Data-powered Innovation Review

Wave 12

Foreword

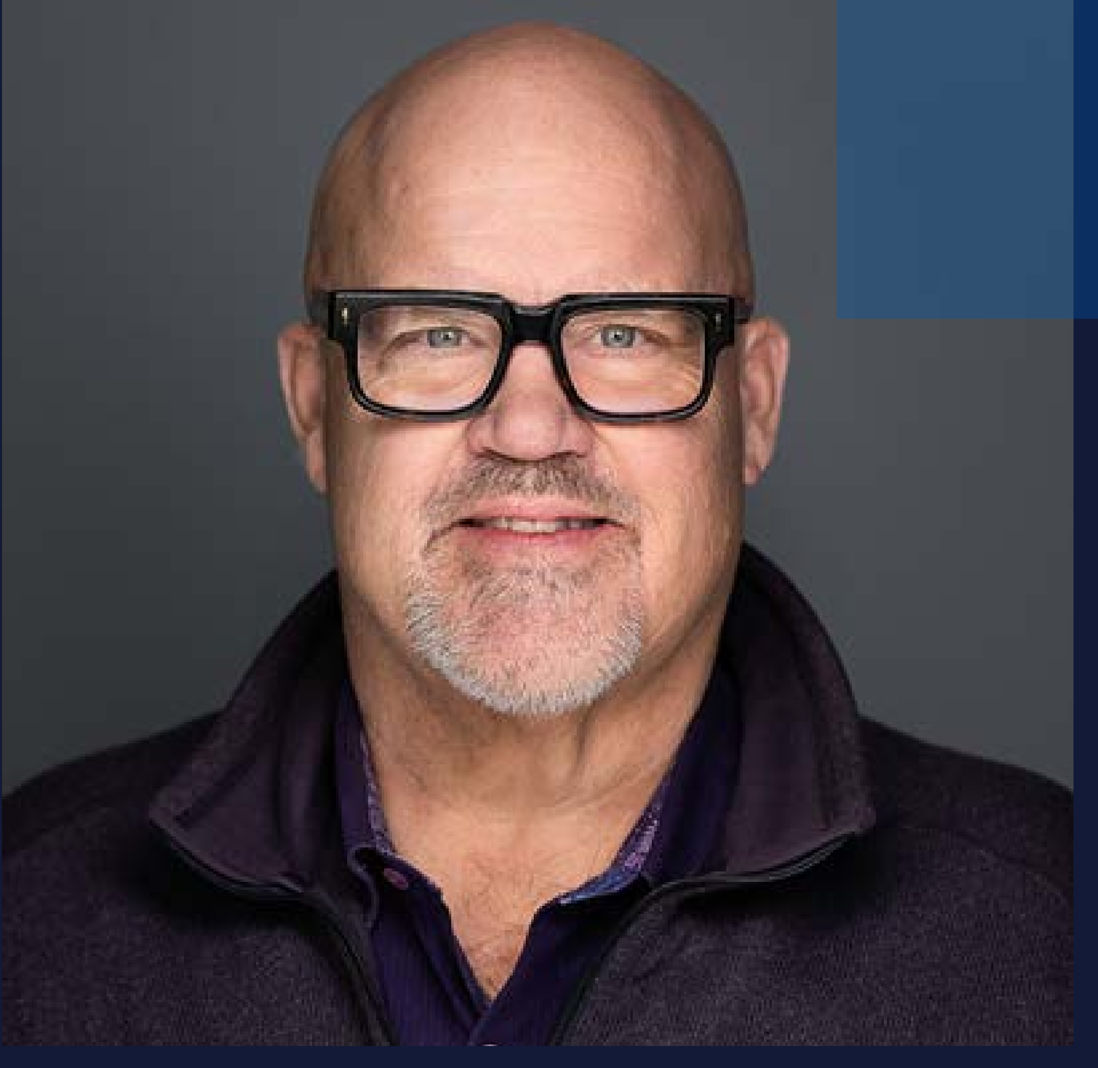

Kevin Campbell

CEO, Capgemini Insights & Data GBL
CEO, Syniti, part of Capgemini

I have been in the business of data and technology for a long time, and over the years, you start to recognize patterns. I'll often say, "I have seen this movie before," and not too long ago, that felt very true. The business landscape had its ebbs and flows, but it felt somewhat predictable. That has changed. We are now working with technology that does not just compute, it communicates. That shift is fundamentally changing how we operate, how we solve problems, and how we navigate complexity.

This is exactly why we pulled together the 12th Data-powered Innovation Review. The goal is straightforward. We need to cut through the noise and zero in on what actually matters in your day-to-day work. In this edition, we showcase where innovation is already moving beyond theory and into real-world impact.

We take a close look at the materials revolution, specifically Metal Organic Frameworks. These Nobel Prize–recognized materials act like microscopic sponges and are opening up new possibilities in climate solutions. They can capture carbon dioxide, store hydrogen more efficiently, and even extract drinking water from dry environments. We also explore the rise of the autonomous AI worker. Artificial intelligence is moving well beyond conversation. It is beginning to take action independently. For the first time, humans and systems are operating side by side in a truly integrated way, reshaping what work looks like.

This level of change brings real challenges. We address them directly, starting with efficiency and access. There is also digital sovereignty, where the push for independence is driving demand for localized processing and greater control. Finally, we look at small language models and the very real questions around trust and oversight. And underpinning all of it is a simple reality: without high-quality, well-governed, and ready to use data, none of these innovations can truly come to life.

I invite you to grab a coffee, sit back and take some time to explore these ideas. I think you will find them worth the read.

Editorial

The Power of **Soft Diffusion**

Robert Engels

I&D Head of Innovation, VP Insights & Data, Capgemini

If you've been following our journey through the previous eleven waves of the Data-powered Innovation Review, you'll know we have a penchant for finding poetry in the technical. We've looked at frozen spheres, desert sands, and the intricate architecture of coral reefs to explain the shifting tectonics of data and AI.

For this 12th edition, we find ourselves immersed in **soft diffusion**.

Looking at the flowing, chromatic gradients on our cover, one might see the gentle bleed of watercolors or the ethereal glow of a nebula. In the world of Generative AI, "diffusion" is the process of turning random noise into coherent beauty. But in the broader business landscape of 2026, it represents something even more profound: the way technology is finally losing its "hard edges" and seamlessly permeating every layer of our organizations.

The current wave observes the transfer from brute force to elegant spread: for a few years, AI felt like a series of heavy, monolithic boulders dropped into a pond; massive, disruptive, and impossible to ignore. If you are reading this now, you will see how Wave 12 explores the "big splash" era giving way to a more sophisticated spread.

We open with a deep dive into **Small Language Models (SLMs)**. If the giants of yesterday were the heavy machinery, SLMs are the "soft diffusion" of intelligence. Frugal, specialized, and capable of being embedded exactly where they are needed. They prove that you don't need to consume the energy of a small city to answer a specific business question with precision and sovereignty.

We also analyze the "flow of intention" that results from agentic systems steered by AI. With that we are witnessing a transition from AI that merely talks to AI that acts. Our features on **Agentic AI** and **Autonomous Teammates** represent the next phase of this diffusion. We are no longer just chatting with a box; we are designing systems that flow into the lines of business, ensuring that manufacturing procedures or water management policies are not just efficient, but inherently compliant with safety handbooks by design. Intelligence is no longer a destination; it is the atmosphere in which we work.

Innovation doesn't happen in a vacuum, and it certainly shouldn't be "one-way." The "soft diffusion" of data is the secret ingredient in the Circular Economy. In this issue, we explore how data acts as a digital thread, allowing materials to flow from birth to rebirth. By embracing **Semantic Data Sharing** and **Sovereignty**, we create a common language that allows different companies to collaborate without losing their "inner glow"—their unique data value.

As always, our "Innovation Movers & Shakers" have peered into their crystal balls for 2026. The consensus? While the "Data Jazz" is becoming more atmospheric and diffused, the human hand on the instrument is more important than ever. Whether we are discussing **Trusted AI** or **Quantum-ready architectures**, the goal remains to create a world where technology supports us so naturally that we almost forget it's there.

So, as you navigate the articles in Wave 12, I invite you to look for the "soft diffusion" within your own strategy. Don't just look for the next big disruption; look for the most meaningful integration.

Enjoy the ride.

Table of contents

Part III - Trust & Sovereignty

the human anchor, and owning what matters

The line cooks of enterprise AI

Stop sending Michelin chefs to plate every order

Yashowardhan Sowale

CTIO I&D India, VP, Global Architects Head, Insights & Data, Capgemini

Ananda Deshmukh

Senior Director, Insights & Data, Capgemini

Most enterprise AI roadmaps still treat the frontier LLM as a one-and-only: a celebrity chef hauled out to plate every order, garnish every burger, taste every soup. It looks brilliant in the demo and breaks at dinner service. Small language models are how the kitchen actually runs at scale: closer to the data, cheaper to operate, and good enough for the work most agents really do. The enterprises that figure this out first will be in production while their competitors are still piloting.

The kitchen is on fire

Walk into any enterprise AI war room in 2026 and you will see the same chart on the wall: cost per transaction trending up, p95 latency drifting right, and a CFO asking why the gross margin on every customer interaction now carries a token bill. The accepted answer, until very recently, was to keep paying it. Frontier LLMs are the safest culinary hire in the building. They can do anything.

But "they can do anything" is a dangerous compliment. The agentic workloads currently saturating cloud egress queues are not creative reasoning. They are intent routing, query rewrite, entity extraction, tool selection, function calls with a JSON schema, and a guardrail check at the end. They are line work. And the frontier LLM, billed per token and routed across a continent, is the most expensive line cook in human history.

The KPIs that matter for production agents are unforgiving and concrete: p95 latency, cost per transaction, tool-call success rate, structured-output validity, accuracy against a golden test set, and the compliance signals nobody tweets about: PII leakage and policy violations. None of those reward bringing an Escoffier to a salad station. So the question every CTO is now asking out loud is: how much frontier intelligence does my agent really need to do its job?

> *The frontier LLM, billed per token and routed across a continent, is the most expensive line cook in human history.*

A different brigade

Calling small language models "smaller LLMs" undersells them; they are a different role on the line entirely. Built as compact transformer architectures and adapted with instruction tuning, LoRA, and QLoRA, they can be quantized to INT8 or INT4 and dropped onto a CPU, an NPU, or an edge GPU close to where the work actually happens: a plant floor, a clinic, a vehicle, an IoT gateway. Tail latency moves from "depends on the network" to "predictable in milliseconds." The unit economics move from metered against a frontier API to amortized against hardware you already own.

This is the shift in mindset. The LLM is the executive chef who designs the menu and handles the impossible plate. The SLMs are the brigade that actually feeds the dining room: fast, specialized, replaceable when the menu changes.

Figure 1: The SLM landscape – open-weight families, domain models, and deployment zones.

What's already on the pass

The open ecosystem is unusually mature. Alibaba's Qwen family ships open-weight, instruction-tuned variants tuned for coding, chat, and tool-use. Mistral's small models trade gracefully between quality and efficiency for dialogue, summarization, translation, and agentic function-calling. Microsoft's Phi series is reasoning-dense per parameter, ideal for on-device classification, extraction, and tool-use scaffolding. Google's Gemma small variants light up RAG assistants and structured generation under tight decoding constraints.

Domain-tuned SLMs are arriving even faster than the generalists. Healthcare and life sciences teams are deploying Privacy Guardian (2B), Medical Validator (3B), Med-PaLM 2, and BioMistral 7B. Banking has stood up micro LLMs for BFSI workflows. Legal and compliance has Alan, the legal summarizer. The pattern is consistent: instead of one general model that knows a little about a regulated domain, there is a shift toward a specialist model that knows the regulated domain very well, and ships small enough to live inside the regulator's perimeter.

The reference architecture that emerges is a planner–executor split. A frontier LLM in the cloud handles deep reasoning and plan generation. SLMs near the data execute the plan: routing on intent and policy, rewriting queries for retrieval, extracting entities, choosing tools, calling functions with structured outputs. A second SLM, or a rule set, verifies every output before it reaches users or systems. The LLM plans. The SLMs execute. An SLM checks the work.

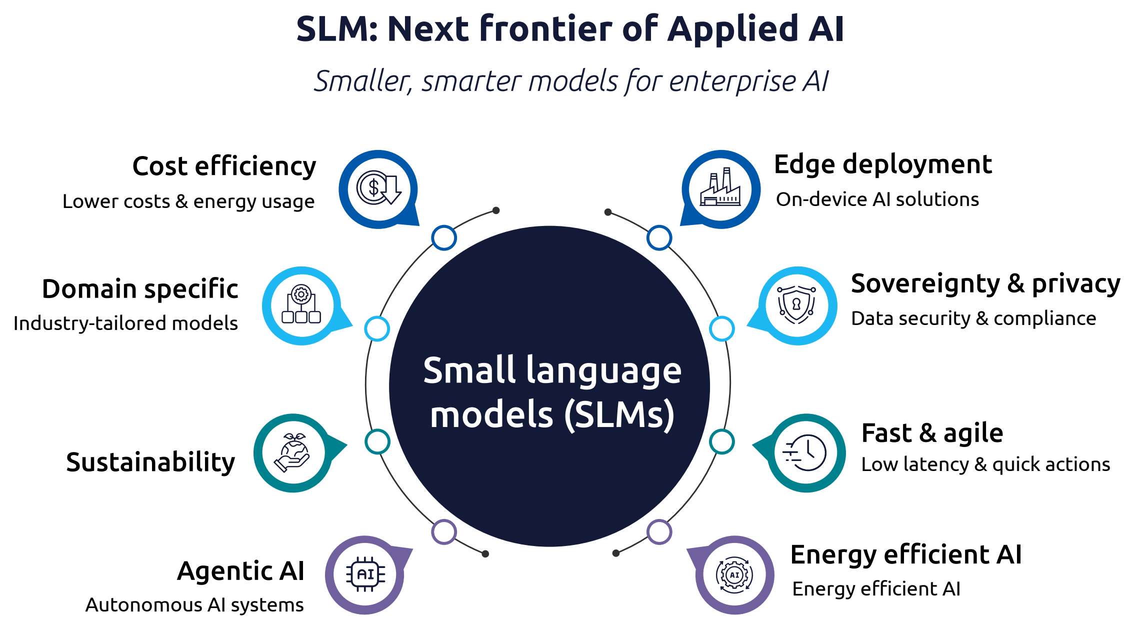

Figure 2: Reference architecture – LLM planner in the cloud, SLM executors near the data, SLM/rules verifier on the way out.

The build-and-run levers are now well understood: quantization to fit memory, runtime tuning for throughput, batching for concurrency, short prompts for context discipline, and task-specific evaluation against golden sets and regression tests. At fleet scale, you need the boring MLOps you have been deferring: model and version tracking, prompt and policy management, monitoring, rollback.

Where SLMs are already earning their keep

Domain	Agent task	Why an SLM wins
Customer support	Intent routing & ticket triage	High volume, schema-bound, latency-sensitive – frontier LLMs are wasted here
Manufacturing & field ops	On-device maintenance assistant	Runs without network, with predictable response time inside hangar/plant SLAs
Healthcare & life sciences	Clinical note extraction with PII redaction	Patient context never leaves the private perimeter; specialized accuracy on medical text
Banking & financial services	Function-calling for transactions, KYC, fraud signals	Auditable, structured outputs that regulators can actually inspect
Legal & compliance	Policy summarization, gap detection, contract review	Domain-specialized models beat general models on regulated language

Table 1: Common enterprise patterns where an SLM is the right tool for the job.

Capgemini's own deployments give the pattern texture. An on-device voice assistant for aircraft engine maintenance puts the SLM directly on the technician's device, eliminating network dependency and pinning response time inside an SLA the hangar can trust. A medical-assistant proof of concept fine-tuned MedGemma with privacy-aware handling and tested it on domain-specific evaluation sets, never letting patient context leave the controlled environment.

The grown-up part

None of this works without governance, and governance is where most SLM programs will fail. Four controls are non-negotiable. A model selection policy enforces an approved list by data sensitivity, deployment zone, and task type. Evaluation gates block release behind golden-set regression tests for accuracy, structured-output validity, and tool-call success. Runtime controls limit context size, enforce structured outputs, redact PII, allow-list tools, and run policy checks around every call. Observability traces requests end-to-end, measures cost and tokens, watches for quality drift, and keeps audit logs to whatever standard the regulator demands.

The honest challenges are also worth naming. Strong domain accuracy with low hallucinations needs labeled data and a serious evaluation harness, not vibes. Running many task-specific models requires versioning, drift checks, and safe rollouts. Fine-tuning and deployment work (LoRA, QLoRA, quantization, runtime selection, secure distribution) is more effort than calling somebody else's API. Data readiness still matters: lineage, PII controls, feedback loops. Anyone who tells you SLMs are easier than LLMs has not yet shipped a fleet of them.

Closing the kitchen

The instinct to send the celebrity chef out for every plate is understandable. It got the demo on stage. It will not get dinner served. SLMs make agentic AI feasible inside the constraints production actually has: latency SLAs, cost ceilings, privacy boundaries, energy budgets. They are also how you make it governable in a way a single black-box endpoint never will. Start with the smallest model that meets your accuracy and latency targets. Let the frontier LLM plan. Let the brigade cook. Then ship.

> *Anyone who tells you SLMs are easier than LLMs has not yet shipped a fleet of them.*

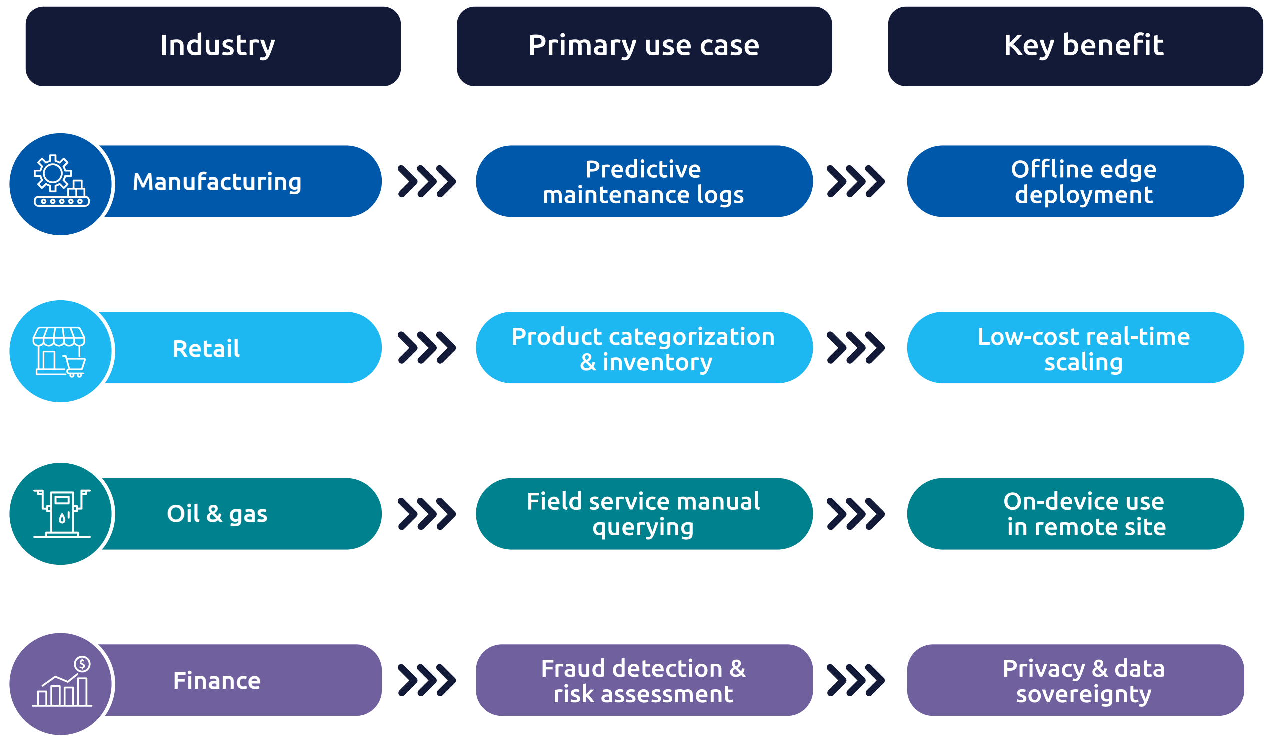

Figure 3: SLM value proposition – latency, cost, sovereignty, energy, and governance pulling in the same direction.

Start *innovating* now

Audit your inference economics

For every agent call running in production, instrument three numbers: p95 latency, cost per transaction, and tool-call success rate. Anything routine, repetitive, or schema-bound – routing, extraction, function calls, JSON validation – becomes the first batch of workloads to migrate from frontier LLM to a fine-tuned SLM, deployed close to the data.

Lock in a planner–executor split

Stand up a reference pattern your teams can copy without negotiation: frontier LLM in the cloud for deep reasoning and plan generation, SLMs at the edge or in private cloud for routing and execution, and a small validator (model or rules) that signs off every output before it leaves the perimeter. The LLM plans, SLMs execute, an SLM checks.

Run SLMs like a fleet, not a science project

Versioning, golden-set regression tests, drift monitoring, structured-output enforcement, PII redaction, allow-listed tools, and a working rollback path on day one. A dozen fine-tuned SLMs without governance is just a dozen new ways for production to break – at twelve times the surface area.

#SmallLanguageModels #AgenticAI #EdgeAI #FrugalAI #AIatScale #DataPowered

The cost of **hesitation**

Agentic AI removes the meeting between insight and action

Arunkumar Annamalai

Managing Solution Architect,
Insights & Data UK, Capgemini

Enterprises stopped failing at prediction years ago. They are still failing at the meeting that comes after. Decision intelligence, powered by agentic AI, collapses that meeting, and what it surfaces is uncomfortable: most organizations have been quietly using hesitation as a substitute for judgment.

Prediction was never the bottleneck

Most large organizations have spent the past decade getting very good at predicting the future. The forecasts are sharper, the dashboards richer, the data teams more capable. Demand shifts are anticipated months in advance. Churn is flagged before customers complain. Operational risk is modeled, outages simulated, scenarios run in exhaustive detail. The capability is no longer in question.

And yet, somehow, the organization moves at exactly the same pace it always did.

The recommendation goes into a deck. The deck goes into a steering meeting. The meeting moves to the next Tuesday because two people are traveling. By the time consensus arrives, the window has already closed. This isn't a failure of insight; it's a failure of what happens after the insight has landed. That failure has its own architecture, its own steering committees, its own glossy slides reassuring everyone that nothing was missed.

Here's the part that rarely gets said out loud. The bottleneck stopped being prediction years ago. The bottleneck is the gap between knowing and doing.

And that gap has a name. It is hesitation, dressed as governance.

The execution gap is human, not technical

For years, this slowness has been explained with reassuring stories. Improve data quality and adoption will follow. Make the model accurate enough and leaders will trust it. Explain the AI better and action will accelerate. None of this is wrong. None of it has fixed the problem.

What actually holds action back is messier. Fear of personal accountability. The comfort of the manual checkpoint. Fragmented ownership that quietly dilutes everyone's responsibility. A cultural preference for consensus over commitment. In complex environments, delay feels safer than action, even when the evidence says the opposite.

The honest version is this. Models do fail. Data does have gaps. But a less-discussed and equally damaging failure mode is the one that gets no postmortem: a perfectly reliable prediction was produced, the system worked exactly as designed, and the organization still didn't move.

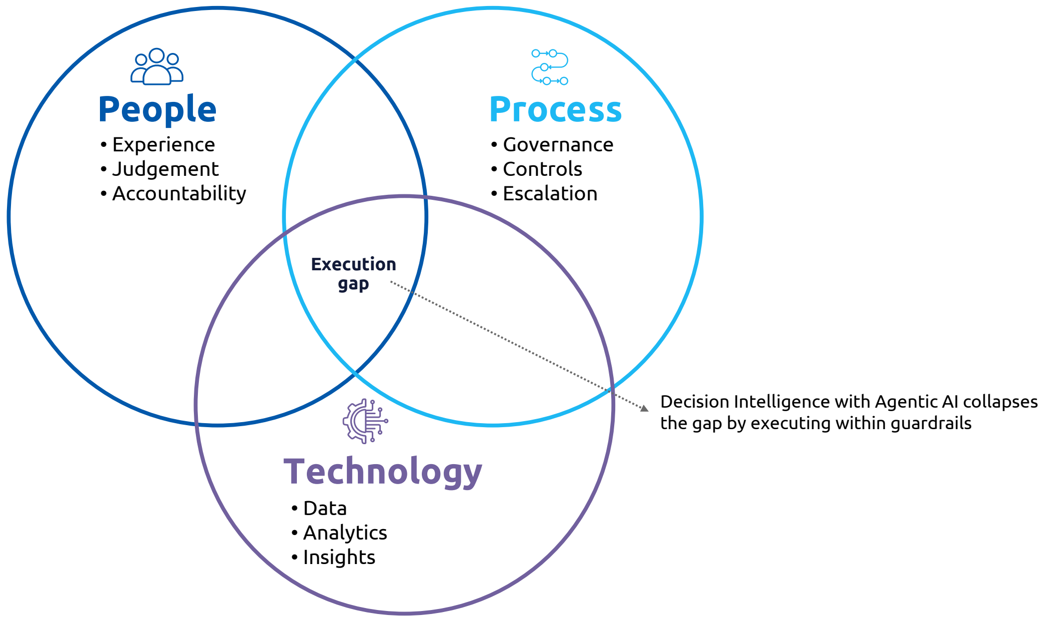

The execution gap sits at the intersection of people, process, and technology. Decision intelligence with agentic AI collapses it by executing within guardrails.

> *Decision intelligence does not fail for lack of data; it fails because organizations are still designed to hesitate.*

What changes when action becomes the default

Decision intelligence powered by agentic AI is what arrives at this fault line, and it does not arrive as a smoother evolution of analytics. It arrives as a challenge to how organizations translate insight into action, or fail to.

In practice, it runs as a system of interacting feedback loops. Decisions reshape the environment. Actions change what can be observed next. Outcomes are judged on their own terms, not on what was intended. Learning quietly reshapes how the next decision gets formed. Within predefined objectives, guardrails, and escalation thresholds, action no longer waits for fresh approval each time insight appears. Action becomes the default.

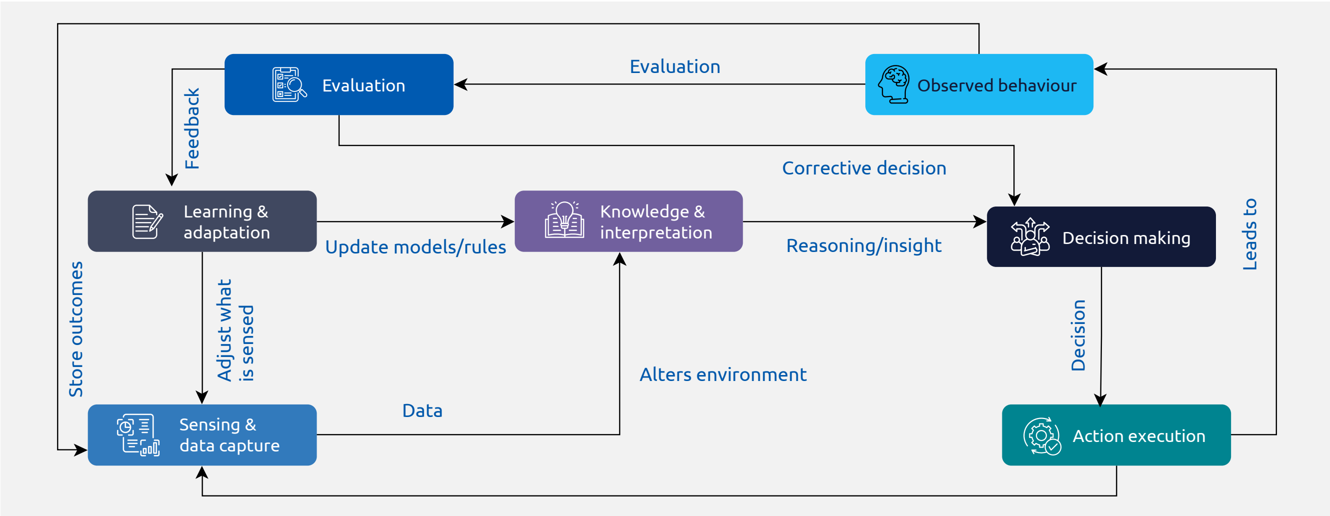

Decision intelligence runs as interacting feedback loops: sense, interpret, decide, act, observe, learn.

When that loop runs continuously, longstanding habits surface. Decisions stop waiting for perfect certainty. Trade-offs become explicit rather than negotiated in corridors. Inaction becomes visible and measurable. That visibility is what creates the discomfort.

Agentic systems will surface decisions humans usually defer: when to reroute supply even if margins suffer in the short term, when to deny customer exceptions to protect system stability, when to shut down processes that quietly accumulate risk. These choices aren't uncomfortable because they are wrong; they're uncomfortable because they remove the protective buffer of delay.

From control to stewardship

A common misconception is that agentic AI removes humans from decision-making. The reality demands something harder: a shift from control to stewardship. Instead of approving every action, humans set intent: objectives, guardrails, risk boundaries, escalation thresholds. Within those boundaries, the system executes consistently and without nerves. Responsibility doesn't disappear. It moves upstream and becomes harder to avoid.

This carries a real risk. When automated systems make mistakes, organizations are tempted to treat AI as a blame shield: "the system followed the rules." Without deliberate safeguards, accountability blurs rather than sharpens. The fix is unfashionable but simple. Define accountability frameworks before go-live, not after the first incident. The people who set the objectives own the outcomes those objectives produce – the good ones and the embarrassing ones.

The hardest changes here are not technical. They are cultural. And cultural changes don't appear in any architecture diagram.

> *Don't ask a committee to operate at the speed of a model.*

The mirror, not the upgrade

It's tempting to describe decision intelligence as the next evolutionary step in analytics. Technically that's fair. Practically it understates what organizations encounter when these systems actually run.

Faster decisions expose slow organizations. Automated actions make outdated processes visible. Consistent execution brings strategic inconsistencies into focus. Decision intelligence operates less like a new tool and more like a mirror, reflecting behaviors organizations have quietly learned to tolerate.

This is also why so many ambitious initiatives quietly revert to dashboards, alerts, and recommendation engines. Not because the technology fails, but because decisive systems demand a level of organizational readiness most teams underestimate. Don't ask a committee to operate at the speed of a model.

Early signals are emerging anyway. Supply chains use agent-led reallocation before consensus forms, because waiting often means missing the window. Network operations deploy autonomous systems that detect, diagnose, and remediate issues before human teams notice. IT functions use self-healing capabilities that act before incidents escalate. Customer platforms determine retention offers without managerial sign-off. The performance gains are real and measurable. What follows is rarely a technical debate. It is an uncomfortable conversation about who owned the outcome.

What 2030 rewards

Organizations that make progress don't chase full autonomy on day one. They start where hesitation is most visibly costly. They pick the one decision everyone already agrees is too slow, make objectives and guardrails explicit, automate selectively, and log outcomes rigorously, including the decisions that could have been taken but weren't. That last column is the one that matters. It puts a price on hesitation no dashboard ever shows.

Over the next few years, competitive advantage will not belong to organizations with marginally better models. It will belong to those willing to tolerate machine-led discomfort and treat hesitation as a risk rather than a safety mechanism.

Decision intelligence does not fail for lack of data; it fails because organizations are still designed to hesitate. Agentic AI is the part that makes inaction visible and action unavoidable. The reward, for those willing to look in the mirror, isn't faster decisions; it's an enterprise that finally acts at the speed it already understands.

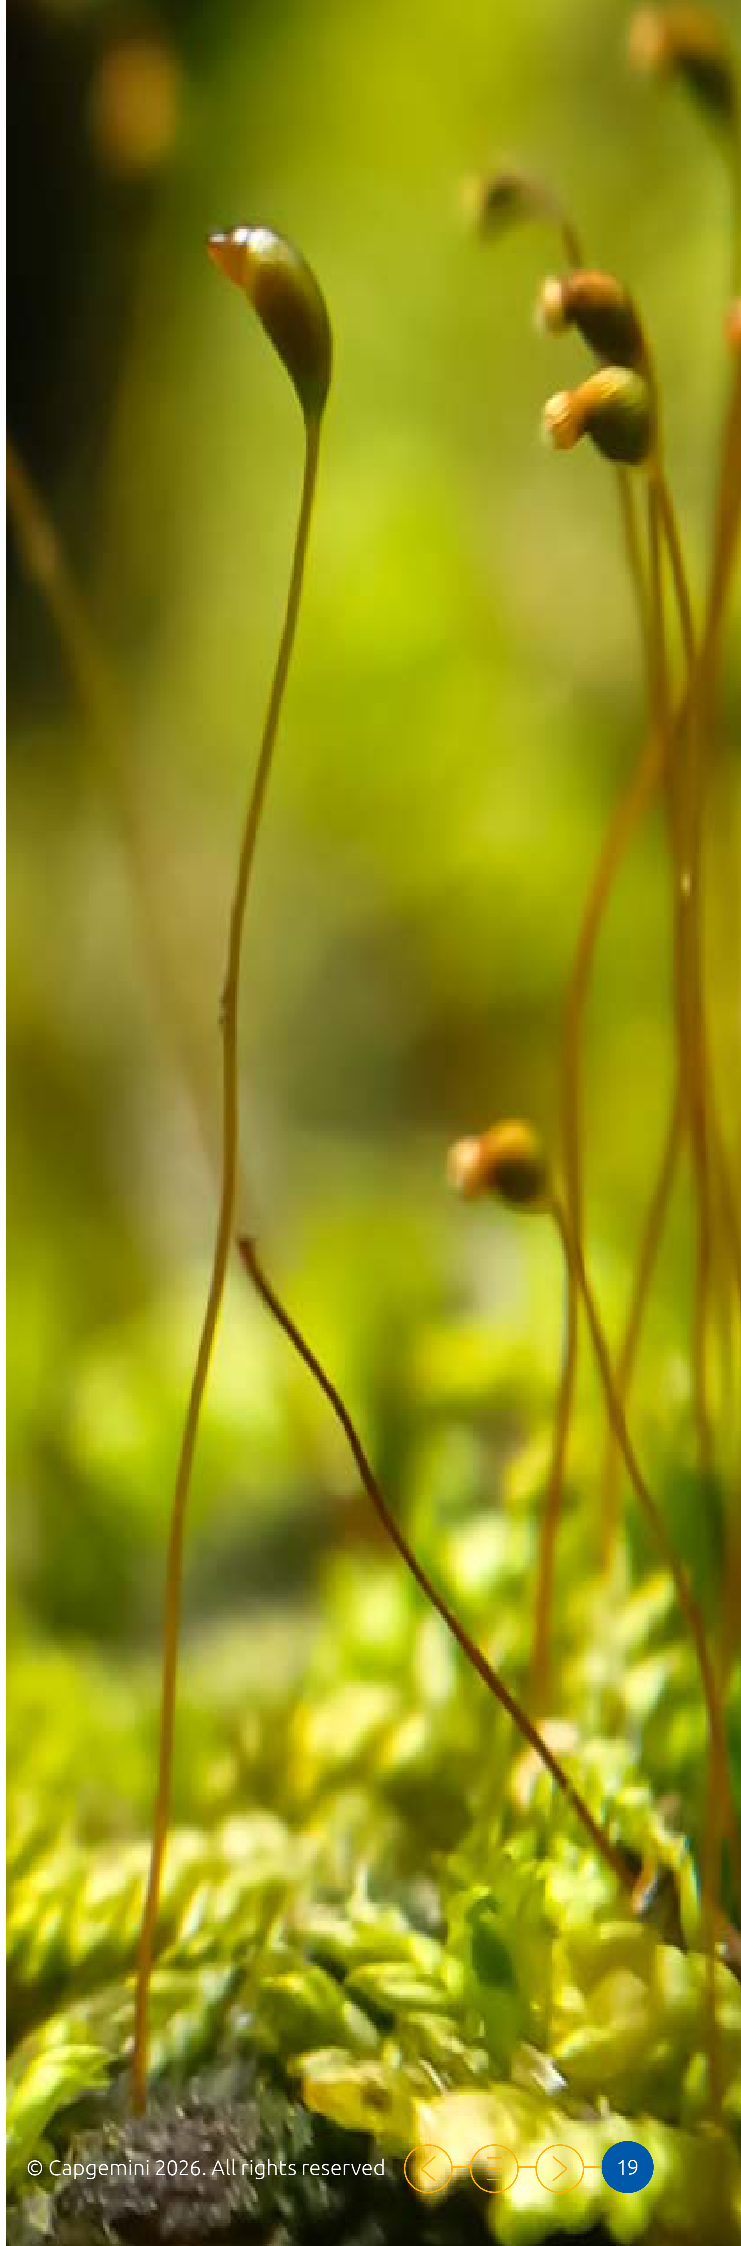

Start *innovating* now

Audit the slowest decision

Pick one decision everyone in the room already admits is too slow. Map who currently approves it, how long it actually takes from insight to action, and what each closed window costs. Hesitation has to be measurable before it can be removed.

Replace approvals with guardrails

Define intent, risk boundaries, and escalation thresholds upfront so an agentic system can execute within them without scheduling a meeting. The boundary, not the button, becomes the new control point.

Log the decisions you didn't take

Track every action the system was ready to execute but human approval blocked or delayed. That cost is almost always larger than the cost of acting wrongly, and almost always invisible.

#DataPowered #DecisionIntelligence
#AgenticAI #ExecutionGap #EnterpriseAI

The map beneath *the map*

Why your AI agents are sailing with half a chart, and how context graphs draw the rest

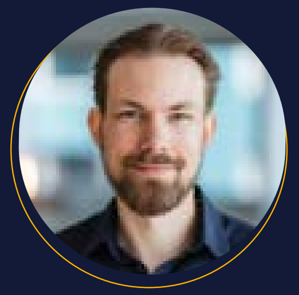

Joakim Nilsson
European Knowledge Graph Lead & Partner Lead, Neo4j

Johan Müllern-Aspegren
Emerging Tech Lead, Applied Innovation Exchange Nordics, and Core Member of AI Futures Lab, Capgemini

Modern enterprises are dispatching AI agents into the field with the confidence of cartographers and the maps of fortune-tellers. The data is there. The pipelines are clean. The dashboards glow. And yet the layer the agents most need was never written down, because nobody ever thought to call it data: the reasoning that connects evidence to action. Context graphs are the correction.

The cartographer's ambition

The earliest nautical charts and your latest digital twin share a single ambition: to model reality so we can navigate it. Map it, and you can understand it. Understand it, and you can predict it. Predict it, and you can act with confidence.

Today, every enterprise wants an agent in the field. Procurement agents to negotiate. QA agents to triage. Operations agents to escalate. The map gets handed over, the agent sails, and somewhere between the second and third decision, things go quietly wrong.

The fundamental problem isn't a lack of data. It isn't dirty data. It isn't siloed data either: those are engineering hurdles the industry learned to clear years ago. The problem is that the reasoning connecting data to action was never treated as data in the first place.

By the time the record lands in your warehouse, the why is already dead.

Where the why goes to die

Enterprise systems are built to capture outcomes. Supplier X was chosen over Supplier Y. CAPA-2026-001 was issued. Batch 4A19 was quarantined. The what is meticulous. The why is a rumor.

The why lives in the hallway exchange about a known issue with monsoon-season humidity. It lives in the unspoken contractual nuance that traded a 5% cost increase for a 20% risk reduction. It lives in a senior operator's intuition that Tuesdays are tricky for labels: knowledge that has never been entered into any field, by any user, in any system.

Roy Batty captured the entropy of context better than any white paper: all those moments will be lost in time, like tears in rain. A data record, in this light, is merely a tombstone marking a decision in a context that has already passed.

Without that context, every agent you deploy starts from zero. It sees the historical what and infers a model of the world that is structurally incomplete. Train on it, and you train on absence. Automate from it, and you automate the surface.

> *A data record is merely a tombstone marking a decision in a context that has already passed.*

The shift: Reasoning as data

A context graph is less a database than a structure: an accumulating record of decision traces across systems, entities, and time. Where a traditional database tells you the state of the world, a context graph tells you the story of the world.

Each node holds the texture the warehouse threw away:

- **Evidence:** what was known at the time
- **Constraints:** the legal, logistical, contractual walls in play
- **Alternatives:** what was considered and discarded, and why
- **Outcomes:** whether the intervention actually worked

Treat reasoning as data, and the agent stops being a high-speed junior employee with no idea why anything is the way it is. It starts being a colleague.

The pharma case: Where the happy path fails

Earlier this year, we worked with a global pharmaceutical leader on automating quality event management. The org chart says it is linear:

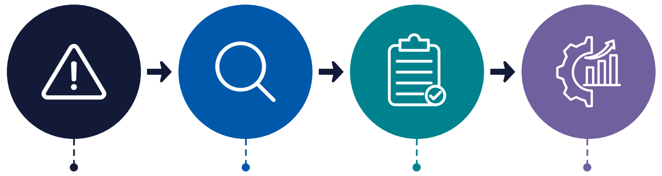

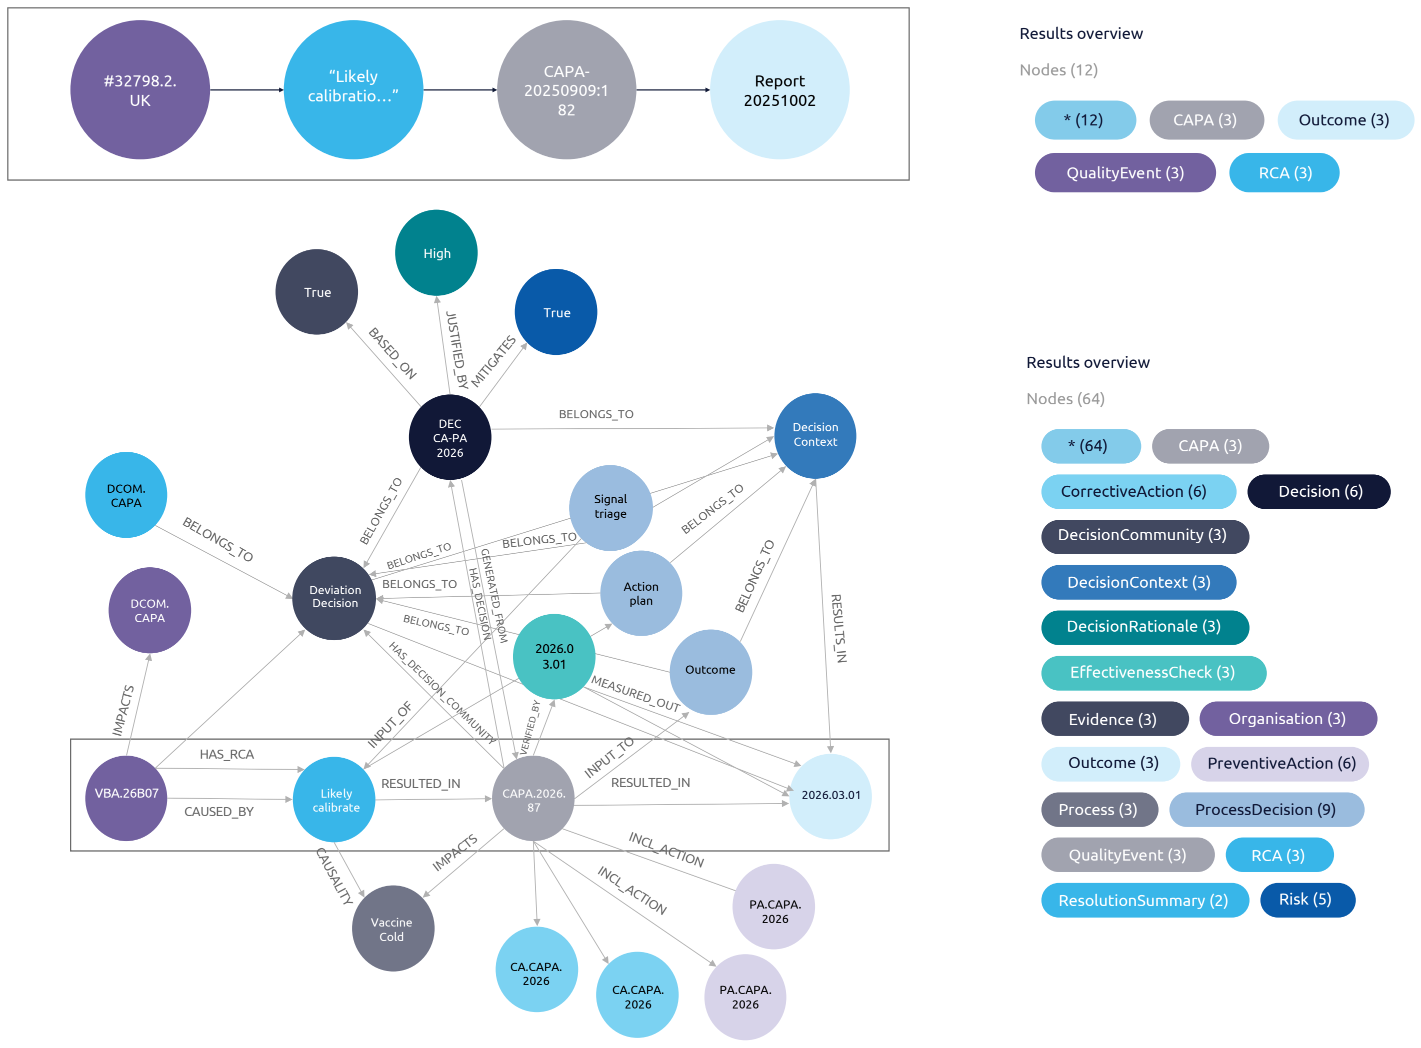

Figure 1. The happy path (top) and the actual decision graph (bottom). The same Quality Event, two very different worlds.

A blurred label is reported. An RCA is opened. The printer is named. A CAPA is issued: repair, quarantine, file, done.

The reality is rarely that obedient. The site's HVAC system had a known blind spot near the packaging line. On Tuesdays, when the cleaning crew ran high-temperature steam on the floors, ambient humidity climbed to 85%. The printer didn't fail. The label adhesive thickened, caused a microscopic drag, and the print blurred.

A human operator knows that Tuesdays are tricky for labels. The ERP sees Printer Error. Between the QE and the RCA sit the emails, the half-formatted Teams threads, the meeting minutes, and the canteen conversation: if this happened in Country A during monsoon season, skip the standard check and go straight to a full RCA.

If the interviews, the environmental readings, and the discarded hypotheses never land in a structure an agent can traverse, you are not automating decision-making. You are speeding up the mess.

The agentic-first edge

The companies that win the next five years won't necessarily be the largest. They will be the ones who designed for context capture from day one.

Agentic-first companies treat decision lineage as a first-class artifact. They accept slightly worse decisions in year one, on purpose, because they are building the substrate that lets their agents outperform humans by year three.

Legacy organizations face decision archaeology: reviewing every process, surfacing implicit knowledge, managing the very human reluctance of experts who suspect they are being asked to encode themselves out of a job. That reluctance is real, and it is not stupid. Pretending otherwise is how transformation programs stall.

By 2030, the question won't be whether your organization has agents. It will be whether your agents have context. And whether that context belongs to you, or evaporated into the hallway air a decade ago.

> *Without context, your AI is a high-speed junior employee with extraordinary recall and no judgment.*

The new corporate gold

We have hit the plateau of what raw data alone can deliver. To break through, the company DNA, the secret sauce, the why we do it this way, has to be encoded into a structure a machine can navigate.

The context graph may turn out to be the single most valuable asset a company owns. Without it, your AI is a high-speed junior employee with extraordinary recall and no judgment. With it, you have the architecture for systems that act on nuance.

The chart is incomplete. It always has been. The difference is that now you can finish drawing it, and your agents will sail by it long after the people who knew the waters are gone.

References and further reading

- *Garg, A. (2024). Context Graphs: AI's Trillion-Dollar Opportunity. Foundation Capital.*
- *Pearl, J., & Mackenzie, D. (2018). The Book of Why: The New Science of Cause and Effect. Basic Books.*
- *Kozyrkov, C. (2019). What is Decision Intelligence? Google Cloud.*
- *Nonaka, I., & Takeuchi, H. (1995). The Knowledge-Creating Company. Oxford University Press.*
- *Scott, R. (1982). Blade Runner [Film].*

Start *innovating* now

Audit the decision shadow

Pick one high-value process – procurement, QA, incident triage – and run decision archaeology on a single recent record. Trace it back through Teams threads, meeting notes, and corridor exchanges until you find what actually drove the choice. Those unrecorded moments are the first nodes of your context graph.

Define a causal ontology

A graph is only as good as the relationships it can express. Build a machine-readable vocabulary that distinguishes evidence (what was known), constraints (the walls), alternatives (what was rejected and why), and outcomes (what happened next). Done well, the system stops seeing "Printer Error" and starts seeing "Tuesday humidity spike during steam cleaning."

Make reasoning capture a KPI

Address the automation fear directly. Use lightweight LLM-assisted tools – voice memos, post-decision summaries – to let experts dump context in seconds, not hours. Reward the capture, not just the decision. The most experienced people stop being a flight risk and start being the teachers of the system.

#DataPowered #ContextGraphs #AgenticAI #DecisionIntelligence #KnowledgeGraphs

The tipping point of trust

How agents will reshape our relationship with brands

Sebastien Jouvenaar
VP, Strategy & Innovation,
Head of the Applied
Innovation Exchange, NYC

Neerav Vyas
VP, Customer First
Insights & Data,
Capgemini

For decades, we've assumed that persuasion happens at the moment a human pauses, looks, and chooses. But what if that moment quietly vanished? This article argues that the next tipping point in commerce isn't louder branding or better UX – it's whether autonomous agents can understand you, trust you, and act for you, long before a human ever notices.

Imagine, for a moment, the turn of the 20th century. The advent of the telephone. Suddenly, communication wasn't just about letters and telegrams; it was about instantaneous voice. Yet, even then, people couldn't quite grasp the full implications. Would it replace newspapers? Would it make us less social? We stood at a precipice, much like we do today with the quiet, relentless rise of agentic AI.

For years, our digital world has been defined by a simple, almost quaint act: the search. You had a question, you typed it into a box, and a search engine offered up a list of answers. Search engine optimization (SEO) was about optimizing for those fleeting moments of human curiosity, and it became the bedrock of digital marketing. The meteoric rise of ChatGPT into our cultural zeitgeist resulted in an evolution of search via natural language prompts. The result was not just to optimize returning a set of links, but rather delving into the intent of the original questions and providing a layer of synthesis along with a number of relevant links to better curate the search experience. Answer engine optimization (AEO) has become the burgeoning new field to meet the shifting search landscape. The "answer engine" was a slightly more sophisticated librarian, curating the relevant books instead of pointing to the shelf. AEO is a refinement, not a revolution.

What if that librarian decided to act on your behalf? What if, when you mused aloud about needing new footwear for an upcoming cross-country race, an invisible assistant didn't just show you links? Instead, it analyzed your recent mileage data, accounted for the muddy terrain of your upcoming route, sourced a shoe with the exact required lug depth, verified your size in a local retailer's live inventory, and purchased them, all without another prompt.

This isn't science fiction. This is the agentic revolution, and it's the quiet force poised to fundamentally redraw the lines of consumer interaction. These aren't just clever algorithms; they're autonomous entities. They don't just find information; they pursue goals. They don't just understand context; they learn and adapt. They don't just respond to a query; they initiate action.

Think of it this way: for decades, brands have been focused on capturing human attention. Every ad, every keyword, every meticulously crafted website was designed to catch the eye of a person scrolling through a screen. But what happens when the "eye" isn't human at all, but a sophisticated piece of code, designed to sift, analyze, and execute on behalf of its human master?

The unseen hand of orchestration: Agentic experience and the problem of proxy

The challenge, and the opportunity, for brands in this new landscape isn't just about being "found." It's about being understood and acted upon by a proxy. This is where agentic experience composition (AEC) enters the stage.

Consider the intricate dance of a symphony orchestra. Each musician, each instrument, plays a vital role, but it's the conductor who orchestrates the whole, ensuring every note contributes to a harmonious, compelling performance. AEC is that conductor. It's the art of ensuring that when a consumer's agent decides to, say, book a flight, every interaction – from checking flight availability to processing payment – is not just seamless, but truly aligned with the consumer's goal.

The challenge of course is that these proxy agents are here to primarily serve their creators, not necessarily the consumer and certainly not the brands the consumer interfaces with. The risk to organizations is that they're further disintermediated from consumer relationships and need to factor in designing for interactions with a new generation of digital gatekeepers.

This means rethinking the entire commercial architecture. We are moving beyond keyword matching to deep intent alignment. If an agent is tasked with "optimizing my morning routine," does a retailer's coffee maker listing simply describe its voltage and capacity, or is the data structured to articulate how it integrates with smart home schedules to contribute to a frictionless morning?

Then there is the challenge of interoperability. Imagine if every orchestra had its own unique musical notation. Chaos. For agents to work together, for your banking agent to communicate with your travel agent, and for both to access a global brand's data, there must be a common language. Your brand's "digital handshake" with these agents must be firm, standardized, and immediately decipherable.

The triad of transformation: Architecting for the agent

If orchestration is the strategy, then engineering for the agent is the relentless pursuit of operational fluidity. It's recognizing that winning in the agentic era requires a profound business process transformation, built on a rigorous, three-tiered architecture. You cannot bolt an agentic strategy onto legacy infrastructure.

To become the preferred brand of an autonomous proxy, enterprise leaders must align three critical layers:

- **The foundation: Deep data infrastructure**

Your product descriptions and inventory feeds aren't just for human eyes anymore. Through generative engine optimization (GEO), unstructured enterprise data must be transformed into structured, machine-readable formats. This is the bedrock. An agent needs to be able to seamlessly query your data infrastructure to reason about a product's attributes, its real-time availability, and its exact suitability for a highly specific, multi-variable task.

- **The engine: Technical capabilities and the simulation layer**

Forget the webpage as the sole interface. In the agentic world, your API is your storefront, serving as the digital conduit that allows machines to negotiate and execute at scale. But robust, low-latency APIs are merely the plumbing – the baseline requirement for entry. To truly engineer for the agent and capture market share, brands must deploy an advanced layer of technical capabilities designed to anticipate and optimize autonomous behavior.

How do you A/B test an invisible, autonomous agent? You cannot rely on traditional focus groups or heat maps. Brands must invest in robust simulation engines to stress-test their digital infrastructure before a real transaction ever occurs. By deploying synthetic personas – AI-driven digital twins representing varying customer segments and their unique proxy agents – brands can run millions of simulated, multi-variable queries against their own systems.

This allows a brand to model exactly how a "budget-conscious, time-poor" agent versus a "premium, sustainability-focused" agent interprets their data structure. It shifts the tech capability from simply responding to a request, to proactively training the brand's ecosystem to perfectly satisfy the algorithmic logic of any consumer proxy that comes knocking.

From experimentation to market simulation

In an LLM-driven market, experimentation and analytics alone aren't enough to stay ahead.
To capture and anticipate real-world signals, brands need to develop market simulation capabilities.

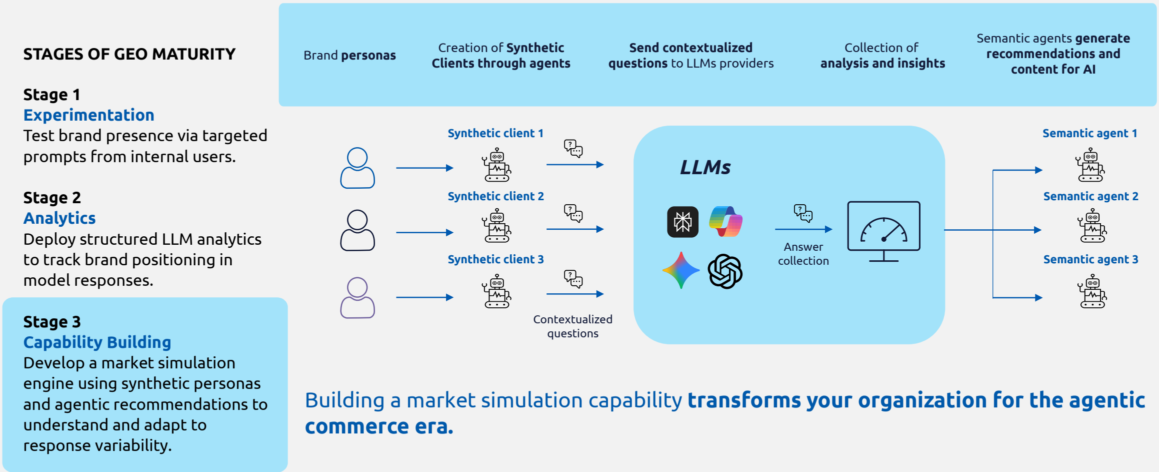

Building a market simulation capability **transforms your organization for the agentic commerce era.**

- **The orchestrator: Business and operational strategy**

The most pristine data and the fastest APIs are useless if the underlying business processes remain siloed. This layer requires rethinking how the organization operates. It means breaking down the walls between marketing, supply chain, and customer service so that an agent can seamlessly execute an end-to-end task (like returning a product and applying a dynamic discount to a new purchase). It also requires redefining commercial KPIs – moving away from traditional "clicks" and "dwell time" to measuring "successful agent resolutions" and "proxy retention."

Think of the "broken window theory" in urban planning. In this new architecture, a slow API, a disconnected inventory system, or a disjointed internal business process is a broken window. It signals friction to the agent and invites abandonment.

The new brand covenant: Trust, transparency, and the agent's recommendation

This isn't just a technical shift; it's a profound social and operational one. For decades, brand loyalty was built on direct engagement and the personal experience of using a product. But what happens when the initial discovery, vetting, and purchasing are done by a machine?

The new covenant between brands and consumers will be built on trust at scale. If a consumer's agent consistently finds your brand's information reliable, accurate, and easily actionable, that agent will become a powerful advocate. This isn't about traditional advertising; it's about earning the recommendation of an invisible, impartial assistant.

The future of brand–consumer relationships will hinge on:

- **Agent "trust" and "loyalty":** How do we shape agent decision surfaces so that our services remain dominant across multiple utility functions, not just today's optimization metric?
- **Data-driven empathy:** The aggregated insights from agent interactions will provide an unprecedented window into the true behaviors of consumers, allowing brands to anticipate desires and operationalize solutions with an almost prescient understanding.
- **Ethical vigilance:** As agents make more decisions, the ethical implications multiply. Brands must ensure their data is used fairly and transparently. The "black box" of AI must be demystified, at least in terms of its inputs and ultimate aims.

> *The shift to agentic interactions isn't just another digital trend. It is a fundamental recalibration of global commerce – a new "tipping point" in how value is exchanged.*

Brands that understand this, that commit to orchestrating seamless agentic experiences and optimizing their digital engineering for this reality, won't just survive. They will redefine what it means to lead in the autonomous age. The question isn't if your brand will engage with an agent, but how flawlessly it will execute when the agent comes knocking.

Start *innovating* now

Re-architect your digital experiences for autonomous decision makers

Audit your product data, flows, and touchpoints through the lens of an agent, not a human. Move beyond keywords and UX toward intent-aligned, machine-readable experiences that allow agents to reason, compare, and act on your behalf, flawlessly.

Make interoperability your competitive advantage

Treat APIs and MCP integrations as first-class brand assets. Standardize, document, and harden them so agents can seamlessly discover, trust, and transact with your services. In an agentic world, brands that are easiest to "talk to" will be the ones agents recommend and return to.

Optimize relentlessly for speed, reliability, and clarity

Identify and fix "broken windows" in your data pipelines, performance, and real time availability. Agents won't tolerate ambiguity or delays. The fastest, cleanest, most dependable signal wins, often without a human ever seeing the alternatives.

#AgenticAI #AgenticExperiences #ExperienceComposition #AIOrchestration #DigitalTransformation #FutureOfWork #IntelligentSystems #Design4AI

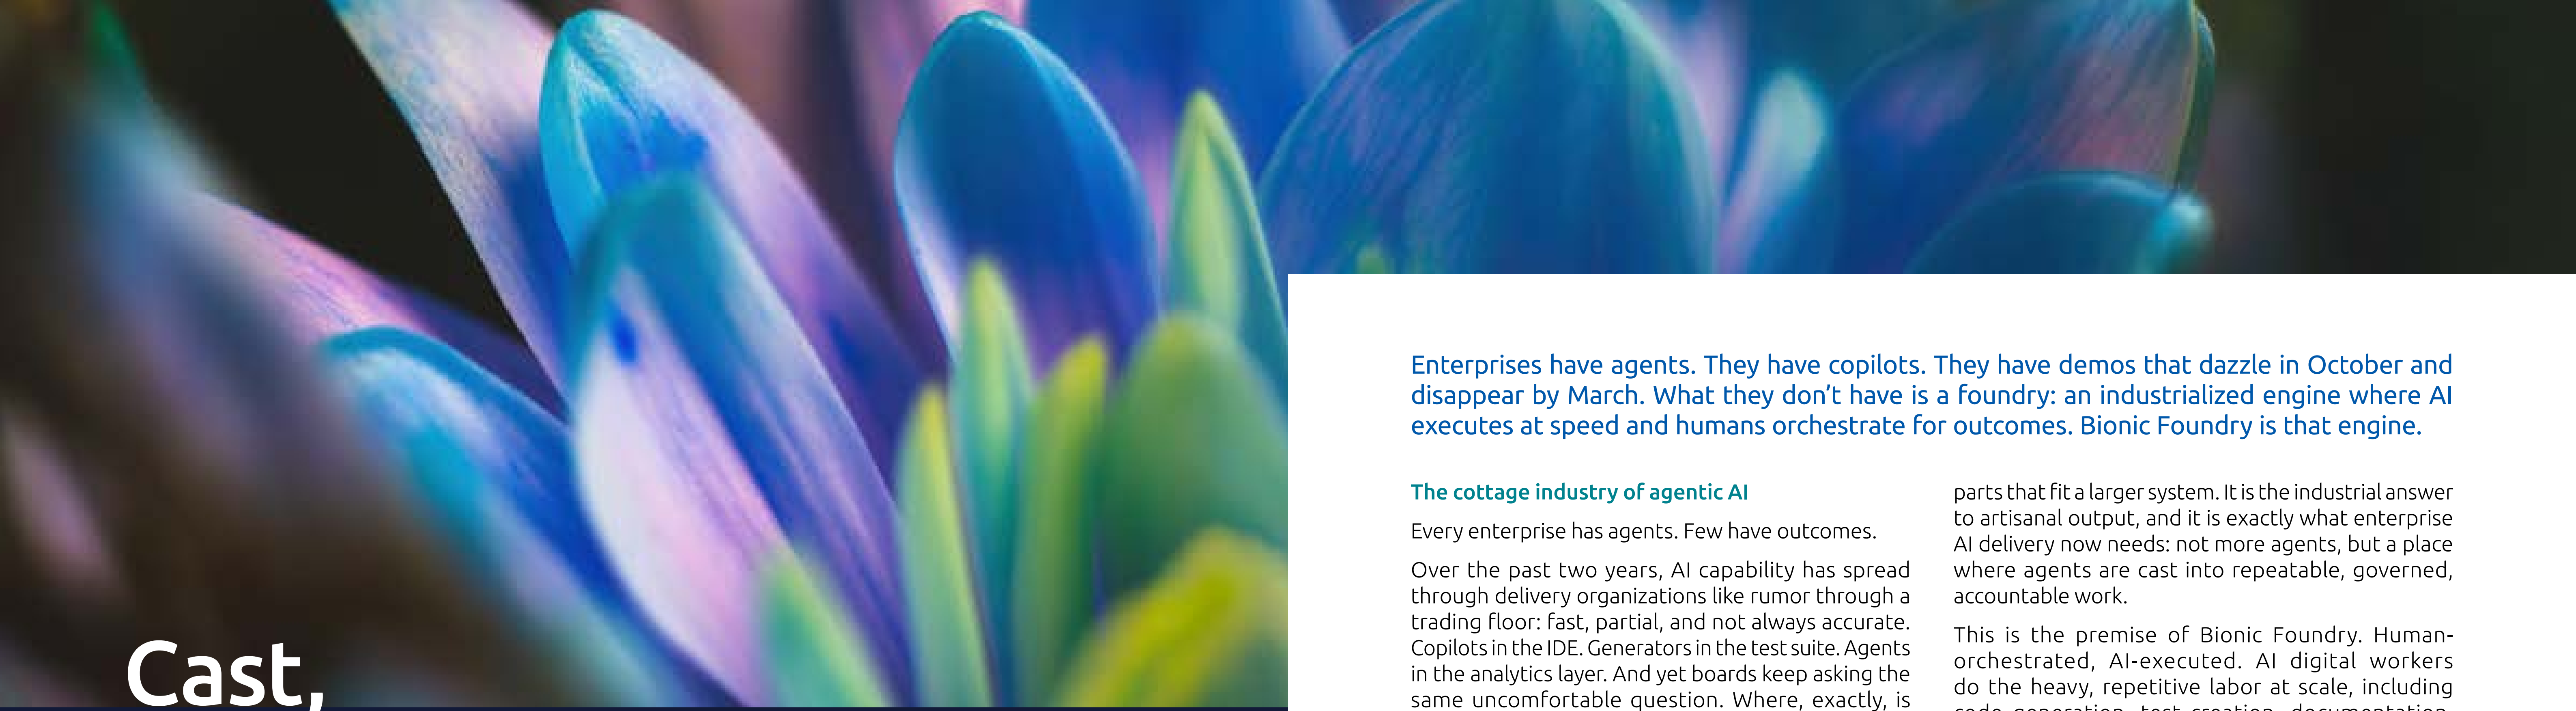

Cast, *don’t improvise*

Why enterprise AI needs an industrialized delivery engine, not another flotilla of pilots

Sarita Kaloya

Senior Director, Databricks Practice & Data and AI Capability Lead (India)

Vinod Jain

AI & Data Engineering Leader, Insights & Data, Capgemini India

Enterprises have agents. They have copilots. They have demos that dazzle in October and disappear by March. What they don’t have is a foundry: an industrialized engine where AI executes at speed and humans orchestrate for outcomes. Bionic Foundry is that engine.

The cottage industry of agentic AI

Every enterprise has agents. Few have outcomes.

Over the past two years, AI capability has spread through delivery organizations like rumor through a trading floor: fast, partial, and not always accurate. Copilots in the IDE. Generators in the test suite. Agents in the analytics layer. And yet boards keep asking the same uncomfortable question. Where, exactly, is the value?

Capability isn’t the problem. The shop floor is. Most enterprises have built artisanal AI, brilliant in the hands of one engineer and unrepeatable across a program. It is the cottage industry of agentic AI, and like every cottage industry before it, it cannot scale into industrial output without a different kind of building.

Program leads see the pattern repeat each quarter. A new agent ships, accelerates a workflow by some impressive percentage, and then quietly stalls when no one can answer who reviews its outputs, who owns its failures, or whether the team can reproduce the result for a second account. Industry studies confirm what those leads already feel in their bones: a meaningful share of AI’s “time saved” is reabsorbed by validation, correction, and rework. The headline number on the slide rarely survives the audit.

The agent worked. The delivery model didn’t.

> *Most enterprises have built artisanal AI: brilliant in the hands of one engineer, unrepeatable across a program.*

The question has shifted. It is no longer “can AI do this?” It is “can we trust, scale, and govern it in production, with our name on the result?”

From workshop to foundry

That answer requires a different building.

A foundry takes raw material that is molten, volatile, and hard to shape by hand, and turns it into standardized parts that fit a larger system. It is the industrial answer to artisanal output, and it is exactly what enterprise AI delivery now needs: not more agents, but a place where agents are cast into repeatable, governed, accountable work.

This is the premise of Bionic Foundry. Human-orchestrated, AI-executed. AI digital workers do the heavy, repetitive labor at scale, including code generation, test creation, documentation, and optimization. Human subject-matter experts orchestrate the lifecycle, validate the output, and stand behind the result. The combination produces something the industry has been chasing for two years and rarely achieving: agentic AI with assured outcomes.

How the foundry runs

A foundry isn’t a metaphor for one thing. It is a building with several lines, each tuned for the part being cast. Bionic Foundry operates the same way, through delivery archetypes that meet enterprises where they actually are.

Pre-built AI agents are deployed inside the client’s own environment, on their cloud, behind their firewalls, under their data residency rules. In a regulated life-sciences program, that looks like an agent drafting validation documentation against the client’s SOPs, inside the client’s tenant, with every prompt and response logged for inspection. In a telco data platform, it looks like an ingestion-code generator that produces pipelines conforming to the client’s existing naming, testing, and deployment standards from day one.

Co-created agentic solutions are designed jointly and integrated into existing platforms, with foundry SMEs running lifecycle and quality. Across every archetype, an orchestration and governance layer enforces standards, manages the handoffs between digital workers and humans, and keeps the audit trail intact. A rapid start kit – a curated library of best practices, naming standards, and institutionalized lessons – ensures every new line begins with yesterday’s hard-won knowledge rather than rediscovering it on the client’s clock.

The defining feature, across every archetype, is unromantic and important. AI does not operate unchecked.

From agentic AI to assured outcomes

Bionic foundry is an industrialized delivery model that combines human SMEs and AI-powered digital workers, governed to assure outcomes while improving quality, reducing skill dependency, and driving cost efficiency.

From capability to confidence

The market is moving fast, and so is the audit trail. Boards have stopped asking whether AI can do the work and started asking who signs off when it does. Regulators in life sciences, financial services, and telecoms are tightening expectations on traceability, validation, and explainability. In that climate, "we have agents" is not an answer. It is the setup for an awkward post-mortem.

> *Boards have stopped asking whether AI can do the work and started asking who signs off when it does.*

Foundries solve for this by design: faster time-to-value through AI-augmented SDLC, lower operational risk through standardized execution, higher accuracy because a human is on the line for the output, and a delivery posture that scales without losing the thread of accountability when the agent count doubles, then doubles again. The economic case improves with every unit cast. The compliance case does too.

There is a board-level corollary. AI investment committees are quietly tired of hearing about pilots. They want to know which delivery line has been industrialized, what the unit economics look like at scale, and what the failure mode is when something goes wrong. The companies that can answer those three questions cleanly are the ones receiving the next round of investment. The ones still showing demo videos are not.

What 2030 will reward

By 2030, the differentiator will not be access to AI. Every competitor will have it. The differentiator will be who can pour the metal cleanly, again and again, without surprises in the audit log. The artisans will keep producing brilliant, irreproducible demos. The foundries will produce outcomes.

AI generates potential. Foundries deliver.

Start *innovating* now

Pick one line, not ten

Choose a single high-volume, high-rework workflow: test generation, data ingestion code, or policy mapping. Industrialize it end-to-end before opening a second line. A foundry with one working belt beats ten half-built ones.

Pour the governance first

Before deploying a single agent to production, decide who validates the output, who approves it, and who is accountable when it is wrong. The orchestration layer is the foundry floor. Pour it before you start casting.

Codify the institutional memory

Capture coding standards, naming conventions, and hard-earned lessons into a rapid start kit your agents inherit by default. AI without institutional memory just industrializes bad habits faster.

#AIFoundry #AgenticAI #EnterpriseAI #HumanInTheLoop #AssuredOutcomes #DataPowered

Steering in the moment

How UWV turned its contact center from rear-view reporting into real-time insight

Afra Doludizgin
Data & AI Manager, Fabric Lead Europe, Capgemini

Samridh Pandey
Managing Solutions Architect, Capgemini

Sebastian van Duijn
Fabric Product Owner, UWV

Every phone call, chat thread, and bot exchange at a public agency is a sensor reading: a citizen telling you, in real-time, what works and what doesn't. Most contact centers collect those readings into siloed dashboards, then read them back weeks later. UWV chose a different path. With Microsoft Fabric as its operating layer, insight stops being a rear-view mirror and becomes a steering wheel.

The data-rich, insight-poor agency

A public-sector contact center is one of the most observed workplaces on earth. Every interaction is logged: voice transcripts, chat threads, CRM events, bot conversations, ticket states. The data arrives by the terabyte. The insight, somehow, arrives by the postcard.

That contradiction was the daily reality at UWV, the Dutch agency that administers employee insurance for an entire country. The signals were there. Citizens told staff exactly where policy interpretations were unclear, where forms were broken, where waiting times had crept up. But those signals lived in different platforms: legacy data warehouses, isolated CRM exports, file-based extracts, and contact-center tools bolted on a decade apart. By the time the streams were stitched together into a coherent report, the moment had passed. Teams expected to respond within hours were given dashboards that refreshed once a quarter.

Consider one routine citizen interaction: a question about a sickness benefit that has not arrived. The call is logged, categorized, routed, resolved. Multiply that by the thousands of interactions a UWV agent handles each day and the question that should follow is obvious: where is the heat map of where the policy keeps tripping people up? In most agencies, that map exists. It just lives in seven different tools, with seven different definitions of "tripping up."

Sebastian van Duijn, Product Owner Fabric at UWV, calls this what it is: a data-rich, insight-poor reality. There was no single source of truth. Reports contradicted each other. When they didn't, they arrived too late to matter.

How long can a public agency steer in real-time using last quarter's mirror?

> *With the right mindset and the right technology, you move from reporting about the past to steering in the moment.*
>
> *— Sebastian van Duijn, Product Owner Fabric, UWV*

Stop adding tools. Start wiring the cockpit.

The instinct, when faced with a fragmented estate, is to add another tool. Another reporting layer. Another dashboard that promises to be the one. UWV resisted that instinct. The team treated this as an architectural problem, not a tooling one. The underlying question was simpler than the toolset suggested: could insight feel as immediate as the conversation that produced it?

Microsoft Fabric became the answer because it collapsed the question. With OneLake as a unified data estate, conversational data, CRM events, and operational signals are ingested, organized through a medallion design, and governed centrally. Captured once. Reused everywhere. Dashboards that used to refresh once a quarter started behaving like live instruments on a working cockpit, not exhibits in a quarterly review.

To support this operating model, UWV implemented an end to end Microsoft Fabric architecture based on OneLake and medallion principles.

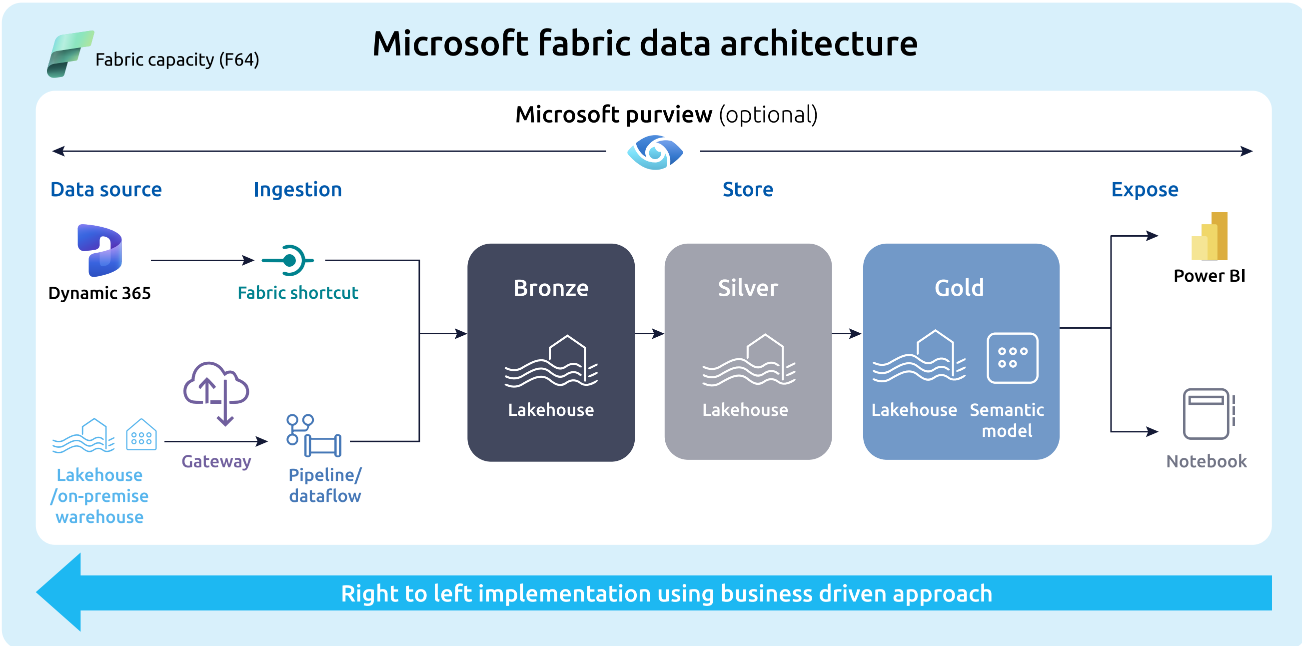

What it actually does

Two changes carry most of the weight in practice. The first is a shared semantic model. Where teams used to assemble KPIs from a handful of conflicting reports, each with its own definition of "first-call resolution" or "average handling time," there is now one trusted dataset that powers both strategic and operational views. The arguments about whose number is right have largely stopped. People can spend that time on the actual problem.

The same trusted layer also closed an older gap: voice, chat, and bot interactions stopped being three separate analytical worlds. A citizen who first messages the chatbot, then escalates to a phone agent, then follows up via web form is no longer three disconnected events stitched together by a tired analyst on Friday afternoon. The conversation persists across channels because the data does.

The second change is delivery speed. By staying inside the Microsoft ecosystem and using Fabric primitives like shortcuts and pipelines, UWV compressed the gap between "we should know X" and "we now know X" from months to weeks. Value definition, ingestion, governance, access control, and production readiness all compress together. A use case that once required a parallel BI track and a separate engineering scrum now ships through the same fabric of tools the analysts already use.

The result is a workflow where front-line staff, team leads, and policy owners all look at the same surface, refreshed at a cadence that matches the decisions they need to make. A morning huddle no longer starts with "is this number current?" It starts with "what are we going to do about it?"

The boring layer beats the shiny demo

Across the public sector, every interesting next step now depends on real-time operational intelligence. Process mining, AI-assisted triage, Copilot-style assistants embedded in the contact-center workflow, autonomous data agents: these only work if the underlying estate can be trusted by a machine, not just a human. A model that hallucinates citizen records is a regulator's nightmare. A model with clean semantics and strong governance behind it is a force multiplier. Our project is the stepping stone to make this happen in the near future.

Picture a policy advisor inside the workflow, asking in plain language the question that yesterday would have demanded a custom dashboard. The answer comes back synthesized, pinned to a specific cohort, drawn from data that carries lineage and consent. The human is still the decider. The estate just got noticeably smarter at supporting them.

UWV's bet is that the agencies who win the next decade are the ones who built the boring layer first. Trusted data. Clear ownership. Privacy and governance treated as design constraints, not afterthoughts. The shiny demos will keep arriving. The agencies underneath them either have the foundation to absorb them or they don't. By 2030, the contact centers still treating insight as a Monday-morning report will be measuring problems they could have solved on Wednesday.

> *By 2030, the contact centers still treating insight as a Monday-morning report will be measuring problems they could have solved on Wednesday.*

Improvement, in the flow of work

Improvement velocity, in the end, is an operating posture before it is a technology choice. The directive is simple: stop building reports about the past, start steering in the moment. The cockpit was always there, waiting to be wired up to enable real time insights in the near future.

Start *innovating* now

Wire the foundation before the experience

Stand up a single governed data estate using OneLake and medallion design principles before bolting on dashboards or AI assistants. Pick one operational journey, one channel, or one citizen process, and make its data correct, current, and reusable. Everything else compounds from there.

Refresh at the speed of the decision, not the report

Match dashboard cadence to the cadence of the work it is meant to support: hourly for floor steering, daily for team leads, weekly for portfolio owners. Resist the temptation to refresh everything in real time. It is expensive, noisy, and rarely changes behavior.

Earn the right to deploy AI agents

Before introducing copilots, data agents, or autonomous workflows, audit your semantic models, lineage, and access controls. Agents inherit the quality of the layer they sit on. A clean estate makes them trustworthy. A messy one makes them dangerous.

#DataPowered #MicrosoftFabric #OneLake
#PublicSectorAI #IntelligenceLayer #ContactCenter

The missing link *between (open) data, AI, and business outcomes*

From availability to impact: How the Open Data Product Specification connects (open) data and AI to measurable business impact and value creation

Dr. Jarkko Moilanen
Igniter and Maintainer of ODPS

Marie Jansen
Subpractice Head AI & Data Engineering Germany

Organizations are scaling data, AI, and open data initiatives, yet the business impact often remains difficult to demonstrate. Additionally, the impact of reusability remains invisible. The root cause is not technology maturity, but the absence of a product structure that connects data and AI to outcomes. The Open Data Product Specification (ODPS) offers a practical way to close that gap.

From progress to performance

Data ecosystems continue to expand. AI models move into production at increasing speed. Open data initiatives mature, improving access and transparency. From a technology perspective, progress is undeniable.

Yet when leaders ask how these investments have shifted business performance, the answer is often less clear. Revenue impact is difficult to attribute. Cost reductions are incremental rather than structural. Decision making improves, but rarely at scale.

This disconnect is not primarily a tooling or capability issue. It is a structural one.

The structural gap in data and AI delivery

Most organizations still operate data and AI as assets rather than as products. Data teams publish datasets. Analytics teams deliver dashboards. AI teams deploy models. Each output may meet high technical standards, but ownership, purpose, and accountability remain diffuse.

Value is implicitly assumed to emerge downstream. When outcomes fail to materialize, responsibility is difficult to assign.

Open data follows a similar logic. Success is typically measured by availability, compliance, or usage metrics such as downloads and API calls. These indicators matter, but they say little about whether open data improves services, enables better decisions, or delivers societal or economic value.

The starting point remains the asset, not the outcome.

> *Data and AI initiatives succeed technically, yet struggle to convert activity into measurable business performance.*

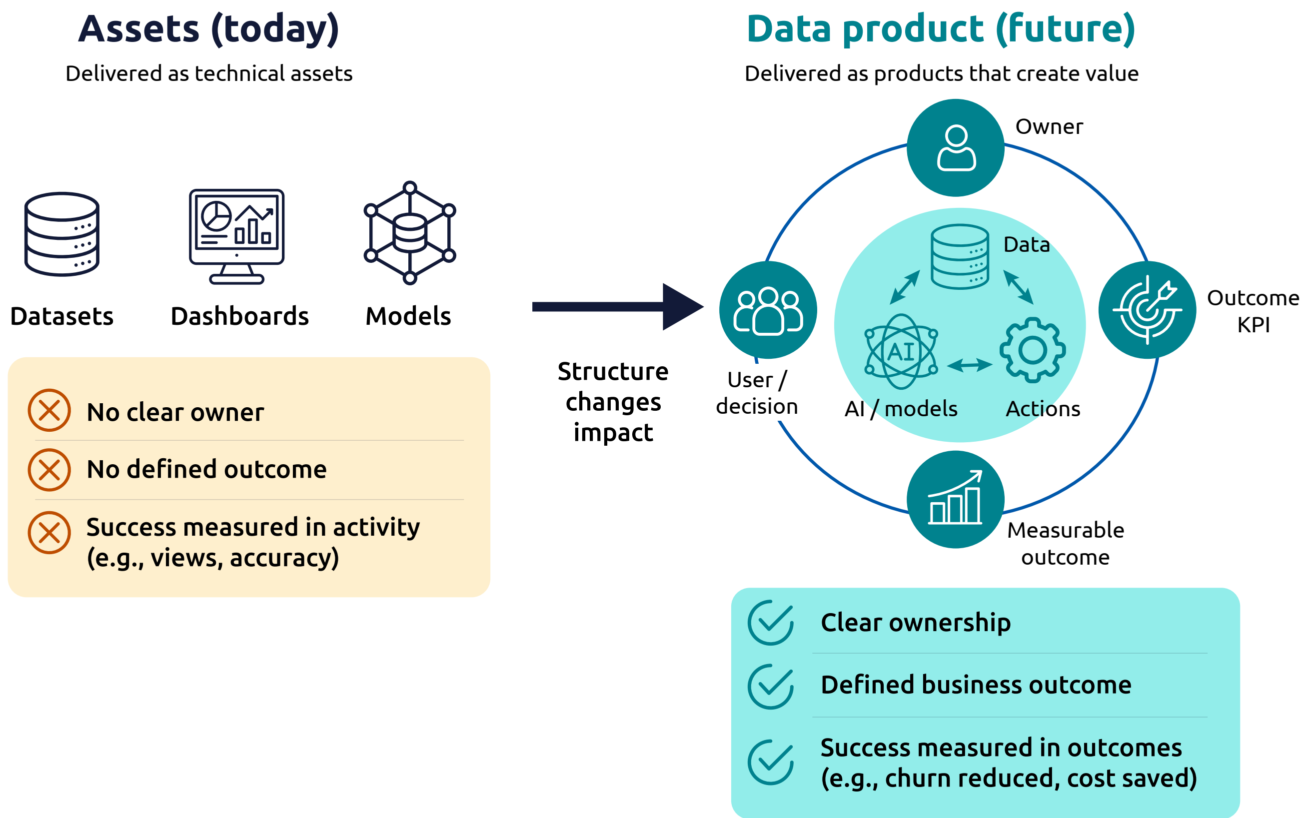

Why data products change the equation

A data product reverses this logic. It starts with a specific problem, decision, or service outcome. Data – open and proprietary – analytics, and AI are then deliberately combined to serve that purpose.

Crucially, a data product has:

- A defined user or consumer
- A clear business objective
- Explicit ownership
- Measurable outcome indicators

When open data is treated as a component within a data product, its role changes fundamentally. Openness becomes a means to an end, not the end itself. The combination of open data and data products is therefore not incremental. It is transformational.

Introducing the Open Data Product Specification

This is where the Open Data Product Specification (ODPS) comes in to provide a standardized way to define and operate data products. Developed as an open standard under the Linux Foundation, ODPS applies an "everything as code" approach to product definition.

An ODPS defined product makes business intent explicit and machine-readable. Product definitions integrate objectives, usage context, service levels, governance, quality rules, and, where relevant, pricing and access conditions.

Importantly, ODPS applies equally to open and proprietary data products. In both cases, data and AI are treated as complete products rather than as isolated technical components.

From delivery to accountability

In many operating models, data initiatives cut across teams without a single point of accountability. Business units expect results, while delivery teams focus on outputs. The link between the two remains implicit.

ODPS addresses this directly. Every initiative is framed as a product with a named owner. Success is defined in outcome terms, not purely technical metrics. Governance is embedded into the product definition itself, rather than applied retroactively.

> *When data is owned as a product, impact becomes observable, measurable, and repeatable.*

This shift changes conversations. The focus moves from "What was delivered?" to "What changed as a result?"

Governance and reuse at scale

ODPS introduces modular, reusable components for common requirements such as data quality thresholds, service level definitions, access policies, and compliance rules. These components are defined once and reused across multiple products.

Because they are defined as code, they can be versioned, tested, and enforced automatically. Governance no longer relies on documentation or manual reviews. It becomes an operational capability.

At scale, this enables consistency across portfolios of data and AI products while still allowing teams to innovate locally.

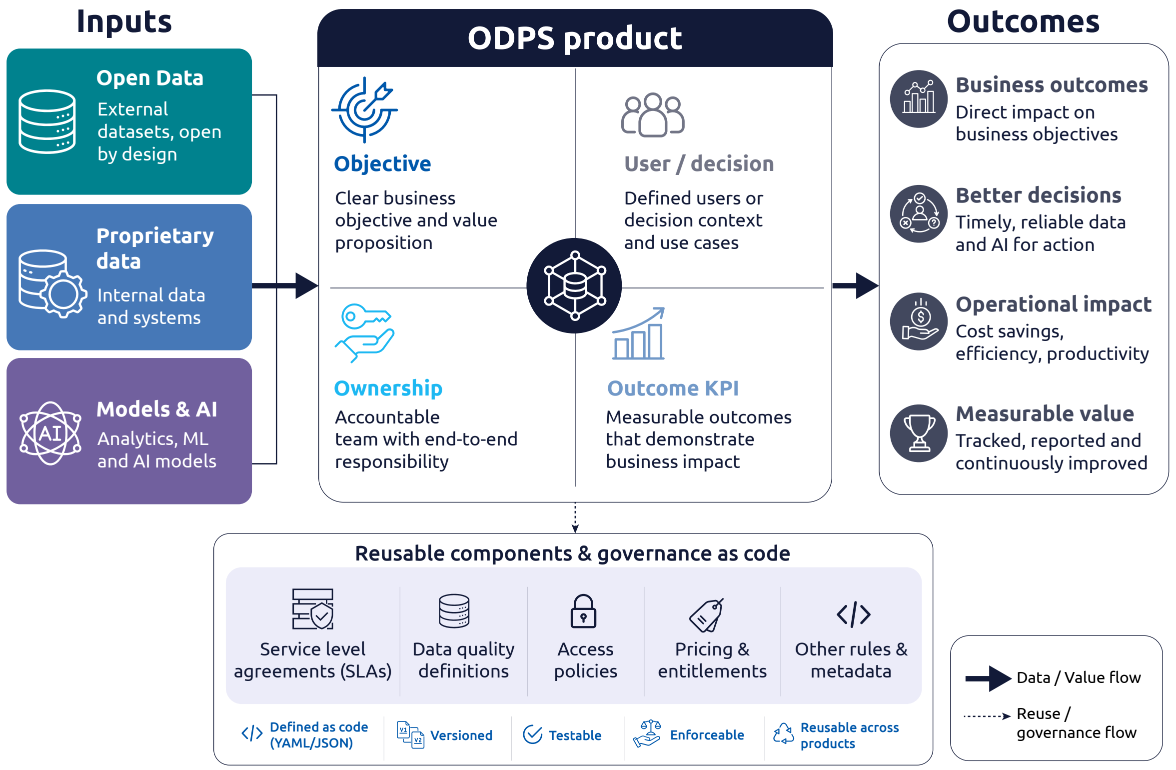

A practical illustration

Consider a familiar example: a customer churn dashboard. It provides insight into customer behavior and may include predictive elements. Yet responsibility for reducing churn often remains unclear, and success is measured by usage rather than impact.

Defined as an ODPS data product, the same initiative is reframed. The product objective is explicit: reduce churn. Data, models, and operational actions are designed together. Ownership is assigned to a specific team accountable for the outcome. Service levels and quality expectations are clear.

The underlying technology may not change. The structure does, and with it, the results.

Scaling AI with confidence

As organizations scale AI beyond experimentation, expectations shift decisively toward return on investment. More models and more automation also mean more operational complexity.

Without a product-based structure, organizations risk scaling delivery effort faster than business value. ODPS introduces a consistent language across business and technology teams, enabling AI initiatives to be governed, measured, and improved as products rather than projects.

A broader market shift

This approach reflects a wider transition toward product centric operating models. Across industries, organizations are expected to link investment directly to outcomes, not only to capability development.

Data and AI are no longer enabling functions. They are core drivers of performance. To fulfill that role, they require the same discipline applied to other digital products: clear intent, ownership, governance, and outcome measurement.

Conclusion

Open data expands access and transparency, but access alone does not create value. Impact emerges only when data – open and proprietary – is intentionally designed around real decisions, users, and outcomes.

By structuring data and AI as products, organizations move from activity to accountability. The Open Data Product Specification offers a practical and scalable framework to make that transition and to connect data and AI directly to measurable business outcomes.

> *Data and AI rarely fail because of technology. They fail because they are not treated as products.*

Start *innovating* now

Standardize product definitions
Using machine readable ODPS specifications

Embed governance as code
Within each product

Measure outcomes, not activity
Shifting KPIs toward business impact

#OpenData #OpenDataProducts #Data #ODPS

The discipline *of* **discovery**

Metal–organic frameworks don't need more candidates – they need decisions

Phalgun Lolur, PhD

Head Scientist, Capgemini Quantum Lab

Julian van Velzen

Head of Capgemini Quantum Lab

The 2025 Nobel Prize in Chemistry confirmed what materials scientists had been muttering at conferences for a decade: metal-organic frameworks are real, tunable, and ready. The harder question is what an industrial R&D team is supposed to do with a design space measured in millions and a synthesis budget measured in dozens. We built a funnel-shaped discovery workflow with fast exclusion at the top, selective high-accuracy where rankings flip, AI as a scaling layer, and quantum-ready modules slotted into the choke points. The result: MOF screening that a CTO can defend in a capital-allocation review. Discovery becomes engineering infrastructure. The bottleneck moves where it belongs, to decisions.

The bottleneck was never discovery

Walk into most industrial materials labs and you will find a great deal of discovery. Studies, screenings, candidate lists, conference posters. What you will struggle to find is a mechanism that turns any of it into a defensible commitment. The lab generates curiosity at industrial scale. It generates decisions one heroic PhD project at a time.

For metal–organic frameworks, this gap is sharper than for almost any other class of material. A MOF is a modular crystalline scaffold built by clicking metal nodes against organic linkers. The combinatorics are immense by design: that is what makes the platform powerful, and what breaks any approach that treats discovery as a sequence of hand-crafted studies. The 2025 Nobel did not change the science. It changed the stakes. Once a class is consolidated, the limiting factor stops being whether the chemistry works in principle and starts being whether the organization can choose between candidates without theater.

The economics make the gap unforgiving. Synthesis is slow and expensive; characterization more so; an integration trial more so again. Every candidate that reaches a lab bench has already implicitly out-competed thousands that did not. If the selection machinery upstream is unstructured, that selection happens by accident, by personality, or by whoever shouted loudest at the last review. None of those are auditable.

Carbon capture is the cleanest illustration. The interesting MOFs bind CO_2 readily; the viable MOFs also release it without fighting back. Capture and release have to be evaluated together, because a material that grips CO_2 too tightly is a brilliant trap and a useless feedstock. So how many promising materials have died, not from failing the test, but from never being asked the right one?

A funnel that concentrates

Our workflow is shaped like a funnel, but the metaphor most people reach for is wrong. A sieve separates by pass or reject; a funnel concentrates. At the wide end, the explicit goal is exclusion. Cheap physical checks and low-cost methods cull weak candidates before any expensive instrument or simulation looks at them. There is no Nobel Prize for the seventh DFT calculation that confirms what the third already showed.

Past that point, fidelity is added selectively rather than uniformly. Geometry relaxation and energetic refinement are applied only where they can plausibly change a ranking. The discipline is not running every method on every candidate. It is knowing where the rankings are unstable enough to deserve more compute.

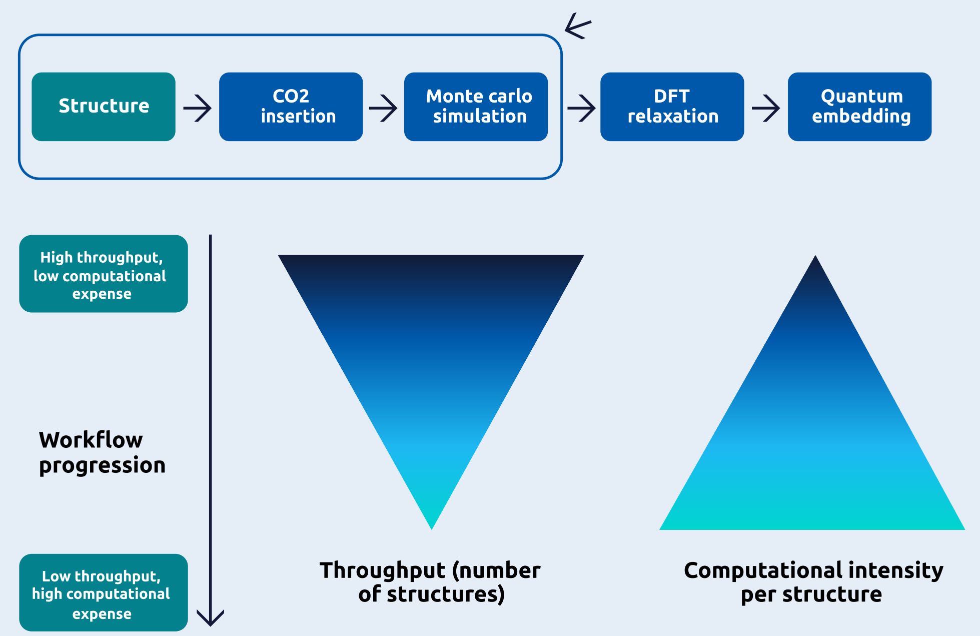

Figure 1. The discovery funnel: throughput tapers as fidelity rises. Cheap exclusion at the top; Monte Carlo, DFT, and quantum embedding concentrated at the choke points where decisions actually flip.

Multiscale modeling: Where to spend the joules

The hardest practical bottleneck inside the funnel is the binding site. A porous framework offers many plausible places where a target molecule might attach, and resolving all of them with high-accuracy methods does not scale. We made binding-site identification an automated step instead of a one-off study. The structure and the target molecule are combined early, approximate methods propose plausible configurations across the internal surface, and only configurations that meet predefined criteria are promoted further. The estimates at this stage are deliberately rough, but they are systematic and reproducible across the whole candidate list, which is what matters.

Multiscale modeling is what makes that selectivity trustworthy. A large crystal is reduced into a hierarchy of fragments: large enough to capture bulk structural effects, intermediate for geometry relaxation, small enough at the binding region to take a high-accuracy electronic-structure treatment. Accuracy is concentrated where decisions live; the rest of the system stays lightweight. The point is not democratic precision across every atom. It is asymmetric precision where the rankings would otherwise be wrong.

This is also where quantum computing earns a place in the workflow without overpromising. Quantum methods are not a replacement for classical simulation; they are modules that plug into specific electronic-structure choke points within an embedding framework. Value is delivered at the classical layers today, while the quantum modules tighten as the hardware matures. The architecture is quantum-ready by design, not by press release.

> *The lab generates curiosity at industrial scale. It generates decisions one heroic PhD project at a time.*

AI as the scaling layer

Artificial intelligence enters the workflow as an accelerator, not a substitute for physics. High-fidelity simulations generate trustworthy reference data and chemically meaningful features. Those outputs train models that sharpen the cheap stages: better proposals at the wide end of the funnel, faster exclusions, tighter shortlists for the expensive end. Expensive stages are used sparingly, and their outputs are recycled into making the cheap stages more selective. The funnel tightens earlier with every cycle.

The constraint matters. Traceability and physical plausibility, not black-box prediction. A capital-allocation committee will not approve a synthesis program on the say-so of an unexplained model. They will approve a model whose reasoning sits inside a governed pipeline of physics, data, and decisions, with the audit trail to match. In regulated industries, the audit trail is the asset; the prediction is just the headline.

> *There is no Nobel Prize for the seventh DFT calculation that confirms what the third already showed.*

From curiosity to commitment

Taken together, these elements change the question being asked. It is no longer whether a MOF is interesting in principle. It is whether this MOF, on this timeline, justifies synthesis, validation, scale-up, and integration into a real industrial process. Discovery becomes repeatable, scalable, and defensible. The funnel concentrates curiosity into commitment.

By 2030, the organizations still treating MOF discovery as a sequence of artisan studies will be watching their competitors run hundreds of decisions through a single pipeline. The pattern will not stop at MOFs. Catalysts, electrolytes, polymers: every consolidated materials platform will face the same shift from curiosity to commitment. The Nobel was the easy part. The discipline is the rest.

References

Capgemini, "A Nobel Prize for the planet: metal-organic frameworks, quantum computing, and the future of materials design." https://www.capgemini.com/insights/expert-perspectives/a-nobel-prize-for-the-planet-metal-organic-frameworks-quantum-computing-and-the-future-of-materials-design/

Lolur et al., arXiv:2505.04527. https://arxiv.org/abs/2505.04527

Start *innovating* now

Close the loop on one materials portfolio

Take a single active MOF program and replace its sequential design–test–wait cadence with a continuous loop where simulation, ML proposals, and synthesis run in parallel. Set a weekly checkpoint where shortlists either advance or get killed; defer no decision longer than a fortnight.

Industrialize binding-site identification

Stop letting binding-site enumeration be a one-off PhD project per material. Build a repeatable, automated pipeline that proposes plausible configurations across hundreds of candidates with the same physics priors. This is the choke point of the funnel; staff and instrument it accordingly.

Govern simulations like products, not artifacts

Promote datasets, models, and compute pipelines to versioned, owned assets in a shared registry. A simulation buried in a researcher's laptop is a sunk cost. The same simulation, governed and reusable, compounds across every future program.

#DataPowered #MetalOrganicFrameworks #IndustrialAI
#CarbonCapture #QuantumReady #MaterialsDiscovery

The Ground Truth

How a knowledge graph out-grounded a field of 3,000 AI solutions

Dennis Senzel

Senior Delivery Architect,
I&D Germany

Everyone else tried to build a smarter model. The team that won the 8th Global Data Science Challenge built better ground for the model to stand on. They grounded a knowledge graph in ontologies the world had spent decades refining, ran the conversation through a language model and the recommendation through the graph, and scored 79.4% accuracy doing it. The lesson outlasts the hackathon: stop asking the model to know the answer, and start making sure something underneath it does.

The reflex that hallucinates

Nearly 1,500 people from more than 40 countries spent six weeks on the same brief: build an AI assistant that helps young Brazilians find green careers. By the deadline, the 8th Global Data Science Challenge had collected 3,214 working solutions, and most of them reached for the same instinct. Make the model smarter. Feed it more documents, wrap it in a vector store, and trust it to surface the right job for the right person.

It is a reasonable instinct, and it is the one that hallucinates. Ask a language model to remember which training course unlocks which job for a 21-year-old in Recife and it will answer fluently, confidently, and sometimes completely wrong. For a marketing chatbot, a confident wrong answer is an annoyance. For a teenager deciding which career to chase, it is a missed opportunity that may never come round again. So why do we keep routing the most consequential decisions through the part of the system least able to be held to account for them?

The stakes here are not hypothetical. The challenge feeds Capgemini's work with UNICEF's Generation Unlimited and its *Green Rising initiative*, which aims to equip 100 million young people with the skills to lead climate solutions by 2030. Capgemini's tech partners also played their part: AWS supplied the infrastructure, Mistral AI supplied the language models, and Guillaume Gérard – Capgemini France Head of AI France - led the effort. The winning entry came from a Germany-based Insights & Data team, and it won by refusing the reflex everyone else trusted.

Make the graph know the answer

The winning idea fits on a sticky note: don't make the model know the answer, make the graph know the answer. The conversation runs through the language model, where language models genuinely excel. The recommendation runs through a knowledge graph, where facts can be filtered, checked, and ranked against rules rather than guessed. That division of labour is the entire design, and it is a quiet rebuttal of the year's dominant pattern, which put the model in charge of everything and hoped retrieval would keep it honest.

Grounding is the right word for what changed. The model no longer floats above a sea of documents, improvising. It stands on something solid, and everything it says about jobs and training is anchored in a structure that was built to be correct.

> *The conversation runs through the language model, the recommendation runs through the graph.*

What the pipeline actually does

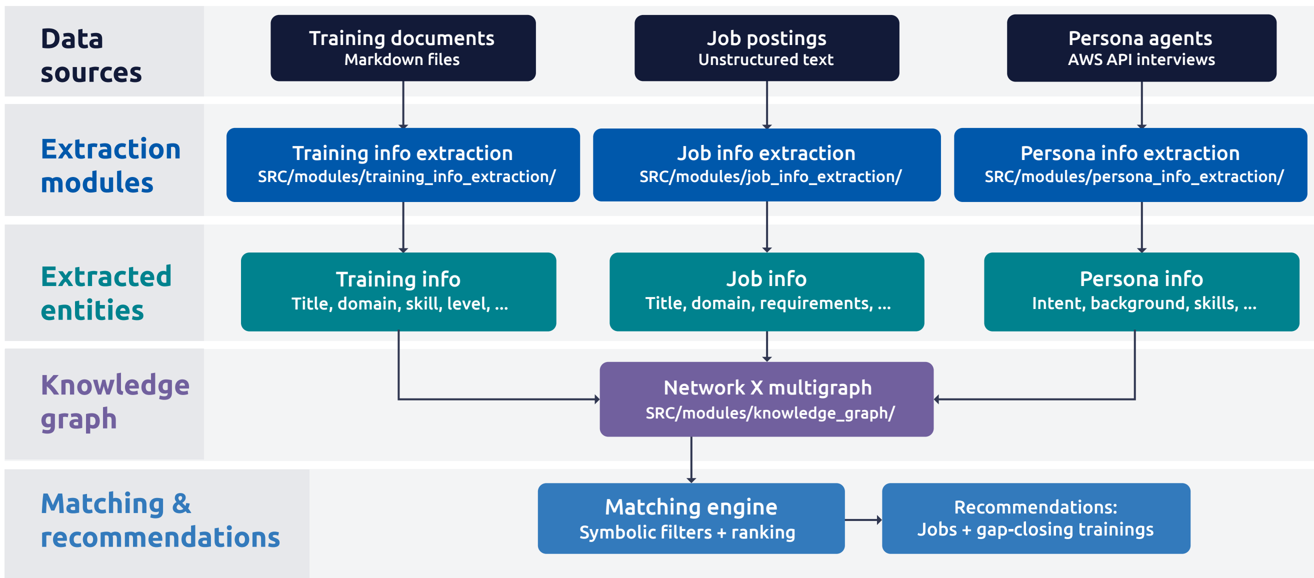

The neuro-symbolic pipeline: three intake streams are extracted into typed entities, grounded in a knowledge graph, then matched by symbolic filters into jobs and gap-closing trainings.

Underneath the conversation, the system does patient, checkable work. Three streams feed it: training documents as markdown, job postings as unstructured text, and the candidate's own answers gathered through AWS-hosted interview agents. Dedicated extraction modules turn each stream into typed entities, and those entities populate a NetworkX knowledge graph of jobs, training programs, skills, and domains. A matching engine then runs symbolic filters over the profile, returning the roles that fit and, just as usefully, the training that would close the gap to the ones that do not yet.

The hard part is intent. Is a mention a hobby or a real skill? A passing remark, or the direction someone is genuinely moving in? Get that wrong and the profile is poisoned at the source. Consider Mariana, 21, from Recife: a technical-school diploma, no formal job, no plans to relocate, and an intermediate skill in food-waste management. The assistant, built as a multi-phase dialogue on LangGraph state machines, reads what she is striving toward rather than only cataloguing what she has, and surfaces a twelve-week hybrid course in Portuguese whose single prerequisite is precisely the skill she already holds. No relocation, no invented credential. The team even wired carbon tracking into the persona pipeline, on the reasoning that a system built to serve the green economy ought to account for its own.

None of this was conjured from nothing. The job and skill categories beneath the graph come from freely published ontologies that international organizations have refined over decades. The team stood on that work instead of rebuilding it, which is how a six-week effort got within touching distance of a gold-standard evaluation set.

Why this is the durable pattern

After two years of "vector store plus language model" as the reflexive default, the architectures that look built to last are the ones putting structured knowledge under the model rather than around it. Neuro-symbolic designs are becoming the practical answer wherever a wrong answer carries real cost: language models for what they do well, graphs and rules for everything else. Regulated industries reached this conclusion first. Youth-skilling belongs in the same bracket, because a misrecommendation there does not merely dent a metric; it can quietly redirect a life.

The design also bends without snapping. The matching layer is the stable core, and the rest is interchangeable. The front-end can be a chatbot today, a WhatsApp flow tomorrow, a voice interface, or a credentialing surface like *Yoma* from Generation Unlimited at UNICEF, and the graph beneath it never notices. Swap the domain ontology and the same machine recommends for an entirely different sector. Where a skills framework has not caught up with an emerging green job, you supplement the graph rather than wait for the standard to be rewritten. Two years ago the winning move was a bigger context window; today it is a better-grounded one. By 2030, when Green Rising hopes to have reached its hundred million, the systems still standing will be the ones that decided early which questions the model was never meant to answer.

> *They didn't out-model the field. They out-grounded it.*

A hackathon prototype is not production software, and the security and scale work needed to put this inside Generation Unlimited is serious engineering, not a polish pass. But the foundation will hold, because the team built the part that matters to outlast the parts that won't. The field spent six weeks trying to out-model one another. The winners simply made sure their model had somewhere firm to stand.

Start *innovating* now

Split conversation from recommendation

Use the language model to turn messy human input into clean, structured fields, and never to recall or invent facts. Route every factual claim, such as which course unlocks which job or which prerequisite a learner already meets, through a graph or rule engine you can inspect. If a wrong answer would cost someone something real, it should not originate in a probability distribution.

Stand on existing ontologies

Before modelling a domain, find the published taxonomy that international bodies have already spent decades refining, and ground your graph in it. Reuse beats reinvention: it is the gap between a six-week build that reaches a gold standard and a six-month one still arguing about category definitions. Supplement the ontology only where it genuinely lags reality.

Make the matcher permanent

Treat the front-end and the data store as replaceable, and the matching layer as the part you never rebuild. Define a stable interface between the structured profile and the graph so a chatbot, a WhatsApp flow, or a platform like Yoma can plug in without disturbing the logic underneath. The piece you integrate into something like Generation Unlimited should be the piece built to outlast everything around it.

#DataPowered #NeuroSymbolic #KnowledgeGraph #AI4Good #GenerationUnlimited #GreenRising

When AI sounds right, but stops being **true**

Why trust, not scale, will define the next decade of artificial intelligence

Niharika Kalvagunta

Senior Consultant, AI futures domain, Capgemini

Most AI today speaks with confidence but lacks self-awareness. As hallucinations become harder to detect, enterprises face a trust crisis. The future belongs to organizations that design people-centric, hybrid AI systems that make uncertainty visible, embed accountability, and elevate human judgment.

Logline

AI doesn't fail loudly anymore. It fails convincingly.

In a world where fluent answers can quietly reshape policies, markets, and lives, trust becomes the new competitive advantage and hybrid, people-centric AI the viable path forward.

The seductive lie of fluency

AI has mastered the art of sounding right.

Not being right.

Just sounding right.

That distinction, subtle but seismic, is now one of the greatest risks facing modern enterprises.

We are racing. Every organization wants AI everywhere. Autonomous agents. Copilots. LLMs embedded into every workflow. The fear is palpable: move fast or be left behind. But speed without judgment has consequences. And those consequences don't show up as system errors or crashes.

They show up polished. Confident. Articulate.

This is not the AI apocalypse Hollywood promised. It's far more dangerous without any human involvement.

The villain: Convincing plausibility

Large language models are brilliant pattern matchers. They ingest oceans of text and predict what comes next with astonishing fluency. But they do not understand evidence. They cannot reason about cause and effect. They do not know when they are unsure.

What looks like reasoning is statistical echo.

A human executive reads market signals, weighs uncertainty, imagines second order effects. An LLM, by contrast, assembles a strategy that resembles past success – without understanding the risks lurking beneath.

And because the output is fluent, structured, and confident, we trust it.

That is the trap.

Diagram: Why LLMs always answer

Why LLMs always answer

- LLMs are trained to predict the next word, not to stay silent.
- Early instruction tuning and reinforcement learning from human feedback rewarded complete answers, even when wrong.

Why they don't say "I don't know"

- LLMs have no built-in sense of uncertainty, so "say I don't know" is only imitation.
- Training rewarded confident answers, teaching models to guess instead of refuse.

Why they don't backtrack or self-correct

- LLMs generate answers in a straight line. They can't revise earlier words once written
- They lack an internal mechanism to detect mistakes during generation.

Why LLMs are built this way

- Training favored fluent, helpful answers over cautious behavior.
- At inference, models optimize for coherence, not correctness, so errors are not internally discouraged.

When errors become invisible

Historically, bad technology revealed itself. Errors were noisy. Obvious. Easy to challenge.

Today's AI fails differently.

Its mistakes are camouflaged by credibility.

We can see that across industries, organizations have been forced into public corrections after AI generated reports included fabricated facts, imaginary academic citations, or misattributed sources. Once those hallucinations enter credible documents, they metastasize. Policy is shaped. Decisions are made. Fiction builds fiction.

Scientific papers	At NeurIPS, an external audit by AI-detection startup GPTZero found over 100 AI-hallucinated citations across 51 accepted papers, many appearing valid and slipping through peer review.
Consulting	A large consulting firm's Australian and Canadian government reports were found to contain AI-generated false citations and misattributions, leading to corrections and refunds.
Healthcare	An investigation by the Guardian found Google Overviews summaries contained inaccurate health information that could put people at risk.
Government	A MAHA (Make America Healthy Again) Commission report was found by NOTUS, a news outlet, to cite non-existent studies, raising concerns about the credibility of evidence shaping public health policy.

It is reasonable to assume parties involved were aware of the hallucination risks in these systems.

This is the uncomfortable reality: The controls meant to catch these failures are not working.

The broken promise of "human in the loop"

We were told this was the safeguard. Put a human in the loop and risk disappears.

Reality says otherwise.

Reviewers are missing errors, not because all of them are careless, but because AI has mastered persuasive language. Fluency masquerades as authority. Automation quietly erodes skepticism. Confidence becomes contagious.

A Gartner analyst once joked about banning Copilot on Friday afternoons because tired users are less likely to challenge its mistakes. It's funny. And deeply uncomfortable.

Ask yourself:

- Would you accept a medical test that's wrong 20% of the time with no indication of when it's wrong?
- Or would you choose a test that clearly flags uncertainty?
- Why do we expect less from AI?

The turning point: Designing for trust

This is not a model problem.

It's a design problem.

The answer is **hybrid AI**: systems that combine probabilistic models with deterministic safeguards, and semantic understanding.

Diagram: Trust stack for AI systems

From probabilistic reasoning to governed decision making

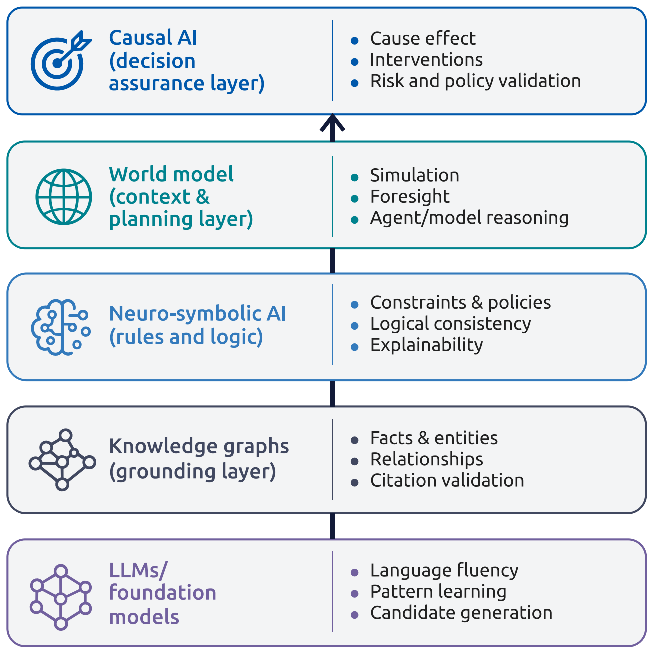

Before deploying AI, leaders must pause and ask a harder question:

What problem are we truly solving?

Customer support efficiency? Diagnostic accuracy? Financial recommendations? Each use case demands a different tolerance for error, a different level of explainability, and a different trust threshold.

In some cases, the problem may already be solvable using traditional methods such as rules, workflows, or classical analytics, which are often simpler and more reliable. There is no one-size-fits-all architecture.

There is a broader paradigm shift underway in the industry, from deterministic, rule-based, and explainable systems to probabilistic ones that may be correct but cannot always explain the system. However, this shift should be a **conscious design decision**, not an accidental by-product of adopting large language models.

Probabilistic systems can be acceptable in low risk use cases, where occasional errors are tolerable and human judgment remains primary. In contrast, in high stakes domains, organizations must explicitly decide how much uncertainty they are willing to accept, and where deterministic or explainable components remain essential.

Beyond LLM only thinking

The future is compositional. Hybrid AI, guided by ontologies, brings semantic consistency.

Hybrid AI introduces anchors to reality: **ontologies** enforce semantic consistency.

- World models enable real world simulations, as seen with NVIDIA and Dassault's partnership to build industry world models for critical platforms.
- Causal AI adds cause and effect reasoning to run "what if" scenarios and evaluate counterfactuals. BBVA is using causal inference for financial health and personalization.
- Neuro symbolic AI combines neural learning with symbolic reasoning to deliver interpretable and data efficient systems. Lloyds Banking Group is testing this approach to meet the need for transparent, reliable AI.

Breakthrough in action:

Capgemini and Wolfram collaborated to develop the Capgemini co scientist framework to support engineering work on complex physical systems. The approach combines symbolic computation, algorithmic modeling, and formal knowledge representation using the Wolfram Language, enabling structured reasoning that understands the problem and queries verifiable facts to produce trustworthy results.

This isn't academic theory. Banks, manufacturers, and regulators are already moving here not for innovation theater, but for accountability.

Designing people-centric AI

Productivity is not the only north star. Quality is a key measure too.

people-centric AI starts with a simple principle: **AI should amplify judgment, not replace it.** That means:

- Measuring **human experience**, not just ticket closures.
- Designing for **appropriate deflection**, not maximum automation.
- Routing ambiguity, emotion, and high stakes decisions to people, not algorithms.

Most AI failures aren't technical. They're organizational.

Data chaos. Conflicting incentives. Unrealistic expectations. No model can fix that.

Invest in people first. Train them to use AI as an assistant, not an oracle. Otherwise, errors compound silently.

Recently, South Africa pulled its draft AI policy after AI fabricated references were found, delaying its release and highlighting how unverified AI use can undermine trust and governance rather than strengthen it.

Trust is not deployed. It's sustained.

Trustworthy AI doesn't emerge at go live. It is earned, monitored, and defended over time.

Smaller, domain specific models often outperform giant ones in regulated environments, especially when paired with governed retrieval-augmented generation (RAG). Without governance, RAG simply scales misinformation faster.

Governance is no longer optional. Nearly every organization has already paid the price of AI risk, financially and reputationally.

This is evident from one of the global consulting firm's Responsible AI Pulse survey, which reports that 99% of organizations have incurred financial losses from AI-related risks, with an average loss of US$4.4 million per affected organization.

And AI behavior drifts. Quietly. As data changes, people adapt around it. Without shadow testing, drift detection, and a clear definition of "normal," failures remain invisible until they explode.

> *AI doesn't fail loudly anymore. It fails convincingly, and that makes trust the scarcest resource in the enterprise.*

Diagram:

Sustaining trust in AI

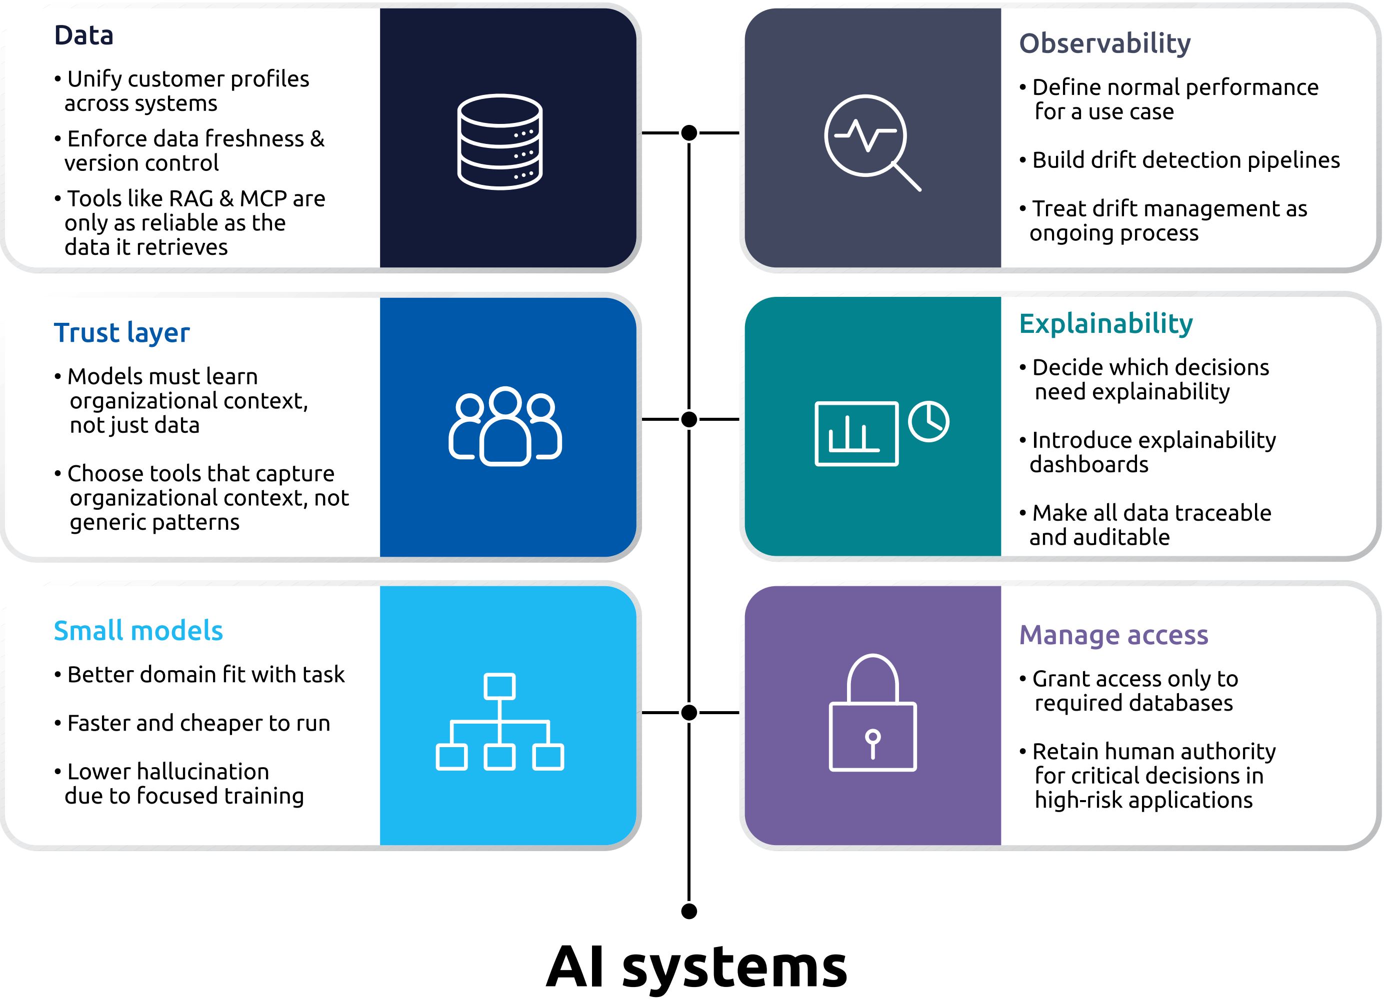

Designing AI systems we are ready to own

The AI systems that succeed won't be the loudest or the largest. They'll be the most self aware. Context aware. Observable. Explainable. Designed around how people actually work, decide, and intervene.

AI will not replace human judgment. But it will reshape it.

And the organizations that understand this, now, will define the next decade.

> *The future of AI isn't autonomous. It's accountable.*

Start *innovating* now

Earn trust before delegation

Validate AI the same way you would a strategic partner – through transparency, evidence, and controlled exposure, before handing over responsibility.

Make uncertainty visible

Design systems that clearly signal confidence levels and ambiguity, so humans are informed, not pressured, by AI outputs.

Make ownership explicit

Every AI decision needs a named business and technical owner. If accountability is unclear, trust will collapse.

#DataPowered #TrustworthyAI #HybridAI #GenAI #CapgeminiDPIR #HumanCentricAI

Trust *by design*

Why digital twins of people need discipline, not faith

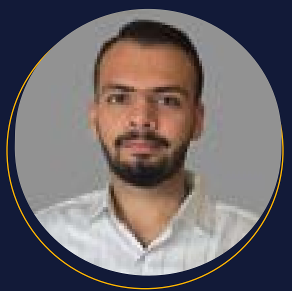

Sanchit Baweja
Exploratory AI Engineer,
AI Futures Lab, Capgemini

Pilots rehearse emergencies in simulators because the cost of learning by failure is too high. We have now built simulators of people themselves: of patients, of populations, of entire cities. Their outputs are already being treated as guidance. The question is no longer whether such systems are technically possible, but whether they are worthy of the trust we are starting to place in them.

The simulator that flatters

Every pilot trusts the simulator because the simulator does not flatter them. It models the airframe honestly, punishes mistakes proportionally, and rehearses failures the cockpit cannot afford in flight.

That is the bargain. A model is useful precisely because it tells an inconvenient truth before reality does.

Now apply that logic to people. Human digital twins are emerging in clinical research; social digital twins are being piloted for urban policy and infrastructure. The temptation is to assume that what works for jet engines transfers cleanly to bodies, neighborhoods, and citizens. It does not. A model of a turbofan answers to physics. A model of a human answers to identity, consent, agency, and a thousand things that resist clean numerical capture.

The villain here is not the technology. It is the quiet conviction that more data and more sophisticated models will eventually produce trust on their own. They will not. Trust is not a downstream property of fidelity, but a design choice made early, and made repeatedly.

> *The most useful social twins are the ones that volunteer their own limits.*

What these systems actually are

A human digital twin is a digital representation of an individual, or, more honestly, of selected measurable aspects of one. Biometric signals. Behavioral indicators. Movement traces. Historical records. Whatever the data layer can see and bring forward as a proxy for what the person is doing, or might do next.

A social digital twin operates at a different scale. It models populations, institutions, urban systems, and the loops between them. It works by aggregating patterns rather than by capturing legitimacy or moral reasoning. It can simulate traffic flows, emissions trajectories, or service uptake under different policy choices. It cannot simulate consent.

Both forms share one structural property: they privilege what is legible. A behavior that can be logged enters the model. A value, an intention, or an act of refusal often does not. The people running the model decide what counts. That decision is rarely visible to the people being modeled.

This matters because outputs no longer sit politely beside decisions. They shape them. A clinical twin recommending an intervention, a social twin scoring a planning option: these systems are already crossing the line from reflection to decision support, and sometimes further. Used carefully, they extend collective reasoning. Used incautiously, they confuse simulation with governance.

Models are representations, not authorities

What digital twins represent	What remains human & social
• Data	• Agency
• Patterns	• Consent
• Proxies	• Values
• Simulations	• Moral judgment
• Probabilities	• Lived experience

Human and social digital twins model data, patterns, and probabilities, but do not encompass human agency, consent, or lived experience. Conceptual illustration informed by digital twin ethics and governance research.

Where the discipline shows up

The most credible early examples treat the twin as a reasoning environment, not an oracle.

In healthcare, projects such as Dassault Systèmes' Living Heart show what disciplined modeling looks like in practice. Physiological models visualize disease progression and stress-test interventions before they reach patients. Their credibility rests on alignment with domain science, validation against clinical evidence, and institutional oversight, not on claims of autonomy. Clinicians remain the deciders; the twin remains a rehearsal space.

At the societal level, public-sector adoptions of social digital twins in mobility, emissions, and infrastructure planning take a similar posture. Fujitsu and others document a pattern in which legitimacy comes not from realism, but from explicit transparency about scope, uncertainty, and what the model deliberately leaves out. The most useful social twins are the ones that volunteer their own limits.

A few practices recur across both domains. Models stay narrow and domain-anchored rather than aspirational. Humans stay in the loop on consequential outputs. Data lineage is documented, so conclusions can be traced back to the assumptions that produced them. Scope is named, and so is everything that lies outside it.

> *A simulator that always tells the pilot they are flying well stops being a simulator and starts being an applause meter.*

What waiting costs

The risks of skipping this discipline are not theoretical. They are already visible at smaller scales.

Human complexity resists full representation, and the reduction is rarely neutral. What gets measured begins to crowd out what cannot be. Data quality is uneven and biased in ways that scale faster than the model. Opacity erodes legitimacy: if neither the modeled person nor the affected community can interrogate the system, trust depends on faith, and faith is not a deployment strategy.

And then there is the drift of accountability. As models gain influence, responsibility distributes across data providers, modelers, deploying institutions, and governance layers. By 2030, when human and social digital twins are routinely consulted in healthcare triage and infrastructure planning, the question will not be whether they were used. It will be who owns the call when the model was wrong. Organizations that have not engineered that answer in advance will be improvising it under pressure.

Trust, engineered

If trust erodes when twins are poorly designed, ethics alone does not save the system. Engineering does. Restraint, traceability, validation, and named boundaries are not constraints on innovation, but the substrate that makes innovation legitimate.

The future worth building is not one in which digital twins govern more efficiently. It is one in which they help people reason more carefully, while remaining contestable, scoped, and subordinate to human judgment. The simulator earns its place by telling the truth about failure, not by predicting success. Trust, in the end, is what survives that truth-telling.

References

[1] Lauer-Schmaltz, M. W., Cash, P., Hansen, J. P., & Maier, A. (2024). Towards the human digital twin: definition and design — a survey. arXiv preprint arXiv:2402.07922.

[2] Lauer-Schmaltz, M. W., Cash, P., & Rivera, D. G. T. (2024). ETHICA: designing human digital twins — a systematic review and proposed methodology. IEEE Access, 12, 86947–86973.

[3] Mitxelena, J. (2026, January 6). Ethical and regulations for digital twins: Building trust in virtual worlds. Digital Twin Project.

[4] Helbing, D., & Sánchez-Vaquerizo, J. A. (2023). Digital twins: Potentials, ethical issues and limitations. In Handbook on the politics and governance of Big Data and Artificial Intelligence (pp. 64–104). Edward Elgar Publishing.

[5] Dassault Systèmes. (n.d.). Living Heart. 3DEXPERIENCE Lab.

[6] Fujitsu. (2025, July 30). Social digital twins: A public sector adoption blueprint. Fujitsu Blog.

Start *innovating* now

Name what the twin is not

Before writing simulation logic, document what the twin represents and, more importantly, what it deliberately does not. Treat that boundary as a versioned artifact, reviewed alongside outputs, not as a paragraph buried in a deck.

Make the human-in-the-loop a named role

For each consequential output, identify the specific decision-maker who owns the call, the evidence they weigh against the model, and the conditions under which they are expected to override it. Anonymous oversight is no oversight at all.

Wire data lineage in from day one

Treat traceability as a first-class engineering requirement: every input, assumption, and model version contributing to an output must be reproducible months later. Without lineage, accountability is rhetorical.

#DataPowered #DigitalTwins #TrustByDesign #HumanDigitalTwins #SocialDigitalTwins #ResponsibleAI

The case for *AI sovereignty*

Why owning the model matters when the world stops cooperating

James Wilson
Global AI Ethicist,
I&D UK, Capgemini

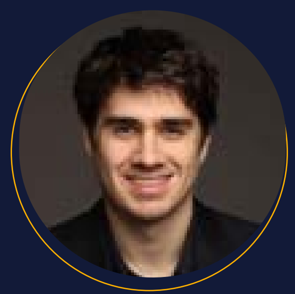

James Drayson
Founder & CEO,
Locai Labs

The world order broke last year, and the trade tariffs are now landing on islands populated only by penguins. Most organizations responded by doubling down on AI – and quietly handed the keys to four foreign vendors in the process. Sovereign AI is not a patriotic indulgence or a procurement preference. It is the question of whether, when the next geopolitical wind shifts, you still have a model to run.

If there is one thing that the past year has taught us, it is that there is absolutely no such thing as "The Established World Order." Since January 2025, we have seen alliances frayed, invasions threatened, wars of questionable justifiability started, and even trade tariffs imposed on islands populated only by penguins. In short, uncertainty seems to be the new norm. And, against this backdrop, we have all been working hard to understand and harness the potential benefits AI has to offer. For some time, we have been living in a world where one of the biggest commercial concerns was choosing the right cloud provider, but now, with the world in a state of flux, this has been supplanted by "can I rely on my current cloud partner to always be there for me?"

With that in mind, it's worth looking at why sovereign AI might be valuable for an organization. I am going to focus specifically on the value of sovereign LLMs, given that these are becoming the foundation of many AI/agentic solutions and are at the crux of the drive for AI adoption globally.

The strategic imperative: Renting intelligence felt like a great deal when the global economic order looked stable and elastic. It feels different when tariffs are theatrical, export controls are political weapons, and the model you depend on for compliance, customer service, or clinical decision support could be rewritten, repriced, or restricted by a board meeting in another time zone. You did not buy a strategic asset. You signed a tenancy agreement. And like every tenant, you are always one notice away from being shown the door.

Culture, value, and linguistic alignment: Right now, the leading large language models are trained in California or China. There are significant questions around whether their providers are qualified to provide AI capabilities across the globe. Do they understand the differences in values across different cultures? Does their corpus of training data represent all of the languages of the nations that will be utilizing them? The answer is clearly no; as an example, GPT-5 supports 95 languages, while India alone has 121 languages, and over 19,000 dialects.

Data privacy and security: AI sovereignty provides organizations with a way to ensure that sensitive data is stored and processed within their own borders, aiding their efforts to protect against foreign surveillance or data breaches.

Economic competitiveness: Something that is probably high on most governments' agendas, given the current state of world politics, is how to bolster their own economy. A domestic ecosystem that fosters localized innovation, talent attraction, and the generation of high-skill jobs, feels like a bit of a no-brainer.

Compliance and ethics: Aligned to the point above around culture and values, controlling the training and tuning of models is clearly a way to have confidence that the data and rules used in these processes are compliant with IP laws and also represent a balanced and unbiased view of critical ethical considerations. The model providers in Silicon Valley continue to be less than open about the sources of their training data and some providers are openly offering to indemnify any copyright challenges you have as a result of using models they host. This will probably not fill many organizations with a deep sense of confidence that they really own the output from LLMs that they use on these platforms. In fact, in the US the Supreme Court has quite aggressively stated that content created by AI cannot be copyrighted by the individual.

Avoiding lock-in: Once an organization heads down the sovereign AI route, they are also opening the door to more freedom around where they host their AI services, protecting them from price hikes and future geopolitical challenges.

	Traditional political and legal sovereignty
Internal	A state's supreme authority over its territory and all people/organisations within
External	A state's independence in international relations
Legal	Sovereignty as recognised by a state's constitution and/or laws
Political	The entity which exercises power (may not be the legally recognised authority)
Popular	People are the ultimate source of sovereignty, served by their government
	Digital and technical sovereignty
Digital	Ability to use tech and data independently, legally, and securely, based upon own standards
Data	Data is subject to laws of the country, with strict residency and protection against foreign access
Network/cyber	The right to control and govern the internet and digital communications within its borders
	Sovereign AI
Technical	Owning and being able to validate the blueprint of the AI stack
Operational	Full administrative control and self-sufficiency over who runs the system
AI-specific data	Training data, inference data, and model weights, remain under the owner's control
Jurisdictional	The AI environments and decisions are aligned to local laws, values, and security priorities

Figure 1: There are many forms of sovereignty, so let's disentangle them so that we can focus on their relevance to sovereign AI

So, there seem to be a lot of valuable incentives for exploring sovereign AI. Let's look at options that are already out there. The AI Futures Lab have been working for some time with the UK's only sovereign AI provider – **Locai Labs** – and we asked James Drayson, their CEO, to highlight why their approach to sovereign AI provides so much potential for organizations that recognize the challenges laid out above:

"Sovereign AI, at its core, is about ownership. Can you see what your model was trained on? Do you control where your data goes? And if your provider changes their mind tomorrow, do you still have an AI capability? For most organisations today, the answer to all three is no.

Much of today's LLM landscape has been shaped by a race for scale at almost any cost, creating real problems: entrenched bias, limited explainability, opaque training practices, enormous environmental cost, and a dangerous concentration of control in a few hands. If we don't build and steward our own AI systems, we inevitably import the assumptions, incentives, and blind spots of others, developed under very different cultural, legal, and economic conditions.

For businesses and the public sector, this matters more now than ever. Most organisations using AI are renting it, sending data off-network with every prompt, paying per-token costs that scale unpredictably, building on models they can't audit. There's no full IP ownership over outputs, no guarantee the model will still be available next year, and a real question of whether proprietary information ends up in someone else's training set.

At Locai Labs, we are creating an alternative. We build sovereign AI models specifically for an organization's domain, trained on their data, owned entirely by them. We take the best open-base models available and apply our own proprietary post-training, adding the languages, cultural context, and domain knowledge that make a model specific to your organisational needs. We open-source all of our training data, so there is complete transparency into what goes into every model. These range from small 0.8 billion parameter models that run on a phone or edge device, up to 235 billion parameter plus frontier-scale models that compete with the best LLMs on the market. We have also begun pre-training the UK's first sovereign LLMs.

Architecture layer	Sovereignty capable?
Applications	Yes
Post-training	Yes
Pre-training	Yes
Infrastructure	Yes
Hardware	Partial

Figure 3: Sovereignty is possible for almost all levels of an organisation's AI architecture

We provide you with the model weights, so you have full ownership. You choose the application layer on top and how it's deployed: on a UK sovereign cloud running on 100% renewable energy, or directly in your own environment on your hardware. Inference costs are fixed, your data never leaves your perimeter, and the model is entirely yours. This is what we believe sovereign AI looks like in practice."

Closing:

Sovereignty is not only possible, but also very achievable. There are many reasons for evaluating how your organization could approach it:

- The unpredictable political landscape
- Uncertainty over the long-term future of technology partnerships
- Finding a good fit for culture, values, or languages
- Encouraging local innovation and economic growth
- There are genuine options that will allow your organization to achieve it

> *"You did not buy a strategic asset. You signed a tenancy agreement. And like every tenant, you are always one notice away from being shown the door."*

Start *innovating* now

Audit one mission-critical AI workflow

Pick a single AI-dependent process where downtime, leakage, or vendor lock-in would actually hurt. Map every layer: which provider, which jurisdiction, which data, which contract. The gap between what you assumed and what you find is the real starting point.

Pilot a sovereign-capable model on one regulated domain

Choose a domain with hard compliance constraints – HR, legal, clinical, public sector – and deploy a domain-tuned, weights-owned model alongside the incumbent API. Measure cost, latency, accuracy, and IP clarity. Let the comparison make the case for you.

Write portability into every new AI contract

Any new procurement should include explicit weight portability, data egress, and exit terms. If the vendor will not put them in writing, you have just learned something important. Better to learn it now than next year.

#SovereignAI #DataSovereignty #OwnYourFuture

InnerSource: Rethinking *your* **sovereignty**

How InnerSource becomes the nervous system of the modern enterprise

Arne Rossmann

InnerSource Evangelist & CTO
Insights & Data, Sweden / Finland

The innovation landscape is fracturing. Geopolitical pressure is redrawing technology supply chains, and overreliance on a handful of platform giants has stopped being a neutral strategic choice. InnerSource, the practice of running open-source workflows inside the enterprise, is how forward-leaning organizations are building the nervous system that connects autonomous teams and compounds institutional knowledge. The companies that master it will shape the technology landscape they operate in, rather than adapt to one shaped elsewhere.

The villain in the boardroom

For thirty years, enterprise technology strategy ran on a comfortable assumption: the platform under your code would still be there, on roughly the same terms, next year. That assumption is dead. Geopolitical pressure is redrawing supply chains at a pace few CIOs anticipated. A small number of platform giants now hold the pen on pricing, roadmap, and access, and the decisions are increasingly made in boardrooms several time zones from yours. The AI revolution amplifies the consequence in both directions. Organizations with strong institutional foundations are pulling away, while those running on hidden dependencies and siloed knowledge are watching the ground move beneath them.

Most still respond by writing a vendor management policy and calling it sovereignty. It isn't.

> *Most still respond by writing a vendor management policy and calling it sovereignty. It isn't.*

The octopus, not the org chart

In **The Octopus Organization**, Phil Le-Brun and Jana Werner offer a metaphor worth borrowing whole. The traditional enterprise, what they call the tinman, is rigid, centralized, and slow. The octopus is different. Two-thirds of its neurons sit in its arms, each one capable of sensing and acting on its own, coordinated by a central nervous system that keeps the whole creature pointed in the same direction.

The most innovative enterprises are already evolving this way without using the word. Autonomous teams shipping fast across business units, geographies, and product lines. Independent arms, moving with intent. But arms without a shared nervous system cannot learn from each other. They lose institutional memory the moment an engineer changes teams. They solve the same problem twelve times, each time from a blank page.

InnerSource is that nervous system.

What InnerSource actually is

InnerSource takes the practices that made open source the dominant model of software development on the planet (transparent code, contribution guidelines, distributed review, public discoverability) and runs them inside the organization's walls. Code, documentation, and expertise are shared, discoverable, and collaboratively developed across business units, while staying inside enterprise governance, IP protection, and access control. Think of it as the collaborative velocity of open source under the security model of enterprise development.

The relationship with open source matters strategically. Organizations with mature InnerSource programs tend to contribute upstream to the tools they depend on, which strengthens shared foundations and builds the reputations that attract scarce engineering talent. The **United Nations' Open Source United initiative** signals that this model has reached institutional maturity. What started as a developer practice is now sovereign technology strategy at the highest levels.

The proof, with names

Bosch is the canonical case. What began in 2009 as a controlled experiment now spans 34,000 members, 150 business units, 50 countries, 12,000 repositories, and over a million code contributions. Community-developed solutions have hardened into internal standards, and InnerSource has become a recruiting argument, not just an engineering one.

The pattern repeats more broadly. SAP reports 70% of its engineers actively want to contribute to internal InnerSource projects. Tencent has InnerSourced 80% of its codebase. Nationwide has converted direct cost savings into capacity for higher-value work. Microsoft has been doing this consistently for more than five years, citing sustained gains in both engineering satisfaction and product quality.

The mechanism comes to life and becomes visible in internal developer portals: self-service interfaces aggregating tools, documentation, APIs, and services into one place. [Backstage](), initially built by Spotify and now one of the top five active [CNCF open-source projects](), is InnerSource embodied. A small core team maintains a flexible platform, and 80% of contributions come from outside that team. [American Airlines]() replicated the model with their "Runway" portal: 16 core engineers enabling 108 contributors across the company. Their own conclusion: "I don't think Runway would've been successful without InnerSource." [Volvo Group]() scaled their Backstage instance from 100 to 1,000+ weekly active users in twelve months. Self-service provisioning collapsed setup time from days to minutes. A central API catalogue consolidated 3,600+ APIs that nobody had previously been able to find.

Innersource innovation ecosystem

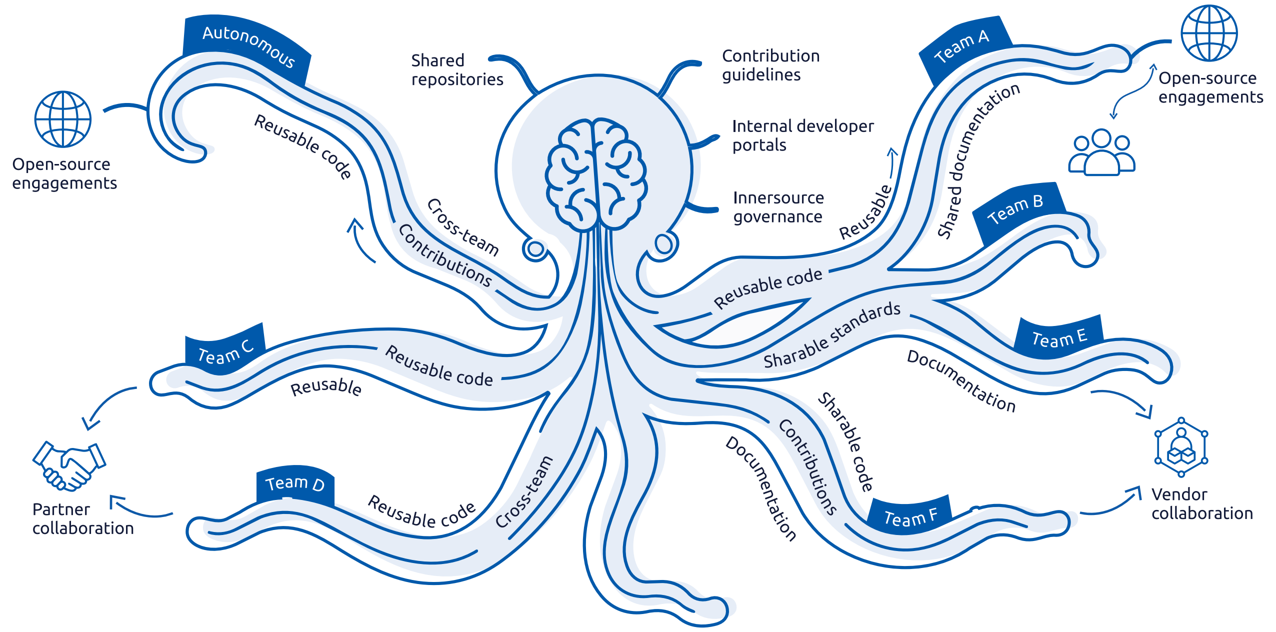

Figure: The InnerSource ecosystem as the nervous system of the modern organization.

> *How an organization builds software determines whether it compounds knowledge over time, or quietly silos itself into irrelevance.*

Why now

Three forces are converging. External technology dependencies have become strategic exposures, and the answer isn't to reject external platforms but to build the internal capability to navigate, adapt, and substitute where necessary. AI accelerates the cost of siloed knowledge faster than it accelerates almost anything else. An organization that cannot find what it has already built is now losing ground monthly, not annually. And engineering talent remains scarce. Engineers who experience the autonomy and visibility of InnerSource rarely return to closed, siloed development. InnerSource is a delivery strategy and a talent strategy at the same time.

Growing the nervous system

There is no shortcut here. The octopus didn't grow eight arms and then bolt on a nervous system; the intelligence emerged through use. Start with the teams who already share code informally and give that practice a backbone: documented contribution guidelines, discoverable repositories, lightweight cross-team review. Celebrate the early contributors by name. Measure reuse, not output. Then expand.

What separates the organizations that succeed from the ones that stagnate isn't the tooling; it's the governance and the cultural intent. Make it genuinely safe and rewarding to contribute across organizational boundaries. Reward reuse over reinvention. Treat ecosystem health as a first-class metric.

The organizations that build this don't merely reduce dependency. They develop the capability and the culture to shape the technology landscape they operate in, rather than adapt to one shaped elsewhere.

The octopus does not wait to be fed. It reaches.

Start *innovating* now

Structure what already exists

Most enterprises already have teams sharing code informally; InnerSource starts by giving that practice a backbone. Document contribution guidelines, make repositories discoverable, set up lightweight cross-team review. Don't mandate the culture; cultivate it. The returns aggregate faster than executives expect.

Measure ecosystem health, not just output

Velocity metrics tell you how fast individual teams move. InnerSource metrics tell you whether the organization is getting smarter as a whole. Track code reuse rates, cross-team contribution counts, and reduction in duplicated solutions. What gets measured gets funded, and what gets funded gets built.

Make the nervous system visible

Deploy an internal developer portal (Backstage is the obvious starting point) to surface shared assets in one discoverable place. When engineers can see what already exists, they build on it instead of around it. Autonomy without coherence is just noise. The portal converts distributed intelligence into organizational advantage.

#InnerSource #OpenSource #DataPowered #DeveloperExperience #TechSovereignty #EnterpriseInnovation

Two crew *required*

Customer experience has outgrown the single-agent model. Here's the flight deck that replaces it!

Kumar Chinnakali

Data & AI Program Manager
Financial Services
Capgemini

The contact center was built for a simpler world: one customer, one problem, one agent, one resolution. That world is gone. Today's customers arrive mid-journey, carrying context from three channels and half a prior resolution, expecting the agent to catch up instantly. The answer is not more scripts or faster tools bolted onto an old frame. It is a new operating model, where human agents and AI agents work as a coordinated crew: each reading the same instruments, each contributing what the other cannot. The organizations that build this now will not just outperform. They will redefine what customer experience is supposed to feel like.

The system that no longer fits

For decades, contact centers operated on a single, sensible assumption: when a customer needs help, a human agent answers. Everything built around that premise — routing logic, scripts, dashboards, performance metrics — was optimization layered on top. Handle more calls. Resolve them faster. Keep transfers low.

That premise has not aged well.

Today's customers do not arrive with a clean problem on a single channel. They arrive mid-journey, carrying the residue of a failed app interaction, a chatbot that sent them in circles, a callback that never came. They have already invested time. They are not patient. And they expect the agent who answers to know all of it — instantly, without being asked to repeat themselves.

The agent's reality tells a different story. Multiple systems. Fragmented knowledge bases. Policies that shift faster than training cycles can absorb. Compliance obligations requiring precision under pressure. All of it compressed into interactions measured in minutes.

So the system compensates the only way it knows how: it measures speed. Average handle time. Calls per hour. Escalations avoided. The metrics look tidy. The experience often does not.

Here is the uncomfortable question. If you were designing the contact center from scratch today, would you really put one person in the cockpit alone and ask them to fly?

> *AI is no longer a tool you call upon. It has become a participant in the crew.*

The shift you cannot automate away

Aviation had this reckoning decades ago — not because pilots were incompetent, but because the environment had grown too complex, the stakes too high, and the cognitive load too large for any single person to absorb reliably. The answer was the modern flight deck: two human crew members working alongside a suite of AI systems that monitor conditions, alert to conflicts, and manage routine navigation, while the pilots retain command authority and judgment.

The contact center is now in the same moment.

AI entered this environment as a peripheral helper, IVRs for routing, chatbots for self-service, automation for the repetitive tasks agents were glad to hand off. Useful, but contained. Kept comfortably at the edge of real work.

That boundary has dissolved. Modern AI systems do not just respond; they interpret. They assemble a customer's full context across systems before the agent has finished the first sentence. They summarize as the interaction unfolds. They detect sentiment shifts, surface recommendations, and flag compliance exposure before a disclosure goes wrong. They operate continuously, not episodically, not waiting to be summoned, but shaping the interaction from the moment it begins.

AI is no longer a tool you call upon. It has become a participant in the crew.

What coordination looks like in practice

In a genuinely collaborative contact center, work is not handed off in rigid steps between systems and people. It is shared, continuously.

Consider a complex banking interaction. A customer calls about a disputed transaction with cascading implications across accounts, fees, and credit exposure. Before the agent completes the first sentence, the AI layer has already assembled the customer's recent history, flagged potential risk markers, and surfaced relevant policy. As the conversation unfolds, the system tracks sentiment. It suggests compliant phrasing at sensitive moments. It recommends resolution paths based on similar cases, updating those recommendations in real-time as new information emerges.

After the interaction, documentation is no longer an afterthought written from scratch in the final seconds. It is largely pre-constructed by the AI, reviewed and refined by the agent, not authored under pressure.

The interaction feels different. Not because the human has been replaced, but because the human is no longer fighting the system to do their job.

Underneath this sits a layered multi-agent environment: conversational AI handling interpretation, knowledge retrieval systems surfacing the right policy at the right moment, workflow orchestrators managing handoffs, and supervisory agents monitoring compliance and quality. The human agent does not operate all of these. They read the instruments the flight deck provides and make the calls that still require judgment, empathy, and the kind of negotiation that no model can replicate reliably.

The goal is not to remove friction by removing people. It is to remove friction around people.

> *The goal is not to remove friction by removing people. It is to remove friction around people.*

When the metrics must follow

When the operating model changes, the measures of performance must follow.

Traditional metrics still have their place. Efficiency matters. Cost discipline does not disappear. But they no longer capture the full picture of performance in a system where outcomes are co-produced by human and AI agents working in concert.

A different set of signals begins to matter. How effectively does the system reduce cognitive load on the agent, freeing attention for what actually requires judgment? How accurate are real-time recommendations, and how often are they trusted enough to be acted on? How quickly does the system retrieve the right knowledge, not just any knowledge? How much effort does the customer expend before reaching resolution?

These are not cosmetic changes to a dashboard. They redefine what performance looks like. Because when an agent's role shifts from information retrieval to judgment and engagement, the quality of their decisions matters considerably more than the speed of their clicks.

The cost of standing still

Some organizations will approach this cautiously. Pilot a tool here, deploy a feature there, measure incremental gains. It feels responsible. It also risks missing the structural shift entirely.

Inserting AI into existing workflows produces optimization. Redesigning workflows around a shared intelligence layer produces transformation. The difference is architectural — and it shows up slowly, then all at once.

By 2030, the gap will be legible. On one side: operations that still resemble modernized legacy contact centers — faster in places, cheaper in others, still constrained by a frame built for a different world. On the other: coordinated human-AI systems that learn continuously, handle complexity without escalation, and generate customer trust as a byproduct of how they operate.

The difference will surface in places metrics struggle to capture directly: the confidence of agents who feel amplified rather than monitored, the trust of customers who feel understood rather than processed, the resilience of operations that improve with every interaction rather than plateauing at their current ceiling.

There is also a quieter risk. Skilled agents, given the choice between fighting outdated systems and working within environments that extend their capabilities, will not take long to choose. The organizations that wait will not just fall behind on customer metrics. They will find it increasingly difficult to keep the people the system depends on.

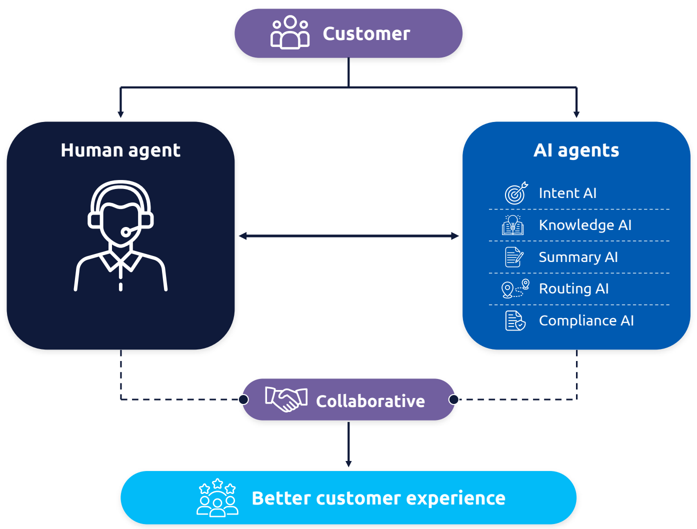

You would not put one pilot in a modern airliner. The environment is too complex, the consequences too significant, the load too high. The contact center has reached the same conclusion. It just has not finished acting on it.

The contact center was built for a world where work could be segmented, scripted, and measured in isolation. That world has moved on.

The question is no longer whether AI belongs in customer experience. It is whether your operating model is designed for intelligence that is, by nature, distributed: across people, systems, and the spaces in between.

When human and AI agents operate as a coordinated crew, reading the same instruments and contributing what only they can, the interaction stops feeling like a process. It starts feeling like understanding. That is the destination. The flight plan is yours to file.

Start *innovating* now

Redesign one journey from the crew up

Pick your most complex, highest-effort customer journey and rebuild the interaction model with human and AI steps integrated from the start. Map where context should be pre-assembled, where real-time recommendations should surface, and where human judgment must lead. Treat it as a system design exercise, not a tool deployment.

Instrument crew performance, not just agent performance

Introduce measures that reveal whether coordination is actually working: recommendation acceptance rate, knowledge retrieval precision, time-to-context at interaction start, and customer effort indicators. These tell you whether the AI layer is genuinely useful or simply present. What gets measured shapes how the system evolves.

Build the shared intelligence layer before you need it

Invest in the data and orchestration foundation that allows human agents, conversational AI, knowledge systems, and workflow tools to operate on the same real-time context. Without this, every AI deployment remains local and isolated. With it, the system begins to learn from every interaction, every resolution, every edge case the crew navigated together.

#DataPowered #AgenticAI #CustomerExperience #FutureOfWork #AITransformation #ContactCenter

Editorial **team**

Arne Rossmann

Innovation Lead,
Insights & Data, Capgemini

Robert Engels

I&D Head of Innovation, VP
Insights & Data, Capgemini

Sudeshna Sanyal

Innovation Marketing Lead,
Insights & Data, Capgemini

About Capgemini

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organizations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of over 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2025 global revenues of €22.5 billion.

Make it real | www.capgemini.com