VIEW POINT



THE EFFORT ILLUSION IS OVER. THE ATTRIBUTION PROBLEM IS NOT.

Abstract

Outcome-based pricing has stalled in BPM — not for lack of appetite, but for lack of proof. BPM has priced effort for three decades, but agentic AI has broken the link between effort and value: When an AI agent resolves work, the effort disappears while the value it produced does not. Outcome-based pricing is the model everyone now agrees on in principle, yet it remains a small share of signed BPM revenue. This paper argues that the constraint is evidentiary, not commercial: Providers cannot yet prove causal attribution to a standard that a CFO will accept or forecast agentic cost predictably enough to underwrite a fixed outcome price. It proposes an integrated agentic FinOps discipline — real-time telemetry, attribution mapping, cost transparency and commercial governance — as the mechanism that makes outcome-based pricing executable rather than aspirational and sets out where to start.

Infosys
Navigate your next



01 An industry at its inflection point

For three decades, BPM sold effort. Headcount multiplied by hours, multiplied by rate. The model worked because effort was an honest proxy for value: More transactions needed more people, and clients could audit spend by counting desks.

Agentic AI has broken that proxy. When an agent resolves an order exception, the effort disappears but the value does not, and the invoice, priced on effort, follows the effort rather than the value. The industry is now caught between an input-based model that no longer describes what we do and an outcome-based model we cannot yet reliably execute.

Consider an illustrative account. A consumer goods manufacturer outsources order management and fulfilment across three geographies. The contract specifies 120 analysts. Every quarterly review is positive: service-level agreements (SLAs) are met at 98%, disputes are closed within the agreed turnaround time, and headcount is fully deployed. Yet the client is dissatisfied: Revenue leakage from unresolved deductions runs at 3% of net revenue, and on-time, in-full (OTIF) performance has not improved in two years. While automation handles the bulk of exception volume, traditional input-based commercial models keep the engagement tethered to headcount metrics, leaving both providers and clients without a contractual framework tied to board-level business outcomes.

The economics: Green on the dashboard. Red in the boardroom. On $400 million of revenue in scope, a 3% leakage translates to roughly $12 million a year, and a contract capturing even a quarter of the recoverable portion would be worth more than the entire analyst fee line beside it.

Value resides in the business outcome that the effort is intended to produce; traditional contracts do not directly reward that outcome.

02 The case for outcome pricing is settled. The mechanism is not.

Recent market reactions illustrate this disconnect: Global service providers reporting strong margin and earnings have faced valuation pressure where traditional booking velocity slowed. It was the steepest one-day fall on record—even as the provider reported strong growth in AI-related bookings. Demand for AI-led work is not the problem. The pricing unit is.

The same structural shift is visible across the market. As agentic AI replaces standardized workflows with context-driven orchestration, process-centric contracts will face increasing pressure to narrow the gap between cost and value. Buyers now open renewals by naming the outcome first.

Yet, outcome-based contracts remain a small fraction of signed BPM revenue — a gap not explained by commercial appetite on either side.

The binding constraint is evidentiary, not commercial. We cannot yet prove, to a standard that a client's or provider's CFO will accept, what share of a business outcome we helped generate, or what it will cost us to keep generating it.

External Document © 2026 Infosys Limited

Outcome pricing asks the provider to accept two risks at once: That the result will materialize, and that the cost of producing it is predictable. Input-based pricing insulated us from both. Agentic delivery has made the first measurable in principle and the second harder in practice, and BPM has built the instrumentation for neither.

03 Four drivers of the complexity

The difficulty is not one problem but four, and they compound

Driver	Why it bites specifically in BPM	What it blocks
Control boundary	**Providers own one slice of an end-to-end chain; upstream master data and downstream logistics dictate the final metric.**	**Accepting unhedged outcome risk on metrics without end-to-end control**
Attribution uncertainty	**Proving causal share to audit standard is difficult due to baseline drift and multi-causal business variance.**	**Defensible gainshare mechanics and clean baseline agreements**
Non-linear cost	**Agentic workflows scale unpredictably across multi-step loops, tool invocations, validation, and re-ranking.**	**Underwriting fixed-outcome pricing against volatile computing/orchestration costs**
Absence of agentic FinOps	**Lack of a unified telemetry ledger that reconciles human FTE cost, gent compute, and busines utcome delivere**	**Managing outcome contracts with realtime margin visibility rather than annua bets**

The fourth driver is worth dwelling on because it is purely a build problem. Cloud computing experienced a version of it: The promise was to stop buying servers and pay for use, but spiky demand forced hyper-scalers to build real-time financial telemetry over infrastructure. Cloud FinOps succeeded because hyper-scalers instrumented cost and usage at the exact consumption grain, giving enterprise buyers verifiable cost telemetry.

BPM did not, and the cost side has become blunt. According to Everest Group, the visible cost of AI is in token consumption, but the larger, more underestimated share is in orchestration, validation, governance, and continuous tuning; hence, AI economics must be measured at the level of business outcomes rather than model usage. Buy-side evidence agrees: One global technology firm exhausted its annual AI coding budget within four months after an agentic coding tool reached roughly 5,000 engineers, with heavy users costing $500 to $2,000 each per month. According to a survey by Benchmarkit and Mavvrik, 85 percent of companies missed AI cost forecasts by more than 10%, and 84% reported AI spend compressing gross margins by more than 6 percentage points.

A provider that cannot forecast its own agentic cost cannot responsibly sell an outcome at a fixed price. Cost predictability is not merely a finance concern within the commercial model; rather, it is a precondition.

04 A framework: outcome-based pricing on integrated agentic FinOps

Outcome pricing in BPM requires a discipline that reconciles cost against a business outcome across a process boundary we only partly control, while human effort and agentic consumption move independently. It comprises four layers, built once and instrumented per tower.

External Document © 2026 Infosys Limited

Layer	What it does	What makes it contract-ready
Real-time telemetry	Instruments the contracted outcome and agent consumption at the same event grain: every exception, every agent call, and every recovered deduction, timestamped in one ledger.	Outcomes measured continuously rather than assembled for the QBR
Attribution	A documented boundary map of steps we control, share, or do not touch, with an agreed error band and a published re-baselining protocol.	Gainshare on the controllable share, with external variance formally carved out
Transparency	Cost per completed transaction, not per token, disclosed to the client, against their commitment to supply the upstream data quality the outcome depends on.	A two-sided contract in which both parties carry a named obligation
Governance	A commercial forum, not a service review: CCO and CSCO in the room, reviewing outcome delivered against blended cost to serve.	Tier movement, re-baselining and dispute resolution as routine mechanics

How the contract actually executes: A two-quarter baseline period establishes the outcome baseline and, critically, the attribution error band, published rather than negotiated afterwards. Commercials then sit on a floor plus a band: a fixed fee covering delivery cost at baseline performance, plus gainshare on improvement above it, applied only to the controllable share the boundary map defines. Agentic cost is governed on cost per completed transaction, with model routing, context caching and capped retry loops as owned engineering discipline, and efficiency gains split rather than silently retained, keeping the client's incentive to expand the scope intact. Re-baselining is scheduled, not triggered by argument.

The trigger to move an account from managed services to outcome pricing should be mechanical and published: two consecutive quarters of a stable, mutually agreed baseline with attribution error inside the agreed band. Name the band, and outcome pricing becomes a commitment with a date rather than an aspiration with a caveat.



External Document © 2026 Infosys Limited



05 Insights, imperatives, and where to start

Three insights follow. First, the constraint is evidentiary, not commercial: Every party wants outcome pricing, and what is missing is proof, which can be built. Second, attribution maturity varies by tower, not by client: OTIF, last-mile performance, quote-to-cash velocity, exception resolution and leakage detection are tractable now, while demand sensing, trade-promotion effectiveness, and complex B2B pricing are not, because no single party controls enough of the causal chain. Outcome pricing should therefore be sold tower by tower within the same account. Third, instrumentation is delivery work, not overhead: Roughly a third of enterprise data is AI-ready and around half of processes still involve manual steps, so the work of making a process measurable and the work of making it agentic are largely the same work, done once. The attribution layer can therefore be funded through managed services delivery rather than competing with it.

Three imperatives follow. First, instrument one tower this quarter — fulfilment, on one live account — as a measurement pilot rather than a pricing pilot, and publish the attribution error band; an honest wide band is more valuable than an unpublished narrow one. Second, build agentic cost telemetry — cost per completed transaction by tower, tracked for two quarters before selling agentic outcomes; without it, every outcome price we quote is an unpriced option against our own P&L. Third, answer two open commercial questions with data rather than debate: Can human effort be priced by skill tier, the way agentic consumption is, so a contract blends both rather than counting FTEs alone; and do the economics hold at a 50-FTE account or only above a threshold.

Return to the account we opened with. Nothing about that engagement required new technology to fix. It required someone to instrument leakage recovery at source, agree with the client which share of it we controlled, publish the error band, and price the recoverable portion. The automation was already built. The measurement was not.

According to a forecast from Gartner, over 40% of agentic AI projects are expected to be canceled by the end of 2027 because of escalating costs, unclear value, and inadequate controls. That is not an argument against this direction. It is a description of what happens to agentic delivery without the FinOps layer underneath it.

External Document © 2026 Infosys Limited

Authors



Gururaj Sagar
Principal Consultant – Solution Design

Gururaj leads global strategic alliances for Infosys BPM's Sales & Fulfilment practice, where he is building the AI-native partner ecosystem making Infosys BPM a pioneer in the AI mandate across sales, commercial, supply chain, and service operations globally. He is responsible for the agentic AI partner stack — from sales and revenue intelligence to order orchestration, control-tower visibility, all-mile execution and post-sale service — stitched into the operating layer that makes them deployable for global enterprises.



Anshuman Kantha
AVP – Strategic Business Practice Head

Anshuman Kantha is Infosys BPM's AI-first supply chain leader who partners with client CXOs to navigate the complexities of digital transformation in today's dynamic business landscape. As the Head of Supply Chain, Anshuman has spearheaded numerous high-impact initiatives across supply chain planning, fulfilment, and aftermarket operations, delivering measurable value in CPG, hi-tech, industrial manufacturing, aerospace and defense, and more. His expertise in large-scale transformation and managed services has empowered clients to achieve significant cost efficiencies, operational resilience, and sustainable growth.

For more information, contact infosysbpm@infosys.com

Infosys
Navigate your next

© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.