



Don't Dissolve It. Steel It.

At scale, dissolving Business Services in the age of Agentic AI is an enterprise's costliest operating-model mistake. There is no version where it works.

June 2026

By Rajiv Gupta and Matt Marchingo



The Argument

A new operating-model thesis is circulating in boardrooms and AI-strategy decks: That agentic AI now makes it feasible for companies to dissolve Global Business Services (GBS) in 18–24 months and “send the work back to the functions.” This thesis is wrong in every dimension that matters: structural, financial, legal, regulatory, operational, and historical. A one-time direct restructuring cost lands at 20–40% of the annual GBS run cost. The recurring annual value cannibalization, once function-owned agent swarms take over, lands at 65–100% of the annual GBS run cost. The permanent unit-cost penalty against top-quartile shared services is 70–110 basis points of revenue. The historical pattern is that every attempt by a large enterprise to dissolve a centralized services backbone in the last 20 years has been reversed within 5–7 years at 2–3x the original program cost. The total expected lifetime cost: 4–7x the annual GBS run cost over a decade. The only defensible play is the exact opposite approach: Repositioning GBS as the enterprise AI Control Tower. Anything else is a future write-down.

1. Why the Argument doesn't hold

The dissolution argument is being presented as a contrarian, AI-enabled strategy. In reality, it is flawed across every dimension that matters: Asset, AI, Functions, History, and Governance.

Here's how:

- **Asset.** Enterprise-scale GBS is a multi-jurisdictional legal-tax-process-technology operating system, worth roughly 1.5–2.5% of revenue in annual run cost. It is not a labor pool you switch off.
- **AI.** Agents are tools. The control tower is management. The dissolution thesis proposes to fire the management and keep the tools.
- **Functions.** Finance, HR, supply chain, and procurement at scale do not run end-to-end services. They consume them. They cannot absorb the work back.
- **History.** Every documented attempt to dissolve a centralized services backbone at large-enterprise scale in the last 20 years has been reversed within 5–7 years at 2–3x the original program cost. There is no counterexample.
- **Governance.** An AI cannot sign a SOX certification. An AI cannot be deposed by a regulator. An AI cannot own a P&L. The dissolution thesis assumes accountability is a software problem. It is not.

No executive should, therefore, sponsor the dissolution thesis.

2. The Scale Reality the Dissolution thesis avoids

The proposition makes sense only if you ignore what GBS is at enterprise scale. A typical GBS global enterprise footprint includes:

- 15–50 captive centers and BPO vendor sites across multiple jurisdictions, with embedded talent, regulatory footprint, and SEZ/tax-holiday positions.
- Long-dated, multi-year BPO master services agreements with exit penalties, ramp-down obligations, and transition tail.
- Transfer pricing and entity design built around the captive and vendor footprint.
- Globally consolidated ERP, HRIS, procurement, and supply-chain platforms configured for centralized delivery.
- Tens of thousands of FTE-equivalents performing transactional and judgmental work across order-to-cash, procure-to-pay, record-to-report, and hire-to-rotate processes.

All of this can neither be switched off nor “sent back” in 18 months. The BPO contract stack alone is rarely renegotiated within that timeframe.

3. The Economics the Dissolution deck does not show

After you strip away the operating-model abstraction, here is what enterprise-scale GBS represents, expressed in magnitudes that scale across companies. (See Figure 1).

The spread alone is decisive. Decentralized finance is 70–110 basis points of revenue more expensive than top-quartile shared services. Dissolution re-introduces this spread - every year, forever, on top of one-time restructuring cost. There is no version of the model in which this is a value-creating move.

4. Five Structural Anchors that make 18-month Dissolution impossible

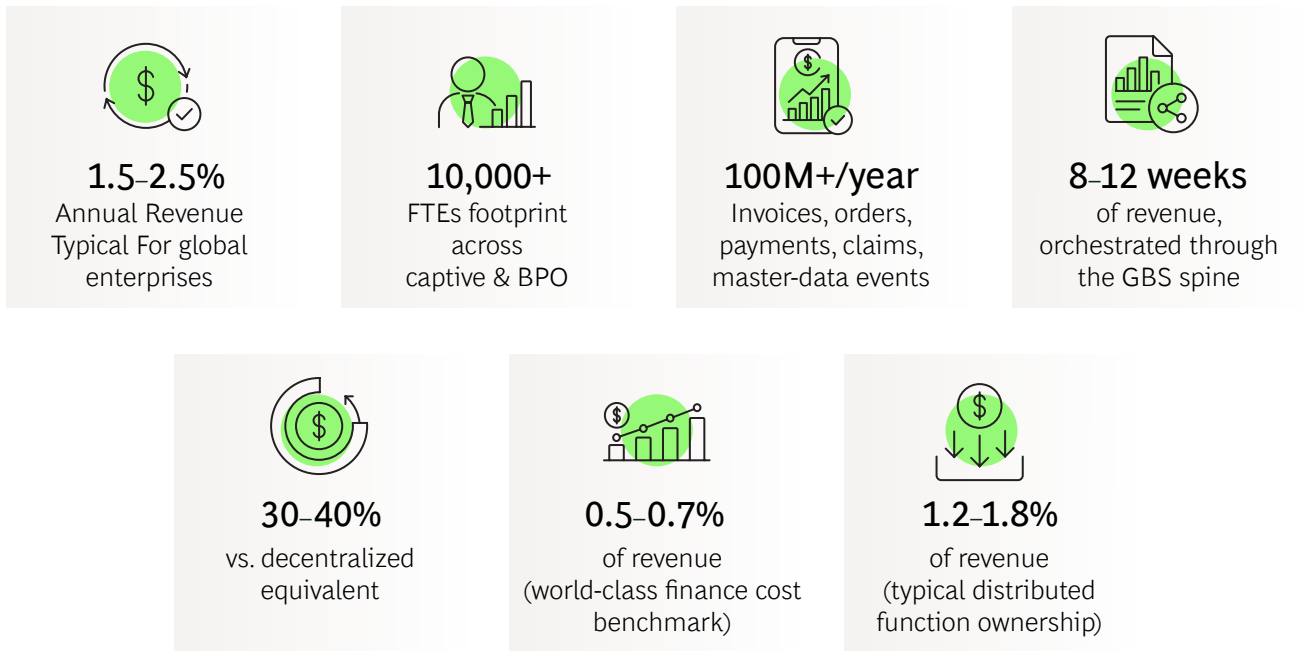
Not “hard.” Not “ambitious.” Impossible.

- **Contractual lock-in.** BPO master services agreements at scale typically carry 12–18-month exit notice periods, 6–12 months of ramp-down obligations, and material exit fees per major contract. Across a multi-vendor stack of dozens

of sites, full disengagement alone is a multi-year initiative, costing a significant percentage of the GBS run cost. The legal calendar does not bend to the case for dissolution.

- **Tax and transfer-pricing dependency.** GBS entities are central to global TP design. Unwinding triggers entity restructuring, re-characterization, profit reallocation, and loss of tax-holiday and SEZ benefits. TP restructuring across 30+ jurisdictions is a 24–36 month, partner-led, regulator-paced track. This is a P&L event – and often a 10–K disclosure event.
- **End-to-end process integration: O2C, P2P, R2R, and H2R cross sales, operations, finance, tax, and HR.** Decentralizing them re-fragments precisely what GBS seeks to consolidate. Industry data consistently shows process unit costs rising 50–100% within 18 months of decentralization, without a compensating platform layer.
- **Platform architecture.** The ERP, HRIS, and procure-to-pay stack are configured for centralized delivery, with – workflow, master data, controls, and SLAs all designed around the hub. Disaggregation is itself an 18–30 month, capital-intensive program.
- **Operational risk.** The framework, ‘pay, close, ship, invoice, file, and report,’ cannot fail. The risk surface of an aggressive 18-month unwind is materially bigger than the savings hypothesis.

FIGURE 1



5. The Functions trap

The dissolution argument’s premise – “send it back to the functions” – assumes the functions can absorb it. They cannot.

Finance, HR, supply chain, and procurement at scale do not run end-to-end services. They consume them. GBS was built precisely because the functions struggled to deliver consistent, controlled, low-unit-cost transactional work across the enterprise. Sending it back means rebuilding what GBS was originally consolidated to replace, that too in a higher unit-cost environment, without the governance, and at a moment when those same functions need bandwidth to lead their own AI agendas.

Decentralizing transactional work in the age of AI is not freedom. It is a tax. The dissolution argument is asking the functions to pay it.

6. The Order-to-Cash cannibalization scenario

Take the dissolution thesis at face value. Dissolve GBS. Send O2C “back to the functions.” Let each function deploy its own agentic swarm to run its slice. The ensuing six swarms, with no end-to-end owner, would include:

- **Sales agents** that own order capture and pricing – optimized for booking velocity, not credit risk.
- **Customer-service agents** that own disputes and credits – optimized for resolution speed; they give away margin to close tickets.
- **Operations agents** that own fulfillment – optimized for service level, not landed cost or invoiceability.
- **Finance agents** that own invoicing and AR posting – working from stale or drifted master data.
- **Credit/Treasury agents** that own credit decisions and collections – with no visibility into upstream dispute patterns.
- **Tax agents** that own VAT/GST classification – over a product master that no one function owns end-to-end.

Every agent would win its local KPI. The enterprise would lose cash.

Here is the cannibalization math at enterprise scale, for O2C alone, expressed in magnitudes that scale across companies. (See Figure 2).

This is the cost of bit-part ownership: six agent swarms, each optimal for its function, none accountable for the enterprise outcome. If you multiply the same pattern across procure-to-pay, record-to-report, and hire-to-retain, the conservative aggregate value cannibalization across the four end-to-end towers: **65–100% of annual GBS run cost destroyed every year.**

The control tower exists precisely to prevent this. Centralized data, single accountability, shared agent governance, end-to-end KPIs. AI does not remove the need for this layer. It increases it. Anyone who has stood up an agentic capability inside an enterprise in the last 24 months knows this.

FIGURE 2

DIMENSIONS	TYPICAL RANGE
DSO EXPANSION	+50% DSO drifts ~40d → ~60d+ working capital trapped equivalent to 15–25 days of revenue
DISPUTE RATE SPIKE	2–3x Disputes rise from ~3% → ~9% deductions and credit leakage of 5–10 bps of revenue
BAD DEBT WRITE-OFFS	+30–60 bps of revenue (rise from typical 0.3% → 0.7–0.9%)
VAT/GST IRRECOVERABLE	5–15 bps of revenue fragmented tax controls recoverability gaps across jurisdictions
TOTAL – O2C ALONE	30–45% of annual GBS run cost, destroyed annually Recurring. From dissolving one process

7. We've seen this Movie – and it always ends the same way

The dissolution-of-the-backbone playbook is not new. It has run, in various forms, across multiple decades, sectors, and management fashions. The pattern is consistent. The reversal is consistent. The bill is consistent. (See Figure 3).

This pattern has played out across CPG, financial services, industrials, telecoms, and now technology. The names change. The structure does not. Centralized backbones that are dissolved tend to come back - with a different branding - because the work they did does not go away. Only the accountability does.

There is no documented case of a large-enterprise shared-services dissolution at this scale that was not reversed. The dissolution thesis is asking enterprises to beat history.

The dissolution argument is not a transformation. It is a future re-centralization at 2–3x the original program cost.

8. What Agentic AI actually does – and what it does not

Agentic AI is real. It will compress automation, accelerate process discovery, reduce manual orchestration, and shrink the labor envelope of GBS materially over the next three to five years.


The effect: A 10–30% effort compression on the transformation timeline itself; a 30–50%+ structural labor compression on transactional work over a 3–5 year horizon.

What AI does not do:

- Negotiate out of a multi-year BPO master services agreement, or eliminate exit fees.
- Restructure transfer pricing across 30+ jurisdictions, or compress a 24–36 month tax track.
- Take fiduciary accountability for SOX–relevant controls, statutory filings, or regulatory reporting.
- Fix the master data quality issues an aggressive unwind would amplify – issues that become the primary cost driver in a bit-part swarm.
- Make functions willing or able to run end-to-end services they have never owned.

The idea that AI dissolves the structure stems from mistaking what it automates with what it owns. This confusion is the basis of the 18-month argument. Acknowledge it, and the thesis falls apart.


FIGURE 3



Year 0

THE ANNOUNCEMENT.


A new operating model promises agility, lower cost, and “ownership at the edge.” Centralized services are broken up. Functions take their work “back.” The CFO is happy.



Year 1–2

THE SPIKE.


Unit costs rise 30–50%. Controls weaken. Audit findings multiply. Vendor management fragments. Data quality drifts. Working capital expands. Functions complain they are running services they were never built to run.



Year 3–4

THE PATCH.

The organization quietly rebuilds capability inside the functions — duplicating cost. Compliance hires double. Shadow shared-services teams appear under different names. Cost trajectory inverts and starts climbing.



Year 5–7

THE REVERSAL.

A new operating-model project is launched under a different banner — Global Operations, Enterprise Services, Digital Backbone, AI Operating Layer — to do exactly what the dissolved unit used to do at 2–3x the original program cost and with significant senior departures along the way.

9. The only Defensible play: GBS as enterprise AI Control Tower

There are two operating-model options on the table. Dissolution is not one of them. Structure, economics, history, and governance rule it out.

The two real options are: (a) the maintenance of GBS in its current form, running it by attrition; and (b) the steeling of it as the enterprise AI Control Tower. Only one of those creates value.

Use AI to reposition GBS as the orchestration brain of the enterprise. (See Figure 4).

In this model, GBS becomes the single source of truth for transactional data, the deployment and governance layer for enterprise agents, the standards-and-controls owner for cross-functional processes, the talent magnet for AI, data, and

process expertise, and the mechanism through which the functions are freed up from being buried under work they cannot do.

The footprint thins (target: 25–35% captive and BPO site reduction by month 18). Transactional headcount compresses 30–50%+ over 3–5 years. But the backbone stays – because the backbone is the moat.

10. The full P&L of getting this wrong

The dissolution argument stops at “savings.” The full P&L does not. (See Figure 5)

Expected value of “dissolve in 18–24 months”: **structurally negative across every defensible scenario.** Expected value of “AI Control Tower in 18–24 months”: **materially positive.** Same calendar. Fundamentally different outcome.

FIGURE 4

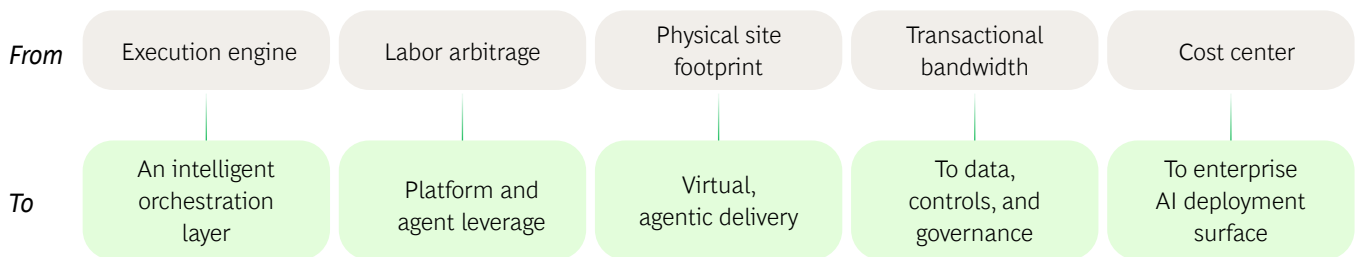
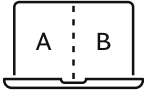



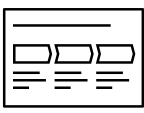
FIGURE 5

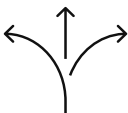
	MAGNITUDE	TIMING
DIRECT RESTRUCTURING COST	20–40% of GBS run cost	ONE-TIME (over 18–24 months)
RECURRING ANNUAL CANNIBALIZATION	65–100% of GBS run cost	EVERY YEAR once function swarms own the work
PERMANENT UNIT-COST PENALTY	70–110 bps of revenue	FOREVER — the structural spread vs top-quartile
FUTURE RE-CENTRALIZATION PROGRAM	1.5–2X GBS run cost	5–7 YEARS OUT — inevitable on the precedent
TOTAL LIFETIME COST (10-YR VIEW)	4–7X GBS run cost	Plus career exposure for sponsoring executives

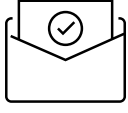
11. The Recommendation

- 

1 Reject the dissolution thesis. It is not a debate between two operating models. It is a debate between a model and an unfeasible thesis. Take it off the table.
- 

2 Mandate “GBS as AI Control Tower” as the explicit operating thesis for the next 18-24 months. Frame it as Phase 1 of a multi-year transformation, not the destination.
- 

3 Lock three Phase 1 outcomes by month 18: (a) agentic automation deployed across $\geq 40\%$ of GBS transactional volume; (b) GBS-hosted enterprise data and agent governance platform live with end-to-end accountability for the four towers (O2C, P2P, R2R, H2R); (c) Captive and BPO footprint rationalized by 25–35% with zero service disruption and DSO held flat or improving.
- 

4 Defer entity and transfer-pricing restructuring to a deliberate, multi-year, tax-led workstream. Do not let an AI narrative drag legal-entity strategy.
- 

5 Re-message GBS internally: From being a cost engine to an AI deployment surface. This unlocks the functions.

The Bottom Line

The dissolution thesis is wrong in every dimension that matters: Asset, AI, Functions, History, Governance, and Economics. At any meaningful enterprise scale, dissolution carries one-time restructuring cost of 20–40% of annual GBS run cost; recurring annual cannibalization of 65–100% of annual GBS run cost; a permanent 70–110 bps unit-cost penalty against top-quartile shared services; an inevitable future re-centralization at 1.5–2x annual GBS run cost; and a 20-year track record of failure with zero documented counterexamples. Total expected lifetime cost: 4–7x annual GBS run cost over a decade. There is no version of the model where this works. The only defensible play is to keep the backbone, steel it with AI, and free the functions through orchestration. GBS at scale is not the demolition target. It is the AI Control Tower.

No executive ought to sponsor the dissolution thesis.

About the analysis

Quantitative anchors in this piece are expressed as magnitudes (percentages of revenue, multiples of annual GBS run cost, basis points) that scale across enterprise contexts. They are drawn from industry benchmarks for shared-services economics, standard order-to-cash metrics (DSO, dispute rates, bad-debt ratios, VAT recoverability), BPO contract economics (exit fees, transition obligations, ramp-down), and 20+ years of large-enterprise reorganization pattern data. Specific values for any given enterprise depend on sector, geography, and operating-model maturity, and require client-specific validation before application.

About the Authors



Rajiv Gupta

Managing Director and Senior Partner
New Delhi
Gupta.Rajiv@bcg.com



Matt Marchingo

Managing Director and Partner
Houston
Marchingo.Matt@bcg.com



Sayan Majumdar

Platinion Director
New Delhi
Majumdar.Sayan@bcg.com

Acknowledgments

We extend our gratitude to the Nopur Nopur from BCG India Marketing Team for managing the marketing. We would also like to thank Saroj Singh, Ratna Soni, and Alivya Saha for their contributions to the editing, design, and production of this report.



Where Strategic Clarity Meets Applied AI

We are navigating an era of unprecedented change and disruption—powered by technology, marked by complexity, where change amplifies at scale. To lead, companies need a partner that can bridge the gap between ambition and outcomes. BCG is built for this moment. We bring strategic clarity, rooted in over 60 years of deep domain knowledge, to ensure leaders make the right choices. We combine it with applied AI, shaped and wielded by our practitioners, teaming shoulder-to-shoulder with your teams to deliver transformative impact at scale. The result? Stronger returns, transferred capabilities, and change that sticks. We are BCG.

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit bcg.com. Follow Boston Consulting Group on [LinkedIn](#), [Instagram](#), [Facebook](#), and [X](#).

