


BEYOND THE PMO

Why PMOs Die and What Should Replace Them — A Twenty-Year Post-Mortem

A Practitioner's Perspective on the Recurring Failure of Programme
Management Offices

Giovanni Leonardi

January 2025



Contents

1. Executive Summary	3
2. The First Wave: Governance by Volume (2003–2008)	3
3. The Second Wave: The Agile Challenge (2009–2015)	4
4. The Third Wave: The Transformation Office (2016–2021)	5
5. The Fourth Wave: Disruption and the Present Moment (2022–2025)	6
6. The Root Cause: A Fundamental Misunderstanding of Purpose	6
7. What Should Replace Them	7
7.1 Decision rights as the foundation	7
7.2 Practitioner credibility	7
7.3 Reporting as a by-product, not the purpose	8
7.4 Value measured by outcomes	8
8. The Choice Ahead	8

Why PMOs Die and What Should Replace Them — A Twenty-Year Post-Mortem

Programme · No. 234

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

“The organisations that sustain their transformation offices are the ones that protect the function’s decision-making remit with the same ferocity they would protect a revenue-generating capability.”

1. Executive Summary

The Programme Management Office has endured one of the most predictable lifecycles in organisational design: established with ambition, loaded with reporting responsibilities, hollowed of decision-making authority, and eventually disbanded when its overhead becomes indefensible. This cycle has repeated across industries for more than twenty years. Each iteration follows a remarkably similar arc — a new PMO is created in response to delivery failure, it accumulates process and reporting obligations, it loses touch with the delivery teams it was meant to support, and it is closed in a cost-reduction exercise, only for the underlying coordination failure to resurface and prompt the creation of another one. This essay traces the pattern across four distinct waves of PMO practice from the early 2000s to the present, examines the structural forces that drive the drift from decision-making to reporting, and argues that the solution is not a better PMO but a fundamentally different construct: a delivery decision function whose value is measured by the quality of decisions it enables, not the volume of status reports it produces.

2. The First Wave: Governance by Volume (2003–2008)

The early 2000s saw the rapid proliferation of PMOs across large organisations, driven by a convergence of pressures. Major programme failures — visible, expensive, and politically damaging — created an appetite for centralised oversight. The maturing of project management as a discipline, codified through frameworks such as PRINCE2 and the PMI’s PMBOK, provided the intellectual scaffolding. And the post-Enron regulatory environment created a broader demand for demonstrable governance and audit trails.

The PMOs of this era were, in most cases, built as reporting functions from the outset. Their founding mandate typically included maintaining programme and project registers, standardising status reporting, consolidating RAG dashboards for senior leadership, and ensuring methodological compliance. In the best cases, they also provided assurance —

independent reviews of programme health, challenge to delivery confidence, and early warning of trouble.

In practice, however, the reporting function consumed the oxygen. PMO teams spent the majority of their time collecting, consolidating, and formatting status information from delivery teams. The assurance function, which required experienced practitioners capable of genuine challenge, was the first to be diluted as headcount pressures grew. What remained was a function that could tell the board what colour each programme was, but could not tell them — with any authority — whether that colour was accurate, whether the underlying risks were being managed, or whether the programme's approach was sound.

The pattern I have observed repeatedly is that the PMO became a post-box: information flowed through it, but it added no analytical value to the flow. Delivery teams saw it as an administrative burden — another set of templates to fill in, another reporting cycle to feed — and senior leaders saw it as a comfort blanket that was, on closer inspection, rather thin. By the end of this period, many organisations had already begun their first cycle of PMO disbandment, citing cost, bureaucracy, and a lack of demonstrable value.

3. The Second Wave: The Agile Challenge (2009–2015)

The global financial crisis of 2008 accelerated a shift in delivery philosophy that had been building for several years. Agile methods, initially confined to software development, began to expand into broader programme and change delivery. The Agile Manifesto's emphasis on working software over comprehensive documentation, and on responding to change over following a plan, posed a direct challenge to the PMO model.

The tension was structural, not merely cultural. PMOs were built around the assumption that delivery follows a predictable sequence — initiation, planning, execution, closure — and that governance operates through stage gates and milestone reviews. Agile delivery operates on fundamentally different assumptions: work is organised in short iterations, scope is managed through prioritised backlogs rather than fixed requirements, and progress is demonstrated through working increments rather than status reports.

The organisations that navigated this tension most effectively were those that recognised the PMO's core problem was not its incompatibility with agile, but its pre-existing drift into reporting. Agile did not break the PMO — it exposed what was already broken.

Many organisations responded by attempting to create "agile PMOs" — lighter-touch functions that maintained portfolio visibility without imposing waterfall stage gates on iterative delivery. The results were mixed. In some cases, the agile PMO genuinely adapted, providing portfolio-level coordination while respecting delivery-level autonomy. More commonly, the rebrand was cosmetic: the same reporting demands were maintained, but delivery teams were now expected to translate their sprint-level progress into the PMO's milestone-based language, creating an additional translation burden without reducing the reporting overhead.

The period also saw the emergence of portfolio management as a discipline distinct from programme management. The argument — compelling in theory — was that the PMO's real value lay not in overseeing individual programmes but in managing the portfolio: prioritising investments, balancing capacity against demand, and ensuring strategic alignment. In practice, portfolio management requires authority over resource allocation and investment decisions, and PMOs were rarely granted that authority. The portfolio management aspiration thus became another layer of reporting rather than a genuine decision-making function.

4. The Third Wave: The Transformation Office (2016–2021)

By the mid-2010s, a new label had emerged: the Transformation Management Office, or TMO. Sometimes it appeared as an Enterprise PMO, a Strategy Execution Office, or a Change Delivery Office. The rebranding was significant — it signalled an attempt to reposition the function at the strategic level, connected to the organisation's transformation ambitions rather than to individual programme delivery.

The TMO model was, in its best incarnations, a genuine step forward. It placed the function closer to the executive team, linked delivery tracking to strategic outcomes, and — in some cases — gave the office a seat at the table when investment and prioritisation decisions were made. Several organisations in financial services, energy, and government adopted TMO models that provided meaningful portfolio-level visibility and genuine challenge to delivery confidence.

The pattern that recurs, however, is that the TMO reproduced the same lifecycle at a higher altitude. The initial mandate was strategic; the steady-state reality was operational reporting. Within two to three years of establishment, most TMOs had accumulated the same burden of status consolidation, RAG dashboard production, and methodological compliance that had defined their PMO predecessors. The strategic intent was diluted by the gravitational pull of the organisation's reporting appetite.

“The organisations that sustain their transformation offices are the ones that protect the function's decision-making remit with the same ferocity they would protect a revenue-generating capability.”

What made this wave particularly instructive was the scale of the transformation agendas it was meant to support. Digital transformation, regulatory change programmes, and large-scale operating model redesigns demanded coordination across multiple interdependent workstreams — precisely the challenge that a well-designed oversight function should address. Yet the TMO, in most cases, was asked to report on these interdependencies rather than to manage them. The decision rights remained with the programme sponsors and delivery leads, and the TMO was left to observe and record.

5. The Fourth Wave: Disruption and the Present Moment (2022–2025)

The disruption of 2020–2021 — the sudden shift to remote working, the acceleration of digital delivery, and the compression of change timescales — created a stress test for every organisational function, including the PMO and its successors. The results were revealing.

Organisations that had invested in genuine delivery decision-making capabilities — whether they called them PMOs, TMOs, or something else — found that they could adapt. The function provided coordination across dispersed teams, maintained portfolio-level visibility in conditions of extreme uncertainty, and supported rapid reprioritisation as business conditions changed week by week. Organisations whose oversight functions were primarily reporting mechanisms found that the reports continued to arrive, but nobody was reading them. The crisis demanded decisions, not dashboards.

The post-disruption period has seen a further evolution. Delivery has become permanently more distributed, more digital, and more iterative. The volume of change that large organisations are managing simultaneously has increased, driven by regulatory pressure, technology modernisation, and competitive necessity. The interdependencies between programmes have become more complex. The need for a coordination and decision-support function is, if anything, greater than it was twenty years ago.

And yet the cycle continues. Organisations that disbanded their PMOs and TMOs during cost-reduction exercises in 2022 and 2023 are now discovering — as their predecessors did in 2005, 2010, and 2017 — that the coordination gap is real and painful. New functions are being established, with new names and new mandates, and the pattern is beginning again. The question that two decades of evidence now demands is whether this iteration can be different.

6. The Root Cause: A Fundamental Misunderstanding of Purpose

The recurring death and rebirth of the PMO is not a failure of execution — it is a failure of conception. Organisations keep building reporting functions when what they need is a decision-making capability.

The distinction is fundamental. A reporting function collects information, consolidates it, and presents it. Its value is measured by the completeness and timeliness of the information it provides. A decision-making capability uses information to drive specific decisions — which programmes to accelerate, which to pause, where to reallocate resources, how to manage interdependencies, when to escalate and to whom. The information is an input, not the output.

The reason PMOs default to reporting is structural, not accidental:

- The authority gap. PMOs are typically established without the decision rights they need to add value. They can observe and report, but they cannot direct. This gap is usually deliberate — senior leaders want visibility without surrendering control — but it ensures the function becomes a spectator rather than a participant in the delivery decisions that matter.

- The reporting appetite. Large organisations have an almost unlimited appetite for information. Every new stakeholder, every new governance forum, and every new regulatory requirement generates a demand for another report, another dashboard, another data point. The PMO, as the function closest to delivery data, becomes the default provider. Over time, the reporting demand crowds out every other activity.
- The staffing model. PMOs are frequently staffed with junior analysts and coordinators rather than experienced delivery practitioners. This is partly a cost decision and partly a reflection of the reporting-centric mandate — collecting and formatting data does not require deep programme management experience. But it ensures the function lacks the credibility and capability to provide genuine challenge or to participate meaningfully in delivery decisions.
- The measurement problem. Reporting is easy to measure: reports produced, dashboards updated, data collected on time. Decision quality is hard to measure: how do you attribute a good portfolio outcome to the quality of the decisions made along the way? This asymmetry means PMOs are evaluated on what is measurable rather than what is valuable, and they optimise accordingly.

These four forces are not independent — they reinforce each other in a cycle that is remarkably consistent across organisations, sectors, and geographies. The authority gap ensures the PMO can only report; the reporting appetite fills its capacity with reporting work; the staffing model provides people suited to reporting rather than decision-making; and the measurement problem rewards reporting activity. Breaking the cycle requires intervening on all four simultaneously.

7. What Should Replace Them

The argument of this essay is not that organisations do not need a function that coordinates delivery across programmes, manages portfolio-level decisions, and provides independent assurance. They do — and the recurring re-establishment of PMOs after each disbandment proves the point. The argument is that the function must be designed differently from the ground up.

The replacement — whether called a Delivery Decision Office, a Portfolio Decision Function, or whatever label avoids the accumulated baggage — must be built on four principles:

7.1 Decision rights as the foundation

The function must be established with explicit decision rights over portfolio prioritisation, resource allocation across programmes, interdependency management, and escalation resolution. These rights must be sponsored at the executive level and defended when challenged. A function without decision rights will inevitably default to reporting, because reporting is the only activity available to a function with no authority.

7.2 Practitioner credibility

The function must be staffed with experienced delivery practitioners — people who have run programmes, managed complex dependencies, and navigated the organisational

politics that determine whether a programme succeeds or fails. They must be capable of providing genuine challenge to delivery teams, not because they have read the status report, but because they understand the delivery dynamics that the status report may be obscuring. This means higher-calibre people, which means higher cost — and organisations must accept that a function capable of genuine decision-making is more expensive than a reporting function, but vastly more valuable.

7.3 Reporting as a by-product, not the purpose

Information is essential, but it is an input to decisions, not an end in itself. The function's reporting should be minimal, decision-oriented, and largely automated. The question every report must answer is not "what is the status?" but "what decision does this require?" If no decision is required, the report should not exist.

7.4 Value measured by outcomes

The function's performance must be measured by the quality and timeliness of the decisions it enables: decision cycle times, the proportion of portfolio resources allocated to strategic priorities, the speed of reallocation when priorities change, the accuracy of delivery forecasts, and the rate at which interdependency risks are identified and resolved before they materialise. These metrics are harder to collect than report-production statistics, but they measure what matters.

The question is not whether organisations need portfolio-level coordination and decision-making — twenty years of evidence confirms that they do. The question is whether they are willing to build a function with the authority, the capability, and the accountability to provide it, and to protect that function from the gravitational pull of the reporting machine.

8. The Choice Ahead

The pattern is not unique to any sector. I have seen it in financial services, in government, in energy, in telecommunications, and in healthcare. The labels differ — PMO, EPMO, TMO, Strategy Execution Office — but the lifecycle is the same. Establishment, drift, disbandment, re-establishment. The consistency of the pattern across such different contexts is itself evidence that the root cause is structural rather than situational.

Twenty years of PMO practice have produced a body of evidence that is difficult to ignore. The coordination need is real, persistent, and growing. The reporting-centric model fails predictably. The cycle of establishment, drift, and disbandment wastes money, erodes trust, and leaves the underlying coordination gap unfilled.

The organisations that break the cycle will be those that make a deliberate choice: to build a function with genuine authority over delivery decisions, to staff it with people who command the respect of delivery teams, to protect its decision-making mandate from the encroachment of reporting demands, and to measure its value by outcomes rather than activity.

In my experience, the PMO does not die because it is unnecessary. It dies because it is built to do the wrong thing. The need it was meant to address — coordination, prioritisation, decision-making across a complex portfolio — is more acute now than at any point in the past two decades. What must change is not the need but the response to it. The next PMO need not repeat the pattern — but only if the organisations that build it are willing to learn from two decades of evidence about why its predecessors failed.