VIEW POINT

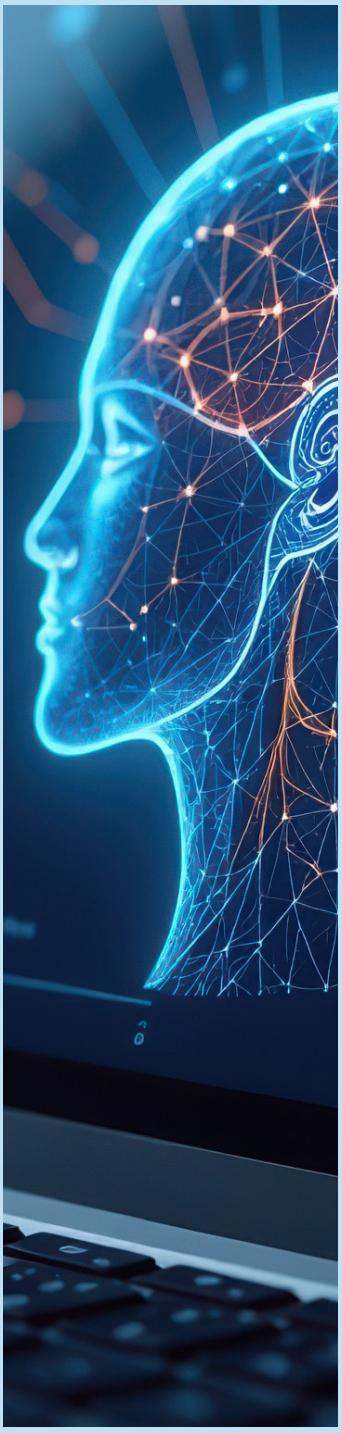

RETHINKING OPERATING MODELS FOR THE HUMAN AND AGENTIC AI WORKFORCE

Infosys | CONSULTING

Executive Summary

AI agents don't fit into yesterday's operating models - they expose their limitations.

Agentic AI, otherwise known as autonomous digital actors capable of evaluating, deciding and executing delegated actions, has moved from an experimental concept to a reality that is hard to escape. This shift isn't about enhancing how work is done but rather redefining who (or what) performs it, how decisions are made and how value is ultimately created.

At its core, this is an economic shift. The cost of intelligence is collapsing, which means that decisions and actions can be carried out at almost no incremental cost. Yet most operating models are still built for a world where human judgement is limited and expensive. This assumption no longer holds and this is where the tension lies.

Most organisations are not constrained by access to AI. They are constrained by their ability to operationalise it.

Against this backdrop, three priorities stand out:

1. **Redesign, don't retrofit**
   Operating models must be rethought end-to-end - across people, processes, technology and governance to enable human and AI collaboration at scale.
2. **Turn capability into performance**
   Early-moving organisations are already realising 25–50% cost efficiencies and materially faster cycle times with outcomes varying by function, maturity and governance design. The differentiator is no longer adoption; it is active execution.
3. **Close the agency gap**
   The primary source of lost value is no longer technical limitation but organisational hesitation - the gap between AI's capability and the level of agency it is given.

This paper sets out a practical, experience-led approach to redesigning operating models for an era where intelligence is abundant, execution is continuous and advantage depends on speed.

Introduction: Why Agentic AI Demands a Rethink of Operating Models

Organisations are at an inflection point.
Agentic AI represents a shift from hierarchical execution models towards more **networked ways of working, where humans and intelligent systems operate alongside each other, often in parallel.**

At the same time, the underlying economics of organisations are changing.

Historically, every additional decision came with a human cost. As a result, operating models were designed to limit how many decisions were made. That dynamic is changing. The cost of executing decisions is falling and the real constraint is no longer intelligence, but how quickly organisations can absorb and act on it.

Traditional operating models, built for predictability and control, are struggling to keep pace with real-time decision environments, continuous execution and more fluid collaboration between people and machines. In selected enterprise functions where agentic patterns have been applied end-to-end, organisations are beginning to observe measurable impact, including decision latency falling by 30–40% and operating costs by over 20–25%. But where AI is layered in without proper integration, the opposite can happen, with increased rework and greater compliance risk.

This transformation affects every dimension of the enterprise:

1. **People:** From execution to oversight, judgment and orchestration
2. **Processes:** From linear workflows to dynamic interaction systems
3. **Technology:** From support layer to operational backbone
4. **Governance:** From accountability of people to accountability of decisions
5. **Location:** From physical constraints to digitally distributed execution

This goes well beyond optimisation – it changes how organisations actually function.

External Document © 2026 Infosys Limited

Drivers for Change:Why Operating Models Cannot Remain Static

Six structural forces are accelerating change:

1. **Structural misalignment**
   Legacy workflows weren't built for autonomous decision-making. As a result, organisations are seeing duplication, delays and gaps in control.
2. **Nonlinear growth potential**
   AI breaks the traditional link between cost and output, enabling scale without simply increasing headcount.
3. **Integration complexity**
   Most operating models are still fragmented. Without alignment across teams and systems, AI struggles to deliver real impact.
4. **Regulatory and ethical pressure**
   Expectations are changing. Decision-making systems increasingly need to be transparent, explainable and auditable.
5. **The cost imperative**
   The focus is shifting from labour arbitrage to the intelligent automation of knowledge work.
6. **The experience standard**
   Real-time, personalised experiences are no longer a differentiator – they are becoming the baseline.

The Agency Gap: Where Enterprise Value Is Constrained

Despite sustained investment into AI, many organisations aren't seeing proportional returns. The issue isn't capability but how that capability is being used.

At the centre of this is what you might call an 'agency gap': the difference between what AI systems can do and what organisations are actually willing to let them do.

Capability is moving fast. Permission isn't.

The result is predictable. High-capability systems are being deployed into environments that still rely on sequential approvals, human bottlenecks and static governance.

Organisations are effectively paying for machine-speed capability while still operating at human-speed throughput.

The Economic Consequence: Stranded Value

In practice, this gap shows up as lost value - slower decisions, weaker returns on AI investment, duplicated oversight and slower response where timing matters most.

If the cost of intelligence is close to zero, then limiting its use is effectively a self-imposed tax on performance.

This shift does not remove human accountability or weaken risk management; it changes where control is exercised - from reviewing individual outcomes to governing the systems, policies and logic that produce them.

Closing the Gap

Closing the agency gap requires a fundamental shift in how organisations operate - from reviewing individual decisions to setting clear boundaries, from controlling execution to governing the underlying logic, and from managing tasks to orchestrating systems.

In an agentic enterprise, performance is no longer defined by how many decisions you can review but by how well the systems that make those decisions are designed.

Who Owns the Agent? The New Organisational Question

As AI moves from a tool to an active participant, a more practical question starts to emerge: **Where do these systems sit within the organisation and who is accountable for them?**

In response, three operating models are taking shape. Some organisations are centralising AI through a dedicated centre of excellence, which brings strong governance and control but can slow execution and create distance from where value is realised.

Others are pushing ownership into the business, gaining speed and closer alignment to use cases but often at the cost of fragmentation and duplication.

Increasingly, a hybrid model is emerging. In practice, this model requires clearly defined accountability for agent behaviour, decision outcomes and lifecycle management, particularly in regulated environments where legal responsibility cannot be delegated to systems.

Core standards, guardrails and governance sit centrally, while execution is distributed closer to the business. This enables organisations to maintain control without sacrificing speed. It is fast becoming the model that scales.

Control and speed no longer need to be trade-offs. With the right governance, organisations can achieve both.

External Document © 2026 Infosys Limited

Kinetic Governance: From Approvals to Policy-Driven Execution

Traditional governance is built around interruption.

Work stops, a human reviews, approval is given and execution resumes. That model begins to break down in environments where decisions and actions are happening continuously.

What's emerging instead is a more dynamic approach to governance - one that's less about stopping decisions and more about shaping how they're made. This model is called Kinetic Governance. Leaders set clear boundaries and policies, agents operate within those guardrails and performance is tracked in real time.

The shift is subtle but important: governance moves away from approving individual actions and toward governing the logic behind them. In effect, you're no longer reviewing each decision - you're designing and overseeing the system that produces them.

This is what allows organisations to scale in practice.

Safeguards in Policy-Driven Execution
In practice, kinetic governance depends on a set of non-negotiable safeguards:

- Clear thresholds for human intervention and override
- Continuous logging and traceability of agent decisions
- Periodic testing and revalidation of decision logic
- Defined escalation paths when outcomes deviate from policy

These mechanisms ensure execution scales without compromising auditability, accountability or regulatory compliance.

In a world where intelligence is abundant, advantage comes down to how quickly you can act, not how many decisions you can review manually.

The Emerging Role of HR: From People Function to Workforce Architect

A more fundamental shift is starting to take shape. AI systems aren't just supporting work - they're beginning to take on characteristics we would typically associate with employees. They execute tasks, make decisions, interact with stakeholders and operate within defined roles.

This is changing how we think about the workforce itself.

In practice, a few models are emerging. In some organisations, these systems are treated as part of the workforce with HR taking on a custodial role. In others, ownership sits within business functions, embedded directly into day-to-day operations. And in some cases, they're managed more like shared platforms - enterprise capabilities that sit alongside core technology.

Whichever model is chosen, the implication is the same: workforce management is expanding beyond people alone. It now includes how organisations deploy, govern and scale cognitive capability.

This creates a very different kind of leverage. Capability can be scaled instantly, without the usual lag of hiring and training. It doesn't degrade over time and it doesn't walk out the door.

This shift increases the importance of deliberate reskilling, role redesign and human capability investment, ensuring that automation augments

organisational resilience rather than displacing institutional knowledge.

In effect, organisations are building a more persistent form of institutional capability - one that sits alongside and increasingly complements human expertise.

Design Criteria: How to Determine Agent Ownership

There is no single model that fits every organisation. The right structure depends on a range of factors - from risk and regulatory exposure to the level of autonomy given to AI; the complexity and criticality of processes; cross-functional dependencies; organisational maturity and data sensitivity.

The model isn't fixed – it evolves over time. Ownership shifts with confidence, scale, and strategy - becoming more distributed as organisations pursue growth, and more centralised as risk and control requirements increase.

External Document © 2026 Infosys Limited

From Workflows to Collaboration Systems

For the first time, employees are working alongside non-human counterparts and that is a significant behavioural change.

The emphasis moves away from doing the work directly to delegating it, orchestrating it and overseeing how it runs. Work becomes less about managing individual tasks and more about shaping the systems behind them.

That brings a different set of capabilities to the forefront - knowing when to trust AI and when to intervene, designing effective human-AI interactions and taking accountability at a system level rather than for individual actions.

This is where the real shift lies.

The Adoption Constraint: Culture, Not Technology

The main barrier to scaling AI is not capability - it is organisational readiness.

In practice, three challenges recur: a lack of trust in the systems, uncertainty around how roles change and a lack of clarity from leadership.

When ownership, accountability and the operating model aren't clearly defined, organisations tend to stall in pilot mode.

Most organisations are not failing to build AI - they are failing to reorganise around it.

The Apprenticeship Paradox: Preserving Future Capability

As agents take on more analytical and repeatable tasks, a less obvious risk emerges.

The erosion of the learning pathways that have traditionally developed future leaders.

Historically, capability was built through doing, through repetition, exposure and gradual progress. That pathway is narrowing.

Responding to this requires deliberate effort. In some cases, that means creating ways to learn by interrogating AI decisions - understanding not just the outcome but how it was reached. In others, it means using simulation environments to build and maintain core capabilities that might otherwise be lost. Without that kind of intervention, there is a risk of optimising for the present while weakening future capability.

Addressing this is a leadership and operating model design responsibility, not simply a byproduct of technology adoption.

A Disciplined Approach to Operating Model Redesign

Transformation needs a clear method, not just experimentation. In practice, that tends to follow a few core steps: aligning on strategic intent, understanding the current state, designing the target model, building a roadmap and business case, and then implementing, governing, and adapting over time.

What matters is each decision ties back to measurable outcomes - whether that's speed, cost, quality or risk.

Critically, this process is iterative, and operating models must adapt as confidence in agent autonomy increases and regulatory expectations evolve.

Advantages of a Reimagined Operating Model

When designed intentionally, agentic operating models can deliver faster, higher-quality decisions, greater efficiency at scale, and more consistent experiences for both customers and employees. They also create conditions for continuous innovation, with governance built in rather than layered on top, and ultimately deliver more measurable financial returns.

At that point, operating model design becomes a primary driver of enterprise value.

Executive Imperatives

For leaders, a few priorities stand out: setting a clear transformation narrative, establishing governance for human-AI systems, investing in capability and culture, designing for adaptability and adopting a phased and disciplined approach to execution.

The differentiator is no longer access to AI, but the ability to operationalise it.

External Document © 2026 Infosys Limited

Conclusion — Design Is The Advantage

Agentic AI is not a future concept – it is already here.

The organisations that pull ahead will not simply be those that adopt it fastest but the ones that reshape how they operate around it.

This goes beyond a technology upgrade. It affects how organisations think, how decisions are made and how work actually gets done.

The defining challenge isn't deploying AI - it's designing around it.

In a world where intelligence is increasingly abundant, AI on its own won't be the differentiator. How an organisation is structured to use it will be.

The organisations that recognise this early and act with intent will not just keep up. They will set the pace.

External Document © 2026 Infosys Limited

Author

Hemang Patel leads global engagements to redesign operating models for the age of agentic AI, helping organisations transition from human-led execution to scalable human–machine co-execution.

As Head of Digital Operations & Service Management (EMEA) within Technology Transformation, he works across operating model design, governance and AI-driven execution, helping senior leaders turn technological capability into measurable enterprise performance.

His work centres on the structural challenges of integrating autonomous systems into the enterprise - particularly decision rights, ownership models and emerging approaches to digital workplace design, with a consistent focus on enabling organisations to operate at machine speed without compromising human judgement.

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.