

BEYOND CONTRACT BOUNDARIES

Multi-Vendor Governance Fails When Accountability Is Divided

Why contract clarity still leaves the programme outcome ownerless



Giovanni Leonardi

July 2005

Multi-Vendor Governance Fails When Accountability Is Divided

Programme · No. 24

Originally written July 2005 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“A programme can distribute work across many contracts, but it cannot distribute ownership of the outcome.”

1. Four Green Reports and One Failed Rehearsal

At eight o'clock on a Monday morning, four suppliers submit four green reports.

The package vendor has completed configuration. The systems integrator has passed its component tests. The network provider has installed the required circuits. The data-conversion specialist has reconciled the extract totals. Each report is supported by an agreed measure, a signed deliverable and a service-level schedule.

At one o'clock the following morning, the first full rehearsal fails.

A batch of 300,000 customer records reaches the new application forty minutes late. The application therefore misses the overnight processing window. The network has remained within its availability target; the conversion file was delivered at the time stated in that supplier's plan; the application performed to its tested capacity; and the integrator can show that the interface specification was approved. No supplier has plainly failed its contract. The programme, however, cannot open for business.

This is not an unusual contradiction in the present wave of outsourcing and multi-sourcing. Organisations have become increasingly skilled at dividing a large programme into purchasable parts. They are far less skilled at preserving ownership of the whole after the division has been made.

The deeper problem is not that suppliers refuse to collaborate, although sometimes they do. It is that the governance architecture usually follows the contract architecture. Work is divided, measures are divided, forums are divided and accountability is divided. The outcome remains stubbornly indivisible.

A programme can distribute work across many contracts, but it cannot distribute ownership of the outcome.

2. The Contract Became the Map

The attraction of multi-vendor delivery is understandable. A single prime contractor can become expensive, slow to challenge and difficult to replace. Specialist firms often bring stronger capability in particular domains. Competition can improve price discipline. Separating application, infrastructure, conversion and business-change work may also prevent one supplier from marking its own homework.

These are serious advantages. The strongest case for multi-sourcing is not fashionable enthusiasm for choice; it is a legitimate refusal to exchange delivery risk for dependency on one commercial party.

Yet the very act of procurement changes how people see the programme. Once the work is expressed through statements of work, acceptance criteria and liability clauses, the contract boundary begins to look like the natural boundary of management. Each supplier establishes a reporting line against its own obligations. Each contract manager protects the rights attached to a particular agreement. Each technical lead concentrates on the component under his control.

The programme then acquires a misleading geometry: several well-defined boxes joined by thin lines. The boxes are governed rigorously. The lines are treated as coordination.

In reality, most failures occur on those lines:

- one supplier's assumption is another supplier's unpriced dependency;
- a data definition satisfies the application design but not the conversion rule;
- a test environment is technically available but lacks representative volumes;
- a change accepted within one contract alters effort in three others;
- an operational procedure belongs to the client, yet depends on knowledge scattered across suppliers.

The contract is essential evidence of obligation. It is a poor map of the outcome.

3. Nobody Is Paid to Own the Space Between

Consider a composite eighteen-month programme with a budget of £42 million. Four principal suppliers are engaged: one for enterprise software, one for integration, one for hosting and networks, and one for data conversion. The client retains business process design and training.

The design catalogue contains 312 interfaces. Each interface has a named technical owner. That sounds complete. But 37 interfaces cross three or more contractual boundaries, and no role owns their end-to-end behaviour in live operation. The systems integrator owns the message leaving the application. The network supplier owns transport. The hosting provider owns receipt at the data centre. The conversion specialist owns the source extract. The client operations team owns the overnight schedule.

When the rehearsal fails, five explanations appear. All are partly correct. The extract arrived at the agreed contractual time but too late for the actual operating window. The network capacity met its specification but the specification used average rather than peak volume. The application processed the file at its tested rate, but the test used one-third of the live record count. Operations had raised the window constraint in a meeting, but it was not translated into an accepted change.

The defect remains open for eleven working days while the parties establish where the obligation sits. During those eleven days, the programme cannot retest the full sequence. The delay is not caused by the technical correction, which takes two days. It is caused by the search for an accountable owner.

This pattern persists because the commercial incentives are rational at the level of each contract. A supplier that accepts an ambiguous dependency may inherit cost without revenue. A contract manager who concedes responsibility weakens a later claim. A

technical team that changes its component without formal instruction risks failing acceptance. What appears from the steering committee as reluctance is often disciplined self-protection.

The system produces the behaviour for which it was designed.

4. Integration Is Not a Work Package

The usual response is to purchase an integration service. One supplier is appointed to maintain the plan, control interfaces, coordinate testing and assemble status reports. This is useful, but it does not necessarily close the accountability gap.

Integration is too often defined as another bundle of activities:

- maintain the integrated plan;
- chair design authorities;
- manage interface documents;
- coordinate defect meetings;
- prepare release and cutover schedules.

Those activities can all be performed competently while the outcome remains ownerless. The integrator may identify a dependency yet lack authority to instruct another supplier. It may escalate a risk yet have no right to trade scope against time. It may coordinate a test yet be unable to require the client business to provide staff. It may report the consequences of a decision without owning the decision itself.

An integrated plan is not integrated accountability. The first records the relationship between tasks. The second gives somebody authority over the choices created by those relationships.

This distinction matters because programmes are not assembled like machinery from stable parts. Requirements move. Data quality becomes visible late. Suppliers discover omissions. Business units resist standardisation. Every such event creates a choice across boundaries: spend more, reduce scope, delay, accept risk or redesign the sequence. Coordination can expose the choice. Only governance can make it.

5. The Client Cannot Outsource the Whole

There is a persistent hope that a sufficiently comprehensive contract will transfer responsibility for the outcome. It cannot.

A supplier can accept responsibility for delivery against specified requirements. It cannot own whether those requirements still serve the organisation, whether business units will adopt the new process, whether operational disruption is tolerable or whether additional expenditure remains justified. Those judgements depend on authority that properly belongs to the client.

This does not mean that every decision should return to a large steering committee. That would replace divided accountability with delay. It means the client must create a small, explicit centre of integrated authority before contracts are signed, not after problems emerge.

That centre needs four things:

1. One outcome owner. A senior client executive must be accountable for the complete operating result, not merely for the portfolio of contracts.
2. One integration authority. A client-side programme director must have delegated power to resolve cross-supplier priorities within defined cost, scope and risk tolerances.
3. One dependency ledger. Critical assumptions, interfaces and hand-offs must be recorded as end-to-end obligations, with an owner for the complete chain rather than only its component parts.
4. One decision route. When a boundary dispute affects the outcome, the issue must move rapidly from contractual interpretation to an explicit programme decision.

The suppliers still own their contractual commitments. Commercial managers still protect the organisation's rights. Independent assurance still tests whether obligations and controls are being honoured. But none of these roles substitutes for the client's ownership of the integrated result.

6. Commercial Discipline and Programme Judgement

The tension is most visible when a problem could become a claim. Commercial discipline says: do not concede liability, preserve evidence and follow the notice provisions. Programme judgement says: restore the outcome while there is still time to do so.

Both are necessary. Treating every issue as collaboration can surrender legitimate rights and reward poor performance. Treating every issue first as a dispute can make the recovery more expensive than the original defect.

The useful distinction is between recovery and attribution. Recovery asks what must happen now to protect the programme. Attribution asks who ultimately bears the cost. They should proceed together, but they should not be allowed to block one another.

In the failed rehearsal, the integrated authority might instruct the conversion supplier to advance the extract, fund a temporary capacity increase and require a full-volume performance test within five days. The immediate cost might be £180,000. Commercial teams can preserve the client's rights and allocate that cost later. Waiting eleven days to determine liability first may consume a rehearsal slot worth considerably more and push the opening date into the next financial quarter.

This is not loose governance. It is governance that recognises time as an economic exposure.

Contract-centred question	Outcome-centred question
Which supplier owns the defect?	What must be restored before the next commitment?
Has the deliverable met acceptance?	Does the end-to-end service work at live volume?
Who should fund the change?	What decision protects the outcome now?
Was notice served correctly?	What evidence must be preserved while recovery proceeds?
Is each work package green?	Is the operating result credible?

7. The Steering Committee Must Govern the Gaps

A multi-vendor steering committee often spends too much time hearing separate supplier reports. This reproduces the fragmentation at the highest level. Each party explains performance against its own plan; the client receives several accurate descriptions and still lacks a view of the whole.

The agenda should instead concentrate on the spaces no contract can govern alone:

- dependencies whose failure crosses commercial boundaries;
- assumptions that no supplier is authorised to validate;
- decisions requiring a trade between scope, time, cost and operational risk;
- commitments that will reduce the client's remaining freedom to act;
- disputes where delayed recovery threatens more value than the claim itself.

Supplier status remains useful, but it becomes supporting evidence. The subject of governance is the integrated outcome and the next decision needed to protect it.

Minutes should record more than actions. They should state what exposure was accepted, which assumption justified it, who owns the consequence and when the decision must be revisited. Without that record, a committee may repeatedly acknowledge the same boundary risk while believing that escalation itself constitutes control.

8. Ownership After Division

The multi-vendor model is neither inherently flawed nor inherently superior. It is a commercial arrangement with a particular governance cost. It can preserve competition, access specialist capability and reduce dependence on a single supplier. It also removes the convenient fiction that one contract naturally contains the whole programme.

That removal can be healthy. It forces the organisation to confront a truth that prime contracting sometimes obscures: ultimate accountability for transformation never left the client.

The question is therefore not whether suppliers should collaborate more, though they often should. Nor is it whether contracts should be tighter, though vague obligations create avoidable disputes. The deeper question is whether the organisation has retained enough authority, knowledge and judgement to govern what lies between the contracts.

When it has not, every supplier can be correct and the programme can still be wrong. Every deliverable can be accepted and the operating result can still fail. Every report can be green and the organisation can still arrive late at an outcome nobody owns.

The craft of multi-vendor governance begins after procurement has divided the work. Its purpose is to put the whole back together without pretending the boundaries have disappeared.