptually difficult.

What is difficult — and what determines whether value assurance survives its first contact with organisational reality — is the cultural precondition that sits beneath all of it.

Value assurance only works where killing an initiative on evidence is a career-neutral act. If the response to adverse data is to question the metrics rather than the initiative, the entire apparatus is decoration.

If the portfolio board treats a recommendation to stop funding as an attack on the sponsor rather than a service to the portfolio, the value assurance function will learn to soften its findings. Within two reporting cycles it will be producing the same consensual fictions that benefits realisation produced, only with better data visualisation.

The pattern I have observed is that this cultural shift does not arrive through exhortation. It arrives through the first hard decision — the first time the investment board actually stops an initiative on the strength of a value assurance assessment, and the organisation discovers that the sky does not fall. That precedent is worth more than any amount of framework documentation. The second time is easier. By the third, the expectation has shifted: sponsors begin to build genuine value cases because they know the hypotheses will be tested, and the quality of the portfolio's investment decisions improves not because anyone became more virtuous, but because the institutional machinery changed what was rational.

6. What This Is Not

Value assurance is not benefits realisation with a new name and a dashboard. It is not a PMO function with upgraded reporting. It is not something that can be delivered by asking programme managers to track their own benefits more carefully.

It is a portfolio-level discipline with independent ownership, operational measurement infrastructure, and a direct connection to funding decisions. We do not ask business units to self-report their financial position and trust the result. We built an entire profession — with institutional independence and statutory authority — to assure the truth of financial claims. Value assurance applies the same logic to investment claims: independent verification, operational data, and consequences for what the evidence shows.

The organisations that build this machinery will make better investment decisions — not occasionally, but structurally, because the architecture of their decision-making will have changed. The organisations that rename their benefits spreadsheet and call it value assurance will continue to wonder why their portfolios underperform. The discipline is available. The question is whether the institution is willing to use it honestly.