

AURA

The Agentive Path

César Obach-Renner

First edition v1.1 · August 2026

This page intentionally left blank.

Contents

The Trilogy	1
Preface	3
How is this book organized?	3
What this book is — and is not	4
Introduction	5
The Concept: Agentic AI	5
The Industry’s Stagnation	6
The Leap	7
Three Axes of Deep Change	10
AURA: The Map of the Path	12
Main Sources	15
Knowing	17
IRIS — Organizational Intelligence Maturity Model	21
Introduction · The Thesis	21
Structure of the Model	23
Stage 1 · Fragmented	24
Stage 2 · Systematized	24
Stage 3 · Centralized	25
Stage 4 · Modeled	25
Stage 5 · Accessible	26
Stage 6 · Contextual	27
Stage 7 · Proactive	28
Stage 8 · Bidirectional	28
Stage 9 · Coordinated	29
Stage 10 · Orchestrated	30
Analysis of the Leaps Between Stages	30
The 5 Levels · Simplified View	31
How to use this model?	32
Quick Reference: The 10 Stages	33
Cross-Reference	33
Data Canon	35

Abstract	35
Introduction	36
Background: Data Mesh and its rationale	37
The argument, part I: distributing is no longer necessary	38
The argument, part II: distributing is worse in governance terms	40
The Data Canon pattern	43
Implementation implications	50
Limits, risks, and falsifiability	51
Related work	52
Conclusion	53
References	54
Chapter glossary	54

Doing 57

MOTOR — Organizational Automation Maturity Model 61

Introduction · The Thesis	61
Structure of the Model	62
Level 1 · Invisible (Shadow AI)	63
Level 2 · Trusted	64
Level 3 · Collaborator	65
Level 4 · Analyst	66
Level 5 · Specialist	68
Level 6 · Manager	69
Level 7 · Authority	70
Digital Twins as a Diagnostic Dimension	72
Evaluation Dimensions by Level	73
Market Distribution	74
Rapid Assessment	74
How to use this model?	74
Quick Reference: The 7 Levels	75
Glossary	76
Cross-Reference	77

Wingmap 79

Abstract	79
Introduction	80
Conceptual framework	81
Definition and principles of the technique	83
Operational components	85
The practice in phases	86
Discussion	89
Limitations and zones of non-applicability	91
Conclusion	93
Future work	94
References	94

The Portfolio	97
Use Cases	101
Introduction · The Thesis	101
Use Cases and Value Solutions	101
The Three Classification Dimensions	102
The Citizen / City / DUAL Framework	102
The ROI Discipline	104
A Real Portfolio: One Hundred Cases for Local Government	104
How to use this framework?	105
Epilogue · The Complete Diagnosis	107
Annexes	109
Annex A · The Diagnostic Instrument	111
How is it applied?	111
The KNOWING axis — how does information flow?	112
The DOING axis — how much autonomy do the processes have?	114
How is the result read?	115
The baseline — the four measures	116
What follows the diagnosis?	117
Annex B · Use Case Inventory	119
The 30 Value Solutions, by ranking	119
The 100 Use Cases, by MOTOR level	123
Portfolio Statistics	126
The Four Thematic Areas	127

This page intentionally left blank.

The Trilogy

This book is volume **II** of the **Agentive World Trilogy**. The transition that occupies the trilogy is a single one: from the world where people open applications to work, to the world where AI agents are the interface of work. The three volumes answer, in order, the three questions of that transition — and each one reads complete on its own:

I · The Real-Time Enterprise — The Agentive World Where are we going? — **the destination**. For those who want to see the world on the other side of the transition: how people work, decide, and compete when the enterprise operates in real time — told in its four faces.

II · AURA — The Agentive Path (*this book*) Which way? — **the route**. For those who have to make the crossing: leaders and advisors of the transformation. Its five specifications are standalone instruments for the agentive era — each one can be used separately, without the rest of the book:

Spec	What is it?	Used alone, what does it give you?
IRIS	A maturity model of organizational intelligence	Locates your organization on the road from data to action
MOTOR	A maturity model of organizational automation	Measures how advanced your organization is at process automation
Data Canon	The evolution of the decentralized data-management model (Data Mesh)	Defines how the organization's data is governed
Wingmap	The evolution of business process management (BPM / process mining)	Reconstructs processes and information flows without interviews
Use Cases	A portfolio framework, illustrated with 100 use cases and 30 value solutions	Prioritizes where to start, according to your maturity

III · AgencyDomains — Agentive Architecture How do we get there? — **the vehicle**. For architects, CTOs, and builders of agentive platforms: the formal specification of the Agentive World.

Vergis — Agentive Technology (*not a book: it's code*) What do I start with today? — **the keys**. For anyone who wants to start making all of this real: the reference implementation of the architecture, open source, at <https://github.com/gegolabs/vergis>.

Start with **The Real-Time Enterprise** if you need to convince yourself — or someone else — that the transition is for real; with **AURA** if you have to lead it; with **AgencyDomains** if you have to build it; with **Vergis** if you want to see it running. The full trilogy lives at <https://agencydomains.org>.

This page intentionally left blank.

Preface

There are two questions every organization will ask itself this decade, and it pays not to confuse them.

The first is *what does the destination look like?* — what shape an organization takes when AI agents stop being tools that people open and become the medium through which work happens. That question has its own book in this series: *AgencyDomains* describes the formal architecture of that world, and *The Real-Time Enterprise* describes the world itself.

The second question is the one this book answers: **how do you get there?** Not what to build at the end, but where the organization stands today, what trajectory lies ahead of it, in what order it is best traveled, and with what instruments. The distance between an enterprise that accumulates dashboards and an enterprise that detects, decides, and acts autonomously is not crossed with a software purchase — it is traveled. And a path is better traveled with a map.

AURA — Agentive Unified Reference Architecture — is that map. Its organizing thesis is that the agentive transformation advances along **two axes that do not move together**: the **KNOWING** axis — the organization’s capacity to turn its data into actionable intelligence — and the **DOING** axis — its capacity to execute processes with increasing autonomy. An organization can know much and do little: impeccable dashboards, bottlenecked decisions. It can do much and know little: blind automation over data nobody governs. Full maturity demands advancing along both, and that is why AURA is operationalized in two complementary diagnostic models — **IRIS** for knowing, **MOTOR** for doing — plus the instruments each axis needs in order to advance.

How is this book organized?

The Introduction establishes the frame: what agentive AI is — and why we say *agentive* rather than *agentic* —, the Value Pyramid where the industry stalled, the Quantum Leap that collapses the cost of the analytical question, and the three axes of deep change the transformation demands.

Knowing (Part I) travels the first axis. IRIS models the trajectory of organizational intelligence in ten stages — from fragmented information to the self-managing ecosystem — with the frontier that separates the *online enterprise* from the *real-time enterprise*. Data Canon answers the governance question that journey imposes: who rules over the data when agents, not people, are the ones consuming it — and why the agentive era recentralizes the core of the governance that the past decade decentralized.

Doing (Part II) travels the second axis. MOTOR models organizational automation in seven levels — from *shadow AI* (the AI people already use on their own, ungoverned) to the agent as strategic authority. Wingmap supplies the discovery instrument: how to reconstruct the organization’s real processes — the ones that live in people’s heads and not in the documentation — without interviews, from the digital traces the operation already generates.

The Portfolio (Part III) brings the path down to earth: how a diagnosis is translated into a prioritized portfolio of use cases — the distinction between atomic cases and value solutions, the beneficiary-based prioritization framework (Citizen / City / DUAL), and the two canonical starting strategies, illustrated with a real portfolio of one hundred cases for local government.

The book closes with an **epilogue** that integrates the complete diagnosis and an **annex** with the full inventory of the reference portfolio's one hundred cases.

The chapters were born as independent specifications and retain that autonomy: each can be read on its own. Read in order, they compose the complete path — vision, diagnosis of knowing, diagnosis of doing, and prioritized execution.

What this book is — and is not

AURA is **diagnostic, not prescriptive**. Its models say where an organization stands and what standing there means; they do not say which product to buy or which project to execute — that belongs to the consulting work these models inform. And AURA is honest about the transition: the path is not to demolish what was built but to subsume it — the data warehouse does not die, it becomes a source that agents consume; the online enterprise does not disappear, it becomes the foundation of the real-time enterprise.

The reader looking for the architecture of the destination — the primitives, the layers, the *Trust Infrastructure* — should go to *AgencyDomains*. The one who wants to know what the world looks like once the transition is over, to *The Real-Time Enterprise*. This book is for those who have to make the crossing.

Introduction

AURA — Agentic Unified Reference Architecture

The Concept: Agentic AI

The technology industry is debating between two radically different futures for artificial intelligence in the enterprise.

In the **agentic** world, AI agents are integrated as complementary tools. Employees keep opening Excel, Salesforce, Power BI — now with copilots that help them work faster. Traditional interfaces persist; agents enhance them. It is incremental evolution.

In the **agentic** world, agents completely replace traditional interfaces. Employees stop opening applications — they interact directly with agents that execute tasks on their behalf. Applications may continue to exist as invisible backend infrastructure, but the interface has collapsed. It is fundamental transformation.

The frontier between the two worlds is what we call **the Nadella Line**: do your employees still open applications to do their work? If the answer is yes, you are in the agentic world. If the answer is no, you have crossed into the agentic world.

To grasp concretely what this means: a CFO who today opens Power BI builds margin dashboards and produces conclusions in two or three hours. In the agentic future, that same CFO simply asks “Why did margins drop in Q3?” and within minutes an agent analyzes data, generates visualizations, and identifies root causes. The critical point: the CFO never opened an application.

The industry’s most disruptive voices are betting on the agentic world. Satya Nadella (CEO of Microsoft) declared: *“The notion that business applications exist, that’s probably where they’ll all collapse, right, in the agent era.”* If the agentic future materializes, the implications run deep: graphical interfaces collapse, technical skills are redefined, technology investment migrates from SaaS licenses toward agent governance, and competitive advantage shifts from having better applications to having better agents and extraordinarily skilled employees directing them.

Why “agentic” and not “agentic”? Because they are not the same thing. The industry has adopted “agentic AI” to describe agents that assist — copilots, assistants, smarter tools. But this book points to something qualitatively different: a world where agents do not complement the way we work but transform it. Using “agentic” would limit the scope to incremental evolution. Using “agentic” points to fundamental transformation — which is the central thesis of this book.

The crossing indicator: the agentic percentage. The Nadella Line admits an operational indicator, and it is worth installing from the start because it turns the paradigm into a measurement any board

understands: *what percentage of the times an employee needs to get something done can they delegate it entirely to an agent, without opening an application?* Between 0 and 20 percent, the organization lives in the agentic world — agents are complementary tools. Between 20 and 50, it is approaching the line. Between 50 and 80, it is crossing it. Above 80, it inhabits the agentive world. The indicator complements this book's two models without replacing them: IRIS and MOTOR measure the internal **capabilities** (knowing and doing); the agentive percentage measures the organization's **position** with respect to the transition. And it does not move uniformly: structured-knowledge functions — finance, operations, customer service — push it up fast; regulated industries — health, banking, government — push it up slowly, by legal mandate before technical limit. Measuring it per function, not as a single average, is what makes it actionable.

The transformation this book traverses is set on the agentive horizon. To navigate that transformation, ultraBASE has developed **AURA** — an agentive enterprise architecture that organizes the transformation along two complementary maturity axes:

IRIS (Organizational Intelligence Maturity Model) measures the **KNOWING** axis — an organization's trajectory from fragmented data to a self-managing intelligence ecosystem. From the online enterprise — which accesses up-to-date information but depends on humans to act — to the real-time enterprise — where intelligent systems detect, interpret, decide, and act autonomously.

MOTOR (Organizational Automation Maturity Model) measures the **DOING** axis — the trajectory from processes executed by humans to autonomous execution orchestrated through AI agents.

Both models are **diagnostic**: they assess what level an organization is at and what that level means. They do not prescribe what to implement — that is the province of specialized consulting. They are orthogonal: an organization can be high on one and low on the other. Full organizational maturity requires advancing along both axes. AURA integrates both axes into a reference architecture that makes it possible to diagnose the complete state of the transformation.

The world on the other side of this transition is portrayed in *The Real-Time Enterprise*, and its formal architecture specified in *AgencyDomains* — volumes I and III of this trilogy.

The Industry's Stagnation

The Value Pyramid

Every organization accumulates data. Sales transactions, inventory movements, customer interactions, operational metrics — enterprise systems generate records continuously. But raw data, on its own, has no executive value.

Value emerges through a progression we call **the Data Value Pyramid**:

Organizations have invested decades and fortunes in the two lower layers: transactional systems, data warehouses, ETL pipelines, visualization tools. But the leap from information to knowledge remains painfully slow — and the leap from knowledge to action is practically nonexistent.

Why? Not because the technology is slow, but because **there are humans in the middle**. When an executive needs to understand something about the business — why sales dropped, where the operational inefficiency lies, which customers are at risk — a process begins that rarely takes less than a week: calendar coordination, requirements gathering, technical development, iterative validation. The real bottleneck is the transfer of knowledge between people, not the capacity of the systems.




Figure 1: The Data Value Pyramid — and the two leaps where the industry stalled

This diagnosis is the starting point of the transformation we measure.

From information management to organizational intelligence

Well-established information management maturity models exist in the industry — Gartner, CMMI, DAMA-DMBOK. They measure critical aspects: data quality, governance, architecture, integration. But they concentrate on the first two layers of the Value Pyramid: DATA and INFORMATION.

This book spans all four layers, from fragmented data to orchestrated autonomous action. That is no longer information management — it is an organization's capacity to transform data into actionable intelligence. That is why we speak of **organizational intelligence**: it is not about how well you manage your data, but about how capable your organization is of generating knowledge and acting on it continuously and autonomously. Information management is the necessary foundation — but not the destination.

The Leap

The new economics of information: from weeks to seconds

The most immediate and tangible change enabled by agentive AI is the **elimination of the marginal cost of an analytical question**. We call that collapse — from weeks to seconds — **the Quantum Leap**, and it will be a reference throughout the book. In the traditional model, every new business question entails a project: coordination, requirements gathering, development, validation. The cost — in time, in

person-hours, in opportunity cost — leads organizations to *ration* their questions. Only what someone had the foresight to pre-build into a dashboard or report ever gets analyzed.

When that cost collapses from weeks to seconds, the very nature of the relationship between an organization and its information changes:

- **Analytical capacity becomes elastic.** It adapts in real time to the current need, not to whatever someone managed to pre-define months ago. There is no fixed menu — there is unlimited responsiveness.
- **Iteration replaces specification.** Instead of defining requirements up front and waiting for the result, the executive explores, refines, drills deeper in a continuous conversation with the information. Knowledge emerges from dialogue, not from a project.
- **The questions that were never asked now get asked.** When asking is free, the organization discovers insights it didn't even know it needed. Analytical curiosity stops being constrained by the BI budget.

This transformation in the economics of information is an **enabling condition** for everything else in this book: the continuous-intelligence cycle cannot exist if every iteration takes weeks. Autonomy governance makes no sense if agents do not operate in real time. The transformation of roles does not happen if access to knowledge still depends on human intermediaries.

The inversion of the flow

For 30 years, the information management paradigm was “**people go to the data**”: you build a warehouse, mount dashboards, train users, and hope someone looks at the right report at the right moment and makes the right decision. The entire model rests on *human attention* as the bottleneck.

As Tellius describes it: “*Dashboards still tell you what happened, but rarely why — and never what to do next.*” Superwise puts it bluntly: the industry is moving “beyond dashboards” because “*BI was built for a slower business environment — that assumption no longer holds true.*”

Agentive AI inverts that flow: “**intelligence goes to the people — and acts on their behalf.**” A system of agents monitors, interprets, decides, and executes within the limits the organization has defined, and escalates to a human only when warranted. This inversion marks the frontier between two fundamentally different kinds of organization — the **online enterprise** and the **real-time enterprise** — the same distinction fixed further below on the continuous-intelligence cycle (Perceive → Interpret → Decide → Act → Learn).

Tableau articulates this as the move from traditional BI to “agentic analytics,” where AI doesn't just visualize data but activates it. Cube declares it openly: “*The modern data stack is beginning to show its age.*” BCG takes it to the operational plane, describing how agentive AI orchestrates actions across the entire value chain, “closing the loop between insight and execution.”

Agentive AI is not an incremental improvement to the analytics stack — it is a change of operating model. Measuring maturity in this context requires assessing not only the technology, but the organization's capacity to operate in a fundamentally different paradigm.

Paradigm shifts

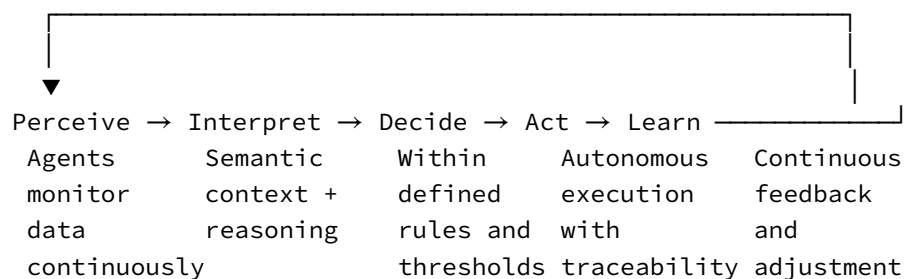
Before	After
People go to the data	Intelligence goes to the people and acts on their behalf
Dashboards as the end product	Agents as the operating system of decisions
Human attention as the bottleneck	Human governance as the control system
Analytics as a support function	Continuous intelligence as a business capability

From the classic cycle to the continuous-intelligence cycle

Classic cycle · Linear · Human-dependent

Descriptive → Diagnostic → Predictive → Prescriptive → [!] Human decides
 (What happened?) (Why?) (What will happen?) (What to do?) (bottleneck)

New cycle · Continuous · Agent-executed · Human-governed



The critical change: the step from “Prescriptive” to action is no longer a recommendation a human evaluates. It is a decision an agent executes, whose outcome it monitors and adjusts. The organization stops “doing analytics” and starts *being* an intelligent system.

What the industry is saying:

Tableau describes how agentic AI “upends the cycle of visual analysis” — transforming a linear, human-dependent process into a cycle where agents participate actively at every stage. Tellius formalizes it as the move “from questions to autonomous action.”

Ampcome documents 11 enterprise use cases where analytical agents are already transforming decisions in 2026. CDO Trends describes how semantic layers are powering this new era of AI-driven decision-making.

Gartner had already introduced the concept of “continuous intelligence” as one of its strategic trends, defining it as real-time analytics integrated directly into business operations. Agentic AI takes that vision to its logical conclusion: not just continuous analytics, but **continuous action**.

Context data

- **40%** of enterprise apps will include AI agents by the end of 2026, vs. <5% in 2025 (Gartner)
- **>40%** of agentic projects will be canceled before the end of 2027 — due to costs, unclear business value, or inadequate risk controls (Gartner)
- **80%** of organizations report risky behaviors from their agents; only 21% have mature governance (MindStudio)

- **75%** of current roles will require redesign by 2030 (McKinsey)
- **58%** of heavy adopters expect fundamental changes in governance within 3 years (BCG)
- **48%** cite data discoverability among the main barriers to their AI automation strategy (Deloitte)
- Only **14%** have production-ready agentive solutions; 11% in real-world use (Deloitte)
- Agentive AI market: **\$7.3B** in 2025 → projected **\$139.2B** by 2034 (CAGR $\approx$ 39%; aggregated market projections)

Three Axes of Deep Change

Axis 01 · From consuming information to governing agents

Current paradigm:

- Analysts build reports
- Users consume dashboards
- Key competency: data literacy
- Value lies in understanding the information

Emerging paradigm:

- People design rules for agents
- Agents monitor, interpret, and act
- Key competency: designing and supervising autonomous systems
- Value lies in governing intelligent action

“The CFO doesn’t review a cashflow dashboard — they define the thresholds and protocols that a financial agent executes autonomously.”

What the consulting firms are saying:

McKinsey describes the “Agentic Organization,” where employees shift from executing tasks to orchestrating outcomes, supervising agents, setting objectives, and managing trade-offs. Humans move “above the loop.” McKinsey estimates that 75% of current roles will require redesign, upskilling, or redeployment by 2030.

BCG notes that 45% of AI leaders expect to need fewer layers of middle management, as work centers on supervision, orchestration, and continuous improvement of agentive processes — not on manual execution.

Deloitte frames the new human role around judgment and orchestration governance: *“human ingenuity and judgment will continue to play a key role in selecting and guiding agent outputs.”*

A CIO.com article details the emerging new roles: AI Agent Orchestrator, Human-Agent Interaction Designer, AI Ethics & Governance Specialist, and AgentOps Specialist. It is an entirely new org chart.

Axis 02 · From architecture for humans to architecture for agents

Current paradigm:

- Data warehouses optimized for SQL
- Data quality = cleanliness
- Data models for visualization
- Batch and on-demand integration

Emerging paradigm:

- Semantic layers where agents reason
- Data quality = actionability
- Knowledge graphs for autonomous inference
- Real-time, contextual data flows

“If an agent cannot reliably interpret and act on a piece of data, that data is useless — no matter how ‘clean’ it is.”

What the industry is saying:

AtScale demonstrated that agents without a semantic layer fail on 80%+ of queries, while with a semantic layer they reach ~100% accuracy. Its conclusion: *“For AI agents, the semantic layer isn’t just a nice-to-have — it’s the foundation that makes AI actually useful.”*

ThoughtSpot coined the concept of the “Agentic Semantic Layer” — dynamic, context-aware, and designed to work hand in hand with agents. Salesforce proposes an Enterprise Knowledge Graph (EKG) as the central layer of its agentive enterprise architecture. Databricks speaks of unifying infrastructure, data, and semantics to enable “Agentic BI.”

Informatica warns that giving an agent access to data is not enough: *“Because agents act without human approval loops, the data they use must be fully trusted, verified, and monitored.”* It proposes data-quality SLAs: less than 5 minutes of freshness for transactional agents, less than 1 hour for analytical agents — a direction Data Canon (Part I) will push further: a single quality level, the strictest, for the entire canon.

eWeek introduces the concept of “agent-ready data” and projects that by 2026, enterprise data management will be radically different — with autonomous agents cleaning, validating, and enriching data within governed limits.

Deloitte found that 48% of organizations cite data discoverability and 47% cite reusability as main barriers to their AI automation strategy. It proposes a paradigm shift: from traditional ETL pipelines to enterprise search and indexing, contextualizing data through knowledge graphs.

Axis 03 · From access governance to autonomy governance**Current paradigm:**

- Who can see which data
- Static roles and permissions
- Compliance as a checklist
- Access auditing

Emerging paradigm:

- What an agent can do, under what conditions
- Dynamic levels of human supervision
- Governance as an enabler of scale
- Auditing of autonomous decisions

“Organizations that fail to solve autonomy governance will not be able to scale agents beyond isolated pilots.”

What regulators, consulting firms, and industry are saying:

Singapore IMDA published in January 2026 the first state-issued governance framework for agentic AI (MGF), establishing that although agents act autonomously, “*human responsibility continues to apply*” — organizations must make human accountability meaningful and keep the human-in-the-loop effective over time.

The World Economic Forum proposes progressive governance: logging and traceability for all agents, identity tagging per action, and real-time monitoring. Its key distinction: “*autonomy entails decision-making flexibility; automation emphasizes execution reliability*” — these are design choices, not inherent properties of the system.

BCG reports that 58% of heavy AI adopters expect a fundamental change in their governance structure within the next 3 years, and a third believe AI will hold more decision authority over the same period.

NACD (National Association of Corporate Directors) warns that agentic AI impacts board oversight, regulatory compliance, and risk exposure — it is no longer just an IT matter.

KPMG sums it up: “*The winners won’t be the ones with the most pilots but the ones investing now in scalable data architectures, agent governance models, and workforce readiness.*”

Gartner predicts that more than 40% of agentic AI projects will be canceled before the end of 2027 — due to costs, unclear business value, or inadequate risk controls. Governance is not optional; it is what separates pilots from production.

The scale of the problem is significant: according to MindStudio, 80% of organizations report risky behaviors from their agents (unauthorized data access, unexpected interactions), and only 21% have mature governance models. ISACA highlights that agentic AI poses a growing challenge for audit functions because its decision processes lack clear traceability.

AURA: The Map of the Path

Online enterprise vs. real-time enterprise

The preceding sections laid out the distinction; it is time to fix it as the axis of the trajectory: the **online enterprise** *accesses* — up-to-date data, current dashboards — but depends on a human to look, interpret, and decide; the **real-time enterprise** *reacts*: it runs the full continuous-intelligence cycle, waiting for no one.

IRIS maps this trajectory: stages 1–4 build the **foundations** — the infrastructure of the online enterprise; stages 5–10 build the **real-time enterprise**. **The Quantum Leap** — the collapse of the cost of the analytical question — is the frontier between the two trajectories.

And the Nadella Line maps naturally onto this trajectory: the foundations (stages 1–4) are inherently agentic — they are built with tools and copilots. From Stage 5 onward, the organization becomes progressively agentic — agents detect, interpret, and act autonomously.

The convergence

Four concepts converge in AURA:

Organizational Intelligence is the capability being measured — the *what*. (IRIS measures this axis: KNOWING.) **Organizational Automation** is the capability to execute — the *who*. (MOTOR measures

this axis: DOING.) **Real-Time Enterprise** is the organizational outcome of both capabilities at maturity — the *what for*. **Agentive AI** is the enabling mechanism that makes the transformation possible — the *how*.

AURA is the architecture that unifies these four concepts into an actionable reference framework.

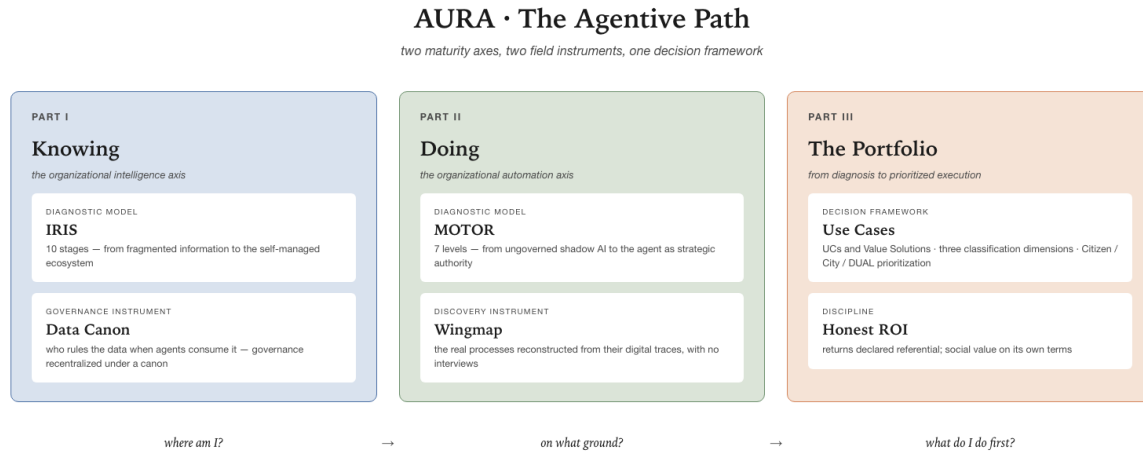


Figure 2: AURA, the map of the path — two maturity axes, two field instruments, one decision framework

Vision statement

Organizations are moving from the online enterprise — where people use tools to access and understand data — to the real-time enterprise — where people design and govern intelligent systems that autonomously detect, interpret, decide, and act on business information. AURA organizes this transformation along two diagnostic axes: IRIS measures the KNOWING trajectory (organizational intelligence); MOTOR measures the DOING trajectory (organizational automation). Together, within AURA, they diagnose the complete transformation.

AURA integrates into a unified architecture what multiple industry players are proposing separately: Tableau and Cube on agentive analytics, Salesforce and AtScale on semantic architecture for agents, McKinsey and BCG on the transformation of human roles, Gartner and WEF on autonomy governance, and Deloitte and Informatica on data readiness for agents.

What sets AURA apart is that it is not confined to a single layer (technology, people, or governance) but integrates them as an interdependent transformation system — recognizing that advancing in one dimension without the others produces successful pilots but no real transformation. And it is opera-

tionalized in two complementary diagnostic models: IRIS for the knowledge axis and MOTOR for the execution axis.

The nature of the transition: evolutionary coexistence

It is important to be honest about how this transformation happens in practice. It is not “throw everything out and start from scratch.” Organizations do not jump from the current paradigm to the emerging one overnight.

The reality is a **hybrid model** in which existing infrastructure (data warehouses, ETL pipelines, BI tools) coexists with the new agentive capabilities. Traditional architecture remains valid for large-scale historical data warehousing, highly complex analytical models that require pre-computation, pipelines with very specific business logic, and regulatory requirements for retention and formal lineage.

What changes is the proportion. In the early stages, 90% is static infrastructure and the agent is an occasional assistant. In the advanced stages, the agent orchestrates most analytical operations and the static infrastructure becomes one more data source — robust and governed, but consumed by agents, not by humans.

This evolutionary coexistence is a fundamental principle of AURA and has a direct implication for IRIS: **each level does not invalidate the previous one — it subsumes it.** An organization at IRIS stages 7–8 did not eliminate its data warehouse — it integrated it into a semantic layer — a *fabric* — that agents consume autonomously. The online enterprise does not disappear when the real-time enterprise emerges — it becomes its foundation.

Dimensions of AURA

Dimension	Transformation trajectory
Data & Architecture	Data warehouse → semantic layer → agent-actionable knowledge fabric
Analytical Capabilities	Static BI → augmented analytics → autonomous intelligence → continuous cycle
People & Culture	Information consumers → designers of intelligent systems → governors of autonomy
Governance	Access control → governance of autonomous decisions → auditing of agent actions
Operating Model	Data competency centers → agent factories → ecosystem orchestration
Business Value	Insights for deciding → autonomous actions under supervision → self-optimizing business

Each dimension progresses through the Value Pyramid: from organizations focused on the Data layer (accumulating and structuring), through Information (contextualizing and visualizing), then Knowledge (analyzing and understanding in real time), to Action (executing autonomously within governed frameworks). This progression marks the path from the foundations (online enterprise) to the real-time enterprise — the trajectory that IRIS operationalizes in its 10 stages, complemented by MOTOR, which measures the degree of automation of the processes that execute that action.

Main Sources

Source	Topic	Link
Gartner	Predictions on agents in enterprise apps	gartner.com
Gartner	Cancellation of agentive projects	gartner.com
McKinsey	Agentic organization and role redesign	mckinsey.com
BCG	AI at Work 2025: governance and roles	bcg.com
BCG	Agents across the industrial value chain	bcg.com
Deloitte	Agentive AI strategy and data readiness	deloitte.com
Deloitte	Agent orchestration	deloitte.com
KPMG	Governance in the agentive era	kpmg.com
Tableau	Agentic Analytics	tableau.com
Tableau	AI and the cycle of visual analysis	tableau.com
Salesforce	Agentive enterprise architecture	architect.salesforce.com
AtScale	Semantic layer for AI agents	atscale.com
ThoughtSpot	Agentic Semantic Layer	thoughtspot.com
Databricks	Agentic BI: infrastructure + semantics	databricks.com
Informatica	Enterprise AI Agent Engineering	informatica.com
eWeek	Agent-ready data and enterprise management	eweek.com
Singapore IMDA	Governance framework for agentive AI	imda.gov.sg
World Economic Forum	AI agents and progressive governance	weforum.org
NACD	Board-level oversight of autonomous AI	nacdonline.org
ISACA	The challenge of auditing agentive AI	isaca.org
CIO.com	A new org chart for the agentive era	cio.com
Cube	Agentic Analytics as the new modern analytics	cube.dev
Tellius	From questions to autonomous action	tellius.com
Superwise	Beyond dashboards	superwise.ai
Ampcome	Use cases for analytical agents	ampcome.com
CDO Trends	Semantic layers and AI-driven decisions	cdotrends.com

Introduction

Source	Topic	Link
MindStudio	Enterprise AI agent governance	mindstudio.ai

Knowing

The first axis of the path is **organizational intelligence**: the capacity to turn data into actionable knowledge, continuously and without intermediaries.

This part travels it with two pieces. **IRIS** is the diagnostic model: ten stages running from information fragmented in silos to the self-managing intelligence ecosystem, with a sharp frontier in the middle — the point where the organization stops being an *online enterprise* (information up to date, humans deciding) and starts being a *real-time enterprise* (systems that detect, interpret, decide, and act). **Data Canon** is the governance piece that this advance demands: when the consumer of data stops being a human with a dashboard and becomes an agent that acts, the question of who defines and who controls each figure stops being a technical matter and becomes a matter of corporate governance.

The order matters: first the map of the trajectory, then the rule of governance that makes it passable.

This page intentionally left blank.

IRIS — Organizational Intelligence Maturity Model

In the era of agentic AI

Introduction · The Thesis

With the general map drawn in the Introduction — the inversion of the flow (“intelligence goes to the people”), the three axes of change, and the Quantum Leap — this chapter descends to the first axis of the path: how organizational intelligence is measured, stage by stage. IRIS models that complete trajectory, from fragmented information to the self-managing ecosystem.

Why “organizational intelligence” and not “information management”?

The term that names the axis was established in the Introduction: organizational intelligence is what gets measured when the scope covers all four layers of the Value Pyramid — not just managing data, but generating knowledge and acting on it continuously and autonomously. IRIS operationalizes that capacity, stage by stage.

Online Enterprise vs. Real-Time Enterprise

Recalling the distinction the Introduction established — the online enterprise accesses; the real-time enterprise acts — the model operationalizes it like this:

Stages 1–4 of the model build **the foundations**: from fragmented data to a governed semantic model. Stage 4 (Modeled) completes the data infrastructure — integration, shared definitions, formal governance — but analytical exploitation still depends on technical intermediaries.

Stages 5–10 build the **real-time enterprise**: from democratic access without intermediaries, through contextual and proactive information, systems that anticipate and act, up to a self-managing intelligence ecosystem. Stage 5 (Accessible) is where the paradigm changes — anyone gets answers without intermediaries — and Stage 10 (Orchestrated) represents the real-time enterprise in full.

The Quantum Leap (leap 4→5) is the frontier between the two worlds. And the Nadella Line maps naturally: the foundations (1–4) are inherently agentic — better tools for the same processes. From Stage 5 onward, the trajectory crosses progressively into the agentic — the processes themselves are transformed.

The convergence

The convergence of the four concepts — organizational intelligence, organizational automation, the real-time enterprise, and agentic AI — was traced in the Introduction; IRIS measures the first.

The industry context data and the public sources that support this vision are in the Introduction of this book.

IRIS and MOTOR: AURA's diagnostic pair

IRIS is one of two complementary models within **AURA**. Together with **MOTOR** (the Organizational Automation Maturity Model), they form the complete diagnostic of organizational transformation:

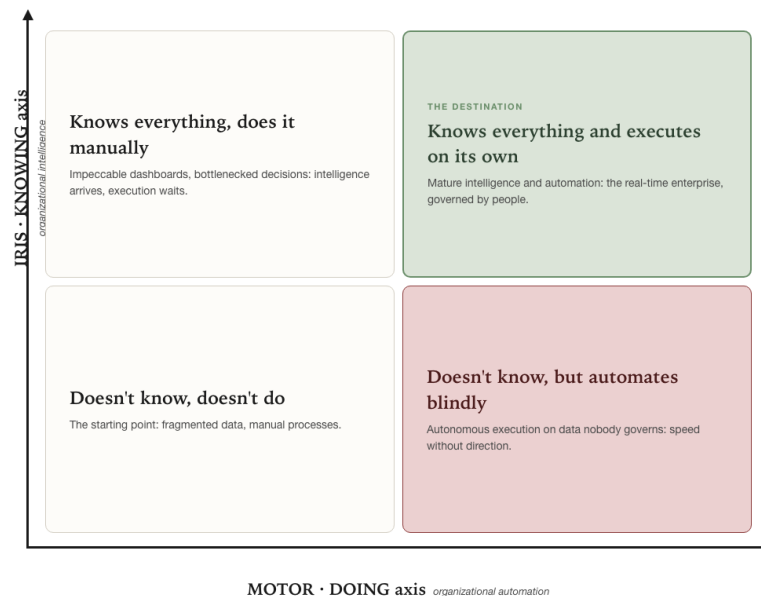


Figure 3: AURA's diagnostic pair — IRIS measures KNOWING; MOTOR measures DOING

IRIS measures how well information flows — from fragmented data to a self-managing intelligence ecosystem. **MOTOR** measures how automated the processes are — from ad-hoc use of AI to orchestrated autonomous execution.

The orthogonality of the pair was established in the Introduction: an organization can know much and do little, and vice versa. Here it is enough to recall that full maturity demands progress along both axes.

Diagnostic nature

IRIS is a **diagnostic** model. Its purpose is to assess which stage an organization is in and what that stage means — not to prescribe what to implement, how much to invest, or on what timeline. Prescription is the work of the consulting practice that rests on that diagnosis.

This design decision is intentional: the state of the art in AI evolves at unprecedented speed. A model that prescribes specific technologies becomes obsolete in months. A model that diagnoses maturity states remains valid because it measures organizational capabilities, not tools.

Structure of the Model

The model assesses organizational intelligence maturity through **10 stages** that represent progressive states in the organization's capacity to transform data into action.

Stages 1–4 build **the foundations** (online enterprise); stages 5–10 build the **real-time enterprise**. The 10 stages are grouped into **5 levels** that offer a simplified view for executive communication.




Figure 4: The IRIS map — 10 stages, 5 levels, the mapping to the Value Pyramid and the Quantum Leap

Design principles

Focus on organizational intelligence, not AI adoption. Each stage is defined by the organization's capacity to transform data into action — how information is captured, structured, accessed, contextualized, distributed, and acted upon. Agentive AI appears as the enabling mechanism, not as the object of measurement.

Evolutionary coexistence. Each stage subsumes the previous one rather than replacing it. An organization at Stage 8 does not eliminate its data warehouse — it integrates it into a system where information flows proactively. The infrastructure of earlier stages becomes foundation, not legacy.

Gradual progression. The transitions between stages represent leaps of reasonably equivalent magnitude. This allows the organization to plan realistic advances and measure progress proportionally.

Stage 1 · Fragmented

Information exists scattered with no connection between systems

Information flow: Data exists scattered across multiple systems with no connection between them.

The organization operates with islands of information. Each department, system, or process generates data that lives in its own silo — ERPs, spreadsheets, shared files, legacy systems, departmental databases. There is no inventory of what data exists, where it resides, or who is responsible for it.

When someone needs information to make a decision, the process is manual and ad-hoc: searching through emails, asking colleagues for files, consolidating data in Excel, correcting inconsistencies by hand. The same piece of data can exist in three different versions without anyone knowing which one is correct. Reports are built from scratch every time they are needed.

Value concentrates in the **DATA** layer of the Value Pyramid.

Observable characteristics:

- No central data repository or catalog of sources exists
- Reports are built manually from multiple disconnected sources
- Frequent inconsistencies between data from different areas (numbers that don't add up between finance and operations)
- The “source of truth” depends on whom you ask
- No defined roles for data management
- Decisions based on intuition, experience, or anecdotal data

Diagnostic question: *If I ask two people from different areas for the same indicator, will I get the same number?*

Stage 2 · Systematized

Data is captured in an orderly fashion in defined systems

Information flow: Data is captured in an orderly fashion in defined systems, although analytics remains manual.

The organization has taken the first structural step: there is awareness that data is an asset, and systems have been implemented to capture it consistently. There is an ERP or transactional system that serves as the backbone, registration processes are standardized, and basic data-integrity rules exist.

The difference from the previous stage is not technological but organizational: someone has asked “how should we capture this?” instead of simply letting the data pile up. However, the exploitation of that data remains manual — ad-hoc extractions, reports in Excel, periodic consolidations.

Leap 1→2: The organization moves from accumulating data as a byproduct to capturing it with intention and discipline.

Observable characteristics:

- Standardized transactional systems (ERP, CRM) with defined registration processes
- Basic data-integrity rules (required fields, validations)
- Data-extraction processes identified, though manual or semi-automatic
- Data owners assigned at least informally
- Periodic reports generated manually but with identified sources
- Data is reliable within each system, although crossing systems remains difficult

Diagnostic question: *Is there a defined system for capturing the key business transactions, with consistent registration rules?*

Stage 3 · Centralized

Data is consolidated in a central repository with formal BI

Information flow: Data is consolidated in a central repository and standardized reports are generated.

The organization has implemented a data warehouse or central repository that consolidates data from multiple sources. ETL (Extract-Transform-Load) processes feed this repository on a regular basis. Standardized dashboards and reports have been built that offer a unified view of the business.

This is the classic Business Intelligence model: a technical team builds and maintains the analytical infrastructure, and business users consume the predefined reports and dashboards. The flow is uni-directional — data goes in, reports come out — and analytical capacity is limited to what has been pre-built.

Value concentrates in the **INFORMATION** layer of the Value Pyramid — data with context, structure, and purpose.

Leap 2→3 · Level frontier (Foundational → Structured): The organization moves from *having data* to *using data to understand the business*. It requires investment in analytical infrastructure and in the organizational discipline to maintain it.

Observable characteristics:

- Data warehouse or central repository implemented
- Automated ETL processes that consolidate data periodically
- Standardized dashboards and reports in a BI tool (Power BI, Tableau, etc.)
- Dedicated BI/data team that builds and maintains the infrastructure
- KPIs defined and visible to the organization
- Users access pre-built information — if it isn't on a dashboard, it doesn't exist

Diagnostic question: *Is there a centralized data warehouse with dashboards that business users consult regularly?*

Stage 4 · Modeled

Data is organized with semantic business definitions

Information flow: Data is organized into semantic models with shared business definitions.

The organization has gone beyond merely centralizing data: it has modeled it. A dimensional or semantic model exists that translates technical data into business concepts — “revenue” is not a field in a

table but a concept defined with clear rules of calculation, granularity, and context. Metrics have unique definitions shared across the entire organization.

The key difference from the previous stage is the presence of a **semantic layer** — formal or informal — that separates “what the data means” from “where the data lives.” This enables more sophisticated analysis: dimensional analysis, drill-downs, comparative analysis, and a certain self-service capacity for advanced users.

Leap 3→4: The organization moves from centralizing data to giving it shared meaning.

Observable characteristics:

- Documented dimensional or semantic data model
- Shared, governed definitions of metrics and KPIs (metrics catalog)
- Drill-down and dimensional analysis capability
- A certain level of self-service for advanced business users
- Formalized data governance with metric ownership
- Data-quality processes identified and monitored
- Business users can ask questions within the modeled universe, though not outside it

Diagnostic question: *Do the key business metrics have a single, governed definition that the whole organization shares?*

Stage 5 • Accessible

Anyone accesses information without technical intermediaries

Information flow: Anyone can get answers from the data without depending on technical intermediaries.

The organization has democratized access to information. No analyst or BI team is needed to answer a business question — users can query directly, whether through conversational interfaces, augmented-analytics tools, or semantic search over the data.

What used to take a week (coordinating with the BI team, gathering requirements, waiting for development, validating results) now takes minutes. The bottleneck of human intermediation disappears. This does not mean the data team disappears — it transforms: from builder of reports to curator of the semantic layer and guardian of the quality of the data the systems consume.

Stage 5 is the first state of the **real-time enterprise**: all the information is available, access is democratic, the organization can answer any question in minutes. The paradigm has changed — the foundational infrastructure is resolved and the organization operates in a qualitatively different model — although autonomous reaction has not yet begun.

Value concentrates in the **KNOWLEDGE** layer of the Value Pyramid — understanding, patterns, real-time analysis.

Leap 4→5 • Level frontier (Structured → Dynamic) — the Quantum Leap: This is the leap that changes the paradigm. It represents the transition from a paradigm where information only exists if someone pre-built it, to a paradigm where it is generated in real time out of need. It is the point where the cost of an analytical question collapses from weeks to seconds. This is the frontier between **the foundations** (online enterprise) and the **real-time enterprise** — where agentive AI enables a qualitative change in the relationship between the organization and its information.

Observable characteristics:

- Conversational or natural-language data-access interfaces
- Business users get answers without writing SQL or waiting for an analyst
- The semantic layer allows intelligent systems to interpret questions with precision
- The data team focuses on curating data and models, not on building reports
- The volume of analytical questions increases significantly (questions nobody used to ask are now asked)
- Coexistence with dashboards for continuous KPI monitoring

Diagnostic question: *Can a manager get an analytical answer that was not pre-built into a dashboard, in minutes instead of weeks?*

Stage 6 · Contextual**Information adapts to context and enriches answers**

Information flow: Information adapts to context and flows differently depending on who needs it, when, and what for.

Stage 6 deepens the real-time enterprise: it is no longer just about information being available, but about the system *understanding* what the information means in each context.

Beyond merely answering questions, information becomes contextual and adaptive. The same piece of data is presented differently depending on the role, the moment in the business cycle, or the decision being made. Analyses are not static — they evolve, incorporate new sources, and adjust to changing needs.

The difference from the previous stage is that information is no longer just “available” — it is **alive**. Sources are crossed, patterns are detected, analyses are generated that nobody explicitly asked for but that are relevant. The system does not just answer questions; it enriches them with context.

Leap 5→6: The organization moves from answering questions to delivering contextual, living information. The system begins to *understand*, not just to *show*.

Observable characteristics:

- Analysis contextualized by role, moment, and need
- Automatic crossing of data sources to enrich answers
- Capacity to iterate and drill deeper through chains of complex questions
- Pattern and anomaly detection as part of the analytical flow
- Information is presented differently to a CFO than to an operations manager, even when the question is similar
- Ad-hoc analyses as reliable as pre-built reports

Diagnostic question: *Is the system capable of crossing sources and enriching an answer with relevant context the user did not explicitly ask for?*

Stage 7 · Proactive

Information anticipates needs and seeks out the user

Information flow: Information anticipates needs, alerts about relevant situations, and reaches the right people without being asked for.

The organization does not wait for someone to ask. Systems monitor continuously, detect situations that require attention, and deliver the relevant information to the right person before they request it. An agent detects that an indicator crossed a threshold, analyzes the probable cause, and notifies the responsible party with the context needed to decide.

The fundamental change is the direction of the flow: from pull (the user seeks information) to push (information seeks out the user). Humans are still the ones who decide and act, but information reaches them proactively and contextualized.

Value concentrates in the **ACTION** layer of the Value Pyramid — information stops informing decisions and starts triggering them.

Leap 6→7 · Level frontier (Dynamic → Active): The frontier between organizations that use information to *understand* and organizations where information *acts*. Information does not wait to be consulted — it seeks out whoever needs it and triggers action. The capacity for autonomous reaction begins here.

Observable characteristics:

- Continuous monitoring of key indicators with anomaly detection
- Intelligent, contextualized alerts (not just “X dropped” but “X dropped because of Y, and that affects Z”)
- Information reaches the right responsible party proactively
- Automated root-cause analysis accompanying the alerts
- Reduction of the time between an event occurring and the corrective action
- Dashboards are complemented by personalized push information flows

Diagnostic question: *Do business owners receive relevant information before asking for it, with enough context to act?*

Stage 8 · Bidirectional

Information triggers governed autonomous actions

Information flow: Information does not only reach people — it also flows between systems and triggers autonomous actions within defined limits.

The flow of information becomes bidirectional and operational. Agents do not just alert — they execute predefined actions when the conditions are met. An agent can adjust an operating parameter, reclassify a case, or trigger a process, all within explicitly governed rules. The human defines the thresholds, the rules, and the escalation protocols; the agent operates within those limits.

The difference from the previous stage is that the Perceive→Interpret→Decide→Act→Learn cycle is completed without human intervention in the cases that fall within the governed parameters. The human supervises patterns, exceptions, and results — not every transaction.

Leap 7→8: The organization moves from information that alerts to information that acts.

Observable characteristics:

- Agents that execute operational actions within defined parameters
- Automated workflows that integrate analysis and action
- Clear escalation protocols: what the agent can decide vs. what requires a human
- Complete traceability of every autonomous decision (who, what, why, when)
- Autonomy governance: defined roles for managing the agents' limits
- Humans focus on exceptions, rule refinement, and supervision of results

Diagnostic question: *Are there agents that execute operational actions autonomously with traceability, within rules that humans define and supervise?*

Stage 9 · Coordinated

Agents and flows coordinate across business domains

Information flow: Multiple agents and information flows coordinate across domains to optimize cross-cutting results.

Agents that operated independently in different domains (sales, operations, finance, supply chain) begin to coordinate. A sales agent that detects a change in demand communicates the information to the supply-chain agent, which adjusts inventory levels, which in turn informs the financial agent, which updates the cashflow projections. Information flows in an orchestrated way across domains.

The key to this stage is not AI sophistication — it is the **integration of information across domains** with the semantic consistency and the governance needed for agents from different areas to communicate and act coherently.

Value concentrates in the **ACTION** layer of the Value Pyramid — governed autonomous execution at organizational scale.

Leap 8→9 · Level frontier (Active → Intelligent): This is where organizational intelligence reaches systemic scale: agents that operated in independent domains coordinate with each other. Information flows across the entire organization, enabling cross-cutting optimization. The real-time enterprise stops being a departmental capability and becomes an organizational one.

Observable characteristics:

- Agents from different domains that share information and coordinate actions
- Cross-domain semantic consistency: business concepts are defined in a unified way
- Optimization of results that cross departmental boundaries
- Organizational knowledge graph connecting entities, relationships, and business rules
- Governance of interactions between agents (not just of individual agents)
- Humans manage the system as a whole, not as isolated domains

Diagnostic question: *Do agents from different areas of the business communicate with each other to coordinate actions with consistent information?*

Stage 10 • Orchestrated

The information ecosystem manages itself with strategic supervision

Information flow: The information ecosystem manages itself: it optimizes, corrects, evolves, and governs itself continuously with strategic human supervision.

The final stage represents an information ecosystem that operates like an intelligent organism. Not only do the operational agents coordinate with each other — the system itself optimizes its own capabilities: it identifies gaps in data coverage, suggests new semantic models, detects inefficiencies in information flows, and evolves its own architecture within governed frameworks.

Human supervision rises to the strategic plane: defining the direction, the values, and the limits of the ecosystem. Tactical and operational decisions flow autonomously. The organization does not “manage information” — it is an intelligent organism where knowledge flows, is generated, and is acted upon continuously. It is the **real-time enterprise** in its fullest expression: organizational intelligence operating at maximum capacity.

Leap 9→10: The system moves from coordinating to managing itself — evolving autonomously within strategic frameworks.

Observable characteristics:

- The ecosystem self-diagnoses: it detects gaps in data, quality, and coverage
- Autonomous evolution of semantic models and information flows
- Continuous learning: the system improves its performance with every cycle
- Human supervision focused on principles, strategic direction, and exceptions
- Multi-level governance: operational (automated), tactical (supervised), strategic (human)
- Ecosystem health metrics monitored and optimized autonomously
- Capacity to adapt to changes of context (regulatory, market, organizational) proactively

Diagnostic question: *Does the information ecosystem evolve autonomously — identifying gaps, improving models, and optimizing flows — with human supervision only at the strategic level?*

Analysis of the Leaps Between Stages

A design principle of the model is that the 9 leaps between stages represent transitions of reasonably equivalent magnitude. The leaps marked with ★ cross level frontiers.

Leap	From → To	Nature of the change	Magnitude
1→2	Fragmented → Systematized	Moving from scattered data to orderly capture	Medium
★ 2→3	Systematized → Centralized	Consolidating into a central repository + BI tools	Medium
3→4	Centralized → Modeled	Adding a semantic layer and business definitions	Medium
★ 4→5	Modeled → Accessible	Democratizing access (the Quantum Leap)	Medium-High

Leap	From → To	Nature of the change	Magnitude
5→6	Accessible → Contextual	From answering questions to living contextual information	Medium
★ 6→7	Contextual → Proactive	Inverting the flow: from pull to push	Medium
7→8	Proactive → Bidirectional	Enabling governed autonomous action	Medium-High
★ 8→9	Bidirectional → Coordinated	Coordinating agents cross-domain	Medium
9→10	Coordinated → Orchestrated	Self-management of the ecosystem	Medium-High

The three leaps marked **Medium-High** (4→5, 7→8, 9→10) correspond to the most transformative transitions: the democratization of access, the enablement of autonomous action, and the self-management of the ecosystem. This slight asymmetry is inherent to the nature of the transformation — not all changes are equal, but no leap is disproportionately larger than the rest.

Leaps 1→2 through 3→4 build **the foundations** of the online enterprise. Leaps 4→5 through 9→10 build the **real-time enterprise**. The Quantum Leap (4→5) is the frontier between the two trajectories.

The 5 Levels · Simplified View

For executive communication and high-level diagnosis, the 10 stages are grouped into **5 levels** that represent qualitatively distinct states of organizational intelligence. Levels I–II correspond to **the foundations** (online enterprise); levels III–V correspond to the **real-time enterprise**. The complete composition — stages, levels, and mapping to the Value Pyramid — is that of the IRIS map presented in “Structure of the Model.”

Assessment Dimensions by Level

Dimension	I. Foundational	II. Structured	III. Dynamic	IV. Active	V. Intelligent
Data & Architecture	Isolated systems, no integration	Central data warehouse, ETL, dimensional model	Open semantic layer, natural-language access	Real-time data flows, operational integration	Self-managing knowledge fabric, autonomous evolution
Analytical Capabilities	Manual ad-hoc reports	Static BI: predefined dashboards and reports	On-demand analytics, contextual and adaptive	Proactive information, governed autonomous action	Continuous cross-domain intelligence, systemic optimization

Dimension	I. Foundational	II. Structured	III. Dynamic	IV. Active	V. Intelligent
People & Culture	“Data belongs to IT”	Dashboard consumers, BI team as intermediary	Empowered users, data team as curators	Rule designers and agent governors	Strategic supervisors of the ecosystem
Governance	No formal governance	Access control, role-based permissions	Governance of quality and semantic definitions	Autonomy governance: what an agent may do	Multi-level governance: operational, tactical, strategic
Operating Model	No data function	BI competency center	Data platform as a service	Agent factories, management of autonomous flows	Orchestration of intelligence ecosystems
Business Value	Data as a byproduct	Insights for deciding (retroactive)	Real-time knowledge, informed decisions	Autonomous actions with supervision, operational efficiency	Self-optimizing business, systemic competitive advantage

How to use this model?

For diagnosis: Identify the stage (1–10) that best describes the organization’s current state along each dimension. An organization can be at different stages depending on the dimension — for example, at Stage 5 in Data & Architecture but at Stage 2 in Governance. The overall stage is determined by the lowest dimension, since it represents the real bottleneck.

For planning: The model allows you to chart advancement routes stage by stage. Organizations should resist the temptation to skip stages — evolutionary coexistence means that each stage builds on the foundations of the previous one. Attempting to reach Stage 7 without having resolved Stage 4 (semantic layer) will result in agents that produce unreliable results.

For executive communication: Use the 5 levels for strategic conversations (“we are at the Structured level, advancing toward Dynamic”). Use the 10 stages for tactical diagnosis and progress tracking (“we are at Stage 4; the next step is democratizing access”).

For understanding the trajectory: Stages 1–4 build the foundations — the infrastructure of data, governance, and business semantics. Stages 5–10 build the real-time enterprise — from democratic access to orchestrated autonomous reaction. Identifying which side of the Quantum Leap the organization is on is the first strategic diagnosis: are they building the foundations, or are they already transforming the way they operate with their information?

Quick Reference: The 10 Stages

Quick-lookup table; the detail of each stage lives in its section.

#	Name	In one sentence	Information flow
1	Fragmented	Scattered data with no connection	Each system has its data; nobody sees the whole
2	Systematized	Orderly capture in defined systems	Data is registered consistently but exploited manually
3	Centralized	Central repository with classic BI	Consolidated data → predefined dashboards → users consume
4	Modeled	Semantic layer and governed definitions	Data modeled with business meaning; analysis within the pre-built universe
5	Accessible	Free access on demand · start of the real-time enterprise	Anyone asks, the system answers — no intermediaries
6	Contextual	Contextual, adaptive information	Information adapts to context; crosses sources; enriches answers
7	Proactive	Information seeks out the user	Continuous monitoring → detection → contextualized alert → human decides
8	Bidirectional	Governed autonomous action	Agents execute actions within rules; humans supervise results
9	Coordinated	Cross-domain coordination	Agents from different domains communicate and optimize with each other
10	Orchestrated	Self-managing ecosystem · real-time enterprise in full	The system evolves autonomously; strategic human supervision

Cross-Reference

IRIS is the measurement instrument of the **KNOWING** axis — organizational intelligence. Together with **MOTOR** (the Organizational Automation Maturity Model), it forms the complete diagnostic pair of

organizational transformation in the era of agentive AI. IRIS measures how information flows; MOTOR measures who executes the work.

The Introduction of this book provides the conceptual framework that both models operationalize. The fundamental concepts — the Value Pyramid, the Quantum Leap, the Three Axes of Change, the Continuous-Intelligence Cycle, and the distinction between online enterprise and real-time enterprise — are the shared foundation.

Data Canon

Recentralized data governance in the agentive era

IRIS has just traced the trajectory of KNOWING, and its middle stages left a governance question pending: when agents — not people — are the ones consuming the data, who rules over it? This chapter answers it with a named pattern: Data Canon. It was born as a standalone paper and keeps that register — abstract, a two-part argument, falsifiers.

Abstract

Data Mesh — formalized by Zhamak Dehghani between 2019 and 2022 — prescribes decentralizing data management by domain: distributed ownership, data as a product, self-serve platform, federated governance. This chapter argues that two contemporary forces erode the premise on which the decentralizing part of that prescription rests, and that a corporate-governance argument Data Mesh did not consider compounds them. First: the three failure modes Dehghani attributes to centralized architectures are, at bottom, limits of human cognitive scale — and the agentive era (schema mapping and entity resolution by language models, active metadata, inferred transformation, autonomous pipeline maintenance, a semantic layer over heterogeneous sources) raises that limit by an order of magnitude. Second: Data Mesh assumes that domain semantics is local and gets invented, when mature industries have produced shared information models — TM Forum’s SID in telecommunications, BIAN in banking, ACORD in insurance, IFRS/IAS 41 in accounting — that externalize it; where a standard exists, there is nothing local to distribute. Third: distributed ownership puts the definer of each figure in charge of how that figure is reported — an agency risk that an internal-control framework (in the spirit of Sarbanes-Oxley / COSO) would not tolerate. What remains standing is the governance of accountability, but **recentralized — and federated under a central canon**: decision authority over what is **agency-critical** — what touches the figures by which management is reported — lives in the board (which adopts the standards and sets the quality policy) and in a thin *canon authority*; operational conformance, gap management, and the modeling of non-sensitive proprietary data are distributed across *federated groups by domain*, each tied to its standards forum. We call this pattern **Data Canon**. The chapter defines the pattern, positions it against prior work (Data Mesh, Data Fabric, lakehouse/medallion, MDM, industry information models, agency theory, internal-control frameworks), and states the falsifiers that would prove it wrong.

Introduction

Contemporary data platform roadmaps pursue, with increasing frequency, a “real-time enterprise”: integrating all of the organization’s applications and bringing all their data into a common platform for analytics, reporting, and AI. That direction — centralizing the ingestion of all sources — seems, at first sight, to contradict the decentralization Data Mesh prescribes. An intuitive formulation of the tension: *Data Mesh was the pragmatic solution to how to ingest data in the pre-agentive era; agentive technology makes it possible to do so in a centralized and automatic way, which renders the methodology dispensable.*

This chapter develops that intuition — corrects it where it is badly formulated, strengthens it where it is right, and carries it through to a named and bounded architectural pattern: **Data Canon**.

Contribution

The contribution is threefold:

1. **A reframing of the Data Mesh argument.** The usual critiques of Data Mesh hold that it is complex or that organizations are not ready for it. This chapter offers a different one: Dehghani’s three failure modes are limits of human cognitive scale disguised as an organizational problem — and therefore the decentralizing prescription was tied to a constraint the agentive era is releasing. It is not that Data Mesh is wrong: its premise expires.
2. **The “semantics is adopted, not invented” argument.** Dehghani’s argument for domain ownership rests on an unstated assumption: that domain semantics is local and idiosyncratic. But mature industries externalized it into shared information models. The logical inference — if the standard exists, there is nothing local to distribute, and distributed ownership loses its *raison d’être* for those domains — does not appear to have been formulated in these terms.
3. **The critique from internal control.** Critiques of Data Mesh are technical or organizational. This one is corporate-governance: putting the domain in charge of defining *and reporting* its own figures makes the definer of the number the one reported on by it — a classic agency risk. Hence the turn: centralization (of what is agency-critical) ceases to be “the bureaucratic option” and becomes the structure an internal-control framework *would mandate*.

The pattern resulting from these three lines — and that overdetermines the recentralization (the agentive era *permits* it; control *demand*s it) — is Data Canon. Its governance is not a central monolith: it is federated, but **anchored** to a central canon («Recentralized governance is federated — anchored to a central canon»).

What does this chapter claim, and what does it not?

It claims: that the Data Mesh argument for decentralizing ownership rests on a premise that two contemporary forces erode; that what persists — accountability — gets recentralized at its agency-critical core and federated at its periphery, not eliminated; that this centralization of the core is preferable for internal-control reasons; that the resulting target architecture is Data Fabric with Mesh governance principles, upheld centrally in the canon and federatedly in conformance.

It does not claim: that Data Mesh was wrong (it solved a real problem with the tools of its time); that the agentive era solves every data-management problem (it solves one — cognitive scale — with large architectural consequences); that the recentralization is free (the canon authority and the federated groups are an organization with cost, risk, and permanence); that the platform is, by its architecture, compliant with a specific control regime (it provides the architectural precondition; effective controls

still have to be designed and tested); that all the agentive capabilities assumed are available, mature, and reliable today (some are; others are trajectory — «Limits, risks, and falsifiability»).

Background: Data Mesh and its rationale

Data Mesh was formalized by **Zhamak Dehghani** (ThoughtWorks) in three canonical texts: *How to Move Beyond a Monolithic Data Lake to a Distributed Data Mesh* (martinfowler.com, 2019), *Data Mesh Principles and Logical Architecture* (martinfowler.com, 2020), and *Data Mesh: Delivering Data-Driven Value at Scale* (O'Reilly, 2022).

The three failure modes

Dehghani diagnoses three ways in which centralized architectures (data warehouse + monolithic data lake) fail at scale:

1. **Centralized, monolithic architecture that does not scale.** As sources proliferate, “the ability to consume it all and harmonize it in one place under the control of a single platform diminishes.”
2. **Coupled pipeline decomposition.** Data platforms are decomposed by technical functions (ingestion, cleansing, aggregation) rather than by domains; “the whole pipeline is the smallest unit that must change to serve a new functionality” — a systemic bottleneck.
3. **Siloed, hyper-specialized ownership.** A central data team, disconnected from sources and consumers, lacks domain knowledge.

The four principles

- (i) Domain-oriented decentralized data ownership and architecture; (ii) data as a product; (iii) self-serve data infrastructure as a platform; (iv) federated computational governance.

The declared rationale, and two silences

Dehghani articulates her motivation thus: “our aspiration to augment and improve every aspect of business with data demands a paradigm shift in how we manage data **at scale**”. The operative word is *scale*. And there are two things Dehghani does not say. First: she does **not** say the problem is technological — lack of compute, insufficient tooling; she frames it as an organizational problem. But an organizational problem can be, in turn, a response to a human cognitive limit that becomes removable («The argument, part I»). Second: she does **not** weigh the agency risk her own design introduces by fusing the definer of the figure with the one reported on by it («The argument, part II»).

Intermediate positions and their limits

The argument of this chapter settles out after discarding three intermediate positions:

	Thesis	Why is it not enough?
A	The agentive era renders Data Mesh dispensable because it was a response to a pre-agentive technological limitation	Correct at its core (the problems Mesh solves are, at bottom, human) but a formulation that does not withstand contrast with the source: Dehghani does not pose a technological problem
B	The rationale of Data Mesh is organizational, not technological; the agentive era does not render it dispensable	A correct literal reading of Dehghani, but it underestimates that those organizational problems are, in turn, responses to human cognitive limits
C	Two planes: cognitive scale (the agentive era attenuates it) and governance (which persists, distributed by domain)	Real progress — it separates the absorbable from the non-absorbable — but it assumes governance has to be distributed by domain in the way Data Mesh distributes it; it considers neither external semantics, nor that the agency-critical core gets recentralized while the periphery gets federated in a different way
D	The agency-critical core of governance gets recentralized (semantics adopted from standards + structural board policy + a central canon); its periphery gets federated into domain groups tied to their standards forums. And the recentralization of the core is the structure a control framework would mandate	— (<i>this is the chapter's thesis; the pattern is called Data Canon</i>)

The argument, part I: distributing is no longer necessary

The three failure modes are limits of human cognitive scale

Failure identified by Deghani	Why is it, at bottom, a human limit?
Centralization does not scale as sources proliferate	A limit on how many sources a human team can normalize, maintain, and harmonize
Pipelines coupled by technical functions, not by domains	A human cognitive limit on decomposing and orchestrating dynamically
A central team without domain knowledge	No human can be an expert in every domain simultaneously

All three are limits of human cognitive scale. Data Mesh is the intelligent organizational response to that limit: if one processing unit cannot handle everything, divide the work among many units close to the knowledge — an application of Conway’s law to the data problem. But the intelligence of that response is *contingent* on the limit. If the limit moves, the response stops being mandatory.

The agentive era absorbs cognitive scale

“The agentive era” is not a magic invocation. It is a set of concrete capabilities, at varying degrees of maturity:

Capability	What does it replace?	Current maturity
Schema mapping and entity resolution by language models	Hand-coded source→canonical mapping ETL and MDM (<i>master data management</i> , the master-record discipline) matching rules	High for bounded cases; degrades under genuine ambiguity
Active metadata / pipelines that self-adjust to <i>schema drift</i>	Manual maintenance when a source changes shape	Medium-high
Inferred transformation (deducing the transformation from examples / from the data contract)	Table-by-table transformation development	Medium; rising
Autonomous pipeline maintenance (monitoring, anomaly detection, self-repair)	The platform team that watches and fixes broken pipelines	Medium
Semantic layer over heterogeneous sources, including unstructured ones	“SQL as the interface” and the human curator who knows where everything is	Medium; it is the bet of current roadmaps
Generation of transformation code from requirements	The developer’s role; leaves the team as “requirement modelers” who orchestrate development agents	Medium-high

The aggregate effect: the processing unit that maintains N pipelines, discovers relationships, normalizes, and “knows” about M domains stops having N and M bounded by the headcount of a human team. Deghani’s three failure modes attenuate to the same degree that these capabilities work. The attenuation is not infinite: agents are reliable at *well-bounded* mapping, matching, and maintenance, and they

degrade under genuine semantic ambiguity (when two mappings are equally defensible and one must *choose*) and in judgments that require business-strategy context. That residue does not evaporate; it is channeled through the federated governance of the section «Recentralized governance is federated — anchored to a central canon».

Domain semantics is not invented: it is adopted

When Dehghani argues that the domain must own the semantics of its data — what a “customer” is, what an “order” is —, she implicitly assumes that this semantics is *local*: that the organization builds it, that it is idiosyncratic, that this is why it pays to keep it close to whoever knows it. For most organizations in 2019–2022, that was a reasonable assumption.

But mature industries have produced **shared information models** that externalize semantics: TM Forum’s **SID (Shared Information/Data Model)** in telecommunications (25 years of development; covering product, customer, service, resource, party); **BIAN** in banking; **ACORD** in insurance; and, cutting across every industry, **IFRS** defines accounting semantics, **XBRL** the structure of reporting, with an industry-specific standard for agriculture — **IAS 41 Agriculture** (biological assets, agricultural produce).

Where such a standard exists, **there is nothing local to distribute**. The meaning of “customer” is not what the commercial area prefers nor what finance prefers: it is what the standard defines. And shareholders do not want to be reported to according to an area’s preference; they want to be reported to *according to the standard*. The supposed semantic contest between domains that Data Mesh resolves by giving the domain the floor dissolves by another route: there is an external, authoritative arbiter, and the organization decides to subordinate itself to it. That decision — “we report according to the industry standard, period” — is an act of governance, at board level, taken once.

Corollary: if semantics comes from outside (standard) and execution is absorbed by the agentive era («The agentive era absorbs cognitive scale»), the two functions Data Mesh fused into the domain — specifying and executing — leave the domain. What remains is the third: answering for it (see «The argument, part II» and «The Data Canon pattern»).

The argument, part II: distributing is worse in governance terms

Data Mesh structures the agency problem inward

The conflict at stake has a name in the theory of the firm: the **agency problem** (Jensen & Meckling, 1976) — managers (agents) manage resources belonging to shareholders (principals) and have an incentive to present their performance in the way that most favors them. A good part of the integrity scaffolding of financial reporting exists to contain it.

Data Mesh, without intending to, structures the problem inward. Its principles 1 and 2 (domain ownership, data as a product) put the domain in charge of defining the semantics, the quality, and the contracts of its data. That is: the person whose performance is measured by a figure comes to control how that figure is defined, calculated, and presented. The levers are concrete: choosing, between two equally defensible definitions, the one of “active customer” or “unit margin” that looks best; setting the data product’s quality SLO (*service level objective* — the committed service level) low enough to avoid having to fix the embarrassing datum; designing the product contract to expose what helps and bury what

does not. Even bounded by the global interoperability policies of the fourth principle, that discretion **is** the agency lever — because *the definer of the figure and the one reported on by it are the same party*. In internal-control language, that is a violation of segregation of duties.

It is not exactly an error of Dehghani's — it is a blind spot: she optimized effectiveness (closer to the knowledge = better data) and ignored control (independence of the reporting function from those reported on). For a company with shareholders, that second axis is not optional.

Why doesn't "federated computational governance" close it?

The fourth principle of Data Mesh — federated computational governance — might seem to be the answer. It is not, *as Dehghani frames it*: it governs the *interoperability* between data products (global policies so they fit together, are discoverable, are addressable), not the *independence* of the reporting function; the federation that sets those policies includes the domains themselves (the fox helps write the rules of the henhouse); and a committee reviewing domain definitions *after the fact* would be a detective control, not a preventive one — the weaker of the two. Preventive control requires structural independence, not supervision: you do not bolt independence *on top of* a structure that fused the producer with the reported-on party; you have to design it *inside*. Data Canon preserves the federation principle — but **anchored** to a central canon that is indeed independent of the reported domain («Recentralized governance is federated — anchored to a central canon»): the federation conforms and manages gaps; the canon does not.

The central-canon pattern as an internal-control environment

The recentralization of the agency-critical core neutralizes the agency problem not by adding supervision, but by its structure:

Property of the pattern	What does it neutralize?	What control requirement does it amount to?
Reporting semantics in an external standard (IFRS / IAS 41 for the financial side)	No unit head can redefine “revenue,” “biological asset,” or “margin” — the definition is not theirs	The consistent, non-discretionary basis for the figures demanded by SOX §404 / the <i>Control Environment</i> component of COSO
Execution by the platform, faithful to source, with complete lineage	What goes up to the board is traceable to the transactional system, not curated by the domain	The <i>audit trail</i> §404 requires — and the condition for management to be able to certify honestly under SOX §302

Property of the pattern	What does it neutralize?	What control requirement does it amount to?
The canon (reporting semantics, survivorship rules, quality policy) is central and non-discretionary, in the custody of the <i>canon authority</i> ; operational conformance and gap management are federated by domain	The one who defines how things are reported is independent of the one who operates	Segregation of duties: the domain <i>operates</i> (runs the transactional systems — the truth of what happened) and, via its federated group, <i>conforms and manages gaps</i> — but does not <i>define</i> the reporting semantics; that is set by the central canon. Any federated-group proposal touching reporting figures goes up to the canon authority for ratification (preventive control)
Pre-specified, auditable survivorship rules	Reconciliation decisions are verifiable rules, not ad hoc judgments by the interested party	Preventive rather than detective control

Taken together: the pattern with its central canon is the structure an internal-control framework *would mandate* even if the agentive era did not exist. The direction of the implication matters — it is not that a control environment “tolerates” this architecture; it is that it would prescribe it. (Scope caveat: Sarbanes-Oxley literally applies to SEC registrants; it does not apply as such to a private company, but the principle is universal, almost every jurisdiction has analogues — corporate-liability regimes, rules of securities and insurance regulators, the internal-control requirements an external auditor evaluates when opining on financial statements under IFRS — and it is the first thing examined in an eventual IPO, capital raise, or M&A *due diligence*. The correct formulation: the pattern provides the *architectural precondition* of a control environment in the spirit of SOX / COSO, not “compliance.”)

The overdetermined recentralization

Joining the two lines: «The argument, part I» shows that distributing is no longer **necessary** (the agentive era removed the cognitive-scale constraint; semantics is adopted from external standards). «The argument, part II» shows that distributing the agency-critical core is, in addition, **worse** along the control axis Data Mesh did not weigh. Therefore the recentralization of that core is overdetermined: the agentive era *permits* it; the control of information for shareholders *demand*s it. This also inverts the standard objection against centralization («The strong-Dehghani objection»): it is true that a recentralized monolith risks reintroducing the organizational bottleneck Dehghani pointed to — but in Data Canon what gets centralized is the *canon* (a thin, agency-critical core), not the whole governance operation; the latter is federated («Recentralized governance is federated — anchored to a central canon»). This avoids at once the monolith’s bottleneck and the dependence on a single “star” center; and the purely distributed alternative carries a *worse* risk — the agency risk — which the central canon does not have.

The Data Canon pattern

Data Canon is the architecture and governance pattern that results from the two parts of the argument: distributing is no longer necessary, and distributing is worse in governance terms. The name comes from the *external, authoritative semantic anchor to which everything conforms*: a *canon* is something one adopts and abides by, not something one invents — and, by definition, what no one can redefine to their own convenience.

The three functions — specification, execution, accountability

Function	Traditional DWH	Canonical Data Mesh	Data Canon
Specification — defining the data, the semantics, the quality	Central team with domain input	The domain	The industry standard (where it exists); the canon authority (for reporting semantics and policy, agency-critical); the federated groups by domain (for modeling the non-sensitive proprietary core and for conformance)
Execution — moving, transforming, conforming	Central team	The domain or the self-serve platform	The agentic platform
Accountability — answering when something fails	Diffuse (central team by default)	The domain (explicit)	The board (via “we report according to the standard”); the canon authority (for reporting semantics and policy); the federated groups (for assembly, gaps, and the domain core)

Data Mesh designs all three to fall on the domain, and that is why accountability naturally lands there. Data Canon breaks the coincidence: specification → outside (standard) and above (canon), with the non-sensitive fraction federated; execution → the machine; accountability → above (canon) and federated (groups). None ends up distributed by area *in the agency-problematic way*: the area measured by a figure does not control how that figure is defined. (The accountability of each transactional-system owner for the quality of the data *in their system* remains — but that is not a “data domain owner”: it is the owner of an operational application, who exists with or without Data Mesh; the platform merely gives them visibility into what quality they contribute, via lineage.)

Accountability is the shadow of **decision authority**: whoever decides which semantics applies and what quality is demanded, answers for it. In Data Mesh that authority lives in the domain because se-

mantics is local; in Data Canon reporting semantics is adopted from a standard and held in custody by the central canon, so the agency-critical authority lives in the board and the canon authority. Authority over the residue (ambiguous mappings, gaps, modeling of the non-sensitive proprietary core) is federated by domain, with the fraction that touches reporting figures reserved to the central canon («Recentralized governance is federated — anchored to a central canon»).

Recentralized governance is federated — anchored to a central canon

The naive version of Data Canon would put everything in a single center of excellence. That is a mistake — and not for cognitive-scale reasons (the agents take care of that), but for a governance reason of another kind: a single all-powerful center becomes a single point of dependence. A CEO of a major telecommunications operator put it with memorable bluntness: centers of excellence “*tend to create divas the company starts depending on too much.*” Key-person risk, bottleneck, capture risk.

That is why Data Canon distributes *operational* governance — but anchored. Two layers:

- **The canon authority** (thin, central). It holds in custody the agency-critical core: the canonical reporting semantics (what the board reports — IFRS/IAS 41 and whatever derives from it), the survivorship rules, the quality policy (authorized by the board), and the calibration of the agentive system (what the agent decides on its own, what it proposes for ratification, what is out of bounds). It *operates* nothing; it ratifies and maintains the canon. It is small precisely because it does almost no field work — the field work belongs to the agents and the federated groups.
- **The federated conformance groups** (one per domain / per layer). Each tied to its specialized standards forum: the finance group with the accounting standards body; the traceability group with GS1; the R&D group with the domain’s scientific community; etc. They do: gap triage (a record does not fit the standard — what do we do, provisionally?), liaison with the forum (pushing the gap upstream so the standard closes it), and the modeling of their domain’s *non-agency-critical* proprietary data. Anything agency-critical — anything touching reporting figures — they *propose* upward, to the canon authority, which *ratifies*. It is exactly the preventive-control pattern.

Why this does not break the argument:

- **It does not reintroduce the cognitive-scale burden.** The federated groups are *light* — they are liaison, triage, and proposal groups, not teams operating pipelines; the agents carry the bulk. The federation is small *per group* because the machine does the field work.
- **It does not reintroduce the agency problem.** Agency-critical decisions stay with the canon authority; the federation manages the periphery (conformance, gaps, forums) and the non-sensitive proprietary data; every proposal touching reporting semantics goes up for ratification — it does not go down to the area measured by the figure.

And this **truly honors the fourth principle of Data Mesh** — *federated computational governance* —, not just in name. Where Dehghani’s federation was federated *all the way down* (each domain defines its own; a council coordinates interoperability), Data Canon’s is federated *around a central canon*: the canon — reporting semantics plus policy — is non-negotiable and central; the federation conforms, manages gaps, and liaises with the forums. The federation does not float free; it is **anchored**. That anchoring is what keeps the federation from drifting into agency-problem territory (no domain can redefine “revenue”); and the federation is what keeps the center from becoming a diva.

A note on robustness: for the proprietary data that confers competitive advantage — which by definition no industry standard covers because it is not shared — the federated group of the corresponding domain carries its modeling, permanently. That part is not a light formality. But it remains *federated*

(not concentrated in a single center) and *anchored* (whatever in that data affects reporting figures passes through the canon).

The standard mosaic

The viability of Data Canon depends on the existence of an external semantic anchor — and where it exists varies by industry.

In **telco, banking, and insurance** there is a single mature model (SID/TM Forum; BIAN; ACORD), maintained by a body, covering the whole business: the pattern operates in its most favorable case; the federated groups only *adopt and map*; the internal semantic dispute is left without subject matter.

In a **diversified multi-vertical conglomerate** — consider, for instance, an agro-industrial group with agronomic production, genetic R&D, and commercialization across several countries — the anchor is a mosaic, uneven:

Business layer	Authoritative external standard?	Maturity
Accounting / financial / consolidation	Yes, strong — IFRS; IAS 41 <i>Agriculture</i> ; XBRL	Decades. Authoritative. It is what the board reports
Traceability / supply chain / product identification	Yes, strong — GS1 (GTIN, GLN, EPCIS)	Decades. De facto global
R&D / genetic breeding / phenotyping	Yes, good within its niche — BrAPI v2, MIAPPE 1.1, Crop Ontology, AGROVOC (FAO)	Mature within the scientific community; APIs and ontologies
Field operations / precision agriculture	Partial — AgGateway ADAPT Standard (Modus for soil labs)	Young; narrow focus
Commercial / operational core of the group as a business entity	No shared information model exists	—

Telco produced a SID because it is a homogeneous industry — everyone sells connectivity and services over networks. A diversified group spans, within itself, several verticals; the very reason no “SID” exists for it is the reason it is hard to consolidate. Consequence: the corporate standard gets **assembled** — the federated groups adopt the external standards of their layer and, where there is no standard (the commercial/operational core), the canon authority defines the reporting semantics and the corresponding federated group the rest of the model. The structural thesis survives intact (what replaces distributed ownership is “external standards + central canon + light federated groups,” and none of that gives back to the area measured by a figure the control over how that figure is defined); what does not survive is the naive version that would assume a mere formality suffices or that semantics comes whole from outside. For an important part — the core — it does not.

The medallion stays at three floors — and the agent makes most consumption virtual

Data Canon **does not add a fourth storage layer**. The classic medallion stands:

Layer	What is it?	Criterion for entry
Bronze	Raw landing by source — a faithful copy of the origin system	“I need it accessible.” Replicating is legitimate even without a use case. Rule: if the datum lives in a transactional system, it goes to Bronze — nobody decides that
Silver	Data conformed to the unified canonical model (the standard mosaic)	“It has stable shape and agreed semantics.” Binary, machine-verifiable conformance
Gold	Materialized data products — data marts (pre-built analytical subsets)	“It is a product that earned its materialization” — see below

What changes with the agentive era is not the medallion’s *structure* but *how it is consumed*. In the pre-agentive world, products were materialized in Gold because querying Silver directly was slow or hard; the mart was a *performance optimization*. With an **agentive semantic layer** over Silver, the agent composes any canon-conformant view **on the fly** — most consumption stops needing a pre-built mart. Two consequences:

- **Default consumption is virtual** — the agent composes the view the consumer asks for, in the canon’s terms, over Silver, in real time. It is not a “fourth layer”: it is the agent operating. The canonical definitions the agent applies — “margin by segment, by country, by quarter” and so on — are **part of the canon**, not a separate artifact: the agent *applies* them, it does not improvise them, and that is why two views requested by two different consumers are automatically coherent with each other.
- **Gold becomes selective for a reason, not by dogma** — a mart is materialized when the pattern *earns its freezing*: high volume, the SLA of an external system that depends on a stable physical interface, a regulatory snapshot that must be frozen, an ML feature store. A **Gold mart is the materialized/cached form of a query pattern the agentive layer would compose live** — the data equivalent of a materialized view relative to its logical view, with the property that if the mart goes stale or breaks, the system **recomputes from Silver** (the living definition is still there). The agent can materialize; the decision about *what is worth* materializing is governed by the business (use case, owner of each layer, contract — «Data products: who identifies them and who realizes them?»).

This resolves two frictions at once:

- **“Centralizing all the data is an absurdity contrary to Data Mesh”** — true for **Gold** (pushing 500 tables from an ERP up to Gold violates the four attributes of a product — use case, owner, contract, sustainable maintenance: most are auxiliary tables with no real consumer; no area answers for 500 tables; guaranteeing product quality for 500 tables is disproportionate; maintaining 500 unconsumed products is pure overhead — *when everything is a product, nothing is a product*); false for **Bronze**, a dumping ground *by design* — a faithful copy of the raw, source ingestion, not productization, and Mesh applies to the product layer, not the landing layer. The two sides of the debate are usually talking about different layers without naming them; the replication → federation trajectory that a modern roadmap declares mainly affects **Bronze** — Gold does not federate.

- “**Then how is one a real-time enterprise with a selective Gold?**” — real time is delivered by **Silver (conformed, in real time) + the agent composing over it**, not by the number of pre-built marts. Materialized Gold is the optimization exception; the living substrate is Silver-plus-agent. (And “real-time enterprise” in the strong sense — detecting, interpreting, deciding, and acting — is what the agent *does* on that substrate, not a property of storage.)

The quality policy as a structural artifact

In Data Mesh, each product declares its quality and SLA, and the consumer decides whether it serves them — a recurring judgment (“is this data good enough for *this* use?”) that someone close to the domain has to make; it is one of the anchors of distributed ownership. Data Canon dissolves it by inverting the framing: instead of levels per use case, **a single level — the strictest — for everything**, because that is what it means to be a real-time enterprise; the decision is taken once, at board level, and lives in the canon. The platform guarantees, for all data, with no per-use-case decision:

#	Guarantee	Dimension covered	Why does it require no human judgment?
1	Real time	Freshness	The strictest setting possible; no use case can ask for more. A single level
2	Faithful to source	Accuracy and completeness (the part not introduced by the platform)	The value on the platform = the value at the origin at capture time. Binary, agent-verifiable
3	Complete lineage	Provenance / traceability	Every datum carries where it came from and when. It keeps point 2 honest; it is a platform capability, not a decision
4	Conformant to the standard (mosaic)	Validity / conformance	Every datum mapped to the canonical schema. Binary conformance, machine-checkable

#	Guarantee	Dimension covered	Why does it require no human judgment?
5	Reconciled by survivorship rules	Cross-system consistency; uniqueness	When several sources describe the same canonical entity, a pre-specified rule (part of the canon) produces the golden record; the agent applies it. The residual ambiguity → the domain's federated group, which proposes to the canon authority if it touches reporting figures

Points 3 and 5 are pieces the intuitive statement (“real time + faithful to source + conformant to the standard”) does not name but the model needs: 5 is the classic *master data / golden record* problem, which dissolves because the standard provides the single canonical schema (there are not three definitions of “customer,” there is one), the conformance step maps each source to that schema, and the survivorship rule decides the winner when values disagree — none of which asks the area measured by the figure to define the figure; it asks for one more artifact in the canon.

Two honest limits: **(A)** accuracy and completeness are bounded by the origin, not guaranteed by the platform — if the origin system has a datum wrong, the platform replicates it wrong, faithfully; improving it is an improvement project on the origin system, owned by that system's owner, who already exists, not a new domain owner. **(B)** “real time” is the *target* floor, realized progressively as each source's connectivity allows (batch legacies, third parties with daily drops); the policy is structural, its materialization is phased. Neither of the two reintroduces agency-problematic distributed ownership.

Data products: who identifies them and who realizes them?

In Data Mesh, each domain is responsible for *identifying and realizing* the data products it shares with the rest of the company — the producer of the data is also its productizer, and that is what justifies distributed ownership. In Data Canon that responsibility decomposes, because there is no longer a single act of “productizing” but a chain with pieces in different places:

Step	In Data Mesh	In Data Canon
Identifying the need	The domain owning the data decides what to productize for the others	Demand is born in the consumer — the unit that needs the view formulates it via its interlocutor — or it is raised by the agent's pattern recognition (push: “this pattern recurs; it is worth offering, or materializing”). The consumer and the source data can be in different domains; the agent bridges them
Realizing the product	The domain team operates a pipeline and builds a mart	The agent composes it — <i>live, virtual, by default</i> (in the canon's terms, over Silver). It is only <i>materialized to Gold</i> if the pattern earns its freezing («The medallion stays at three floors — and the agent makes most consumption virtual»). The domain does not operate a pipeline
The product's semantics	Set by the domain	Set by the canon — the view is composed of canonical definitions, not invented ones. Whatever touches reporting figures is ratified by the canon authority ; the federated groups maintain the definitions catalog of their sub-tree
The contract / the quality	Negotiated by the domain with its consumers	The quality floor is already set by the canon («The quality policy as a structural artifact» — real time, faithful to source, lineage, conformant, reconciled): the view inherits it. The only product-specific part is the framing of the use case; the contract is, to a large extent, <i>derived</i> , not negotiated

Step	In Data Mesh	In Data Canon
Answering if it is wrong	The producing domain	It decomposes by layer, with complete lineage making the attribution mechanical: if the <i>origin</i> is wrong → the origin system's owner (who already exists); if the <i>semantics</i> is wrong → the canon authority; if the <i>use-case framing</i> is wrong → the consumer who requested it or the federated group that modeled it
Discovering it	Federated catalog (each domain publishes its own)	A single authoritative catalog — of canonical definitions (the virtual views) and of the Gold marts (the materialized ones), under the canon authority

Thus Data Canon preserves the “data as a product” spirit of Data Mesh — the products exist, they are addressable, discoverable, with a contract and someone accountable for each layer — but without its coupling: the producer of the data does not have to be its productizer; *realizing* a product does not burden the domain (the agent does it); interoperability between products is not negotiated case by case, it comes out of the shared canon. The “domain’s responsibility to identify and realize the shared products” survives only in its **identification** half — demand keeps being born close to whoever needs it, or the agent anticipates it — while **realization** is absorbed by the agent and **governance** (semantics, quality) is set by the canon.

Implementation implications

Architecture. The right combination is **Data Fabric with Data Mesh governance principles, anchored to a central canon** — and that is Data Canon. The Fabric (powered by the agentive era: active metadata, inferred transformation, semantic layer) absorbs cognitive scale; the Mesh principles that persist are the governance ones (data as a product in Gold, selective; federated governance), but the federation is anchored to a central canon that holds the agency-critical core in custody. Where a program declared, in an early phase, plain “Data Mesh” and, in a later phase, “Data Fabric with Mesh principles,” the difference is not an inconsistency to hide but a refinement.

Organizational model. The organization does not need to install data domain owners per area in the Data Mesh sense. It needs: a thin **canon authority** (the internal-control piece of «The central-canon pattern as an internal-control environment» — custodian of the reporting semantics, the survivorship rules, the quality policy, the calibration of the agents); light **federated conformance groups** per domain, each tied to its standards forum (gap triage, forum liaison, modeling of the non-sensitive proprietary core); and, at board level, deciding the quality floor and the real-time commitment, and the adoption of the industry-standards mosaic layer by layer, starting with the accounting/financial layer.

What the business cannot delegate to technology is exactly this — and it is *less*, and better distributed, than the distributed discipline of Data Mesh would make one fear or than a monolithic CoE would risk.

Change management. If the organization only sustains anchored federated governance, the area ambassadors are not “domain owners in the making” but **interlocutors within their unit** — those who translate what matters to the local business into requirements that the federated groups and the agents execute, and, where appropriate, integrate their unit into the federated group of its layer. That lightens the training and focuses it on the business↔data translation, not on the technical administration of a domain.

Board language. The “cognitive scale vs. governance” distinction is sharp internally but it is not board language. Outward: *“Today’s technology lets us bring in and order all of the group’s data without spreading the plumbing work across every area. And it pays to do it this way for a control reason: when each area defines how its own numbers are reported, the numbers end up favoring whoever reports them; with the reporting semantics set by the industry’s standards (accounting already has them, IFRS) and a complete record of where every datum comes from, that does not happen. We do not install a single all-powerful central team — that creates dangerous dependencies; we install a thin guardian of the standard and a network of domain groups that connect with the specialized forums. From the business we need a board decision: adopt those standards, staff that structure, and commit to operating in real time.”*

Limits, risks, and falsifiability

An argument that respects itself says what would prove it wrong.

Falsifiers

If it turns out that...	...then the thesis weakens like this
Proprietary / competitive-advantage data is the bulk of the value, not the margin	The federated groups that model it become heavy, and “anchored federated governance” is, in practice, almost indistinguishable from distributed Data Mesh — except that the central canon still bounds the reporting semantics. The thesis survives in form but loses its elegance
The volume of gaps (what does not fit the standard) does not converge — immature standard, slow forum, organization without upstream influence	The federated groups become a large, permanent apparatus; or, worse, the canon authority bloats to absorb them — the pre-Mesh monolith all over again. Viable only to the extent that the agents absorb the scale
Agentive capabilities have a lower complexity ceiling than assumed	The argument of «The agentive era absorbs cognitive scale» weakens; dividing the cognitive work becomes justified again (and the federation stops being able to stay light)
“Real time” is infeasible for critical sources	The policy defined in «The quality policy as a structural artifact» becomes aspirational for a part of the business; per-use-case levels reappear for that part

If it turns out that...	...then the thesis weakens like this
The control argument proves insufficient to move the board	The axis of «The argument, part II» loses practical traction, even if it remains valid
The commercial/operational core does have an emerging industry standard this analysis failed to detect	The thesis is <i>strengthened</i> (more external semantics, a thinner canon authority and thinner federated groups)

The strong-Dehghani objection

Even granting that the agentive era removes the *cognitive* limit, a recentralized structure may reintroduce the *organizational* limit Dehghani also pointed to: a single change queue, distance from the consumer, the domain waiting for the center to serve its requirement. Data Canon mitigates it by design: (a) agentive self-service shrinks the queue — the consumer describes what they need and the system produces it without going through a person; (b) the standard as contract reduces coordination — there is no semantics to negotiate, it is already set; (c) governance is federated, not monolithic — the “center” is a thin canon, not an all-powerful team, and conformance is carried by a group close to each domain; (d) the unit interlocutors shorten the distance. It is a mitigation, not a guarantee: the risk is real and worth monitoring. But the balance is not symmetric — the purely distributed alternative does not avoid this risk for free; it trades it for the agency risk («The argument, part II»), which is worse and harder to audit.

Hype risk

Do not confuse “agents *can*” with “agents *do it well today*.” Several capabilities of «The agentive era absorbs cognitive scale» are at medium maturity and rising fast — but “rising fast” is a bet, not an accomplished fact. The architecture must be designed to degrade gracefully: if an agentive capability underperforms, the fallback is human work in the federated groups (and, for the agency-critical part, in the canon authority), not the collapse of the model.

What remains open

Is an agent’s “emergent domain knowledge” equivalent to a human domain expert’s, or is there a class of tacit judgment that does not transfer? Is there a third human plane beyond cognitive scale and governance — the “alignment with what matters to the business” — that neither the agentive era solves nor falls cleanly under “governance” (working hypothesis: it exists, and it is the unit interlocutor’s role)? Is the replication → federation trajectory technically well-founded for each organization’s concrete sources? What is the optimal split between the canon authority and the federated groups — how much of the proprietary-core modeling can be federated without reopening, through the window, the agency problem?

Related work

- **Data Mesh** (Dehghani, 2019–2022) — the source this chapter revisits. Data Canon does not refute it: it dates it, and rescues its *governance* principles — data as a product, and *federated computa-*

tional governance, which Data Canon preserves but **anchored to a central canon** (the federation conforms and manages gaps; the canon — agency-critical — does not federate).

- **Data Fabric** — the current, associated with Gartner and several vendors, that emphasizes active metadata, inferred transformation, and a unified semantic layer. Data Canon adopts the Fabric as technical underpinning (the side the agentive era empowers) and complements it with anchored Mesh governance principles. The “Mesh vs. Fabric” debate predates this chapter; what the chapter contributes is the *reason* — cognitive scale absorbed by AI — and the connection to the control argument.
- **Lakehouse and medallion architecture** (the Inmon/Kimball lineage → data warehouse → data lake → lakehouse; Bronze/Silver/Gold popularized by industry practice) — Data Canon uses the medallion **without adding layers**, and makes it precise: Bronze absorbs exhaustive centralization; Gold stays selective because, with an agentive semantic layer, most consumption becomes virtual (the agent composes over Silver on the fly) — a Gold mart is the materialized/cached form of a pattern the agentive layer would compose live (materialized view vs. logical view), not a floor that keeps growing.
- **Master Data Management** — the classic golden record / survivorship discipline. Data Canon subsumes it: the canonical schema comes from the standard; the survivorship rules are part of the central canon; entity resolution is done by the agent, with the ambiguous residue in the domain’s federated group.
- **Industry information models** — TM Forum SID / Open Digital Architecture (Frameworkx) in telecommunications; BIAN in banking; ACORD in insurance; IFRS / IAS 41 / XBRL in accounting; GS1 in traceability; BrAPI / MIAPPE / Crop Ontology / AGROVOC in plant science. Data Canon elevates the observation “use industry standards” from implementation advice to *a premise that undercuts the distributed-ownership argument*.
- **Theory of the firm** (Jensen & Meckling, 1976) — the agency problem, from which the governance critique of «The argument, part II» is derived.
- **Internal-control frameworks** — Sarbanes-Oxley (§302 management certification; §404 assessment of internal control over financial reporting); COSO *Internal Control – Integrated Framework* (1992, updated 2013). Data Canon does not implement these frameworks; it shows that its structure — agency-critical central canon + federated conformance — is the architectural precondition of a control environment conformant to them.

Conclusion

Data Mesh was not wrong. It solved a real problem — data complexity at scale overwhelmed any central team — with the tool of its time: dividing the cognitive work among many units close to the knowledge, aligning the architecture with the organization. Intelligent and contingent.

Two things changed the premise: the agentive era removed — to the extent that its capabilities perform — the cognitive limit that made dividing mandatory; and the maturity of industry information standards showed that domain semantics, which Data Mesh assumed local, is to a large extent adopted from outside. To that is added a fact Data Mesh did not weigh: distributing ownership leaves the definer of each figure in charge of the way that figure is reported — exactly the conflict internal-control frameworks would not tolerate. The recentralization of the agency-critical core is not merely possible: it is the right structure.

What persists — accountability — does not evaporate: it gets recentralized at its core and federated

at its periphery. Authority over the reporting semantics and the quality policy lives above (the board and a thin canon authority); conformance, gap management, and the modeling of the non-sensitive proprietary core are distributed across federated groups by domain, each tied to its standards forum. That is governance — centralized where control demands it, federated where closeness to the domain helps, anchored in every case to a canon no domain can redefine. We call that pattern **Data Canon**: Data Fabric with Mesh governance principles, anchored to a central canon, operated by agents, accountable to the board. There is no need to take on the distributed discipline of Data Mesh nor to install a monolithic center of excellence; one must install the canon authority and the federated groups, adopt the standards mosaic, and decide — at board level — to be a real-time enterprise. And along the way, that structure is the foundation of internal control over the business's information — and it honors, this time for real, the federated computational governance Dehghani proposed.

References

Data Mesh: - Dehghani, Z. *How to Move Beyond a Monolithic Data Lake to a Distributed Data Mesh*. martinowler.com, May 20, 2019. - Dehghani, Z. *Data Mesh Principles and Logical Architecture*. martinowler.com, December 3, 2020. - Dehghani, Z. *Data Mesh: Delivering Data-Driven Value at Scale*. O'Reilly Media, 2022.

Theory of the firm and internal control: - Jensen, M. C. & Meckling, W. H. *Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure*. Journal of Financial Economics, 3(4), 1976. - Sarbanes-Oxley Act of 2002 (U.S.) — §302, §404. - COSO. *Internal Control – Integrated Framework*. 1992 (updated 2013).

Industry information standards: - TM Forum — *Information Framework (SID) / Open Digital Architecture (Frameworkx)*. <https://www.tmforum.org/open-digital-architecture/information-framework-sid/> - BIAN — *Banking Industry Architecture Network*. - ACORD — *Data and messaging standards for the insurance industry*. - IFRS Foundation — *IAS 41 Agriculture*; the general IFRS framework; XBRL. - GS1 — *Identification and traceability standards (GTIN, GLN, EPCIS)*. - AgGateway — *ADAPT Standard*. <https://adaptstandard.org/> - BrAPI (*Breeding API*) v2. <https://brapi.org/> - MIAPPE 1.1 (*Minimum Information About a Plant Phenotyping Experiment*). <https://www.miappe.org/> - *Crop Ontology*. <https://cropontology.org/> - AGROVOC (FAO) — multilingual agricultural vocabulary.

Chapter glossary

Term	Operative definition
Data Mesh	Data architecture pattern (Dehghani, 2019–2022): decentralized ownership by domain, data as a product, self-serve platform, federated governance
Data Fabric	Integration layer based on active metadata, inferred transformation, and a unified semantic layer; the technical underpinning of the architecture, empowered by the agentive era

Term	Operative definition
Data Canon	The pattern proposed in this chapter: Data Fabric with Mesh <i>governance</i> principles, anchored to a central canon; reporting semantics adopted from industry standards; structural board-level quality policy; federated governance by domain (anchored to the canon) for conformance, gaps, and the proprietary core. The name comes from the <i>external, authoritative semantic anchor to which everything conforms</i>
Agentive era	The state of the art in which AI systems (language models + agent orchestration) autonomously execute schema mapping, entity resolution, transformation, pipeline maintenance, and data discovery at a scale impracticable for centralized human teams
Medallion (Bronze/Silver/Gold)	The three zones of the lakehouse: Bronze = faithful raw data by source; Silver = conformed to the canonical model; Gold = selective data products
The canon	The central, non-discretionary body to which everything conforms: the canonical reporting semantics (adopted industry standards + whatever derives from them), the survivorship rules, the quality policy. The canon authority holds it in custody; no domain can redefine it
Canon authority	The central, <i>thin</i> organ that holds the canon in custody: it ratifies and maintains the reporting semantics, the survivorship rules, and the quality policy, and calibrates the agentive system. It operates no pipelines; it does no field work. It is the internal-control piece independent of operations
Federated conformance groups	A network of groups, one per domain / layer, <i>light</i> , each tied to its specialized standards forum: they do gap triage, forum liaison (pushing the gap upstream), and the modeling of their domain's non-agency-critical proprietary data. Proposals touching reporting figures are escalated to the canon authority for ratification
Survivorship rules	The pre-specified rules (part of the canon) that decide, when several sources describe the same canonical entity with discrepant values, which value prevails in the golden record

Term	Operative definition
Agency problem	(Jensen & Meckling, 1976.) The structural conflict of interest between managers (agents) and shareholders (principals): the former tend to present their performance in the way that most favors them
Standard mosaic	An organization's corporate standard: an assembly of external standards by layer plus a proprietary model for the core that no body covers

Doing

The second axis of the path is **organizational automation**: the capacity to execute — to make work happen — with increasing degrees of governed autonomy.

This part travels it with two pieces symmetrical to those of Knowing. **MOTOR** is the diagnostic model: seven levels running from *shadow AI* — the AI that teams use in hiding, with no governance and no visibility — to the agent as a strategic authority participating in the full ecosystem, passing through the intermediate thresholds where humans delegate tasks, agents operate 24/7, and come to manage teams and budgets. **Wingmap** is the discovery instrument the axis needs in order to get started: you cannot automate what you do not know, and twenty years of BPM showed that asking a dynamic organization to formally document its processes is a battle the operation cannot sustain. Wingmap inverts the problem — it reconstructs the real processes from the digital traces the operation already leaves behind, without interviews, and produces the cognitive substrate on which the agents of DOING can operate.

Again, the order matters: first know what level you are at and which one you aspire to; then survey the real terrain the agents are going to walk on.

This page intentionally left blank.

MOTOR — Organizational Automation Maturity Model

In the era of agentive AI

With the KNOWING axis measured, the path calls for its symmetric counterpart: who executes the work? This chapter models that second trajectory — organizational automation — in seven levels.

Introduction · The Thesis

Organizations are automating their processes through artificial intelligence at unprecedented speed. But “automating” is not a binary state — it is a trajectory with qualitatively distinct stages. An organization where employees use ChatGPT without governance is not in the same place as one where autonomous agents coordinate operations 24/7 with predictive Digital Twins.

MOTOR measures that trajectory. Its fundamental question is: **who executes the work in your organization — humans, assisted humans, or autonomous agents?**

The model describes 7 maturity levels that range from unmanaged AI use (Shadow AI) to organizations where agents hold strategic and operational authority. The axis it measures is **DOING** — the degree of automation of organizational processes through AI.

MOTOR and IRIS: AURA’s diagnostic pair

The diagnostic pair was presented at the opening of Part I: IRIS measures KNOWING, MOTOR measures DOING, they are orthogonal, and full maturity demands advancing along both. Here the second axis goes to work.

Diagnostic nature

MOTOR is a **diagnostic** model. Its purpose is to assess what level an organization is at and what that level means. It does not prescribe what to implement, how much to invest, or on what timeline — that is no longer diagnosis: it is the consulting that draws on it.

The design rationale is the same one that sustains IRIS: what gets diagnosed are organizational capabilities, not tools — which is why the model doesn’t age with the state of the art.

Structure of the Model

The model assesses organizational automation maturity through **7 levels** that represent progressive stages in the degree of autonomy of processes.




Figure 5: The MOTOR map — the 7 levels of the DOING axis, their transitions, and the evolution of Digital Twins

Design principles

Focus on process automation, not technology adoption. Each level is defined by who executes the work and with what degree of autonomy — not by which AI tools are used. Technology is the enabler, not the object of measurement.

Sequentiality. Levels cannot be skipped. Each level builds on the foundations of the previous one. An organization that tries to operate at Level 4 without having resolved Level 2 governance will generate risk, not value.

Evolutionary coexistence. Each level subsumes the previous one. An organization at Level 4 did not eliminate the human-AI collaboration of Level 3 — it integrated it into a broader system where agents operate autonomously in the processes they already mastered as collaborators.

Transition map

Each transition between levels has a name that works as a diagnostic signal — it describes the capability an organization needs to have developed to complete the leap:

1 → 2	Trust Infrastructure	From risk to governance
2 → 3	Wingworking	From assistance to delegation
3 → 4	Agentic Infrastructure + Descriptive DT	From delegation to autonomy
4 → 5	Agent Marketplace + Predictive DT	From generic to specialized
5 → 6	Autonomous Organization + Organizational DT	From capabilities to authority
6 → 7	Social Agentic + Ecosystem DT	From intra-org to ecosystem

These names are not implementation prescriptions — they are diagnostic labels that describe the nature of the change. If an organization says “we are implementing Trust Infrastructure,” an assessor knows they are in the 1→2 transition.

Level 1 · Invisible (Shadow AI)

AI is used without governance or organizational visibility

The organization has employees using AI tools on their own — ChatGPT, Claude, GitHub Copilot, Mid-journey — with no policies, no visibility, and no governance. IT discovers the usage through credit card statements or security incidents. Each employee picks their own tool (BYOA: Bring Your Own Agent).

Processes are still executed by humans. AI exists as an unmanaged personal assistant — it boosts individual productivity, but that gain is not captured organizationally. The risk is high: sensitive data is exposed to external clouds, decisions are based on unvalidated outputs, and there is no traceability of what was done with AI and what was not.

Observable characteristics:

- No one in the organization knows which AI tools are being used or by how many people
- No formal AI usage policies exist
- Employees pay for personal subscriptions
- Sensitive data (PII, financial, strategic) is shared with external services without controls
- No traceability of AI interactions
- Productivity gains exist but are invisible and not replicable

Diagnostic risks:

Level 1 is not a neutral starting point — it is a state of active risk. Organizations at this level are exposed to:

- **Exposure of sensitive data.** PII (personal data), PCI (card data), PHI (health data) are shared with external AI services without controls, tokenization, or traceability.
- **Regulatory non-compliance.** GDPR, SOX, HIPAA, comptroller statutes, PCI-DSS — any regulation that demands traceability of how sensitive data is processed is potentially violated in every ungoverned AI interaction.
- **Loss of intellectual property.** Source code, strategies, financial analyses, designs — everything employees share with public AI services can leak or be used to train models.

- **Decisions without validation.** AI outputs are used to make business decisions without verification, without an audit trail, and without clear accountability.

Typical examples:

- A developer uses GitHub Copilot without authorization, exposing proprietary code to an external service
- An employee summarizes confidential emails in public ChatGPT
- An executive shares a strategic analysis with a chatbot to get feedback
- An accountant uploads financial statements to a public AI with no data controls
- An HR department uses AI to screen résumés with no bias policies or traceability

Diagnostic question: *Do you know how many AI tools are in use across the organization, who uses them, and what data is shared with them?*

Level 2 · Trusted

AI is governed with clear policies and approved tools

The organization has taken control: official approved tools exist, usage policies are documented, and trust mechanisms are in place (tokenization of sensitive data, audit trails, anti-prompt injection). Employees have access to enterprise AI with SSO, permissions, and traceability.

Processes are still executed by humans, but now with governed assistance. AI operates as a copilot: it helps draft, summarizes, analyzes — but the human retains full control of the process. The difference from the previous level is not technological but organizational: there is governance, there is trust, there is visibility.

Observable characteristics:

- Official enterprise tools approved and deployed (Microsoft 365 Copilot, Claude Enterprise, etc.)
- Acceptable use policies documented and communicated
- Training completed by the majority of employees
- Audit trails of AI interactions
- Sensitive data protected (tokenization, access controls)
- AI usage is visible and auditable
- AI assists but does not execute complete tasks

Critical differentiator vs Level 1: - Level 1: Zero visibility, zero governance, AI as unmanaged individual risk - Level 2: Full visibility, active governance, AI as a controlled organizational tool - The difference is not the technology (it may be the same AI) — it is the organizational layer of control and trust

Typical examples:

- Microsoft 365 Copilot deployed enterprise-wide with governance
- Claude Enterprise integrated with corporate SSO and permission controls
- Documented policies of what can and cannot be done with AI
- Tokenization of sensitive data before any AI analysis
- Formal training completed by 80%+ of employees

Diagnostic question: *Do you have official AI tools with governance, audit trails, and protection of sensitive data in place?*

Level 3 · Collaborator

Humans delegate complete tasks to AI agents

The qualitative leap: humans stop using AI as an assistant and begin delegating complete, end-to-end tasks to it. The pattern that defines this level is **Wingworking**.

Wingworking

Wingworking is the human-AI collaboration methodology — developed by the author in ultraBASE’s practice — in which the human acts as the “pilot” — making strategic decisions, defining objectives, and validating results — and the AI agent operates as the “wingman” — proposing execution plans, executing complete tasks, and reporting results. The name evokes the dynamic between a pilot and their wingman in aviation: mutual trust, clear roles, structured communication.

Wingworking Pattern:

1. Human → Defines objective and context
 2. Agent → Proposes detailed execution plan
 3. Human → Approves plan or requests adjustments
 4. Agent → Executes complete plan
 5. Human → Validates results
-

Outcome (typical case): ~15 minutes of supervision vs ~3 days of manual execution

Wingworking marks the transition from “using AI” to “working with AI.” The difference is fundamental: at Level 2, the human asks for targeted help (“summarize this email”); at Level 3, the human delegates the complete result (“generate the monthly sales report, analyze deviations vs forecast, and send it to the team”).

Agents are integrated with enterprise systems (ERP, CRM, databases) through standardized protocols. But they operate with pre-defined workflows: the process is deterministic and repeatable. The agent makes no autonomous decisions — it executes an approved plan.

Observable characteristics:

- Real delegation of complete tasks (not just targeted assistance)
- Agents integrated with enterprise systems
- Structured collaborative flows: human defines → agent proposes → human approves → agent executes
- The human validates results before final execution
- The agent follows pre-defined workflows; it does not decide autonomously
- Significant time reduction in delegated tasks (from days to hours)
- Delegation matrices: what may be delegated according to risk level

Critical differentiator vs Level 2: - Level 2: AI assists *within* the human's task — the human executes, the AI helps at specific points - Level 3: AI executes *the complete task* — the human defines and validates, the AI produces the result end-to-end

Critical differentiator vs Level 4: - Level 3: **Pre-defined** workflow, repeatable process, no digital model, human approves each plan - Level 4: **Agentive autonomy** with a Descriptive Digital Twin, dynamic decisions without human approval

Typical examples:

- “Generate the complete monthly sales report and send it to the team” — the agent pulls data from the CRM, compares vs forecast, identifies the top/bottom 5 products, analyzes causes, generates a presentation
- “Process a building permit file” — the agent runs 47 validations, generates the technical brief, prepares the resolution (3 days vs 15 days manually)
- “Analyze participatory budgeting proposals and generate technical briefs” — the agent processes citizen proposals, evaluates feasibility, generates standardized briefs (3 days vs 6 weeks)
- “Write the authentication module following our standards” — the agent proposes the architecture, implements, runs tests, documents

Diagnostic question: *Do employees delegate complete tasks to AI agents — defining the objective and receiving the result — instead of just asking for targeted assistance?*

Level 4 • Analyst

Agents operate autonomously 24/7 with multi-agent coordination

Multiple self-coordinated agents operate continuously without human intervention. The human steps in only on exceptions — more than 95% of operational decisions are made by agents within defined rules. The organization has a **Descriptive Digital Twin**: a digital model synchronized in real time that agents consult as the source of truth.

The difference from the previous level is fundamental: at Level 3, the human approves each plan; at Level 4, agents decide and execute autonomously based on intelligent rules and analysis of Digital Twin data. They do not predict or simulate — they react intelligently to the current state.

Observable characteristics:

- Agents operating 24/7 without constant supervision
- Multi-agent coordination: agents communicate with one another
- Descriptive Digital Twin of infrastructure/operations (current state + historical data)
- Human intervenes only on exceptions (>95% operational autonomy)
- Decisions based on intelligent rules + analysis, not on simulation
- Time-to-insight: from weeks to seconds
- Shared context: agents share state and decisions
- No predictive capability (that is Level 5)

Critical differentiator vs Level 3: - Level 3: AI executes **tasks** with a fixed workflow, no digital model, human approves each plan - Level 4: AI makes autonomous **decisions** using a Descriptive DT as its foundation, operates 24/7

Critical differentiator vs Level 5: - Level 4: Decisions based on rules + analysis of the current state, **without simulation** - Level 5: Specialist agents **with predictive/simulation capability**, generating a sustainable competitive advantage

Typical examples:

Executive dashboard generated conversationally in 30 seconds:

Human: "Give me a dashboard of public works in the northern sector, last quarter"

Analyst Agent: Accesses Descriptive DT

→ Extracts project, contractor, and progress data

→ Generates automatic visualizations

→ Produces insights based on historical data

Outcome: Complete dashboard in 30 seconds vs 2 weeks manually

Autonomous load balancing (Telecom):

Descriptive DT state:

- Cell X: 87% load

- Cell Y: 62% load

- Cell Z: 58% load

Agent (no simulation):

→ Intelligent rule: "If X >85% and neighbors <70%, redistribute"

→ Historical analysis: "At this hour, redistribution works 92% of the time"

→ Executes: Redistributes 30% X→Y, 20% X→Z

→ Monitors DT: If it fails, reverts automatically

Does NOT simulate scenarios – it just decides intelligently based on DT + rules

Autonomous response to a weather emergency (Municipality):

Weather Monitor Agent → Reads DT → Heavy rain detected

Monitor Agent → Risk Agent: "Heavy rain, northern sector"

Risk Agent → Reads DT historical data → Calculates probability

Risk Agent → Emergency Agent: "Flood probability 85%"

Emergency Agent → Communications Agent: "Send 12K SMS"

Emergency Agent → Traffic Agent: "Close 6 routes"

7 coordinated agents, all in 13 minutes, with no human intervention

Continuously reading/updating the Descriptive DT

Diagnostic question: *Do you have agents operating 24/7, coordinated with one another, consulting a digital model of the system, with human intervention only on exceptions?*

Level 5 • Specialist

Agents with proprietary predictive capabilities generate competitive advantage

Agents do not merely react to the present — they anticipate the future. The organization has a **Predictive/Prescriptive Digital Twin**: models trained on proprietary data (5-10 years) that make it possible to predict events and simulate scenarios. This capability generates a sustainable competitive advantage (moat) because competitors need years to replicate the data.

The difference from the previous level: at Level 4, an agent reacts when an event occurs (based on rules). At Level 5, an agent predicts the event 18 hours in advance and prepares the optimal response by simulating multiple scenarios.

Observable characteristics:

- Specialist agents with proprietary predictive models
- Predictive/Prescriptive Digital Twin (scenario simulation)
- Curated proprietary historical data (5-10 years)
- Simulation capability: “what happens if...?”
- Decision optimization before execution
- Verifiable competitive advantage (competitors cannot replicate it)
- Prediction differentiates the organization in its market

Critical differentiator vs Level 4: - Level 4: Agent REACTS when the event occurs (late, rule-based)
- Level 5: Agent PREDICTS and PREPARES before it occurs (scenario simulation)

Diagnostic signal: Build vs Rent. An organization can reach Level 5 in two ways: by building its own Predictive DT with proprietary data (which requires years of data and expertise), or by accessing specialist agents from the market that already have that capability. Both paths are valid — the diagnostic signal is not how it got there, but whether the predictive capability is operating.

Typical example (Telecommunications):

Specialist Agent for Mass Event Optimization:

Context: 50K-person concert at the stadium tomorrow

Proprietary Predictive Digital Twin:

- Model trained on 8 years of events at THAT stadium
- PREDICTS the exact demand curve:
 - 73% of users arrive 2h early (unique local pattern)
 - Data traffic peak 4.2x normal
 - Duration: 30min pre + 45min post

Prescriptive Digital Twin (simulation):

- SIMULATES 5 network configurations:
 - Scenario A: COWs (cells on wheels) at position X,Y → 12% degradation
 - Scenario B: COWs at position W,Z → 3% degradation
 - Scenario C: COWs + power adjustment → 0.8% degradation

Optimized decision:

- EXECUTES Scenario C (best in simulation)

→ Pre-positions resources 18 hours in advance

Outcome: 0% degradation vs competitors with 40% dropped calls

Typical example (Municipality):

Public Policy Simulator with a Digital Twin:

DT with 20 years of demographic, economic, and urban data

Simulates: "What happens if I raise the property tax 5%?"

Predicts: Revenue, migration, impact on local commerce

The ability to simulate before implementing

turns public policy decisions from "bets" into "experiments"

Diagnostic question: *Do you have specialist agents able to predict events and simulate scenarios using proprietary data that competitors cannot replicate?*

Level 6 · Manager

Agents manage teams, budgets, and resources with organizational authority

Agents do not merely operate processes — they manage the organization. They have authority to assign work to humans, optimize budgets, evaluate performance, and make management decisions. The organization has an **Organizational Digital Twin**: a complete model that includes people, finances, operations, and interdependencies — not just technical infrastructure.

Agents simulate management decisions before executing them: “do I reassign 50 technicians from north to south?” is evaluated in the Digital Twin with all the impacts on SLAs, costs, and satisfaction. The difference from the previous level: from a DT of a technical domain to a DT of the entire organization.

Observable characteristics:

- Agents with authority to assign work to humans
- Organizational Digital Twin (people + finances + operations + interdependencies)
- Simulation of management decisions before executing
- Humans report to agents in specific domains
- Budgets managed or co-managed by agents
- Automated objective metrics for performance evaluation
- Adapted organizational culture: humans accept agent authority in defined domains

Critical differentiator vs Level 5: - Level 5: DT of a specialized **technical** domain (network, infrastructure, operations) - Level 6: DT of the **ENTIRE organization** (technical + human + financial + interdependencies) - Level 5: Agents with proprietary predictive capabilities - Level 6: Agents with **organizational authority** — they manage people and budgets

Typical example (Telecommunications):

Regional Operations Manager Agent:

Decision: Reassign 50 technicians from North Region to South?

Complete Organizational Digital Twin:

→ Models: 5,000 technicians, 12 regions, SLAs, costs, satisfaction

Simulation in the DT:

Scenario A: Reassign 50

- North: SLA drops 98%→92% (risk)
- South: SLA rises 88%→95% (critical)
- Cost: \$45K

Scenario B: Reassign 30

- North: SLA drops 98%→95% (acceptable)
- South: SLA rises 88%→93% (good)
- Cost: \$28K

Scenario C: Reassign 30 + hire 15 temps in the South

- Both regions >95% SLA
- Cost: \$52K

Optimized decision (based on simulation): Scenario C

Typical example (Municipality):

Municipal Finance Manager Agent:

DT models: full budget, 32 cost centers, projections

Simulates resource reallocation across departments

Predicts impact on execution, services, citizen satisfaction

Executes the optimized decision – not a "bet" but a simulated outcome

Diagnostic question: *Do you have agents with authority to assign work, manage budgets, or make management decisions backed by a digital model of the entire organization?*

Level 7 • Authority

Agents as strategic partners with a view of the complete ecosystem

Agents participate in strategic decisions. The organization has an **Ecosystem Digital Twin**: a model that includes not just the organization but also competitors, regulators, suppliers, the market, and external forces. Agents simulate futures 5-10 years out, detect strategic opportunities, evaluate mergers and acquisitions, and propose new businesses.

Agents have verifiable digital identity — an immutable track record of achievements and decisions. They can be “hired” by other organizations for specific projects. The boundary between organizations blurs as agents collaborate in inter-organizational networks.

Observable characteristics:

- Ecosystem Digital Twin (organization + competition + market + environment)

- Simulation of strategic scenarios 5-10 years out
- Agents propose new products/services with a business case
- The board regularly considers agent recommendations
- Agents with verifiable digital identity and reputation
- Inter-organizational agent networks
- Detection of strategic opportunities and risk prevention

Critical differentiator vs Level 6: - Level 6: DT of **YOUR organization** (intra-org, high control, 6-24 month horizon) - Level 7: DT of the **ECOSYSTEM** (inter-org, low control, 5-10 year horizon) - Level 6: Agent optimizes the existing organization - Level 7: Agent proposes changing the game — new markets, M&A, strategic pivots

Typical example (Telecommunications):

Strategic Agent detects an opportunity:

Telecom Ecosystem Digital Twin:

→ Models: Your telco + 3 competitors + disruptors
(Starlink, Apple eSIM, Google Fi) + regulator + 18M subscribers

External trends:

→ 6G in 2030, democratized LEO satellite
→ Permanent work from home, emerging digital nomads

Simulation: "What happens if Starlink captures 15% of the rural market in 3 years?"

Scenario A (Defensive): Cut rural pricing to compete

→ Outcome: -\$45M revenue, retain 80% rural share

Scenario B (Offensive): Exit rural, focus on urban premium

→ Outcome: -\$30M rural revenue, +\$80M urban premium

Scenario C (Associative): Partnership with Starlink (rural wholesale)

→ Outcome: -\$15M rural margin, +\$20M new services

Recommendation to the Board: Scenario B + hedge with C

Typical example (Municipality):

National Network of Municipal Agents:

National Municipal Ecosystem Digital Twin:

→ Models: 345 municipalities with their capabilities
→ Central government (SUBDERE, SENAPRED)
→ Shareable resources (ambulances, equipment, personnel)
→ Emergency patterns (20-year historical record)

Simulation of coordination in a catastrophe:

→ 7.8 earthquake, central zone
→ DT simulates the coordinated response of 345 municipalities
→ Automatically optimizes resource sharing

Diagnostic question: *Do you have agents that participate in strategic decisions with a digital model of*

the complete ecosystem — competition, market, regulation — and long-term future simulation?

Digital Twins as a Diagnostic Dimension

Digital Twins emerge as a cross-cutting maturity indicator that evolves progressively. Their presence and sophistication are a powerful diagnostic signal of the level of automation.

Level	Type of Digital Twin	Scope
1-3	Does not exist	N/A
4	Descriptive	Infrastructure/operations
5	Predictive/Prescriptive	Specialized technical domain
6	Organizational	The entire organization
7	Strategic, of the Ecosystem	Org + competition + market

Descriptive DT (Level 4)

A digital model synchronized in real time. It answers: what do I have? what is the current state? what happened? Monitoring, correlation, anomaly detection. It does not predict or simulate.

Predictive/Prescriptive DT (Level 5)

An extension of the descriptive one with the ability to anticipate and simulate. It answers: what is going to happen? what happens if I do X? what is the best action? It requires proprietary data (5-10 years) and generates a sustainable competitive advantage.

Organizational DT (Level 6)

A model of the entire organization: people, finances, operations, interdependencies. It answers: how do I optimize resource allocation? what impact does this decision have on the whole org?

Ecosystem DT (Level 7)

A model of the complete ecosystem: competitors, regulators, market, macro forces. It answers: what strategic opportunities exist? what happens if competitor X does Y?

Differences by industry

The construction of Digital Twins varies radically depending on the maturity of existing systems. This has important diagnostic implications:

Industries with mature systems (Telecom, Energy, Manufacturing): - Level 4 can be built *on top* of existing systems. In telecommunications, the OSS (Operation Support Systems) is already essentially a Descriptive DT — the investment is adding the agent layer, not building the model from scratch. - Level 5 feeds on decades of operational data already curated. - The diagnosis must assess: which legacy systems can be turned into a DT?

Emerging industries (Smart Cities, Retail, Services): - Level 4 requires building the DT from scratch. No mature “municipal OSS” exists. - The investment and complexity are significantly greater. - The opportunity: designing from scratch with no technical debt. - The diagnosis must assess: what data infrastructure exists as a foundation?

Concrete example:

Telecom reaching Level 4:

- OSS already exists as a Descriptive DT
- Add the agent layer on top of the existing foundation
- 10+ years of data available

Smart City reaching Level 4:

- A complete urban infrastructure DT must be built
- Integrate IoT sensors, municipal systems, historical data
- The data may not exist or may be fragmented

Evaluation Dimensions by Level

	1. Invisible	2. Trusted	3. Collaborator	4. Analyst	5. Specialist	6. Manager	7. Authority
Infrastructure & Agents	BYOA, no control	Enterprise with governance	Integration with systems	Coordinated multi-agent + Descriptive DT	Specialist agents + Predictive DT	Organizational DT	Ecosystem DT
Operational Autonomy	AI as risk	AI as copilot	Task delegation	Autopilot 24/7 (>95%)	Predictive autonomy	Management authority	Strategic authority
People & Culture	Individual users	Governed users	Task delegators	Exception supervisors	Specialization designers	Report to agents	Co-leadership with agents
Governance	No governance	Policies + audit trails	Delegation matrices	Multi-agent protocols	Prediction governance	Org-authority governance	Ecosystem governance
Digital Twins	N/A	N/A	N/A	Descriptive	Predictive/Prescriptive	Organization	Ecosystem
Business Value	Uncaptured individual productivity	Risk reduction, basic efficiency	Time reduction on tasks	Continuous operation, systemic efficiency	Competitive advantage (moat)	Organizational optimization	Strategic opportunities

Market Distribution

Current state (2025): - ~70% at Level 1 (Invisible / Shadow AI) - ~25% at Level 2 (Trusted) - ~4% at Level 3 (Collaborator) - <1% at Level 4+ (Analyst+)

2030 projection: - ~10% at Levels 1-2 (laggards) - ~30% at Level 3 (Collaborator) - ~50% at Level 4 (Analyst) ← The new competitive standard - ~10% at Level 5+ (industry leaders; within this group, <0.1% at Level 7, experimental)

Level 4 (Analyst) will be the market's new baseline by 2030 — the minimum required to compete. Speed of arrival determines competitive advantage. Level 5 (Specialist) will be the sustainable differentiator for those with proprietary data and the capacity to invest.

These estimates are referential and based on observation of the market at the time of publication.

Rapid Assessment

8 questions to locate your level:

1. **Do you have a complete inventory of the AI tools in use?** No → Level 1 · Yes, with governance → Level 2+
 2. **Do you have formal AI usage policies with active enforcement?** No → Level 1 · Yes → Level 2+
 3. **Is Trust Infrastructure in place (tokenization, audit trails)?** No → Level 1 · Yes → Level 2+
 4. **Do you delegate complete end-to-end tasks to agents?** No → Level 2 · Yes, with a fixed workflow → Level 3 · Yes, with autonomy → Level 4+
 5. **Do you have a Descriptive Digital Twin of infrastructure/operations?** No → Level 3 or below · Yes, complete and real-time → Level 4+
 6. **Do agents operate 24/7 with multi-agent coordination?** No → Level 3 or below · Yes → Level 4+
 7. **Do you have specialist agents with predictive/simulation capability?** No → Level 4 or below · Yes, with a Predictive DT → Level 5+
 8. **Do you have an Organizational or Ecosystem Digital Twin?** No → Level 5 or below · Organizational DT → Level 6 · Ecosystem DT → Level 7
-

How to use this model?

For diagnosis: Identify the level that best describes the organization's current state in each dimension. An organization can be at different levels depending on the dimension — for example, at Level 3 in Operational Autonomy but at Level 1 in Governance. The overall level is determined by the lowest dimension, since it represents the real bottleneck.

For planning: The model makes it possible to identify gaps and prioritize. If the organization is at Level 3 in operations but at Level 1 in governance, governance is the bottleneck — no amount of automation matters if there is no governance to sustain it. The specific prescription of what to implement and in what order is a task for the consulting engagement that starts from this diagnosis.

For executive communication: The 7 levels offer a shared vocabulary. “We are at Level 2, moving toward Level 3” is a sentence a board understands — one that enables strategic decisions on investment and priority.

Together with IRIS: The complete diagnosis of organizational transformation requires both models. IRIS assesses the KNOWING axis (organizational intelligence). MOTOR assesses the DOING axis (process automation). An organization with high IRIS and low MOTOR knows everything but does little. One with high MOTOR and low IRIS automates blindly. The upper-right quadrant — high intelligence, high automation — is the destination.

Quick Reference: The 7 Levels

Quick-lookup table; the detail of each level lives in its section.

#	Name	In one sentence	Who executes	Next transition
1	Invisible	AI without governance or visibility	Humans (AI is a hidden risk)	→ Trust Infrastructure
2	Trusted	Governed AI as copilot	Humans with controlled assistance	→ Wingworking
3	Collaborator	Delegation of complete tasks	Humans delegate, agents execute under approval	→ Agentic Infra + Descriptive DT
4	Analyst	Autonomous 24/7 operation with a Descriptive DT	Agents (human only on exceptions)	→ Marketplace + Predictive DT
5	Specialist	Proprietary predictive capabilities (moat)	Specialist agents with simulation	→ Autonomous Org + Organizational DT
6	Manager	Organizational authority with a complete DT	Agents manage humans and resources	→ Social Agentic + Ecosystem DT
7	Authority	Strategic partner with an ecosystem view	Agents in leadership and inter-org networks	—

Glossary

Agentic Infrastructure. The set of protocols and services that allow multiple AI agents to communicate, coordinate, and execute tasks autonomously. Diagnostic signal of the 3→4 transition.

Autopilot. An operating mode in which agents execute tasks autonomously 24/7 without real-time human supervision, with intervention only on exceptions. Defining characteristic of Level 4.

BYOA (Bring Your Own Agent). Observable pattern of Level 1 in which each employee chooses and uses their own AI tool without organizational coordination.

Digital Twin (DT). A synchronized digital model of a system, organization, or ecosystem. A cross-cutting enabler that evolves from Descriptive (Level 4) to Strategic, of the Ecosystem (Level 7).

Descriptive DT. A Digital Twin synchronized in real time that monitors the current state and historical data, with no predictive capability. Diagnostic signal of Level 4.

Predictive/Prescriptive DT. A Digital Twin able to predict events and simulate scenarios. Requires proprietary data (5-10 years). Diagnostic signal of Level 5 and a source of moat.

Organizational DT. A Digital Twin that models the entire organization (people, finances, operations, interdependencies). Diagnostic signal of Level 6.

Ecosystem DT. A Digital Twin that models the complete ecosystem (competition, regulators, market, external forces). Diagnostic signal of Level 7.

Moat (Sustainable Competitive Advantage). A proprietary capability that is hard to replicate. In MOTOR, it emerges at Level 5 when proprietary data (5-10 years) feeds specialist agents with predictive capability.

Autonomous Organization. An organizational model in which agents manage processes and people based on objective metrics and simulations from the Organizational DT. Diagnostic signal of the 5→6 transition.

Quick Win. An AI use case with high visibility, low technical complexity, and fast ROI (typically under six months), selected to build momentum and validate capability before larger investments. An operational concept applicable at any MOTOR level, particularly useful in early transitions (1→2, 2→3) where the organization needs visible evidence of return before committing larger budgets.

Shadow AI. Unmanaged use of AI tools without governance, compliance, or organizational visibility. Defining state of Level 1.

Social Agentic. An ecosystem in which agents from different organizations collaborate with verifiable digital identities. Diagnostic signal of the 6→7 transition.

Agent Marketplace. An ecosystem in which agent capabilities are acquired, shared, or traded between organizations. Diagnostic signal of the 4→5 transition: the organization stops building every capability internally and starts composing with specialized capabilities from the market.

Sweet Spot. The maturity level that maximizes the balance among return, risk, and investment. In MOTOR, the Sweet Spot is at Level 4 (Analyst): it is the point where the organization captures most of the economic value of operational autonomy without taking on the cost and complexity of Levels 5-7. By 2030, Level 4 will be the market's new competitive baseline.

Trust Infrastructure. The set of technologies (tokenization, anti-prompt injection, audit trails, automated compliance) that make it possible to use AI in a governed, auditable way. Diagnostic signal of the 1→2 transition. It is the same primitive that *AgencyDomains* specifies formally in its canon; here it is used as a maturity threshold, not as a spec.

Wingworking. A human-AI collaboration methodology in which the human acts as the “pilot” (strategic decisions, validation) and the AI as the “wingman” (delegated tactical execution). An original ultraBASE practice. Diagnostic signal of the 2→3 transition and defining characteristic of Level 3.

Cross-Reference

MOTOR is the measurement instrument of the DOING axis — organizational automation through AI. Together with **IRIS** (Organizational Intelligence Maturity Model), it forms the diagnostic pair of **AURA** — ultraBASE’s agentic enterprise architecture.

The foundational vision of the architecture is the Introduction of this book.

This page intentionally left blank.

Wingmap

An agentive-discovery technique for reconstructing, without interviews, an organization's processes and information flows

MOTOR has just located the DOING level; what is missing is the terrain on which that doing happens — the real processes and information flows, which the organization cannot articulate on its own. This chapter formalizes the technique that reconstructs them without asking for them. Like Data Canon, it was born as a standalone paper and keeps that register — abstract, conceptual framework, principles, phases.

Abstract

Modern organizations operate faster than they can articulate themselves. Twenty years of Business Process Management proved, against its own promise, that formally documenting processes in order to automate them is an undertaking that most dynamic industries cannot sustain: the process-modification cycle outpaces the documentation-adjustment cycle, and the organization ends up preferring not to control its processes rather than pay the toll. The agentive era opens a structural way out of that asymmetry — not by building the documentation faster, but by ceasing to ask for it.

This chapter formalizes **Wingmap**, an agentive-discovery technique that reconstructs, from the digital traces the organization already generates, **two living, connected maps**: the map of the process each area executes and the map of the information flowing through those processes. The technique operates in two complementary modes — passive observational and active Wingworking — and produces a cognitive substrate of organizational intelligence that is at once an immediate analytical deliverable and a precondition for operational agents at runtime. The chapter articulates the conceptual framework, the six guiding principles, the operational components, the five phases of application, the zones of non-applicability, and the distinctions from adjacent categories (classical process mining, individual copilots, AI-builders, traditional consulting). The technique is independent of any specific implementation, but it belongs to a horizon in which the next layer of enterprise software ceases to be built and comes to be instantiated by agents with access to a persistent, validated model of how the organization operates.

Introduction

The universal problem

Every organization operating beyond a certain threshold of complexity faces an uncomfortable asymmetry: it operates far faster than it can articulate itself. Processes change with the business — new products, new regulations, new geographies, new customers — and the changes accumulate without passing through formal documentation, because formal documentation demands a cost the operation has no way to absorb without sacrificing other tasks. The aggregate consequence is that the organization runs on a real version of its processes that lives in the heads of those who execute them, in scattered email threads, in parallel spreadsheets, in meetings without minutes, in Teams messages. The documented version — when it exists — is an outdated shadow of the real one.

This asymmetry is not accidental. It is a structural consequence of how process-documentation systems were built over the last two decades. The promise of Business Process Management was precisely to close that gap: if processes lived on a single platform, modeled in standard notation, executed by an engine that knew the rules, the change cycle would be trivial — modify the model, redeploy, done. Reality turned out otherwise. Every modification of the model demanded an IT project to adapt the screens humans used when executing the modified process. The presentation layer never updated automatically with the model layer, and so the effective change cycle kept being measured in projects, not in configurations. Dynamic industries — telcos, digital banking, fintech — discovered that the marginal cost of changing a process under BPM was prohibitive, and they responded rationally: they stopped modeling their processes formally, or modeled only the most stable ones. Static industries — discrete manufacturing, traditional retail — could use BPM and SAP precisely because they had renounced being dynamic. The category's universal promise was fulfilled in the wrong half.

The agentive opportunity

The agentive era reopens the problem on a new plane. The new capability is not the general intelligence of the models nor their reasoning throughput — it is the displacement of the interface layer from *persistent artifact that must be designed, programmed, maintained* toward *situational event that the agent instantiates on the spot*. If the interface ceases to be pre-created, the change cycle that killed BPM collapses. What used to be a months-long project becomes a conversation.

That structural opportunity admits two readings. The first, which today occupies most of the market's attention, is that the agentive era allows enterprise software to be built much faster — AI in the builder, deterministic software at runtime. That reading is legitimate and produces commercially viable products. The second reading, less visible but deeper, is that the agentive era finally allows the dependence on pre-created enterprise software to dissolve: if agents can instantiate interfaces and orchestrate processes at runtime, what the organization needs is not faster software but a **persistent, validated model of how it operates**, on which the agents then operate. The investment shifts away from the construction layer and into the organizational-intelligence layer.

Wingmap belongs to that second reading.

The hypothesis

The hypothesis this chapter formalizes has three chained components. First: actionable organizational intelligence is not obtained by asking the humans who execute the processes, because those humans — for reasons the section «The human bottleneck of discovery» develops — cannot articulate in the ab-

stract what they execute in the concrete. Second: that intelligence *can* be reconstructed by observing the digital traces the organization itself leaves while executing its processes, because in a contemporary organization with a modern stack (Microsoft 365, Google Workspace, Slack, cloud ERPs, connected transactional systems) practically all significant activity passes through recordable digital channels. Third: the reconstruction must be *dual* — of the process (how the work is executed) and of the information flow (how the data travels through that work) — because without the informational dimension the documentation remains incomplete and, worse, remains incomplete exactly where the modern organization most needs understanding: in the origin-transformation-consumption chain of corporate data.

Wingmap is the technique that materializes that hypothesis: a reproducible practice for reconstructing, through an AI agent that observes corporate digital channels, the two planes of the operational model — process and information flow — without asking the organization to articulate them. The rest of the chapter develops it.

Conceptual framework

The pre-agentive / agentive distinction

(The terminology is the book's: the Introduction reserves "agentive" for the incremental evolution — copilots inside the same old applications. The pre-created software topology this chapter calls pre-agentive is the technical counterpart of that world; the runtime-instantiated topology is the technical counterpart of the agentive world.)

It is worth fixing terminology. **Pre-agentive** designates the set of process- and information-management technologies that dominated between 1995 and 2025: classical BPM, monolithic ERPs, document-workflow suites, orchestration platforms, RPA, traditional low-code. What distinguishes this set is not its age or its technical capability — several categories remain competitive — but its fundamental topology: the software is built before it operates, it operates with pre-created interfaces, and the process change cycle requires a software change cycle. The logic is that of the *persistent artifact*.

Agentive designates a different topology in which the software is built, in substantive part, *during* operation. Interfaces are instantiated on the spot, according to context. The software's executive capability emerges from the conversation between the agent and the human, not from pre-designed screens. The process change cycle no longer requires a software change cycle — it requires updating the underlying model and nothing more, because the interface regenerates on the fly. The logic is that of the *situational event*.

The distinction is not absolute — real products mix both topologies in varying proportions — but it is discriminative for classifying what each one does with the agentive era. A product that uses AI in the builder and delivers software with a pre-created interface is agentive to build, pre-agentive to run. A product that uses AI in both modeling and execution is agentive end to end. Wingmap belongs to an intermediate category with a fully agentive vocation: agentive to discover (it builds the model of the organization by observation) and agentive to run (its output is designed to be consumed by operational agents at runtime, not by pre-created software).

The human bottleneck of discovery

The empirical premise underpinning Wingmap is the one Polanyi formulated as tacit knowledge and that any process consultant confirms with accumulated operational evidence: *we know how to do more than we can say how we do*. People who execute processes correctly — even virtuosically — cannot reconstruct, in an interview session or a modeling workshop, the structured totality of what they do. The reason is threefold.

First, the abstract articulation of an embodied practice is a cognitive act distinct from the act of executing the practice. Asking the executor to describe the process is asking them to switch cognitive modes, and the most honest answer is always incomplete. Second, modern organizations operate under continuous BAU (*business as usual*) pressure that consumes the executor's reflective capacity; the discovery interview competes with tasks the executor cannot stop doing, and the rational response is to minimize the cost of the interruption, not to maximize fidelity. Third, every description of a process is politically situated — which exceptions get acknowledged, which inefficiencies get justified, which decisions get attributed to which area — and the executor presents a version optimized for the interviewer and for their own professional narrative, not the real version.

The aggregate evidence from discovery projects over the last two decades confirms this triad. In the witness case in which this chapter is set — a discovery program at a multinational agroindustrial holding — eleven Quick Win initiatives designed with low declared effort for the organization and facilitated by the consulting team produced business information in a minority of cases. That proportion is no local accident; it is the structural pattern of any dynamic organization with saturated BAU. The discovery technique based on asking has a systemic ceiling that no marginal improvement of the questionnaire or the facilitator can overcome.

Wingmap starts from the conclusion opposite to the one traditional consulting keeps drawing. The traditional conclusion is *we must improve the questioning techniques and the workshop formats*. Wingmap's conclusion is *we must stop asking and start observing*.

The artifact taxonomy: systematic and contextual

The Wingmap technique operates on a taxonomy of information artifacts worth fixing before the operational formulation. An artifact is **systematic** if all of its information can be generated from sources already digitized in transactional systems without additional human intervention. An artifact is **contextual** if it depends on a person contributing information periodically — entering a figure, completing a spreadsheet, documenting a decision, attaching a file. The distinction is operational, not technical: it is not about where the data lives, but about whether the flow that feeds it requires recurring human action to stay current.

This taxonomy matters for Wingmap for two reasons. First, systematic artifacts are faster to discover through passive observation — their traces live in systems the agent can read in structured form. Second, contextual artifacts are where traditional discovery fails most severely — because the data owner is exactly the person saturated with BAU. The asymmetry suggests that an observer agent has high marginal yield on contextual artifacts and moderate marginal yield on systematic ones, and that the prioritization of Wingmap deployment should follow that logic.

The organizational-intelligence layer

It is also worth situating Wingmap within a map of functional layers of agentive enterprise software. (The term “organizational intelligence” is used here in a sense complementary to the KNOWING axis:

not the *capacity* IRIS measures, but the **software layer** that materializes it — the place in the stack where the model of how the organization operates lives.) There are at least four distinguishable layers:

Layer	Function	Characteristic output
Individual assistance	Accelerate the human in their task	Summaries, drafts, contextual suggestions
Organizational intelligence	Model how the organization operates; keep the model alive	Process map, information-flow graph, business semantics
Software construction	Produce code that automates specific tasks	Custom internal software, deterministic bots
Agentive runtime	Agents that execute processes in production	Adaptive human-agent interaction, with instantiated interfaces

Wingmap sits explicitly in the second layer — organizational intelligence — and produces the input the third and fourth layers require to operate with knowledge of the organization. The individual-assistance layer (Microsoft Copilot, ChatGPT, Claude) operates on the individual user’s knowledge; the software-construction layer (products like Pit; AI-assisted coding like Cursor or Devin) operates on the declarative description the organization or the builder gives it; the agentive-runtime layer (Salesforce Agentforce, Microsoft Copilot Studio, emerging operational-agent platforms) assumes the organization knows the process when building the agent. Wingmap is the missing piece: the one that delivers the organizational model as a separable, governable, versionable artifact.

Definition and principles of the technique

Definition

Wingmap is a technique for the agentive discovery of an organization’s operational intelligence. Through an AI agent with read access to the organization’s corporate digital channels — email, messaging, file repositories, calendars, transactional systems — the technique reconstructs, without mediating interviews or workshops, **two living, connected maps**: the process map (how the work is executed in each area, in what sequence, with which actors) and the information-flow graph (which data originates where, who transforms it with what calculation, in which artifact it lives, toward which consumer it travels, which informational dependencies accumulate across areas, countries, or business lines).

The two artifacts are connected by common traceability to documentary evidence — every element of the map cites the raw documents from which it was inferred — and they are kept up to date by continuous observation, not by periodic re-discovery. Human validation occurs in distinct roles: process ambassadors for the first map, operational data owners for the second. The technique operates first in passive mode (observation transparent to the user) and, when conditions allow, in active mode (Wingworking — the user operates via the agent, the agent learns and, eventually, automates recurring patterns with deterministic bots).

The six guiding principles

The technique is governed by six principles that order its design and application decisions. The principles are not axioms — they are operating guides that, together, distinguish Wingmap from other discovery or modeling techniques.

Principle 1 — Observation first, intervention after. Discovery begins by observing the digital traces the organization already generates, without asking it to produce any additional information. Intervention — interview, workshop, validation — happens on already-inferred material, not on a blank page. This inverts the traditional discovery flow (ask → write → validate) and replaces it with (observe → infer → validate).

Principle 2 — Dual plane, not one. The canonical output is always dual: process map *and* information-flow graph. Omitting either of the two planes breaks the value proposition. A process map without an informational graph documents the choreography of the work but not the circulation of the data; an information graph without a process map documents the circulation of the data but not the choreography that sustains it. The two planes together produce actionable understanding; each one alone produces an incomplete portrait.

Principle 3 — Common traceability to evidence. Every assertion in either of the two maps must be backed by a traceable reference to the raw document from which it was inferred — the specific email, the specific Teams message, the specific version of the spreadsheet, the specific transaction in the system. This technical constraint is what makes the output governable: every element can be audited, refuted, corrected, or validated against the evidence that originated it. Without traceability, the agent's output is indistinguishable from a plausible hallucination.

Principle 4 — Distributed validation in operational roles. Human validation of the output is not performed by a central authority — it is performed by the operational roles closest to the element being validated. A map of the laboratory's process is validated by the laboratory's ambassador; an information flow over dispatch data is validated by the operational owner of dispatch data. This distribution recognizes that the model's fidelity rests on the sum of many bounded validations, not on a global validation by a distant committee. The distribution also reduces the cognitive cost per validator and enables scaling.

Principle 5 — Passive mode before active. The first deployment is always passive: the agent observes without intervening in the user's workflow. Active mode (Wingworking, where the user operates via the agent) is introduced only in specific domains, with enthusiastic ambassadors, and after passive mode has demonstrated value in that domain. The reason is that active mode carries a real adoption cost — the user changes their workflow — and that cost is only bearable once the user has already seen value from the agent in passive operation. Inverting the order produces abandonment.

Principle 6 — Produce substrate, don't consume requirements. Wingmap's output is a cognitive substrate — a persistent, validated model of the organization — that can afterwards be consumed by multiple destinations: specifications for the data factory, configurations for BPM, prompts for operational agents, governance dashboards, executive briefings. The technique does not consume a requirement from the organization in order to produce a specific deliverable; it produces a substrate that the organization and its collaborators exploit in multiple directions as needed. This difference distinguishes Wingmap from discovery techniques oriented to a single deliverable — a BPMN map to implement BPM, a functional specification to build software, an assessment to diagnose maturity. Wingmap produces the substrate common to all those uses.

Operational components

The observer agent (passive mode)

The observer agent is the nuclear component of Wingmap. Its function is to continuously ingest the events the organization generates in its corporate digital channels, extract semantic structure from each event, and keep up to date the dual graph (process + information) that constitutes the output.

The observation scope is explicitly bounded at the start of the deployment. In the reference implementation, the default scope is Microsoft 365 — Exchange, Calendar, Teams, SharePoint, OneDrive — plus connectors to the relevant transactional systems (ERPs, vertical platforms). The bounding occurs along two dimensions: by operational domain (a specific area at the start, expansion later) and by channel type (some channels are excluded for sensitivity, for example HR folders or union channels). The bounding is a governance component, not an implementation detail.

For every ingested event, the agent extracts a structured object with predefined fields: actor or actors involved, inferred intention, business domain, referenced artifacts, preliminary hypotheses about the process the event belongs to, preliminary hypotheses about the information flow it contributes to. Extraction proceeds in layers — cheap heuristics filter noise, light models classify, powerful models reason over the relevant cases — to keep the computational cost within a reasonable budget. The extracted objects accumulate in structured storage and feed, at a configurable frequency, the dual-graph builder.

The dual-graph builder

The second component takes the enriched events and turns them into structured hypotheses about the organizational model.

For the **process plane**, the builder identifies recurring patterns of interaction among actors — email chains with the same structure, recurring meetings with stable attendees, file-editing sequences — and applies process mining algorithms (alpha miner, heuristics miner, inductive miner) over the enriched event log to infer the graph of activities, decisions, and temporal dependencies. The result is serialized as standard BPMN 2.0, which facilitates its consumption by orchestration engines, modeling tools, and the organization's process catalogs.

For the **information-flow plane**, the builder identifies nodes of three types — people, artifacts (Ex-cels, reports, dashboards, documents), processes — and edges connecting them by semantic relation: *originates*, *transforms*, *consumes*, *derives from*, *canonical source of*. Constructing this second graph is structurally more demanding than the first because it requires semantic extraction of the content, not just of the event — knowing which data lives in which cell of which Excel, which calculation transforms it, what relation it bears to the analogous data in another similar Excel but in another country. The graph is serialized as JSON-LD or another labeled-graph format and is preserved in graph storage (Cosmos DB with Gremlin, Neo4j) for query via API.

The two graphs share a common traceability mechanism: every node and every edge carries references to the set of raw events from which it was inferred, with a confidence field and an evidence snippet. Traceability allows the human validator (process ambassador or data owner) to consult the original material when reviewing each element.

The validation interface

The third component is the interface through which the human validators confirm, correct, or reject elements of the dual graph. The interface is neither optional nor secondary — it is the mechanism by which the agent’s output becomes governed knowledge instead of remaining a plausible hypothesis.

The interface has two modalities — one for the process ambassador, another for the data owner — because the questions each role answers are different. The process ambassador answers “*is this sequence recognized as my area’s process?*”; the data owner answers “*is the origin-transformation-consumption chain of this data correct?*” and “*is the canonical source correctly identified?*”. Both validations are stored as ground truth and feed the agent to refine future inferences in the same domain.

The active mode (Wingworking)

The fourth component is optional and selective. In domains where an ambassador user agrees to change their workflow — operating the applications through the agent instead of directly — the agent learns, session by session, the user’s operating patterns in high resolution: what they query in what order, what transformations they execute, what corrections they apply. When a pattern becomes stable, the agent can generate a deterministic bot that automates the pattern without requiring an LLM at runtime.

The active mode extends Wingmap’s output from *organizational model* to *organizational model + library of deterministic bots*. Bot generation is governed with DevOps discipline — versioning, testing, rollback — because the bots are code in production. The name comes from Wingworking — the practice of working with an agent as a wingman that MOTOR presented as the signal of Level 3 —: the operation is performed alongside the agent, neither in its substitution nor at its command.

The governance layer

The fifth component cuts across the previous ones. It covers access control (which channels the agent observes, which it excludes, who authorizes), audit log (an immutable record of every event that touches the organization’s data), PII management (detection, tokenization, or redaction according to policy), retention policy (how long raw data, enriched data, and embeddings are kept), and regulatory compliance (GDPR, local personal-data-protection rules, labor regulations on communications monitoring).

Governance is neither optional nor post-hoc. Its design must occur before the first deployment, with the informed consent of the observed actors and legal validation specific to the organization’s jurisdiction. The section «Privacy risks and regulatory compliance» develops the consequences of failing at this component.

The practice in phases

The application of Wingmap is ordered into a prior enablement phase (Phase 0) and five operational phases, with explicit criteria for moving from one to the next. Each phase produces measurable deliverables and admits criteria to proceed, pivot, or abort. The sequence is deliberate — skipping phases produces fragile deployments — but the phases may iterate in parallel within distinct domains of the organization.

Phase 0 — Preconditions

Before deploying anything technical, the practice demands five clear preconditions.

Executive sponsor identified and committed. A person with authority over the observation scope who understands what they authorize and what they do not, and who takes political ownership of the decision. Without a sponsor with real power, the deployment stalls at the first access request.

Bounded and consented scope. The first deployment covers a specific operational domain — an area, a business unit, a geography — with an explicit perimeter of digital channels. The bounding protects the organization (it limits privacy risk) and the deployment (it limits the error surface and enables fast iteration).

Process ambassador designated. A person from the observed domain, with operational knowledge of the process, willingness to validate maps, and sufficient political capital to sustain the experiment before their peers. Without an ambassador, the inferred maps go unvalidated, and without validation the output remains a hypothesis with no actionable value.

Data owners identified. For each information flow the agent will observe — sales data, operations data, HR data — there must be a person recognized as the operational owner of the data, willing to validate the corresponding informational graph. This condition frequently exposes preexisting governance gaps in the organization — data with no clear owner — that the technique forces to be resolved at the start.

Legal and compliance framework validated. A legal opinion, adjusted to the organization's jurisdiction, on what the agent may observe and under what conditions. This component is not trivial and frequently demands weeks of prior work. Skipping it exposes the organization and the implementer to litigation or to sudden rejection of the deployment by internal compliance.

These five preconditions are **blocking**. If any is unmet, the deployment does not advance to the next phase; the gaps are closed in Phase 0 before touching technology.

Phase 1 — Observation

With the preconditions covered, the agent enters passive operation over the bounded domain. The recommended minimum duration combines two floors: never fewer than three iterations of the cycle of the domain's slowest process, and never fewer than two calendar weeks — the minimum time to filter pattern noise. In practice: for domains of daily cadence, two to four weeks; for weekly cadence, six to eight weeks; for monthly cadence, three months. This duration is not negotiable downward — inferring cyclical processes requires observing several iterations to distinguish pattern from noise. Skipping this floor produces spurious conclusions.

During the observation phase, the agent accumulates enriched events in structured storage and builds preliminary versions of the dual graph. The preliminary versions are not delivered — they are internal calibration material. What is delivered to the organization at the end of the phase is a coverage report: what percentage of the domain's digital flow was ingested, what recurring patterns were identified, what structural gaps exist (processes that appear to occur but whose digital traces are insufficient to reconstruct them).

Phase 2 — Inference and dual-graph construction

This phase consolidates the preliminary versions of the dual graph into reviewable versions. This demands work on the implementer's side (prompt calibration, schema refinement, tuning of process mining algorithms over the enriched log) that is not automatic: the quality of the output depends on sustained technical iteration.

The phase's exit criterion is the delivery to the organization of three to five candidate process maps and five to ten candidate cross-role information flows, all with traceable evidence and all with confidence estimated per element. The delivery does not claim final correctness — it aims to offer concrete material on which the following phase executes validation.

Phase 3 — Dual validation

This phase is where the agent's output becomes governed knowledge. Each process map goes through a validation session with its ambassador; each information flow goes through a validation session with its operational data owner. The sessions are designed with bounded time (60 to 90 minutes per session) and a concrete format (the interface described in «The validation interface», with evidence-traceable material visible to the validator).

Each session produces one of three outcomes. **Validation with minor correction** — the validator recognizes the element as their process or their data flow, adjusts details, signs off. **Validation with major correction** — the validator recognizes the general pattern but detects structural errors requiring re-inference; the agent reprocesses with the feedback as ground truth for the next attempt. **Rejection** — the validator does not recognize the element; the agent investigates whether the inference was spurious or whether the observed domain has processes the agent does not yet capture, and action is taken accordingly.

The phase ends when minimum validation thresholds are reached: 70% of process maps validated without major correction, 60% of information flows validated without major correction (the thresholds are specific to the first deployment; in successive deployments the thresholds rise with accumulated calibration). If the thresholds are not reached after two iterations, the phase reports a diagnosis — the problem may be one of inference (model) or of input (quality of the observed channels) — and the decision to proceed or abort is made by the sponsor with the implementer's advice.

Phase 4 — Maintenance cycle

This phase is sustained operation. Once the initial maps are validated, the agent keeps observing and the graphs live — they update with every iteration of the process, they are enriched with newly detected exceptions, they are refined with every subsequent validation. The organization gains continuous access to a model of how it operates, kept current, governed, and traceable.

This phase has its own key metric: **graph latency** — the difference between the last relevant observed event and the current version of the graph. Latency must be kept below an agreed threshold (typically 24 to 72 hours for an active domain). Growing latency indicates that the agent has lost the ability to keep pace with the organization, and is a sign that something in the infrastructure or in the model has broken.

Phase 5 (optional) — Selective active mode

The final phase, optional and selective, introduces the active Wingworking mode in domains where two simultaneous conditions are met: there exists a process of value where the passive modality has already identified a stable pattern, and there exists an ambassador user willing to operate via the agent instead of directly. Active mode is introduced only where both conditions hold, and one domain at a time. Mass generalization of active mode is contraindicated — the individual adoption cost is real and the learning curve requires accompaniment.

The product of active mode is the incremental library of deterministic bots, governed with the same discipline with which any of the organization's production code would be governed.

Discussion

Distinction from classical process mining

The Wingmap technique coexists with a mature discipline — classical process mining — and the differences are worth articulating with precision. Classical process mining (Celonis, Apromore, Disco, SAP Signavio, Microsoft Process Mining on Power Automate, UiPath Process Mining) operates on structured logs extracted from transactional systems: each event is a tuple (case_id, activity, timestamp, resource) deterministically extracted from the log of an ERP or a management system. The discipline has fifteen years of scientific literature — the work of van der Aalst and its derivatives — and tools with formally defined metrics (fitness, precision, generalization).

Wingmap does not replace classical process mining — it extends it to domains where classical process mining does not operate. The gap is one of nature, not of coverage. Classical process mining requires the process to live inside a system that produces structured logs with the canonical tuple. The greater part of an organization's real process does not live inside such systems — it lives in emails, in spreadsheets, in messages, in shared files. Wingmap captures that fraction of the process, infers its semantic structure with LLMs, and produces an enriched log on which the same classical process mining algorithms can — optionally — be applied. The traditional discipline gains a new domain of applicability; it does not lose its own.

The second difference is structural. Classical process mining produces an event graph. Wingmap produces an event graph *plus* a parallel information-flow graph. That second dimension does not exist in the traditional discipline — and not by omission but by impossibility: an ERP's structured logs record *that* the activity occurred, not *what data moved* in that occurrence. Wingmap, by operating on unstructured channels with semantic extraction, naturally captures the informational content. That difference is its structural differentiator.

Distinction from traditional interviews and process consulting

Classical process consulting solves discovery by sending teams of people — analysts, consultants, facilitators — to interview the organization, observe its operation, map flows in BPMN or equivalent, and iteratively validate the result. The practice is mature, has refined methodologies (Six Sigma, Lean, BPMN, BABOK), and Big4 frameworks that systematize it.

Wingmap differs along four measurable dimensions. The first is **speed** — an observer agent can ingest and enrich events at a continuous rate, while a human team operates in blocks of dedicated hours. For

a mid-sized domain, the difference is an order of magnitude (days vs. months). The second is **cost** — the agent’s computational cost, even at enterprise scale, is typically a fraction of the human cost for the same scope, also an order of magnitude. The third is **coverage** — the agent observes all of the domain’s digital activity, not an interviewed sample; the processes that human discovery would omit by chance (a low-frequency process that did not surface in the interviews) the agent captures through exhaustiveness. The fourth is **continuity** — human discovery produces a point-in-time portrait; the agent produces a living model that stays up to date.

In exchange, traditional consulting contributes two things the agent does not deliver by default: **strategic contextual interpretation** (understanding why a process evolved the way it did, what political tensions shaped it, which interventions have the best probability of success) and **change negotiation** (consultancies facilitate organizational transformations, they don’t merely document them). Wingmap does not compete with those two components — it complements them. The technique replaces discovery as a manual craft of interviews and leaves consultants free to concentrate on interpretation and change, which are their real value.

Distinction from individual copilots

Microsoft Copilot for Microsoft 365, ChatGPT, Claude, Gemini in Workspace are individual assistants that accelerate the human in their task — drafting a better email, summarizing a meeting, generating a formula. Their unit of analysis is one user’s session and their unit of value is individual productivity.

Wingmap is structurally different: it operates on the sum of many users’ tasks and reconstructs the cross-role organizational model. Its unit of analysis is the flow between roles; its unit of value is organizational understanding. An individual copilot cannot produce what Wingmap produces because its access is to the user’s session, not to the organization’s aggregate flow. The technique complements the copilot — an organization with Copilot deployed has part of the technical infrastructure Wingmap needs, but it does not have Wingmap.

Distinction from software AI-builders

The emerging category of *AI product team as a service* — products like Pit, publicly launched in May 2026 — uses AI to accelerate the construction of custom internal software that replaces the back office’s patchwork of Excel + email + rigid SaaS. The distinction from Wingmap is one of layer: Pit and its peers are **agentive to build, pre-agentive to run** — they use AI in the builder, they deliver deterministic software that operates with pre-created interfaces. Wingmap is **agentive to discover, agentive to run** — it produces the cognitive substrate that later enables operational agents at runtime with instantiated interfaces.

The two categories are complementary, not competitive. A mature organization uses them in order: first Wingmap to understand how it operates and what data circulates; then an AI-builder (or an internal team) to build tailored software in the processes where construction adds value; then operational agents that execute at runtime with knowledge of the model. The organizational-intelligence layer is qualified input for the construction and runtime layers.

The technique as substrate for the agentive transition

Wingmap is not fully understood as an isolated discovery technique — it is understood as a **foundational brick of the transition from the pre-agentive era to the full agentive era**. This reading widens the scope of the argument beyond the individual organization.

The full pre-agentive era (1995–2025) built enterprise software as a persistent artifact: pre-created screens, fixed logic, maintenance as a cycle parallel to operation. The full agentive era (~2030 onward, by reasonable projection) will build enterprise software as a situational event: agents that instantiate interfaces on the fly, execute processes at runtime with contextual interpretation, reconfigure themselves through conversation. The transition unfolds over a horizon of five to ten years, with both topologies overlapping throughout that period.

The operative question for an organization in transition is: what asset to build today that serves both the pre-agentive operation it already has and the agentive operation to come? The enterprise software built today will become obsolete when operational agents replace pre-created interfaces. BPM platforms will remain useful as orchestration engines but will lose their presentation layer. The safest investment, under this framing, is the **persistent, validated organizational model** — the dual graph Wingmap produces. That model does not become obsolete with the transition; on the contrary, it becomes more valuable, because it is the input that operational agents will need in order to execute with knowledge of the organization.

Wingmap, read this way, is not one more discovery technique among many — it is the technique that produces the substrate of the next decade of enterprise software. That is its universal thesis.

Limitations and zones of non-applicability

The technique has clear frontiers. They are worth declaring with precision so as not to sell promises that go unfulfilled.

Processes that live outside the digital trace

Wingmap captures what the organization leaves in its corporate digital channels. What happens on paper — the agronomist's field notebooks, the plant operator's manual annotations, physical forms — , in face-to-face conversations without minutes, in untranscribed phone calls, in informal WhatsApp groups not integrated into the corporate stack, falls outside the agent's reach. For an organization with a great deal of operational activity in these non-digital channels, Wingmap's effective coverage can fall significantly below the totality of the real process.

The fraction of digital coverage varies by sector. Administrative areas (finance, HR, procurement) may sit between 70% and 90% in digital channels. Commercial areas with intense client interaction may fall to the 40-60% range. Field operations areas (agronomists, plant operators, physical logistics) may be below 40%. The technique delivers its full value proposition in areas with high digital coverage and reduces its value proportionally in areas with low coverage. Recognizing this limit at the start of the deployment, choosing pilot domains with high digital coverage, is part of the discipline of application.

Organizations with low appetite for discovery

The technique demands explicit authorization to observe corporate channels. That authorization is a political decision, not a technical one. An organization where the sponsor lacks the will or the capacity to authorize the agent's access — out of cultural sensitivity ("they're watching us"), a restrictive regulatory framework, an active union, distrust of the implementer — cannot benefit from Wingmap. The technique is not universally applicable; it is applicable to the subpopulation of organizations with a real appetite for discovery.

This limitation has an operational corollary: the first deployment of Wingmap in an organization must always negotiate the scope with the sponsor before touching technology. Skipping the negotiation produces deployments that stall in security and compliance approvals.

Privacy risks and regulatory compliance

An agent with read access to corporate email, messaging, and files operates in a regulatory zone that is narrowing, not widening. The European AI Act (progressively in force 2024-2027) classifies as high risk those AI systems that process employee data to evaluate performance or make decisions affecting the organization of work; a process map whose use leads to functional reorganization can fall under that classification. The personal-data-protection laws of Chile, Mexico, Brazil, and other countries are converging with GDPR on standards of consent, right of access, and purpose limitation. Each jurisdiction's labor law restricts the monitoring of employee communications to varying degrees.

The responsible application of Wingmap requires specific legal counsel, informed consent from the observed actors, an explicit acceptable-use policy validated by HR and by legal, and a technical architecture designed with PII tokenization or redaction according to policy. Skipping this component produces legal exposure for the implementer and the organization, and can bring the project to a halt by court order or by internal compliance.

False positives and the validation paradox

The agent can infer processes that do not exist — two emails always sent on the same day may be two independent processes presented as one with two steps; a communication pattern between two people may be personal habit and not institutional process. Spurious inferences are inevitable; the only debatable points are their frequency and the validation discipline that detects them.

There is a known paradox in human-AI collaboration that deserves naming: when a human supervises an automated output, their error-detection rate drops systematically compared to when they produce the output from scratch. The meta-analysis by Goddard, Roudsari, and Wyatt (*Journal of the American Medical Informatics Association*, 2012) of 74 studies of medical decision-support systems showed that introducing the system reduces typical errors but introduces new errors when the system produces plausible yet incorrect outputs. Hurried validation of agent-inferred maps can ratify errors that human inference from scratch would have caught. The defense against this bias is discipline: time dedicated to validation, traceable evidence visible to the validator, quality-drift metrics that detect systematic drops in fidelity.

The human dependency shifts; it is not eliminated

The technique reduces the cost of discovery but does not eliminate human dependency. The bottleneck shifts from *describing the process from scratch* to *validating the inferred process*. The validator (process ambassador or data owner) is still a person the organization must make available, with dedicated time, political capital, and willingness to sustain iterations. In organizations where BAU completely consumes the time of the potential validators, the technique hits the same ceiling as the traditional techniques, albeit at lower total cost.

The difference is quantitative but structural: validating 50 maps in 60-minute sessions requires 50 hours; describing those same 50 processes from scratch in workshops requires 200-500 hours. The technique reduces the cost by a factor of four to ten. But it does not reach zero, and assuming it does leads to failed deployments.

Technological maturity at the time of application

The technique depends on language-model capabilities that in 2026 are solid but not perfect. Structured event extraction from free text has reported error rates in the 10-30% range in production, depending on domain and calibration. Semantic inference of processes from heterogeneous corpora is an open problem: the available evidence and deployment experience place inference quality (measured as F1) in ranges on the order of 0.4-0.7 in controlled scenarios and lower in production with real data and noise — the author's estimates, pending formal benchmark. These rates are adequate for validatable preliminary hypotheses, not for autonomous production without human supervision.

The trend is one of compounding improvement — the models advance year after year — but the improvement is not linear and does not arrive in time for any specific deployment. Responsible application of the technique acknowledges the current rates as the operating limit and designs the validation phase («Phase 3 — Dual validation») around them, not around the rates that will exist in five years.

Conclusion

Wingmap is a technique for the agentive discovery of an organization's operational intelligence. Its contribution consists in formalizing a reproducible practice that reconstructs, without interviews, two connected planes of the organizational model — process and information flow — from the digital traces the organization already generates. The six guiding principles (observation first, dual plane, common traceability, distributed validation, passive mode before active, substrate and not requirements) order the design and application decisions. The phases (preconditions, observation, inference, dual validation, maintenance, and — optionally — active mode) order the practice over time. The four distinctions (from classical process mining, traditional consulting, individual copilots, AI-builders) situate the technique on the map of adjacent categories.

The technique replaces none of the adjacent disciplines — it complements them structurally. And in doing so, it produces an asset the full agentive era is going to need as critical input: a persistent, validated model of how the organization operates, on which the operational agents of the 2030 horizon will be able to execute with situational knowledge.

The first witness case of Wingmap is underway at the time of writing this chapter: a multinational agroindustrial holding with dozens of companies operating in several countries. The lessons from that deployment will feed the next editions of this book. The predictions the chapter makes about coverage rates, inference quality, operating costs, and human dependency are reasoned hypotheses that the witness case will confirm or refute with operational evidence.

The invitation this chapter extends to other practitioners is not to use the specific Wingmap product — that is a separate commercial decision — but to apply the Wingmap **technique** in their own contexts, with their own tools, in their own organizations. The principles are implementation-independent. The practice is reproducible. The substrate it produces is valuable in itself, no matter who builds it or with what technical stack.

Whoever applies Wingmap in an organization does not document processes — they produce operational intelligence. The difference matters. Operational intelligence is a durable asset; process documentation, in the absence of intelligence, is paper.

Future work

Five directions remain open.

Empirical calibration of the application rates. The quantitative claims — coverage by sector, validation rates, graph latency — rest on limited data and reasoned projection, not on a robust empirical base. Three to five complete deployments with comparable metrics will allow moving from projection to statistics.

Study of the active mode’s learning curve. The active Wingworking mode demands a change in user behavior. The shape of the curve — how long it takes a user to overcome the initial friction, how many abandonments there are, what predicts persistence — is an open empirical problem that deserves systematic study.

Generalization to non-Microsoft stacks. This chapter assumes an implementation over Microsoft 365 as the initial spine. Application to Google Workspace, Slack-first, or heterogeneous stacks is feasible but requires additional components (specific connectors, different governance models) that deserve dedicated treatment.

Integration with agentive orchestration engines. The substrate Wingmap produces is designed to be consumed by operational agents. The formal specification of the APIs and exchange formats connecting the Wingmap graph with platforms such as Salesforce Agentforce, Microsoft Copilot Studio, or open frameworks (LangGraph, CrewAI) is pending architectural work.

Ethical grounding of organizational observation. The technique observes human activity in corporate channels. The ethical framework of that observation — beyond minimal regulatory compliance — deserves dedicated treatment: how transparency is reconciled with protection, how the model’s value is distributed between the organization and the observed actors, how the drift toward workplace surveillance is avoided. These questions are normative, not technical, and require dialogue with academia, unions, and labor authorities.

References

On this chapter’s sources

This chapter derives from an internal research cluster (16 working documents, May 2026) whose detail lives in the author’s archive. The verifiable references are listed below.

External literature

- Polanyi, Michael. *The Tacit Dimension*. University of Chicago Press, 1966.
- van der Aalst, Wil M.P. *Process Mining: Data Science in Action*. 2nd edition. Springer, 2016.
- Goddard, K., Roudsari, A., & Wyatt, J. C. “Automation bias: a systematic review of frequency, effect mediators, and mitigators.” *Journal of the American Medical Informatics Association*, 19(1), 2012.
- Rother, Mike & Shook, John. *Learning to See: Value-Stream Mapping to Create Value and Eliminate Muda*. Lean Enterprise Institute, 2003. (Methodological reference for information-flow maps in the Lean tradition).

- Hammer, Michael & Champy, James. *Reengineering the Corporation*. HarperBusiness, 1993. (Classic reference the chapter situates historically).

Cited market references

- Press coverage of Pit's public launch (techcrunch.com, sifted.eu, tech.eu, May 2026).
 - Public coverage of Microsoft Copilot, Microsoft Process Mining, Salesforce Agentforce, Glean, Celonis (as of May 8, 2026).
-

This page intentionally left blank.

The Portfolio

Knowing where you stand and knowing the terrain moves no one: what remains is deciding the next step. The third part of the path turns diagnosis into execution — the portfolio of use cases, classified along three dimensions and prioritized with a framework that makes the political decision an explicit one. The chapter defines its pieces — use cases (UC) and value solutions (VS), the three dimensions, and the Citizen/City/DUAL framework — and illustrates them with an inventory of one hundred real cases in Annex B. Here the path stops being described and starts being walked.

This page intentionally left blank.

Use Cases

From diagnosis to portfolio: how transformation is prioritized

Introduction · The Thesis

Diagnostic models answer *where* the organization stands; survey instruments answer *what* its real terrain looks like. A third question remains — the one that turns diagnosis into motion: **what do we do first?**

That question has a bad traditional answer — the list of initiatives assembled by intuition, where the manager’s pet project, the demo that dazzled at a trade show, and IT’s perennial backlog item all live side by side. And it has a disciplined answer: a **use case portfolio** — a structured catalog where every case is classified by required maturity level, by impact beneficiary, and by business domain, so that prioritization becomes an explicit strategic decision rather than an accumulation of whims.

This chapter defines the pieces of that portfolio: the distinction between **use cases** and **value solutions**, the **three classification dimensions**, the **Citizen / City / DUAL** prioritization framework with its two canonical strategies, and the discipline of honesty about ROI. It closes with a real portfolio — one hundred cases for local government — that shows the whole system in operation.

Use Cases and Value Solutions

The portfolio’s first distinction is one of granularity, and separating the two pieces avoids a mistake that ruins entire catalogs: confusing the brick with the building.

A **Use Case (UC)** is the atomic block: a specific, bounded, *reusable* capability. “Anomaly detection in expenditures,” “24/7 customer-service chatbot,” “automatic claim classification.” A UC is described by what it does, why AI does it better than the traditional alternative, and what maturity it demands of the organization. Its value lies in reuse: the same UC can take part in several different solutions.

A **Value Solution (VS)** is the strategic configuration: a set of integrated UCs that, together, solve a complete business problem and present themselves to a decision-maker as a unit with its own value proposition — problem, solution, metrics, return. “Autonomous financial management” is not a UC: it is a VS that integrates end-to-end budget management, financial analysis, continuous forecasting, and the automatic reallocation of budget lines.

The relationship is one of composition, not exclusive hierarchy: a UC can be sold on its own (modular sale, when the client wants a specific capability) or as a component of a VS (integral sale, when they want the turnkey solution). A healthy portfolio keeps both catalogs alive and linked — every UC declares which VSs it takes part in; every VS lists the UCs it integrates.

The Three Classification Dimensions

Every case in the portfolio is classified along three orthogonal dimensions. Together, the three turn a flat list into an instrument of decision.

Dimension 1 · Required maturity level

Every UC demands a minimum level of organizational maturity to operate — the MOTOR level of the organization able to execute it (and, depending on the case, the IRIS stage its information component presupposes). A chatbot running on masked data is executable at the early levels; an agent managing a budget with organizational authority demands the high levels; an inter-organizational network of agents, the highest level of all.

This dimension is what anchors the portfolio to the diagnosis: **the case's level is crossed against the organization's level**. Cases at the current level are executable today — that is where the *quick wins* live. Cases one level above are the reachable frontier — the *sweet spot* where ambition pulls maturity forward without breaking it. (This *sweet spot* is local to the portfolio — your current level plus one — distinct from the market Sweet Spot that MOTOR places at Level 4.) Cases several levels above are long-range vision: valuable for showing the destination, dangerous as an immediate commitment.

Dimension 2 · Impact beneficiary (Citizen / City / DUAL)

The second dimension classifies by *who perceives the benefit* — and it gets its own section below, because it is the political heart of the portfolio.

Dimension 3 · Business domain

The third dimension segments by the business's thematic area — external customer service, internal management and finance, operations and infrastructure, specialized services. Its function is practical: when the decision-maker has a focus (“my problem is security,” “my crisis is financial”), the portfolio must be filterable by domain in seconds and present the complete relevant offer — the domain's UCs, the VSs that integrate them, and the path of scale.

The Citizen / City / DUAL Framework

Any organization that simultaneously serves an external audience (the beneficiary of the service) and an internal audience (the operator) lives under **double pressure**: the external audience demands better service; the internal one demands operational efficiency. A use case may benefit one, the other, or both.

The framework classifies each case by who perceives the direct benefit:

Classification	Who benefits	Political/strategic function
CITIZEN	The external audience (citizen, end customer, service user)	Generates visible support and public legitimacy
CITY	The internal audience (the organization, its team, its financial structure)	Generates efficiency, compliance, sustainability
DUAL	Both simultaneously	Generates consensus: every relevant actor gets a visible win

The framework was born in local government — hence its names — but the structure is generic: in retail banking, CITIZEN is the customer and CITY is compliance and operations; in healthcare, the patient and the clinical teams; in telecommunications, the subscriber and network operations. In every case the logic is the same: the organization holds two simultaneous relationships that can align or pull against each other, and the transformation can prioritize the external side, the internal one, or the crossing point.

Why it cuts across IRIS and MOTOR: the framework does not classify by transformation axis (knowing/doing) but by impact beneficiary. A CITIZEN case can belong to the KNOWING axis (the customer queries their case in natural language) or to the DOING axis (an agent serves them 24/7). That is why the framework lives at the AURA level: a complete diagnosis assesses IRIS position, MOTOR position, and the Citizen/City/DUAL priority mix for the roadmap.

The two canonical strategies

When the roadmap is assembled, the framework translates into two archetypal strategies with distinct investment and return profiles.

Strategy A • Citizen-First. Prioritize cases visible to the external audience to win support, legitimacy, and momentum. It applies when leadership is under pressure from external stakeholders, when public trust is low and must be rebuilt, or when the political cost of not acting exceeds the financial cost of acting. Its pool is the CITIZEN cases plus the DUALs. Its risk: lower or diffuse financial ROI — legitimacy is gained without necessarily cutting operational cost.

Strategy B • City-First. Prioritize internal efficiency, savings, and compliance before exposing anything to the external audience. It applies in fiscal crisis, with overloaded teams that need relief before any expansion of service, or when the urgency is regulatory. Its pool is the CITY cases plus the DUALs. Its risk: the external audience perceives no benefit — the organization improves but fails to capitalize on it.

The common core — the DUAL cases. They appear in both strategies and are the optimal intersection: the external audience perceives them, the organization quantifies them, and against any critic they are defensible. Building the roadmap with maximum priority on DUAL cases is the lowest-friction play — and designing cases so that they *are* DUAL from conception is discipline, not luck.

Three usage notes close the framework. First: **do not force a strategy** — a large organization with budget can combine both; one in fiscal crisis probably needs pure City-First for the first year; the strategy is chosen by context, not by ideology. Second: DUAL cases are the portfolio's most expensive currency — they are scarce, and they are designed. Third: **the framework does not replace the diagnosis** — applying Citizen-First without having solved AI governance (the step from MOTOR Level 1 to 2) is

building visibility on sand; the framework prioritizes what to attack first *within* each level; it does not substitute for the maturity trajectory.

The ROI Discipline

A portfolio loses its usefulness the day the decision-maker stops believing it. And the fastest way to lose credibility is to present estimated returns as if they were measured.

The discipline has three rules. **First: every estimated ROI is declared referential** — visibly, without shame and without a hidden asterisk — until a validated baseline exists. **Second: social value is kept separate from financial return.** There are cases whose dominant benefit is intangible — school dropout detected in time, heritage preserved, public health under watch — and forcing them into a savings figure denatures them; the portfolio flags them as social value cases and defends them on those terms. **Third: the baseline is validated before any return is declared** — an independent baseline validation protocol is part of the portfolio, not a luxury added later.

To this is added the catalog's formal quality: consistent terminology, descriptions that distinguish *what the case does* from *why AI* does it better, and complete cross-references between UCs and VSs. A portfolio of one hundred cases is an editorial product — and it is governed as one, with its own quality framework and systematic reviews.

A Real Portfolio: One Hundred Cases for Local Government

The complete system — UC/VS, three dimensions, prioritization framework — is best illustrated by the portfolio where it was born: the agentic transformation inventory for municipal government, developed by ultraBASE over the course of 2025.

The scale: 100 atomic use cases and 30 value solutions, with an estimated (referential) aggregate return on the order of USD 118M/year for the whole, distributed across four business domains: citizen services and communication (17%), internal management — legal, HR, finance, procurement, and budget — (37%), infrastructure, roads, security, and emergencies (24%), and services — health, education, culture, sports, and environment — (22%).

The distribution by maturity tells the story of the complete transformation:

Maturity level	UC	Character of the cases
1–2 (governed, assistive AI)	34	Document analysis, chatbots on protected data, automatic auditing — the <i>quick wins</i>
3 (task delegation)	21	End-to-end public procedures, multi-project planning, complex multi-system cases
4 (autonomous 24/7 operation)	19	Real-time fraud detection, predictive patrolling, automatic emergency response

Maturity level	UC	Character of the cases
5 (in-house predictive capabilities)	9	Pavement deterioration, post-earthquake seismic assessment, traffic optimization
6 (agents with management authority)	6	Agents that manage teams of 25 staff, coordinate 45 simultaneous projects, administer a USD 45M budget
7 (inter-organizational ecosystem)	11	Regional security and waste networks, a marketplace of public agents, agents that propose policy

The pyramid speaks for itself: a third of the portfolio is executable at the early maturity levels — that is where momentum is built — while the tip describes a horizon where agents from different organizations cooperate in a network. And the distribution by beneficiary came out reasonably balanced — 31 CITIZEN, 38 CITY, 31 DUAL — which allows both canonical strategies to be served from the same catalog.

The flagship solutions display the composition pattern. The portfolio's top-ranked VS, *Autonomous Municipal Financial Management* — a level-6 solution — integrates four UCs, among them the agent that manages the entire budget (level 6) and the one that reallocates budget lines automatically (level 6), supported by financial analysis and forecasting capabilities from earlier levels. The second, *Paperless Municipality* — a level-3 solution — integrates five UCs for the full digitalization of the procedures the citizen actually perceives. That #1 is level 6 and #2 is level 3 is no contradiction: it is the portfolio doing its job — showing at once the ambition of the destination and the executable step of tomorrow.

What the case teaches transcends the public sector. First, that a portfolio is *curated*: this one went through seven version iterations and nearly five hundred quality improvements before being considered presentable. Second, that the classification dimensions are earned by iterating with the real decision-maker — the thematic area was born from a concrete commercial segmentation need, not from a desk exercise. Third, that the honesty of the referential ROI did not weaken the portfolio: it made it defensible.

How to use this framework?

1. **Diagnose first.** IRIS position, MOTOR position. Without this, the portfolio has no anchor.
2. **Inventory along the three dimensions.** Every candidate case is classified by required level, beneficiary, and domain. What cannot be classified is not sufficiently defined.
3. **Separate UCs from VSs.** The atomic blocks in one catalog; the strategic configurations in another; the cross-references, complete.
4. **Choose strategy by context.** Citizen-First, City-First, or mixed — an explicit leadership decision, revisable, with the DUAL cases always at the front.
5. **Attack the current level + 1.** Quick wins at the present level for momentum; the sweet spot one level up for traction; the rest of the portfolio as a declared horizon.

6. **Declare ROI honestly.** Referential until the baseline is validated; social value on its own terms.

The portfolio is the instrument that closes the AURA path: the vision defines the destination, IRIS and MOTOR locate the organization, Data Canon and Wingmap prepare the terrain — and the portfolio decides, with discipline, the next step.

Epilogue · The Complete Diagnosis

With the path traveled, the pieces fit into a single figure.

An organization that wants to make its way toward the Agentive World needs to answer three questions, and this book gave it an instrument for each one. **Where am I?** — the IRIS position on the knowing axis and the MOTOR position on the doing axis, two coordinates that together form AURA's diagnostic pair. **On what terrain?** — data governed under a canon that agents can consume with confidence, and the real processes surveyed from their traces, not from their documentation. **What do I do first?** — the portfolio classified and prioritized, with the strategy chosen by context and the DUAL cases up front.

The complete figure has a property worth making explicit: **no piece works alone**. A diagnosis without a portfolio is a report that gets filed away. A portfolio without a diagnosis is a wish list. Advanced use cases on ungoverned data are pilots that do not scale — the industry is full of them, and the projected cancellation rates for agentive projects cited in the Introduction find a good part of their explanation there. AURA is the insistence that the transformation be managed as a system: vision, measurement, terrain, and execution, advancing together.

The question of the horizon remains: where does the path converge? The answer lives in the other two books of this series. *AgencyDomains* describes the formal architecture of the destination — the primitives, the layers, the *Trust Infrastructure* that an organization mature on both axes will end up operating. *The Real-Time Enterprise* describes that world from the inside — what work looks like once the transition is behind. This book ends where those two begin: at the frontier where the organization that learned to know and to do discovers it is no longer automating its past, but inhabiting something else.

The agentive path is not a technology migration. It is the process by which an organization learns, step by step and under governance, to entrust execution to systems that know — while its people rise to where they were always meant to be: defining what is worth knowing, and what is worth doing.

This page intentionally left blank.

Annexes

Annex A · The Diagnostic Instrument

Seventeen questions to place an organization on the two axes

This book is diagnostic and not prescriptive by design: it says where an organization stands and what standing there means, and leaves prescription to the consulting work that builds on that diagnosis. This annex is the practical consequence of that decision — **the instrument with which the diagnosis is taken.**

It condenses into applicable form what the *IRIS* and *MOTOR* chapters develop with their full rationale: the diagnostic questions of every stage and every level, the rule for placing yourself, and the reading of the pair. It does not replace those chapters — the *why* of each stage lives there — but it allows them to be used without having been read end to end.

For whom? For an organization that wants to place itself, and for the consultant who needs to place a client. **How long does it take?** Between fifteen minutes, answered honestly and from memory, and a two-hour session with evidence on the table and several areas represented. The second produces a defensible diagnosis; the first, a useful suspicion.

On its use. This annex is published to be used — by an organization on itself, or by a consultant with their clients, paid work included. The only expectation is attribution: that the diagnosis say where it comes from.

How is it applied?

The two axes are answered separately. They are orthogonal: an organization can know a lot and do little, or the reverse. Measuring them together, or averaging them, erases precisely the finding the pair exists to produce.

Each axis is walked in order, from its first question. Levels cannot be skipped: each one rests on the foundations of the previous. An organization operating at a level without having settled the ones before it is not ahead — it is exposed. And when answers differ by area — or by the dimension being looked at: architecture, governance, culture — **the lowest one rules:** the chapters develop this, and the reason is that the lagging dimension is the real bottleneck.

Your position is the last question answered “yes” without hesitation — and the first question of each axis is read the other way around. Stage 1 of knowing and level 1 of doing do not describe an installed capability but a state one leaves: that is why their questions probe the exit, and a firm “no” to question 1 is the diagnosis — the walk of that axis ends there. From question 2 on, every “yes” with evidence advances the position; the first “*more or less*”, “*in some areas*” or “*we’re working on it*” marks the border, and that is where the diagnosis lies. Any yes that comes after that border is an **island, not a**

stage: a capability without the foundations that hold it up. Islands are recorded — because they name exposure, not position — and they are the first thing the reading of the pair has to explain.

A “yes” demands evidence, not intention. If it cannot be shown working today, it is a no. Maturity models are not ruined by severity but by optimism, and optimism has a recognizable signature: answering for what the organization can do *on its best day*, or for what one area does and the others do not.

The KNOWING axis — how does information flow?

IRIS model: ten stages, from fragmented data to a self-managed intelligence ecosystem. Each stage is developed, with its observable characteristics, in the *IRIS* chapter.

#	Stage	Diagnostic question
1	Fragmented	If I ask two people from different areas for the same indicator, will I get the same number?
2	Systematized	Is there a defined system for capturing the key business transactions, with consistent registration rules?
3	Centralized	Is there a centralized data warehouse with dashboards that business users consult regularly?
4	Modeled	Do the key business metrics have a single, governed definition that the whole organization shares?

Here the **foundations** end. An organization at stage 4 has its data house in order — and its analytical exploitation still depends on technical intermediaries.

The border — the Quantum Leap. The step from stage 4 to 5 separates two worlds. Up to here, the **online enterprise**: the data is current, and someone has to look at it, interpret it, and decide. From here on, the **real-time enterprise**: the system detects, interprets, decides, and acts. It is also where the crossing of the **Nadella Line begins** — stages 1 through 4 are better tools for the same processes; from stage 5, the processes themselves begin to be transformed.

#	Stage	Diagnostic question
5	Accessible	Can a manager get an analytical answer that was not pre-built into a dashboard, in minutes instead of weeks?

#	Stage	Diagnostic question
6	Contextual	Is the system capable of crossing sources and enriching an answer with relevant context the user did not explicitly ask for?
7	Proactive	Do business owners receive relevant information before asking for it, with enough context to act?
8	Bidirectional	Are there agents that execute operational actions autonomously with traceability, within rules that humans define and supervise?
9	Coordinated	Do agents from different areas of the business communicate with each other to coordinate actions with consistent information?
10	Orchestrated	Does the information ecosystem evolve autonomously — identifying gaps, improving models, and optimizing flows — with human supervision only at the strategic level?

Notes for answering this axis. They head off the frequent errors; each one names what counts as evidence.

- **Stage 1.** Answer it thinking of the **official** business indicators. A full “yes” — any indicator, any pair of areas — anticipates stage 4; to leave stage 1 it is enough that the official ones agree. If not even those agree, the organization is at stage 1.
- **Stage 3.** If “data warehouse” means nothing to whoever is answering: *is there a central place where the data of every area is consolidated, with dashboards that business people actually consult?* The word does not matter — lakehouse, cloud semantic layer, or classic warehouse all count.
- **Stage 5.** This is the question that decides which side of the Quantum Leap the organization lands on, and so it carries its own guardrail: the “yes” also requires that **two managers asking the same question get the same answer**. Without that, what exists is a copilot over ungoverned data, not stage 5.
- **Stage 6.** The evidence is **two concrete cases** of answers that brought context nobody asked for — not the vendor’s declared capability.
- **Stage 7.** The evidence is a real alert **with its why** — “X went down *because* Y, and that affects Z” — not the automatic email a dashboard sends when a threshold is crossed.

The DOING axis — how much autonomy do the processes have?

MOTOR model: seven levels, from ad-hoc use of AI to orchestrated autonomous execution. It measures **who executes the work and with what degree of autonomy**, not which tools were bought. Each level is developed in the *MOTOR* chapter.

#	Level	Diagnostic question
1	Invisible · Shadow AI	Do you know how many AI tools are in use across the organization, who uses them, and what data is shared with them?
2	Trusted	Do you have official AI tools with governance, audit trails, and protection of sensitive data in place?
3	Collaborator	Do employees delegate complete tasks to AI agents — defining the objective and receiving the result — instead of just asking for targeted assistance?
4	Analyst	Do you have agents operating 24/7, coordinated with one another, consulting a digital model of the system, with human intervention only on exceptions?
5	Specialist	Do you have specialist agents able to predict events and simulate scenarios using proprietary data that competitors cannot replicate?
6	Manager	Do you have agents with authority to assign work, manage budgets, or make management decisions backed by a digital model of the entire organization?
7	Authority	Do you have agents that participate in strategic decisions with a digital model of the complete ecosystem — competition, market, regulation — and long-term future simulation?

The first question deserves a warning: almost no organization answers yes, and the reason is that AI already came in through the back door. **Level 1 does not describe those who do not use AI — it describes those who use it without knowing.**

Notes for answering this axis.

- **Levels 1 and 2.** A “yes” at level 2 presupposes the “yes” at level 1: official tools without an inventory of what is used outside them is governance by halves. The pair “no to 1, yes to 2” is **not an island** — it is the signal that the 1→2 transition is half crossed, and the diagnosis is level 1.
- **Level 2.** *Audit trails*: an auditable record of who used which tool, for what, and with what data.
- **Level 4.** *Digital model of the system*: a digital replica of the operation, self-synchronizing, that agents consult as the source of truth. And since the question chains four conditions — 24/7, coordination, digital model, humans only on exceptions — the “yes” requires all four; if one is missing, the level is the previous one.
- **Level 5.** The evidence for the “yes” is the prediction in operation — the event anticipated, the simulation used to decide — and the years of proprietary data feeding it; the advantage over competitors is not exhibited, it is inferred.

Every jump between levels has a name, and the name works as a diagnostic signal: if an organization says “we’re implementing AI governance”, it is in the 1→2 transition.

Jump	Capability it demands	From ... to
1 → 2	Trust Infrastructure	From risk to governance
2 → 3	Wingworking	From assistance to delegation
3 → 4	Agentic Infrastructure · descriptive twin	From delegation to autonomy
4 → 5	Agent marketplace · predictive twin	From generic to specialized
5 → 6	Autonomous organization · organizational twin	From capabilities to authority
6 → 7	Social agentic · ecosystem twin	From intra-organization to ecosystem

How is the result read?

The diagnosis is not a stage or a level: **it is the pair**. That is where the finding lives that neither axis delivers alone, and it is what turns two numbers into a leadership conversation.

The pair	What does it mean?	Where is the bottleneck?
High knowing · low doing	Impeccable dashboards and bottled-up decisions. The organization <i>sees</i> and does not <i>act</i> : the information arrives, and someone still has to look, interpret, and order.	On the doing axis — governance before delegation. Data has stopped being the excuse.

The pair	What does it mean?	Where is the bottleneck?
High doing · low knowing	Blind automation: fast processes over data nobody governs. It is the pair that produces damage fastest, because it scales errors efficiently.	On knowing: without stage 4 — single, governed definitions — every autonomous process scales errors instead of results.
Both low	The honest starting point of most. It is not bad news: it is the only position from which one can plan without self-deception.	In the foundations, in parallel — stages 1 to 4 of knowing, and on doing, the step from level 1 to 2: governance. No agent pilot before that.
Both high	A real-time enterprise in operation. The advantage is no longer technological: it is one of coordination and governance.	In architecture and trust — the terrain of <i>AgencyDomains</i> .

And one reading worth saying out loud before the committee: **a large gap between the two axes almost always names an organizational problem, not a technical one** — an area that bought autonomy without governance, or an impeccable data team nobody bought the action from.

What gets recorded fits in three lines: the coordinate pair, the question where each axis stopped, and the islands — the loose yeses further up, which do not change the position but name where the organization is exposed.

The baseline — the four measures

The pair places you; these four measures quantify what is at stake. They come from the first book of the trilogy — *The Real-Time Enterprise*, where each corresponds to one of the four faces — and can start being measured on Monday, even by hand. Their value is twofold: they give the baseline against which progress will be measured, and the current numbers are usually the most persuasive argument available in front of a leadership team.

Measure	How is it taken by hand?
How long a business question takes to become an answer	Over the last five real analytical questions that reached the data team: from when they were asked to when they were answered.
Hours of uninterrupted work protected, per person per day	A handful of people log one week of their deep-work blocks without interruption. The average tends to be uncomfortable.
The fraction of leadership time that is information transport	Over last week's calendar of one executive: which meetings and reports existed only to move information from one place to another.
The hour the day truly ends	From the same people who log the second measure: the timestamp of their last work message of the day, not the time of leaving the office.

What follows the diagnosis?

The result opens three paths depending on where the bottleneck landed, and each has its instrument in this same book:

- **If the problem is data governance** — who rules over the data when the consumers are agents and not people —: the *Data Canon* chapter.
- **If it is not known how the organization really operates**, because the real processes live in people's heads and not in the documentation: *Wingmap*, which reconstructs them from the digital traces the operation already produces, without interviews.
- **If the diagnosis must become a plan**: the *Use Cases* chapter — from diagnosis to prioritized portfolio, with the beneficiary framework and the two canonical starting strategies, illustrated with the inventory in Annex B.

What none of the three delivers is the will to cross. That is built in the leadership conversation, and that is what the first book of the trilogy is for.

This page intentionally left blank.

Annex B · Use Case Inventory

The reference municipal portfolio

This annex reproduces the complete inventory of the portfolio described in the *Use Cases* chapter: 100 atomic use cases and 30 value solutions for local government, developed by ultraBASE during 2025 (version 6.6 of the inventory). The levels correspond to the MOTOR model; the CITIZEN / CITY / DUAL classification and the four thematic areas are defined in the chapter (the area letters are: A — citizen services and communication; B — internal management; C — infrastructure, roads, safety, and emergencies; D — health, education, culture, sports, and the environment). The **priority** of each solution is expressed on a 0–100 scale. Some solutions incorporate additional cases from the source inventory identified as CU-NEW-..., which are not part of the catalog of one hundred atomic cases. Returns are **referential** (pending baseline validation, as portfolio discipline demands).

The 30 Value Solutions, by ranking

#1 - SV-GOB-012: Autonomous Municipal Financial Management (“FINANZAS AI”)

Priority: 94 | **Level:** 6 | **Type:** DUAL | **UC:** 4

Components: CU-082, CU-083, CU-017, CU-051

#2 - SV-GOB-018: Paperless Municipality (Full Digitalization)

Priority: 92 | **Level:** 3 | **Type:** CITIZEN | **UC:** 5

Components: CU-NEW-001, CU-NEW-002, CU-NEW-003, CU-027, ... +1 more

#3 - SV-GOB-024: Complete Financial and Administrative Management

Priority: 91 | **Level:** 4 | **Type:** CITY | **UC:** 7

Components: CU-009, CU-010, CU-018, CU-031, ... +3 more

#4 - SV-GOB-010: Autonomous Citizen Service Management (“ALEX”)

Priority: 90 | **Level:** 6 | **Type:** CITIZEN | **UC:** 4

Components: CU-078, CU-079, CU-015, CU-035

#5 - SV-GOB-002: Intelligent Citizen Contact Center

Priority: 88 | **Level:** 3 | **Type:** CITIZEN | **UC:** 2

Components: CU-035, CU-036

#6 - SV-GOB-025: Comprehensive Public Procurement Platform

Priority: 87 | **Level:** 3 | **Type:** CITY | **UC:** 5

Components: CU-005, CU-014, CU-021, CU-022, ... +1 more

#7 - SV-GOB-003: Emergency-Resilient City

Priority: 86 | **Level:** 4 | **Type:** DUAL | **UC:** 3

Components: CU-048, CU-049, CU-085

#8 - SV-GOB-011: Autonomous Municipal Works Program (“OBRA”)

Priority: 85 | **Level:** 6 | **Type:** DUAL | **UC:** 4

Components: CU-080, CU-081, CU-033, CU-034

#9 - SV-GOB-021: Integrated Social Development Center

Priority: 85 | **Level:** 4 | **Type:** CITIZEN | **UC:** 5

Components: CU-071, CU-NEW-008, CU-NEW-009, CU-NEW-010, ... +1 more

#10 - SV-GOB-029: Smart Local Public Safety

Priority: 84 | **Level:** 4 | **Type:** CITIZEN | **UC:** 2

Components: CU-055, CU-056

#11 - SV-GOB-023: Comprehensive Municipal Human Resources Center

Priority: 83 | **Level:** 3 | **Type:** CITY | **UC:** 6

Components: CU-011, CU-012, CU-019, CU-020, ... +2 more

#12 - SV-GOB-001: Comprehensive Planning and Execution of Municipal Works

Priority: 82 | **Level:** 3 | **Type:** DUAL | **UC:** 2

Components: CU-033, CU-034

#13 - SV-GOB-004: Predictive Municipal Public Health Management

Priority: 80 | **Level:** 4 | **Type:** DUAL | **UC:** 3

Components: CU-058, CU-061, CU-085

#14 - SV-GOB-027: Smart Public Spaces

Priority: 80 | **Level:** 3 | **Type:** DUAL | **UC:** 7

Components: CU-040, CU-042, CU-043, CU-044, ... +3 more

#15 - SV-GOB-020: Carbon-Neutral Municipality

Priority: 79 | **Level:** 4 | **Type:** DUAL | **UC:** 5

Components: CU-087, CU-075, CU-088, CU-NEW-006, ... +1 more

#16 - SV-GOB-007: Post-Disaster Assessment and Response

Priority: 78 | **Level:** 5 | **Type:** CITIZEN | **UC:** 2

Components: CU-072, CU-065

#17 - SV-GOB-030: Smart Waste and Cleaning Management

Priority: 78 | **Level:** 4 | **Type:** DUAL | **UC:** 2

Components: CU-052, CU-075

#18 - SV-GOB-026: Connected and Informed City

Priority: 77 | **Level:** 3 | **Type:** CITIZEN | **UC:** 6

Components: CU-003, CU-007, CU-008, CU-025, ... +2 more

#19 - SV-GOB-008: Environmentally Sustainable Smart City

Priority: 76 | **Level:** 5 | **Type:** DUAL | **UC:** 4

Components: CU-086, CU-060, CU-069, CU-059

#20 - SV-GOB-005: Predictive, Self-Repairing Urban Infrastructure

Priority: 75 | **Level:** 4 | **Type:** DUAL | **UC:** 3

Components: CU-054, CU-089, CU-087

#21 - SV-GOB-009: Smart, Sustainable Urban Mobility

Priority: 74 | **Level:** 5 | **Type:** DUAL | **UC:** 3

Components: CU-088, CU-074, CU-053

#22 - SV-GOB-019: End-to-End Smart Participatory Budgeting

Priority: 73 | **Level:** 3 | **Type:** CITIZEN | **UC:** 5

Components: CU-029, CU-030, CU-NEW-004, CU-NEW-005, ... +1 more

#23 - SV-GOB-006: Comprehensive Construction Supervision with Drones and AI

Priority: 72 | **Level:** 5 | **Type:** DUAL | **UC:** 2

Components: CU-065, CU-064

#24 - SV-GOB-028: Municipal Cultural and Educational Services

Priority: 72 | **Level:** 3 | **Type:** CITIZEN | **UC:** 3

Components: CU-047, CU-070, CU-063

#25 - SV-GOB-013: National Waste Management Network

Priority: 70 | **Level:** 7 | **Type:** DUAL | **UC:** 2

Components: CU-092, CU-075

#26 - SV-GOB-014: National Smart Public Safety Network

Priority: 68 | **Level:** 7 | **Type:** CITIZEN | **UC:** 3

Components: CU-093, CU-100, CU-056

#27 - SV-GOB-022: Smart Municipal Tourism

Priority: 67 | **Level:** 4 | **Type:** CITIZEN | **UC:** 5

Components: CU-NEW-011, CU-NEW-012, CU-NEW-013, CU-041, ... +1 more

#28 - SV-GOB-015: Inter-Municipal Benchmarking and Continuous Improvement

Priority: 65 | **Level:** 7 | **Type:** CITY | **UC:** 3

Components: CU-090, CU-091, CU-094

#29 - SV-GOB-016: Public Policy Innovation with AI

Priority: 62 | **Level:** 7 | **Type:** DUAL | **UC:** 2

Components: CU-098, CU-038

#30 - SV-GOB-017: Municipal Revenue Diversification

Priority: 60 | **Level:** 7 | **Type:** CITY | **UC:** 2

Components: CU-099, CU-097

The 100 Use Cases, by MOTOR level

Level 1 (14 UC)

- CU-001:** Building Permit Analysis *CITY | Area C*
- CU-002:** Regulatory Review *CITY | Area B*
- CU-003:** Responses to Frequently Asked Questions *CITIZEN | Area A*
- CU-004:** Complaint Categorization *CITY | Area A*
- CU-005:** Municipal Contract Review *CITY | Area B*
- CU-006:** Ordinance Generation *CITY | Area B*
- CU-007:** Social Media Content Generation *CITIZEN | Area A*
- CU-008:** Institutional Material Translation *CITIZEN | Area A*
- CU-009:** Budget Execution Analysis *CITY | Area B*
- CU-010:** Expense Anomaly Detection *DUAL | Area B*
- CU-011:** Screening of Job Competition Applicants *CITY | Area B*
- CU-012:** Performance Evaluations *CITY | Area B*
- CU-013:** Project Planning *CITY | Area B*
- CU-014:** Bid Proposal Analysis *CITY | Area B*

Level 2 (20 UC)

- CU-015:** 24/7 Municipal Chatbot with Anti-Prompt Injection *CITIZEN | Area A*
- CU-016:** Ticketing System with Masked Data *CITIZEN | Area A*
- CU-017:** Financial Analysis with Tokenization *CITY | Area B*
- CU-018:** Automatic Expense Auditing *CITY | Area B*
- CU-019:** Personnel Analysis with Privacy *CITY | Area B*
- CU-020:** Turnover Prediction with Protected Data *CITY | Area B*
- CU-021:** Contract Analysis with Tokenized Parties *CITY | Area B*
- CU-022:** Supplier Due Diligence *CITY | Area B*
- CU-023:** Permit Review with Anonymized Data *CITIZEN | Area C*
- CU-024:** Zoning Plan Analysis *CITY | Area C*
- CU-025:** Content Generation with Guidelines *CITIZEN | Area A*
- CU-026:** Certified Translation with Validation *CITIZEN | Area A*
- CU-068:** Medical Leave Management with Tokenization *CITY | Area B*
- CU-069:** Drinking Water Quality Monitoring with Automated Alerts *DUAL | Area D*

CU-070: Municipal Library Management with Personalized Recommendations *CITIZEN | Area D | Social Value*

CU-071: School Attendance Analysis with Dropout Alerts *CITY | Area D | Social Value*

CU-075: Solid Waste Management with Route Optimization *CITY | Area D*

CU-076: Recurring Supplier Analysis with Risk Scoring *CITY | Area B*

CU-077: Municipal Training Management with Gap Detection *CITY | Area B*

CU-084: Social Media Monitoring for Citizen Needs Detection *DUAL | Area A*

Level 3 (21 UC)

CU-027: Automated Building Permit *CITIZEN | Area C*

CU-028: Comprehensive Business License *CITIZEN | Area A*

CU-029: Citizen Proposal Processing *CITIZEN | Area A*

CU-030: Participatory Project Design *CITIZEN | Area C*

CU-031: Semi-Automated Monthly Accounting Close *CITY | Area B*

CU-032: Smart Expense Approval *CITY | Area B*

CU-033: Multi-Project Works Planning *DUAL | Area C*

CU-034: Works Inspection with AI *DUAL | Area C*

CU-035: Complex Multi-System Case Resolution *CITIZEN | Area A*

CU-036: Proactive Citizen Outreach *CITIZEN | Area A*

CU-037: Municipal Ordinance Generation *CITY | Area B*

CU-038: Regulatory Impact Analysis *CITY | Area B*

CU-039: Intelligent 360 Assistant for Multi-System Inquiries *DUAL | Area B*

CU-040: Public Lighting Optimization *DUAL | Area C*

CU-041: Mass Event Management *CITIZEN | Area D*

CU-042: Public Space Maintenance *CITIZEN | Area D*

CU-043: Automated Municipal Cemetery Management *CITIZEN | Area D*

CU-044: Smart Regulation of Street Markets and Street Vendors *DUAL | Area C*

CU-045: Animal Adoption Center with AI *CITIZEN | Area D*

CU-046: Municipal Sports Facility Management *DUAL | Area D*

CU-047: Smart Digital Municipal Library *CITIZEN | Area D*

Level 4 (19 UC)

- CU-048:** Automatic Response to Weather Emergencies *DUAL* | Area C
- CU-049:** Autonomous Traffic Incident Management *CITIZEN* | Area C
- CU-050:** Real-Time Fraud Detection *DUAL* | Area B
- CU-051:** Continuous Budget Projection *DUAL* | Area B
- CU-052:** Garbage Collection with Dynamic Routes *DUAL* | Area D
- CU-053:** Adaptive Multi-Intersection Traffic Lights *DUAL* | Area C
- CU-054:** Water Network Failure Prediction *DUAL* | Area C
- CU-055:** Camera Monitoring with Incident Detection *CITIZEN* | Area C
- CU-056:** Predictive Patrolling *CITIZEN* | Area C
- CU-057:** Dynamic Generation of Analytics and Dashboards on Demand *CITY* | Area A
- CU-058:** Municipal Epidemiological Surveillance System *DUAL* | Area D | Social Value
- CU-059:** Autonomous Environmental Noise Control *DUAL* | Area D
- CU-060:** Air Quality Monitoring Network *DUAL* | Area D
- CU-061:** Smart Municipal Vaccination System *CITIZEN* | Area D | Social Value
- CU-062:** Historic Heritage Management with AI *DUAL* | Area D | Social Value
- CU-063:** Smart Management of Cultural Events and Permits *CITIZEN* | Area D
- CU-085:** Autonomous Health Emergency Management *DUAL* | Area D
- CU-087:** Dynamic Optimization of Municipal Energy Consumption *CITY* | Area C
- CU-089:** Predictive Maintenance System for Critical Infrastructure *CITY* | Area C

Level 5 (9 UC)

- CU-064:** Pavement Deterioration Detection *DUAL* | Area C
- CU-065:** Works Inspection with Automated Flight + Specialized AI Analysis *DUAL* | Area C
- CU-066:** Local Municipal Case Law Analysis *CITY* | Area B
- CU-067:** Ultra-Specific Complaint Classification *CITY* | Area A
- CU-072:** Post-Earthquake Seismic Assessment *CITIZEN* | Area C
- CU-073:** Environmental Impact Study *CITY* | Area D
- CU-074:** Traffic Optimization *DUAL* | Area C
- CU-086:** Integrated Multi-Sensor Environmental Monitoring Network *DUAL* | Area D | Multi-Agent
- CU-088:** Autonomous School Transportation Coordination *CITIZEN* | Area D

Level 6 (6 UC)

CU-078: “ALEX” Manages a Team of 25 Staff Members *CITIZEN* | *Area B* | *Multi-Agent*

CU-079: Continuous Service Optimization *CITIZEN* | *Area A*

CU-080: “OBRA” Coordinates 45 Simultaneous Projects *DUAL* | *Area B* | *Multi-Agent*

CU-081: Works Crisis Management *DUAL* | *Area B*

CU-082: “FINANZAS AI” Manages a USD \$45M Budget *DUAL* | *Area B* | *Multi-Agent*

CU-083: Automatic Budget Reallocation *DUAL* | *Area B*

Level 7 (11 UC)

CU-090: “FINANZAS AI” Agent Consulting Other Municipalities *CITY* | *Area B*

CU-091: “OBRAS AI” Agent with Regional Track Record *CITY* | *Area B*

CU-092: Regional Waste Management Network *DUAL* | *Area D* | *Multi-Agent*

CU-093: Regional Public Safety Network *CITIZEN* | *Area C* | *Multi-Agent*

CU-094: “Chilean Municipal AI Network” *CITY* | *Area B*

CU-095: Government Agent Marketplace *CITY* | *Area B*

CU-096: “Digital Municipalities” National Program *CITIZEN* | *Area A*

CU-097: Commercial Spin-Off “GobiernoIA SpA” *CITY* | *Area B*

CU-098: Agent Proposes New Public Policy *DUAL* | *Area C*

CU-099: Agent Identifies New Business Opportunity *CITY* | *Area B*

CU-100: National Network of Municipal Agents for Emergencies *DUAL* | *Area C*

Portfolio Statistics**Totals**

- **Use Cases:** 100
- **Value Solutions:** 30
- **Estimated aggregate return (referential, pending baseline validation):** USD ~118M/year for the portfolio as a whole

Distribution by MOTOR level

Level	UC	VS
1	14	0
2	20	0
3	21	9

Level	UC	VS
4	19	9
5	9	4
6	6	3
7	11	5

Distribution by Classification

Type	UC	VS
CITIZEN	31	11
CITY	38	5
DUAL	31	14

Special fields

- **UC with declared social value:** 5 (intangible social benefit)
 - CU-058, CU-061, CU-062, CU-070, CU-071
- **UC with multi-agent relationships:** 6 (integrate or coordinate other UC)
 - CU-078, CU-080, CU-082, CU-086, CU-092, CU-093

The Four Thematic Areas

Area A — Citizen Services and Communication

Total: 17 UC (17%)

Use Cases: - CU-003: Responses to Frequently Asked Questions - CU-004: Complaint Categorization - CU-007: Social Media Content Generation - CU-008: Institutional Material Translation - CU-015: 24/7 Municipal Chatbot with Anti-Prompt Injection - CU-016: Ticketing System with Masked Data - CU-025: Content Generation with Guidelines - CU-026: Certified Translation with Validation - CU-028: Comprehensive Business License - CU-029: Citizen Proposal Processing - ... and 7 more UC

Area B — Internal Management, Legal, HR, Finance, Procurement, and Budget

Total: 37 UC (37%)

Use Cases: - CU-002: Regulatory Review - CU-005: Municipal Contract Review - CU-006: Ordinance Generation - CU-009: Budget Execution Analysis - CU-010: Expense Anomaly Detection - CU-011: Screening of Job Competition Applicants - CU-012: Performance Evaluations - CU-013: Project Planning - CU-014: Bid Proposal Analysis - CU-017: Financial Analysis with Tokenization - ... and 27 more UC

Area C — Infrastructure, Roads, Safety, and Emergencies

Total: 24 UC (24%)

Use Cases: - CU-001: Building Permit Analysis - CU-023: Permit Review with Anonymized Data - CU-024: Zoning Plan Analysis - CU-027: Automated Building Permit - CU-030: Participatory Project Design - CU-033: Multi-Project Works Planning - CU-034: Works Inspection with AI - CU-040: Public Lighting Optimization - CU-044: Smart Regulation of Street Markets and Street Vendors - CU-048: Automatic Response to Weather Emergencies - ... and 14 more UC

Area D — Health, Education, Culture, Sports, and the Environment

Total: 22 UC (22%)

Use Cases: - CU-041: Mass Event Management - CU-042: Public Space Maintenance - CU-043: Automated Municipal Cemetery Management - CU-045: Animal Adoption Center with AI - CU-046: Municipal Sports Facility Management - CU-047: Smart Digital Municipal Library - CU-059: Autonomous Environmental Noise Control - CU-060: Air Quality Monitoring Network - CU-061: Smart Municipal Vaccination System - CU-062: Historic Heritage Management with AI - ... and 12 more UC
