

BEYOND THE STATUS REPORT

The Rise of the PMO — Reporting Function or Decision Engine?

A Practitioner's Perspective on What the Programme Office Is Actually For



Giovanni Leonardi

February 2002

Contents

1. Executive Summary	3
2. The Disorder the PMO Was Born to Cure	3
3. Two Identities	4
4. Why Gravity Pulls Toward Reporting	5
5. What a Decision Engine Actually Does	5
6. Whose Choice It Is	6
7. Where This Leaves the Programme Office	7

The Rise of the PMO – Reporting Function or Decision Engine?

Programme · No. 7

Originally written February 2002 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“A programme office that only reports is a very expensive way of finding out, slightly late, what you could not change.”

1. Executive Summary

The programme management office has become one of the fastest-spreading fixtures of large organisations. A decade ago few outside government or defence had heard the term; today it is difficult to find a major change initiative that does not have a PMO attached to it, or an organisation that is not standing one up. The chastening of the past two years – the collapse of so many over-ambitious technology investments as the boom unwound, and now a deepening crisis of confidence in how large enterprises report on themselves – has only accelerated the trend. When money is tight and trust is scarce, the instinct to impose order on programmes through a dedicated office is overwhelming.

Yet as the PMO has spread, a quiet disappointment has spread with it. Sponsors who established a programme office expecting it to give them control over their programmes find, a year on, that they have gained something less: a well-run reporting function that tells them, accurately and on time, what has already happened, while the programme continues to go wherever it was going to go. The office produces; it does not steer. It has become a mirror held up to the programme rather than a hand on the wheel.

This essay is about that fork in the road. Every programme office faces a choice between two identities – to be a reporting function that observes and documents, or a decision engine that shapes what the programme does. Most, I will argue, drift into the first without ever consciously choosing it, pulled there by forces that are worth understanding because they are structural rather than accidental. The essay examines why the gravity runs toward reporting, what a decision engine actually looks like in practice, and why the choice between them is not the PMO's to make alone but the sponsor's. It is written in the conviction that the difference between the two is the difference between a programme office that earns its cost and one that merely adds to it.

2. The Disorder the PMO Was Born to Cure

To understand what a programme office is for, it helps to remember the disorder that calls it into being. A large programme is, before anything is done to it, a scene of structured chaos. A dozen projects run in parallel, each with its own plan, its own view of the truth, its own dependencies on the others that it half-understands. No single person can hold the whole picture. Decisions are taken locally that make sense within one project and quietly wreck another. Risks are managed in a dozen registers that never speak to one another. The sponsor, sitting above all this, is presented with a different story by every project manager and has no means of reconciling them.

Into this the programme office arrives, promising coherence. It will impose a common language and a common plan. It will consolidate the dozen views into one. It will make the dependencies visible, the risks comparable, the reporting consistent. And in its early months it usually delivers exactly this, which is why sponsors value it and why the model has spread so fast. The first thing a good PMO does is to make the programme legible — to turn the chaos into something that can be seen whole. This is a genuine achievement and not to be dismissed.

But legibility is where the trouble starts, because legibility is so satisfying, and so measurable, that it is easily mistaken for the whole job. Once the programme can be seen clearly, the PMO has an obvious and endlessly renewable task: to keep seeing it clearly, week after week, in ever finer detail. And that task — the production of an accurate, consistent, up-to-date picture — will happily consume every hour the office has, and every hour it will ever be given. The reporting function is not a corruption of the PMO's purpose. It is the PMO's first success, over-extended until it crowds out everything else.

3. Two Identities

Let me draw the distinction as sharply as I can, because the whole argument rests on it.

	Reporting function	Decision engine
Core question	What has happened?	What should we do?
Output	An accurate picture of the past	A better decision about the future
Relationship to the programme	Observes it	Shapes it
Measure of success	Is the reporting timely and correct?	Did the decision change, and for the better?
What it does with a problem	Records and escalates it	Frames the choice and drives it to resolution

A reporting function answers the question what has happened? Its product is an accurate picture of the past, delivered up the chain. It is judged on whether that picture is timely, consistent and correct. It observes the programme, and observation is the limit of its ambition.

A decision engine answers a different question: what should we do? Its product is not a picture but a better decision. It takes the same raw material — the plans, the risks, the dependencies, the slippages — but it does not stop at describing them. It frames the choices they imply, marshals the analysis needed to resolve them, forces those choices in front of the people who can make them, and drives them to a conclusion that changes what the programme actually does. It is judged not on whether its reports were accurate but on whether its interventions improved the outcome.

The test is brutally simple. If your programme office vanished overnight, would the programme make worse decisions — or merely have worse reports? If only the reports would suffer, you have built a reporting function and called it a PMO.

Both identities are legitimate; an organisation may reasonably want only accurate reporting. But the two are routinely confused, and the confusion is expensive, because sponsors pay for a decision engine — that is what they think control means — and receive a reporting function, without anyone ever having decided that this is what would be built.

4. Why Gravity Pulls Toward Reporting

If the decision engine is what sponsors actually want, why do so few programme offices become one? Because three structural forces pull relentlessly toward reporting, and only a deliberate, sustained effort resists them.

The first is that reporting is safe and deciding is dangerous. A programme office that produces an accurate report has done its job and cannot easily be blamed for anything; the decisions, after all, belong to others. A programme office that inserts itself into decisions takes on risk — it can be wrong, it can tread on the authority of project managers and sponsors, it can be blamed when a choice it drove turns out badly. Faced with this asymmetry, the office that wishes to survive learns quickly that the reporting role is the comfortable one. It can always defend a correct report. It can never fully defend a contested decision.

The second force is that reporting is measurable and deciding is not. It is easy to judge whether a report was delivered on time and free of error, and so a programme office is naturally held to that standard, because that is the standard that can be held. It is very hard to judge whether the office improved a decision — the counterfactual is unknowable, the credit is diffuse, the timescale is long. What gets measured gets managed, and what gets managed gets done; the office optimises for the reporting it is measured on and lets the deciding, which no one measures, quietly lapse.

The third force is the most human: producing reports feels like work, and it fills the day. There is always another status pack to compile, another register to update, another dashboard to refine. The office that spends its time this way is visibly busy and can point to a mountain of output. The harder work of a decision engine — thinking through what a slippage really implies, framing an uncomfortable choice, sitting with a sponsor until a decision is actually taken — is slower, quieter, and produces far less that can be shown. Under the pressure of a busy programme, the office gravitates to the work that visibly fills the hours, and the reporting always will.

*“An office that is always busy producing reports has an alibi
against the accusation that it is not actually deciding
anything.”*

These three forces are not failings of particular people. They act on every programme office, everywhere, all the time. Left alone, they guarantee the drift to reporting. Only a sponsor and a PMO leader who understand them, and who consciously push the other way, will end up with a decision engine.

5. What a Decision Engine Actually Does

It is easy to call for a PMO that "drives decisions" and much harder to say concretely what that means. Let me try, because the abstraction has done enough damage.

- It converts information into a framed choice. A reporting function tells the sponsor that a project is six weeks late. A decision engine tells the sponsor that the six-week slippage forces a choice between three options — re-sequence and protect the go-live at the cost of scope, hold the scope and move the date, or add resource at a stated cost — and sets out what each implies. The raw fact is the same; the decision engine has done the work of turning it into a decision that can actually be taken.
- It owns the decision until it is made. A reporting function escalates a problem and considers its duty discharged. A decision engine treats an unmade decision as its own open item, chases it, convenes the people who must resolve it, and does not let it drift. The escalation is the beginning of its work, not the end.
- It brings analysis, not just aggregation. A reporting function consolidates what the projects report. A decision engine interrogates it — tests whether the dependency everyone is relying on is real, models what happens if the optimistic assumption fails, finds the second-order effect that no single project could see. It adds understanding the programme did not previously have.
- It protects the scarce attention of decision-makers. A reporting function sends everything upward and lets the sponsor sort it out. A decision engine decides what the sponsor must engage with and what can be resolved below, spending the sponsor's limited attention only on the choices that genuinely need it. It is a filter with judgement, not a pipe.

Notice that every one of these requires the office to take a position, to exercise judgement, and therefore to accept the risk of being wrong. That is the price of being a decision engine, and it is precisely the price a reporting function exists to avoid paying.

6. Whose Choice It Is

Here is the point I most want to land. The choice between these two identities is usually treated, if it is treated at all, as the PMO's own to make — a matter of how the office chooses to operate. This is a mistake. A programme office cannot make itself a decision engine by wishing to be one, because the thing that separates the two identities is authority, and authority can only be granted from above.

A decision engine intervenes in decisions. It can only do that if the sponsor has given it standing to do so — has made clear to the project managers and the wider programme that the office speaks with the sponsor's backing, that its framing of a choice is to be taken seriously, that its chasing of an unmade decision carries weight. Without that grant of authority, an office that tries to act as a decision engine is simply an irritant with no mandate, easily ignored, and it will retreat to reporting not from cowardice but from necessity, because reporting is the only thing it has the standing to do.

So when a sponsor complains that their programme office "only produces reports," the honest response is often to ask what authority they ever gave it to do more. The reporting-only PMO is frequently not a failure of the office but a faithful reflection of the mandate it was handed. If you want a decision engine, you must build one deliberately, and the first brick is not a process or a tool but a grant of authority that only the sponsor can lay.

7. Where This Leaves the Programme Office

The PMO will go on spreading; the conditions that drive its adoption are only intensifying, and in a climate that prizes control and accountability above almost everything, no large programme will soon be without one. The question is not whether organisations will have programme offices but what those offices will be for. And on current trends, the honest answer is that most will be reporting functions — accurate, diligent, busy, and quietly beside the point — because that is where the gravity leads and few will consciously resist it.

The ones that become something more will not get there by accident. They will get there because a sponsor understood that legibility was the beginning of the job and not the whole of it, because someone named the forces pulling toward reporting and pushed deliberately against them, and because the office was granted the authority to shape decisions and then held to account for whether it did. That is a harder thing to build than a reporting function, and a rarer one. But it is the only version of the programme office that repays what it costs, because a programme office that only reports is a very expensive way of finding out, slightly late, what you could not change.