VIEW POINT



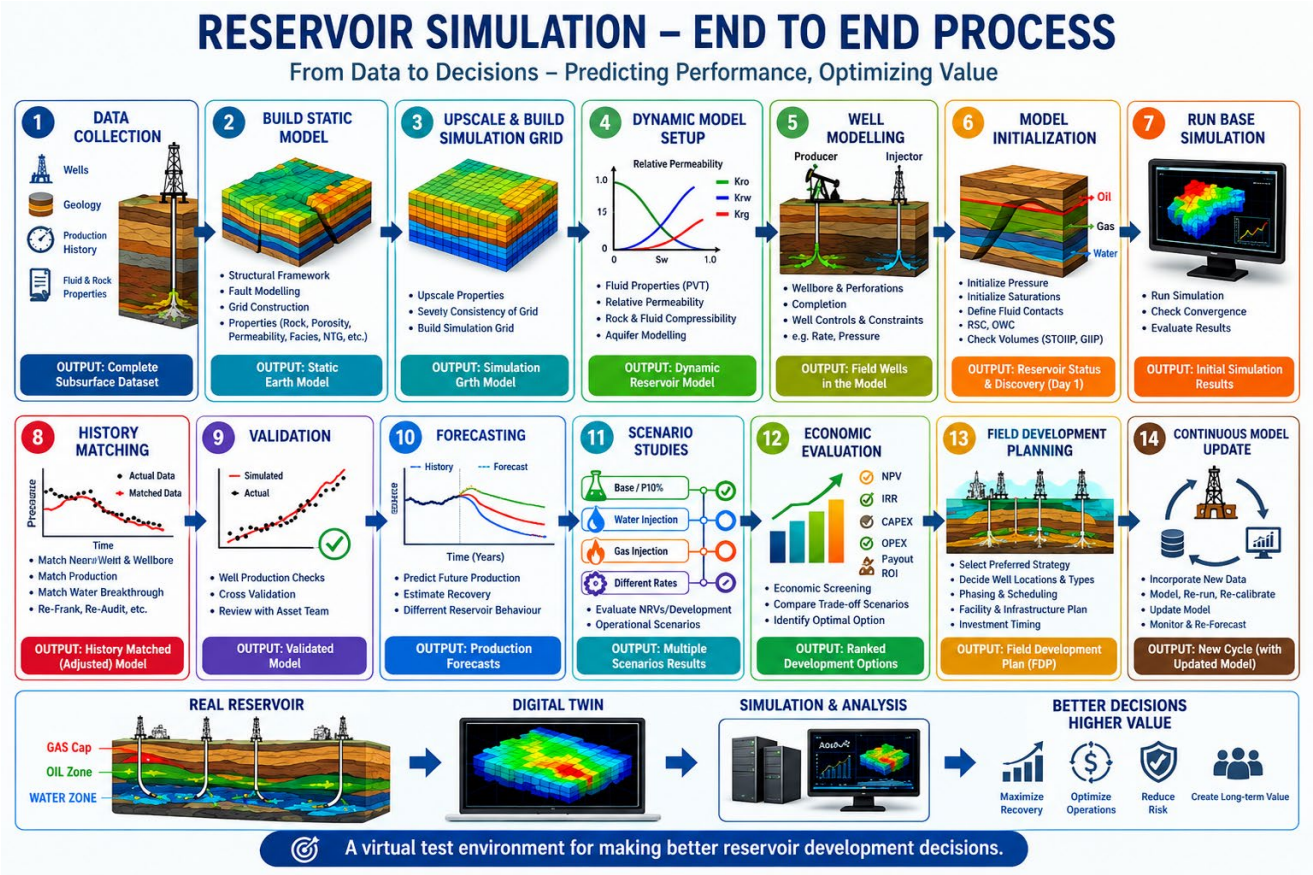
FROM MIGRATION BACKLOGS TO MODEL-READY CONFIDENCE: AGENTIC AI FOR RESERVOIR SIMULATION

Infosys | CONSULTING

What Reservoir Engineering Actually Involves

Before getting into migration, let's understand what reservoir simulation is and where it sits? Reservoir simulation is a computer run of a reservoir model over time to examine the flow of fluid within the reservoir and from the reservoir. Engineers calibrate the simulator against historic pressure and production data, a process called history matching, then use it to forecast performance under development scenarios such as new wells, injection schemes or stimulation.

Reservoir engineering, the discipline this sits inside, is the connective layer of the upstream value chain: the reservoir model's outputs feed a well model, a surface-network model, a facilities model, and ultimately an economic model that decides whether a field development goes ahead. It is, in effect, a virtual test environment — a way to de-risk a capital-intensive decision before an operator commits real money to it.



Reservoir simulation, end to end — from raw subsurface data to a living digital twin that supports development decisions.

That model — and the simulator it sits in — is rarely static. It moves between simulators, between platforms, and between companies. That movement is where this point of view begins.



External Document © 2026 Infosys Limited

Where Model Migration Becomes a Boardroom Conversation

Reservoir model migration stops being a back-office task and becomes a business decision in a handful of recurring situations, each already playing out across the industry today:

01 **Vendor platform change.** Operators standardize on a single simulator, or a vendor's product line is superseded, forcing every asset built on it to move. Multiple commercial platforms are in active use across the industry today, including SLB Eclipse, Halliburton Nexus, SRT's ECHELON and Rock Flow Dynamics tNavigator and many more, none of which is guaranteed to remain an asset's platform of choice indefinitely.

02 **Cloud and HPC modernization.** Operators are moving simulation workloads from on-premise clusters to containerized, cloud-native execution to capture pay-as-you-go economics and elastic scale-out for large history-matching ensembles — a shift documented in SPE technical literature on cloud-native reservoir simulation(IPTC-20237-MS)

03 **Mergers, acquisitions, and portfolio carve-outs.** When assets change hands, incoming or outgoing models typically need to be reconciled onto the acquirer's, or the standalone entity's, standard platform and workflows.

04 **New model builds for emerging use cases.** Particularly CCS and large-scale CO2 storage certification, where operators are standing up new static and dynamic models rather than migrating legacy ones. ADNOC and AIQ, for example, project their ENERGY ai agents will accelerate the building of detailed geological models for CO2 storage planning by up to 75%.

05 **A thinning bench of experienced engineers.** The IEA's World Energy Employment 2025 report finds that, across the energy sector in advanced economies, there are 2.4 workers nearing retirement for every entrant under 25 — a ratio the report expects to worsen through 2035. This figure is economy-wide rather than oil-and-gas specific, but it raises the stakes on any manual, tribal-knowledge-dependent activity, migration included.

06 **Recurring model refresh cycles.** Mature and brownfield assets require periodic model updates to support field development plans and reserves bookings — each a smaller-scale version of the same migration and validation problem.

For a firm like ours, every one of these is an entry point for a conversation about how migration is done, not just when.

The Pinch Points Nobody Has Fully Tapped

Automated conversion has already done its job. Deterministic translators and keyword dictionaries have compressed syntax translation from months to days, and the industry has absorbed that gain already. What has not moved is everything on either side of it. Across the migrations we have delivered, four gaps keep resurfacing, largely untouched by scripted tooling:

Understanding what a legacy model actually contains — dependencies are frequently undocumented, hidden in INCLUDE files, or held only in one engineer's memory.

Translating logic with no direct target-simulator equivalent — custom well-control strategies, field-development logic, and bespoke functions a keyword dictionary cannot map.

Proving the migrated model behaves like the original — today, thousands of well-by-well, rate, and pressure comparisons done largely by hand.

Explaining why it does not, when it does not — root-cause analysis that currently depends on scarce, senior reservoir-engineering time.

External Document © 2026 Infosys Limited

The table below reflects pain points and ranking the gaps that combine the highest frequency with the highest business impact:

Pain Point	Frequency	Business Impact	Priority
Unsupported functions & custom logic	High	Very High	Critical
Root-cause analysis of output mismatches	Very High	Very High	Critical
Manual QC and validation	Very High	High	Critical
Repeated learning not reused across projects	Very High	High	Critical
Field-development logic recreation	Medium	Very High	Critical

None of this is a scripting problem. Each requires reasoning across model structure, simulator behavior, prior migration experience, and engineering judgement — which is exactly why these gaps have resisted automation, and exactly why agentic AI, which can reason and adapt rather than execute a fixed rule set, is now relevant to closing them.



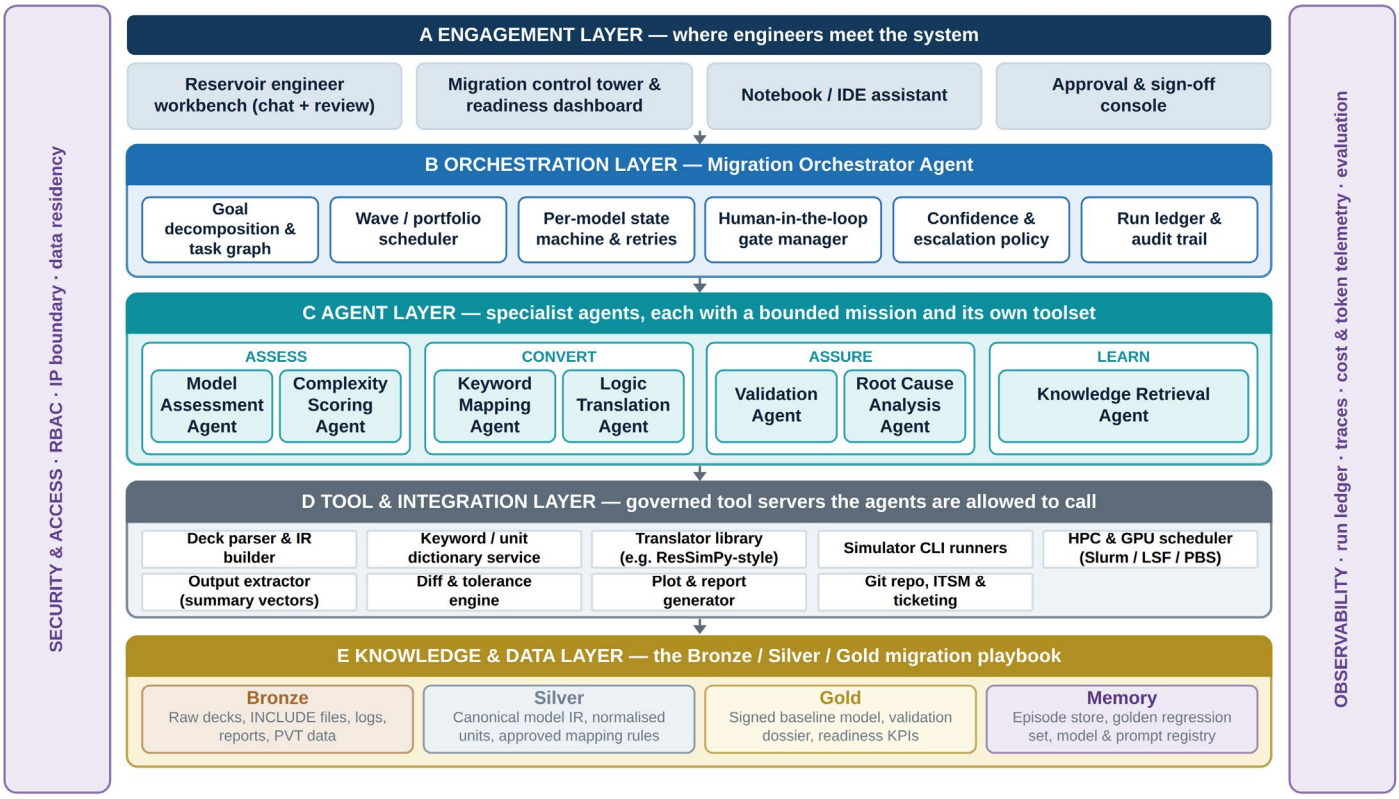
An Agentic Solution: How It Adds Value

Agentic AI differs from earlier automation in one crucial respect: instead of executing a fixed script, specialized agents reason about a construct, propose an option, evaluate it against evidence, and escalate only when they cannot resolve it themselves. This pattern is already delivering results at industry scale.

Example : ADNOC and AIQ's ENERGYai platform, built on agents trained on decades of ADNOC's proprietary subsurface data, is projected to compress development-planning cycles from one to two years down to weeks, and AIQ states the platform can lift operational efficiency in targeted upstream workflows by as much as 300%.

Reservoir simulation migration — with its large volumes of structured data and effort-heavy transformation and validation steps — is the kind of workflow this pattern of gains applies to.

We propose a reference architecture of five layers. An engagement layer is where engineers see, question, and approve the work. An orchestration layer owns the migration plan, sequences it into waves, and enforces mandatory human sign-off gates. An agent layer houses specialist agents grouped into Assess, Convert, Assure, and Learn. A tool and integration layer is the only route through which an agent can touch a real system, so every action is a call to a governed, permissioned tool rather than freehand generation of simulator syntax. And a knowledge layer is organized into Bronze (immutable, as-found), Silver (canonical, rule-governed), and Gold (signed, decision-ready) zones, so every migration adds to a reusable asset rather than starting from a blank page.



Reference architecture: five layers, with security and observability applied across all of them.

Seven specialist agents cover discovery, complexity scoring, keyword mapping, logic translation, validation, root-cause analysis, and knowledge retrieval. An eighth — the orchestrator — does no domain itself, and instead decides which agent acts next, and when a human must act instead. Agents propose; they do not quietly decide. Non-trivial logic translations, tolerance changes, and final sign-off all pass-through named human gates, and every parameter change, tool call, and approval is written to an auditable run ledger.

External Document © 2026 Infosys Limited

The value shows up in the same four places identified above

Dimension	Traditional Migration	Agentic AI-Enabled Migration
SME effort	Weeks of keyword-by-keyword mapping and manual cross-checks	Engineer time redirected to reviewing flagged exceptions and approving outputs
Validation	Full manual re-runs and comparisons	Structured, partially automated comparison that narrows effort to genuine discrepancies
Traceability	Decisions live in individual engineers' notes	Every mapping, flag, and decision logged with a clear audit trail
Reusable knowledge	Each migration starts from a blank page	Mapping rules and QA patterns accumulate as a reusable asset

Migration is also, deliberately, the proving ground for something larger. The same agents that learn to reconcile a legacy model's structure and dependencies are, by design, learning to reconcile raw subsurface data itself — PVT reports, logs, and core data arriving from different vendors and formats. Once that reconciliation is trustworthy, the natural extension is not just migrating an existing model but assembling a static model directly from source data, with engineers reviewing rather than manually importing. This point of view is intentionally scoped to migration; model generation is the logical next conversation.



Exit Note

Reservoir simulation migration has never really been about simulator syntax. It has been about time — SME time, elapsed time, and the time it takes an organization to trust a new number. Agentic AI does not remove the reservoir engineer from that picture. It removes the parts of migration that never needed one: chasing keyword equivalence, comparing thousands of curves by eye, and re-explaining last year's decisions because nobody wrote them down.

For operators weighing a simulator change, a cloud migration, a portfolio integration, or simply a backlog of models due for refresh, the question worth asking is not whether migration can be made faster, but whether the next migration can make every subsequent one cheaper.

Sources

1. SLB Energy Glossary — "reservoir simulation" definition — https://glossary.slb.com/terms/r/reservoir_simulation
2. SLB Oilfield Review — The Defining Series: Reservoir Engineering — https://www.slb.com/resource-library/oilfield-review/defining-series/defining-reservoir-engineering
3. Rescale — "Reservoir Simulation Moves to the Cloud" (commercial simulator landscape) — https://rescale.com/blog/reservoir-simulation-moves-to-the-cloud/
4. IPTC/OnePetro — "Transitioning a Legacy Reservoir Simulator to Cloud Native Services" — https://onepetro.org/IPTCONF/proceedings/20IPTC/3-20IPTC/D031S086R003/154791
5. AIQ Newsroom — "ADNOC and AIQ Developing First-of-a-Kind Agentic AI Solution for Global Energy Transformation," Nov 4, 2024 (1–2 years to weeks; up to 75% faster geological model building for CO2 storage) — https://aiq.ae/newsroom/news-and-press-releases/ADNOC-and-AIQ-Developing-First-of-a-Kind-Agentic-AI-Solution-for-Global-Energy-Transformation
6. IEA — World Energy Employment 2025, Executive Summary (2.4:1 retirement ratio, advanced economies) — https://www.iea.org/reports/world-energy-employment-2025/executive-summary
7. AIQ — ENERGYai product page ("boost operational efficiency by up to 300%") — https://aiq.ae/products/AgenticAI/energyai



About the Authors



Pulkita Rohilla
Principal, Infosys Consulting, Mumbai, India

Pulkita Rohilla is an Energy Consultant with over 15 years of experience leading upstream oil & gas operations and driving AI-powered digital transformation across the energy sector. A woman in energy with hands-on field leadership across both onshore and offshore operations — including vessel operations and 100+ field jobs — she brings a rare combination of ground-level operational command and applied AI expertise, honed over a career built on both rigs and reservoirs. She holds a B.Tech in Petroleum Engineering from PDEU Gandhinagar and an MBA (Gold Medalist) from IIM Indore.

Her broader portfolio spans predictive asset monitoring, production forecasting using time series techniques, computer vision-based automation, and sentiment-driven crude price forecasting — reflecting her leadership in embedding AI methods into core energy-sector decision-making and her continued push to modernize how the industry operates.



Raghavendar Kannan
Engineering Lead, Infosys Limited, London, UK

Engineering Lead SME, product analyst with 14+ years of experience across KBE, IoT, and digital transformation, currently leading technical tracks and product analysis for BP PLC, UK. Expert in Digital Twin, ALIM, IDMS, APM, EIM migrations, and AI-powered solutions using ETL/ELT and Power Platform. Domain expertise spans Aerospace, Oil & Gas, Automotive, and Manufacturing; recognized as Infosys MVP and Microsoft Certified Cloud Expert.



Soumitra Bhaskar Nande
Senior Consultant, Infosys Consulting, Pune, India

Soumitra Nande is a Reservoir Engineer, Data Scientist, and Consultant with over 13 years of global upstream experience spanning reservoir engineering, digital transformation, and business consulting. He currently leads strategic requirement definition and Agile delivery for large-scale digital migrations and BI ecosystems in the energy sector, with prior experience across reservoir engineering assignments in the Gulf of Mexico, North Sea, and Australia. He holds a Ph.D. in Petroleum Engineering, specializing in deep learning-based reservoir characterization, and works at the intersection of subsurface physics and machine learning to modernize upstream decision-making. He is an active researcher and published author in SPE and Scopus-indexed journals, with a continued interest in applying AI to solve real-world subsurface challenges.

Reviewers



Pranav Gupta
Associate Partner, Infosys Consulting,
London, UK



Ramkumar Gunasekaran
Senior Principal, Infosys Consulting,
London, UK

ABOUT INFOSYS CONSULTING

Infosys Consulting is a next-generation consulting partner that bridges strategy and execution. With an AI-first mindset, deep industry knowledge, and the combined strengths of business and technology consulting, it helps enterprises turn bold vision into tangible outcomes, faster, smarter, and at scale. Infosys Consulting is helping some of the world's most recognizable brands transform and innovate. Our consultants are industry experts that lead complex change agendas driven by disruptive technology. With offices in 20 countries and backed by the power of the global Infosys brand, our teams help the C-suite navigate today's digital landscape to win market share and create shareholder value for lasting competitive advantage.

For more information, contact consulting@infosys.com

Infosys | CONSULTING

© 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.