


BEYOND THE DEMO

Multi-Agent Orchestration in Enterprise Delivery — The Gap Between the Demo and the Delivery Room

A Practitioner's Perspective on Complex Agent Systems in Programme
Delivery

Giovanni Leonardi

June 2024



Multi-Agent Orchestration in Enterprise Delivery — The Gap Between the Demo and the Delivery Room

Programme · No. 261

Originally written June 2024 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“The agent can draft the report. It cannot read the room.”

The Promise on the Screen

The demonstration is always compelling. Multiple AI agents, each with a defined role, coordinate to produce an output that no single agent could achieve alone. One agent researches, another drafts, a third reviews, a fourth refines. The whole sequence completes in minutes. For anyone who has spent years managing complex delivery programmes — coordinating dozens of workstreams, managing dependencies, chasing status updates — the appeal is immediate and visceral.

The pattern I have observed in 2024, however, is that the distance between this demonstration and anything resembling enterprise delivery orchestration is vast, and the organisations now attempting to close that gap are discovering problems that the technology community has not yet seriously addressed.

Where the Analogy Breaks

The multi-agent paradigm borrows heavily from human organisational design. Agents are given roles, responsibilities, and reporting lines. The language is deliberately familiar: supervisor agents, worker agents, quality assurance agents. This familiarity is seductive and, in my experience, dangerous.

Human delivery teams function because of context that is never explicitly stated. A senior programme manager does not need to be told that the finance director's concerns about quarterly reporting will shape every steering committee decision. A lead architect does not need written instructions to recognise that the integration with the legacy payments platform is the real risk, regardless of what the risk register says. Delivery teams operate on institutional memory, political awareness, and professional judgement that has been built over years.

Multi-agent systems have none of this. Each agent operates within the boundaries of its prompt and its context window. The orchestration layer — the component that coordinates agents, manages handoffs, and resolves conflicts — is typically the thinnest

and least sophisticated part of the architecture. It is precisely the opposite of how effective delivery works, where the orchestration function — the programme manager, the delivery lead, the integration authority — is the most experienced and contextually rich role in the room.

The orchestration layer in multi-agent systems is typically the thinnest part of the architecture. In effective delivery, the orchestration function is the most experienced and contextually rich role in the room. This inversion explains most of the failures.

The Dependency Problem

Enterprise delivery is fundamentally a dependency management challenge. The value of a programme manager is not in tracking tasks — any tool can do that — but in understanding which dependencies are real, which are political, which are technical, and which are artificial constructs of organisational history. A competent delivery professional can look at a dependency map and immediately identify which constraints are genuinely immovable and which can be negotiated, escalated, or designed around.

Multi-agent systems, as currently conceived, treat dependencies as data. They can be told that Task B depends on Task A, and they will sequence accordingly. What they cannot do is recognise that the dependency between the data migration workstream and the regulatory approval process is actually a proxy for a territorial dispute between two divisions — and that the real path forward is a conversation between two specific people who have worked together before and trust each other.

This is not a limitation that will be solved by larger context windows or more sophisticated prompting. It is a category difference between information processing and situated professional judgement.

What Multi-Agent Orchestration Can Actually Do

None of this is to say that multi-agent approaches have no place in enterprise delivery. The pattern I have seen work — and it is genuinely useful — is narrower than the vision but more immediately practical.

Multi-agent orchestration is effective for well-bounded, repeatable processes where the inputs are structured, the quality criteria are explicit, and the orchestration logic can be fully specified in advance. Document generation pipelines, compliance checking workflows, test scenario generation, status report compilation — these are legitimate and valuable applications.

- Structured document production — where multiple perspectives need to be synthesised into a single output, and the synthesis rules can be articulated.
- Parallel analysis — where the same data set needs to be examined through several lenses simultaneously, and the results aggregated.

- Quality assurance chains — where one agent's output is systematically reviewed against defined criteria by another.

These are not trivial applications. They can save significant time and improve consistency. But they are automation applications, not orchestration in the sense that a programme delivery professional would recognise.

The Risk of Premature Adoption

The risk I see most clearly in 2024 is premature adoption driven by executive enthusiasm. The demonstrations are impressive. The vendors are persuasive. The pressure to show AI-driven productivity gains is intense. And so organisations are attempting to apply multi-agent orchestration to genuinely complex delivery challenges — benefits realisation, stakeholder alignment, risk response planning — where the technology is categorically unsuited.

The result is predictable: the multi-agent system produces outputs that look plausible but miss the point. A risk assessment that identifies technical risks accurately but is blind to the political and organisational dynamics that will actually determine whether the programme succeeds or fails. A status report that aggregates data correctly but cannot distinguish between a workstream that is genuinely on track and one that is reporting green while quietly heading for a cliff.

In my experience, the organisations that will extract real value from multi-agent orchestration are the ones treating it as a tool for specific, bounded tasks within a delivery framework that is still fundamentally governed by experienced human judgement. The ones reaching for it as a replacement for delivery capability are storing up problems that will become visible only when the programme hits genuine difficulty — which, in complex transformation, it always does.

The agent can draft the report. It cannot read the room.