The Large Deal Paradox

Why the best solution rarely wins — and what actually does

Reshmi Nandy | The Unfiltered Architect | reshminandy.com | March 2026

There is a question I have been asked in every large deal pursuit I have worked on, across three decades and four organizations: what does a good solution look like? The answer, it turns out, is entirely different depending on which side of the table you are sitting on. And closing that gap — between what the customer means by 'good' and what the vendor means by 'good' — is where most large deals are won or lost before the first slide is presented.

This article is my attempt to put that gap on record. Not to assign blame, and not to perform cynicism. But because I have watched the same pattern repeat too many times to believe it is accidental — and because the AI era is about to force this conversation into the open whether the industry is ready or not.

What is a large deal, exactly?

A large deal — in the IT services world — is a bundled contract for the operational management of an enterprise's production environment. It typically covers application support across the full landscape, infrastructure operations, user support, and may extend to continuous development, enhancements, and transformation projects on demand. Depending on the organization's IT maturity, the commercial appetite of the enterprise, and the strategic priorities of the moment, any of these layers may be included or excluded from the work package.

The scale threshold for a 'large deal' varies by firm, but the structural complexity — multi-tower, multi-geography, multi-stakeholder — is consistent. It is not a project. It is a long-term operating model transfer. And that distinction matters enormously for understanding why solutions fail.

What a good solution means to the CIO issuing the contract

From the enterprise side, a good solution resolves three anxieties simultaneously, and a vendor who understands all three — rather than addressing only the technical one — is already ahead of most of the field.

The first anxiety is cost containment. The CIO needs IT and infrastructure operations to stay within a defined percentage of revenue — one that sits at or below the peer benchmark for their sector. The metric the board cares about is not uptime; it is the ratio of IT spend to business output, and any solution that cannot speak to that ratio will never make it past the CFO's desk.

The second anxiety is risk absorption. The enterprise wants a partner who carries the operational risk — who has enough skin in the game to behave like an owner when something goes wrong, not a contractor who escalates and waits. This shapes everything from penalty structures to the composition of the delivery team. The CIO is not just buying a service; they are buying a transfer of accountability.

The third anxiety is commercial flexibility. Large enterprises operate in unpredictable environments. Headcount shifts, regulatory changes, acquisitions — any of these can make a fixed-rate, fixed-scope contract a liability within eighteen months. A solution that offers commercial elasticity — the ability to scale consumption without renegotiating the entire contract — is not a nice-to-have. It is, increasingly, the deciding factor.

> *The CIO is not just buying a service. They are buying a transfer of accountability — and any solution that cannot speak to that distinction will be outmaneuvered by one that can.*

What a good solution means to the service company delivering it

The vendor's definition of a good solution, in 2026, has been substantially rewritten by ITIL 4 and the observability wave. The modern service company is no longer proposing application-by-application support. They are proposing value stream management — organising support around the business processes that generate revenue, not the technical components that underpin them.

In practice, this means several things happening in parallel. Observability platforms — Dynatrace, AppDynamics, ServiceNow — replace manual monitoring. Automation handles business service restoration before a human even reads the alert. Self-service channels and intelligent categorizations reduce the noise at the front of the queue. Preemptive incident detection means failures are resolved before the user reports them.

The KPIs that sit above all of this — cost reduction year-on-year, NPS improvement, value stream performance against a shrinking cost base — are the commercial proof points the vendor needs to protect margin, and the customer needs to report to the board. A modern solution brings a quantitative dashboard that holds all these parameters in a single view: commercial, operational, and experiential.

The honest observation here is that these good solutions are no longer rare. They are table stakes. Most vendors of any credible scale can produce this framework. Which raises the more interesting question: if the solutions have converged, what actually differentiates the winner?

The homogenization problem — and the confidence premium

Here is something that is rarely said directly in the industry but is understood by everyone inside it: the solution architectures submitted by a $5B services firm and a $60B services firm for the same opportunity are, in substance, largely identical. The ITIL 4 value stream model, the observability layer, the automation roadmap, the gain-share commercial structure — these have been standardized across the market. The differentiation, such as it is, has migrated almost entirely to the presentation layer.

The $60B firm presents the same solution with what I can only describe as structural confidence. Not arrogance, not fabrication — but the quiet assurance that comes from institutional scale, OEM endorsements, and a reference list that includes the client's peer set. The $5B firm presents with competence but without that gravitational pull.

This matters because the CIO, standing in front of a board that will scrutinize their vendor selection, needs a decision that is defensible, not merely correct. And a decision backed by the institutional credibility of the largest player in the market is almost always more defensible than one backed by a technically superior but less visible alternative.

> *When solutions converge, the decision migrates to confidence. And confidence, at the enterprise level, is built not from proposals but from institutional gravity — analyst endorsements, OEM partnerships, peer references.*

The market intelligence that informs this dynamic — the benchmarking of rates, the mapping of competitive positions — happens through both formal channels (Gartner, ISG, third-party advisors) and informal ones: professional networks, industry conversations, the accumulated signal of years of market exposure. Neither is illegitimate. Together, they form the ecosystem within which a vendor's credibility is calibrated before any formal response is submitted.

Why incumbent vendors hold the structural advantage — and when they do not

The fear of switching vendors is real, and it is rational. It mirrors, at enterprise scale, the same calculation any individual makes when changing a long-standing service provider. The known failure modes of the incumbent are at least understood. The failure modes of a new vendor are not. For a CIO managing a $500M IT operation, that asymmetry is never trivial.

What follows from this is that the incumbent vendor, in most large deal renewals, holds a structural advantage — provided they have met the minimum bar of the key decision-maker's expectations. Not exceeded it. Met it. The threshold for retaining a large deal contract is often lower than the industry assumes, which is why delivery excellence matters less in renewal cycles than relationship equity.

The incumbent loses that advantage under one condition: when the delta between expectation and experience has accumulated to a point where the CIO can no longer defend the relationship to the board. This is the window that challengers must operate within. And it is narrower than most challengers realize — because the CIO will often attempt to reset the relationship with the incumbent before opening the door to a replacement, even when the failure is systemic.

The delivery gap — and why winning is the easier half

What happens after the win is where the industry's most honest conversations happen, and they do not happen in public.

Large deals, almost structurally, run on compressed margins. The pricing pressure of the pursuit — driven by benchmarking, competitive tension, and the client's commercial sophistication — leaves very little room for the delivery organization to operate comfortably. The solution that won was optimized for persuasion. The organization that delivers it is optimized for efficiency. These are not the same objective, and the gap between them is where most large deal failures originate.

The client partner who closed the deal will maintain visibility with the customer. Behind them, a delivery organization of hundreds will carry the weight of the commitment — often with insufficient context about the win themes, the commercial pressures, or the client's unstated expectations. The Sufficiency Gap — which I wrote about separately — lives here. The solution document that won is not the operating manual that delivers.

> *The solution was optimized for persuasion. The delivery organization is optimized for efficiency. These are not the same objective — and the gap between them is where most large deal failures begin.*

This is not an indictment of any organization. It is a structural reality of how large deal economics work, and it persists because the incentives that drive winning are not aligned with the incentives that drive delivery. Changing that alignment — building commercial models where the same team that pursues also owns the margin consequence — is the reform the industry is slowly, reluctantly beginning to consider.

What the Agentic AI shift changes — and what it does not

The arrival of Agentic AI at enterprise scale will change the economics of large deals more fundamentally than any development of the past decade. The unit of delivery — the human FTE — is being replaced, in the automation layer at least, by an AI agent that can handle L1 and L2 resolution, documentation, categorization, and pattern detection at a fraction of the operating cost.

For service companies, this is both an opportunity and a structural threat. The opportunity is obvious: the same scope, delivered at lower cost, yields higher margin. The threat is equally obvious: once customers understand the new cost basis of AI-native delivery, they will renegotiate contracts accordingly. The 20% productivity gain that AI delivers in the greenfield case becomes, in the large deal context, a negotiation lever in the next renewal cycle.

What Agentic AI does not change is the relationship premium. The CIO who trusts their partner — who has built equity through consistent delivery and honest communication — will still give that partner the first conversation about the AI-enabled future. The CIO who does not will go to market. Technology commoditizes the execution layer. It does not commoditize the judgement layer — and the judgement layer is where the architects, the advisors, and the relationship managers still hold irreplaceable value.

The large deal of 2027 will look different from the large deal of 2024. The tower structure will compress. The headcount pyramid will flatten. The observability and automation layer will be assumed, not differentiated. What will remain as the primary arena of competition is exactly what has always mattered most: the ability to understand what the customer is actually trying to accomplish — and the courage to say so, even when the answer is inconvenient.