Synarchy

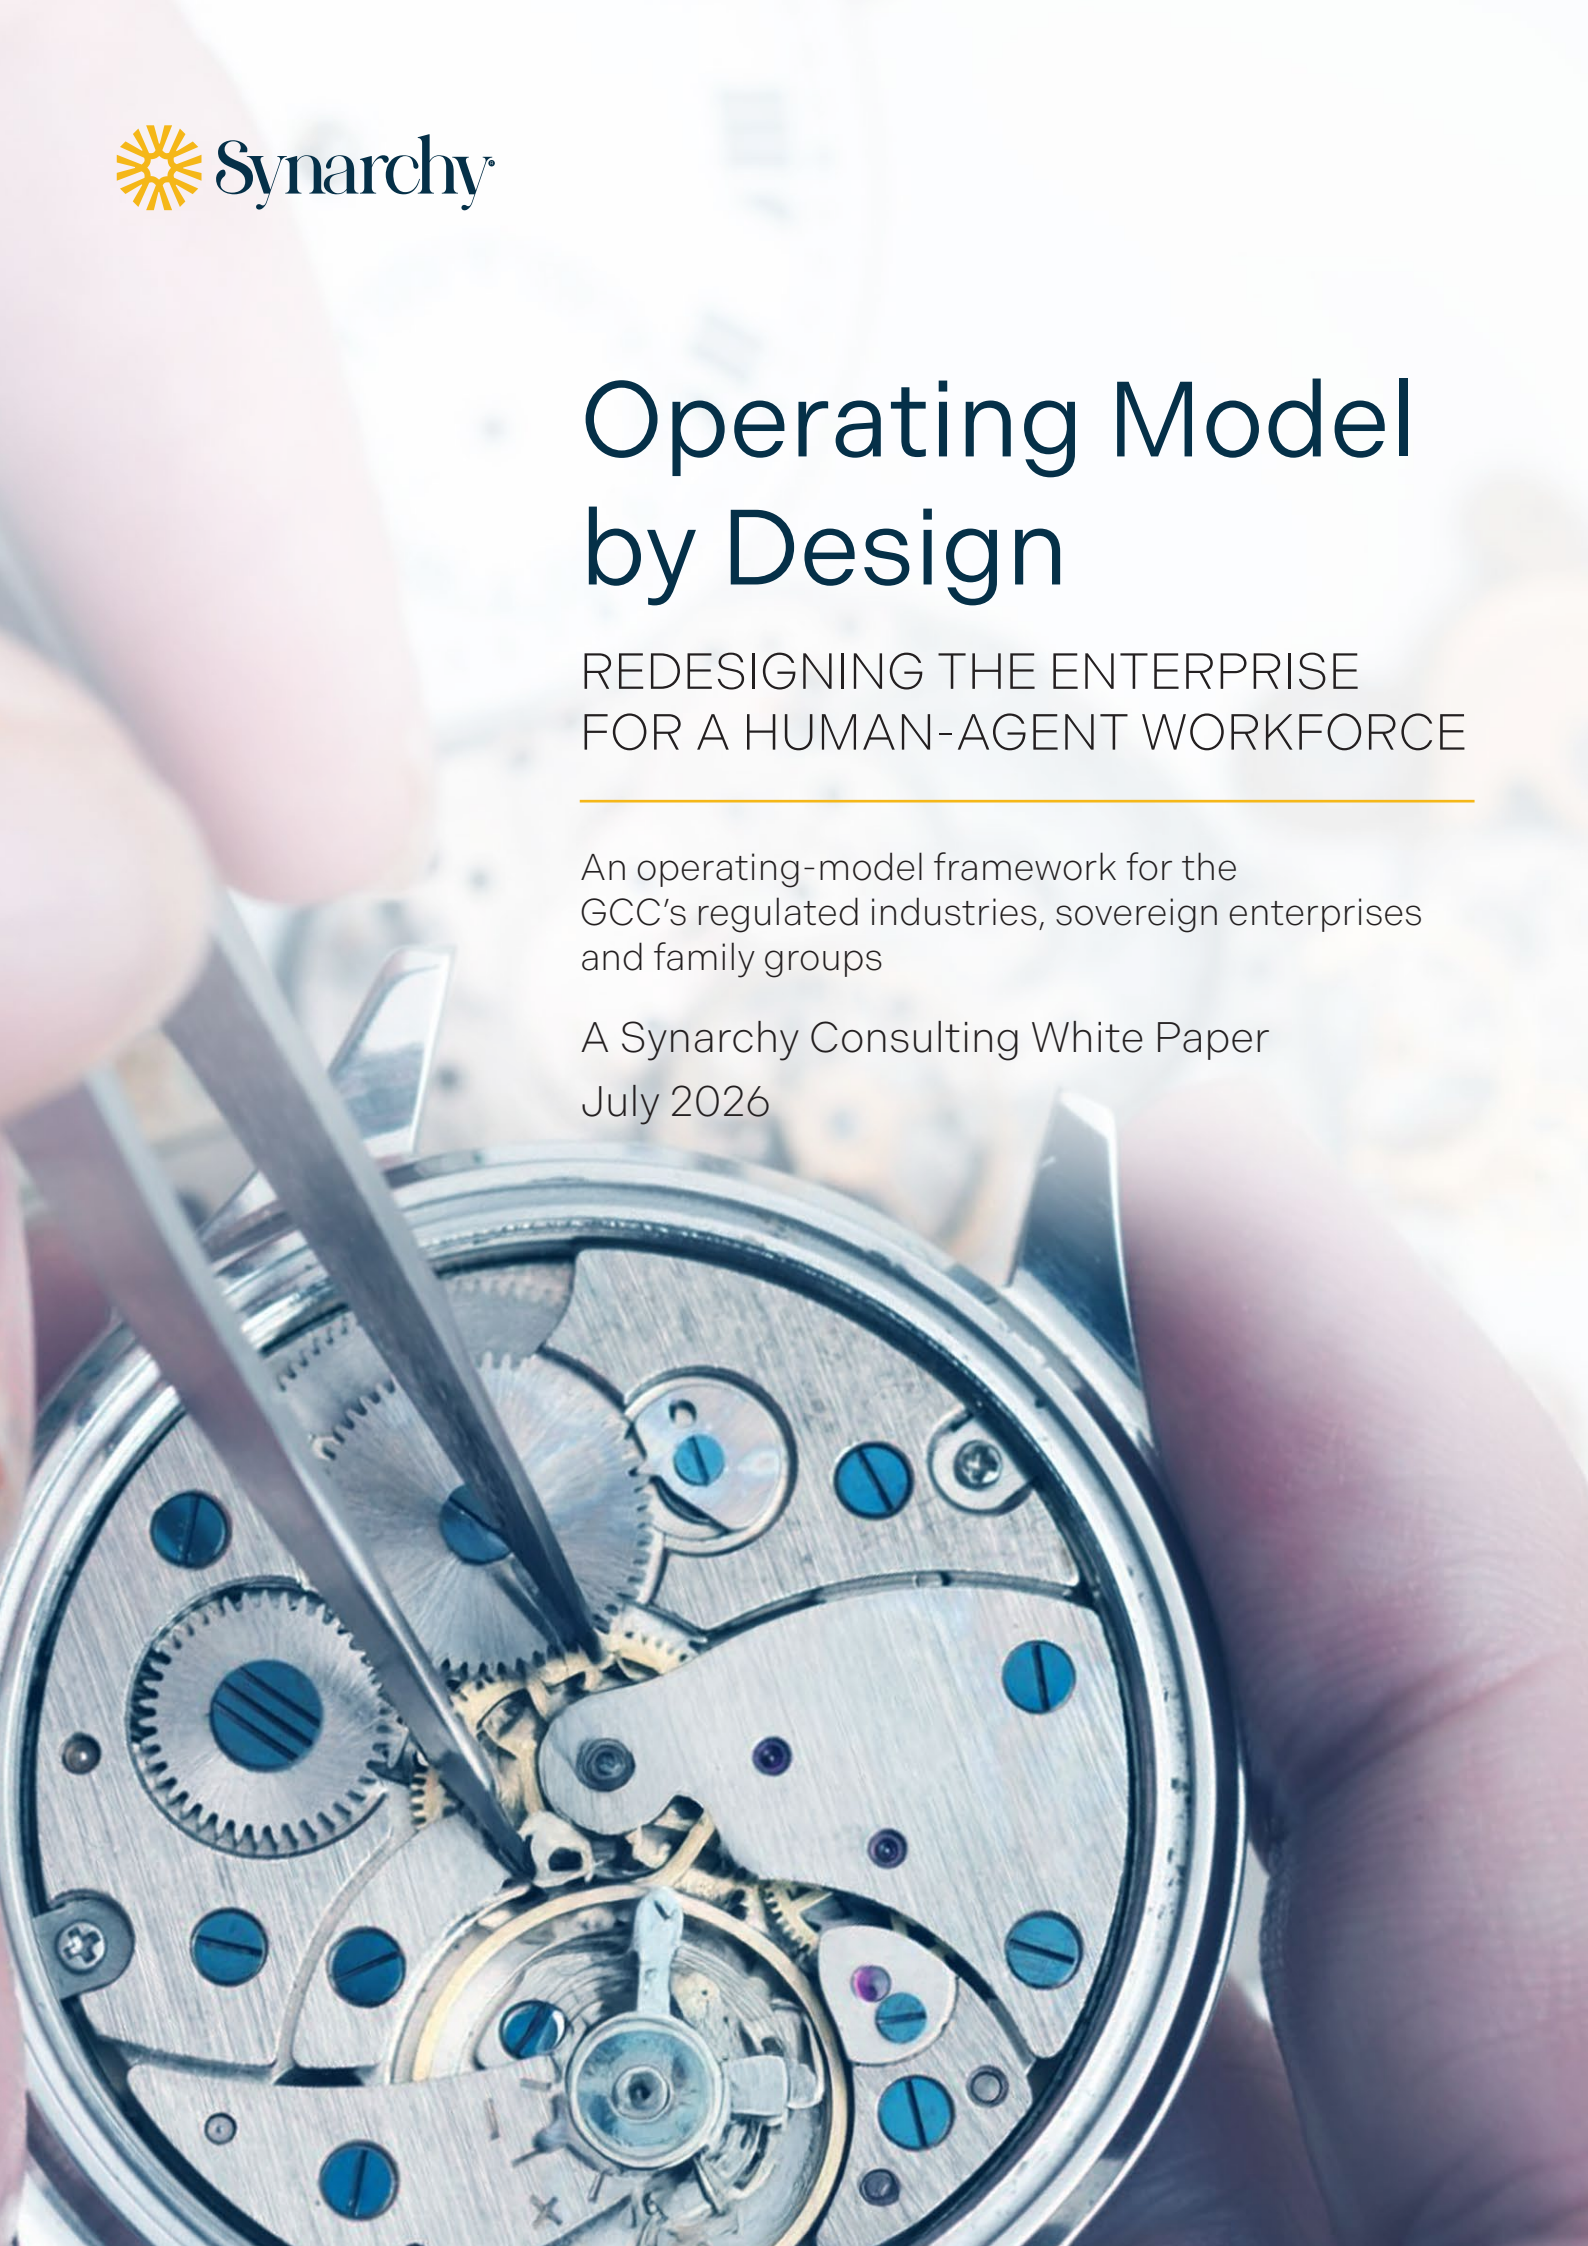
Operating Model by Design

REDESIGNING THE ENTERPRISE FOR A HUMAN-AGENT WORKFORCE

An operating-model framework for the GCC's regulated industries, sovereign enterprises and family groups

A Synarchy Consulting White Paper

July 2026



Executive Abstract

Many features of the traditional operating model, its layers, handoffs, approval chains and spans of control, were shaped partly by the cost and limits of human coordination. They were also shaped by specialization, scale, regulation, risk, segregation of duties and capital allocation, and those reasons do not go away. Agentic AI, software that can plan, reason and carry out multi-step work, changes the coordination constraint in parts of the enterprise without automatically changing the structures built around it. As agents take on more execution, accountability has to become more explicit: every material outcome needs a clear, traceable owner when an agent acts.



This helps explain why many pilots stall on the way to scale. Technology is rarely the whole of it. More often, enterprises install agents into an operating model whose decision rights, controls and ownership were designed for people, and the operating model becomes one of the factors limiting scale. In the GCC's regulated industries, sovereign enterprises and family groups, that limit is sharpest where it matters most: who answers for a decision once an agent has made it.

The practical response is to redesign the operating model deliberately. Operating Model by Design anchors on the enterprise's value logic, set by strategy, and redesigns how work is structured, decided and owned for a workforce of people and agents. At its center is the Accountability Gap, a portfolio view of where work has been delegated to an agent without a traceable human owner of the outcome. Two plain management views sit alongside it: how far delegation has gone in material work, and how much of that delegation is traceable to an accountable person.

Exhibit 1. As agents take on more execution, accountability must become more explicit

The argument of this paper, in brief

01 The org chart is only one part of the operating model

The org chart shows reporting lines; the operating model also defines how value is created, work is organized, decisions are made and accountability is exercised.

02 AI removes some of the coordination friction organizations were built to manage

Coordination that existed because execution was expensive loses value; controls and judgement that stay valuable when execution is cheap do not.

03 Agentic execution changes where decision rights need to sit

Decision rights drawn for a workforce of people do not fit work an agent can execute end to end.

04 As agent autonomy rises, accountability has to become explicit

When an agent can act, a clear, traceable owner of the outcome becomes a first-order requirement, most of all in regulated and family-owned settings.



THE ARGUMENT OF THIS PAPER, IN ONE LINE

Many features of the traditional operating model were built around the limits of human coordination. As coordination and execution become easier to scale, accountability becomes more important.

The paper is written for the CEO, the chief operating officer, the board and the owners of the region's private enterprises, family groups and corporates. It closes where design should begin, with one question: of the material work agents now perform, how much has a clearly accountable human owner?



1 Agentic AI is advancing faster than the operating model around it

Pilots are moving to production faster than the operating model is changing

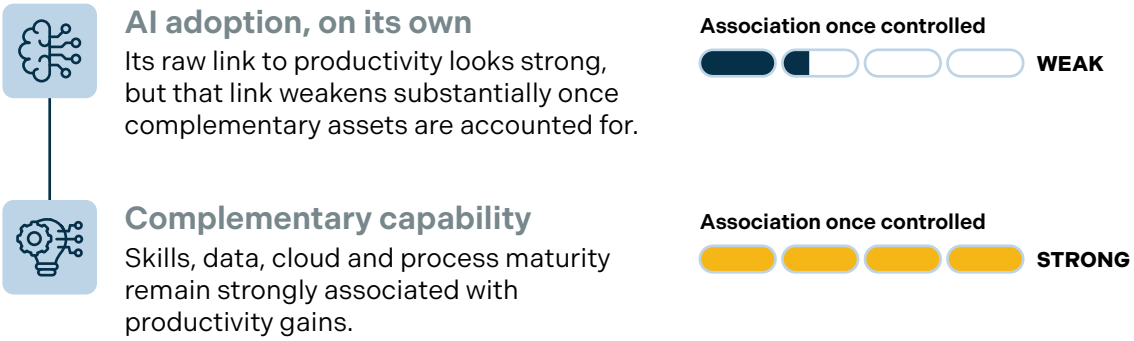
Across the region, enterprises are moving AI pilots from controlled environments into production, and many struggle at that step. McKinsey's 2025 survey of senior executives and board directors in the GCC found that 84 percent of respondents said their organizations use AI. Only 31 percent said it had been scaled or fully deployed (McKinsey and GCC BDI, 2025). The figures are self-reported. Technology and talent can hold back an individual deployment, but scaling also requires the surrounding operating model, its decision rights, ownership, controls and work design, to adapt to agents that act.

Part of the reason is structural. Transaction-cost theory holds that a firm organizes work internally when doing so is cheaper than coordinating it through the market (Coase, 1937). One implication is that layers, spans of control, handoffs and approval chains were shaped partly by the cost and limits of human capacity to decide and coordinate, alongside specialization, regulation, risk and control. Agentic AI lowers the cost of some of that coordination. Where software takes on coordination and execution, some of the assumptions behind the structure no longer hold, while others, such as segregation of duties and regulatory checkpoints, remain as valid as before.

Evidence on what produces value supports this. Across nine advanced economies, the OECD takes account of a firm's complementary assets: its skills, data, cloud and process maturity. Once it does, the raw statistical link between AI use and productivity weakens substantially, while those assets remain strongly associated with productivity (OECD, 2024). Scaling AI is therefore an operating-model question as much as a technology rollout. Value comes from the operating capability around the model, and in many enterprises that capability is still built for the older constraint.

Exhibit 2. Complementary capability, not AI adoption alone, tracks productivity

Association with productivity, once other factors are accounted for (OECD, 2024)



Illustrative of the direction of the OECD (2024) findings; qualitative, not to scale.

Source: *OECD (2024), across Belgium, Denmark, France, Germany, Israel, Italy, Japan, Korea and Switzerland. Once complementary assets are accounted for, the raw AI–productivity association weakens substantially while complementary assets remain strongly associated with productivity.*

Earlier firm-level data point the same way. An analysis of older World Bank Enterprise Survey microdata found adoption predicted by research intensity, managerial experience and digital readiness rather than by access to the technology (Udomsaph, 2025). Access to the underlying technology is now widespread. The capability to turn it into value is not, and it is unevenly distributed.

The more important question is which assumptions in the inherited operating model no longer hold as agents take on more execution, and what must be redesigned in their place. An operating model built for the older constraint and left unchanged does not fail loudly. It holds the enterprise at the pilot stage, busy and unable to scale. Four tensions matter most.



2 AI changes coordination, but not every control should disappear

AI changes the cost of coordination, but it does not remove the need for control and accountability.

1. The org chart is only one part of the operating model

The org chart shows reporting relationships: who manages whom. The operating model also defines how value is created, how work is organized, how decisions are made and how accountability is exercised. Structure is shaped by many forces, and one of them is the cost of coordination: firms organize work internally when internal coordination is cheaper than market exchange (Coase, 1937). Research on management practices shows that how a firm is organized and managed is itself a productive technology. It explains a large part of productivity differences between firms and countries (Bloom, Sadun and Van Reenen, 2012; 2016). When the cost of coordination changes, the operating model should be redesigned deliberately rather than simply automated, and that redesign reaches well beyond the boxes on the chart.

2. AI removes some of the coordination friction organizations were built to manage

Many enterprises are organized, in part, around coordination friction: layers, handoffs and approval chains that once reduced coordination cost and can themselves become the cost when software takes over the coordination. Transaction-cost logic implies that technology which lowers coordination and search costs reduces the value of structures built only to manage those costs (consistent with Coase, 1937). The distinction that matters is between coordination that exists because

execution was expensive and control or judgement that stays valuable when execution is cheap. Some layers exist to manage risk, preserve challenge, segregate responsibilities, assure quality or bring specialist judgement, and those do not lose value. In the GCC's regulated industries some friction is also mandated. Anti-money-laundering monitoring requires documented human review of alerts and does not transfer accountability to the system that generates them (BCBS, 2020). Supervisors expect human approval or dual authorization for higher-value automated transactions (FSB, 2026). Structures built to manage coordination cost may lose value as that cost falls; the checkpoints a regulator requires, and the controls that protect judgement, remain.

Exhibit 3. AI erodes coordination friction; controls that protect judgement and regulatory checkpoints stay

What to automate, what to keep, and how to design around each

AI can reduce	Regulation / ownership preserves	Design implication
Manual coordination	Human review of alerts	Automate preparation; keep the accountable review
Information gaps	Human-in-command	Improve information flow; keep the decision with a person
Handoffs and rework	Auditability and redress	Redesign the workflow; keep traceability
Queuing and approval delay	Dual authorization	Automate routing; keep the second approval above the threshold

The right-hand column turns a classification of friction into an operating instruction.

3. Agentic execution changes where decision rights need to sit



Decision rights in most enterprises were drawn for a workforce of people: who may approve, who must escalate, who signs. Agentic work does not always fit those lines. An agent can carry out a multi-step task that once passed through three roles and two approvals, which moves a decision that the org chart still shows sitting with a person. This is one reason pilots can stall on the way to scale: the decision rights, ownership and controls around the model were drawn for a different kind of worker. The evidence is consistent with this. Value from AI comes through complementary organizational assets (OECD, 2024). McKinsey's practitioner work on agentic organizations argues that when systems act, decision rights move with them, so organizations have to decide explicitly who is accountable (McKinsey, 2025). The evidence does not isolate how much of a stalled deployment is organizational rather than technical. It does show that complementary organizational capability matters to value creation. The pattern is consistent enough to act on: where agentic work meets decision rights drawn for people, those rights can become a binding constraint before the technology does.

4. As agent autonomy rises, accountability has to become explicit



Boards rightly worry about data, architecture, cybersecurity, talent and cost, and each can constrain agentic scaling. When an agent can act rather than only advise, one more requirement moves to the front: a clear, traceable owner of the outcome. Supervisory frameworks are converging on this. The FSB's proposed sound practices ask boards to set out prohibited uses, including fully automated decision-making in critical business areas. They also ask boards to define clear roles across the three lines of defense and to keep people in charge of highly autonomous systems (FSB, 2026, consultation report). The NIST framework treats accountability as requiring documented roles and executive-level responsibility (NIST, 2023). Outsourcing the work does not outsource the answerability. A bank that automates monitoring keeps full regulatory accountability (BCBS, 2020), and a Gulf central bank has codified this expectation for its supervised institutions through binding IT-governance requirements (SAMA, IT Governance Framework). McKinsey's practitioner work reaches the same conclusion from the management side: autonomous workflows make explicit human accountability essential, both for using a workflow and for building and configuring it (McKinsey, 2025). As agents move from advising to acting, accountability becomes a first-order operating-model constraint. It tends to bind early in regulated and concentrated-ownership settings, where the interface between owners, boards and management is itself varied and not standardized (Martínez-García, Basco and Gómez-Ansón, 2021).



Scaling agentic AI requires redesigning the operating model

Automation makes the existing model faster without changing what it is. Scaling agentic AI therefore requires a corresponding redesign of delegation boundaries, decision rights, controls and accountability when people and agents share the work.



3 Design the operating model around value, decisions and accountability

For Synarchy, an operating model is the system through which strategy becomes work. It covers how value is created, where decisions sit, who does the work, how technology takes part, how control is exercised and how the whole is governed over time. The org chart is one expression of it. Operating Model by Design anchors on the enterprise's value logic, which strategy sets and periodically revisits. It then redesigns the rest of the system (structure, decision rights, work design, control and governance) for a workforce of people and agents. It does not discard the organization, the plan or the controls; it puts accountability where agents now act. Five disciplines, guided by one design principle, form a chain:

- Map where value is created
- Decide which decisions matter and who owns them.
- Divide the work between people and agents.
- Design the controls around that division.
- Set the rhythm that keeps the model current.

Each rests on the one before it.

The design principle: import proven methods, design the governance locally

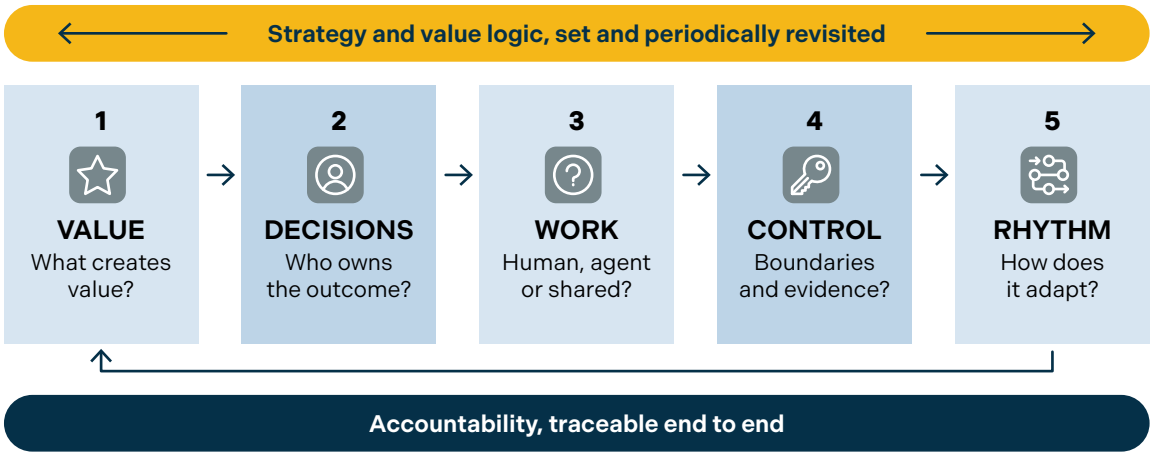
One rule shapes how the disciplines are adopted in the region. Management practices are portable: they transfer across borders and explain much of the productivity gap between firms and countries (Bloom, Sadun and Van Reenen, 2016). The governance that gives those practices authority is not. In the GCC, the effect of family and state ownership on performance depends on context. The interface between ownership and management is where decision rights are actually settled (Martínez-García, Basco and Gómez-Ansón, 2021; Al-Hadi et al., 2022). So import the method, and design the decision rights, the board or family interface and the accountability logic for the enterprise's own circumstances. This is a design choice, not a deficiency.

Why the GCC raises the stakes

The stakes are high here for reasons specific to the region. Regulated industries answer to supervisors who hold the institution accountable whatever the technology, and a Gulf central bank has made IT and cloud governance a binding requirement for its supervised institutions (SAMA). Sovereign enterprises carry a national mandate that can turn an unowned agent decision into a public matter as well as a commercial one. Family groups concentrate decision rights at the owner-manager interface, which agents cannot navigate without an explicit accountability map. Conglomerate structures also carry familiar internal-capital-market risks, including cross-subsidy between units (Scharfstein and Stein, 2000), and in concentrated-ownership settings the governance of those decisions matters even more. Each raises the cost of an Accountability Gap, which Operating Model by Design is built to close.

Exhibit 4. Five disciplines form one chain, from value to rhythm, on a base of traceable accountability

The Operating Model by Design framework



Each discipline rests on the one before it.

IN ONE LINE

Anchor on value; redesign the work around it.

Discipline 1 | Value-loop mapping.

What it is

Naming the small number of value loops through which the enterprise creates value. A value loop is an end-to-end sequence of work and decisions that produces a material customer, financial or stakeholder outcome. In a bank, a value loop runs from origination to underwriting, approval, servicing and monitoring; in a family group, from investment thesis to capital allocation, portfolio governance and exit. Mapping them makes the current value logic explicit and gives the anchor against which work, decisions and controls are redesigned.

Why it matters

You cannot redesign a model you have not made explicit. Most enterprises can draw their org chart but not their value loops, which is why automation lands on tasks rather than on the logic that matters. Mapping the loops is what separates real redesign from tinkering.

How to implement

- Draw the three to five loops that create the enterprise's value; if it takes more than five, the model is being confused with the org chart.
- For each loop, mark where value is earned from judgement and trust, and where it is earned from friction that AI will expose.

Failure mode
Automating tasks inside loops nobody has mapped.

Quick win
Map one value loop end to end and mark every point where an agent could act and who would answer for it.

Discipline 2 | Decision and accountability architecture.

What it is
The explicit assignment of who owns each material decision and its outcome once agents share the work, expressed as a spectrum from agent-executable to never-delegable.

Why it matters
This is the discipline the moment demands and the one imported templates get wrong. When an agent can act, accountability does not distribute itself; it must be designed. In regulated and family-owned contexts the never-delegable core, the decisions that must keep a named human owner, is both larger and less negotiable than a generic template assumes.

How to implement

- Place every material decision on the spectrum: agent-executable, agent-assisted, human-reserved, never-delegable.
- For the never-delegable core, name the accountable human and the evidence trail that proves the outcome was owned.

Failure mode
Assuming accountability follows the agent; it does not.

Quick win
Take one high-materiality decision an agent now touches and write down, in one line, who answers for the outcome.

Discipline 3 | Agent-human work design.

What it is
The deliberate division of each loop into what agents execute, what they assist, and what humans own, with the boundaries drawn by materiality and reversibility rather than by convenience.

Why it matters
Agents are strongest at sensing, retrieval and bounded execution; human ownership remains critical where judgement, context and answerability are material. Designing the division deliberately, rather than letting it emerge, is what lets the enterprise use machine speed without ceding the decisions that carry accountability.

How to implement

- Divide each mapped loop into agent-executable, agent-assisted, human-reserved and never-delegable segments, bounded by materiality and reversibility.
- Reserve material, irreversible and high-judgement decisions for accountable humans; delegate bounded, reversible work to agents where controls and escalation are explicit.

Failure mode
Confusing what an agent can do with what it should own.

Quick win
Redesign one loop so the agent runs the reversible steps and a human owns the irreversible one.

Discipline 4 | Control and assurance by design.

What it is
Building the traceability, logging and human-oversight checkpoints into the model from the start, so that assurance is a property of the design rather than a control bolted on after a pilot succeeds.

Why it matters
In regulated and sovereign contexts, an agent action that cannot be traced to a decision and an owner creates an assurance and accountability exposure alongside any productivity benefit. Agentic systems demand a fuller record than earlier automation. They need a complete decision-and-action trail: the sequence of inputs, tool calls, actions, approvals, exceptions and outputs, rather than the internal reasoning of the model. They also need controls such as agent-level certification and audit trails of agent transactions (FSB, 2026).

How to implement

- Require every agent action in a material loop to leave a decision-and-action trail that resolves to a named owner.
- Build oversight checkpoints, dual authorization and thresholds into the design, not into a remediation project later.

Failure mode
Bolting controls on after the pilot works, when the Accountability Gap is already live.

Quick win
Instrument one agent workflow so its full decision-and-action trail and its human owner are visible before it scales.

Discipline 5 | Operating rhythm.

What it is
The recurring governance, cadence and ownership that keep the operating model current as agents, tasks and risks evolve, rather than designed once and left to drift.

Why it matters
Because agentic capability changes monthly, a model designed once is out of date within a cycle. The rhythm converts a one-time redesign into a standing capability, and it is where the board's assurance role lives, feeding what agents actually did back into how the model is designed.

How to implement

- Run a standing cadence that reviews the agent-human boundary and the accountability map against what agents actually did.
- Feed assurance findings back into the value-loop map, so the model keeps pace rather than drifting out of date.

Failure mode
Treating the redesign as a project with an end date.

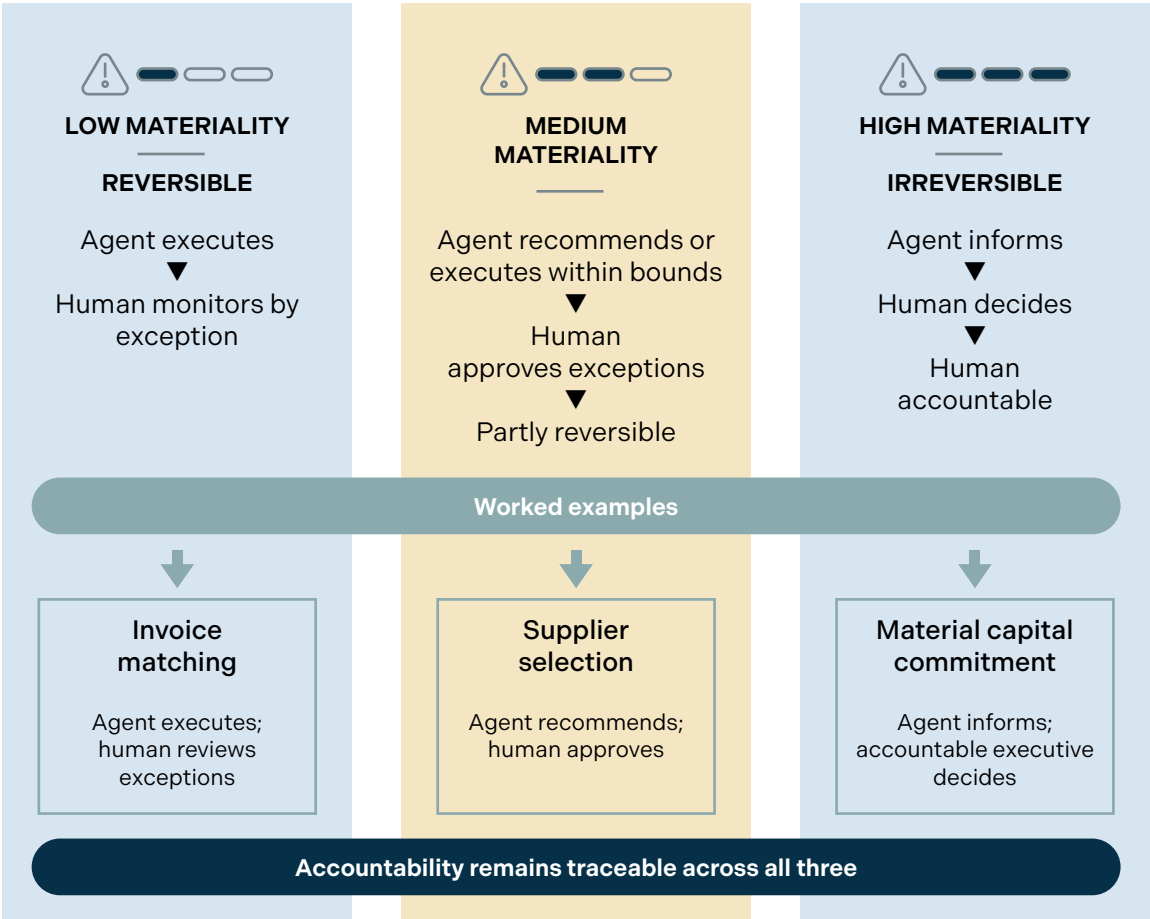
Quick win
Stand up a monthly operating-model review for one business unit and run it for a quarter before scaling.

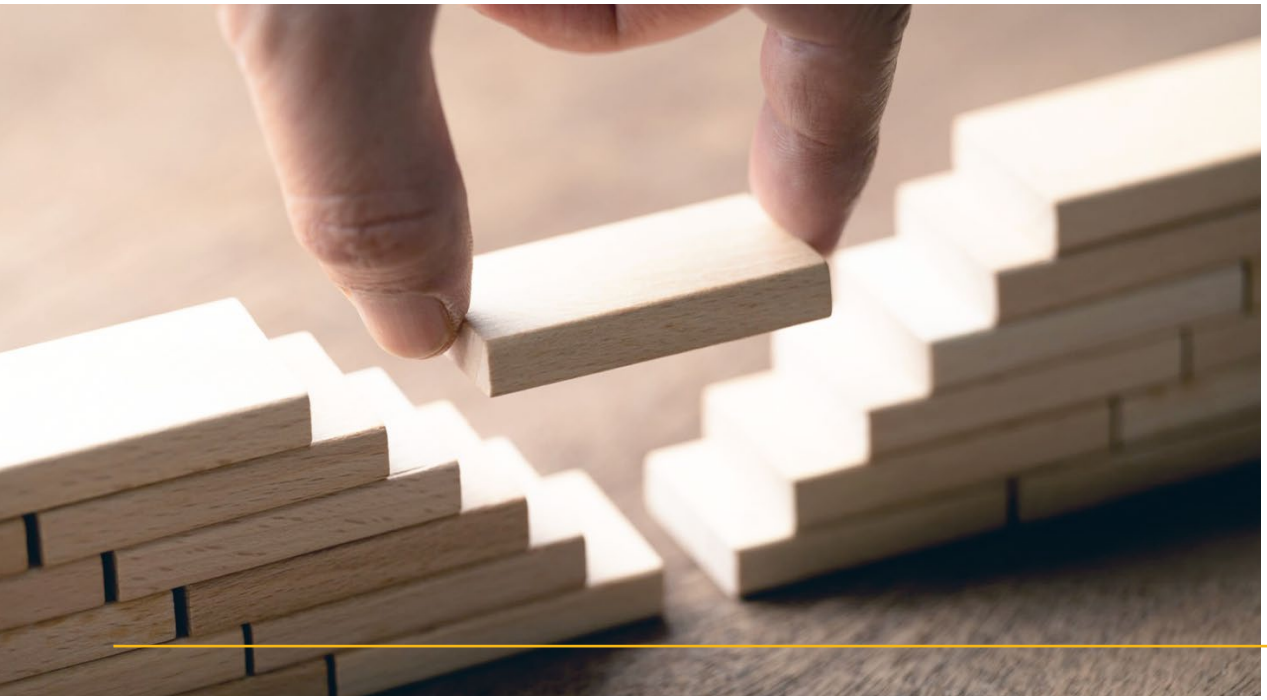
Where agents execute, and where a person is accountable

Beneath the disciplines sits one loop, and it is worth being exact about where machine speed helps and where human answerability must stay. Agents are strongest at the front, sensing and interpreting, and the appropriate human role rises with the materiality and irreversibility of the decision. The loop then learns, feeding what happened back into the design. Keeping a person in every step for its own sake is not the aim. Low-materiality, reversible, bounded work can sit with an agent under explicit controls, provided every material agent action resolves to a person who owns the outcome. This allocation by reversibility and risk, rather than by seniority, is the core of the design.

Exhibit 5. Human involvement should rise with the materiality and irreversibility of the decision

Human involvement by materiality and reversibility





4 Measure where agent autonomy has outrun accountability

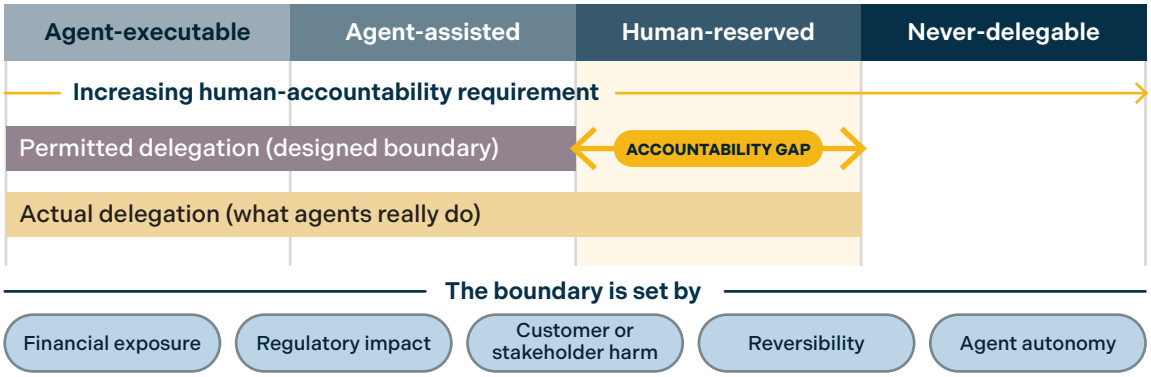
The next question for management is where the enterprise falls short today: where agent delegation has moved ahead of clear ownership and traceability.

The spectrum, and the gap

Every material workflow sits somewhere on a spectrum of delegation. At one end is work that is agent-executable: routine, reversible and low in materiality. Next is agent-assisted, where the agent proposes and a person decides. Then human-reserved, where judgement, discretion or the risk of customer harm keeps the decision with a person. At the far end is the never-delegable core: the decisions that, in a regulated, sovereign or family-owned enterprise, must always keep a named human owner. The Accountability Gap is simple to state and harder to measure: material work delegated to an agent beyond the boundary designed for it, with no traceable human owner of the outcome. Defining which workflows are material needs explicit criteria, set by the enterprise, before the gap can be read.

Exhibit 6. The Accountability Gap opens when agent autonomy moves past the designed boundary

Where work sits, and where the gap opens



Accountability is not manual execution

Agent. May perform the work.

Decision owner. May approve, or reserve, a material decision.

Accountable executive. Answers for the outcome, whoever or whatever performed it.

Assurance. Verifies the system and the decision-and-action trail.



Two supporting views, read as a portfolio

Two management views make the gap visible.

Delegated material work shows how far delegation has gone: the material workflows in which an agent can act independently.

Accountability traceability shows how much of that delegation is owned: the delegated workflows that resolve to a named accountable person and a reviewable record of decisions and actions.

The Accountability Gap is the delegated material work that lacks that complete traceability, everything an agent can do independently that no one can yet be held to account for.

These views are a portfolio picture rather than a precise ratio: materiality is a judgement, traceability is rarely all-or-nothing, and a single capital-allocation decision and a routine reconciliation are not equal workflows.

These are management views, not validated thresholds, and they are presented as such. Their value lies in the act of measurement. Most enterprises have never asked, of the work an agent now touches, how much of it has someone who can be held to account for the result. The answer is usually uncomfortable and immediately actionable. The gap should be read by materiality rather than by count. Materiality is judged on financial exposure, regulatory impact, harm to customers or stakeholders, reversibility and the degree of agent autonomy, so that the gap shows where unowned delegation actually carries risk.

Exhibit 7. The never-delegable line, and the gap, fall differently across GCC enterprise types

How the gap differs across regulated, sovereign and family enterprises

	Regulated	Sovereign	Family Group
Accountability anchors to	Regulator and board	Mandate, board and state shareholder	Owner-principal and board
Never-delegable core	Compliance, credit, customer harm	Major capital, mandate, public value	Capital allocation, related-party and succession-related principal decisions
Where agents create most tension	Regulated decisions and controls	Delegated agent action without a clear mandate owner	Informal decision rights not written into the operating model
What an unowned agent action puts at risk	Regulatory breach, customer harm and supervisory action	Public accountability and the national mandate	Family capital, owner relationships and succession
Design priority	Traceability	Mandate alignment	Explicit owner-management decision rights

The five disciplines are the same in every setting. What changes is who accountability answers to, which decisions are never delegable, where agents create the most tension, what an unowned agent action puts at risk, and which design priority closes the gap first. Illustrative archetypes; individual enterprises will differ.



5 Business leaders own the model; assurance tests it

Seeing the gap is only the first step; governance closes it and keeps it closed. The operating model belongs to management, and the design lasts when accountability for it is assigned and reviewed, so that it does not depend on one sponsor. The CEO and COO own the operating model and its value logic. Business and process owners own the outcomes and workflows. The CIO or CTO owns the agent platform and its technical controls. Risk and compliance set policy and risk boundaries and provide challenge. Internal audit gives independent, periodic assurance. The board oversees material agentic risk and accountability. This matches the direction of supervisory practice. The FSB's proposed practices, a consultation rather than a binding standard, ask for board oversight, clear roles across the three lines of defense and a deliberately chosen operating model (FSB, 2026). Four commitments turn this into named ownership.

Four commitments the board can hold

The first commitment is to protect the value logic. Name the small set of value loops the enterprise genuinely owns and the boundaries that must not move, so that everything else can be redesigned without losing the center. This belongs to the CEO, and it is reviewed as a structured judgement of how closely the model stays tied to owned value.

The second is to design the decision and accountability architecture: place every material decision on the spectrum and name the owner of each outcome, most rigorously for the never-delegable core. This belongs to the executive owner of the operating model, most often the COO, with business and process owners accountable for their workflows and risk and compliance providing challenge rather than owning the first-line decision. It is read through delegated material work and accountability traceability.



The third is to make control and assurance part of the design: require the record of decisions and actions, and human oversight, before an agent workflow scales, not after. This sits with the CIO or CTO and the risk function.



The fourth commitment is to review the operating model continuously as agent capabilities and risks change: run the operating rhythm that checks the boundary between agent and human work against what agents actually did. Management runs this review; internal audit tests it periodically and independently; and the board oversees the Accountability Gap over time.

Together the four turn a one-time redesign into a capability the enterprise runs.

Exhibit 8. Management owns the operating model; assurance tests it; the board oversees it

Who owns each part of the operating model, and how it is checked

Role	Commitment	Owns	Checked through
Board	Oversees all four	Oversight of material agentic risk and accountability	Accountability Gap trend; exceptions
CEO	1 Protect the value logic	The value logic and the value loops the enterprise owns.	Value-loop outcomes
COO	2 Decision and accountability architecture 4 Runs the operating rhythm	The operating model and its decision and accountability architecture	Traceability and escalation
Business and process owners	2 Own each workflow outcome	Outcomes and workflows, including who answers for each agent action	Delegated material work in their loops
CIO or CTO	3 Control and assurance by design	The agent platform and technical controls, including the record of decisions and actions	Trail completeness before scale
Risk and compliance	Challenge on 2 and 3	Policy, risk boundaries and challenge	Boundary breaches; control exceptions
Internal audit	Tests 4 independently	Independent, periodic assurance of the design and its operation	Audit findings



6 What this looks like in practice

What does this change in practice? The contrast below is an illustrative composite: a diversified enterprise in a regulated GCC sector, built from the mechanisms documented in this paper. It describes no single company. It sets what happens when an enterprise automates inside the old model against what happens when it redesigns the model.

Exhibit 9. Redesigning the model closes the gap without adding automation
An illustrative before-and-after for the composite enterprise

BEFORE
Automation inside the old model

- 28 delegated material workflows
- 9 fully traceable
- **19 not traceable**

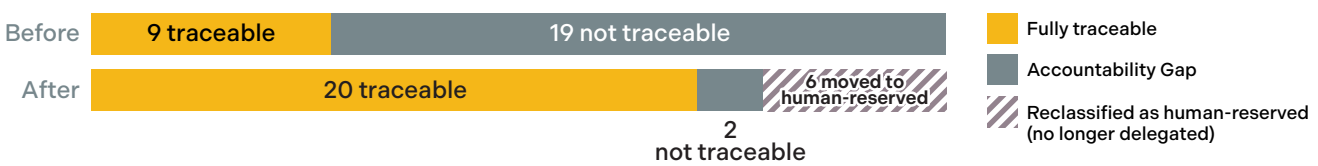
THE REDESIGN
Same agents, redesigned model

- 6 workflows reclassified as human-reserved
- Accountable owners named
- Dual-authorization controls added
- Logging and escalation built in
- Board reporting in place

AFTER
Operating model redesigned

- 22 delegated material workflows
- 20 fully traceable
- **2 not traceable, governed as remaining exposure**

The Accountability Gap: delegated material workflows with no traceable owner



The redesign does not add agent capability. It changes the operating model around it: workflows that had drifted past their designed boundary are reclassified as human-reserved, accountable ownership is established, controls are built in, and the remaining delegation becomes traceable. Unowned exposure falls materially without reducing the role of agents in reversible work.



7 Five questions for leadership

An enterprise can locate itself with a few questions about value, delegation, ownership and traceability. A useful starting question is: of the material work an agent now touches, how much resolves to a named human who can be held to account for the outcome? If the answer is unclear, the operating model warrants as much attention as the technology and talent supporting the deployment.

A five-question self-assessment

Value logic
Can you name the three to five loops through which your enterprise actually creates value, distinct from your org chart?

Accountability
For the material decisions an agent now touches, can you name, for each, the human who owns the outcome?

Never-delegable core
Have you defined the decisions that must always keep a named human owner, and is any agent currently acting past that line?

Traceability
For agent actions in your material loops, can you trace the decision-and-action trail to a decision and an owner before they scale?

Rhythm
Does your executive team revisit the operating model on a standing cadence, or was it designed once and left to drift?



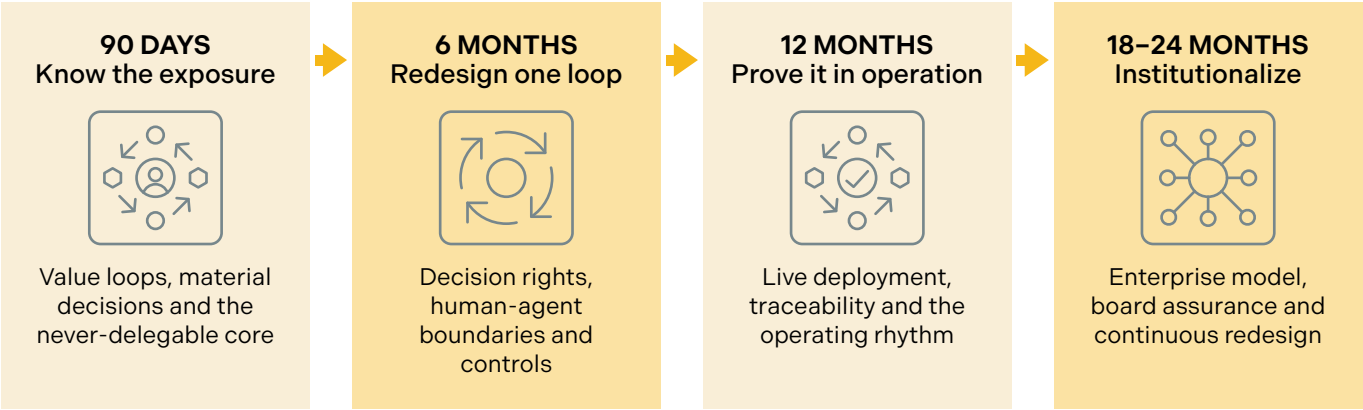
Where to start

Start with one material value loop. Map the decisions agents already touch, identify where ownership is unclear, and redesign that loop before scaling the same pattern. Run the loop with its new decision rights and controls for a quarter, then extend the design across the enterprise as the gap closes.

As an illustrative pathway, implementation should produce visible operating changes early rather than defer value until an enterprise-wide rollout is complete. It can begin with value-loop mapping and the never-delegable core in the first ninety days. By six months, the enterprise redesigns one loop's decision rights, human-agent boundaries and controls. By twelve months, it proves that design in live operation, with traceability and the operating rhythm running in a lead business unit. Enterprise-wide institutionalization typically follows over eighteen to twenty-four months, depending on the enterprise's starting point.

Exhibit 10. A first accountability baseline in ninety days; the model institutionalized over eighteen to twenty-four months

From first map to standing capability



Visible operating changes emerge within the first quarter.

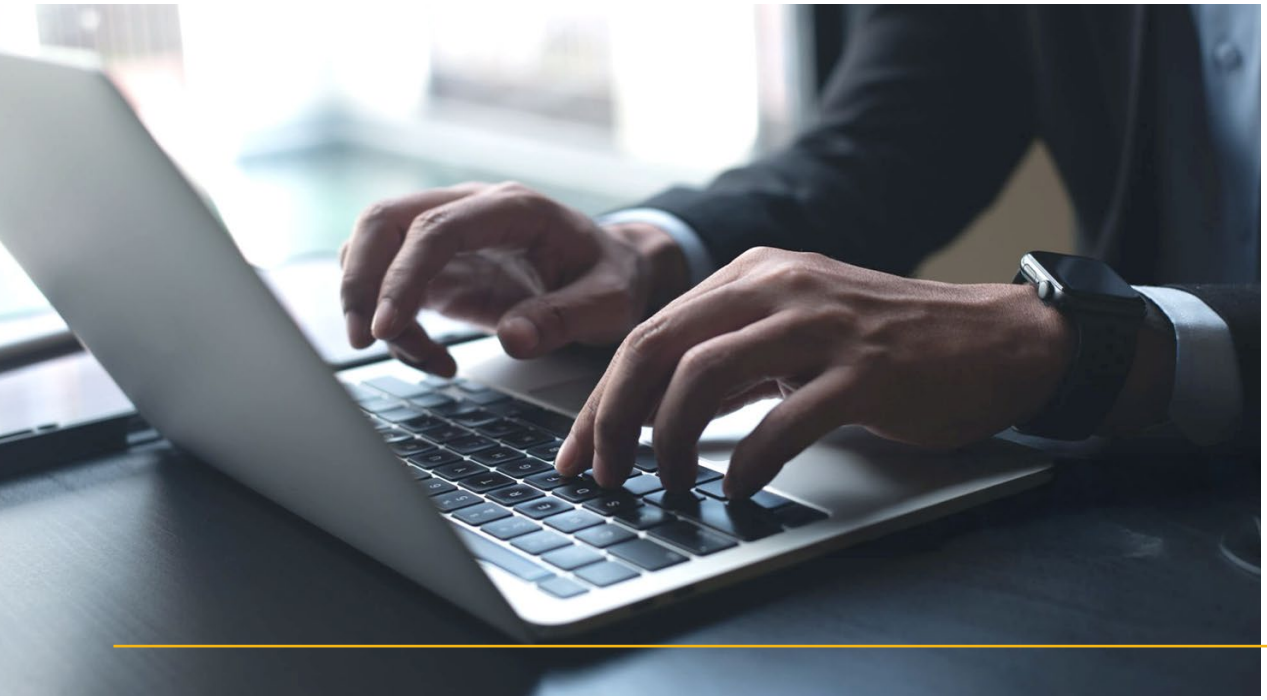
THE POINT OF VIEW

The operating model extends well beyond the organization chart.

As agents take on more work, decision rights, controls and accountability must move with it.

Operating Model by Design redesigns that system deliberately, and the Accountability Gap shows where delegation has moved ahead of ownership.

AI is changing how much work an enterprise can delegate. The operating-model question is how far authority and accountability can move with it, and the enterprises that answer it deliberately will scale agentic AI with greater speed, control and confidence.



APPENDIX · REFERENCES

A Frequently Asked Questions

What is Operating Model by Design?
A designed operating model that anchors on the enterprise's value logic, set by strategy, and redesigns how work is structured, decided and owned for a workforce of people and agents. It puts accountability where agents now act, and the Accountability Gap shows where it is still missing.

What is the Accountability Gap?
Material work delegated to an agent beyond the boundary designed for it, with no traceable human owner of the outcome. It shows whether an enterprise has automated faster than it has designed the accountability to carry the automation.

Why do our AI pilots stall before they scale?
Technology is often only part of the explanation. The evidence is consistent that realized value depends on complementary organizational assets, and that once those are accounted for the raw link between AI use and productivity weakens. Pilots can stall when ownership, decision rights and traceability have not been redesigned for scaled agent deployment.

How is this different for regulated and family-owned enterprises in the GCC?
The never-delegable core is often larger, because of regulatory, ownership and mandate considerations. Supervisors hold the institution accountable whatever the technology, and a Gulf central bank has made IT and cloud governance binding for its supervised institutions. Family groups concentrate decision rights at the owner-manager interface, which agents cannot navigate without an explicit accountability map. The methods can be imported; the decision rights and accountability have to be designed locally.

Does this mean keeping a human in every step?
No. It means every agent action, however autonomous, must resolve to a human who owns the outcome. Agents run the reversible, low-materiality work at machine speed; humans own the irreversible decisions and the assurance. The goal is traceable accountability, not universal manual review.

References

Al-Hadi, A., Eulaiwi, B., Duong, L., Taylor, G. and Dutta, S.K. (2022) 'Family power and corporate investment efficiency', Emerging Markets Finance and Trade. Available at: https://doi.org/10.1080/1540496x.2022.2088348.

Basel Committee on Banking Supervision (2020) Sound Management of Risks Related to Money Laundering and Financing of Terrorism (d505). Bank for International Settlements. Available at: https://www.bis.org/bcbs/publ/d505.pdf.

Bloom, N., Sadun, R. and Van Reenen, J. (2012) 'Americans do IT better: US multinationals and the productivity miracle', American Economic Review, 102(1), pp. 167–201. Available at: https://www.aeaweb.org/articles?id=10.1257%2Faer.102.1.167.

Bloom, N., Sadun, R. and Van Reenen, J. (2016, revised 2017) Management as a Technology? Harvard Business School Working Paper 16-133 (revised October 2017). Available at: https://www.hbs.edu/ris/Publication%20Files/16-133_64fd57c1-5f76-415a-9567-f1c0d310aff3.pdf.

Coase, R.H. (1937) 'The nature of the firm', Economica, 4(16), pp. 386–405. Available at: https://doi.org/10.1111/j.1468-0335.1937.tb00002.x.

Financial Stability Board (2026) Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Consultation report, 10 June 2026. Available at: https://www.fsb.org/uploads/P100626.pdf.

Martínez-García, I., Basco, R. and Gómez-Ansón, S. (2021) 'Dancing with giants: contextualizing state and family ownership effects on firm performance in the GCC', Journal of Family Business Strategy, 12(4), article 100373. Available at: https://doi.org/10.1016/j.jfbs.2020.100373.

McKinsey & Company (2025) Accountability by design in the agentic organization, The Organization Blog (practitioner article). Available at: https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/the-organization-blog/accountability-by-design-in-the-agentic-organization.

McKinsey & Company with the GCC Board Directors Institute (2025) The state of AI in GCC countries: In pursuit of scale and value (survey of senior executives and board directors; self-reported). Available at: https://gccbdi.org/sites/default/files/2025-11/The-state-of-AI-in-GCC-countries-In-pursuit-of-scale-and-value.pdf.

National Institute of Standards and Technology (2023) Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1. Available at: https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf.

OECD (2024) Fostering an Inclusive Digital Transformation as AI Spreads Among Firms. Available at: https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/11/fostering-an-inclusive-digital-transformation-as-ai-spreads-among-firms_cd50d324/5876200c-en.pdf.

Saudi Central Bank (SAMA) Information Technology Governance Framework. SAMA Rulebook. Available at: https://rulebook.sama.gov.sa/en/information-technology-governance-framework.

Scharfstein, D.S. and Stein, J.C. (2000) 'The dark side of internal capital markets: divisional rent-seeking and inefficient investment', The Journal of Finance, 55(6), pp. 2537–2564. Available at: https://doi.org/10.1111/0022-1082.00299.

Udomsaph, C. (2025) AI Adoption in MENAAP: Evidence from Enterprise Surveys. Working paper using World Bank Enterprise Survey microdata. Available at: https://thedocs.worldbank.org/en/doc/e8ee6db45dc7385cb334c646072791a7-0280032026/original/87-udomsaph20250920-MENAAP.pdf.

About the Author



Ramki Jayaraman,
Managing Partner, Synarchy Consulting.

Ramki Jayaraman is the Managing Partner and founder of Synarchy Consulting. With over 25 years of consulting experience spanning India, the Gulf and Africa, he specializes in transforming complexity into clarity, harmonizing strategy, technology and operations to unlock sustainable growth for governments, enterprises and family businesses. He has led large-scale transformation programs for sovereign enterprises, family groups, industrial groups and government bodies across the GCC and sub-Saharan Africa, combining the strategic depth of top-tier advisory with the pragmatism required to deliver in complex, multi-stakeholder environments.

About Synarchy Consulting

Synarchy Consulting is a Dubai-based strategy and advisory firm that helps organizations simplify complexity and unlock strategic value. Serving governments, enterprises, family businesses, and growth ventures across the Middle East and Africa, the firm works in strategy execution, digital transformation, AI and data advisory, people and organization design, and operational excellence. Synarchy combines deep regional context with a network of senior practitioners, each bringing two to two and a half decades of domain expertise and deployed on a project-specific basis, delivering the capability of a top-tier firm with the agility and accessibility of a specialist partner.