



AI GOVERNANCE SERIES | ADVISORY PERSPECTIVE

Agentic Orchestration Is a Compelling Idea.

Here Are the Questions Your Organization Should Be Asking Before It Becomes Your Strategy.

By Greg Aldrich | Global CIO & Strategic Advisor | April 2026

UiPath recently published their [Definitive Guide to Agentic Orchestration](#). It's a serious document, and it describes a real problem.

Enterprises are drowning in disconnected automation tools. AI pilots are failing to produce ROI at scale. Agents are being deployed without coordination, creating complexity rather than reducing it. The market data they cite is credible: 81% of application leaders currently use up to six overlapping automation tools, and 80% of GenAI initiatives report no material return on investment.

Their proposed solution — a unified orchestration layer that coordinates AI agents, automation robots, and human workers across end-to-end processes — is architecturally sound. The “Agents think. Robots do. People lead.” framing is a useful abstraction. The emphasis on governance, auditability, and human-in-the-loop controls is consistent with how I think about responsible AI deployment.

I'm not here to argue with the direction. I'm here to sharpen the conversation — because the questions this document raises are more important than the answers it offers.

“Orchestration platforms don't solve transformation by themselves. They expose whether your organization is ready for it.”

The Problem Is Real. The Framing Deserves Scrutiny.

UiPath correctly identifies that the root cause of AI pilot failure isn't model quality — it's system-level fragmentation. Agents without coordination create sprawl. Tools without integration create islands. Neither adds up to transformation.

But the document makes an implicit leap that deserves examination: it frames fragmentation primarily as a tooling problem and positions a unified platform as the solution.

In my experience, fragmentation is almost always a symptom of something deeper. Organizations don't fail to scale AI because they lack an orchestration platform. They fail because:

- Processes aren't well-defined before automation is applied.
- Ownership of outcomes is unclear across functions.
- Decision rights are fragmented at the organizational level.
- Data is inconsistent, ungoverned, or not ready to support agent decision-making.

A platform is able to amplify organizational clarity. It cannot create it. And when a sophisticated orchestration layer is deployed on top of an organization that hasn't solved these foundational problems, it doesn't simplify complexity — it centralizes it.

“The risk isn't that orchestration fails. It's that it succeeds at coordinating processes that were never properly defined in the first place.”

The Questions That Matter Before You Commit

The following questions are not criticisms of UiPath's approach. They are the questions I would ask any organization considering agentic orchestration as a strategic investment — questions the document doesn't fully address, and that your leadership team should be able to answer before a platform decision is made.

On Organizational Readiness

- *Are your processes defined clearly enough to be orchestrated? Orchestration requires explicit handoff points, decision rules, and exception paths. If your processes are ambiguous when humans execute them, they will be even more ambiguous when agents execute them at scale.*
- *Who owns the orchestration layer operationally? The platform can coordinate agents, robots, and humans — but someone in your organization has to own the design, governance, and evolution of that coordination model. Is that role defined? Is it funded? Is your data ready?*
- *Agentic orchestration assumes clean, accessible, well-governed data flowing between systems. Most organizations aren't there. Orchestration doesn't solve a data readiness problem — it amplifies it.*

On Governance and Control

- *Where are the decision boundaries defined? The document describes governance conceptually — human-in-the-loop, safe AI by design, and auditability. But where specifically are the boundaries drawn? What decisions can agents make autonomously? What requires human approval? What is never delegated? These boundaries must be defined before deployment, not discovered in production.*
- *How are conflicting agent outputs resolved? In a multi-agent environment, agents will occasionally reach different conclusions or recommend incompatible actions. Who resolves that conflict? What is the escalation path? What is the governance record of that resolution?*
- *What happens when orchestration fails mid-process? The document emphasizes durable execution and resilience. But failure modes need to be explicitly designed for. What is the recovery path when a long-running process fails at step seven of twelve? How are partial states handled? How are affected parties notified?*

On Value and ROI

- *What specific cost or revenue line improves — and by how much? The real-world results cited are encouraging. But the ROI model for your organization needs to be built from your baseline, processes, and data—not from benchmark cases. What assumptions drive your ROI projection? Have those assumptions been stress-tested?*
- *What are the full costs of adoption? Orchestration introduces new expenses that are easy to undercount: platform licensing, integration effort, governance overhead, change management, and the organizational redesign required to operate a fundamentally different model of work. Have those been quantified against the projected benefits?*
- *Where has this replaced human coordination rather than introducing a new layer? The goal is operational simplification. If orchestration adds a governance and management layer on top of existing tools and processes without replacing anything, the TCO calculation changes significantly.*

On Strategic Risk

- *What lock-in are you accepting at the control layer? If your orchestration platform becomes the coordination hub for enterprise-wide processes, the switching costs become extreme over time. This is not an argument against the investment — it is an argument intended for understanding it clearly before you make it. What exit options exist? What is the dependency profile in three years?*
- *Does this reduce complexity or reorganize it? The document argues that orchestration reduces complexity. I would argue it moves complexity — from manual coordination and fragmented tools into orchestration design, business rules, and governance models. That trade may well be worth making. But ‘less complex’ and ‘differently complex’ are not the same thing, and the skills required to manage the new complexity are different from those required to manage the old.*
- *What is your vendor dependency strategy? UiPath is a credible platform with a genuine enterprise track record. Yet a strategic dependency on any single orchestration vendor at the control layer of enterprise operations is a risk worth naming explicitly. What is your diversification posture? What open standards are you relying on to preserve optionality?*

What Good Looks Like Before You Orchestrate

Organizations that will get the most from agentic orchestration share a common profile. They have already done the harder, less visible work:

- Processes are defined, documented, and owned — not just performed.
- Data governance is in place, with clear accountability for data quality.
- Decision rights are explicit — humans know what they decide, and what agents can decide.
- The organizational structure has been designed to operate a hybrid workforce, not just accommodate one.
- Failure modes have been thought through — resilience is designed in, not hoped for.

If your organization hasn't done this work, orchestration won't fix it. It will make the gaps more expensive.

The Bigger Point

UiPath is right that enterprises need a coordination layer for AI agents, automation, and human workers. The architectural direction is sound. The governance emphasis is welcome. The real-world results are real.

But the most important question in the document isn't answered by the platform. It's answered by your organization.

Google Cloud CEO Thomas Kurian said it plainly at Google Cloud Next this month: “The experimental phase is behind us, and now the real challenge begins.” He's right. The real challenge is not deploying more capable AI. It is building the operating model, governance architecture, and organizational design that make AI deployment controllable, auditable, and sustainable.

Orchestration platforms are a deployment infrastructure. They are necessary. They are not sufficient.

The organizations that will lead in the agentic era are those that treat orchestration as an organizational design challenge first and a technology selection decision second. The platform follows the operating model. It should never precede it.

“If your organization is not already process-defined, data-governed, and decision-accountable — agentic orchestration will not fix that. It will expose it.”

Questions for Your Leadership Team

If you are evaluating agentic orchestration — or if your organization is already in a pilot — these are the questions worth bringing to your next leadership conversation:

- Are we process-ready, or are we automating ambiguity at scale?
- Do we have a named owner for the orchestration layer — not the platform, but the governance model that runs on it?
- Have we quantified the full cost of adoption, including organizational change?

- Do we understand the strategic dependency we are creating, and have we designed for optionality?
- What does failure look like — and have we designed for it explicitly?

The technology is real and the capability is compelling. The questions are what separate organizations that will scale AI effectively from those that will simply scale their existing problems faster.

This article is part of an ongoing series on AI governance and organizational readiness. [Part 1: The AI Governance Tightrope](#) | [Part 2: The New Control Layer](#) | [Part 3: The Operating Model for AI Governance](#)

About the Author

Greg Aldrich is a Global CIO and Strategic Advisor with 30+ years of experience helping boards, executive teams, and C-suite stakeholders navigate IT strategy, AI governance, and digital transformation. He serves as Senior Strategy Advisor at Blue Tree Technology Group and Senior Transition Architect at SDS Consulting, and currently serves as CIO at Andrew Wommack Ministries & Charis Bible College, where he chairs both the AI Committee and IT Steering Committee. He has advised organizations across financial services, gaming, healthcare, higher education, logistics, and nonprofit sectors.

Connect: [linkedin.com/in/galdrich/](https://www.linkedin.com/in/galdrich/)