



From Operating Model to Operating Reality

Why the Intelligence-Centred Enterprise Stalls, and Where to Focus First


A diagnostic framework for CIOs, CFOs, and Chief AI Officers

Mark Cameron
Vijayan Seenisamy
Vinod Bijlani
Sharma Madiraju

| CEO, Alyve
| Enterprise Agentic AI Systems Delivery
| AI Practice Leader, HPE
| Enterprise AI Strategist

May 2026

TABLE OF CONTENTS



THE AUTHORS	3
EXECUTIVE SUMMARY	4
SECTION 1: THE QUESTION EVERY READER ASKED	5
SECTION 2: THE OPERATING MODEL GAP	6
SECTION 3: FIVE WAYS THE METABOLIC LOOP BREAKS	7
SECTION 4: FOUR LAYERS, FOUR CONSTRAINTS	15
SECTION 5: THE MISSING LAYER: DELIVERY DISCIPLINE	19
SECTION 6: ICE DELIVERY READINESS ASSESSMENT	22
SECTION 7: FROM DIAGNOSIS TO ACTION	25

The Authors



Mark Cameron

CEO and founder, Alyve

Mark has over 27 years of experience advising government and enterprise organisations on practical, value-driven AI adoption. His work spans strategy, leadership education, and hands-on implementation. He is a former Business Review Weekly digital strategy columnist and a current contributor to Forbes, where he writes on AI strategy, organisational change, and the realities of deploying AI at scale. Mark teaches Digital and AI Transformation Leadership as part of the Deakin University Executive MBA program, bridging the gap between emerging technology, business strategy, and human-centred change.

[linkedin.com/in/markcameron](https://www.linkedin.com/in/markcameron)



Vijayan Seenisamy

Enterprise AI Role Strategist, Woolworths

Vijayan created the AI Role Operating Framework™ (AI ROF™) and authored *The AI Delivery Manager Blueprint* (2025) and *The Pilot Trap: The Category Error Behind Enterprise AI Failure* (2026), the first books defining the discipline of enterprise agentic AI systems delivery. His work addresses why the vast majority of enterprise AI pilots fail to reach production, and what organisations must build to close that gap: governance structures, observability systems, and the human operating model for AI in production. Across retail, banking, telecom, public sector, and technology, including Accenture, Telstra, Deloitte, Bupa, and Woolworths Group, he has guided enterprises from agility to autonomy, specialising in AI delivery systems, governance frameworks, and operating model redesign.

<https://www.linkedin.com/in/vijayan-seenisamy/>



Vinod Bijlani

AI Practice Leader, HPE

Vinod is a distinguished inventor and thought leader with a career spanning over 22 years in data science and solution architecture. He holds 25 patents in AI/ML and IoT technologies and has architected smart city and intelligent transportation projects globally. A contributor to the World Economic Forum and the Forbes Technology Council, Vinod is widely recognized for his work on sovereign AI models and the development of sustainable, industrial-scale smart ecosystems.

<https://www.linkedin.com/in/vinodbijlani/>



Sharma Madiraju

Enterprise AI Strategist and Evangelist.

Ranked among the region's leading CIOs, Sharma is a distinguished technology executive with deep expertise in enterprise technology strategy and architecture consulting. He has advised C-suite leaders across financial services, defence, retail, logistics, and professional services, helping them build scalable, resilient, and future-ready technology foundations. Sharma's career is defined by his ability to engineer high-velocity decision systems and transform complex, legacy technology estates into adaptive, data-driven ecosystems that unlock value through technology-led transformation.

<https://www.linkedin.com/in/madiraju/>

This white paper reflects the personal views and professional perspectives of the authors. The opinions expressed do not necessarily represent the views of the organisations with which the authors are affiliated.



Executive Summary

Paper 1 in this series introduced the Intelligence-Centred Enterprise: an operating model built around a metabolic loop of Sense, Reason, Act, Learn. The response was immediate and consistent. Senior leaders across industries accepted the thesis, but asked the same question: how do we actually do it?

This paper answers that question with a diagnostic framework, not a prescription. Through analysis of enterprise AI programmes across financial services, healthcare, infrastructure, public sectors and retail, we have identified five predictable ways the metabolic loop breaks. Each maps to a specific failure mode: Loss of Situational Awareness (when Sense breaks), Decision Quality Collapse (when Reason breaks), Execution Bottleneck (when Act breaks), Organisational Amnesia (when Learn breaks), and Fragmented Intelligence (when no one owns the loop end to end).

The diagnosis reveals four organisational layers that must work together for ICE to become operational: Leadership and Governance, Enterprise Architecture, Infrastructure and Platforms, and a Delivery Discipline for AI systems at enterprise scale. The first three have established disciplines behind them. The fourth does not exist yet, not as a defined body of knowledge, not as a professional discipline with shared standards, not as something an organisation can hire for, train for, or certify against.

Each component is *healthy*. The system is *broken*.



This is the pattern we see in every organisation we work with. Individual functions perform well against their own metrics.

The enterprise outcome is dysfunction. The metabolic loop never actually runs because nobody owns it end to end.


The paper provides an **ICE Delivery Readiness Assessment** : a practical diagnostic that leadership teams can complete in 30 minutes. It covers four dimensions, one per layer.

The pattern is consistent : most organisations score reasonably on the first three dimensions, then score critically low on Delivery Discipline.

The dimension they have invested least in is the one that determines whether the other three produce results

After reading this paper, you can no longer attribute AI underperformance to technology failure. The failure modes are structural. The layers are identifiable. The gaps are diagnosable.

What remains is a leadership decision: fund the disciplines that do not yet exist or continue investing in technology your operating model cannot absorb



S E C T I O N 1

The Question Every Reader Asked

Paper 1 introduced the Intelligence-Centred Enterprise: an operating model built for learning velocity,

human purpose, and machine-scale scenario exploration. At its core is the metabolic loop, Sense, Reason, Act, Learn, the way an organisation converts information into improved performance. The argument was clear: every process in your organisation was designed before AI existed, and those processes are structurally incapable of absorbing what this technology can do.

The response to Paper 1 was immediate and consistent. Senior leaders across industries, from financial services to healthcare, from infrastructure to retail, asked the same question: How do we actually do it?

Not as a philosophical exercise. As an operational one. They had accepted the thesis. They could see the metabolic loop. They recognised the gap between how their organisations currently operate and what the technology now makes possible. What they wanted was a way to diagnose where, specifically, their organisation is stuck.

They had accepted the thesis. They could see the metabolic loop. They recognised the gap between how their organisations currently operate and what the technology now makes possible. What they wanted was a way to diagnose where, specifically, their organisation is stuck.

Paper Series Progression

1



Vision

What to build



2



Diagnostic

Why it stalls



3

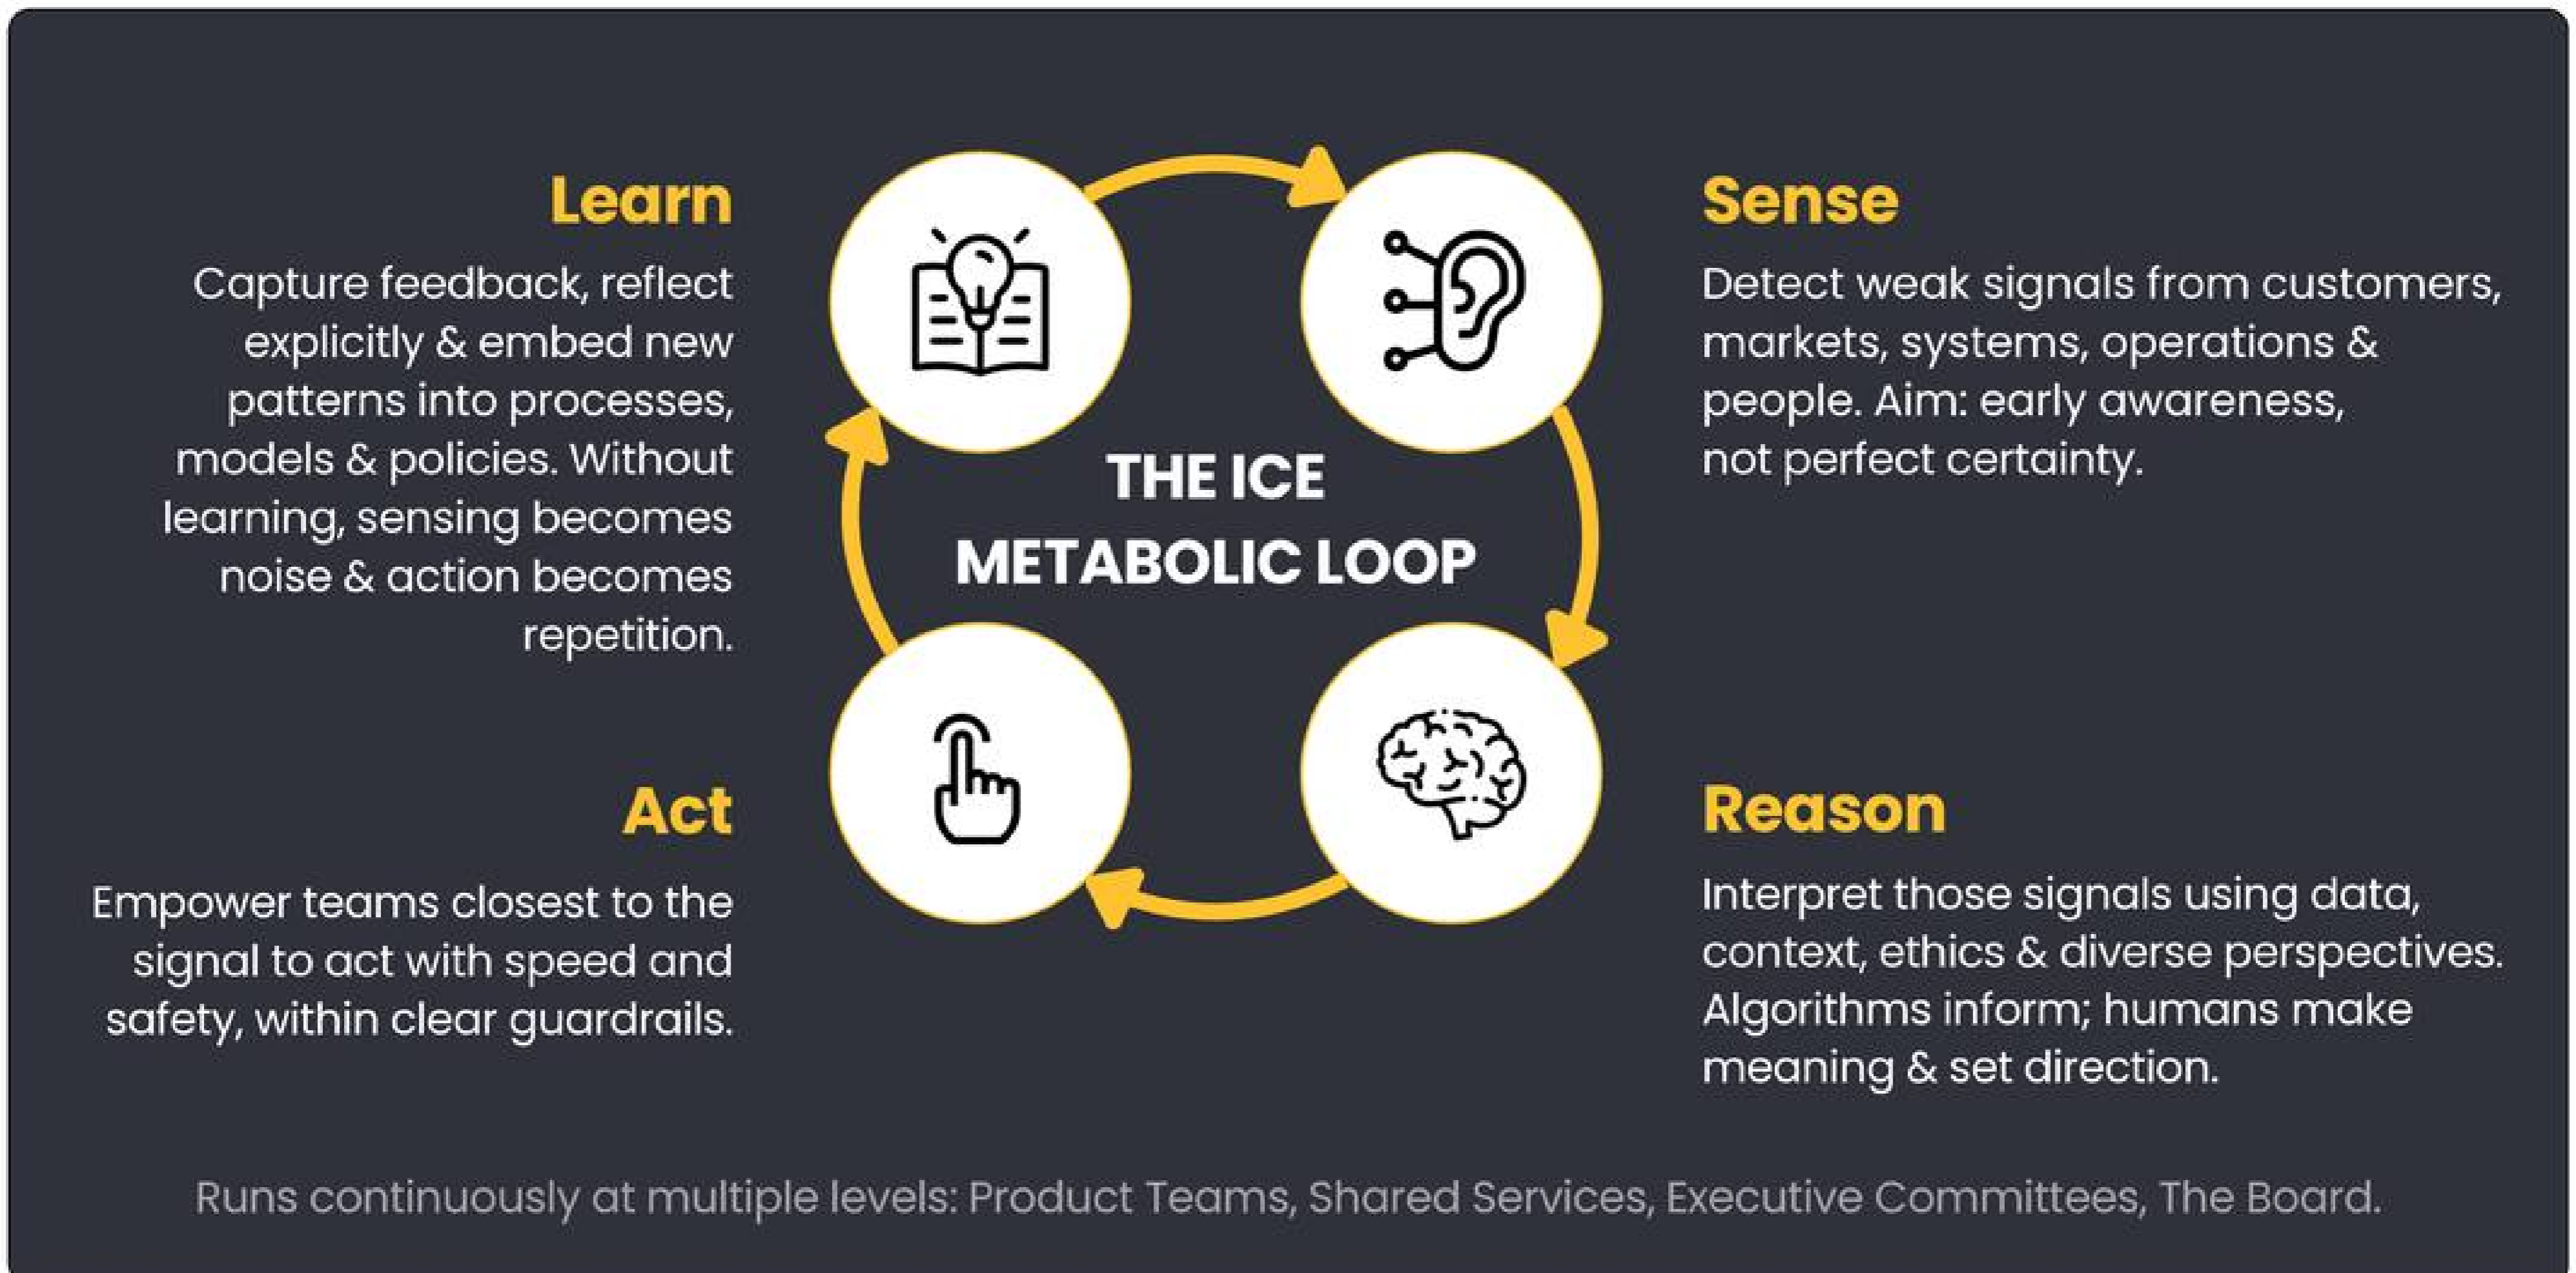


Implementation

How to build and run it

▲ YOU ARE HERE

This paper does not offer a prescription. The organisations asking this question are too varied, their contexts too specific, for a single answer to be credible. What it offers instead is a diagnostic framework: a structured way to identify where the metabolic loop breaks, which organisational layers are responsible, and what readiness looks like.



Whether you are the CIO wondering why your AI programmes stall despite sound technology choices, the CFO asking why significant AI investment has not produced visible returns, or the Chief AI Officer discovering that the role you were hired for has no operating model underneath it, the answer is the same: your organisation’s metabolic loop is broken in predictable places, and no one in your structure owns the full cycle.

This paper identifies the five predictable ways the loop breaks, maps the four layers required to fix it, and gives you a diagnostic to assess where your organisation stands today.



S E C T I O N 2

The Operating Model Gap

ADAPT research shows Australian CFOs estimate that 39% of technology capability spend is wasted. Not because the technology was wrong, but because it was pushed through structures designed for a different era. Capabilities are purchased, piloted, occasionally celebrated, and then quietly shelved. The waste is not in the buying. It is in the absorbing.

The structural barriers are well documented. Fewer than half of large enterprises can access operational data in real time. Half report that their data remains inaccessible to the teams that need it most. Forty percent are still running core processes on legacy systems that were never designed to integrate with modern AI capabilities. These are not technology problems waiting for better tools. They are architecture problems waiting for a different operating model.

McKinsey's analysis of digital transformation programmes puts the failure rate at 70%. That number has remained stubbornly consistent for over a decade, across industries and geographies, through multiple generations of technology. The technology improved. The failure rate did not. This is the clearest possible signal that the constraint is structural, not technical.

The consequences are visible at every level of the enterprise. Strategy is disconnected from the signals that should inform it. Governance checks compliance boxes but never asks what the organisation learned. AI programmes launch with energy and executive sponsorship, only to stall when they encounter the frozen middle: the layer of the organisation where innovation goes to be absorbed, diluted, and quietly deprioritised. Billions are spent on technology that sits on top of processes hostile to intelligence.

The instinct in most organisations is to respond by applying existing methods more rigorously. Better project management. Tighter governance. More oversight committees. This instinct is understandable. It is also the wrong response.

The methods are not failing because they are being applied poorly. They are failing because they were designed for a fundamentally different type of system. A governance framework built to approve capital expenditure cannot evaluate whether an AI system is learning. A project management methodology built for deterministic software delivery cannot manage a system whose behaviour changes after deployment. The tools are precise. They are precisely wrong.

We will return to this argument in Section 5, where we examine why a new delivery discipline is required. For now, the point is simpler.

A dark, dystopian scene with a cracked ceiling and glowing AI interface elements. In the center, a glowing cylinder contains the letters 'AI'. Surrounding it are several floating screens displaying various data visualizations, including network graphs, bar charts, and flowcharts. A person in a dark suit stands in the foreground, looking towards the AI interface. The overall atmosphere is one of technological decay and control.

***This is not an AI
problem.***

***It is an operating
model problem
that AI has made
impossible to ignore.***



S E C T I O N 3

5 Ways the Metabolic Loop Breaks

At the heart of ICE is a simple loop: Sense, Reason, Act, Learn. Most organisations over-rotate on Act. They are very busy, but not necessarily very intelligent.

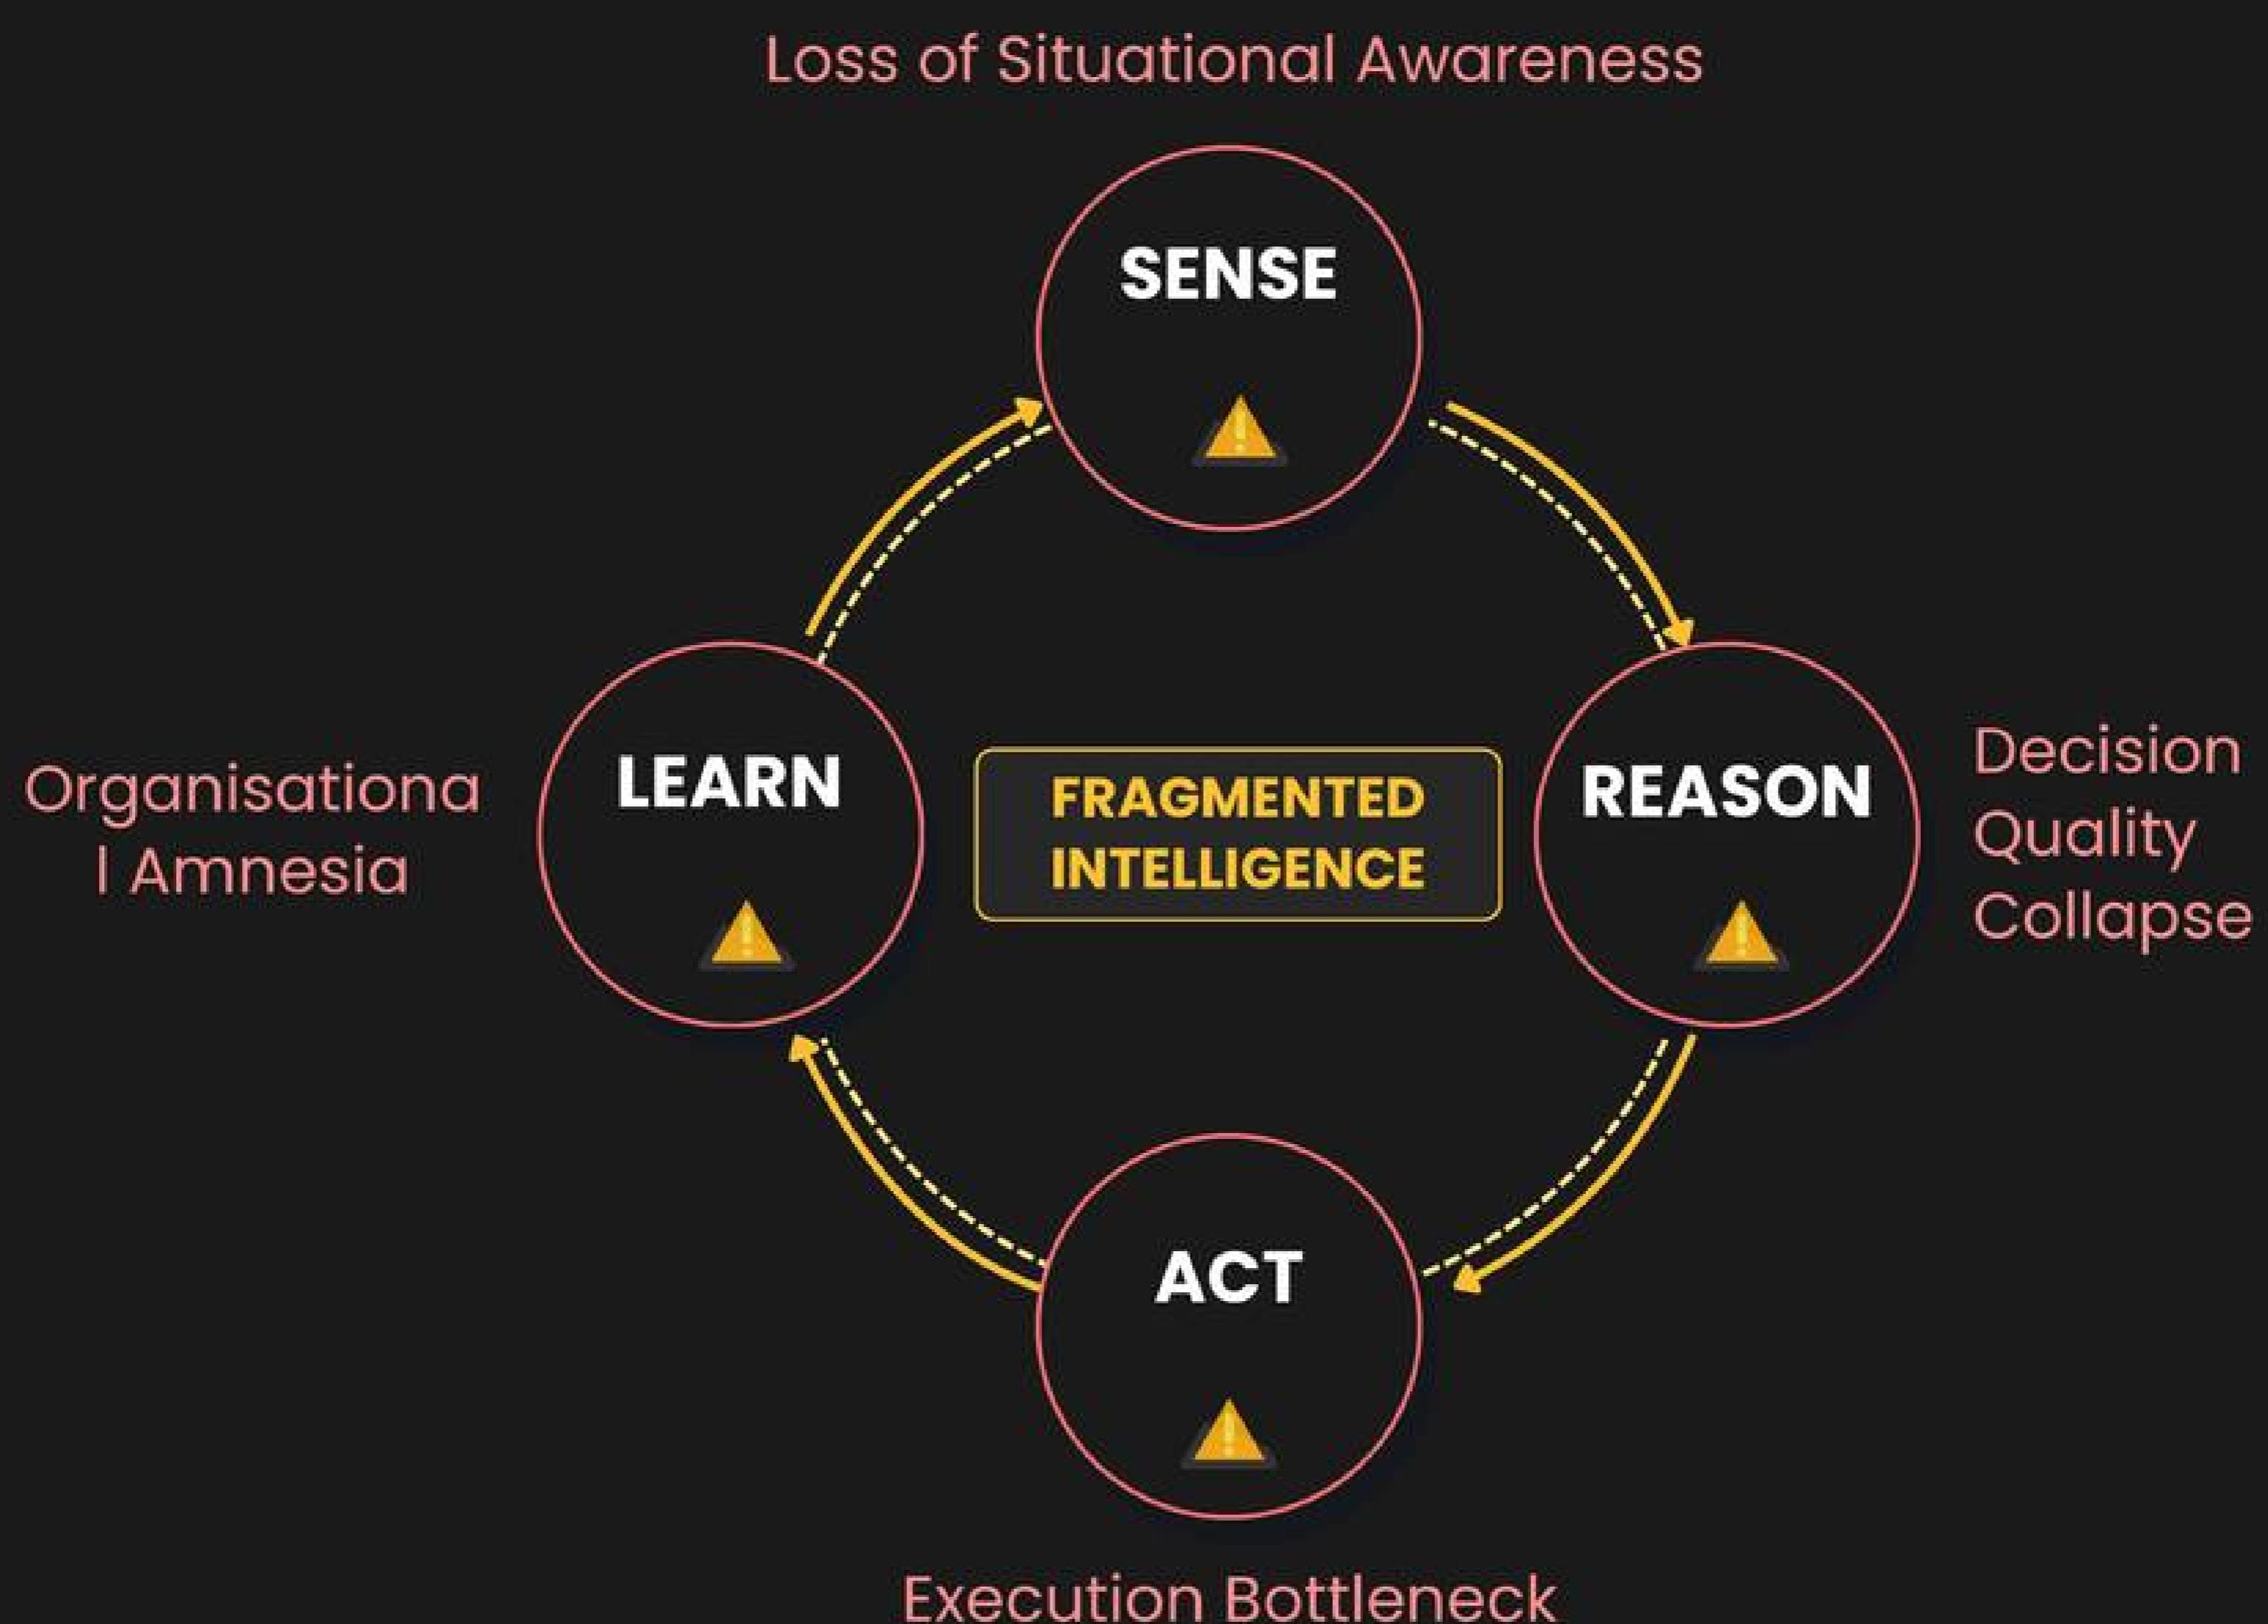
Paper 1 introduced the loop. This paper examines what happens when it breaks.

Through our combined experience across board advisory, enterprise architecture, infrastructure delivery, and AI programme leadership, we have identified five predictable failure modes.

Each maps to a specific break in the metabolic loop. Each produces recognisable symptoms. And each reinforces the others in ways that make diagnosis difficult unless you know what to look for.

The Metabolic Loop: Five Predictable Failure Modes

The loop breaks in five predictable places. Which ones are breaking in your organisation?



***When no one owns the full cycle,
each component is healthy.
The system is broken.***

Failure Mode 1: Loss of Situational Awareness

When Sense breaks



The organisation loses the ability to detect weak signals. Customer shifts, market moves, and operational anomalies go undetected because feedback loops are too slow, too filtered, or do not exist. The data may be there. The signals may be there. But the organisation's nervous system cannot carry them to the people who need them.

What this looks like in practice is leaders making decisions based on quarterly reports while the world moves weekly. Customer complaints surface months after the root cause began. Competitive moves are discovered through news articles, not internal intelligence.

A frontline team spots a pattern that could save millions, but the signal dies three management layers before it reaches anyone with authority to act on it.

The cost extends beyond missed opportunities. It is strategic decisions made on information that was already obsolete when it reached the boardroom. For the CFO, this appears as investments that looked sound at approval but failed because the market had already moved. For the CIO, it appears as technology programmes that solve yesterday's problem. For the CAIO, it appears as AI systems trained on data that no longer reflects reality

Failure Mode 2: Decision Quality Collapse

When Reason breaks




When was the last time your executive committee made a decision that surprised anyone in the room? If the answer is never, Reason is broken. The organisation has data but cannot make meaning from it. Scenario planning does not happen because the governance model has no mechanism for it. Decisions default to hierarchy and habit rather than evidence and exploration. The 60-slide deck is reviewed, discussed, and used to justify the conclusion that was reached before the meeting started. This failure mode is insidious because it is invisible to the people inside it. Decision quality collapse does not feel like a crisis. It feels like stability. The executive committee believes it is making sound decisions because it follows rigorous process. But rigorous process applied to incomplete information, filtered through organisational politics, evaluated without diverse challenge, produces decisions that are internally consistent but externally wrong.

In a board-level case observed through advisory work, an AI proposal was forced through the same decision template as a major systems upgrade: fixed ROI, fixed scope, vendor chosen up front. The committee spent months debating assumptions, then approved a “safe” pilot; it worked in the lab but died in the business because data access, customer consent and frontline workflow changes were never owned in the decision rights. The technology wasn’t the blocker, the process optimised for certainty and control, so it starved the initiative of learning and accountability.

AI compounds this failure mode. Organisations invest in systems that generate insights at machine speed, then route those insights through governance structures designed for quarterly review. The intelligence exists. The organism cannot metabolise it.

Failure Mode 3: Execution Bottleneck

When Act breaks



A financial services firm invested \$22 million in AI over eighteen months. Fourteen initiatives launched. Three reached production. The technology in all fourteen was sound. The operating model could only absorb three.

A metropolitan transport authority invested in an AI programme that combined real-time traffic prediction with an agentic response layer, a system designed not just to forecast congestion but to automatically generate and dispatch response plans: signal timing adjustments, incident crew deployment, passenger information updates.

The prediction capability was delivered and accuracy rates exceeded the business case assumptions; the operations centre could see the value clearly.

The response automation never reached production.

Legal counsel could not resolve who carried liability when an automated response plan contributed to an incident. The programme was scaled back to prediction-only. It is not the outcome the business case was written for, and the gap between the two represents a significant portion of the original investment

The most common failure mode, and the most visible. Initiatives succeed in controlled environments. Leadership celebrates. The initiative then enters the broader organisation, where it collides with legacy processes, competing priorities, and a middle management layer that was never consulted and has no incentive to accommodate it.

Paper 1 named this the frozen middle: the organisational layer where innovation goes to be absorbed, diluted, and quietly deprioritised. The frozen middle is not malicious. It is rational. Treat it as the organisation's immune system: unless incentives and safety around change are redesigned, the immune system will do its job and neutralise the foreign body.

Middle managers are measured on stability, efficiency, and predictability. They are being asked to adopt systems that are, by nature, unstable, expensive to learn, and unpredictable. The incentive structure makes resistance the logical response.

The direct cost is stranded investment. The indirect cost is worse: talent attrition. Your best people leave because they cannot get things done. They join organisations that have figured out how to absorb innovation, and the gap widens.



Failure Mode 4: Organisational Amnesia

When Learn breaks



The same integration mistake happens in three consecutive AI deployments. A risk identified and documented in one programme reappears in the next because the knowledge was never transferred. Post-implementation reviews are scheduled but never happen or happen and produce reports that nobody reads.

This is what it looks like when an organisation cannot learn.

Governance checks compliance but never asks the question that matters: what did we learn?

The organisation generates enormous amounts of data about its own performance but has no mechanism to convert that data into institutional learning. Knowledge stays trapped in teams and individuals. When those individuals leave, the knowledge leaves with them.

Every repeated mistake has a direct cost. The compounding cost is worse. An organisation that cannot learn falls further behind with every cycle. Its competitors, the ones that have built the learning infrastructure, accelerate. The gap between them does not close. It compounds.

Failure Mode 5: Fragmented Intelligence

When no one owns the loop



The data team delivers excellent insights. The strategy team makes reasonable decisions.

The project teams execute competently. The governance function monitors risk diligently and yet the enterprise outcome is dysfunction.

The insights do not reach decision-makers in time

Decisions do not translate into execution priorities.

Execution results do not feed back into strategy.

Governance reviews what happened but cannot influence what happens next. Each function performs well against its own metrics. The metabolic loop never actually runs..

Each component is *healthy*.

The system is *broken*.

This is the integrating failure, and it is the hardest to diagnose because no single function's dashboard will reveal it. The CIO sees healthy programme delivery metrics. The CFO sees approved business cases tracking to plan.

The CAIO sees capable AI systems producing valid outputs. The failure is between them, in the white space of the organisation chart, where nobody is accountable for the loop running end to end.

The Chief AI Officer was hired to solve this.

But she has no operating model to work from, no governance framework designed for it, and no delivery discipline that connects the layers. She was given a title and told to make intelligence work across the enterprise. She was not given the infrastructure to do it.

The cost of fragmented intelligence is invisible in any single function's metrics but catastrophic at the enterprise level. It is the difference between an organisation that learns and one that merely accumulates experience.

The Reinforcing Spiral

These five failure modes are not independent. They reinforce each other.

An organisation that cannot sense will make poor decisions. Poor decisions create execution bottlenecks. Bottlenecks prevent learning. And without learning, sensing degrades further. The loop does not just break in one place. It spirals downward. This is why isolated fixes fail. Investing in better data infrastructure (Sense) without addressing decision-making structures (Reason) produces more data that nobody acts on. Streamlining execution (Act) without building learning mechanisms (Learn) produces faster repetition of the same mistakes.

Hiring a Chief AI Officer to own the loop (Fragmented Intelligence) without building the governance, architecture, and delivery discipline to support that role produces an expensive new title with no operating model underneath it.

Failure Mode Summary

Sense breaks

FAILURE MODE

WHAT IT LOOKS LIKE

WHAT IT COSTS

Loss of Situational Awareness

Leaders decide on quarterly reports while the world moves weekly. Weak signals die before reaching anyone who can act.

Strategic decisions made on information already obsolete when it reached the boardroom.

Reason breaks

FAILURE MODE

WHAT IT LOOKS LIKE

WHAT IT COSTS

Decision Quality Collapse

The executive committee reviews a 60-slide deck and makes the same decision they would have made without it.

Slow, low-quality decisions compound. Within two years, the decision architecture reflects the past, not the present.

Act breaks

FAILURE MODE

WHAT IT LOOKS LIKE

WHAT IT COSTS

Execution Bottleneck

14 AI initiatives launch, 3 reach production. The frozen middle blocks scaling. The technology in all fourteen was sound.

*Stranded investment.
Talent attrition : Best people leave because they cannot get things done.*

Learn breaks

FAILURE MODE

WHAT IT LOOKS LIKE

WHAT IT COSTS

Organisational Amnesia

The same integration mistake happens in three consecutive deployments. Knowledge stays trapped in teams, never embedded in systems.

*Every repeated mistake has a direct cost.
The compounding cost : Competitors accelerate while you repeat.*

No loop owner

FAILURE MODE

WHAT IT LOOKS LIKE

WHAT IT COSTS

Fragmented Intelligence

Each function optimises locally. Insights do not reach decision-makers. Decisions do not translate into priorities. Results do not feed back

***Each component is healthy.
The system is broken.***



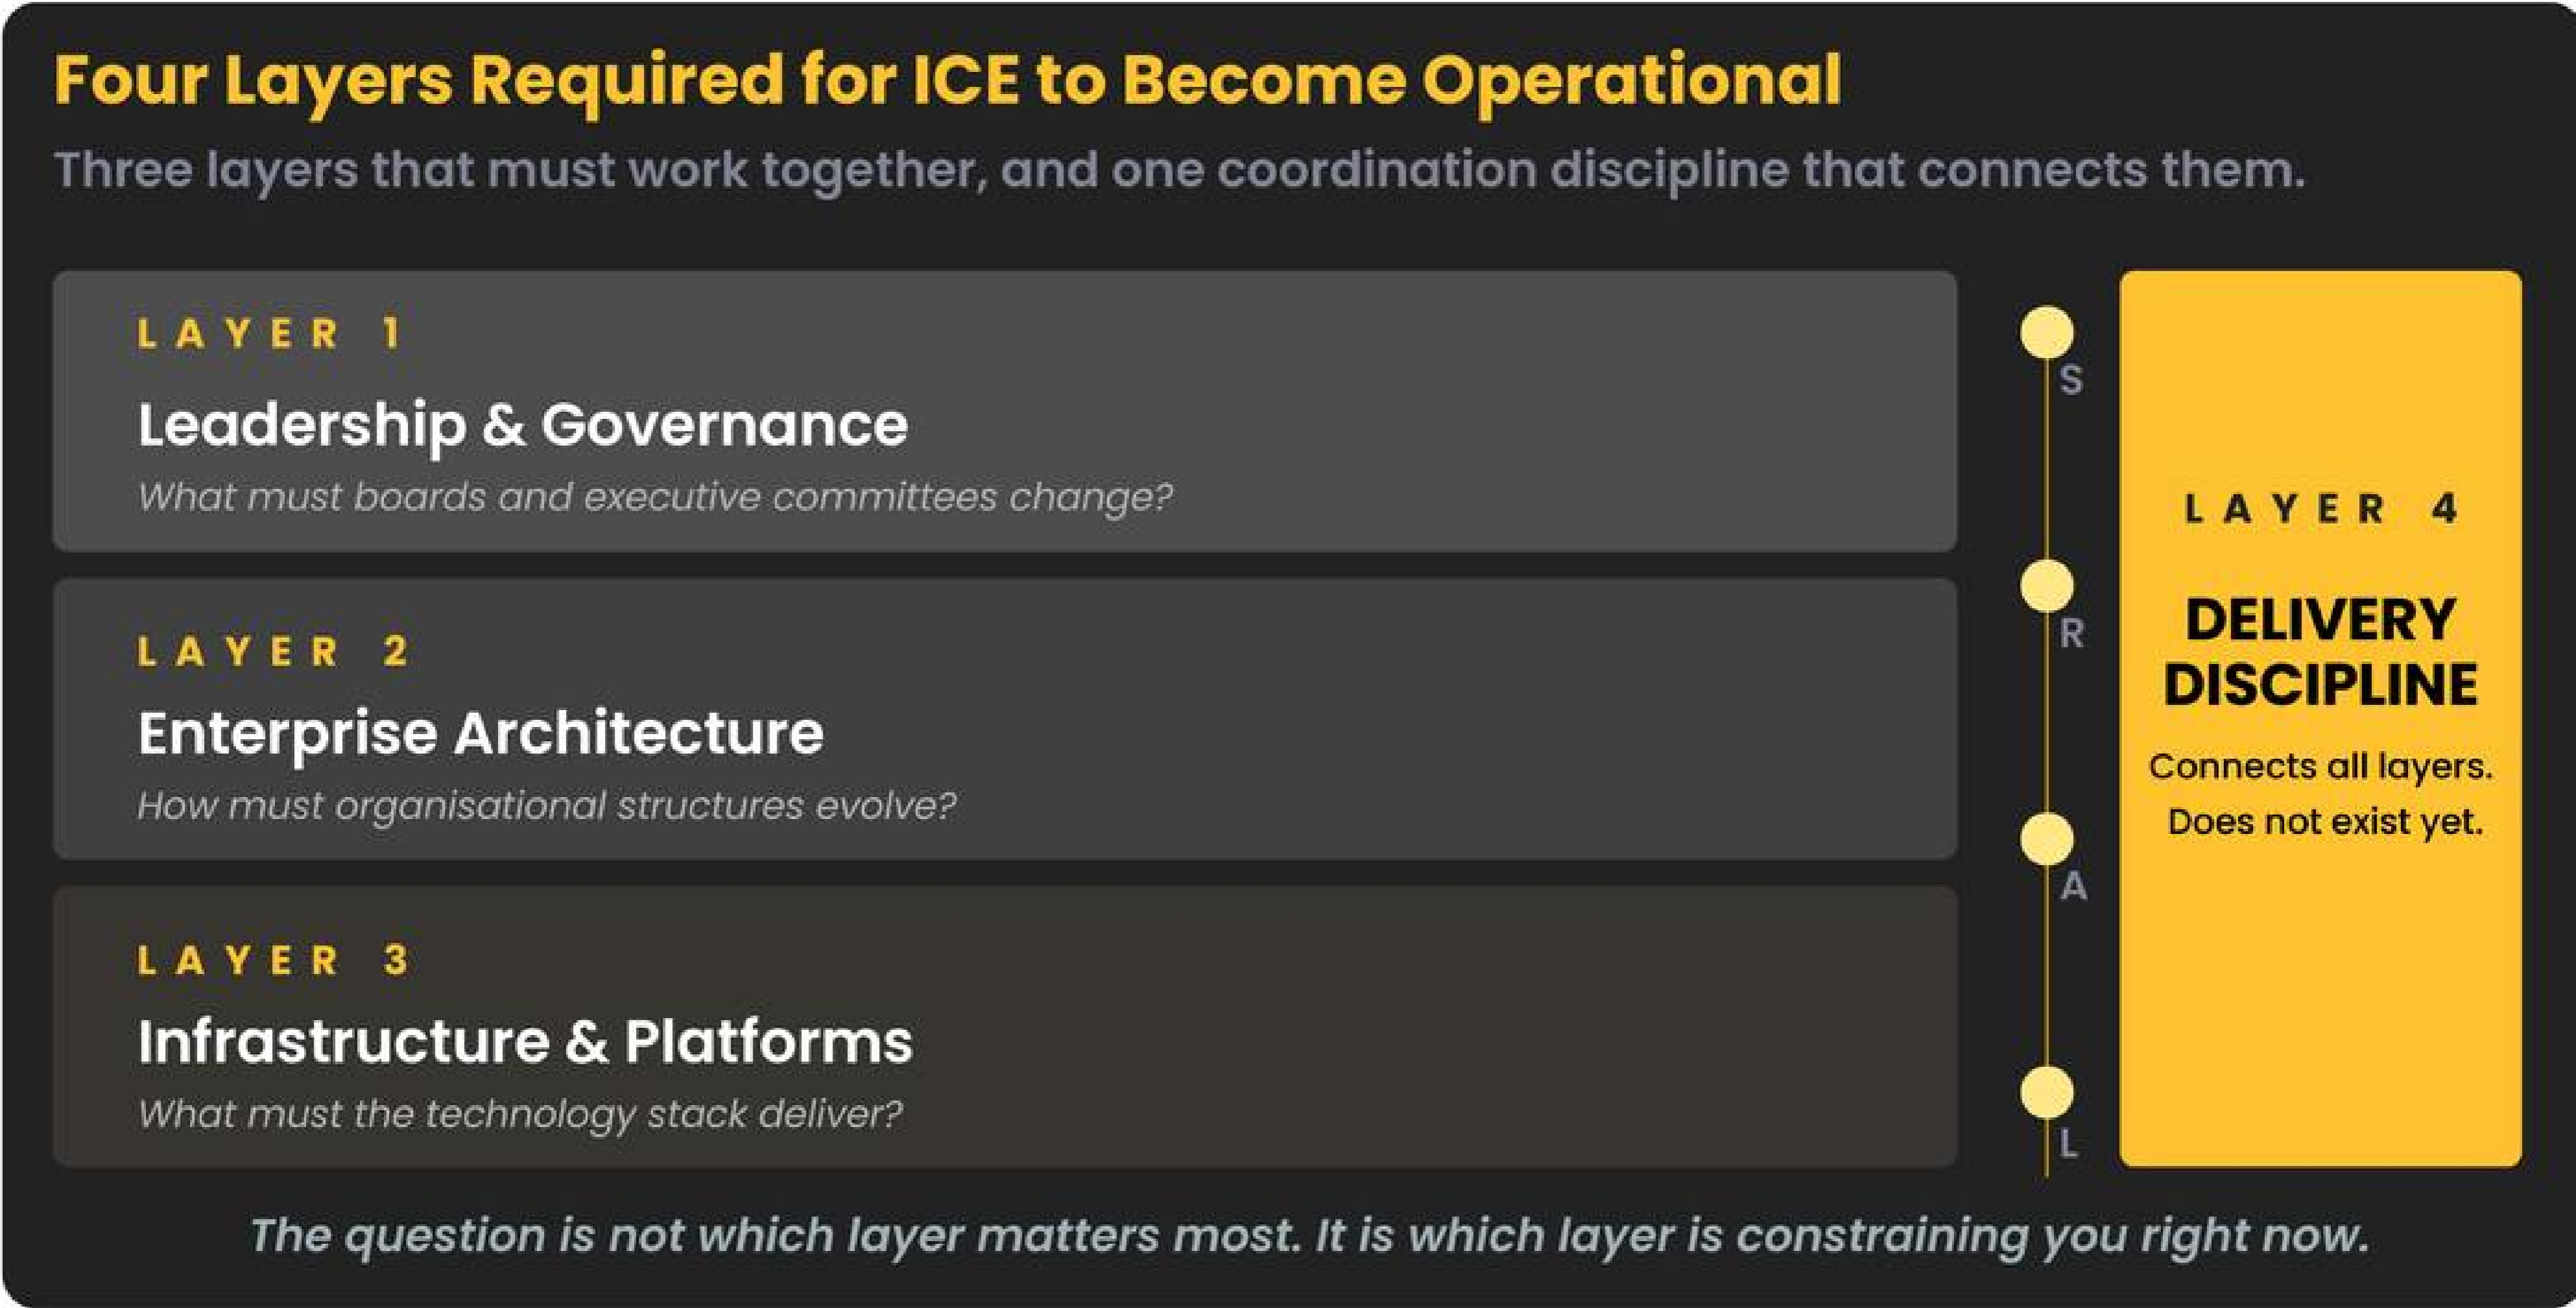
S E C T I O N 4

Four Layers, Four Constraints

The metabolic loop breaks because four layers that should work together do not. Each layer, when missing or misaligned, creates a specific constraint that prevents ICE from becoming operational. No single layer is sufficient. All four are necessary.

What follows is not a capability maturity model.

It is a constraint analysis. For each layer, we describe what happens when it is absent. The framing is deliberate: leaders respond to the cost of inaction faster than they respond to the promise of a capability.



Layer 1 : Leadership and Governance

Most boards review AI the same way they review ERP implementations: on time, on budget, within scope. This is the wrong lens for a system that changes its own behaviour after deployment. Without leadership and governance redesign, AI investment is governed by rules designed for capital expenditure.

Approval gates that worked for infrastructure projects create bottlenecks for initiatives that need to iterate weekly. Risk frameworks that categorise systems as “compliant” or “non-compliant” cannot evaluate a system whose compliance status shifts with each model update.

The board believes it has oversight. It has oversight of the wrong things.

What this layer requires is a fundamental shift in governance posture. Boards that ask “what did we learn?” before “did we follow the plan?” Executive committees that operate as learning systems, not approval gates. Risk functions that can evaluate probabilistic systems, not just deterministic ones. And an explicit accountability structure for AI portfolio health, beyond individual AI project delivery.

What “good” looks like

In one organisation, the board shifted AI oversight from quarterly project updates to a monthly AI portfolio health review. Every initiative had to show three things: adoption (who is actually using it), drift and risk (what changed since last review), and learning (what was tried, what was stopped, and what will change next). That one change cut the time from insight to decision and made it safe to stop low-value work early without blame.

The key is recognising that well implemented AI should be focused on improving decision velocity, by increasing transparency and accountability across the board.

What happens when it doesn't work well

Where that shift didn't happen, AI was governed like ERP: annual funding, stage-gates, and green status reporting. Pilots were declared successful, then sat in a queue for risk, procurement and architecture approvals until sponsors moved on and the operating environment changed; by the time go-live arrived, nobody owned adoption. The board saw a long list of “delivered” projects and very little operating impact, because governance measured delivery, not learning or behaviour change.

THE HONEST CHALLENGE

Governance transformation requires leaders to relinquish control they have spent careers building. The frozen middle, named in Paper 1, exists because middle management was optimised for a control-based operating model. Asking those same managers to enable speed, experimentation, and tolerance for uncertainty is not a training problem. It is a structural incentive problem that most organisations have not yet confronted.

Layer 2: Enterprise Architecture

Paper 1 introduced the concept of **ICE Pods**: cross-functional teams organised around value streams rather than functional silos. Claims Resolution Pods, Student Experience Pods, Chronic Care Pods. The operating model makes sense on a whiteboard. The question is whether the organisation's architecture can actually support it.

In most large enterprises, the answer is no. Technology capability models still mirror the organisational chart. Data flows follow departmental boundaries. Integration patterns were designed for batch processing, not the real-time signal flows that the metabolic loop requires. Intelligence stays trapped in functional silos because the architecture was built to keep it there.

What this layer requires is a transition architecture that bridges legacy and target state simultaneously. Not a five-year transformation programme that replaces everything. A pragmatic approach that identifies where cross-functional pods can form now, which integration constraints must be resolved first, and how the organisation operates in both modes during the transition.

One of the most technologically advanced transportation authorities in Asia set out to build an integrated incident response platform, a system designed to detect disruptions across all modes simultaneously and coordinate a unified response in real time.

The organisation operated one of the region's most complex urban mobility networks. What the programme exposed was an architectural reality that had never needed to surface before: traffic management, bus operations, and rail systems had each been built as sovereign technology domains. They shared a city but not a data model, not an integration layer, and not a common definition of what an "incident" even meant.

Each mode had optimised for its own operational excellence. Nobody had ever needed to view them as a single network, until AI made that a requirement. The transition architecture work that followed was not a technology problem. It was an organisational one: aligning three operationally independent entities, each with its own governance, budget, and performance metrics, around the premise that their data needed to behave as if they were one.

Enterprise restructuring is the slowest layer to change. It involves every function, every reporting line, every budget allocation. This is why transition architecture matters more than target-state architecture. An organisation that waits until the target architecture is complete before operationalising ICE will wait indefinitely.

Layer 4: The Delivery Discipline

The fourth layer is the one that does not currently exist anywhere in the industry. It is the delivery discipline for taking AI systems from pilot to reliable production at enterprise scale. Without it, the other three layers may each be individually sound, but the system as a whole has no coordination, no measurement framework designed for probabilistic systems, and no one accountable for the full cycle. Section 5 develops this argument in full.

The Interdependence Problem

These four layers are not independent workstreams to be managed in parallel. They are interdependent constraints.

Leadership that governs without understanding infrastructure will make uninformed decisions. Architecture that designs without a delivery discipline will create structures nobody can operate. Infrastructure that builds without governance will create capability nobody controls. And a delivery discipline without leadership support will remain an unfunded aspiration.

The question is not which layer matters most. It is which layer is constraining you right now.

LOOP BREAK	FAILURE MODE	PRIMARY LAYER	SUPPORTING LAYERS
Sense breaks	Loss of Situational Awareness	Infrastructure & Platforms	Enterprise Architecture
Reason breaks	Decision Quality Collapse	Leadership & Governance	Enterprise Architecture
Act breaks	Execution Bottleneck	Enterprise Architecture	Leadership & Governance; Infrastructure
Learn breaks	Organisational Amnesia	Delivery Discipline	Leadership & Governance
No loop owner	Fragmented Intelligence	Delivery Discipline	All three layers

Delivery Discipline is the primary layer for both learning failures and integration failures. This is why its absence is catastrophic.



S E C T I O N 5

The Missing Layer: Delivery Discipline

Of the four layers mapped in Section 4, three have established disciplines behind them. Leadership and governance has decades of board effectiveness research and frameworks. Enterprise architecture has TOGAF, capability modelling, and organisational design methodologies refined over thirty years. Infrastructure has platform engineering, MLOps, and cloud architecture patterns that are well documented and widely taught.

The fourth layer, the delivery discipline for AI systems at enterprise scale, does not exist yet.

Not as a defined body of knowledge. Not as a professional discipline with shared standards. Not as something an organisation can hire for, train for, or certify against. Individual practitioners are inventing it inside their organisations, but there is no shared language, no common framework, and no agreed set of capabilities. This is not a gap that will close by itself. It is a structural absence that explains why the other three layers, even when individually strong, fail to produce enterprise outcomes.

The Category Error

In Section 2, we noted that most organisations respond to AI delivery failure by applying existing methods more rigorously. It is worth understanding precisely why this is the wrong response.

Organisations are treating AI delivery as a project management problem. It is a system stewardship problem. This is a category error.

Traditional delivery assumes a defined end state. Define requirements, build to specification, test against specification, deploy. The success metric is on-time, on-budget, to-specification. This works well for deterministic systems, where the same input reliably produces the same output, where failure is visible and reproducible, and where the system behaves the same way after release as it did during testing.

AI systems do not work this way. They drift. They produce different outputs from the same inputs depending on context, data state, and upstream conditions. They degrade gradually rather than failing cleanly. Their behaviour changes as the data they learn from changes. There is no defined end state because the system is never finished. It is always becoming.

Applying project management to this type of system is not just ineffective. It creates a false confidence that is more dangerous than having no method at all. The project closes on time. The governance review signs off. The system enters production. And then it begins to change, in ways nobody is watching, using metrics nobody defined, with consequences nobody owns.

From Project Management to System Stewardship

If the category error is applying project management to a system that does not behave like a project, the question becomes: what discipline should replace it?

We use the term system stewardship: the ongoing practice of monitoring, evaluating, governing, and adapting AI systems in production. The distinction from project management is not cosmetic. It reflects a fundamentally different set of assumptions about what the work is. Project management assumes a defined end state. You plan, build, test, and deploy. The project closes. The team moves on. Success is measured at the point of delivery.

System stewardship assumes a living system. There is no defined end state because the system is never finished. It is always becoming. Deployment is not the finish line. It is the first day of school, not graduation. The system will drift. Its costs will fluctuate. Its outputs will vary. The data it depends on will change. The regulations that apply to it will evolve. Success is not measured at the point of delivery. It is measured continuously, across multiple dimensions, for as long as the system operates.

This is not a subtle distinction. It changes everything about how the work is structured. The question shifts from “did we deliver on time and on budget?” to “is the system healthy, safe, economically viable, and learning?” The accountability shifts from a project manager who owns a timeline to a steward who owns system health. The governance shifts from milestone-based review to continuous evaluation. The economics shift from project budget to unit economics that must be tracked at the level of individual tasks, not annual allocations.

For the CIO, this means the delivery function must be redesigned, not retrained. Adding AI to a project manager’s responsibilities does not create stewardship. It creates a project manager with an impossible job.

For the CFO, this means the investment model must change. AI systems do not have a capital expenditure profile with a defined return period. They have an operating expenditure profile that must be governed continuously. The business case does not close at go-live. It must be re-evaluated as the system’s behaviour and economics change.

For the Chief AI Officer, this is the operating model she has been looking for. System stewardship is the discipline that connects the other three layers: it tells leadership what to govern, architecture what to support, and infrastructure what to measure. Without it, each layer optimises independently. With it, the metabolic loop has a chance of actually running.

The concept is clear. The question is what it requires in practice.

What This Discipline Must Include

We are not proposing a rebrand of existing project management. The delivery discipline required for enterprise AI is structurally different. At minimum, it must address five requirements that traditional delivery does not.

First, continuous evaluation rather than checkpoint-based testing. Deterministic systems can be tested at milestones. Probabilistic systems must be evaluated continuously because their behaviour changes between milestones. The discipline must define what to measure, how often, and what constitutes an unacceptable change.

Second, economic governance that operates at task level, not project level. A traditional project has a budget. An AI system has a cost-per-task that fluctuates with usage patterns, model updates, and upstream conditions. A system that was economically viable at launch can become unviable within weeks. The discipline must make this visible before the business case inverts.

Third, governance that adapts to the risk profile of the system, not a static compliance framework. A customer-facing AI system that makes autonomous decisions requires different governance from an internal summarisation tool. The discipline must match governance rigour to system autonomy and must evolve that governance as the system's capabilities change.

Fourth, role clarity for a type of work that does not yet have defined roles. Who owns the ongoing health of an AI system in production? Not the development team, which has moved on. Not the operations team, which was not built for probabilistic systems. Not the business owner, who lacks technical visibility. The discipline must define who is accountable for system health after deployment, and what authority they carry.

Fifth, coordination across all four layers simultaneously. This is the integrating function that Section 3's fifth failure mode, Fragmented Intelligence, demands. The delivery discipline does not sit beneath the other three layers. It runs through them, ensuring that leadership decisions, architecture choices, and infrastructure capabilities are connected to each other and to the metabolic loop.

AN HONEST ASSESSMENT

We should be direct about the current state. This delivery discipline is being invented in real time. No single organisation, consulting firm, or technology provider has codified it fully. What exists are fragments: individual practitioners building local solutions, emerging frameworks being tested in specific contexts, and a growing body of evidence about what works and what does not.

What we can say with confidence is that the discipline is needed, that its absence is the primary driver of the Fragmented Intelligence failure mode, and that its core requirements are becoming clearer through practice. Paper 3 in this series will move from diagnosis to implementation, proposing specific structures, roles, and governance patterns for this discipline.

For now, the practical question for any CIO, CFO, or Chief AI Officer reading this paper is not whether they need this discipline. It is whether they have built even the foundations for it. Section 6 provides a tool to assess exactly that.



S E C T I O N 6

ICE Delivery Readiness Assessment

The preceding analysis identifies the failure modes and the layers required to address them. This section provides a practical diagnostic: where does your organisation stand today?

The assessment below covers four dimensions, one for each layer. Each dimension has five statements. Score each statement from 1 (not started) to 5 (embedded and functioning). The tool is designed to be completed by a leadership team in 30 minutes.

Have each member score independently, then compare. The disagreements are more valuable than the scores. Where leaders see the same gap, you have consensus to act. Where they disagree, you have found a blind spot that needs investigation.

ICE Delivery Readiness Assessment

D1 Leadership & Governance

/ 25

- Our board regularly reviews AI portfolio health, not just AI project status 1 2 3 4 5
- Our executive committee asks "what did we learn?" as a standard agenda item 1 2 3 4 5
- Decision rights for AI systems are explicitly defined and documented 1 2 3 4 5
- Our risk function can evaluate probabilistic systems, not just deterministic ones 1 2 3 4 5
- The frozen middle has been actively addressed through role redesign or incentive change 1 2 3 4 5

D2 Enterprise Architecture

/ 25

- We have cross-functional pods or teams organised around value streams, not just functions 1 2 3 4 5
- Our technology capability model reflects AI-native requirements 1 2 3 4 5
- We have a transition architecture that lets us operate in both legacy and target modes 1 2 3 4 5
- Our organisational structure supports rapid reallocation of talent to emerging AI opportunities 1 2 3 4 5
- Role definitions have evolved beyond static job descriptions toward capability-based models 1 2 3 4 5

D3 Infrastructure & Platforms

/ 25

- Our data infrastructure supports real-time signal detection, not just batch reporting 1 2 3 4 5
- We have observability across our AI systems (performance, cost, compliance) 1 2 3 4 5
- Our compute architecture can scale AI workloads without unsustainable cost increases 1 2 3 4 5
- We have reduced dependency on legacy platforms for mission-critical AI workflows 1 2 3 4 5
- Our infrastructure team understands AI-specific requirements (latency, throughput, cost-per-inference) 1 2 3 4 5

D4 Delivery Discipline

/ 25

- Someone in our organisation is explicitly accountable for AI systems in production (not just projects) 1 2 3 4 5
- We measure AI system health continuously, not just at deployment 1 2 3 4 5
- We have governance frameworks that account for AI's probabilistic nature (drift, cost variance, unexpected behaviour) 1 2 3 4 5
- Our scaling process is defined: we know what evidence is needed before moving from pilot to production 1 2 3 4 5
- Feedback from production AI systems is systematically captured and used to improve future deployments 1 2 3 4 5

1 Not Started 2 Aware 3 Developing 4 Established 5 Embedded

Score Interpretation

- ✔ **20-25: Embedded** - Functioning focus on sustaining and optimising
- ✔ **15-19: Developing** - Layer exists but is not yet reliable. Close specific gaps
- ✔ **8-14: Emerging** - Early efforts, no systematic approach. This layer is constraining ICE adoption
- ✔ **5-7: Not started** - This is a critical gap. Begin here

The Pattern You Will See

- Most organisations score between 12 and 18 in the first three dimensions.
- They typically score between 5 and 8 on Delivery Discipline.
- The dimension they have invested least in is the one that determines whether the other three produce results.

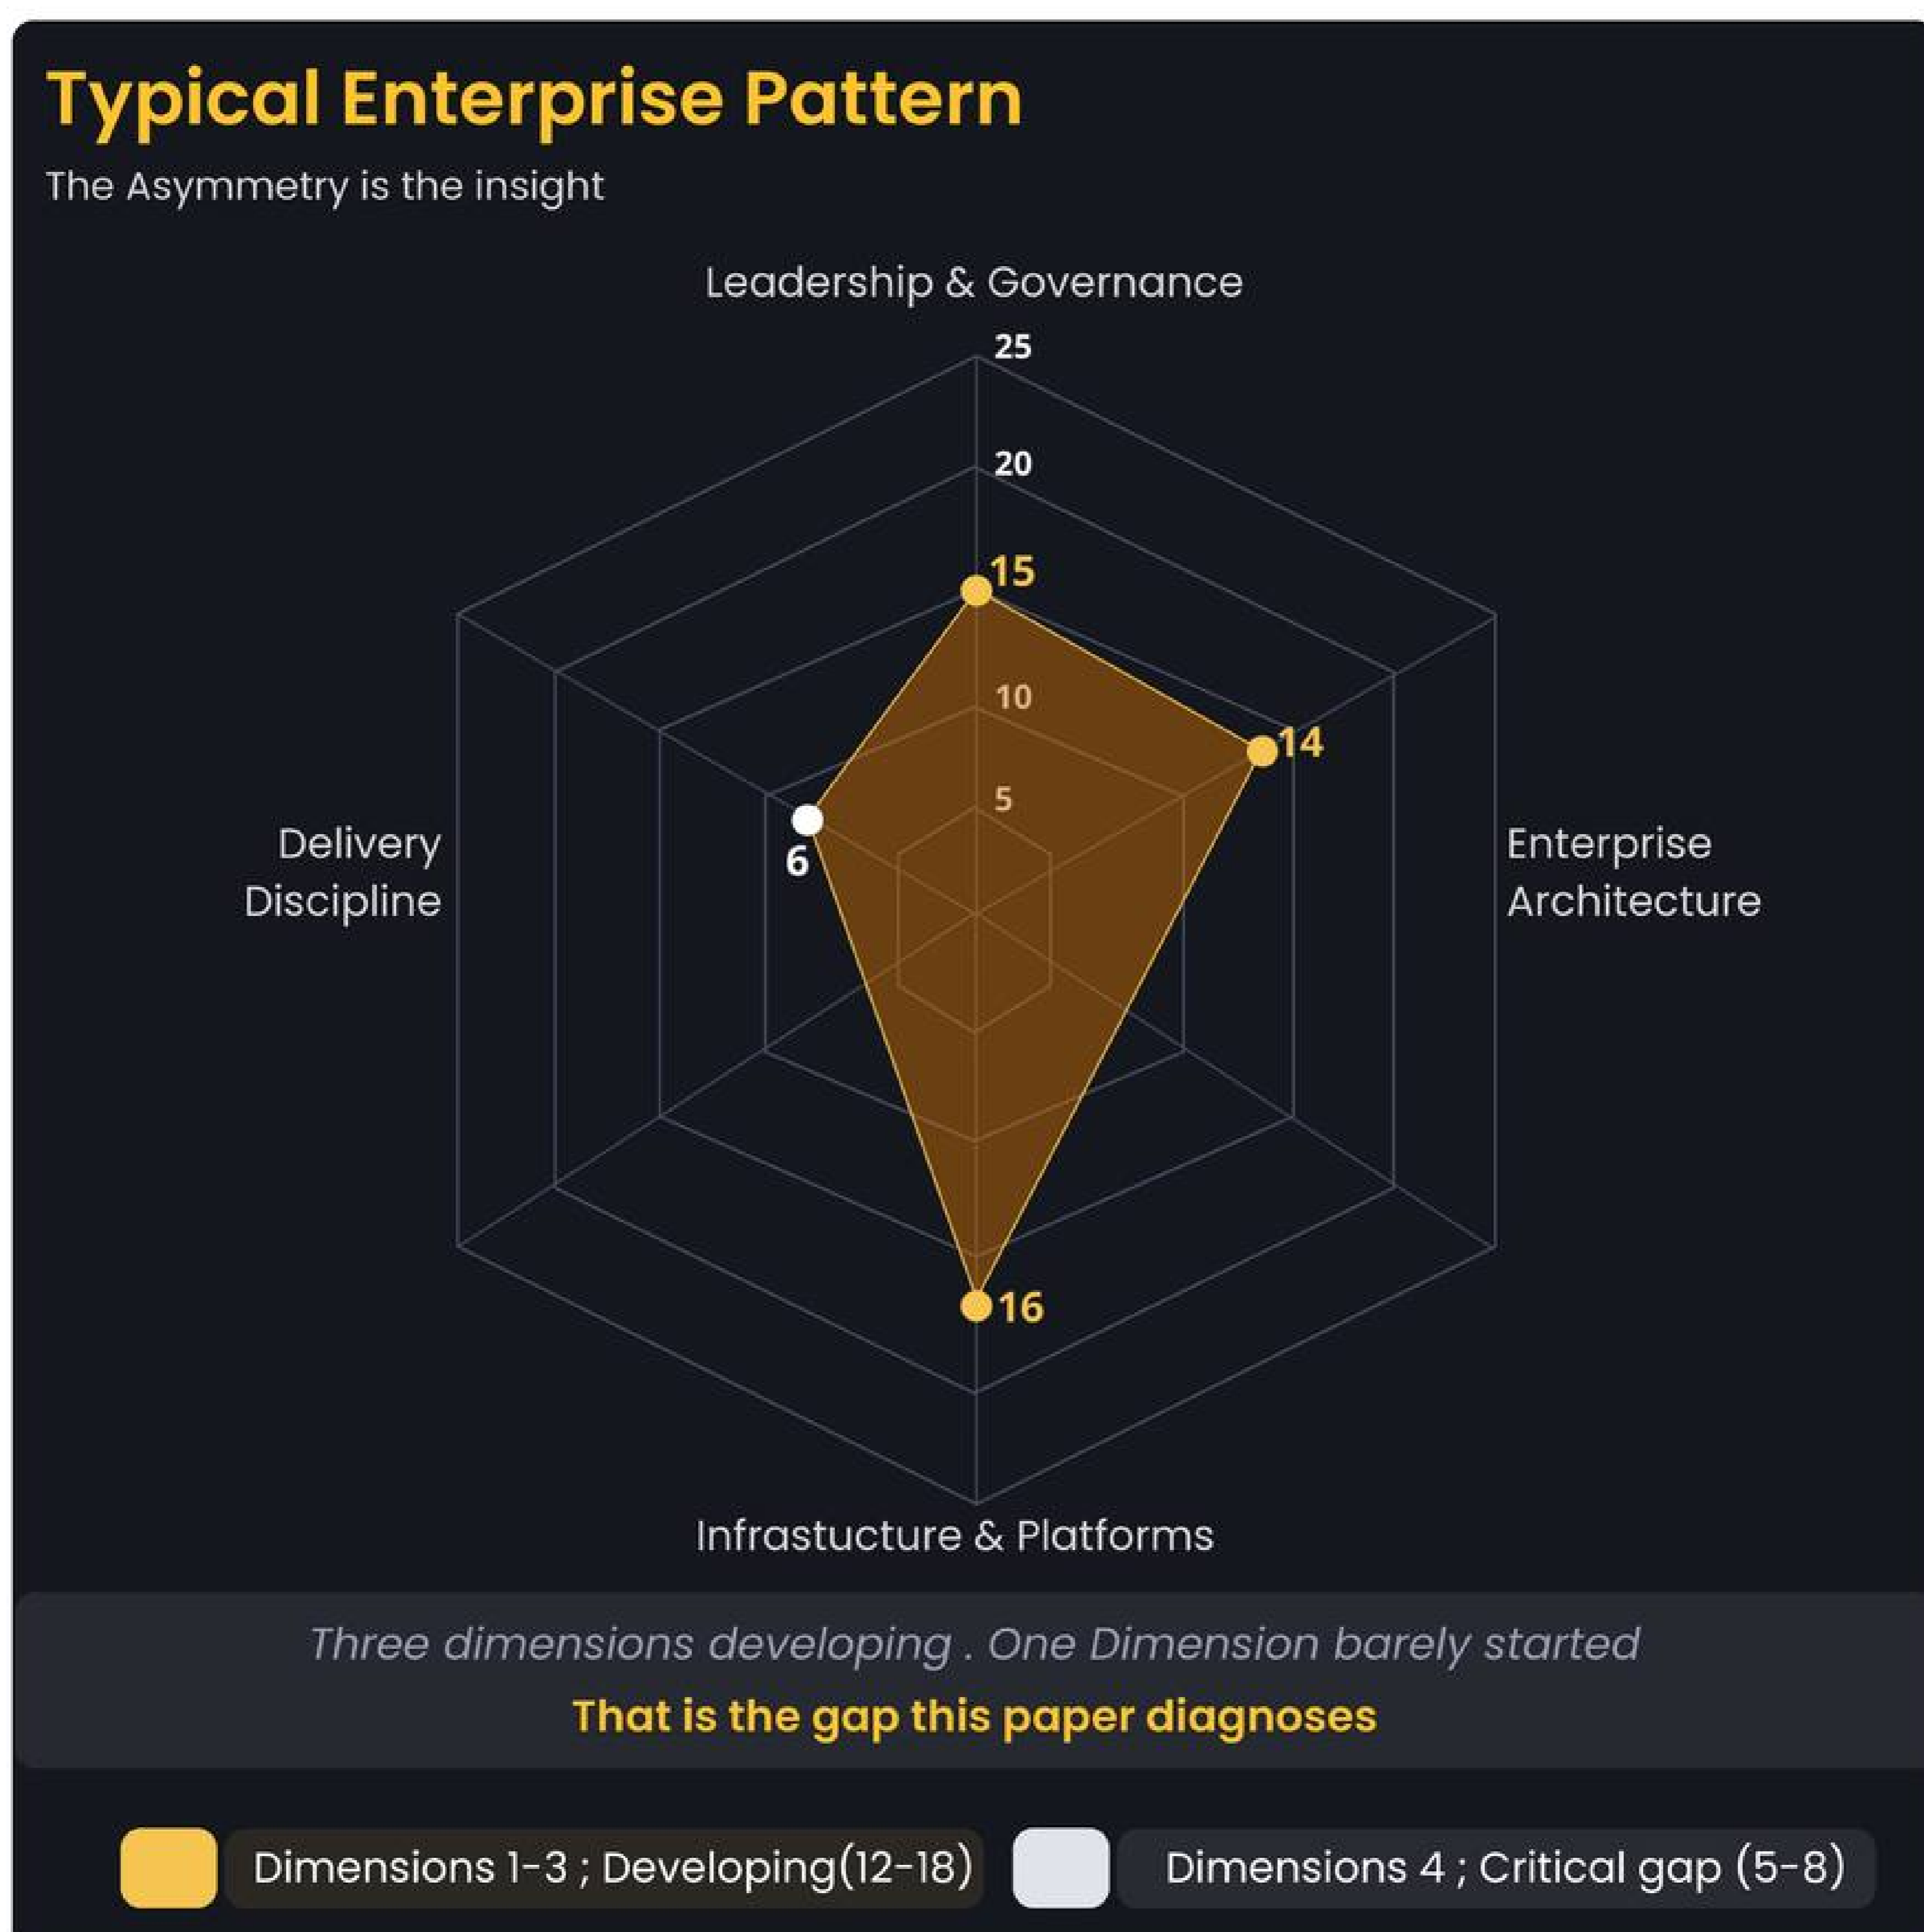
The Pattern You Will Likely See

When leadership teams complete this assessment, a consistent pattern emerges. Most organisations score between 12 and 18 on the first three dimensions. They have invested in governance conversations, started architecture modernisation, and built platform capabilities. Then they score between 5 and 8 on Delivery Discipline.

The dimension they have invested least in, the one most have never even considered as a distinct discipline, is the one that determines whether the other three produce results.

This is not a coincidence. The first three layers have established disciplines behind them. Organisations know how to hire governance consultants, enterprise architects, and infrastructure engineers. The fourth layer has no established discipline, no standard job title, no accepted methodology, and no training programme. It is the gap that explains why organisations with strong leadership intent, sound architecture, and modern infrastructure still cannot get AI programmes to production.

You are building capability you cannot govern. The readiness assessment makes that visible.





S E C T I O N 7

From Diagnosis to Action

This paper diagnosed the five predictable ways the metabolic loop breaks and mapped the four layers required to address them. The readiness assessment reveals where your organisation stands today.

Diagnosis is necessary but insufficient. Knowing that your organisation suffers from Decision Quality Collapse does not, by itself, fix the decision-making structure. Identifying an Execution Bottleneck does not dissolve the frozen middle. Scoring 6 out of 25 on Delivery Discipline does not create a discipline where none exists.

What the diagnosis does is change the conversation. Instead of “why is our AI programme not delivering?” the question becomes “which failure mode is constraining us and which layer do we address first?” That is a question a leadership team can act on.

The Pattern You Will Likely See

For the CIO




Use the readiness assessment to identify which layer is your binding constraint. If your infrastructure scores are strong but your delivery discipline is absent, that explains why sound technology choices are not producing enterprise outcomes. The constraint is not what you are building. It is how you are governing, coordinating, and sustaining what you build.

For the CFO



The 39% waste figure from ADAPT is not a technology problem. It is an operating model problem that the readiness assessment can help you diagnose. If your leadership and governance scores are low, your investment approval process is optimised for the wrong type of system. The returns will not improve until the governance model changes.

For the Chief AI Officer



If you scored 5 to 8 on Delivery Discipline, you now have language for the gap you have been experiencing since you took the role. The discipline that your role requires does not yet exist as a codified body of knowledge. That does not mean you cannot act. It means you are building the discipline as you go, and the readiness assessment gives you a framework for explaining to your board what must be built and why.

What We Do Not Yet Know

We should be honest about the limits of this paper. We have proposed a diagnostic framework, not a proven methodology. The delivery discipline for enterprise AI is being invented in real time, across industries and geographies, by practitioners who are learning as they go. What we offer is the clearest map we have.

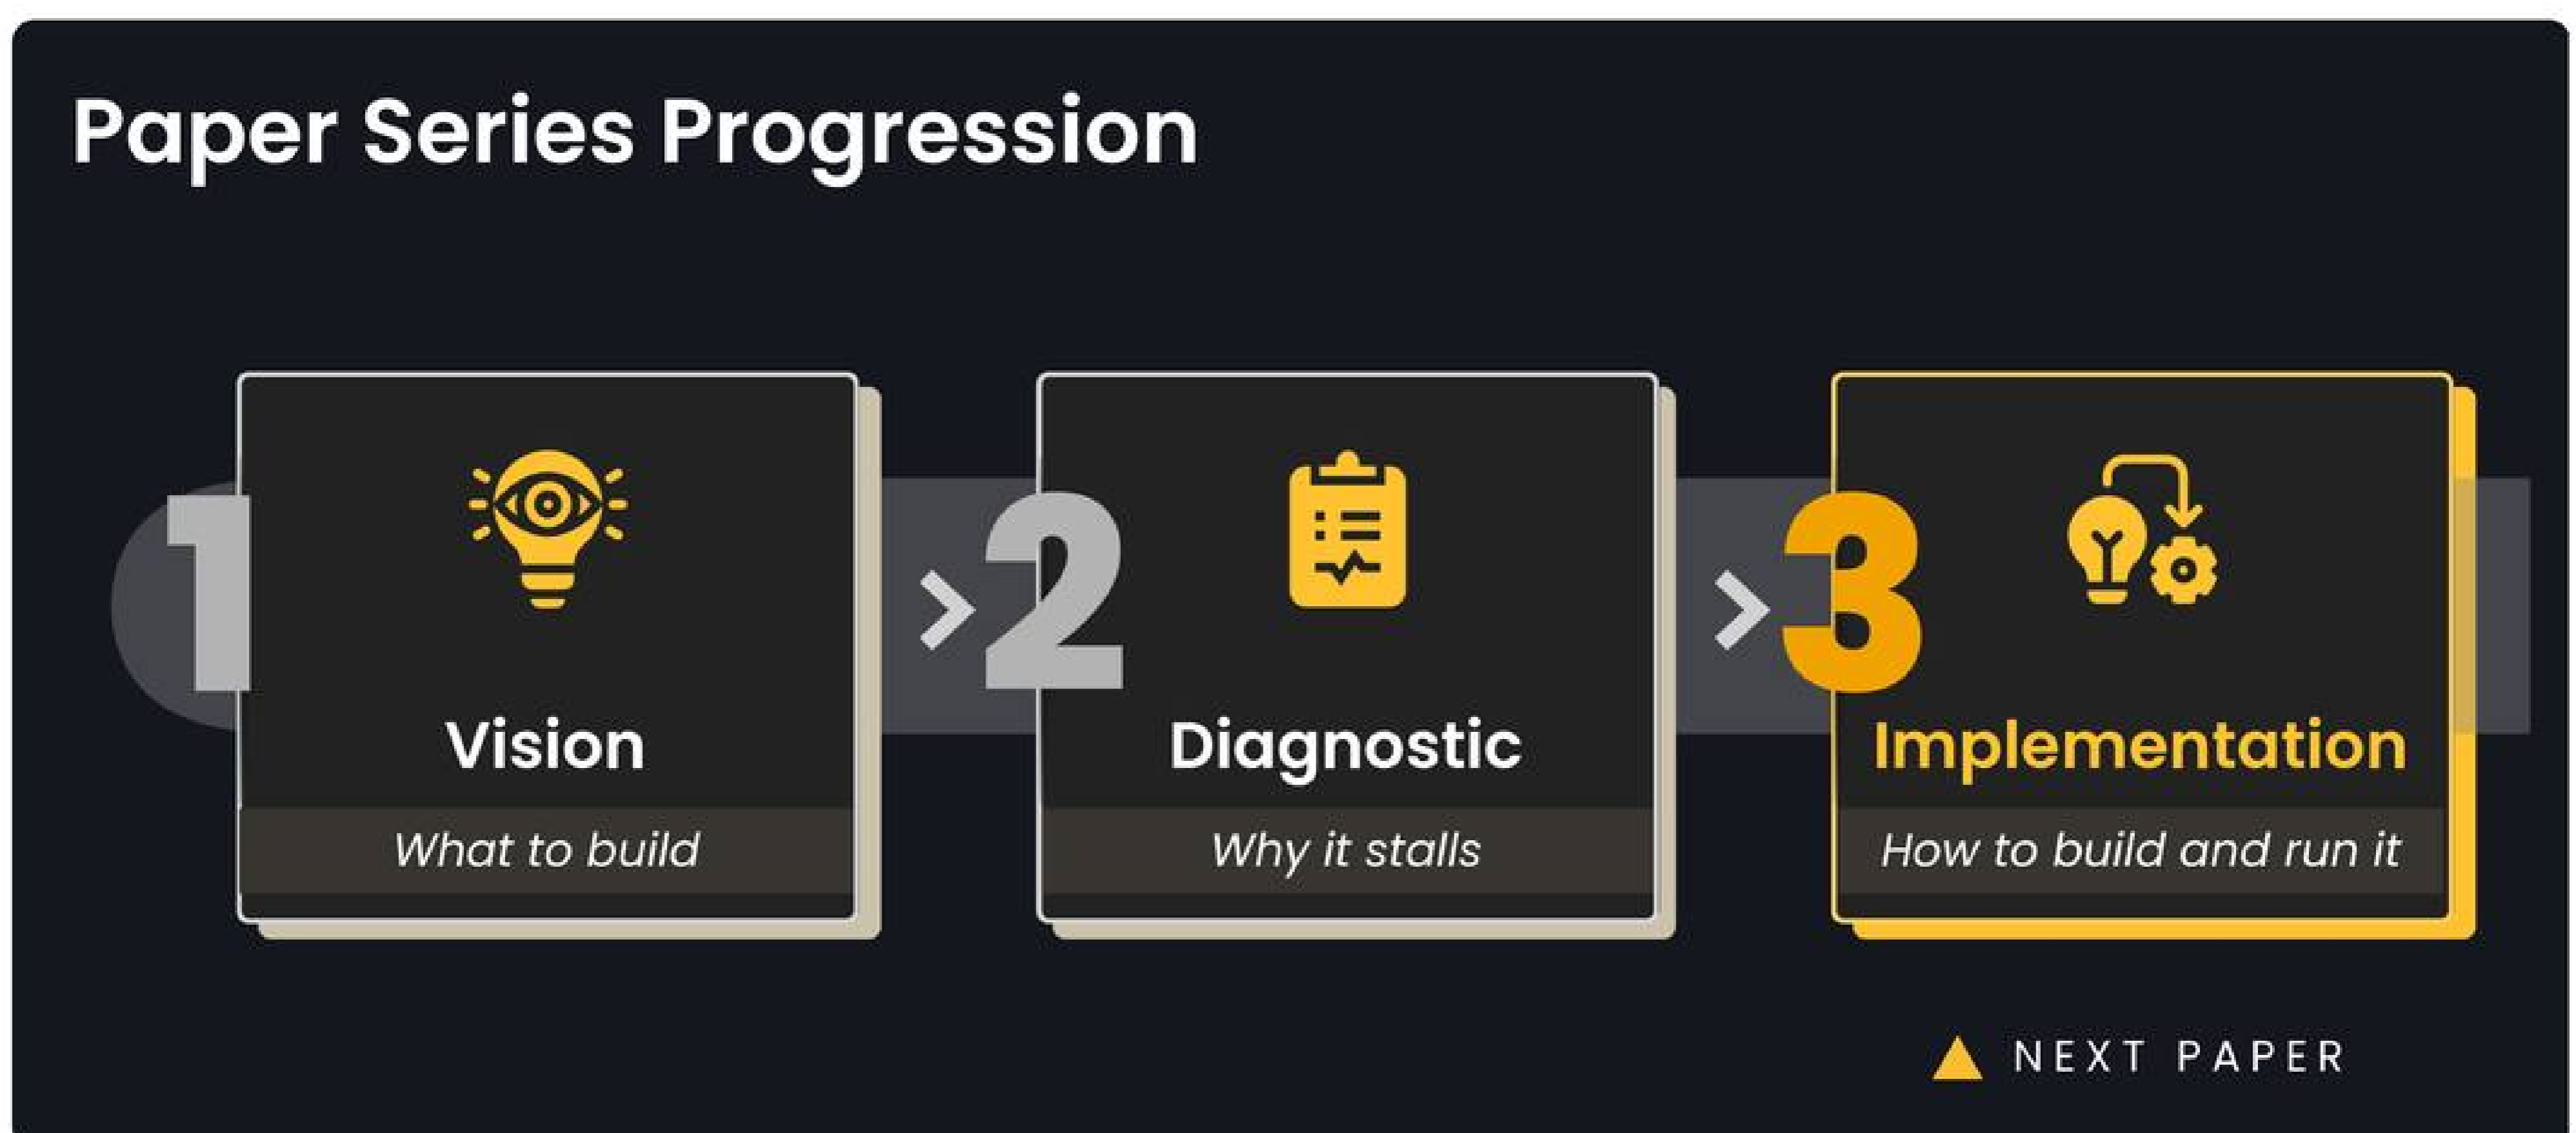
Paper 3 in this series will move from diagnosis to implementation. It will propose specific structures, roles, governance patterns, and delivery methodologies for operationalising ICE. It will draw on the evidence gathered from the organisations that have begun this work, including the failures that have taught us the most.

The Leadership Decision

After reading this paper, you can no longer attribute AI underperformance to technology failure. The failure modes are structural. The layers are identifiable. The gaps are diagnosable. What remains is a leadership decision: fund the disciplines that do not yet exist or continue investing in technology your operating model cannot absorb.

The metabolic loop will not fix itself.

The delivery discipline does not exist yet. We are building it. The question is whether your organisation will help shape it, or wait until someone else defines it for you.



References

- Cameron, M., Seenisamy, V., Bijlani, V., & Madiraju, S. (2025). The Intelligence-Centred Enterprise: Leading the AI-Centred Enterprise of the Future. ICE Whitepaper Series, Paper 1.
- ADAPT. (2025). Australian CIO and CFO Edge Research: Technology Capability Spend and Data Readiness.
- McKinsey & Company. (2024). The State of Digital Transformation: Why 70% Still Fail.
- ADAPT. (2025). CFO Edge: 39% Technology Capability Spend Waste.
- ADAPT. (2025). CIO Edge: Real-Time Data Access, Legacy Platform Dependencies, and AI Infrastructure Readiness.
- Gartner. (2025). Agentic AI Project Cancellation Forecast: 40% by 2027.
- MIT Sloan Management Review. (2025). State of AI in Business Report.
- S&P Global. (2024). Enterprise AI Adoption and Failure Rate Analysis.



From Operating Model to Operating Reality

Mark Cameron
Vijayan Seenisamy
Vinod Bijlani
Sharma Madiraju

| CEO, Alyve
| Enterprise Agentic AI Systems Delivery
| AI Practice Leader, HPE
| Enterprise AI Strategist

June 2026