



Portfolio-Level Resource Planning

WHITE PAPER

Allocating Scarce Program Talent Across Parallel Trials and Workstreams

Abstract

Every PMO leader running more than one active program at a time faces the same quiet crisis: the same three or four experienced program managers keep getting pulled onto whatever trial is currently on fire, while the rest of the portfolio runs thin. It rarely shows up as a formal staffing gap. It shows up as burnout, missed milestones nobody predicted, and a resourcing conversation that only happens after a program is already behind. Recent research puts a number on the problem: forty-seven percent of organizations lack real-time visibility into capacity and demand, and only nine percent fully trust the resource data they do have. On-time project delivery across professional services has fallen to just over seventy-three percent, down from over eighty percent five years ago. This white paper applies established program and portfolio management best practices to portfolio-level resource planning specifically, addressing the dimension most PMOs discover is their weakest only after a resource conflict has already cost them a milestone.

The Resourcing Crisis That Never Gets a Formal Name

Ask a PMO leader whether their organization has a resourcing problem and most will say no, at least at first. Ask whether the same handful of senior program managers seem to end up on every high-visibility program regardless of what else is on their plate, and the answer changes fast. That gap between the formal answer and the honest one is where portfolio-level resource planning actually lives, and it rarely gets addressed as its own discipline. It gets treated as a side effect of individual program management, solved one program at a time by whichever program manager shouts loudest in a resourcing meeting.

The scale of the underlying visibility problem is well documented outside life sciences specifically, and there is no reason to think pharmaceutical and biotech PMOs are meaningfully better positioned. Forty-seven percent of organizations report lacking real-time visibility into capacity and demand, and only nine percent say they fully trust the resource data they have, with sixty-three percent blaming inconsistent data inputs across systems for that distrust.¹ Average resource utilization sits around seventy-two percent, below the eighty percent benchmark most resourcing models target, while utilization that climbs past one hundred twenty-five percent reliably produces the kind of project delays that eventually surface as a missed milestone.² Professional services on-time delivery overall has fallen to just under seventy-four percent as of this year, down from over eighty percent five years ago, a decline that tracks closely with the growing complexity of resourcing decisions organizations are being asked to make with worse visibility than they had before.³

Life sciences adds its own layer of difficulty. Clinical research staffing has genuinely improved over the past two years, with clinical research associate retention climbing to roughly eighty percent in 2024 from about seventy-three percent the year before, but that recovery has been uneven.⁴ Some research sites saw research nurse turnover exceed forty percent during the pandemic era, and shortages remain concentrated in specific geographies and specialties rather than resolved evenly across the industry.⁵ A PMO planning resources across parallel trials is not drawing from an infinite, interchangeable talent pool. It is drawing from a specific, currently recovering, unevenly distributed one.

Why Program-Level Resourcing Cannot See the Real Problem

Most organizations resource at the program level because that is where the budget and the accountability sit. A program manager fights for the headcount their program needs, wins or loses that fight largely on the strength of their own escalation, and moves on. That approach works reasonably well when programs run in isolation. It breaks down the moment two programs need the same senior biostatistician in the same quarter, because nothing in a program-level resourcing process is designed to see that conflict coming.

The failure mode is predictable. A skilled, senior program manager becomes the default answer to every difficult assignment, not because a portfolio-level process deliberately chose to concentrate risk on that person, but because program-level resourcing has no mechanism to notice it is doing so. That person's calendar becomes the true bottleneck constraining how many programs the organization can actually run well at once, and nobody discovers this until the person leaves, burns out, or simply cannot be in two critical meetings at the same time.

Portfolio-level resource planning exists to make that pattern visible before it becomes a crisis, not to strip individual program managers of authority over their own staffing. The goal is a shared, accurate picture of where



capacity is genuinely available across the whole portfolio, so a resourcing decision on one program can be made with honest knowledge of what it costs every other program sharing the same talent pool.

Applying Best Practices: A Model for Portfolio Resource Planning

A structured PMO maturity model treats technology and data enablement, and talent and change leadership, as two of the five dimensions separating a tactical PMO from a strategic one. Portfolio-level resource planning sits squarely at the intersection of those two dimensions, and it is worth building out a specific, practical model for this discipline rather than treating resourcing as a generic governance topic. Four elements make up that model.

1. **One Capacity Model, Not One Per Program.** A shared view of every program's resource demand and every individual's availability, built on the same data model across the portfolio, rather than each program manager maintaining a private spreadsheet with its own assumptions about who is available when.
2. **Named-Resource Risk Scoring.** Explicit identification of which individuals represent a portfolio-level single point of failure, the senior specialists whose unavailability would stall more than one active program, and a deliberate plan to build depth behind each of them rather than discovering the gap when they leave.
3. **Forward-Looking Demand Signals.** Resourcing decisions made against a rolling forecast of upcoming program milestones, not just current-quarter staffing requests, so a resourcing conflict six months out gets addressed with a hiring or training decision now instead of a crisis reallocation later.
4. **A Portfolio-Level Arbitration Path.** A defined, transparent process for resolving genuine resource conflicts between programs of differing priority, so the outcome depends on portfolio strategy rather than which program manager has the most political capital in a given quarter.

None of this replaces program-level resourcing authority. It gives program managers something they currently lack: an honest, shared picture of the constraint they are all actually competing for, so the decisions made about where to place a scarce senior specialist reflect the whole portfolio's priorities rather than whichever program escalated most recently.

What Portfolio Resourcing Buys, and Where It Gets Hard

Advantages

- **Fewer avoidable crises:** Resource conflicts get surfaced and resolved while they are still a planning conversation, rather than discovered when a program manager reports a milestone slip they saw coming for weeks but had no forum to escalate.
- **Protection for key talent:** Explicitly naming the individuals a portfolio is over-relying on creates the visibility needed to build depth behind them, rather than letting burnout or attrition make the decision by default.
- **Better-informed portfolio strategy:** Leadership making prioritization calls with an honest resourcing picture can make real trade-offs instead of assuming, incorrectly, that every priority program can be fully staffed at once.
- **A defensible answer for the board:** When a program does slip for resourcing reasons, a portfolio-level process can show the trade-off that was consciously made, rather than leaving leadership to explain an outcome nobody saw coming.



Risks and Barriers

- **Data quality dependency:** A shared capacity model is only as good as the inputs feeding it, and the sixty-three percent of organizations citing inconsistent data inputs as their core resourcing problem is a reminder that this is usually a data discipline problem before it is a planning problem.¹
 - **Program manager resistance:** Program managers who have learned to protect their own resourcing through informal escalation may see a transparent portfolio process as a loss of leverage rather than a fairer system, and that resistance needs to be managed directly, not assumed away.
 - **Over-engineering the model:** A capacity model built with more granularity than the organization can actually maintain becomes stale within a quarter and gets abandoned. The right level of detail is the level the organization will actually keep current.
 - **Arbitration without real authority:** A conflict-resolution process only works if it has genuine executive backing to make an unpopular call. A portfolio arbitration path with no teeth just adds a meeting without changing the outcome.
-

Case in Point: The Biostatistician Bottleneck

A mid-size biotech running four concurrent Phase II programs relied on a single senior biostatistician for the adaptive design work that gave the company's trials their competitive edge. Each program manager knew this individually. None of them, nor the PMO, had a shared view of how many programs were simultaneously depending on that one person's calendar in a given quarter. When two programs needed the same statistician for overlapping interim analyses, the conflict surfaced only when both program managers separately escalated to the same portfolio lead within the same week, each unaware the other was making the identical request.

The immediate fix, reprioritizing one program's analysis timeline by three weeks, was manageable. The more useful outcome was what the PMO did afterward: it built a named-resource risk register identifying every individual whose unavailability could stall more than one active program, and started a deliberate cross-training plan to build a second qualified adaptive-design statistician over the following year. The near-miss became the forcing function for exactly the kind of forward-looking capacity planning that would have caught the conflict months earlier if it had already existed.

Recommendations for PMO and Portfolio Leaders

1. **Build one shared capacity model across the portfolio.** Retire the program-by-program spreadsheet approach in favor of a single source of truth for who is available, on what programs, and when.
2. **Identify named-resource risk explicitly.** Name the specific individuals whose unavailability would stall more than one program, and treat closing that gap as a standing portfolio priority, not an occasional talent-development nice-to-have.
3. **Resource against a rolling forward forecast.** Plan against milestones six to twelve months out, not just current requests, so hiring and training decisions happen ahead of the conflict instead of in reaction to it.



4. **Give arbitration real teeth.** Ensure the executive sponsoring the portfolio resourcing process has genuine authority to make an unpopular reallocation call, and that program managers see it exercised, not just documented.
 5. **Fix the data before scaling the model.** Address inconsistent resource data inputs before investing in more sophisticated planning tools built on top of them. Sophisticated planning built on bad data is worse than simple planning built on good data.
 6. **Treat near-misses as planning inputs.** Every resourcing conflict that gets resolved at the last minute is a preview of a future crisis. Capture it, and use it to update the capacity model rather than moving on once the immediate fire is out.
-

Measuring Success: KPIs for Portfolio Resource Planning

- **Resource conflict lead time:** how far in advance a resourcing conflict is identified relative to when it actually affects a program milestone, tracked to confirm the shift from reactive to forward-looking planning.
 - **Named-resource concentration:** the number of active programs depending on any single individual, tracked as a direct measure of portfolio risk concentration.
 - **Resource data trust score:** a periodic survey of program managers on whether they trust the shared capacity model's data, addressing the trust gap the broader industry data suggests is common.¹
 - **Utilization variance from target:** how far individual and team utilization sits from the organization's target range, flagging both under- and over-allocation before either produces a downstream problem.
 - **Arbitration resolution time:** how long a genuine cross-program resource conflict takes to resolve once escalated, and whether the resolution reflected documented portfolio priority.
-

Conclusion

Every PMO running more than one active program already has a resourcing strategy, whether or not anyone would call it that. Most of the time, that strategy is simply whoever escalates loudest gets the help, and whichever senior specialist is most reliable gets asked to do the most. That is not a strategy so much as an accumulation of individually reasonable decisions with no one checking their sum, and it is exactly the pattern that turns a trusted, high-performing program manager into the single point of failure an entire portfolio did not know it had.

Applying disciplined best practices to a deliberate portfolio resource planning practice does not require a large technology investment or a wholesale reorganization. It requires a shared, honest capacity model, explicit attention to where risk is concentrated in specific people, and a resourcing conversation that happens before the conflict rather than during the escalation. Portfolios that make that shift stop discovering their resourcing gaps in a crisis and start finding them in a routine planning review, which is exactly where a gap like this is cheapest to close.

About OmniTek & Contact Information

OmniTek Consulting is a trusted partner for organizations seeking to enhance governance, operational agility, and decision-making. We specialize in modernizing business processes with a blend of human-centered design, technology expertise, and hands-on change management.



For PMO leaders managing scarce talent across parallel trials and workstreams, OmniTek applies proven best practices to build a practical portfolio resource planning model: a shared capacity view, explicit named-resource risk management, and a forward-looking forecast that catches conflicts while they are still a planning conversation. Whether your organization is trying to protect a handful of irreplaceable specialists or simply wants better visibility before the next resourcing fire drill, OmniTek can help you build a portfolio resourcing practice your program managers can actually trust.

Contact us today to explore how OmniTek can support your transformation journey and drive lasting results.



McLean, VA | Serving
clients nationwide



info@omnitekconsulting.com



omnitekconsulting.com

[Click here to book a 15-minute call to see how OmniTek can help your team put these insights into practice.](#)

Endnotes

ⁱ Runn, "Resource Management Statistics: What the Numbers Reveal," citing the 2025 SPI Professional Services Maturity Benchmark, 2026, available at <https://www.runn.io/blog/resource-management-statistics>

ⁱⁱ PharmaVoice, "A pandemic-era staffing crunch in clinical trials is easing, but trouble spots remain," 2025, available at <https://www.pharmavoice.com/news/clinical-trials-staffing-crunch-hiring-cra/708614/>

