


BEYOND THE FUNGIBLE RESOURCE

The Great Resignation and Programme Delivery Capability

A Practitioner's Perspective on What Leaves When People Leave, and How Programmes Survive It

Giovanni Leonardi

July 2021



Contents

1. Executive Summary	3
2. The Capability We Assumed Was Fixed	3
3. Why the Exits Landed So Hard	4
4. The Illusion of the Fungible Resource	5
5. What Actually Walks Out the Door	5
6. Designing for Departure	6
7. The Intent–Reality Gap	7

The Great Resignation and Programme Delivery Capability

Programme · No. 237

Originally written July 2021 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“A programme does not lose a headcount when a key person leaves; it loses a set of decisions it no longer remembers making.”

1. Executive Summary

The wave of departures now moving through knowledge organisations — the phenomenon a few commentators have begun calling the great resignation — is usually discussed as a human resources problem: a matter of retention, reward, and engagement. Viewed from the delivery side of a major programme, it is something more specific and more alarming. It is a stress test of an assumption most programmes were built on without ever examining it, namely that delivery capability is a property of the organisation rather than of the particular people who happen to be doing the work this quarter.

This essay argues that the current turnover is exposing how much of our programme delivery capability was never institutional at all. It lived in individuals — in the delivery lead who knew why a decision was taken two years ago, in the analyst who understood which of the requirements were real, in the architect who carried the shape of the system in their head. When those people leave in ones and twos, a programme does not lose a headcount it can backfill. It loses continuity of judgement, and that is far harder to replace than a pair of hands. The organisations that come through this well will be the ones that stop treating delivery capability as a resourcing line and start treating it as something that must be deliberately built to outlast the people who carry it.

We planned our programmes as though people were interchangeable and knowledge was written down. This year is teaching us that neither was true.

2. The Capability We Assumed Was Fixed

For as long as I have watched programmes staffed and resourced, the underlying model has been remarkably crude. Capability is expressed as a number of people of a given grade and skill, arranged against a plan. A programme needs so many delivery managers, so many analysts, so many developers; the resourcing function supplies them; and when

one leaves, another of the same grade is found to take their place. The model treats people as units of a standard type, and it treats the standard type as the thing that matters. It has the great virtue of being easy to plan with, and the great flaw of being false.

Under ordinary conditions the flaw stays hidden, because turnover is low enough and gradual enough that the tacit knowledge held by individuals transfers informally before it is lost. Someone leaves, but they have spent months alongside the person who will inherit their work, and the important things pass across in corridors and handovers without anyone noticing that a transfer of genuine value has taken place. The model survives not because it is right but because reality has been forgiving.

This year reality has stopped being forgiving. When departures rise and accelerate at the same time, the informal transfer mechanism breaks. People leave faster than their knowledge can migrate. They leave in numbers that overwhelm the corridors and handovers that used to carry the load. And crucially they leave from every part of the programme at once, so there is no stable core to absorb the loss. The resourcing model says the gaps can be filled; the programme discovers that the new people, however capable, arrive without the one thing that was actually keeping delivery on the road — the accumulated context of why things are the way they are.

3. Why the Exits Landed So Hard

It would be comforting to treat the current turnover as a passing squall, a consequence of a strange year that will settle once the world does. I think that reading is complacent. Several forces are pushing people to move, and most of them are structural rather than temporary. Understanding them matters, because a programme that assumes the pressure will subside will under-invest in exactly the defences it most needs.

- The past year revealed how portable the work had become. Once delivery could be done from anywhere, the local employer lost its geographic monopoly on a person's labour. The talent market widened for everyone at once, and a programme that used to compete for people within a city now competes with every organisation prepared to hire remotely.
- A period of enforced stability released a backlog of moves. Many people who would have changed roles in the normal course of things stayed put through the most uncertain months. That deferred motion has not been cancelled; it has been postponed and is now arriving all at once.
- The experience of the crisis changed what people want from work. Having proved to themselves that the old arrangements were not immutable, many are unwilling to return to them, and will move to find the flexibility, the purpose, or the conditions they now consider non-negotiable.
- Demand for delivery talent is rising precisely as supply becomes restless. Organisations are relaunching the programmes they paused, all reaching for the same scarce delivery skills at the same moment, which bids up the market and hands the initiative to anyone willing to move.

These forces do not cancel each other out and they will not resolve on a convenient timetable. The prudent assumption for anyone accountable for a programme is that elevated turnover is the new operating condition, not a temporary disturbance to be

waited out. That assumption changes the design of the programme, and it is to that design that the rest of this essay turns.

4. The Illusion of the Fungible Resource

At the heart of the resourcing model sits a single seductive idea: that a person of a given grade and skill can be swapped for another of the same grade and skill with only a short loss of productivity while they get up to speed. This is the illusion of the fungible resource, and it is worth being precise about why it fails, because the failure is not obvious and the remedy depends on understanding it.

What a new joiner can acquire quickly is the explicit content of the role — the technologies, the processes, the documented decisions, the formal structure of the work. What they cannot acquire quickly, and often cannot acquire at all without the person who left, is the tacit layer beneath it: which of the documented requirements actually matter and which are dead letters, why a particular design path was chosen over the obvious alternative, which stakeholder's stated position differs from their real one, where the bodies are buried in the legacy system nobody wants to touch. This is the knowledge that makes delivery efficient rather than merely possible, and almost none of it is written down anywhere a successor could find it.

“The organisation chart shows the roles. It is silent on the judgement, and the judgement was doing most of the work.”

The consequence is that the true cost of a departure is systematically underestimated. The resourcing model sees a gap of a few weeks while a replacement onboard. The programme experiences something quite different: a prolonged period in which decisions that used to take an afternoon take a fortnight, in which old ground is re-litigated because no one remembers why it was settled, in which mistakes are made that the departed person would have caught without thinking. The number on the resourcing plan is unchanged. The capability of the programme has fallen, and it will not recover until the new person has, slowly and expensively, rebuilt the context their predecessor took with them.

5. What Actually Walks Out the Door

If we are to defend against this loss, we have to be honest about what is actually being lost, because it is not what the plan says. When a key person leaves a programme, three distinct things depart with them, and each requires a different response.

The first is decision memory: the record, held almost entirely in people's heads, of what was decided, when, and above all why. A programme runs on a long chain of choices, most

of them reasonable only in light of the circumstances at the time. Strip out the person who remembers the circumstances and the choices look arbitrary, inviting the successor to reopen them and the programme to relive its own history.

The second is relationship capital: the trust and the working understanding built up with stakeholders, suppliers, and other teams over months of contact. A new person inherits the contact list but not the credibility, and much of the friction that follows a departure is simply the cost of rebuilding trust that the organisation had already paid for once.

The third is system intuition: the hard-won feel for how the thing being built actually behaves — where it is fragile, where it can bear weight, what will break if touched. This is the knowledge that separates a confident change from a reckless one, and it is acquired only by living with the system through its failures. A newcomer has read the documentation; they have not been woken at three in the morning by the part of it that the documentation gets wrong.

What the plan tracks	What actually leaves	Why it is hard to replace
A headcount by grade	Decision memory	Reasons were never recorded, only the outcomes
A named contact	Relationship capital	Trust is earned in time, not transferred on a list
A skill set	System intuition	Learned through failure, not documentation

Naming these three losses precisely is not an academic exercise. It tells us what to build. A programme that wants to survive turnover must find ways to hold decision memory, relationship capital, and system intuition at the level of the team rather than the individual — to make them institutional before the individual leaves, not scramble to reconstruct them afterwards.

6. Designing for Departure

The conventional response to rising turnover is to fight it: to improve retention, sharpen reward, and work to keep people in their seats. That work is worth doing, but as a delivery strategy it is insufficient and slightly self-deceiving, because it assumes a level of control over people's choices that no organisation actually has. The more robust posture is to accept that people will leave and to design the programme so that departure, when it comes, costs less. In my experience this rests on a handful of deliberate moves.

1. Make knowledge capture a condition of the work, not an afterthought. The reasons behind decisions must be recorded as decisions are made, in a form a successor can actually use. This is unnatural and unpopular, and it will not happen unless it is built into how the programme runs and treated as a deliverable in its own right.
2. Refuse single points of knowledge. Wherever only one person understands something that matters, that is a risk on the programme's register, not a convenience. Pairing, rotation, and shared ownership are slower in the moment and cheaper across the life of the programme, because they spread the tacit layer across more than one head.

3. Stage relationships deliberately. Key stakeholder and supplier relationships should never rest on a single individual. Bringing a second person into the important conversations before a departure is the only reliable way to transfer trust rather than merely a contact.
4. Treat handover as a designed process, not an exit formality. The value in a departing person's head does not transfer in a two-week notice period through goodwill alone. Where a handover matters, it should be planned, resourced, and started early — ideally before anyone has resigned, as a standing discipline rather than a reaction.
5. Resource for resilience, not just for the plan. A programme staffed to exactly its planned need has no capacity to absorb a loss. A small, deliberate margin — in people, in overlap, in slack — is not inefficiency; it is the premium paid for continuity, and in a high-turnover market it is among the best value a programme can buy.

7. The Intent–Reality Gap

There is a wider point beneath all of this, and it is really a point about how we have thought about programmes for years. We have planned them as machines: define the work, break it into roles, staff the roles, and the output follows. The machine metaphor is comforting because it makes delivery look controllable, and it is wrong in exactly the way that matters now. Programmes are not machines whose parts are interchangeable. They are more like living systems that carry a great deal of their essential information informally, in the relationships and memories and intuitions of the people inside them — and living systems are wounded, not merely reconfigured, when their parts are removed.

The gap between transformation intent and transformation reality has many causes, but this year is exposing one we have long preferred to ignore. Our intent was expressed in plans that assumed capability was a fixed, institutional property we could rely on. Our reality is that much of that capability was personal, tacit, and quietly walking out of the door. The current turnover did not create this vulnerability; it revealed one that was always there, hidden by the good fortune of a stable workforce. The organisations that treat the revelation as a passing inconvenience will keep rebuilding the same fragile machine and being surprised each time it breaks. The ones that treat it as a design lesson — that build delivery capability to survive the loss of any individual who carries it — will find they have bought something more valuable than retention. They will have bought a programme that no single resignation can knock off course.