Design for Coherence

Turning AI Potential into Enterprise Value

A short book for the executive team

Author: Erlend Hollebosch

November 2026

DESIGN FOR COHERENCE

Turning AI Potential into Enterprise Value

Erlend Hollebosch

First edition — November 2026

© 2026 Erlend Hollebosch · grow-faster.be

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, without the prior written permission of the author.

Cover artwork and interior figures were designed for this edition, as part of the accompanying executive handbook series.

CONTENTS

HOW TO USE THIS BOOK 4
INTRODUCTION 5
CHAPTER 1 — YOUR AI PILOT WORKS. WHY IS THE ORGANISATION NOT CHANGING? 14
CHAPTER 2 — WHAT COHERENCE ACTUALLY LOOKS LIKE 21
CHAPTER 3 — WHAT INCOHERENCE COSTS 25
CHAPTER 4 — THE COHERENCE DIAGNOSTIC 29
CHAPTER 5 — THE SYSTEM AI ENTERS 33
CHAPTER 6 — ENTERPRISE & AI TRANSFORMATION 37
CHAPTER 7 — FIVE STEPS, FIVE GATES 42
CHAPTER 8 — WHAT EACH ROLE CONTRIBUTES 46
CHAPTER 9 — THE VALUE CASE 51
CHAPTER 10 — TRAINING & CAPABILITY 55
CHAPTER 11 — HOW THIS FAILS: THE SEVEN SIGNALS AND THEIR COUNTER-MOVES 59
CHAPTER 12 — THE FIRST FIFTEEN WORKING DAYS 63
APPENDIX A — THE NINE ROLES, IN BRIEF 67
APPENDIX B — GLOSSARY 69
APPENDIX C — THE INSTRUMENTS AT A GLANCE 72

HOW TO USE THIS BOOK

This is a book to be used, not admired. It was written for the executive who has read enough about artificial intelligence and suspects — correctly — that the missing discipline is not technological. Its argument runs in one direction, from recognition to calendar: what incoherence is, what it costs, how to find it, and what to do in the first fifteen working days after you have decided to stop tolerating it. Read it in order the first time; the chapters are short by design and each builds the vocabulary the next one spends.

Three features earn their space. The figures are the working instruments of the method — the heatmap, the gates, the loop positions, the calendar — so that what you read here and what you show your board are the same pictures. The reflection that closes each chapter is written for the executive team rather than the individual reader; take the questions into the room, because they only work when asked aloud among the people who can act on the answers. And the appendices condense the book's roles, terms, and instruments into reference form, for the day someone asks you a hard question in a corridor.

If you have time for only one chapter, read the introduction and answer its three questions honestly. If you have time for two, add Chapter 4 and schedule the diagnostic. The rest of the book will be waiting when the heatmap tells you where to begin.

INTRODUCTION

THE OPENING SCENE

The pilot worked. That is precisely the problem.

Somewhere in your organisation, let's say in customer service, in the planning department or in finance, a team connected an AI application to a real process. The demo was convincing. Processing time went down, quality went up, the team was proud, and the board was pleased; someone used the word *potential* at least four times in the debrief. That was six months ago. Today the application is still a pilot. The productivity gain is real, yet invisible in the profit and loss account. The licence costs, on the other hand, are very visible. Between the demo and the P&L, the value leaked away — through meetings, integrations and approvals that never quite concluded.



Figure 1: The pilot succeeded — and six months later the organisation did not move.

RECOGNITION

Nine voices around one table

Ask around the executive table why the pilot never scaled, and you will hear nine different answers, each delivered with complete sincerity. The CEO leads nine sensible people toward an outcome that serves no one and wonders how to lead this at all. The CFO watches licences precede evidence, the gap between AI cost and verified value widening month by month. The COO knows the operation cannot pause for a transformation, and remembers that every transition so far dipped the performance figures he is held accountable for. The CIO sees shadow AI spreading and agents acting without owners, on data flows she can no longer audit. The HR director watches people fear replacement and doubt their own competence — a resistance no communication plan has reached. The Sales and Marketing director sells promises that a fragmented organisation struggles to keep. Procurement watches every vendor sell intelligence, while lock-in is signed before value is proven. Supply Chain removes buffers from a

forecast nobody trusts, adding fragility in the name of efficiency. And the business unit managers see the transformation take their best people for months, while the benefits are claimed centrally.

Nine explanations, nine valid concerns — and one shared root cause that none of them can see from where they sit.



Figure 2: Nine voices around one table — every concern locally true, the root cause shared.

THE ROOT CAUSE

It is not an AI problem. It is incoherence.

That root cause has a name, and it is not an AI problem, a data problem, a culture problem or a leadership problem. It is incoherence: the state in which the parts of an enterprise each function reasonably well, but the connections between them no longer form a working whole. You can recognise it by three signals. Decisions keep waiting — several people are involved in every significant choice, but no one owns the trade-off. Work repeatedly changes hands — each department completes its part correctly, yet the shared outcome stays out of reach. And priorities compete — initiatives ask for the same people, the same systems and the same information at the same time, so everything is priority one and nothing is. None of this was created by AI. Complexity has been

accumulating for decades, in layers of governance, in systems that do not connect, in interfaces nobody designed. AI did not create this condition. AI exposes it and then amplifies it.

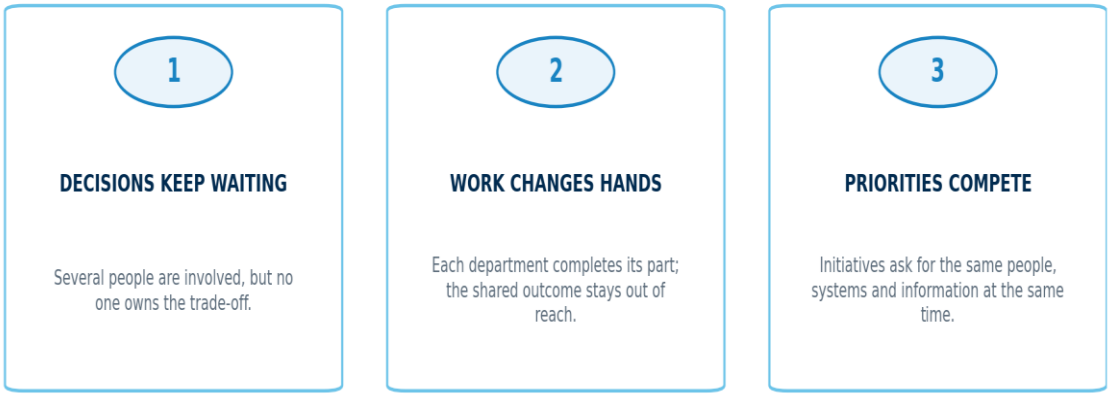


Figure 3: The three signals of incoherence — and the condition that predates AI.

THE CENTRAL FACT

AI amplifies the system it enters

The amplification works in both directions, and that is the central fact this book is built on. AI does not arrive with a way of working of its own; it accelerates the way of working it finds. Deploy it into a coherent enterprise — explicit decision rights, connected systems, one priority list — and it compounds value: faster decisions, shorter cycles, freed capacity. Deploy the identical technology into a fragmented enterprise and it compounds the fragmentation: faster confusion, more interfaces, more output that nobody owns, costlier disappointment. The technology is an amplifier. The enterprise is the signal. Which is why the sequence of a successful transformation is not a matter of preference but of arithmetic: design for coherence first, then amplify.

Coherent enterprise	**+ AI**	Compounding value — faster decisions, shorter cycles, freed capacity
Fragmented enterprise	**+ AI**	Faster confusion, more interfaces, costlier disappointment

The technology is an amplifier. The enterprise is the signal.
This book is about designing the signal.

Figure 4: Same AI, different system: how value develops over thirty-six months in a coherent versus a fragmented enterprise.

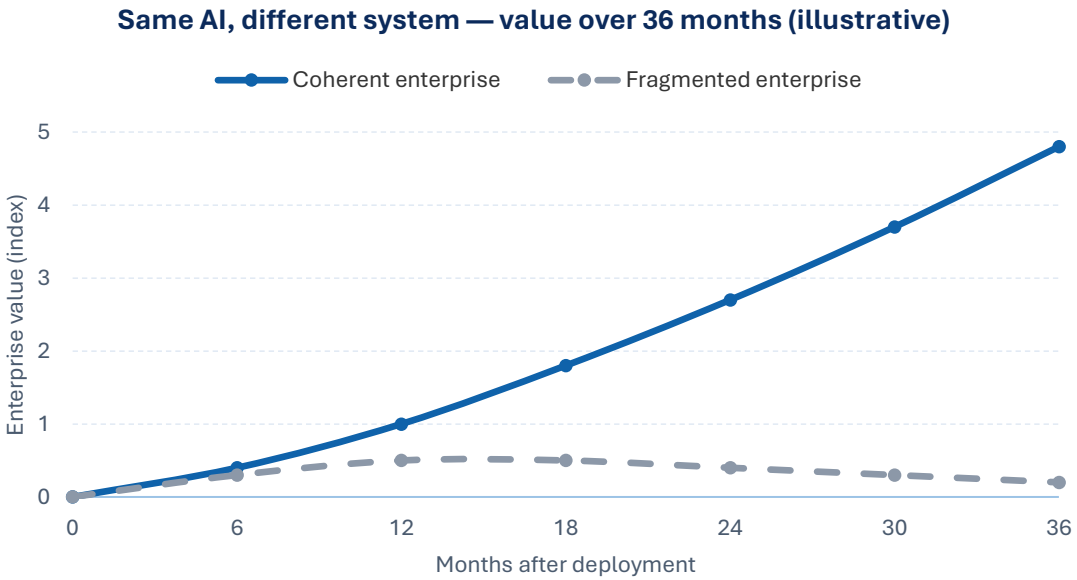


Figure 5: AI amplifies the system it enters — same AI, different system, different curve.

THE FOUNDATIONS

Four foundations, kept deliberately compact

This book gives you the means to work in that sequence. It stands on four foundations, kept deliberately compact. The first is enterprise coherence itself — the operating condition in which strategy, people, process, structure, infrastructure and data reinforce one another rather than undermine one another. The second is systems thinking — the

discipline of seeing the enterprise as one system, in which a local choice always has system-wide consequences. The third is enterprise and AI transformation — the understanding that agentic AI does not merely answer but plans, decides and acts, and therefore needs ownership, boundaries and control designed in advance. The fourth is the 7S lens — McKinsey's seven interdependent elements, strategy, structure, systems, style, staff and skills held together by shared values, which turns the word incoherence from an accusation into a map the executive team can actually discuss.

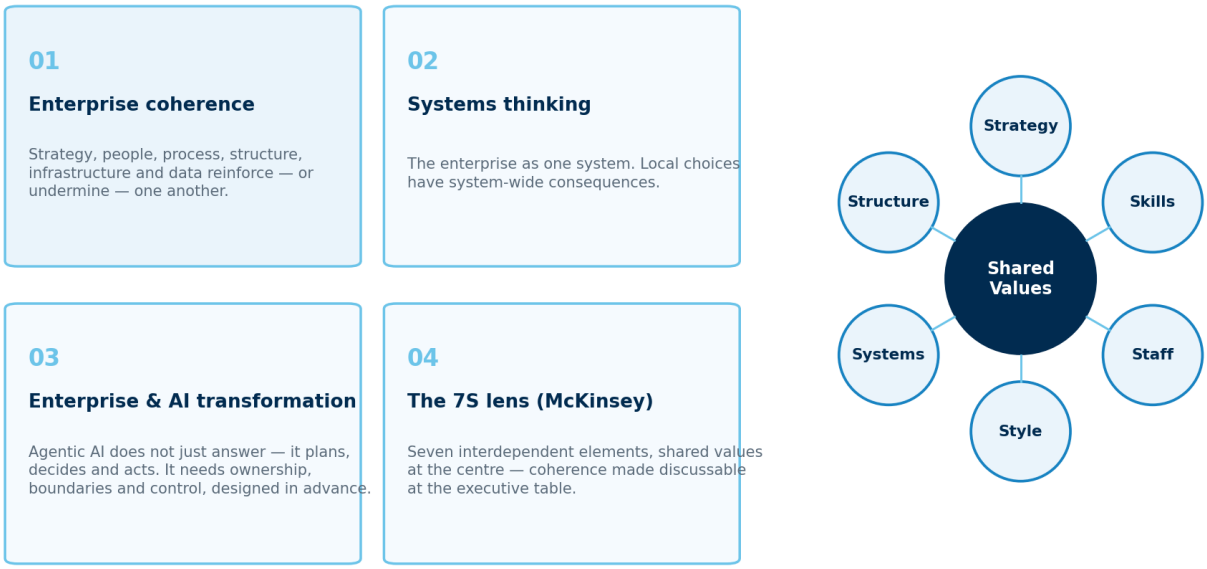


Figure 6: Four foundations — with the 7S lens as the shared language of the executive table.

THE METHOD

Five steps, five gates — one rhythm

On these foundations rests a five-step approach, and it is deliberately a rhythm rather than a programme. It begins with mandate and diagnosis — which business outcome must change, and what blocks it today, established on evidence in fifteen working days. It continues with readiness — an honest assessment across business, data, technology, governance and people. It then sets direction — a North Star AI vision and a portfolio sequenced on impact and feasibility. Only then does it activate workflows, redesigned around human–agent collaboration with the performance dip managed rather than denied. And it ends in anchoring and scaling — governance, gated funding and a curriculum that keeps the capability in-house. Each step is timeboxed, and each closes

with a gate that has a named decision owner: the CEO opens, the CIO co-signs, the CEO sets the course, the COO owns the workflow, the CFO funds. Not a programme that runs for years — a rhythm you can start within fifteen working days.

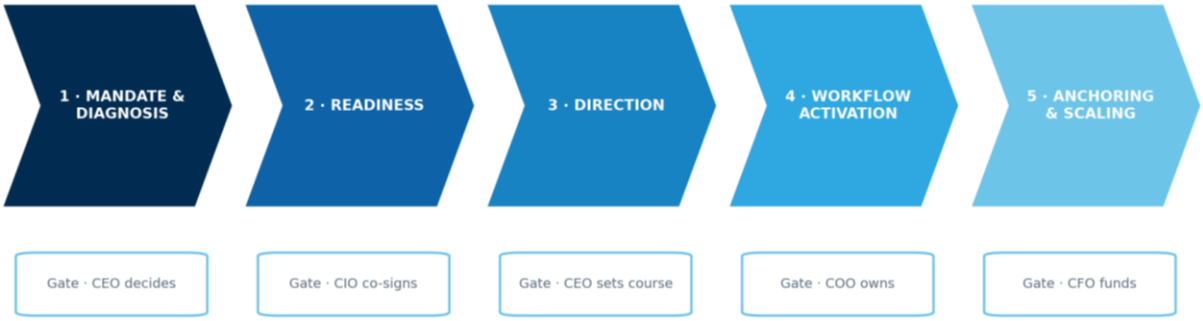


Figure 7: Five steps, five gates — the rhythm of the coherent transformation.

THE VALUE CASE

What the P&L can expect — and when

A book for executives owes its readers an answer about money, and this one does not evade it. In the first six months the value is avoidance and speed: misinvestments stopped before signature, less rework, faster executive decisions. Between six and eighteen months the redesigned workflows deliver verified results — cycle times thirty to forty percent shorter in the chosen workflows, lower licence and coordination costs. From eighteen months onward coherence compounds: a shared operational backbone lowers structural operating cost, and each next transformation costs less than the last. One rule governs all of it: enterprise funding flows only to initiatives demonstrating at least one and a half times verified annual value — verified by Finance, not claimed by the project.

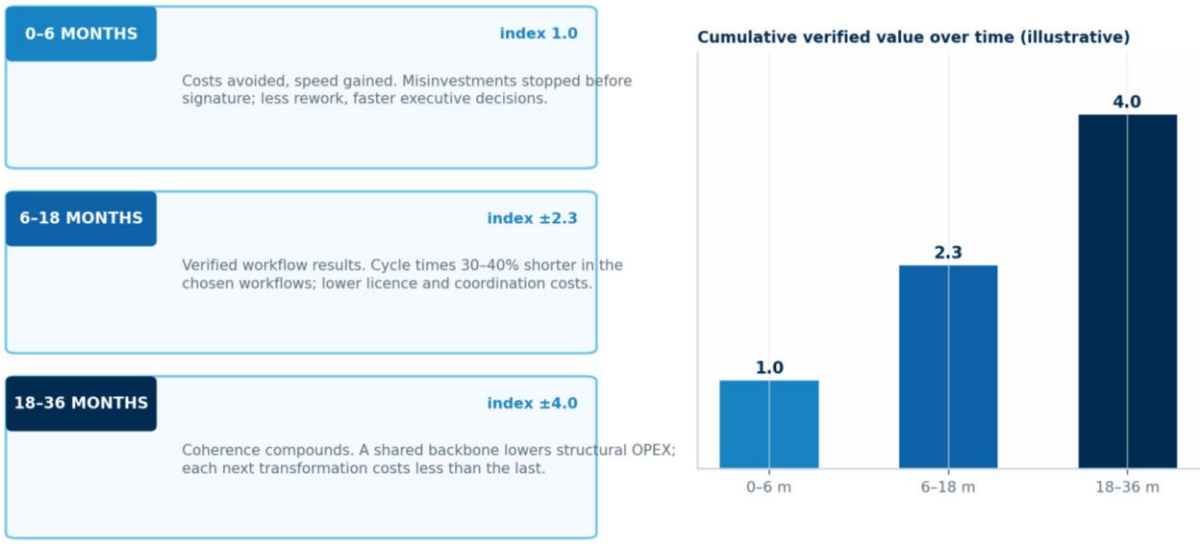


Figure 8: What the P&L can expect, and when — cumulative verified value across the three horizons (illustrative indices).

READING GUIDE

How to read this book in 45 minutes

You do not have to read this book in order. It is built for the way executive teams actually read: each role starts with the one-page summary at the back and follows its own track through the chapters, ending at the decision that role owns. The CEO reads the paradox, the five steps and the value case. The CFO reads what incoherence costs and adopts the funding rules. The COO and CIO read the system chapters and take workflow ownership. The extended team reads its own contributions and the failure modes. And the whole executive team reads the first part together — and then decides on the first fifteen working days. A shared yes, or a shared no.

ROLE	YOUR 45-MINUTE TRACK	THE DECISION YOU TAKE
CEO	Role summary → Ch. 1 (the paradox) → Ch. 7 (five steps) → Ch. 9 (value case)	Mandate and gate decisions
CFO	Role summary → Ch. 3 (what incoherence costs) → Ch. 9 (value case)	Funding gates — ≥ 1.5× verified value
COO / CIO	Role summary → Ch. 6 (the system AI enters) → Ch. 8 (your contribution)	Workflow ownership · readiness and control
HR, Sales, Procurement, Supply Chain, BU leaders	Role summary → Ch. 8 (what each role contributes) → Ch. 11 (how this fails)	Adoption, benefits delivery, supplier gates
Whole executive team	Part I together — then decide on the first fifteen working days	A shared yes, or a shared no

Figure 9: How to read this book in 45 minutes

CLOSING THE INTRODUCTION

One promise, one caution

One promise, then, and one caution. The promise: everything in this book is designed to be executable. Every claim about value comes with a way to measure it, and a gate that must be passed before further money is committed; nothing asks for belief where verification is possible. The caution: this book will not tell you that AI is easy, or that your organisation is readier than it is. The winners are not the enterprises with the most pilots — they are the ones that did the less glamorous work first. The problem is not out there in the technology, and not down there in the teams. It is here, between us — in how this enterprise holds together, or does not.

Design for coherence. Then amplify.

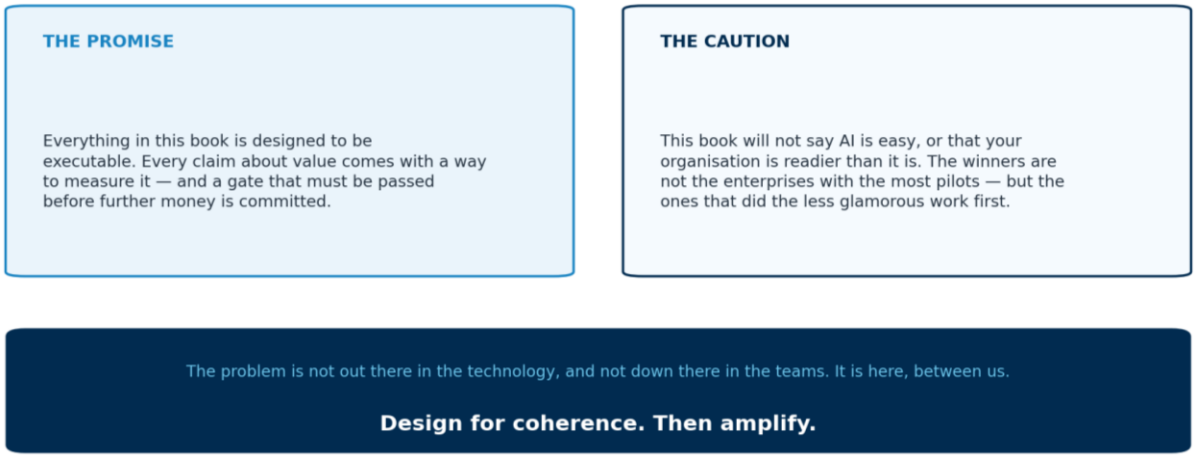


Figure 10: One promise, one caution.

REFLECTION FOR THE EXECUTIVE TEAM

Before you read on, take the three signals into your next management meeting — not as a slide, but as a question asked in your own voice. Which decision in this organisation has been waiting longest, and why does everyone in this room already know which one it is? Where did a piece of work change hands twice this week that should have moved once? And which of our stated priorities would survive if we asked the floor, anonymously, to name our top three? The value of the introduction is not in agreeing with its argument but in testing it against the organisation only you can see. If the signals come back quiet, close the book and keep it as a reference — you are rarer than you know. If they come back loud, then the uncomfortable recognition you feel is not a problem. It is the beginning of the diagnosis, and it cost you nothing but honesty.

CHAPTER 1 — YOUR AI PILOT WORKS. WHY IS THE ORGANISATION NOT CHANGING?

THE PILOT THAT NOBODY CAN EXPLAIN

The demo was convincing. Somewhere in your organisation — in customer service, in the planning department, in finance — a team connected an AI application to a real process. The numbers were strong. Processing time down, quality up; the team was proud, the board was pleased, and someone used the word *potential* at least four times in the debrief. That was six months ago. The application is still a pilot. The productivity gain is real, but it is invisible in the P&L, while the licence costs are very visible indeed. And when you ask why the pilot has not scaled, you get an answer that depends entirely on whom you ask.

Ask the CFO, and the problem is evidence: licences were signed before value was proven, and the gap between AI cost and verified return widens every month. Ask the COO, and the problem is continuity: the operation cannot pause for a transformation, and every transition so far has dipped performance on the numbers he is held accountable for. Ask the CIO, and the problem is control: shadow AI is spreading, agents act without owners, and the data flows she is supposed to govern have become impossible to audit. Ask the HR director, and the problem is fear: people expect replacement, doubt their own competence, and resist in ways no communication plan has reached. Ask the Sales and Marketing director, and the problem is promises: he sells commitments that a fragmented organisation struggles to keep, and AI has only raised customer expectations further. Ask Procurement, and every vendor now sells intelligence, while lock-in is signed before value is proven. Ask Supply Chain, and every buffer removed in the name of efficiency has added fragility to a forecast nobody trusts. Ask your business unit managers, and the transformation takes their best people for months and then claims the benefits centrally.



Figure 11: Month zero, month three, month six — the trajectory every executive team recognises.

Nine explanations. All of them sincere, all of them supported by something real — and none of them sufficient. Because if any single explanation were the whole story, the organisation would have fixed it by now. This is the moment to consider a different possibility: the pilot is not failing. The organisation around the pilot is not built to absorb it.

EVERY EXPLANATION IS LOCALLY TRUE — AND THAT IS THE TRAP

Notice something about the nine voices. Each one describes the situation accurately from where that person sits. The CFO really does see costs precede evidence. The COO really does see performance dips. The HR director really does see fear. None of them is wrong — and that is exactly what makes the problem so persistent. When every diagnosis is partially correct, every intervention is partially effective. Finance tightens the funding rules, and the pilots slow down. IT tightens the governance, and shadow AI moves further underground. HR launches an engagement programme, and the fear does not move, because the fear was never the root. Each fix treats a symptom in one corner of the enterprise, and the condition in the middle stays untouched.

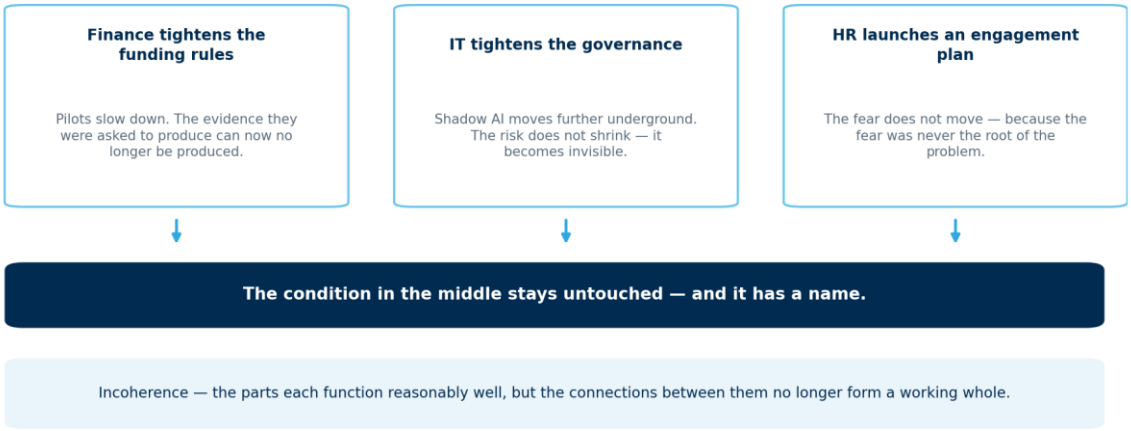


Figure 12: Local fixes, applied sincerely — and the condition in the middle untouched.

There is a name for that condition, and it is not an AI problem, a data problem, a culture problem or a leadership problem. It is incoherence: the state in which the parts of an enterprise each function reasonably well, but the connections between them — decisions, handovers, priorities, information — no longer form a working whole. You will recognise its three signals immediately. Decisions keep waiting: several people are involved in every significant choice, but no one owns the trade-off, and questions travel upward, sideways and back down, returning as new questions. Work repeatedly changes hands: each department completes its part correctly, yet the shared outcome stays out of reach, and value leaks away in interfaces nobody designed. And priorities compete for the same capacity: initiatives ask for the same people, the same systems and the same information at the same time, so everything is priority one and therefore nothing is. If two or three of these sound familiar, the pilot was never your bottleneck. The system was.

WHAT INCOHERENCE COSTS — BEFORE AI EVER ENTERS THE PICTURE

Incoherence predates artificial intelligence by decades. It accumulated quietly: in layers of governance added after each incident, in systems bought per department and connected afterwards, in reorganisations that moved boxes on the chart without moving decision rights. None of this was irrational at the time. Each layer solved a local problem. Together, they built an enterprise that is too complex to change quickly.

The cost is real, and it falls on the P&L in three places. The first is decision latency — the time between *we should* and *we did*. In most established organisations, a cross-functional decision takes weeks of alignment that nobody experiences as waste, because it is distributed across calendars. The second is interface friction — the rework, reconciliation, double entry and alignment meetings that exist only because the parts do not fit; it rarely appears as a line item, but it consumes a structural share of capacity. The third is initiative congestion — too many programmes competing for the same scarce people, each slowing all the others, the organisation paying full price for every initiative and receiving a fraction of each one's value. Here is the uncomfortable arithmetic: these costs were already there before your first AI pilot. The pilot did not create them. It simply made them visible, because a technology that promises speed runs into an organisation designed for waiting.

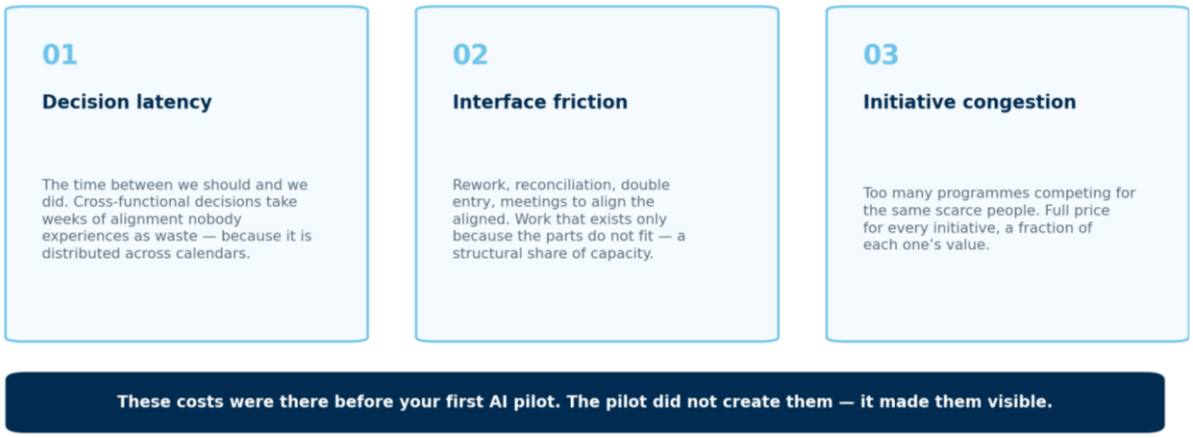


Figure 13: The three costs of incoherence — present long before the first AI pilot.

WHY AI MAKES THIS URGENT NOW

For decades, incoherence was expensive but survivable. Competitors carried the same burden, and markets moved slowly enough to absorb the friction. AI changes that equation in two ways. First, AI is an amplifier: it does not arrive with a way of working of its own, it accelerates the way of working it finds. In a coherent enterprise it compounds value; in a fragmented one it compounds the fragmentation — faster confusion, more interfaces, more output nobody owns, costlier disappointment. The technology is an amplifier; the enterprise is the signal.

Second, agentic AI raises the stakes from advice to action. A chatbot that answers questions can tolerate ambiguity; the cost of a wrong answer is a sigh and a retry. An agent that plans, decides and acts — placing orders, adjusting schedules, moving money, contacting customers — cannot tolerate ambiguity at all. It needs to know who owns the outcome, where its authority ends, what it may never do, and who is accountable when it acts on the boundary. In an incoherent enterprise, those questions have no answers, because they have no answers for humans either. This is why so many agentic initiatives stall between demo and production: the technology is ready, and the organisation has not yet decided what it is allowed to decide. This is the central fact of this book: AI did not create your incoherence. It exposes it — and then amplifies it. Which fixes the sequence. Design for coherence first. Then amplify.



Figure 14: An amplifier and an actor — why AI turns an old weakness into a new risk.

A FIRST, HONEST TEST

Before you read on, take ten minutes — alone, or better, with your executive team — and hold three questions against your own organisation. Not as a survey, but as a conversation. First, the waiting decision: name one decision that has been open for more than a month, and ask who owns the trade-off it requires; if the honest answer is *several people* or *it depends*, you have found decision latency. Second, the travelling work: pick one outcome that matters — an order delivered, a claim settled, a product launched — and ask how many hands it passes through and who owns the whole rather than a part; if the answer is *everyone and therefore no one*, you have found interface friction. Third, the crowded calendar: list your active initiatives and the five people each cannot succeed

without, and count how often the same names appear; if your best people are allocated at one hundred and eighty percent, you have found initiative congestion. Write the answers down. They return in Chapter 4, where the Coherence Diagnostic turns this first impression into an evidence-based picture, and in Chapter 7, where the same three signals determine where your transformation starts.

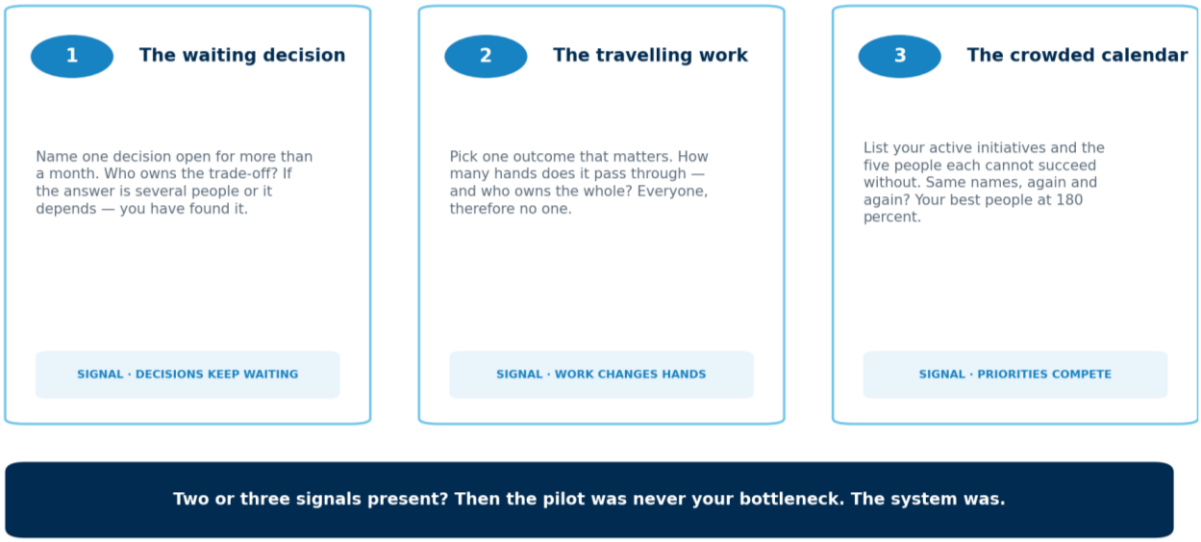


Figure 15: The ten-minute test — three questions, three signals, one verdict.

WHAT THIS MEANS FOR YOU, SPECIFICALLY

This book assigns no spectator seats. If you are the CEO, this chapter is your mandate: the recognition that the bottleneck is not technological, and that no one below you can fix a condition that lives between the parts you lead together. If you are the CFO, the three costs above are your opening position: decision latency, interface friction and initiative congestion are measurable, and Chapters 3 and 9 show how to price them and how to release funding in gates, only against verified value. If you are the COO, your concern about the performance dip is legitimate and addressed head-on: this book redesigns one workflow at a time, manages the dip openly, and never asks the operation to pause. If you are the CIO, the governance you need cannot be built around AI after the fact — it must be designed into the operating model first, which is exactly what Chapters 5 and 6 cover. And if you lead HR, Sales, Procurement, Supply Chain or a business unit: the pilot's failure was not your failure, and the fix is not something done to your

department. It is something your department does, with a named part and a visible return.

THE FIRST GATE

Every chapter of this book ends the way a coherent transformation proceeds: with a decision, not a summary. The decision at the end of this chapter belongs to the CEO, and it is deliberately small. Not *shall we transform* — that question is too large to answer honestly at page twenty. The question is this: are the three signals — waiting decisions, travelling work, competing priorities — present in this enterprise, yes or no? If the answer is no, close this book with our compliments; you are the exception, and Chapter 2 will test that confidence. If the answer is yes — and in our experience, in established organisations, it is yes — then the pilot's lesson is clear. You do not have an AI problem to solve. You have a coherence to build. The next chapter shows what that coherence actually looks like: not as a slogan, but as an operating condition you can see, measure and design.

REFLECTION FOR THE EXECUTIVE TEAM

The trap this chapter describes is expensive precisely because it feels like progress. Ask yourselves where, in the last budget cycle, amplification was approved on the strength of a demonstration rather than a diagnosis — and ask what the demonstration was actually demonstrating: the capability of the tool, or the coherence of the workflow it would enter. Then ask the harder question the chapter ends with. If an agent were dropped into your three most critical processes tomorrow, unannounced, which of them would it make faster, and which would it make faster at being wrong? The answer is already inside your organisation; the people who carry those processes know it. What the chapter asks of you is only this: that you hear it from them before you hear it from the incident report. The pilot you are proudest of deserves the most honest version of this question, because pride is the filter through which the trap is hardest to see.

CHAPTER 2 — WHAT COHERENCE ACTUALLY LOOKS LIKE

NOT A SLOGAN — AN OPERATING CONDITION

Ask ten executives to define a coherent enterprise and you will get ten admirable abstractions: aligned, agile, customer-focused, one team. None of them survives contact with a Tuesday morning. You cannot inspect "aligned." You cannot put "one team" on an agenda and decide anything. Coherence, as this book uses the word, is not a value to aspire to. It is an operating condition — a set of properties an enterprise either has or does not have, visible in how work actually moves. You can see it, measure it, and design for it. That is what makes it useful at an executive table: coherence is not something you believe in, it is something you check.

Five properties define the condition, and none of them is exotic — it is their combination that is rare. The first is that decision rights are explicit: for every significant decision, one named role owns the trade-off. Not a committee, not a process — a role. Others contribute input; one decides. You can test this in five minutes: pick any open decision and ask the room who owns it. In a coherent enterprise the answer takes one second, and everyone gives the same name. The second is that work is owned end-to-end: every outcome that matters — an order delivered, a claim settled, a product launched — has one owner for the whole, however many departments contribute parts. Hand-offs still exist, but they are designed, counted and managed, not discovered after the fact. The third is that priorities live on one list: initiatives compete on a single, ranked portfolio owned by the executive team as a team — not three lists maintained in three offices and reconciled in corridor negotiations. When two initiatives want the same scarce people, the list decides, not the loudest sponsor. The fourth is that information is shared by default: the data a decision needs flows to it without bilateral negotiation. Departments do not own information; they maintain it. This is as much an architectural property as a cultural one, which is why Chapter 5 treats it as a backbone question. The fifth is that boundaries are designed, not assumed: everyone — including every AI agent — knows what falls inside their authority, what requires escalation, and what is never allowed. Boundaries are written down where the work is done, not implied by job titles.

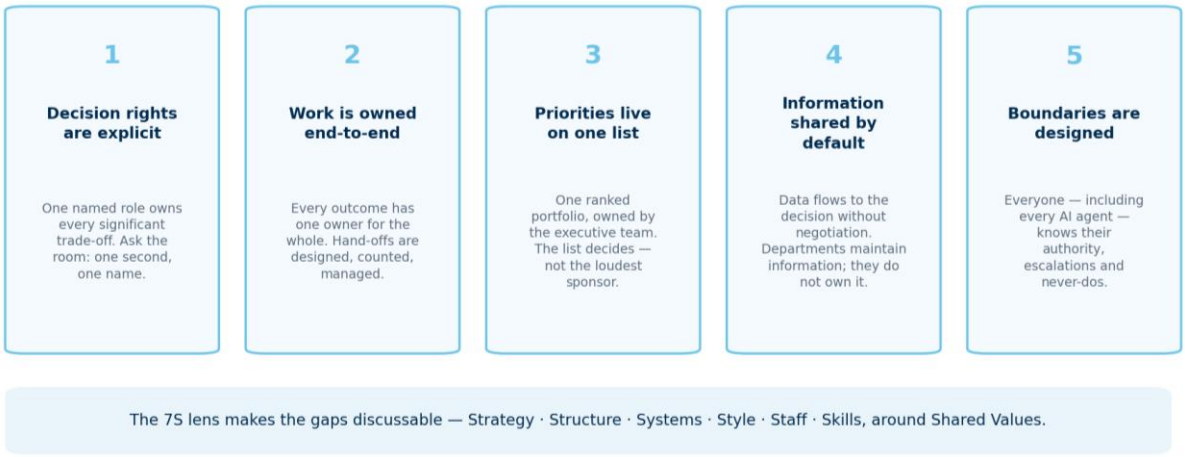


Figure 16: The five observable properties of a coherent enterprise.

Walk through your own organisation with these five. Where a property holds, work flows. Where it fails, you will find one of the three signals from Chapter 1 — waiting decisions, travelling work, competing priorities. The mapping is not subtle. It is mechanical.

A COMPOSITE EXAMPLE

Consider a mid-sized industrial company — call it a composite of several we have worked with — rolling out AI-assisted demand forecasting. The forecast itself improved within weeks. But procurement kept ordering on its own safety logic, production kept planning on last year's pattern, and sales kept promising delivery dates from a spreadsheet nobody else could see. Every department was locally correct. The enterprise had no single owner for the forecast-to-delivery chain, no shared information layer, and no boundary stating which decisions the model could take and which it could only recommend. The AI did not fail. The condition it entered was not an operating condition at all. Six months later the same company redesigned one workflow end-to-end — forecast to purchase order — with one owner, shared data and explicit agent boundaries. The model was unchanged. The outcome was not: forecast adoption rose, safety stock fell, and the AI investment finally appeared in the numbers. Same technology. Different signal.

COHERENCE IS NOT ALIGNMENT, AND NOT STANDARDISATION

Two confusions must be cleared away early, because they send well-intentioned programmes in the wrong direction. Alignment is agreement at a point in time — the offsite where everyone nods. Coherence is what remains three weeks later, when the nods meet reality. Alignment is an event; coherence is a property. You can be perfectly aligned and perfectly incoherent: everyone agrees, and the work still does not flow. Standardisation, likewise, is one way to achieve coherence, not coherence itself. Some enterprises buy coherence with uniformity — one process, one system, everywhere — and it works until the business needs variety, which it always does. The mature position, developed in Chapter 5, is more precise: standardise the backbone, differentiate the edge. Coherence is about knowing the difference and designing both deliberately.

THE 7S LENS — COHERENCE MADE DISCUSSABLE

There is one more instrument this chapter adds to your table, because executives need a shared vocabulary before they can redesign anything. The McKinsey 7S model, older than most of the software in your building, remains the most practical lens: strategy, structure, systems, style, staff and skills — six elements held together by shared values at the centre. Its value here is not academic. The 7S lens turns "we are incoherent" from an accusation into a map. When a transformation stalls, the question becomes specific. Is the blockage in systems — the tools and processes do not connect? In structure — nobody owns the outcome end-to-end? In skills — the organisation cannot yet do what the strategy requires? In style — leadership behaviour contradicts the announced priorities? Or in the centre — do the shared values people actually observe match the ones on the wall?

Agentic AI stresses all seven at once. A strategy that delegates decisions to agents changes structure — who supervises? It changes systems — what may the agent touch? It changes staff and skills — who designs the boundaries? It changes style — do leaders trust verified agent output? And ultimately it touches shared values — what do we optimise for when speed is free? The enterprises that treat AI as a systems purchase discover, expensively, that the other six S's were part of the price. Throughout this book the 7S lens does quiet duty: every diagnostic question in Chapter 4, every role contribution in Chapter 8, every gate in Chapter 7 can be located on it. It is the chessboard on which coherence is played.

WHAT THIS CHAPTER CHANGES FOR YOUR ROLE

For the CEO, this chapter provides a checkable definition: the five properties let you ask, in any meeting, which property is failing here — a far better question than why it is taking so long. For the COO, the second property is yours: end-to-end ownership of outcomes is an operating-model decision, and Chapter 8 gives you the mandate to take it. For the CIO, the fourth and fifth properties are architectural before they are cultural: shared-by-default information and designed boundaries are built, not announced. For HR, the 7S lens shows why training alone never sticks — skills are one S of seven, and the other six either reinforce or undo everything the classroom teaches. And for the CFO, the five properties are the leading indicators behind the costs of Chapter 3: where they fail, latency, friction and congestion follow — with invoices.

GATE 2 — CHOOSE ONE WORKFLOW

The decision at the end of this chapter belongs to the CEO and the COO together, and it is concrete: select the one workflow — one outcome that crosses at least three departments — through which the rest of this book will be demonstrated. Not the most broken one, and not the most strategic one. The most typical one. Write it down. It returns in the diagnostic of Chapter 4, in the design work of Chapter 7, in the value case of Chapter 9 and in the first fifteen working days of Chapter 12. Every concept in this book lands on that one workflow, so that by the final chapter you will have seen the whole method applied to something real — your own.

REFLECTION FOR THE EXECUTIVE TEAM

Five properties sound few until you try to attest to them. A useful exercise for your next offsite: take each property in turn and ask not "do we have this?" but "who in this room could prove we have this, and with what?" Direction is proven by a decision made faster than it used to be, not by a strategy document. Decision rights are proven by a transfer that stuck, not by an organogram. Information flow is proven by a meeting that no longer exists because the data removed its reason to meet. Process design is proven by a handoff that was deleted, not documented. And amplification readiness is proven by a workflow whose rules a newcomer — human or machine — could execute without asking anyone anything. Wherever the proof is thin, the property is aspirational, and an aspirational property is a red cell waiting for a heatmap to notice it.

CHAPTER 3 — WHAT INCOHERENCE COSTS

THE CFO'S CHAPTER

Chapter 1 named three costs: decision latency, interface friction, initiative congestion. Naming is not enough for a funding decision. This chapter turns the three costs into numbers your Finance function can produce within days — not with a study, but with an audit of what already exists. The method is deliberately rough. A defensible order of magnitude beats a precise illusion, and incoherence is expensive enough that precision is unnecessary to act.

COSTING THE WAITING

Decision latency is the elapsed time between *we should* and *we did*, and measuring it requires no new data — only the calendars and minutes you already have. Take the ten most significant cross-functional decisions of the last six months: investment approvals, product changes, supplier switches, reorganisations. For each, reconstruct three dates: when the question was first articulated, when it was first scheduled for decision, and when it was actually decided. The gaps are your latency. Each waiting decision carries two costs. The direct cost is the value of the delayed outcome — margin not captured, cost not removed, risk left open. The compound cost is everything downstream that waited with it: in our composite company, a pricing decision that travelled for eleven weeks delayed a campaign, two supplier negotiations and a system change — a latency multiple of roughly four times the decision itself. A working rule many CFOs find sobering: multiply the number of open cross-functional decisions by their average age in weeks. If that product exceeds the number of weeks in a quarter, your enterprise decides slower than it plans — and every plan inherits the drag.

COSTING THE FRICTION

Interface friction is the work that exists only because the parts do not fit: rework, reconciliation, double entry, alignment meetings about alignment meetings. It is the largest of the three costs and the least visible, because it hides inside normal jobs. The audit is physical. Take the one workflow chosen at Gate 2 and walk it — not the process map, the actual work — counting three things: how many hands touch it, how many

times information is re-entered or re-explained, and how much elapsed time the work spends waiting between hands versus being worked on. In typical established organisations the ratio is startling: value-adding time is often less than a tenth of elapsed time. Then estimate the loaded cost of the people involved and multiply by the share of their effort spent on coordination, correction and re-entry rather than on the outcome itself. Teams that do this honestly usually land between fifteen and thirty percent of the workflow's total cost. That share is your friction tax — structural, recurring, and payable every single quarter.

COSTING THE CONGESTION

Initiative congestion prices itself through the allocation map. List your active change initiatives; for each, name the five people it cannot succeed without; then count, per person, how many initiatives claim them. Where the same names appear three, four or five times, you have found your congestion — your most experienced people allocated at one hundred and eighty percent. Congestion does not show up as a cost line; it shows up as slippage. Every initiative in a congested portfolio runs longer and delivers less. If your portfolio of ten initiatives is staffed by the same fifteen people, you are not running ten initiatives — you are running a queue, and paying full price for the illusion of parallel progress. The corrective is not heroic prioritisation workshops. It is the third property from Chapter 2: one ranked list, owned by the executive team, where congestion is visible by construction.

THE INCOHERENCE STATEMENT

The three audits fit on one page, and that page has a structure worth keeping: decision latency, sourced from the minutes and calendars of the ten largest decisions, typically finding six to twelve weeks of average delay; interface friction, sourced from the hand-off count on one typical workflow, typically finding fifteen to thirty percent of workflow cost; initiative congestion, sourced from the allocation map of scarce people, typically finding key people at one hundred and fifty to two hundred percent. Total the three. In our experience with established mid-sized and large organisations, the sum lands between three and eight percent of revenue — before anyone mentions AI. This is the number that reframes the entire transformation discussion: the enterprise is already paying for incoherence. The question is never whether you can afford coherence. You are paying for its absence.

Three audits Finance can run within ten working days — only artefacts that exist.

COST	AUDIT SOURCE	TYPICAL FINDING
Decision latency	Minutes + calendars of the ten largest decisions	6–12 weeks average delay
Interface friction	Hand-off count on one workflow walk	15–30% of workflow cost
Initiative congestion	Allocation map of scarce people	Key people at 150–200%

3–8%

of revenue

The typical annualised cost of incoherence — before anyone mentions AI.

Figure 17: The incoherence statement — three audits, one page, one number per line.

NOW ADD AI — THE MULTIPLIER

Here is where this chapter connects to the board's agenda. Every euro of incoherence cost is amplified by AI in one direction or the other. Deploy AI into friction and you automate the friction: faster wrong forecasts, quicker misaligned hand-offs, agents confidently acting on data nobody owns. Deploy the same AI after the three costs are reduced, and the savings compound: shorter decision cycles mean the AI's speed is actually cashed; clean interfaces mean agent output flows without rework; a single priority list means AI capacity concentrates where the return is. This is why the sequence from Chapter 1 is not a preference but an economic law: design for coherence first, then amplify. The incoherence statement is therefore also your AI business case baseline. Whatever your AI programme promises, its realised value will be discounted by exactly the three costs on that page.

WHAT THIS CHAPTER CHANGES FOR YOUR ROLE

The CFO owns the incoherence statement: it takes days, not months, and it converts every later funding gate of Chapter 9 from an act of faith into a comparison against a priced baseline. For the CEO, the statement is the mandate in numbers — three to eight percent of revenue is not a transformation budget request, it is a recovery claim. The COO gains the performance-dip argument inverted: the friction audit shows the operation is already dipping, structurally, every day. The CIO learns that friction measured in re-entry and reconciliation is, in large part, an architecture invoice — and Chapter 5 shows how the backbone pays it off. And for every other role, the statement

answers the question of why the transformation touches your department: your part of the three costs is on the page.

GATE 3 — COMMISSION THE STATEMENT

The decision belongs to the CFO: commission the one-page incoherence statement, delivered within ten working days, using the three audits exactly as described — minutes and calendars for latency, one workflow walk for friction, the allocation map for congestion. No consultancy, no survey programme, no new system. Three audits, one page, one number per line. That page becomes the baseline against which every later chapter — the diagnostic, the design, the value case — is measured. Without it, the transformation will argue about opinions. With it, the enterprise argues about money it is already losing.

REFLECTION FOR THE EXECUTIVE TEAM

Three to eight percent of revenue is a range wide enough to hide in, and the natural executive reflex is to assume the low end. Resist it with a small act of arithmetic. Take last year's revenue, take three percent, and write that number where your management team must look at it. Then ask: if this sum were a business unit — one that produces nothing, consumes management attention daily, and has never once been asked to justify its budget — would we tolerate it for another quarter? Incoherence is that business unit, and it is wholly owned. The point of pricing it is not precision; no heatmap cell will ever be exact to the euro. The point is that an unpriced problem belongs to everyone and therefore to no one, while a priced problem can be assigned, attacked, and retired. The CFO who reads this chapter will already be calculating. Let that calculation happen in the open, and let it be conservative — a conservative number that is believed will move more money than an aggressive one that is debated.

CHAPTER 4 — THE COHERENCE DIAGNOSTIC

FROM SIGNALS TO EVIDENCE

Chapter 1 gave you three signals and a ten-minute self-test. Chapter 3 gave you a priced baseline. This chapter turns both into an instrument: the Coherence Diagnostic — a structured, evidence-based assessment that shows where the enterprise is incoherent, how much it matters, and what to address first. Three design principles govern the diagnostic, and they are non-negotiable. The first is evidence over opinion: every finding traces to an artefact — a decision log, a workflow walk, a data lineage, an allocation map — because interviews provide hypotheses and artefacts provide findings. The second is a timebox of fifteen working days: a diagnostic that takes a quarter becomes part of the problem it measures, and fifteen working days is long enough for evidence and short enough to keep executive attention. The third is that it ends in a gate, not a report: the diagnostic concludes with a decision session in which the executive team accepts, adjusts or rejects the findings and commits to the sequence. A report nobody decides on is interface friction in paper form.

THE FIVE DOMAINS

The diagnostic assesses coherence across five domains, corresponding deliberately to the 7S lens from Chapter 2, so that every finding lands somewhere the executive team can act. The business domain examines decision rights, end-to-end ownership and the single priority list, sourcing its evidence from decision logs, governance charters and the initiative portfolio, and asking one core question: can we name, for every outcome that matters, who owns it and who decides? The data domain asks whether information is shared by default, sourcing from the data lineage on the chosen workflow, the count of manual re-entries, and the catalogue of private spreadsheets holding the enterprise together; its core question is whether the data a decision needs reaches it without negotiation. The technology domain asks whether the system landscape connects the parts or partitions them, sourcing from the integration map, the count of point-to-point connections, the shadow-IT inventory and the AI tools already in use — sanctioned or not. The governance domain asks whether boundaries are designed, and here the diagnostic looks specifically at decision rights for AI: who may deploy an agent, what it

may touch, who audits it, where it must stop. In most enterprises this domain returns the emptiest page, and that emptiness is itself the finding. The people domain, finally, covers skills, capacity and the human side of the 7S wheel, sourcing from the allocation map of Chapter 3, the skills inventory measured against what the transformation requires, and the honest question HR already knows the answer to: do our people experience this as something done with them, or to them?

THE HEATMAP

Findings land on a five-by-three heatmap: the five domains against the three signals — waiting decisions, travelling work, competing priorities. Each cell is scored red, amber or green, but only with an artefact attached. A red without evidence is a grievance; an amber without evidence is a hope. The discipline of attaching artefacts is what separates this diagnostic from a maturity survey. The heatmap has two uses. First, it concentrates attention: in practice, three to five cells carry most of the cost, and they are rarely the cells anyone expected. Second, it sequences the work: the reddest cell with the highest cost linkage — traceable to the incoherence statement from Chapter 3 — becomes the first address of the transformation.

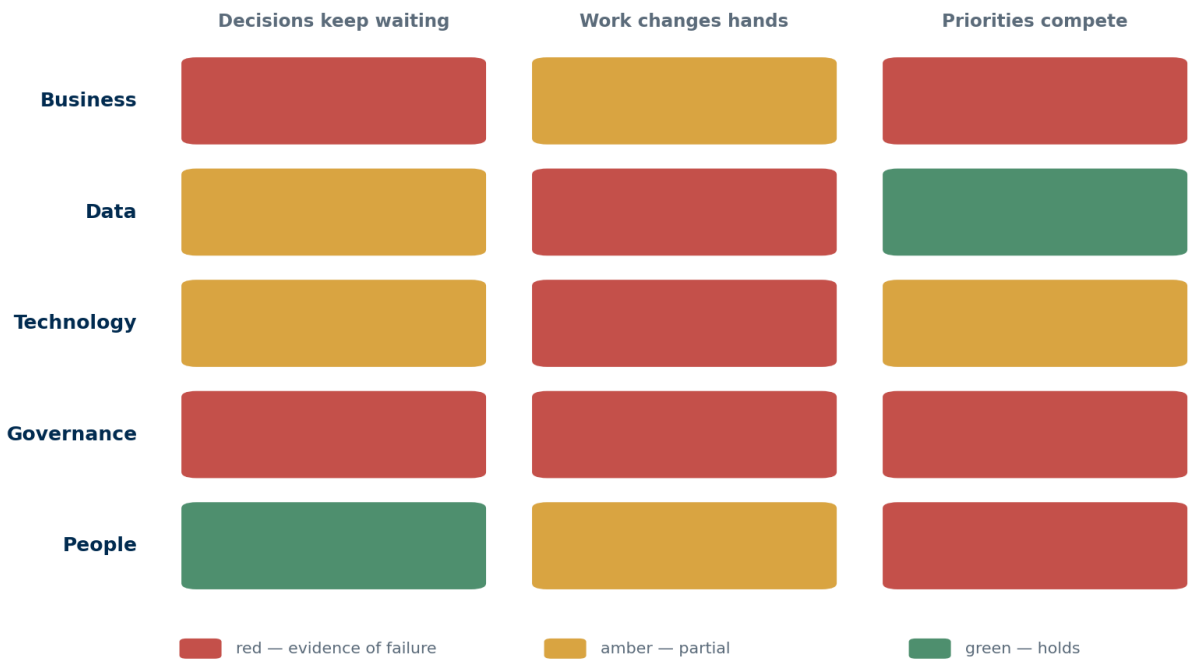


Figure 18: The coherence heatmap — five domains against three signals, every cell scored on evidence.

THE CAPABILITY AND GAP MAP

The heatmap says where it hurts; the capability map says what is missing. For each domain, the diagnostic records current capability and required capability for an AI-amplified enterprise — and the distance between them. Three gaps recur so reliably that we name them here. The ownership gap: outcomes with many contributors and no end-to-end owner. The boundary gap: AI capabilities deployed without designed limits — who may the agent act for, and where must it stop? And the capacity gap: the same scarce people claimed by every initiative at once. Your composite company from Chapter 2 had all three; most enterprises do. The diagnostic's job is not to surprise you with exotic weaknesses — it is to make the familiar ones undeniable, priced and sequenced.

FIFTEEN WORKING DAYS

The timebox is structured, and it holds because each day has a defined output. Days one to three are mandate and artefacts: the executive sponsor confirmed through the Gate 1 decision, the Gate 2 workflow selected, and artefact collection started — decision logs, the portfolio list, the integration map, the allocation data. Days four to eight are evidence gathering: workflow walks, data lineage traces and structured interviews in all five domains, with interviews testing hypotheses against artefacts, never the reverse. Days nine to eleven are scoring: heatmap cells scored with attached evidence, the capability and gap map drafted, cost linkages to the incoherence statement estimated. Days twelve and thirteen are the mirror session: findings played back to the executive team, unedited — and expect resistance in exactly the domains that score reddest, because that correlation is itself diagnostic. Days fourteen and fifteen produce the executive gate pack and the decision: a one-page summary per domain, the heatmap, the gap map, the proposed sequence, and the decision on the table — proceed to design, and with which first workflow.

WHAT THE DIAGNOSTIC IS NOT

Three boundaries keep the instrument honest. It is not an audit: nobody is graded, and findings describe the system, never individuals. It is not a strategy study: it does not ask where the enterprise should go, only whether it can move at all. And it is not repeatable bureaucracy: it is run once to start the transformation, after which the gates of Chapter 7 carry the measurement forward. If the diagnostic starts breeding quarterly re-runs, it has become the congestion it was designed to expose.

WHAT THIS CHAPTER CHANGES FOR YOUR ROLE

The CEO chairs the mirror session and the gate; your presence is the signal that this is not another assessment that dies in a drawer. The CFO receives every red cell with a cost linkage to the incoherence statement, and validates the numbers before anyone acts on them. The COO opens the doors for the workflow walks — your operation, your people, no surprises. The CIO gains, for the first time, an evidence-based map of shadow AI and unbounded agents through the technology and governance domains; this is the mandate for Chapters 5 and 6. HR becomes co-designer of the transformation through the people domain, rather than its communications department. And the business unit leaders should know that the diagnostic runs with your operation, on your real work — its findings will name your constraints, and remove them.

GATE 4 — ACCEPT THE EVIDENCE

The decision belongs to the whole executive team, and it is binary in form though not in content: accept the diagnostic findings and the proposed sequence — adjusted if the evidence supports adjustment — or state openly which finding you dispute and commission the one additional check that would settle it. What the gate forbids is the silent slide: nodding in the meeting, disputing in the corridor. An executive team that cannot decide on its own evidence will not decide on AI agents, transformation budgets or anything else that follows. This gate is, in that sense, the first real exercise of the coherence the rest of the book builds.

REFLECTION FOR THE EXECUTIVE TEAM

The diagnostic's fifteen days are not the obstacle; the decision to run it honestly is. Before you schedule one, settle three questions among yourselves. First, scope: which single value stream would you choose if you knew the heatmap would be shown to your board unedited — the one where you expect to look good, or the one where you suspect you will not? Choose the second; the first teaches nothing. Second, access: are you prepared to let the Diagnostic Lead speak to the people who carry the work, without a manager in the room, and to hear what they say without rebuttal? Third, the exit: can you genuinely accept "do nothing here" as a successful outcome, or will someone in your team feel compelled to find a case regardless? A diagnostic commissioned to confirm is worse than none, because it lends the appearance of evidence to a decision already made. If you can answer all three honestly, the calendar in Chapter 12 is ready for you.

CHAPTER 5 — THE SYSTEM AI ENTERS

YOUR BUSINESS IS TOO COMPLEX TO BE DIGITAL

Jeanne Ross of MIT CISR spent years studying why digital transformations disappoint, and her conclusion is uncomfortable in its simplicity: most enterprises are too complex to be digital. Not too small, not too old-fashioned, not underfunded — too complex. Decades of local optimisation built layers of processes, systems and governance that nobody designed as a whole. Dropping digital or AI capability on top of that complexity does not transform it. It accelerates it. Ross's remedy is the discipline this chapter adopts: simplification before digitalisation — and at its heart sits the concept of an operational backbone, the deliberately standardised, shared core of processes, data and systems on which the enterprise runs, surrounded by the variety the business genuinely needs.

SYSTEMS THINKING, IN ONE PAGE

Systems thinking sounds abstract until you state its working rule: local choices have system-wide consequences, and a decision in one place changes the work everywhere else. The planning department that optimises its own forecast degrades procurement's ordering. Procurement's safety stock degrades the warehouse's layout. The warehouse's constraints degrade the delivery promise sales just made. Nobody decided anything wrong. The system produced the outcome it was built to produce. Three habits make systems thinking operational rather than philosophical. First, optimise the flow, not the station: measure the whole workflow — the Gate 2 workflow — end to end, and accept that one department's numbers may look worse while the enterprise's outcome improves. Second, find the constraint before adding capacity: in any workflow one constraint governs throughput, and adding AI anywhere else speeds up the queue, not the flow. Third, expect delayed feedback: system interventions show their effects later than local ones, and the executive team that judges a coherence programme at week six will kill it at the exact moment it starts working.

THE OPERATIONAL BACKBONE

The backbone is the answer to a question every executive team should be able to answer in one sentence: what do we do the same way everywhere, so that everything else may differ? Its components are unglamorous and decisive: one set of master data —

customers, products, suppliers, employees, defined once and maintained once; a small number of shared core processes — order-to-cash, procure-to-pay, hire-to-retire; and the integration layer that connects them so that information is shared by default, the fourth property of Chapter 2 now given an address. Around the backbone sits the edge: the differentiated processes where the business competes — the special customer journeys, the local service models, the innovation. The rule that keeps the enterprise coherent is asymmetrical by design: the backbone standardises ruthlessly; the edge differentiates deliberately. Most enterprises do the reverse — a fragmented core drowning in variants, surrounded by innovation that cannot get data. That inversion is incoherence given architectural form.

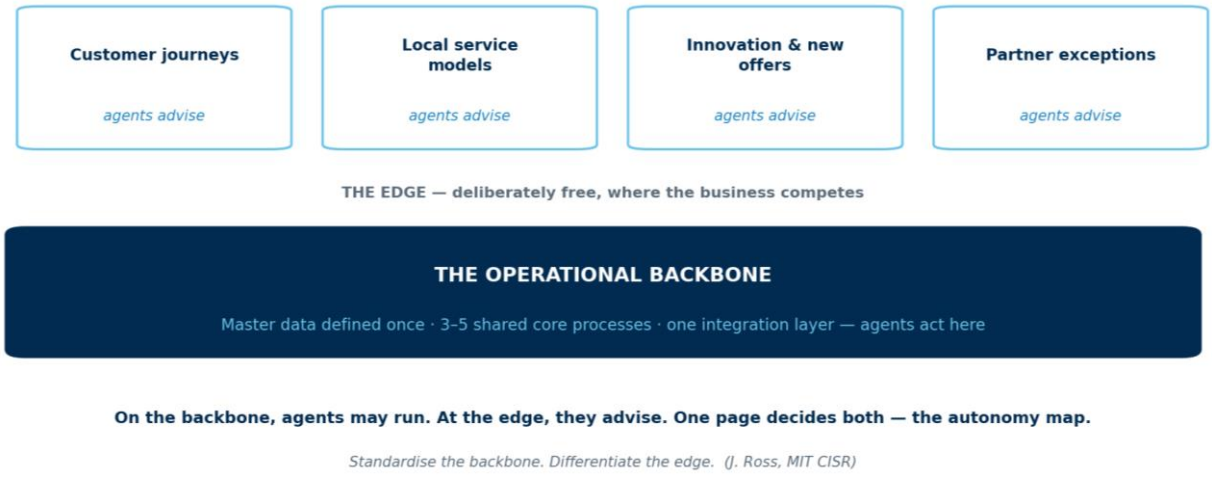


Figure 19: Standardize the backbone, differentiate the edge — the autonomy map for agents.

WHERE AI PLUGS IN

The backbone is what makes AI deployable at scale, for three reasons. Data first: an agent is only as reliable as the master data it acts on, and one customer definition, one product truth and one supplier record form the difference between an agent that works and an agent that confidently corrupts. Boundaries second: the backbone defines where standard rules hold, which is exactly where agents may act autonomously, while the edge — where judgement and variety live — is where agents recommend and humans decide; the backbone/edge distinction is, quite literally, the autonomy map for your AI. And reuse third: an AI capability built on shared data and shared processes is built once and reused everywhere, while one built on departmental fragments is built ten times and maintained forever — the AI hidden tax that Chapter 9 prices. Ross's maturity

research adds the sequence: enterprises that build the backbone first, then expose it through reusable digital services, and only then scale AI, reach value repeatedly. Enterprises that buy AI first spend their transformation budget discovering they needed a backbone — while paying vendors to work around its absence.

SIMPLIFICATION: THE WORK NOBODY VOLUNTEERS FOR

Simplification is the least celebrated and most valuable activity in this book. It means retiring process variants that exist because "we once had a customer who…", decommissioning systems kept alive by one report, and dissolving governance layers added after incidents nobody remembers. Every removed layer is decision latency, friction and congestion removed with it — the Chapter 3 costs melting at the source. Two rules keep simplification honest. Retire before you add: no new AI capability enters a workflow until the workflow's redundant variants are named and scheduled for removal. And the business decides what is standard: simplification imposed by IT alone fails politically and factually, because the enterprise knows which variants carry real customer value, and it is not the IT department.

WHAT THIS CHAPTER CHANGES FOR YOUR ROLE

For the CIO, this chapter is the strategic mandate: the backbone is not an infrastructure programme but the platform on which every AI ambition in this book stands — own it, and own the sentence that answers what we do the same way everywhere. The CFO funds the backbone as value infrastructure, not as IT cost, because its return arrives through every future initiative that no longer needs to rebuild it. The COO co-defines the core processes on the backbone with the CIO — standardisation is an operating decision, not a technical one. The CEO arbitrates the backbone/edge line, because only the executive team can decide where the enterprise will be uniform and where it will compete on difference. And the business unit leaders should hear the deal clearly: the edge is yours, and backbone discipline is what buys you the freedom — and the data — to differentiate.

GATE 5 — NAME THE BACKBONE

The decision belongs to the CIO, co-signed by the COO and the CEO: produce the one-page backbone definition — the master data domains, the three to five shared core processes, and the explicit statement of what is edge and stays free. One page is

deliberate. If the backbone cannot be stated on a page, it does not exist yet — and that finding, uncomfortable as it is, tells the executive team exactly where the transformation's first structural work lies. Every AI deployment decision that follows in Chapters 6 and 7 references this page: on the backbone, agents may run; at the edge, they advise.

REFLECTION FOR THE EXECUTIVE TEAM

Simplification before digitalisation is easy to endorse and hard to enforce, because simplification has no vendor and no launch event. Test your own organisation with the backbone question: if you drew the five processes your enterprise truly cannot live without, how many versions of each would you find running across divisions, countries, or acquisitions — and who, today, has the authority to retire the duplicates? If the honest answer is "no one," then your backbone is aspirational and every edge investment you make is negotiating with that fact. The political economy of this chapter deserves a moment of candour in the room: standardisation transfers discretion from the many to the design, and the many will experience it as loss even when it is plainly gain for the whole. That is why backbone work is executive work and cannot be delegated to architecture teams. The question for your table is not whether you agree with the principle — it is which of you will spend the political capital to make one process, one way, actually stick.

CHAPTER 6 — ENTERPRISE & AI TRANSFORMATION

FROM TOOLS TO AGENTS

The first wave of enterprise AI answered questions. The wave now arriving acts. Agentic AI — systems that plan, decide and execute across applications — changes the transformation question fundamentally. A chatbot tolerates ambiguity: a wrong answer costs a retry. An agent that places orders, adjusts schedules, reassigns work or contacts customers does not tolerate ambiguity at all. It needs an owner, a mandate, boundaries and an audit trail — the very things incoherent enterprises have never made explicit, because their humans navigate by experience and goodwill instead. This is the deepest reason AI transformations stall between demo and production: the agent forces the enterprise to write down what it has always left implicit. Who may commit a supplier? At what amount must a human approve? What may never be automated, regardless of accuracy? Enterprises that cannot answer these questions for an agent usually cannot answer them for their own organisation either. The stall is not technological. It is the moment incoherence becomes undeniable.

THE WORKFLOW IS THE UNIT OF TRANSFORMATION

McKinsey's ongoing State of AI research, surveying thousands of executives across industries, keeps returning the same finding: the companies capturing real value from AI are the ones that redesign workflows — not the ones that deploy the most models. The workflow, not the algorithm and not the department, is the unit at which transformation happens. The reason is structural. A workflow is where AI output meets human work meets enterprise outcome. Deploy a model into an unchanged workflow and you get a faster local step inside an unchanged system — the pilot paradox of Chapter 1. Redesign the workflow around what humans and agents each do best, and the same model produces a different enterprise outcome. The redesign discipline has four moves, always in this order: define the outcome — what changes for the customer or the P&L, in one sentence with a number in it; allocate the work — what the agent does autonomously on the backbone, what it prepares and a human decides at the edge, what stays human entirely; design the boundaries — authority limits, escalation triggers, never-do rules, and the audit trail that makes every agent action reconstructable; and rewire the

measures — because the workflow's KPIs must reflect the new division of labour, or the organisation will manage the old workflow and quietly undo the new one.

OWNING AN AGENT LIKE YOU OWN A PROCESS

Reference architectures for agentic AI — the composable platforms of capabilities and primitives now emerging across the industry — agree on one governance constant: every agent needs what every process needs. An owner accountable for its outcomes. A documented scope of action. Guardrails that prevent runaway autonomy — the agent that optimises its local goal at the enterprise's expense. And auditability: the ability to answer, for any action, what did it do, why, and on whose authority. Notice what this means. Agent governance is not a new discipline to be invented by the AI team. It is the enterprise's existing coherence properties — explicit decision rights, designed boundaries, end-to-end ownership — applied to a new kind of worker. Enterprises with those properties will govern agents almost for free. Enterprises without them will buy governance tooling and discover that tooling cannot supply what the organisation lacks.

THE HUMAN IN THE LOOP IS A DESIGN DECISION

"Human in the loop" is often treated as a safety slogan. It is actually the central design variable, with three positions. In the first, the human commands and the agent executes — for irreversible, high-value or regulated actions, where the agent prepares and the human commits. In the second, the agent executes and the human supervises — for reversible, bounded actions on the backbone, where the agent acts, exceptions escalate, and sampling audits quality. In the third, the agent acts autonomously — for high-volume, low-risk, fully reversible micro-decisions on verified data, where a human still owns the outcome and the stop-button. Mature organisations place each workflow step in one of these positions deliberately, review the placement as reliability grows, and — critically — earn autonomy with evidence. An agent starts supervised and graduates; it never starts free. This single policy prevents both the recklessness that creates incidents and the timidity that creates shelf-ware.

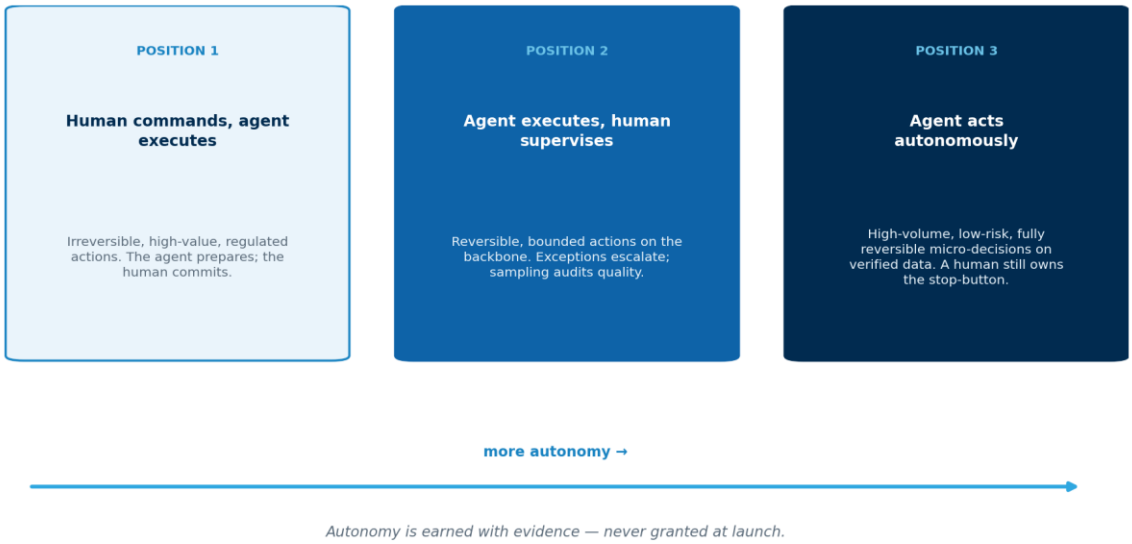


Figure 20: Three positions for the human in the loop — autonomy is earned with evidence.

THE PERFORMANCE DIP, MANAGED HONESTLY

Every workflow redesign carries a performance dip: the weeks in which the new division of labour produces worse numbers than the old routine, because people are learning, exceptions are surfacing, and measures are recalibrating. The dip is not a risk of transformation; it is the transformation, visible. What kills programmes is not the dip but the surprise of it. So the dip is managed as a planned quantity: announced in advance to whoever owns the numbers, bounded in duration, measured weekly, and exited by explicit criteria rather than by morale. The COO who can say "this workflow will dip six percent for eight weeks, here is the recovery curve, and here is the trigger that stops the experiment" converts the board's deepest fear into a managed variable. Chapter 7 makes this a standard step; Chapter 9 shows the dip's cost is typically recovered within the first year of verified value.

COGNITIVE CAPITAL, NOT COST CUTTING

A final framing, because it determines what your AI programme dares to attempt. Agentic systems are best understood as cognitive capital — capacity that reasons, not just capacity that types. The enterprises that treat it as ten-percent-cheaper labour get ten-percent outcomes and a disappointed board. The ones that treat it as new capability — analyses nobody could afford, response times nobody could staff, personalisation

nobody could operate — find the ten-times value. The difference is not the technology budget. It is whether the enterprise is coherent enough to absorb capability: one priority list to aim it, a backbone to run it on, boundaries to trust it, and owners accountable for what it does. Everything in the earlier chapters converges here. Coherence is not the prerequisite that delays AI value. It is the multiplier that creates it.

WHAT THIS CHAPTER CHANGES FOR YOUR ROLE

The CEO sets the ambition at cognitive capital, not cost cutting — and insists that every agent has a named owner before it has a budget. The COO owns workflow redesign and the performance dip; Chapter 7 hands over the method, and this chapter the argument. The CIO builds agent governance into the platform, not around it — it is the architecture expressing the enterprise's decision rights. The CFO treats the dip as a planned, bounded investment line: fund it explicitly and it stops being a surprise; leave it implicit and it becomes the reason programmes die. HR recognises the three human-in-the-loop positions as job design — role profiles, skills and trust all move with them, and Chapter 10 builds the curriculum. And for Risk and Legal, the boundaries, audit trails and never-do rules of this chapter are your language: this is where you become an enabler rather than a brake.

GATE 6 — THE AUTONOMY POLICY

The decision belongs to the CIO and the COO jointly, ratified by the CEO: adopt the one-page autonomy policy — the three human-in-the-loop positions, the rule that autonomy is earned with evidence, the ownership requirement for every agent, and the never-do list for the enterprise. With that page signed, the enterprise can finally say yes to agents at speed, because it knows exactly what it is saying yes to. Without it, every agent proposal re-opens every question — decision latency in its most expensive form, now multiplied by the number of vendors at the door.

REFLECTION FOR THE EXECUTIVE TEAM

Where the human sits in the loop is the most consequential design choice in this book, and the one most often made by default. In your next review of an agent deployment, ask three things. By whose analysis was this position chosen — and if the answer is "the vendor's reference architecture," treat that as a red flag, because the vendor's incentive is throughput and yours is judgement. What does the reviewer actually see when they

review: the case and the reasoning, or a score and a button? A reviewer who cannot see the reasoning is not in the loop; they are a signature on it. And what happens when the reviewer disagrees — is the disagreement captured as data that improves the system, or does it evaporate as friction? The mature version of this discipline feels slower at first and is faster forever after, because every returned case teaches the system the boundaries it could not be told. The organisations that get this right will be identifiable by a small marker: their reviewers can explain, in one sentence, why the human sits exactly where they sit.

CHAPTER 7 — FIVE STEPS, FIVE GATES

ONE RHYTHM, FROM MANDATE TO SCALE

Everything so far — the signals, the costs, the diagnostic, the backbone, the agent discipline — converges in this chapter: the working method of the coherent transformation. Five steps, each timeboxed, each ending in a gate with a named decision owner. Not a programme that runs for years: a rhythm you can start within fifteen working days. The gates are the design's core. In an incoherent enterprise, transformations fail silently — activities continue while commitment evaporates. Gates make failure impossible to hide: every step ends with a decision, by a named owner, on evidence. A gate can pass, adjust, or stop. What it cannot do is drift.

STEP 1 — EXECUTIVE MANDATE & DIAGNOSIS

The first step establishes the evidence base and the executive commitment that carries everything after it, and it is timeboxed at fifteen working days. The executive team confirms the mandate through Gate 1, selects the demonstration workflow through Gate 2, and commissions the incoherence statement through Gate 3 together with the Coherence Diagnostic of Chapter 4 — the five domains, the heatmap with artefacts, the capability and gap map. The output is deliberately compact: one page of costs, one heatmap, one gap map, one selected workflow. The gate belongs to the CEO, who decides to proceed on this evidence, with this workflow, in this sequence.

STEP 2 — READINESS & CAPABILITY

The second step, ten working days, determines what the enterprise can absorb — honestly. It produces seven deliverables: the readiness assessment across business, data, technology, governance and people; the refined heatmap; the capability and gap map; the governance and responsible-AI baseline, consisting of the backbone page of Gate 5 and the autonomy policy of Gate 6; the workforce readiness picture; the risk register; and the action pack for the next step. The output is a readiness verdict per domain: green means proceed, amber means fix-in-parallel with a named owner, red means fix first — and the discipline to honour that. The gate belongs to the CEO with the CIO's co-signature: the enterprise is ready enough, with named conditions.

STEP 3 — DIRECTION

The third step, again ten working days, aims the transformation at value, in the right order. It begins with the North Star AI Vision: one page stating which enterprise outcomes AI will change within three years, expressed in the language of the P&L, not in technology. It then sequences candidate initiatives on two axes only — impact, verified against the Chapter 3 statement, and feasibility, verified against the readiness verdict of Step 2. High impact and high feasibility starts now. High impact and low feasibility builds the missing capability first — that is what readiness gaps are for. Low impact and high feasibility is politely declined, because easy pilots are how portfolios drown. The output is one North Star page, one sequenced portfolio, one funded first move. The gate belongs to the CEO: the vision and the sequence are the enterprise's, not IT's.

STEP 4 — WORKFLOW ACTIVATION

The fourth step runs per workflow, six to twelve weeks each, and it banks the value. It applies the four moves of Chapter 6 to the selected workflow: the outcome defined, the work allocated across the three human-in-the-loop positions, the boundaries and audit designed, the measures rewired. The performance dip is announced, bounded and tracked. Training lands in the workflow, as Chapter 10 describes, not before it. And value is measured by Finance against the baseline, not by the project team against its own enthusiasm. The output is one workflow running on the new division of labour, with verified numbers. The gate belongs to the COO: the workflow exits its dip on evidence — or it is adjusted, or it is stopped.

STEP 5 — ANCHORING & SCALING

The fifth step is continuous, and it makes the gains permanent and repeatable. Governance moves from project to line: agent audits, boundary reviews, backbone stewardship. Funding moves from creation to curation: most AI capacity goes to running and improving what works, not launching new experiments. The portfolio scales along the backbone — each next workflow reusing the data, agents and patterns of the last, which is why each next transformation costs less than the previous one. The capability curriculum keeps the skills in-house. The gate belongs to the CFO: enterprise funding flows only to initiatives at one and a half times verified annual value or better — validated by Finance, at the gate, every time.

STEP	THE GATE DECIDES	OWNER	TIMEBOX
1 · Executive Mandate & Diagnosis	Proceed on this evidence, this workflow, this sequence	CEO	15 working days
2 · Readiness & Capability	Ready enough — with named conditions	CEO + CIO	10 working days
3 · Direction	North Star vision and sequence are the enterprise's	CEO	10 working days
4 · Workflow Activation	Exit the performance dip on evidence — or adjust, or stop	COO	6–12 weeks / workflow
5 · Anchoring & Scaling	Enterprise funding only at ≥ 1.5× verified annual value	CFO	Continuous · quarterly

Figure 21: Five steps, five owners, one rhythm — the gates rotate deliberately.

WHY THE GATES WORK

Look at the ownership pattern: CEO, CEO with CIO, CEO, COO, CFO. The gates deliberately rotate across the executive team. Coherence cannot be delegated to a transformation office, because the condition being fixed lives between the offices. When the CEO opens, the CIO verifies, the CEO aims, the COO banks and the CFO compounds, the executive team practises — in miniature, every few weeks — exactly the cross-functional decision behaviour the whole enterprise is learning. The rhythm also protects against the two classic failure modes. Against pilot theatre: no workflow enters Step 4 without North Star linkage and a value baseline. Against big-bang exhaustion: no step lasts longer than twelve weeks without producing either evidence or a stop-decision.

WHAT THIS CHAPTER CHANGES FOR YOUR ROLE

The CEO carries Gates 1 and 3 — two decisions, well prepared, set everything else in motion, and this method is designed so the top job carries the mandate rather than the management. The CIO's co-signature on Gate 2 makes readiness a shared executive fact instead of an IT opinion. The COO's Gate 4 puts the transformation where it belongs: in the operation, banking operational value. The CFO's Gate 5 and its 1.5× rule end the era of AI funding by enthusiasm — Chapter 9 supplies the measurement rules. And HR is present at every gate, because every gate carries a people condition: readiness, training-in-the-workflow, capability retention.

GATE 7 — ADOPT THE RHYTHM

The decision belongs to the whole executive team: adopt the five-step rhythm as the enterprise's single AI transformation method — these steps, these gates, these owners, these timeboxes. One method matters more than a perfect method. An enterprise with five competing transformation approaches is initiative congestion in its purest form. The signature on this gate is the moment the executive team stops discussing coherence and starts operating it.

REFLECTION FOR THE EXECUTIVE TEAM

Five gates exist to make stopping cheap, and their value is only ever tested at the moment stopping is warranted. Before your next gate review, examine the last three initiatives your organisation continued past the point where the evidence had cooled, and ask what made continuation easier than closure: the sponsor's reputation, the sunk budget, the absence of kill criteria written before anyone cared, or the simple fact that no one's calendar contained a meeting whose purpose was to decide. A gate without kill criteria is a milestone, and milestones only measure passage. The counter-intuitive discipline this chapter asks of the executive team is to celebrate a well-stopped pilot as visibly as a successful one — to say, in the town hall and in the review, what was learned and what was saved by stopping. Organisations that punish honest closure do not get bolder pilots; they get quieter evidence. Which behaviour your next pilot team chooses will be determined long before the pilot, by what they saw happen to the last one that stopped.

CHAPTER 8 — WHAT EACH ROLE CONTRIBUTES

NO SPECTATORS

The premise of this book is that coherence is built by the executive team as a team. This chapter makes that concrete: for each of nine roles, what you contribute, what you receive, and which gate carries your name. Read your own role first — then read the others, because coherence fails precisely where roles stop understanding each other's part.

THE CHIEF EXECUTIVE — OWNER OF THE CONDITION

The CEO contributes the mandate: the Gate 1 decision to proceed on evidence, the Gate 3 decision that sets the North Star course, and the continuous message that coherence is the enterprise's operating condition, not a project. The CEO chairs the mirror session and protects the single priority list against the politics of sponsorship. What the CEO receives is an executive team that decides on evidence, a transformation explainable to the board on one page, and an organisation whose speed no longer depends on personal escalation. The risk, if the CEO is absent, is that every gate degrades into a meeting — coherence cannot be delegated, because it lives between the seats only the CEO can connect.

THE FINANCIAL DIRECTOR — KEEPER OF VERIFIED VALUE

The CFO contributes the incoherence statement, the value measurement rules of Chapter 9, and Gate 5 — funding that flows only against one and a half times verified annual value. It is the CFO who defines what "verified" means: measured against baseline, net of the performance dip, confirmed by the finance function rather than claimed by the project. What the CFO receives is an AI portfolio defensible line by line, the end of licence-leak, and the incoherence recovery — three to eight percent of revenue — as the largest cost-reduction programme available without a restructuring. The risk of absence is the return of funding by enthusiasm, and with it the pilot graveyard.

THE OPERATIONS CHIEF — BANKER OF OPERATIONAL VALUE

The COO contributes end-to-end workflow ownership, the workflow walks, and Gate 4 — the decision that a redesigned workflow exits its performance dip on evidence. The COO co-designs the backbone's core processes with the CIO and manages the dip as a planned quantity instead of a surprise. What the COO receives is cycle times thirty to forty percent shorter in the chosen workflows, an operation that absorbs change without pausing, and the end of being the excuse every failed initiative hid behind. Without the COO, the transformation remains a headquarters exercise and dies at the first real workflow.

THE TECHNOLOGY CHIEF — ARCHITECT OF COHERENCE

The CIO contributes the operational backbone, the autonomy policy, the co-signature on readiness, and an architecture in which information is shared by default and agent boundaries are built in rather than audited afterwards. What the CIO receives is a mandate that finally matches the accountability, the end of shadow AI as an unsolvable problem — it becomes visible, priced and governed — and a platform whose value compounds with every reuse. Without the CIO, AI lands on fragmentation, and every incident becomes the CIO's to explain.

THE HR DIRECTOR — BUILDER OF CAPABILITY AND TRUST

The HR director contributes the people domain of the diagnostic, the role-based curriculum of Chapter 10, job design for the three human-in-the-loop positions, and honest work on the fear — replacement anxiety is addressed by redesigning work with people, not by reassuring them about it. What HR receives is a transformation people experience as growth rather than threat, a skills inventory that finally matches strategy, and a seat at every gate, because every gate has a people condition. Without HR, the resistance catalogued in Chapter 11 arrives on schedule, wearing the mask of change fatigue.

THE COMMERCIAL DIRECTOR — VOICE OF THE PROMISE

The Sales and Marketing director contributes customer outcomes for the North Star vision, honesty about what the operation can currently keep, and the discipline of selling what the redesigned workflows can actually deliver — then more, as coherence compounds. What this role receives is promises the enterprise keeps, response times

competitors cannot match, and an end to selling on personal heroics. Without it, customer expectations — raised by your own AI-powered competitors — outrun an operation nobody redesigned.

PROCUREMENT — GATEKEEPER OF THE VENDOR FLOOD

The Procurement director contributes the supplier gate: no AI contract without an owner, a value hypothesis and an exit clause; consolidation on the backbone rather than ten departmental tools; lock-in resisted before signature, when it is still cheap to resist. What Procurement receives is leverage — an enterprise with a backbone and a method negotiates from architecture, not from need. Without it, the AI hidden tax of Chapter 9 arrives through Procurement's own contracts.

SUPPLY CHAIN — GUARDIAN OF THE FLOW

The Supply Chain director contributes the end-to-end view nobody else has — forecast-to-delivery crosses more boundaries than any other flow — and the realism about buffers: efficiency that removes all slack removes resilience with it. What Supply Chain receives is forecasts the organisation actually uses, planning cycles AI genuinely shortens, and resilience designed in rather than begged for after the disruption. Without it, AI optimises the forecast while the chain keeps guessing — faster, more confident, wrong.

THE BUSINESS UNIT MANAGERS — OWNERS OF REALITY

The business unit managers contribute the workflows, the people and the truth. The diagnostic walks their floors; the redesign happens in their operation; the dip lands on their numbers — which is why they receive the recovery curve in advance and a voice at the gate. What they receive is their best people back from initiative congestion, benefits that stay visible in their own results instead of being claimed centrally, and tools that finally fit the work. Without them, the transformation designs for an enterprise that exists on slides, not in the units.

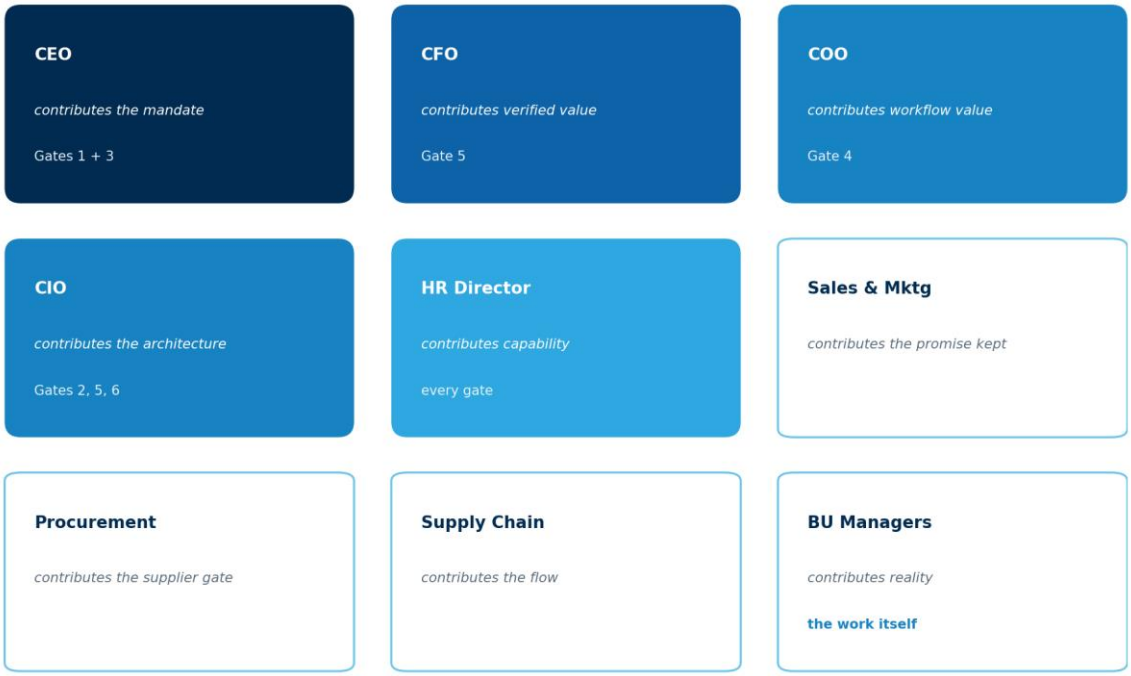


Figure 22: Nine roles — each contributes authority, each receives value, each owns a gate.

THE PATTERN UNDERNEATH

Nine roles, one structure: each role contributes authority the others lack, receives value the others cannot give, and owns exactly the gate where its accountability already lies. That is coherence applied to the executive team itself — and it is why this chapter may be the most photocopied in the book. When a transformation stalls, the question is no longer who is resisting, but which role's contribution is currently empty.

GATE 8 — SIGN YOUR ROLE

The decision belongs to each executive individually, in one session: read your role's contribution aloud, amend it where reality requires, and sign it. The amended contributions become the executive team's working contract for the transformation. This takes one meeting and is worth more than most kick-offs: it is the first time the team states, in each other's presence, what each member is actually responsible for. Enterprises that skip this conversation spend the whole transformation having it anyway — in corridors, at double the price.

REFLECTION FOR THE EXECUTIVE TEAM

Nine roles is not an argument for nine hires; it is an argument that nine accountabilities exist whether or not anyone owns them. Take the list into your next executive meeting and run a simple exercise: for each role, write the name of the person in your organisation who holds it today. Where the same name appears four times, you have found a bottleneck with a diary. Where no name appears, you have found the accountability your next failure will be traced back to. Two cautions belong with the exercise. Do not assign the roles as additional duties to your already-strongest people without removing something in return — coherence work added to an intact calendar is coherence work that will be done at midnight, badly. And do not confuse the role with the title: the Value Analyst who reports to the programme they are measuring is not a Value Analyst but a press office. The test of the whole design is singular: when something goes wrong at two in the morning, does everyone in this room know whose phone rings? If the answer takes a discussion, it is a no.

CHAPTER 9 — THE VALUE CASE

WHAT THE P&L CAN EXPECT, AND WHEN

Every executive who has followed the argument to this page asks the same question: what is it worth, and when? This chapter answers with three horizons, four value types and five measurement rules. The numbers are deliberately structured as a model, not a promise — your Finance function fills it with your figures from the incoherence statement of Chapter 3 and the first workflow of Chapter 7. What the model guarantees is not the amount but the shape: when value arrives, in what form, and under what conditions it compounds.

THE THREE HORIZONS

In the first horizon, months zero to six, the returns arrive before any AI scales, because coherence work removes cost directly. Misinvestments are stopped: the single priority list and the supplier gate kill initiatives and contracts that would have consumed budget without return, and avoided cost is real cost — it simply needs Finance to count it. Decision latency shortens as explicit decision rights cut the cycles measured in Chapter 3, so delayed outcomes arrive earlier and optionality stops expiring. And licence and coordination costs consolidate, as shadow AI and duplicate tools surface in the diagnostic and move onto the backbone. The typical shape is modest, real, and bankable within two quarters; index the value at this horizon as 1.0, the baseline against which everything after it is compared.

In the second horizon, months six to eighteen, the redesigned workflows come out of their performance dips and produce the numbers. Cycle times run thirty to forty percent shorter in the chosen workflows — the range we consistently see when hand-offs are designed rather than discovered and agents carry the coordination load. Capacity is released as the friction tax of Chapter 3 — fifteen to thirty percent of workflow cost — converts into freed hours: the same people producing more outcomes, or the same outcomes with structural cost removed. And quality and auditability rise together, because designed boundaries and logged agent actions reduce rework, errors and compliance exposure at the same time. Cumulative verified value at this horizon typically indexes at two to two and a half times the first — call it 2.3 — with the dip costs netted out, because the measurement rules below insist on it.

In the third horizon, months eighteen to thirty-six, coherence compounds. Each next workflow costs less than the last, because backbone data, agent patterns, training materials and governance are reused — the second redesign typically takes half the effort of the first. The backbone lowers structural operating cost permanently, as shared master data and consolidated platforms retire the reconciliation work and licence overlap. And optionality returns: an enterprise that decides fast and deploys fast can respond to markets, regulations and technology shifts as opportunities rather than threats — a benefit that fits no spreadsheet cell, and that boards feel most. The cumulative value index reaches 4.0 and keeps rising, because the mechanism compounds rather than depletes.

THE FOUR VALUE TYPES

Finance will ask which column of the ledger each benefit touches, and the discipline has four answers. Cost avoided lands in operating expense immediately — stopped initiatives, consolidated licences, retired systems. Cost reduced lands in operating expense structurally — the friction tax removed, coordination capacity released. Revenue enabled arrives on the top line, lagged — faster quotes, kept promises, offerings that were impossible before. And risk reduced lands in expected loss — bounded agents, audit trails, resilience by design. A value claim that cannot name its type is not yet a value claim. This one rule eliminates most transformation theatre at birth.

THE FIVE MEASUREMENT RULES

The credibility of the entire case rests on five rules. Baseline first: no workflow claims value without its Chapter 3 baseline measured before redesign, because value without a baseline is a story. Finance verifies: the project proposes and Finance disposes, because self-certified value is marketing. Net of the dip: performance-dip costs are charged to the initiative that caused them, because honest net numbers build the trust that funds horizon three. Creation to curation: value is re-verified quarterly after go-live, and benefits that decay get fixed or get removed from the books, because claimed-once value that silently evaporates is how AI programmes lose board confidence. And the 1.5× rule: enterprise-wide funding — the expensive, structural kind — flows only to initiatives demonstrating at least one and a half times verified annual value; below that threshold,

initiatives may continue at experimental scale, but they do not get to rebuild the enterprise around themselves.

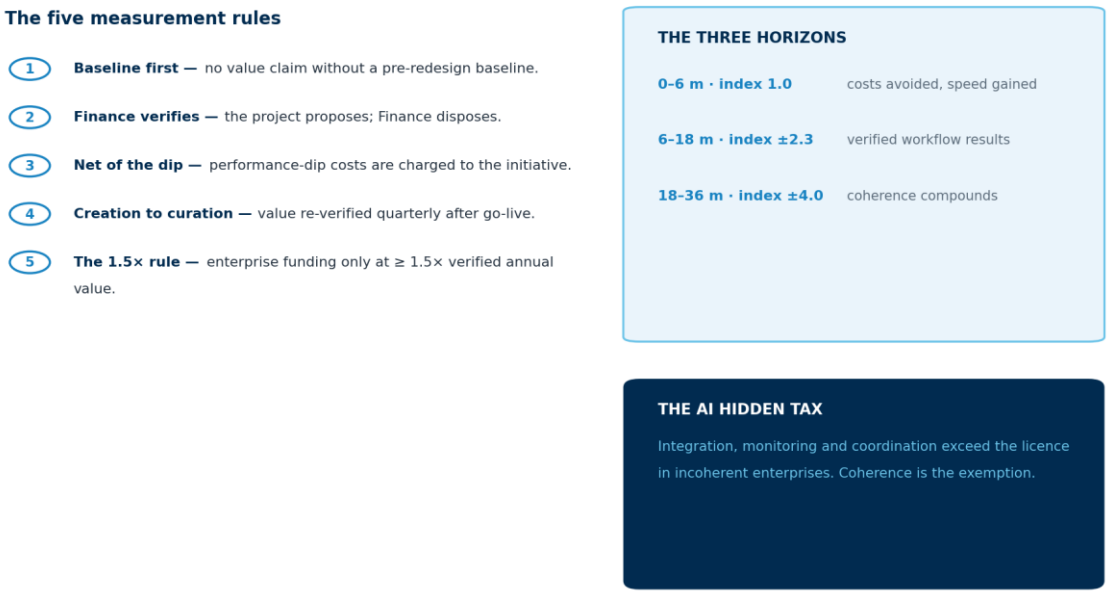


Figure 23: The five measurement rules, the three horizons — and the exemption from the hidden tax.

THE AI HIDDEN TAX

One cost line deserves its own warning, because it ambushes even disciplined programmes: the total cost of AI is structurally larger than the licence. Integration into fragmented data, prompt and boundary engineering, evaluation, monitoring, retraining, vendor switching costs, and the coordination overhead of unowned agents — in incoherent enterprises these overheads routinely exceed the licence itself. That is the hidden tax. Coherence is the tax exemption. Shared master data removes most integration cost. Designed boundaries remove most re-prompting and error-handling. Ownership removes the coordination overhead. The same AI, run coherently, costs a fraction of the same AI run on fragments — which is why the backbone of Chapter 5 keeps appearing in this chapter's arithmetic.

WHAT THIS CHAPTER CHANGES FOR YOUR ROLE

The CFO owns this chapter: the three horizons are the narrative to the board, the five rules are the defence against enthusiasm, and the 1.5× rule is the gate. The CEO presents

the horizons as a shape, never as a guarantee — boards trust executives who promise a mechanism over executives who promise a number. The COO delivers horizon two: the thirty to forty percent cycle improvements are operational facts before they are financial claims. The CIO finds in the hidden tax the argument for the backbone, stated in money. And every role learns the same discipline: every benefit your department claims must survive the five rules — uncomfortable at first, and precisely what makes your later claims believed without discussion.

GATE 9 — ADOPT THE MEASUREMENT RULES

The decision belongs to the CFO, ratified by the executive team: adopt the five measurement rules as enterprise policy — baseline first, Finance verifies, net of the dip, quarterly re-verification, and the 1.5× funding threshold. With this gate, the transformation's economics are locked before the serious money moves. Every subsequent funding conversation becomes arithmetic instead of advocacy — which is, in the end, what a CFO-friendly transformation actually means.

REFLECTION FOR THE EXECUTIVE TEAM

Of all the disciplines in this book, honest measurement is the one that will be resisted most politely. Nobody will argue against the counterfactual in principle; the resistance arrives as practicalities — the baseline is hard to reconstruct, the pilot period was atypical, the freed capacity is "strategically redeployed" in ways that resist quantification. Your role as an executive team is to notice that every one of these practicalities bends the number in the same direction. Set the tone early with two commitments. Commission the baseline before the pilot, when everyone still expects success and no one has an interest in a low starting point. And when the first verified result comes in below the promised ratio, receive it as the system working — the gate doing its job — rather than as a performance to be explained. The alternative is the arithmetic of hope, and hope compounds worse than interest: each generously counted pilot raises the evidentiary standard the next one is quietly excused from meeting. A board that trusts its numbers can be brave; a board that suspects them can only be loud.

CHAPTER 10 — TRAINING & CAPABILITY

THE CLASSROOM IS NOT WHERE TRANSFORMATION HAPPENS

Every failed transformation has a training certificate. The courses were delivered, attendance was logged, satisfaction scores were high — and Monday morning looked exactly like the previous one. The problem was never the teaching. It was the location: capability was built in a room, while the work stayed on the floor. This book takes the opposite position, and Chapter 7 already smuggled it in: training lands in the workflow, not before it. Capability is built by doing redesigned work, with support, on real cases — what the best programmes call progressive exposure: people meet AI first as observers of their own process, then as users with a safety net, then as supervisors of agents, and finally as designers of the next workflow. Four phases, each earned by the last, none skippable.

THE FOUR PHASES OF PROGRESSIVE EXPOSURE

In the first phase, people see it. Before anyone is trained on anything, the team watches the redesigned workflow operate — or a faithful simulation of it. The goal is not skill; it is orientation. People who have seen agents handle their own invoices, their own forecasts, their own customer emails stop arguing from imagination. Most resistance evaporates here, because most resistance was fear of an abstraction.

In the second phase, people use it, with a net. Team members operate the workflow with the agent as an assistant: it prepares, they decide. Every action is reversible, and a coach — a colleague, not a consultant — is present. The human-in-the-loop positions of Chapter 6 are experienced physically, not read from a slide.

In the third phase, people supervise it. As reliability evidence accumulates, users become supervisors: agents act autonomously within boundaries, humans handle exceptions, audit samples, and tune escalation triggers. This is the new core skill of the agentic enterprise — supervision: knowing when to trust, when to check, when to stop. It is teachable, and it is taught here.

In the fourth phase, people design it. The most experienced practitioners join the next workflow redesign as co-designers. The flywheel closes: the people who run the

transformed work become the people who extend it. Capability stays in-house — by design, not by hope.

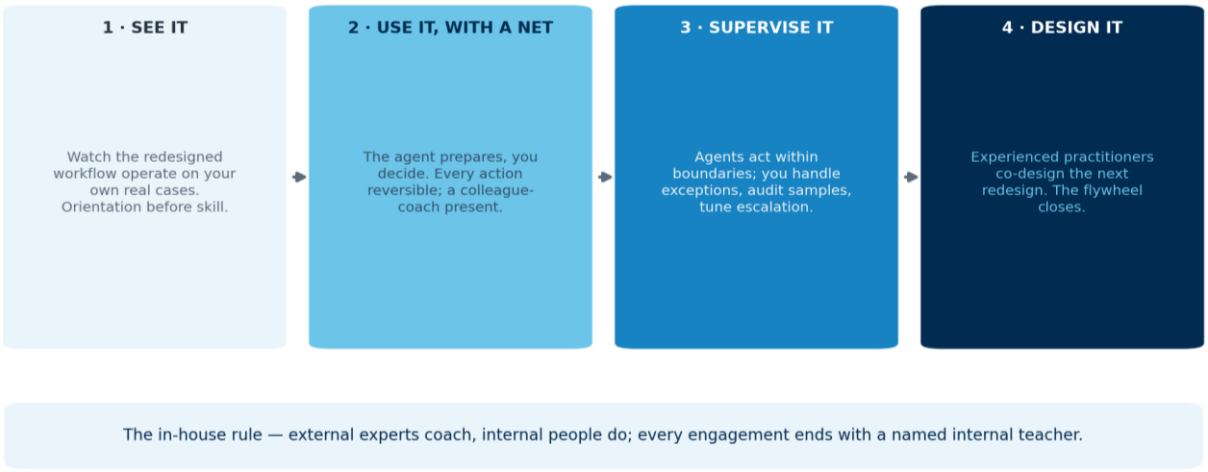


Figure 24: Progressive exposure — four earned phases, from seeing to designing.

THE CURRICULUM, BY ROLE

Progressive exposure is the spine; the content differs per seat at the table. Executives need three things: the vocabulary of this book's concepts in half a day, the economics of Chapters 3 and 9 in another half, and gate-craft — how to decide on evidence, run a mirror session, and stop a drifting initiative — in one day, practised on a real gate. Workflow owners and team leads need the redesign method of Chapter 6, dip management, and the measurement rules: two days, then apprenticeship on a live redesign. Professionals in the workflow need the first three phases on their own work — hours, not days, and always on real cases. The CIO organisation and data teams need the deep track: agent architecture, boundary engineering, evaluation and audit — the technical curriculum behind the autonomy policy. And HR needs everything above, one level up: enough to design roles, profiles and assessments for an enterprise of humans supervising agents. Notice what is absent: generic AI awareness training for everyone. Awareness without a workflow is trivia. People learn what their work requires, when their work requires it.

TRAIN THE WORKFLOW, NOT THE PERSON

The unit of training mirrors the unit of transformation: the workflow. When a workflow is redesigned, its whole team passes through the phases together — including its manager, its adjacent departments, and its executive sponsor for the first phase. Three reasons. Shared language: the team that learns together argues about the work, not about the words. Mutual visibility: everyone sees what the agent does, what colleagues decide, and where the boundaries are — trust is a by-product of watching the system behave. And no orphan skills: individually trained enthusiasts who return to an untrained team lose their skills within weeks. The workflow either changes together or not at all.

KEEPING CAPABILITY IN-HOUSE

Consultancies build dependency by default: they design, they configure, they leave — and the knowledge leaves with them. The counter-design has three rules. External experts coach, internal people do: every deliverable in this method — the heatmap, the backbone page, the redesign, the value case — is produced by your people with expert support, never by experts with your people watching. Every engagement ends with a teacher: the explicit exit deliverable of any external support is a named internal person who can teach the skill onward. And the curriculum is an asset: materials, cases and assessments live on your backbone, versioned like code, improved with every cohort. The goal is not self-sufficiency as ideology — specialists will always have their place. The goal is that the enterprise owns its own capacity to change, because in an AI-accelerated market, the speed of learning is the one advantage competitors cannot simply purchase.

CULTURE, THROUGH TWENTY-ONE CONVERSATIONS

Culture programmes fail when they announce; they succeed when they structure conversation. A practical instrument: a curated series of short internal talks — twenty-one over the transformation's first year is a proven cadence — where leaders and practitioners discuss real cases: a decision that waited, a boundary that held, a dip that recovered. Not success theatre: the failures teach more. Twenty-one honest conversations, each ending with one question to the room, move an enterprise's shared values further than any values poster — because shared values, as the 7S lens insists, sit at the centre and hold everything else.

WHAT THIS CHAPTER CHANGES FOR YOUR ROLE

The HR director owns this chapter — the four phases, the role curricula, the in-house rule — and it makes HR the architect of the enterprise's most durable asset. The COO hosts the training-in-the-workflow inside the operation, in the dip already being managed: the same weeks, the same people, one integrated design. The CIO staffs the deep technical track and keeps it current, because agent supervision ages quickly. The CFO funds capability as the compounding asset it is in the value case: measured, maintained, never the first line cut. And the CEO opens the twenty-first conversation personally — what the top of the house attends, the enterprise takes seriously.

GATE 10 — ADOPT THE CAPABILITY RULE

The decision belongs to the HR director, ratified by the executive team: adopt progressive exposure as the enterprise's single capability method — training in the workflow, in phases, by role, with the in-house rule for all external support. With this gate, every later initiative inherits its capability plan by default. Skills stop being the excuse discovered at go-live — they become the engine that makes each next transformation cheaper, which is horizon three's arithmetic in human form.

REFLECTION FOR THE EXECUTIVE TEAM

Progressive exposure asks the organisation to spend its scarcest asset — the time of its best people — before any return is visible. That is precisely why it works, and precisely why it will be the first discipline sacrificed when the quarter tightens. Two questions will tell you how your organisation really treats it. When a phase of training collides with an operational peak, which calendar wins — and who has the authority to make that call? And when your most experienced expert finishes the fourth phase, what have you built: one certified judge, or a certified judge who has already begun observing the next observer? The difference between those two answers is the difference between training and capability. There is also a question of respect embedded in this chapter that your people will read accurately whatever you say: an organisation that asks experts to judge machine output without teaching them how the machine reasons has announced what it thinks expertise is worth. The reverse announcement — that judgement is valuable enough to warrant four phases of investment — is one of the few retention strategies that does not appear in any remuneration survey.

CHAPTER 11 — HOW THIS FAILS: THE SEVEN SIGNALS AND THEIR COUNTER-MOVES

A book that promises you coherence owes you a chapter on how coherence is lost in practice, because it will be attempted on your watch, and the failure modes are not exotic. They are ordinary, familiar, and — this is the uncomfortable part — they usually arrive dressed as success. The pilot that was green everywhere stalls anyway. The training that tested brilliantly never reaches the line. The dashboard that was supposed to end the argument becomes the next thing the meeting argues about. None of these accidents is random. Each is a known signal, and each has a counter-move that can be taken before the signal hardens into a failure you will have to explain to the board.

SIGNAL	COUNTER-MOVE
1 · Pilot theatre	North Star linkage test — workflow, baseline, horizon-2 value, or no budget line
2 · Governance after the fact	Autonomy policy before the next deployment — boundaries designed, not excavated
3 · Training without workflow	No course without the redesign it serves; fund embedded learning instead
4 · Funding without gates	The five measurement rules — the CFO becomes referee, not obstacle
5 · The invisible dip	Announce every dip in advance — size, duration, exit criteria, stop-trigger
6 · Central benefit harvesting	Benefits stay visible where the work happens; corporate reports the total only
7 · The leadership hand-over	The method is enterprise policy with role-owners — gates survive rotations

Figure 25: The seven failure signals and the counter-move that belongs to each one.

The first signal is the **hidden veto**: the pilot continues, metrics are met, the team is praised — and the person whose decision rights the design actually challenges quietly starves it of cases, calendar, and people. Nothing is refused, so nothing can be confronted; the work simply never gets the input it needs. The counter-move belongs to the sponsor, and it is structural rather than personal: name the decision rights that change on day one, have the owner sign that transfer in the Gate 1 review, and instrument the flow so that starved cases are visible as facts rather than felt as sabotage. A veto that has to happen in the open usually does not happen at all.

The second signal is **mispriced value**: the pilot was scaled because everyone liked it, not because anyone verified it, and six months later the CFO asks where the money is — and the honest answer is that nobody priced it at the start. Enthusiasm is not a business case, and a pilot without a priced counterfactual cannot be defended when the budget cycle turns. The counter-move is the discipline of Chapter 9: price the case before it begins, in hours, euro, and risk, against the method of Chapter 4. A case that cannot survive honest pricing at the gate should not enter the pilot.

The third signal is the **stalled champion**: the one person who believed is promoted, rotated, or poached, and the transformation discovers that it was a person, not a capability. When coherence lives in an individual, the organisation has not changed; it has borrowed someone. The counter-move is to make the capability plural early — a second trained team before the first wave closes, roles documented rather than remembered, and training embedded in the operating rhythm rather than in the champion's diary. If your pilot cannot survive its sponsor's promotion, it was never an organisational capability.

The fourth signal is **skipped training**: under quarter-end pressure, the progressive exposure of Chapter 10 is compressed into a week of slides, and people are asked to judge outputs they were never taught to read. Judgement unaccompanied by training is not oversight; it is a rubber stamp with a person's name on it, and the first error that slips through will be charged to the technology rather than to the skipped step. The counter-move is to treat the four phases as a gate, not a guideline: no one signs a judgement who has not passed through observation and review, and the Training Lead's calendar is protected by the COO with the same firmness as the pilot's budget.

The fifth signal is **automation theatre**: the agent demonstrably does the work, and the organisation then adds the work back — a second check beside the first, a summary of the summary, a meeting to discuss what the dashboard already says. The old effort persists next to the new, the promised capacity never materialises, and the CFO concludes, reasonably, that the whole thing costs money instead of making it. The counter-move is in the job letter and in the measurement: the hours an agent takes over must be visibly redeployed, the rules the agent executes must be visibly retired from human checklists, and the Value Analyst reports both. Automation that does not remove work somewhere is decoration.

The sixth signal is **gravity**: the coherence heatmap cools, attention moves to the next initiative, and within two quarters the old coordination patterns reassert themselves — the meetings creep back, the interfaces re-multiply, the escalation paths regrow like hedges. Coherence is not a state you reach; it is a discipline you hold, and it decays the moment it stops being measured. The counter-move is rhythm: the heatmap re-run on a fixed cadence, the Gate 5 audits unannounced enough to be honest, and one named executive — the Chief Coherence Officer by whatever title — whose standing agenda item is the temperature of the whole.

The seventh and most dangerous signal is the **false negative**: an agent that was genuinely capable is deployed into an incoherent environment, fails against the confusion, and the organisation concludes that "AI doesn’t work here." The verdict is wrong — it was the environment that failed, not the agent — but the verdict sticks, and it costs the organisation years. This is the signal the entire book exists to prevent. The counter-move is the argument of Chapter 3 made operational: diagnose before you buy, fix the workflow before you amplify it, and let no agent enter a process whose heatmap is red. The false negative is only possible when amplification is allowed to precede coherence.

Read together, the seven signals share a single property: none of them is a technology failure. Every one is a governance failure — a right not transferred, a value not priced, a person not trained, a measurement not kept. That is the chapter’s real message. The discipline that prevents them is not more sophisticated tooling but the mundane executive craft of gates, owners, rhythms, and honest numbers — the same craft this

book has asked of you since page one. Failures of this kind are optional. They happen only where the discipline was agreed in principle and waived in practice.

REFLECTION FOR THE EXECUTIVE TEAM

Read the seven signals a second time, more slowly, and this time keep score privately: how many are visible in your organisation right now, not as risks but as facts? Most executive teams find two or three, and almost always include the hidden veto among them — because the veto works by remaining deniable, and deniability is a skill senior people practise for decades. The useful move is not to hunt the vetoer but to make the veto expensive: decision rights transferred in writing, cases instrumented, starvation visible as a number on a page. Treat this chapter as a checklist for your next quarterly review rather than as literature. For each signal, ask the same three questions in order: is it present, who would know first, and what would make them tell us? The third question matters most, because every one of these signals is visible to someone long before it is visible to you — and the distance between those two moments is a precise measure of your information flow. A organisation that hears bad news early has one problem. An organisation that hears it late has two.

CHAPTER 12 — THE FIRST FIFTEEN WORKING DAYS

Every argument in this book converges on a calendar. Fifteen working days is short enough that no executive can plead calendar poverty and long enough to produce what no amount of further reading will: a heatmap of your own organisation, a priced shortlist, and a board-ready decision package. This closing chapter walks the calendar as a narrative — what happens each week, who is in the room, and what must exist on paper before the next week may begin. The picture below is the same calendar at a glance; the prose is the version you can act from.



Figure 26: The fifteen-working-day calendar — three weeks from commitment to a board-ready decision.

Days one and two belong to the sponsor. Before any diagnostic work begins, the chief executive does two things that cannot be delegated. The first is to name the scope: one value stream or business unit where the pain is felt, the data exists, and the political cost of honesty is survivable — the selection logic of Chapter 4 applied to the diagnostic itself. The second is to open the gate in writing: a short mandate that names the

Diagnostic Lead, grants access to the people and systems the diagnosis will touch, and states in one sentence that the outcome may be "do nothing here," so that the team is not implicitly commissioned to find a case whether or not one exists. An executive who cannot write that sentence is not ready for the diagnostic; better to discover that on day one than on day fourteen.

Days three through seven are the diagnostic itself, executed exactly as Chapter 4 describes. The Diagnostic Lead interviews the owner of the chosen scope and the people who actually move its work; the five domains are walked in order — direction, decision rights, information flow, process design, and amplification readiness — and each is scored against the three signals you met in the introduction: decisions that keep waiting, work that keeps changing hands, priorities that keep competing. By the end of day five the heatmap exists, and it is almost never the picture the leadership team expected. Days six and seven are spent on the two readings that turn the heatmap from a picture into an argument: where the red cells concentrate — concentration is where the money is — and what the green cells imply about where amplification would be safe today. The Value Analyst begins the pricing in parallel, using the time-and-motion baseline, so that no red cell is discussed without a number attached.

Days eight and nine produce the shortlist. The heatmap's red cells are candidate coherence projects; its green-adjacent workflows are candidate pilots. The team reduces both to a short list — rarely more than three of each — and prices every entry: what incoherence costs there today, per the method of Chapter 3, and what a pilot there could plausibly return, with the ratio honestly stated. This is where the discipline of Chapter 9 enters early: an entry that cannot be priced does not make the list, because an unpriceable case on day nine becomes an indefensible pilot on day ninety.

Days ten and eleven are the rehearsal. Before anything reaches the board, the package is defended in front of the people who will have to live with it: the COO reads it for operational truth, the CIO for architectural truth, the CFO for financial truth, and the works council or employee representatives for human truth. This is not courtesy; it is quality control. Every objection surfaced in the rehearsal is an objection the board will not be able to surprise you with, and every amendment accepted here makes the package more yours and less the consultant's.

Day twelve is the board package itself — short by design. The heatmap on one page. The priced shortlist on one page. The recommended first move — almost always a coherence fix in a red domain, occasionally a pilot in a green one — with its gate design: the owner named, the decision rights that will transfer on day one, the kill criteria stated before anyone is attached to the outcome, and the ratio and horizon stated the way Chapter 9 insists they be stated. Everything else is appendix. A board that needs more than three pages to decide will not decide.

Day thirteen is the decision, and it is binary. The board either opens Gate 1 for the recommended case — sponsor named, owner named, gate calendar set — or it closes the matter with reasons minuted. What it must not do is the third thing boards love: commission further study. Further study is how the three-hour decision of the introduction becomes the three-week decision and then the three-month one; the diagnostic you have just completed *is* the study, and its shelf life is short, because the organisation it describes does not stand still.

Days fourteen and fifteen close the loop. The decision — whichever way it went — is communicated to the people who were interviewed, in the sponsor's own voice, because a diagnostic that disappears into a boardroom teaches the organisation that honesty is harvested and never returned. The heatmap goes onto the standing agenda of the executive team, with a re-run date already fixed; the Diagnostic Lead's playbook is written down while it is fresh; and the date of Gate 1 for the approved case is in the calendar of everyone this book has given a role to.

What you hold on day fifteen is not a transformation. It is something more valuable and more rare: an honest picture of your own coherence, a priced set of choices, a decision actually taken, and a calendar that has already begun. Everything else in this book — the gates, the roles, the training, the measurement — now has somewhere to land. The method scales from here exactly the way the failure chapter warns it must: never by enthusiasm, always by rhythm.

And that is the whole argument. Organisations do not fail at amplification because amplification is hard. They fail because they amplify before they cohere. Diagnose first. Fix the workflow before the technology touches it. Put a trained human in the loop where judgement is irreplaceable. Price the case before you build it, and audit the result after you bank it. Give each role its page, each gate its authority, each quarter its

heatmap. None of this is heroic, and that is the point: coherence is not a vision to be proclaimed but a property to be built — cell by cell, gate by gate, decision by decision — until the enterprise can finally hear itself think. Only then does it deserve to be amplified. Design for coherence. Then amplify.

REFLECTION FOR THE EXECUTIVE TEAM

Fifteen working days is three weeks, and three weeks is the length of time your organisation routinely spends scheduling a meeting about scheduling a meeting. Hold the calendar up against your own diaries and ask what, honestly, stands between today and day one. The answers will be familiar — a budget cycle, a reorg in progress, a leadership transition, the sense that the timing is not quite right — and every one of them deserves the same test: is this a reason, or is it gravity dressed as a reason? There is one legitimate cause for delay, and it is worth naming so the others lose their cover: if the executive team cannot currently agree that incoherence is worth pricing, run the thirty-second version first — the three signals, asked in the open — and let that conversation decide. But if the agreement exists, then each week of postponement is not caution; it is the cost of incoherence, still accruing, still at three to eight percent, now with a number you have read. The calendar fits in any quarter. The question this book leaves you with is simpler than any in it: which Monday?

APPENDIX A — THE NINE ROLES, IN BRIEF

The transformation this book describes is carried by nine roles. Chapter 8 presented them as a design; this appendix compresses each into the essentials an executive would want before putting a name against one.

The Chief Coherence Officer

Owns the whole. This is the executive — often the COO or a direct delegate of the CEO — in whose name the heatmap is re-run, the gates are chaired, and the word "no" is spoken when amplification is proposed ahead of coherence. The role's instruments are the diagnostic and the gate calendar; its success measure is the temperature of the heatmap over time, not the number of pilots launched.

The Diagnostic Lead

Owns the truth at the start. This person runs the fifteen-day diagnostic, interviews without flinching, and delivers a heatmap the leadership team would rather not see — and can defend every cell of it in the room. Analytical independence matters more here than seniority; the worst diagnostic leads are the ones who know in advance what the board wants to hear.

The Process Owner

Owns one workflow end to end and signs for its coherence. Before any agent enters a process, this is the person who can state its decision rights, its handoffs, and its information flows as they actually are — not as the procedure manual claims. The role predates this book and will outlast it; amplification simply makes its absence unaffordable.

The Decision Rights Architect

Owns the map of who decides what. Working with the Process Owner, this role drafts the transfer of rights that Gate 1 signs into effect, so that an agent's authority — and the human's retained judgement — are written down before the first case runs. Where the hidden veto of Chapter 11 is the risk, this role is the structural answer.

The Human-in-the-Loop Designer

The Human-in-the-Loop Designer owns the placement of judgement. From the three positions of Chapter 6, this role chooses where the human sits for each agent deployment — before, beside, or above — and writes the judgement criteria the reviewer will actually use. The role's failure mode is the rubber stamp; its success is a reviewer who can say, precisely, why a borderline case was returned.

The Training Lead

The Training Lead owns progressive exposure. This role builds and protects the four-phase calendar of Chapter 10, certifies that no one signs judgements untrained, and treats the training time of experts as capital expenditure rather than overhead. The role reports on capability, not attendance.

The Value Analyst

The Value Analyst owns the numbers at both ends. Before the pilot, this role prices the case by the method of Chapters 3 and 9; after it, the Value Analyst verifies the banked result against the promised one, including the honest treatment of capacity that was freed but not redeployed. The role's loyalty is to the counterfactual.

The Gatekeeper

The Gatekeeper owns the integrity of the five gates. Distinct from the executives who sit at the gates, this role administers them: evidence complete before the meeting, decisions minuted, kill criteria recorded before anyone is attached to the outcome. A gate without a gatekeeper becomes a ceremony.

The Communications Lead

The Communications Lead owns the story the organisation tells itself. This role closes the loop after every diagnostic and every gate — what was decided, why, and what happens to the people whose work changes — because the silence of Chapter 12's day fourteen, once skipped, is never forgiven by the floor.

APPENDIX B — GLOSSARY

Amplification readiness

The degree to which a workflow's rules, data, and decision rights are explicit enough that an agent amplifying them produces value rather than faster confusion. The fifth domain of the diagnostic.

Coherence

The property of an organisation in which direction, decision rights, information flow, and process design agree with one another, so that the work behaves the way the strategy says it should.

Coherence Diagnostic

The fifteen-working-day instrument of Chapter 4: five domains scored against three signals, producing a heatmap, a priced shortlist, and a board-ready decision package.

Counterfactual

What would have happened to the measured value had the intervention not occurred. The Value Analyst's standard of proof; a pilot without one has an anecdote, not a result.

Decision rights

Decision rights. The explicit allocation of authority to decide — who may decide what, at what threshold, with whose countersignature. The second domain of the diagnostic and the ground on which the hidden veto is fought.

Five gates

Five gates. The staged commitment structure of Chapter 7: mandate, design, pilot, scale, audit — opened in turn by the CEO, CIO, CEO, COO, and CFO. No gate opens on enthusiasm; each opens on evidence specified in advance.

Gravity

Gravity. The tendency of a cohered organisation to decay back toward its old coordination patterns once attention moves on. Countered by rhythm, not by effort: a re-run heatmap and unannounced audits.

Heatmap

Heatmap. The five-by-three picture of the enterprise's coherence — five domains, three signals — in which red cells mark where incoherence is costing money and green cells mark where amplification would be safe.

Hidden veto

Hidden veto. The passive obstruction of a change by a party whose decision rights it challenges — never refused, only starved. The first and most common of the seven failure signals.

Human-in-the-loop positions

Human-in-the-loop positions. The three placements of human judgement relative to an agent's work: before (approval of what may run), beside (review of what is running), and above (oversight of what has run). Chosen per deployment, never by default.

Incoherence statement

Incoherence statement. The sentence a coherent diagnosis makes writable: *we lose this much, here, because these domains disagree.* If the sentence cannot be written, the diagnosis is not finished.

Kill criteria

Kill criteria. The conditions, written before the pilot begins, under which it will be stopped. Their purpose is to make stopping cheap precisely when attachment makes it expensive.

Mispriced value

Mispriced value. Value claimed without a priced counterfactual — the pilot everyone liked and nobody costed. Discovered, always, by the CFO, at the worst moment.

Progressive exposure

Progressive exposure. The four-phase training discipline of Chapter 10: observe, review, judge with support, judge alone. No one signs a judgement who has not passed through all four.

Ratio

Ratio. Promised value divided by pilot cost, honestly priced. The board's single number at Gate 1, stated together with its horizon.

Three signals

Three signals. The observable symptoms of incoherence: decisions that keep waiting, work that keeps changing hands, priorities that keep competing. What the diagnostic scores and what the floor already knows.

Value horizons

Value horizons. The honest treatment of time in the business case: what returns this quarter, what returns this year, and what returns only across the transformation — indexed, in the book's running example, as 1.0, 2.3, and 4.0.

APPENDIX C — THE INSTRUMENTS AT A GLANCE

The book leaves you four working instruments. The **three-signal test** — waiting decisions, wandering work, competing priorities — is the thirty-second version, usable in any meeting, needing no consultant. The **coherence heatmap** is the fifteen-day version: five domains against three signals, scored with the people who do the work, producing the picture the board must see. The **five gates** are the commitment discipline: mandate, design, pilot, scale, audit, each with its named opener, its evidence standard, and its kill criteria written in advance. And the **role set** — nine roles, each with an owner, an instrument, and a success measure — is what turns the other three from a project into a property of the organisation. Use the first this week. Schedule the second within the month. The third and fourth are what make sure you never have to start over.



THE COHERENT ENTERPRISE

Why complexity is quietly breaking modern organisations — and how to move with coherence.

A field guide for executives: diagnose coherence before you amplify, price what incoherence costs, place human judgement where it matters — and move from first signal to board decision in fifteen working days.

ERLEND HOLLEBOSCH

GROW-FASTER.BE