VIEW POINT

REWIRING THE REVENUE ENGINE: A NEW MANAGED SERVICES OPERATING MODEL FOR TELECOM B2B REVENUE OPERATIONS

Abstract

Core challenges in European telecom B2B revenue operations are structural, not technological, arising from fragmented and manual processes across lead to quote, quote to order, and order to cash. Despite major investments in CRM, CPQ, ERP, and BSS platforms, operators continue to face slow quote cycles, high order fallout, billing complexity, revenue leakage, and extended cash conversion timelines because their operating models remain siloed and labor-intensive. The paper proposes a new managed services operating model built around agentic AI, outcome accountability, and process redesign rather than traditional labor arbitrage. Across each revenue stage, it outlines how intelligent automation, human-in-the-loop governance, and continuous process intelligence can reduce exceptions, improve accuracy and speed, and create scalable, auditable operations. It concludes that successful transformation depends not on technology alone, but on rigorous process standardization, exception-led design, portable architecture, measurable baselines, and commercial models aligned to business outcomes.

Infosys®
Navigate your next

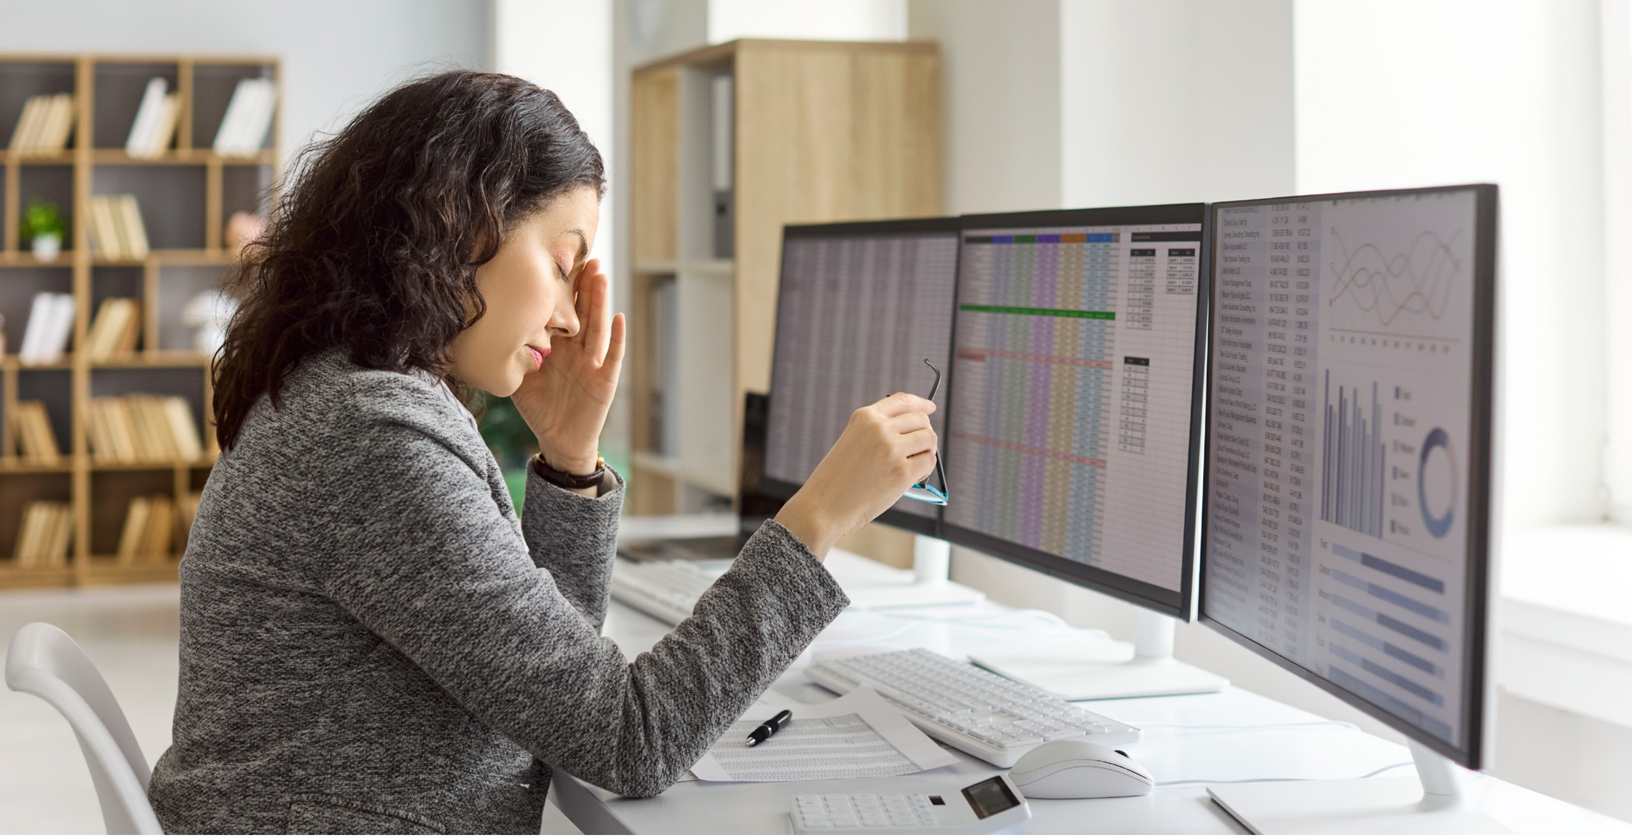

The structural problem in telecom B2B revenue operations

European B2B telecoms operators face a compounding operational problem that is structural rather than cyclical. The three process domains that constitute the commercial spine of every B2B telco business — lead to quote (L2Q), quote to order (Q2O) and order to cash (O2C) are failing in consistent, measurable and expensive ways. The root cause is not insufficient technology investment. Most large operators have spent heavily on CRM, CPQ, ERP and BSS platforms over the past decade. The problem is the operating model built on top of those platforms: Largely manual, deeply siloed and optimized for a simpler commercial environment than the one these businesses now operate on.

This paper examines why that gap exists, what a structurally different operating model looks like across each of the three process domains, and how managed services providers — when designed around outcome accountability and agentic AI rather than labor arbitrage — can close it.

65%
Of fully manual B2B O2C across the industry

300K+
Order fallout incidents per year at a representative mid-size European telco

20–28
Parallel bill cycles typical in a multi-brand telco operation

6 weeks
Average complex bid cycle time, driven entirely by manual pricing iterations

20K+
Back-office exception cases per month generated by a single telco operation

42 days
Average B2B DSO across the industry versus 25 days at top-quartile operators

These figures are drawn from operational assessments of live European telco B2B environments, not industry surveys. They represent the baseline from which transformation must start — and the scale of value that a structurally different operating model can address.

Why traditional managed services did not solve this

The managed services model that most telcos have applied to B2B operations was built on a single economic premise: Move labor to a lower-cost location, standardize what can be standardized, and pass the saving on. This model delivered cost reduction — but it also embedded the same broken processes in a cheaper wrapper. The fundamental design of the operations did not change. Exception volumes remained high. Manual handoffs persisted. The operating model continued to scale in a linear fashion with transaction volume, providing no structural relief as complexity grew.

A different model is now available — one that uses agentic AI as the primary execution mechanism rather than as a productivity tool layered onto human workflows. Let's examine what that model looks like in each of the three critical process domains.

The gap is not platforms — it is the operating model. Technology investments without operating model redesign produce a more expensive version of the same problem.

Lead to quote: Where revenue is won or lost before the order

The L2Q process is where B2B revenue is created or conceded before a single order is placed. It is also the domain where the gap between current performance and achievable performance is widest — because the activity mix combines high-volume, rules-driven work that is straightforwardly automatable with genuinely complex judgment work that requires specialist human input. Most telco B2B operations do not distinguish between these two categories. They apply the same manual process to both.

What we observe in the field

Process assessments of European telco B2B commercial operations consistently surface the same structural weaknesses in L2Q:

- **Lead management fragmentation:** Lead intake operates through shared inboxes, with manual qualification steps that introduce delays of 20–30 minutes per case before a lead reaches the sales team.
- **Spreadsheet-based pricing:** Bid and quote management in most mid-to-large telco B2B operations still relies on manual spreadsheet iterations for pricing. Standard priority deals take five to seven working days. Complex deals involving cross-functional coordination or multi-country pricing extend to four to six weeks.
- **Informal approval chains:** Discount and margin exception approvals operate through email, adding two to five working days to deal cycles with no audit trail, no escalation management and no visibility into where approvals are outstanding at any given moment.

The intelligent L2Q model

A managed services provider operating an intelligent L2Q model deploys three categories of agentic capability in sequence.

- A lead qualification and routing layer replaces shared-inbox intake
- A catalogue-aware quote assembly layer handles the 85–90% of opportunities that are standard
- An approval orchestration layer automates the approval matrix

The design principle is consistent across all three: Separate the rules-driven majority from the judgment-intensive minority and apply automation to the former so that human expertise can be concentrated on the latter. The consequence is 2X improvement in deal velocity.

In multi-country telco B2B environments — where pricing, tax treatment and contract terms vary significantly across markets — the agentic rules fabric is particularly valuable. Country-level variations that are currently managed through per-country manual workarounds, institutional memory and bespoke configurations can be encoded directly into the agent framework, making the process consistent, auditable, and scalable without adding headcount as the business grows into new markets.

The test of an intelligent L2Q model is not whether it can handle the standard case — any automation can do that. The test is whether it accurately identifies the non-standard case and routes it to the right human at the right moment, with full context pre-populated.

Quote to order: Closing the fallout gap

The Q2O process is where a commercial agreement becomes an operational commitment — and in B2B telecoms it is where the volume and complexity of exceptions is greatest. Order fallout — the failure of an order to progress through the fulfilment chain without manual intervention — is the defining operational problem in this domain. In the European telco B2B environments that have been assessed in depth, fallout rates of 15–20% are not uncommon on complex product orders. At the operational scale of a mid-to-large telco, that translates to hundreds of thousands of manual interventions per year.

The anatomy of order fallout

Practitioners who have mapped fallout patterns across multiple telco environments find that the large majority of incidents are attributable to five predictable causes: Incomplete or inconsistent data passed from the quoting system to the order management system; credit validation issues that were not caught during the quote process; product catalogue mismatches between CRM and OMS that generate invalid order configurations; provisioning capacity exceptions that were not flagged at the point of order creation; and contract terms that conflict with billing system parameters. All five are detectable before an order enters the fulfilment chain. All five can be addressed through pre-validation logic applied at the handover point between quote acceptance and order entry.

The agentic Q2O architecture

- **Pre-validation and case scoring.** An intelligent validation layer sits between quote acceptance and OMS entry. It performs multi-field validation against the accepted quote, and either passes it for straight-through processing or flags specific issues for targeted human resolution.
- **Zero-touch order entry.** Validated orders are mapped directly to the OMS and ERP via API integration, with the sales order record created and the provisioning workflow triggered without manual data re-keying.
- **Real-time jeopardy management.** An SLA monitoring agent tracks order progress against committed timelines at short polling intervals, surfacing risk to the relevant operational team before a breach occurs rather than after. This converts the exception management function from reactive firefighting to proactive intervention.
- **Rules fabric for multi-country complexity.** Customer-specific and country-specific configuration — pricing tiers, tax treatments, billing type selection, approval thresholds

External Document © 2026 Infosys Limited

The provisioning and delivery layer

Beyond order entry, the Q2O process extends into service provisioning — the operational steps that convert a booked order into a delivered service. The managed services model for this layer applies a structured delivery architecture: standardized, language-neutral provisioning tasks handled from global delivery centers; language-dependent and customer-facing work managed from regional locations; and technically complex provisioning coordinated by onsite or nearshore specialists. The AI orchestration layer connects all three tiers, maintaining a single unified order tracker and routing exceptions to the appropriate resource automatically rather than relying on human coordination across delivery locations.

The critical design principle is that the tier structure follows the nature of the work — not the cost of the location. Functions that are rule-based and language-neutral sit offshore regardless of their apparent sensitivity. Functions that require cultural proximity, client relationship management or technical judgment remain close to the client. Conflating these two categories — either by over-retaining transactional work onshore or by pushing relationship-critical work to low-cost locations — is the most common failure mode in Q2O managed services design.

Order fallout is not an irreducible operational reality. It is the predictable output of a process that lacks pre-validation logic. The practitioner question is not how to manage fallout more efficiently — it is how to prevent the conditions that generate it.

Order to cash: Closing the revenue conversion gap

Order to Cash is the process that converts delivered service into recognized revenue. It is also where value destruction in B2B telecoms is most persistent, most measurable and — in the aggregate — most significant. Best-in-class operators achieve Days Sales Outstanding of 25 days or fewer. The industry average runs to 42 days. For a telco with two billion euros of B2B managed revenue, the working capital difference between these two positions represents hundreds of millions of euros in cash that is either locked in receivables or lost to billing errors, disputes and leakage.

The convergent billing problem

The O2C challenge in European telecoms is compounded by the legacy of product line-specific billing architecture. Operators who have grown through acquisition or who operate across fixed, mobile and digital product lines typically run multiple rating and billing systems in parallel — each optimized for its original product domain and none designed for the convergent offer structures that B2B customers now expect. The operational consequences are significant: separate bill cycles for different product components, manual reconciliation between usage data and billing records, split billing complexity for customers with mixed product portfolios, and a dispute surface that is far larger than a unified architecture would generate.

In telco B2B environments where this pattern has been assessed in detail, parallel bill cycle counts in the range of 20–28 are not unusual. Each cycle represents a reconciliation workload, a dispute risk and a potential leakage point. The aggregate of these workloads is a material operational cost that is almost entirely structural — it does not reduce as transaction volume grows, and it cannot be eliminated through incremental process improvement within the existing architecture.

External Document © 2026 Infosys Limited

Six capabilities that define the intelligent O2C model

- **Convergent rating and billing:** A single AI-native rating engine replaces fragmented bill cycles and the manual exception handling that accompanies them. It supports fixed-mobile-digital bundle pricing, split billing for complex B2B account structures and real-time convergent rating — delivering first-run billing accuracy at a level that makes disputes exceptional rather than routine. The target benchmark for well-implemented deployments is 99% or better first-run accuracy.
- **Continuous revenue assurance:** An AI-powered revenue assurance capability monitors usage-to-cash flows continuously, applying machine learning anomaly detection to identify leakage patterns in real time rather than through periodic manual reconciliation. The practical significance is that leakage which would previously accumulate undetected between reconciliation cycles is surfaced and corrected within hours. A target leakage rate below 0.5% of managed revenue is achievable where data quality is sufficient.
- **Autonomous dispute resolution:** Standard dispute types — misapplied pricing, incorrect usage rating, contract term mismatches — follow sufficiently predictable patterns to be resolved autonomously in the majority of cases. An agentic resolution layer handles these automatically, escalating genuinely complex disputes to specialists with full case context pre-populated. Dispute volume reductions of 40–50% within the first year of deployment are consistent with what practitioners observe in well-structured programs.
- **Predictive collections:** Collections effectiveness in B2B telecoms depends on the ability to differentiate treatment strategies by account type, risk profile and relationship value. Propensity models trained on payment history, product type, account behavior and external economic signals segment the collections queue and assign the right intervention automatically — concentrating human collector effort on the accounts where relationship management changes the outcome, and automating the remainder.
- **Intelligent cash application:** Cash application — matching incoming payments to outstanding invoices — is high-volume, rules-intensive and directly impacts both DSO and the accuracy of the receivables ledger. AI-driven remittance processing consistently achieves match rates of 97–99% in deployments where remittance data quality is adequate, against a manual benchmark that typically runs at 72–80%. The financial consequences of this improvement compound: unallocated cash decreases, effective DSO shortens, and the analyst capacity previously consumed by manual matching is freed for higher-value work.
- **Process intelligence as a continuous improvement engine:** A process mining layer across all O2C workflows monitors conformance, identifies deviation patterns and feeds root-cause insights back into agent configuration. This creates a feedback loop in which operational performance data continuously improves the intelligence of the agents that run the operation — without requiring periodic re-engineering to sustain the improvement trajectory

The O2C domain is where the compounding effect of intelligent operations is most visible. Each improvement in billing accuracy reduces dispute volume. Each reduction in dispute volume improves cash application accuracy. Each improvement in cash application accuracy reduces DSO. These are not independent gains — they reinforce each other.

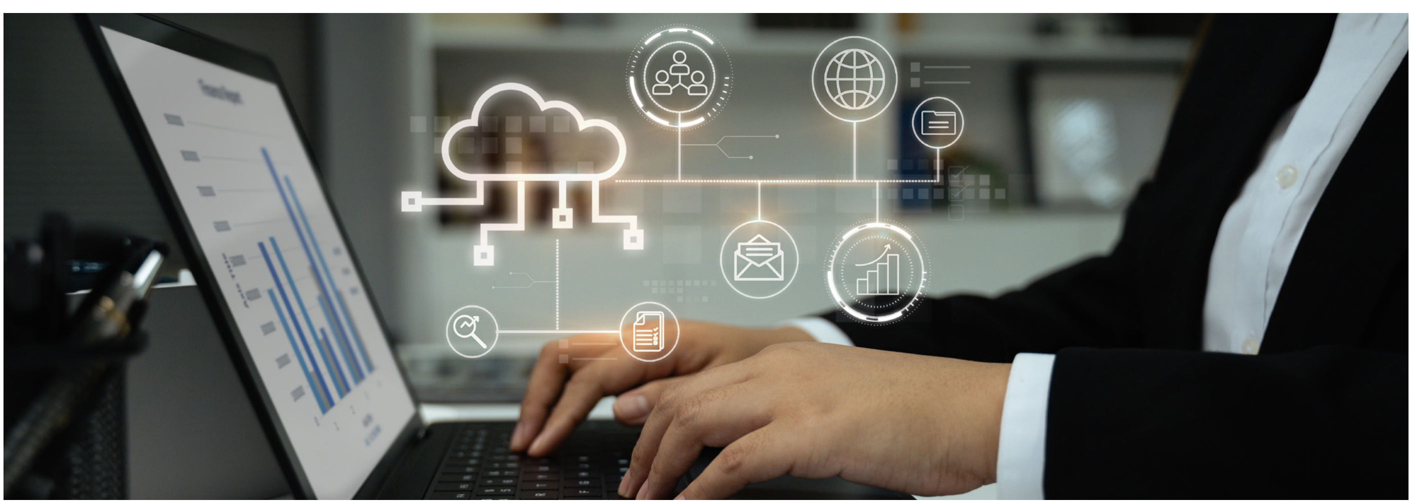

External Document © 2026 Infosys Limited

The intelligent operations blueprint

The capabilities described in the preceding sections are not useful in isolation. Their value depends on an operating model that integrates them — defining how AI agents, human specialists and continuous improvement mechanisms work together, and how accountability for outcomes is assigned and measured. This is where most managed services programs fail: not in the quality of the technology deployed, but in the design of the operating model around it.

Four interdependent design layers

- **Process standardization:** Before any AI can be effectively deployed, the actual process — not the documented one — must be understood, standardized and governed. In most European telco B2B environments, years of growth, acquisition and tactical workaround have produced a process landscape that bears little resemblance to any process map that exists on paper. The managed services provider's first substantive task is to map what is actually happening, identify the minimum standardization required to make automation safe, and establish the governance mechanisms that will maintain that standard as the operation evolves.
- **Agentic intelligence tiered by complexity**: The agent layer operates in two tiers. Rule-based tactical agents handle the high-volume, deterministic process steps that account for 70–80% of transaction volume — order validation, quote assembly, case routing, SLA monitoring, payment matching. These agents operate continuously, at scale, without the variation that characterizes human execution of repetitive tasks. Cognitive AI capabilities — large language model reasoning, anomaly detection, predictive modelling — handle the medium-complexity tier: exception classification, dispute analysis, renewal propensity scoring, revenue leakage identification. The two tiers together address the large majority of operational volume without human intervention.
- **Human-in-the-loop governance:** The residual volume — genuinely complex exceptions, high-value commercial decisions, regulatory determinations — is handled by specialist human teams working with AI-generated context and recommendations. The governance design is explicit about which decision types require human sign-off, at what threshold, and with what escalation path. This is not a residual — it is a deliberate design that ensures accountability is maintained where it matters most.
- **Continuous improvement through process intelligence:** Process mining across all operational activity measures conformance, surfaces SLA deviations and identifies the root causes of recurring exceptions. This feeds directly into agent configuration updates, approval matrix revisions and process redesign — creating a closed improvement loop that does not require periodic re-engineering to sustain.

External Document © 2026 Infosys Limited

The commercial model that makes it work

The commercial structure of an outcome-linked managed services engagement moves progressively from a base consumption fee — covering the platform, people and process foundation — to outcome-linked components tied to specific operational KPIs. As automation rates rise and the manual headcount profile reduces, the cost savings generated are shared between the telco and the managed services provider. This creates a structural incentive for the provider to drive transformation as rapidly as the operation can absorb it — the opposite of the incentive structure in a traditional FTE-based contract, where provider economics improve when headcount grows.

The managed services model that works is the one in which the provider's commercial interest and the telco's operational interest are genuinely aligned. Outcome-linked commercial structures are the mechanism that creates that alignment — not a pricing variation, but a fundamental redesign of the engagement.

Design principles for telco B2B managed services transformation

This paper closes with the design principles that, in the practitioner's experience, most consistently separate successful transformation programs from those that deliver incremental improvement but fall short of structural change. These principles apply regardless of the specific managed services provider, the technology platform or the process domain in scope.

Principle 1 — Start with the process, not the technology

The process assessment / process mapping what is actually happening, not what the process documentation says should happen — is not a prerequisite to be completed before the real work begins

Principle 2 — Design for the exception, not the average

The average case in any B2B telco process is straightforward. The challenge is the exception. An intelligent operating model is designed around how exceptions are identified, classified, routed and resolved — not around how the standard case flows.

External Document © 2026 Infosys Limited

Principle 3 — Separate the tier structure from the cost logic

Work should be placed in the delivery location that matches its intrinsic characteristics — not in the cheapest location that can plausibly execute it. Rule-based, language-neutral, high-volume work is appropriate for global delivery centers. Language-dependent, customer-facing work requires regional capacity.

Principle 4 — Build outcome accountability into the contract, not the governance

Outcome accountability that exists only in governance documents — review meetings, performance dashboards, SLA penalty clauses — is not effective. The only mechanism that reliably aligns provider's behavior with client outcomes is a commercial structure in which the provider's economics are directly linked to the outcomes being sought.

Principle 5 — Plan for portability from day one

Telco B2B technology landscapes change. CPQ platforms are replaced, BSS stacks are migrated, ERP systems are consolidated. A managed services model that encodes its intelligence into platform-specific configurations will need to be substantially rebuilt each time the underlying platform changes. An intelligent managed services model encodes its intelligence into a platform-agnostic agentic layer that connects to underlying systems via APIs rather than being embedded within them. When the system of record changes, the agent layer is rewired — not rebuilt.

Principle 6 — Measure from baseline, not from aspiration

Transformation programs are routinely compromised by the absence of a credible operational baseline. Without an objective measurement of current performance — order fallout rate, quote cycle time, billing accuracy, DSO — it is impossible to demonstrate improvement, impossible to calibrate outcome-linked commercial terms and impossible to make rational decisions about where to invest next. The baseline assessment is not a governance formality. It is the foundation on which every subsequent claim of improvement either stands or falls.

The managed services model that transforms telco B2B revenue operations is not defined by the AI capabilities it deploys. It is defined by the rigor of its process design, the honesty of its baseline measurement, the clarity of its outcome accountability and the commercial alignment that makes it in everyone's interest to succeed.

Data points and operational benchmarks referenced in this paper are drawn from process assessments of European telco B2B environments, TM Forum industry research on B2B operations maturity, Gartner O2C and fulfilment benchmarks, IDC Future of Operations research and Hackett Group O2C performance data. Operational data from assessed environments has been anonymized. Outcome ranges reflect observed results across multiple program deployments rather than single-point case study data.

Authors

Karthikeyan Shan
Supply Chain Practice Leader – EMEA & APAC, Infosys BPM

Karthikeyan leads the sales and fulfilment practice engagement for EMEA & APAC with more than two decades of experience in business and digital transformation. Having joined Infosys BPM in 2017, he has worked with multiple telecom clients over past 5 years in shaping their transformation and crafting the AI led transformation journey.

Karthikeyan has a MS in supply chain consulting from BITS Pilani and a degreed in production engineering from University of Madras

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.