

CEO BRIEFING PAPER

Redesigning the Enterprise for the Agentic Era

A CEO's Operating Framework

Vijayan Seenisamy

Enterprise Agentic AI Systems Delivery

Creator, AI ROF™

Author of The AI Delivery Manager Blueprint and The Pilot Trap

The enterprise AI value gap is not a technology problem. It is an operating model problem.

The same models, the same vendors, the same talent pool are available to every enterprise. The outcomes differ by orders of magnitude. The variable that explains the difference is how the enterprise is organised to capture value, not what it has bought.

The strategic response is to run two operating models in parallel. Path A maintains productivity within the current structure. Path B redesigns the structure itself. Path A funds Path B. Path B eventually absorbs Path A.

Only the CEO has the authority to authorise both. Only the CEO has the authority to design the integrated operating model Path B requires. This work cannot be delegated to a CAIO, CTO, or CIO without delegating the outcome itself.

This paper provides the framework. The book that follows provides the operating manual.

1. The Gap

Most CEOs reading this paper are sitting on the same quiet contradiction. The dashboards say the AI agenda is working. The board updates have shown progress for six straight quarters. The CIO reports productivity gains. The Chief AI Officer reports new deployments. The CFO, however, cannot find the value in the numbers. The General Counsel is uneasy about how the audit trail will hold up under the next regulatory review. The C-suite peers each defend their own corner of the agenda, and no one is solving the whole thing.

Behind the contradiction sits a financial reality the dashboards do not yet show. The enterprise is accumulating an operating tax — the invisible compounding cost of monitoring, evaluation, governance, recovery, vendor change, and regulatory adaptation — that the original business cases did not model. By the second year of any consequential AI investment, this tax determines whether the agent is actually profitable in the system, regardless of how it appeared as a component in Year 1. The CEOs who can see this curve are early. Most cannot.

If this is the situation, it is not personal. It is the pattern across most of the Fortune 500, the FTSE 100, and the ASX 30 in 2026. MIT NANDA reports that 95 percent of enterprise AI pilots fail to reach reliable production ^[1]. ADAPT's most recent CFO survey finds that only four percent of finance leaders describe their organisations as very effective at generating AI value ^[2]. Those two numbers, read together, are the financial signature of an enterprise running AI inside an operating model that was not built to receive it.

The technology is not the problem. Every enterprise has access to the same frontier models, the same agentic platforms, the same vendors, the same talent pool. Two enterprises that buy the same stack produce outcomes that differ by orders of magnitude. The variable that explains the difference is not the technology. It is the operating model.

Enterprises that produce value from AI at scale share a structural characteristic. They have redesigned how the business is organised, governed, and operated around a class of system that behaves probabilistically rather than deterministically. Enterprises that have not produced value share the inverse characteristic. They have deployed AI into operating models designed for a different kind of technology and are now paying an invisible compounding cost for the mismatch.

This paper argues that the AI value gap is not a technology problem. It is an operating model problem. It is a CEO problem. And it is solvable.

The argument has three parts. First, the structural conditions of agentic-era enterprise transformation, named as principles. Second, the strategic response those principles require: the two-model duality. Third, what the CEO must lead personally, with a diagnostic to apply this

quarter and a frame for the next twenty-four months. The book that follows extends each part into the operating depth required to execute. This paper is the spine.

2. The Eight Principles

The principles below are the structural conditions of agentic-era enterprise transformation. They hold whether or not the enterprise acknowledges them. Each has been pressure-tested against capability progression, against the AI-native startup model, against the regulated enterprise context, and against the existing competitive landscape of strategy writing. They are not best practices. They are not advice. They are the constituent physics of operating agentic AI at enterprise scale.

Principle 1. AI is probabilistic, not deterministic.

Traditional software produces the same output every time given the same input. AI systems based on probabilistic generation do not. Variance is the substrate. It is not a flaw to engineer out.

This holds through every capability progression on the horizon. Improvements reduce variance around the mean. They do not eliminate it. Everything downstream cascades from this: measurement must reason about distributions, governance must include drift detection, accountability must include statistical reasoning, capital allocation must price in variance. Treating AI as deterministic, even implicitly, is the category error behind most enterprise AI failure.

"Variance is the substrate, not a flaw to engineer out."

Principle 2. Agentic AI is constitutive, not additive.

AI can form the operating layer of the business, not just augment existing processes. The positioning is a spectrum. At one end, AI sits on top of existing workflows as a productivity tool. At the other end, the business is reorganised around agentic AI as the constitutive logic of operations.

The two ends produce structurally different enterprises after five years. The additive enterprise has roughly the same people doing roughly the same work, slightly faster. The constitutive enterprise has fewer humans, organised differently, doing different work, governed differently, financed differently, regulated differently.

The positioning a CEO chooses on this spectrum, deliberately or by default, determines outcomes more than any other single decision in the AI agenda.

"Additive AI gives you a faster version of the business you had. Constitutive AI gives you a different business."

Principle 3. Unit economics of agentic AI are non-monotonic over time.

Traditional software has monotonic economics. Build once, deploy widely, marginal cost approaches zero, value compounds. Agentic AI does not work this way.

Costs compound through the operating tax: monitoring, evaluation, governance, recovery, vendor change, regulatory adaptation. Value erodes against edge cases through drift and hallucination. The two curves are not parallel. The cost curve rises. The value curve declines or plateaus. The two curves can cross. The point at which they cross is the Inversion Point.

Past the Inversion Point, the agent destroys value rather than creating it. Traditional ROI math, designed for monotonic systems, cannot see this coming. The investment cases approved in Year 1 with positive returns invert in Year 2 if the operating tax was not modelled. This is the most expensive single thing an enterprise can get wrong about AI.

"By Year 2, the operating tax decides whether the agent is profitable in the system, regardless of how it looked in Year 1."

Principle 4. Accountability is distributed across six functions and concurrent in time.

Every consequential AI decision involves six functions: the line of business that owns the use case, the privacy office that owns the data, the finance function that owns the value claim, the compliance function that owns the regulatory record, the risk function that owns the audit trail, and the CISO that owns the security perimeter.

Distribution alone is shared with deterministic software. The differentiator for AI is concurrency. Deterministic systems can be governed periodically because they are stable. AI systems are not stable. They generate new decisions every day. Drift can produce material changes in risk profile within weeks. The six functions must engage continuously, not at change events. Enterprises with sequential governance frameworks fail at AI governance because their cadence is wrong.

"Six functions own every AI decision. None of them can govern it alone."

Principle 5. Information flow constraints persist.

Hayek's dispersed knowledge. Coase's transaction costs. Weber's coordination at scale. Three sets of structural insights about how information moves through complex social systems. They are not artifacts of pre-AI technology. They are properties of distributed cognition, competing interests, and contextual judgment.

AI compresses some flows. It cannot dissolve the constraints. Tacit knowledge resists codification. Political legitimacy cannot be adjudicated by inference. Contextual judgment cannot be aggregated centrally without loss. The Silicon Valley narrative that AI will dissolve middle management and produce flat AI-native organisations mistakes a property of small organisations for a property of large ones. Even the most-cited examples of flattened AI-native structure retain the substantive functions of middleware. They consolidate those functions into fewer roles. They do not eliminate them.

"AI can compress information flow. It cannot dissolve the constraints that produce it."

Principle 6. Operating model determines value, not technology.

Two enterprises can access the same models, deploy the same agentic platforms, hire from the same talent pool, use the same vendors. Their outcomes vary by orders of magnitude. The variable that explains the difference is the operating model.

As technology, talent, and capital commoditise, the operating model becomes the binding constraint on value. This is counter-intuitive for technology-trained leadership. The technologist's intuition is that better models solve more problems. The structural reality is the opposite. Better models, available to everyone, make the operating model gap more visible, not less. When everyone has access to frontier capability, the difference between enterprises is how they are organised to capture it.

This is also the principle that frees the CEO from vendor narrative dependency. Every vendor pitches their technology as differentiating. The principle says no technology is. The operating model is.

"Better models, available to everyone, make the operating model gap more visible, not less."

Principle 7. Time-to-coherence beats time-to-deployment.

Speed of individual deployments produces fragmentation. Coherence across deployments produces compounding value. In agentic enterprise systems, deployments interact through shared trust, governance, and unit economics. An incoherent second deployment does not just add value of its own. It degrades the value of the first.

Ten incoherent agents do not produce ten times the operating tax of one. They produce closer to twenty-five times, because incoherence multiplies the monitoring, evaluation, and recovery surface area. The right tempo for AI is not maximum speed. It is the speed at which coherence holds.

"The right tempo for AI is the speed at which coherence holds."

Principle 8. Only the CEO has authority over the integrated operating model.

The operating model integrates business model, capital allocation, organisational structure, governance, accountability, capital structure, talent strategy, vendor strategy, regulatory engagement, and board reporting. One role has authority across all of these. Every other role has authority over their function.

For an AI redesign that requires changes across all of these, only the CEO can authorise the integrated change. The CIO cannot impose accountability on the CFO. The CAIO cannot impose policy on the General Counsel. The General Counsel cannot redirect capital. Each function operates inside its remit. The integration is the CEO's work.

This produces the most common failure pattern in enterprise AI today. The CAIO is appointed in a press release. Eighteen months later the board asks why the transformation has produced productivity tools and no operating model change. The CAIO is removed. A new CAIO is appointed. The pattern repeats. The diagnosis is always the same: the role was given the responsibility without the authority. Only the CEO holds the authority, and the authority cannot be transferred without transferring the role.

"Operating model is the one thing the CEO cannot delegate without delegating the outcome."

These eight principles produce one strategic implication that follows from them operating together. That implication is the centre of gravity of this paper, of the book that follows it, and of the redesign work every enterprise must now undertake.

3. The Two-Model Duality

Because the positioning of AI is a spectrum (Principle 2), because the economics compound non-monotonically (Principle 3), because the operating model determines value (Principle 6), and because only the CEO can authorise both states (Principle 8), enterprises in the agentic-era transition must run two operating models in parallel.

This is not a strategic preference. It is the structural consequence of the principles operating together.

Path A: Productivity within the current operating model

Path A is the enterprise as it is today, layered with additive AI productivity tools. Microsoft Copilot for the workforce. ChatGPT or Claude enterprise access. Point automations inside existing workflows. Retrieval over existing knowledge stores. Chatbots on top of existing customer service. Agents that sit on top of the existing structure. The structure itself does not change.

Path A produces real value. Productivity gains are measurable. The workforce becomes more capable in narrowly defined tasks. The CFO sees the lift in efficiency metrics. The board sees activity. None of this is wrong. It is necessary.

Path A is also not the destination. The operating tax compounds across the portfolio of incoherent additive deployments. Accountability remains sequential where it needs to be concurrent. The operating model has not been redesigned, so value capture is bounded by the pre-AI structure. Path A produces a ceiling that no amount of additional Path A investment can lift.

Most enterprises today are operating somewhere along Path A and treating the productivity gains as the AI transformation. They will hit the ceiling before they recognise it as a ceiling.

Path B: The redesigned operating model

Path B is the enterprise as it must become. The operating model itself is redesigned around agentic AI as a constitutive layer rather than an additive tool. Six structural components define Path B operationally.

First, a shared evaluation framework that every AI deployment must pass before reaching production and continuously while in production. Second, a single accountability map that names the six functions for every AI decision rather than reinventing accountability per deployment. Third, a single governance cadence that engages all six functions continuously rather than at change events. Fourth, a single unit economics view that tracks cost-per-outcome across the portfolio with the operating tax made explicit. Fifth, a single trust monitoring system that watches drift across all

agents, not just within each one. Sixth, standardised vendor and integration patterns so each new deployment absorbs into the shared system rather than fragmenting it.

These six components are not coherence. They are the operating model. The principle that coherence beats speed is the same principle as the principle that the operating model determines value. They are two ways of stating the same structural truth.

Path A versus Path B at a glance

	Path A	Path B
Position	Additive: AI sits on top of existing processes	Constitutive: AI is the operating layer of the business
Capital	Allocated per deployment, by function	Allocated to the integrated operating model first, deployments second
Accountability	Reinvented per deployment, sequential governance	Single map across six functions, concurrent governance
Unit economics	Reported by deployment, operating tax invisible	Reported across the portfolio, operating tax explicit
Talent and roles	Existing structure with new AI tools added	Redesigned C-suite ensemble, new roles for system stewardship
Vendor strategy	Best-of-breed per use case, fragmented contracts	Standardised contracts, common integration patterns
Board reporting	Deployment counts, productivity gains, AI tool adoption	Operating model maturity, operating tax, Inversion Point, convergence
Outcome by Year 3	Productivity ceiling. Operating tax compounds. Board patience erodes.	Compounding value. Operating tax bounded. Capability extensible to next wave.

Path A is necessary. Path B is the destination. Most enterprises today run Path A only and treat productivity gains as the AI transformation. The framework in this paper argues that both must run in parallel, that Path A funds Path B, and that Path B eventually absorbs Path A as functions migrate into the redesigned structure.

How this actually unfolds: a worked case

The following is a composite case drawn from a pattern visible across several mid-cap financial services enterprises in 2024 to 2026. The composite is illustrative. The structure is real.

A mid-cap regional bank with \$50 billion in assets, mid-market and SME lending, a small insurance subsidiary, and a consumer customer base of approximately 1.5 million. By mid-2024, the bank had deployed agentic AI across three functions. The customer service agent, on vendor A, handled first-line inquiries and product recommendations. The claims processing agent, on vendor B, served the insurance subsidiary. The underwriting risk-scoring agent, built in-house, supported commercial lending. Each deployment was approved as a discrete business case. Each reported to a different executive. Each ran on a different vendor stack.

Through 2024 and into the second quarter of 2025, each agent reported positive metrics against its own success criteria. Customer service handle time was down. Claims processing throughput was up. Underwriting cycle time had compressed. The CEO took these numbers to the board quarterly. The board approved continued investment.

By the third quarter of 2025, the picture had shifted. The Chief Risk Officer noticed that the customer service agent was recommending products to customers whom the underwriting agent would later decline. The customer experience was incoherent. When those declined customers complained, the customer service agent had no record of why the underwriting agent had declined them. The CFO could not produce a unified view of unit economics across the three agents because each had its own FinOps trail with different cost categorisations. The compliance team, preparing for the bank's first EU AI Act audit, discovered that each agent had a different audit trail format and that none met the high-risk system requirements.

In late 2025, the CEO authorised what this paper calls a Path B redesign. Operations continued on the existing structure. In parallel, a redesigned operating model was built. The Chief Risk Officer was made chair of a new AI Operating Council, with standing members from each of the six functions. A shared evaluation framework was specified. A unified observability and unit economics platform was implemented. Vendor contracts were renegotiated to common terms. The accountability map was published across the bank as a single artefact.

The initial redesign sprint took approximately nine months. Migration of the three legacy agents into the new operating model took a further five to seven months in parallel. During the migration, two new agents shipped inside the redesigned structure: a fraud detection agent and a document review agent for commercial lending. Both shipped in weeks rather than months because the operating infrastructure absorbed them. By the end of the migration, the bank had five agents running inside a coherent operating model rather than three running in parallel structures.

The bank's CEO did not stop Path A during the transition. The legacy three agents continued to serve customers throughout. The cost of running both paths in parallel was higher than continuing Path A alone would have been. The trajectory of value capture in the redesigned state, measured at eighteen months, was materially steeper than Path A alone would have produced, and structurally extensible to the next wave of deployments without re-architecting.

Two observations on timing are worth naming. First, the specific durations in this case will vary substantially across enterprises depending on sector, use case complexity, governance maturity, regulatory exposure, and the political dynamics of the C-suite ensemble. The point is not the exact month. The point is that the curve exists and that most enterprises in production have not yet modelled it. Second, the parallel running cost is a real number that has to be defended to the board. The CEO who cannot frame it as the cost of a transition that produces compounding value will struggle to authorise it.

This is the pattern. Path A funds and stabilises. Path B redesigns and absorbs. The CEO authorises both and decides the convergence point. The two-model duality is not a strategic option. It is the structural answer to the question of how an enterprise transitions from an operating model built for deterministic systems to one built for probabilistic ones.

Path A is necessary. Path B is the destination. The CEO who authorises both, and sequences the convergence, leads the redesign. The CEO who authorises only Path A approves a productivity gain and calls it a transformation.

4. The Operating Tax and the Inversion Point

Every AI agent in production carries an operating tax. The components are well-defined and recurring.

Monitoring	Continuous observation of inputs, outputs, latency, cost, error rates, and downstream effects.
Evaluation	Structured testing against drift, hallucination, edge case performance, and behaviour against changing input distributions.
Governance	Continuous engagement of the six functions, escalation paths, audit-ready documentation.
Recovery	Detection of failure, root cause analysis, rollback, retraining, re-deployment.
Vendor change	Model deprecation, vendor migration, prompt re-engineering, re-evaluation.
Regulatory adaptation	EU AI Act compliance, state legislation, sector-specific rules, the regulatory cadence of the next decade.

Each component is real. Each is recurring. Each scales with deployment depth, not against it. As an enterprise adds more agents into more workflows, the operating tax does not divide across the portfolio. It grows superlinearly because the integration surface area grows.

The operating tax is not on the business case. It does not appear in the total cost of ownership model the CFO is using. TCO models in current use were built for software that, once deployed, runs at near-zero marginal cost. They do not account for the continuous evaluation, drift monitoring, governance engagement, and recovery work that agentic AI requires throughout its lifecycle. The operating tax accumulates out of sight. By Year 2, it determines whether the agent is profitable in the system, regardless of whether it appeared profitable as a component in Year 1.

Three cases, three faces of the same pattern

Klarna deployed an AI customer service agent in early 2024 and reported strong cost savings through 2024. In May 2025, CEO Sebastian Siemiatkowski announced the company was rehiring human agents because customer service quality had dropped ^[9]. The agent's apparent unit economics in Year 1 had inverted by Year 2. Once escalation, rework, and human-handled exceptions were accounted for, the cost per successful outcome had moved above the human baseline the agent was meant to replace. The business case had inverted in roughly fourteen months. The operating tax was visible in retrospect. It was invisible at investment time. Timing will

vary considerably across enterprises, sectors, and use cases. The principle is that the curve exists. The discipline is to model it before it arrives.

Air Canada experienced a parallel inversion through liability rather than escalation cost. In February 2024, the Civil Resolution Tribunal of British Columbia ruled in *Moffatt v. Air Canada* that the airline was liable for misinformation given by its chatbot to a customer enquiring about bereavement fares^[4]. Air Canada had argued the chatbot was a separate legal entity. The tribunal found otherwise. The case is small in dollar terms. The precedent is not. It establishes that the operating tax includes regulatory and legal exposure that was previously assumed to sit with the vendor.

McDonald's pulled its IBM partnership for AI drive-through ordering in June 2024 after viral videos showed edge-case failures^[5]. The agent could not handle the long tail of real-world ordering. The operating tax there was operational and reputational rather than financial. Three different failure surfaces. The same underlying structure.

The Inversion Point as an operating concept

For every AI deployment in the enterprise portfolio, there is a month at which the cost-per-outcome curve crosses the human or legacy-system baseline. That month is the Inversion Point. It may be twelve months out. It may be twenty-four. It is rarely never. The exact timing depends on sector, use case, governance maturity, vendor stability, and regulatory pressure.

The CEO question this quarter is concrete and answerable. For the three largest AI deployments in the enterprise, what is the projected Inversion Point? If the CFO cannot model it, that is the first finding. If the modelled point is closer than the current business case assumes, that is the second finding. If there is no plan to either push the point further out through evaluation and recovery investment, or to retire the agent before the curves cross, that is the third finding.

The Inversion Point is the most useful single concept to take to the board this year. It reframes the AI conversation from deployment activity to financial reality.

[visual: operating tax curve crossing value baseline — to be produced by Brandon Daniel]

5. Self-Diagnostic

Ten questions. Score each from 1 to 3. One is clearly no. Two is partial. Three is clearly yes. Total possible: 30. Score honestly. Apply with the C-suite ensemble for triangulation.

1. The accountability map

Is there a documented accountability map naming roles in each of the six functions — line of business, privacy, finance, compliance, risk, CISO — for every AI deployment in production?

Score: ☐ 1 ☐ 2 ☐ 3

2. Concurrent governance

Do those six functions engage with each AI deployment continuously (at least monthly), rather than periodically at change events?

Score: ☐ 1 ☐ 2 ☐ 3

3. The Operating Tax

Has the CFO produced Year-2 operating tax estimates for the top three AI investments, with the six components broken out separately?

Score: ☐ 1 ☐ 2 ☐ 3

4. The Inversion Point

For each production AI deployment, is the projected month at which cost-per-outcome crosses the human or legacy baseline known and documented?

Score: ☐ 1 ☐ 2 ☐ 3

5. Coherence across deployments

Do all AI deployments in the enterprise share an evaluation framework, governance cadence, accountability map, and unit economics view?

Score: ☐ 1 ☐ 2 ☐ 3

6. Two-model clarity

Can each member of the C-suite articulate the distinction between Path A and Path B, and name which deployments belong to which?

Score: ☐ 1 ☐ 2 ☐ 3

7. Personal AI fluency

Does the CEO personally use AI tools for at least three substantive pieces of work per week?

Score: ☐ 1 ☐ 2 ☐ 3

8. The non-delegable list

Has the CEO defined what is non-delegable in this transformation, written it down, and shared it with the C-suite as shared expectation?

Score: ☐ 1 ☐ 2 ☐ 3

9. Board engagement reframed

Does the board hear about AI in operating model terms rather than deployment activity terms?

Score: ☐ 1 ☐ 2 ☐ 3

10. The convergence plan

Is there a sequenced plan for Path B absorbing Path A across functions, with decision points and milestones?

Score: ☐ 1 ☐ 2 ☐ 3

Interpreting the score

24 to 30	The enterprise is likely already running Path B in some form. The framework names what has been built and accelerates the convergence work.
17 to 23	The enterprise is in transition. Path A is functioning. Path B is partial. The diagnostic results identify where to begin.
10 to 16	Path A is being treated as the destination. The Inversion Point on current deployments is closer than dashboards suggest. The redesign needs to start this quarter.

The diagnostic is not an audit. It is a structured conversation starter for the C-suite. Run it personally first. Run it with each direct report individually. Compare the scores. The pattern of disagreement is often more informative than any single score. When the CEO scores 3 and the CIO scores 1 on the same question, that disagreement is the diagnostic.

6. What the CEO Cannot Delegate

Most of this transformation can be delegated. Implementation. Vendor management. Day-to-day governance. Talent strategy execution. Operations. These are the work of the C-suite ensemble and their teams. The CEO does not do them personally.

Six things cannot be delegated. Each follows from a principle. Each is the CEO's own work. Each, when delegated, produces a specific and predictable failure mode.

1. Authorising and sequencing the dual operating model

Follows from Principle 8.

Only the CEO can authorise Path A and Path B running in parallel, allocate capital across both during the transition, and decide the triggers and timing of convergence. This is a sequencing decision. Which function migrates first. What evidence triggers convergence. When the parallel cost becomes unsustainable. The CFO can model. The COO can advise. The decision is the CEO's, because only the CEO holds authority over both models simultaneously.

2. Designing the integrated accountability map

Follows from Principle 4.

This is a design decision, not a sequencing one. The CEO is the only role with authority across the six functions to name and operate the accountability structure as a coherent whole. The CIO cannot impose an accountability framework on the CFO. The CAIO cannot impose one on the General Counsel. Each peer designs within their function. The integration is the CEO's work and cannot be delegated to a coordinating role.

3. Redesigning the governance cadence

Follows from Principle 4.

Only the CEO can rebuild the enterprise's governance rhythm from periodic to concurrent. This touches the board calendar, the C-suite committee structure, the audit and risk committee cadence, and the operating reviews. Every other role runs their function inside the cadence. The cadence itself is set by the CEO.

4. Maintaining personal AI fluency

Follows from Principle 8, behavioural.

The CEO cannot delegate their personal experience of the tools. Conviction in this transformation comes from direct use, not from briefings. A CEO who has been briefed on AI but has not used it at depth cannot lead a redesign that depends on understanding what AI actually can and cannot do.

Three specific practices distinguish the CEO who has done this work from the one who has not. First, use the tools weekly to draft a substantive document the CEO would otherwise have asked a deputy to draft. The point is not productivity. The point is calibration of what the tools actually produce. Second, use the tools weekly to challenge a position the CEO has already formed. The point is to discover what AI surfaces that the CEO's own judgement missed. Third, use the tools weekly to learn something about a function the CEO does not personally run. The point is to understand the levelling effect AI has on functional expertise. The C-suite peers can tell the difference between a CEO who has done this work and a CEO who has been briefed on it. The board can tell. The conviction shows or it does not.

5. Owning the board narrative

Follows from Principles 6 and 7.

Only the CEO can reframe the AI conversation with the board from deployment activity to operating model maturity. The CFO supports with the financial layer. The General Counsel supports with the regulatory layer. The reframe itself is the CEO's. Boards trained on deployment-count narratives will continue to ask for those numbers until the CEO leads them somewhere else.

The right board update is not the count of agents in production. It is four numbers. Where the enterprise sits on the additive-constitutive spectrum. The Year-2 operating tax estimate on the top three deployments. How many of the six functions are continuously engaged in governance across the portfolio. How close convergence is between Path A and Path B. Four numbers, quarterly, until convergence. That single shift in board reporting changes the questions the board asks, the capital they authorise, and the patience they show.

6. Leading the C-suite ensemble redesign

Follows from Principle 8.

Every C-suite peer's role changes in the redesigned operating model. The CFO operates differently. The CIO operates differently. The General Counsel operates differently. The COO operates differently. Only the CEO can lead that redefinition with each peer individually and the ensemble collectively. This is the most political and most personal piece of the redesign. It cannot be delegated to HR. It is leadership work in the original meaning of the term.

Approve and delegate are the failure modes. Lead is the only mode that works.

7. The Next 24 Months

By 2028, the gap between Path A enterprises and Path A plus Path B enterprises will be visible at the level of the operating accounts. The Path A enterprises will continue to report productivity gains. They will also continue to report unit economics inversions, fragmented governance findings during regulatory audits, and growing operating tax that the CFO cannot quite explain. Their AI agendas will have plateaued. The board narrative will have shifted from optimism to defensiveness.

The Path A plus Path B enterprises will look structurally different. Fewer deployments, but each absorbing into a coherent operating system. Lower operating tax per agent. Unit economics that improve over time rather than degrade. Regulators easier to satisfy because the audit trails follow a single pattern. Customers experiencing AI as a coherent capability rather than a fragmented set of touchpoints. Boards engaged in operating model maturity discussions rather than deployment counts.

The difference will not be primarily a matter of capital spent. Path B enterprises will have spent roughly the same on AI, often less. The difference will be a matter of how the capital was deployed and what operating reality it built.

Three specific shifts between now and 2028 should focus the CEO's attention this year.

First, the EU AI Act compliance regime moves into full enforcement through 2026 and 2027. The Act entered into force in August 2024 with phased application ^[6]. General-purpose AI obligations and most provisions apply from August 2026. High-risk system requirements under Article 6 enter full enforcement through 2026 and 2027. Enterprises operating in or selling into the EU face audit and disclosure requirements that sequential governance frameworks cannot satisfy. The accountability map and the cadence model become regulatory necessities, not just operating preferences.

Second, the unit economics inversion of the 2024 deployment cohort begins arriving in 2026 across most large enterprises. The Klarna case in May 2025 was an early signal. The pattern repeats across the cohort that deployed customer service, claims, content, and code-generation agents in 2024. CFOs in multiple industries are beginning to see the inversion in their Year-2 numbers. Boards will start asking harder questions. The CEOs who have already established the Operating Tax and the Inversion Point as operating concepts will be ready for those conversations. Those who have not will be caught.

Third, the regulatory and liability landscape extends beyond the EU. State-level legislation in the United States creates parallel obligations. The Colorado AI Act, effective February 2026, requires risk management programs and disclosure obligations for high-risk AI systems ^[7]. The *Moffatt v. Air*

Canada precedent establishes that vendor distance does not transfer liability. Cyber insurers are now requiring AI security riders with documented red-teaming and risk assessments ^[8]. The exposure surface expands during 2026 and 2027. The operating model that respects the principles is the operating model that survives this expansion.

In short, the next twenty-four months are when the gap between Path A only and Path A plus Path B enterprises becomes visible, measurable, and decisive. The CEOs who lead the redesign now are positioned for the gap. The CEOs who defer the redesign find themselves on the wrong side of it.

One specific action this quarter

Run the self-diagnostic in Section 5 personally. Then run it with each member of the C-suite individually. Then run it with the full C-suite as an ensemble. Compare the scores. The pattern of disagreement across the C-suite is the most useful information the diagnostic produces. The disagreement is the diagnostic.

From there, the framework gives the sequence. Name the accountability map. Redesign the cadence. Model the operating tax and the Inversion Point. Name the convergence plan. Authorise both paths visibly. Lead the C-suite ensemble through the redesign with the discipline that only the CEO can bring.

The book that follows this paper extends each step into the operating depth required to execute it. The framework here is the spine. The book is the operating manual. Both are the discipline that the enterprise transformation of the agentic era demands.

AI value does not announce itself. It accumulates in enterprises redesigned to receive it, and evaporates in those that were not.

Vijayan Seenisamy

Enterprise Agentic AI Systems Delivery | Creator, AI ROF™
Author of The AI Delivery Manager Blueprint and The Pilot Trap
vijayan@aideliverydiscipline.com

Notes and Sources

Other data referenced in this paper:

- [1] MIT NANDA Initiative, *The GenAI Divide: State of AI in Business 2025*, August 2025. Based on 150 interviews, a survey of 350 employees, and analysis of approximately 300 public deployments.
- [2] ADAPT, CFO survey on AI value realisation, 2025-2026. The same survey reports that 77 percent of CFOs say their organisations are not effectively realising value from AI investments.
- [3] Klarna public statements and CEO Sebastian Siemiatkowski announcements, May 2025. Klarna had previously reported significant cost savings from its AI customer service deployment through 2024 and into early 2025.
- [4] Civil Resolution Tribunal of British Columbia, *Moffatt v. Air Canada*, 2024 BCCRT 149, February 2024.
- [5] McDonald's and IBM joint announcement, June 2024, on the conclusion of the Automated Order Taking partnership in McDonald's restaurants.
- [6] Regulation (EU) 2024/1689 of the European Parliament and of the Council on harmonised rules on artificial intelligence (the EU AI Act), entered into force 1 August 2024, with phased application. General-purpose AI obligations apply from August 2026. High-risk system requirements under Article 6 enter full enforcement through 2026 and 2027 depending on classification.
- [7] Colorado Senate Bill 24-205, Consumer Protections for Artificial Intelligence, effective February 2026.
- [8] Wilson Sonsini AI regulatory developments brief, January 2026, on the requirement by leading cyber insurers for AI security riders with documented red-teaming and risk assessments.

Additional references

- Camunda / Coleman Parkes, 2026 State of Agentic Orchestration and Automation, January 2026 (1,150 IT decision makers). Reports that only 11 percent of agentic AI use cases have reached production despite 71 percent of organisations claiming to use AI agents.
- S&P Global Market Intelligence, Enterprise AI POC abandonment survey, 2025. Reports that 42 percent of companies abandoned most AI initiatives in 2025, up from 17 percent in 2024.
- BCG, Enterprise AI value realisation report, September 2025. Reports 74 percent of companies seeing no value from AI investments.
- Gartner, agentic AI project cancellation prediction, June 2025. Forecasts that 40 percent or more of agentic AI projects will be cancelled by the end of 2027.
- McKinsey QuantumBlack, The State of AI 2025, November 2025 (1,993 participants across 105 nations).
- LangChain, State of Agent Engineering, December 2025 (1,300 professionals).