

From PoC to Industrialisation: Leading AI Change Management in European SMEs · MATIA Method™ · June 2026 Edition

Paul-Antoine TUAL

June 2026

*The question for 2026 is no longer “should we do AI?” but “**how do we move from PoC to industrialisation?**”. **Only 46% of AI PoCs reach production** (IDC-Lenovo, CIO Playbook 2026): a wall that is 70% human, not technological. This white paper sets out the method: the seven frictions of the “last mile”, the five levers of change management, AgentOps, and a 12-month roadmap. It follows on from **AI Maturity of French SMEs 2025-2026**.*

Sourcing rule. Every claim is linked to a dated source. At most one third of the links point to internal content; at least two thirds point to external reference sources (Eurostat, BCG, Deloitte, KPMG, IDC-Lenovo, HBR, Stanford HAI, ADP, the European Commission, and others). **June 2026 edition: all AI-market sources have been refreshed to under six months;** where no fresher data exists (Eurostat, regulations, foundational academic evidence), the wording remains dated and attributed.

1. June 2026: the year the question changed

For three years, the question SME leaders asked was: “Should we do AI?” That question is dead. Across Europe, **Eurostat** records **20.0% of enterprises using AI in the EU27** (2025 data, published 11 December 2025), up from 13.5% a year earlier. Everywhere, in every boardroom, the same question:

“We have a dozen PoCs. Some of them work. But nothing has scaled. How do we industrialise?”

This is **the** question of June 2026. According to IDC-Lenovo’s *CIO Playbook 2026* (January 2026), **only 46% of AI PoCs reach production**: half of them stay stuck at the pilot stage. **HCLTech** (May 2026, 467 executives at firms with over \$1 billion in revenue) estimates that **43% of large AI initiatives are doomed to fail**, through an execution gap more than a technology one. On the agent side, **Deloitte** counts, as of January 2026 (*State of AI in the Enterprise 2026*, 3,235 leaders surveyed across 24 countries), **23% of companies already using agentic AI at least moderately** and almost three quarters planning to do so within two years. Yet **only 21% have a mature governance model** for those agents.

The temptation is to read this as a technology gap. The BCG “10-20-70” rule, reaffirmed in 2026 in *AI Transformation Is a Workforce Transformation*, disproves it: technology accounts for **10%** of success, infrastructure and data for **20%**, and **people and processes for 70%**. Moving from PoC to industrialisation is, by more than two thirds, a problem of **change management** and, where agents are concerned, of **AgentOps**.

Go deeper: The MATIA Method™ Scale: the five levels of AI maturity.

2. The European picture: 20% adoption, a skills wall

Eurostat puts **enterprise AI adoption at 20.0% across the EU27** (2025 data, published 11 December 2025, latest available; next update December 2026), with a dizzying gap by company size: **55% of large companies against 17% of small ones**. The geographic divide is just as stark: according to **rankings compiled from Eurostat data**, Denmark leads at **42%**, ahead of Finland (37.8%), Sweden (35%) and Belgium (34.5%), while Romania (5.2%), Poland (8.4%) and Bulgaria (8.5%) bring up the rear.

The real message comes down to a single figure: among the companies that considered AI but did not adopt it, **Eurostat** finds that **70.9% cite a lack of skills**, ahead of regulatory uncertainty (52.5%) and data protection (48.8%). The **BCG AI at Work 2026** study (3 June 2026, 11,749 workers across 14 markets) confirms the human tension: **72% of workers say AI has already changed the skills their role requires** and **61% think agents could do at least half their job within three years**.

The technology is mature. What is missing is **acculturation, trust and coordination**.

Go deeper: The end of prompt engineering.

3. The valley of death is not technical: the “last mile”

When a PoC fails to reach production, we look for the cause in the wrong place. That is the thesis of the landmark 9 March 2026 article by Karim Lakhani, Jared Spataro and Jen Stave in *Harvard Business Review*, *The “Last Mile” Problem Slowing AI Transformation*: the obstacle “is rarely model quality or data availability, but the *last mile* where technical capability must meet organisational design.” Companies are **“pilot-rich and transformation-poor.”**

Deloitte finds that **only 30% of organisations redesign their core processes around AI** and that **37% use it superficially**, changing nothing at the core (*State of AI in the Enterprise 2026*). Yet it is precisely that redesign that delivers the strongest bottom-line impact. On the individual-usage side, the **Microsoft Work Trend Index 2026** documents a persistent *Shadow AI*: employees adopt AI en masse in their personal capacity before the company formalises its use.

The diagnosis is global. The **Stanford AI Index 2026** (20 April 2026) captures the gap in a single line: **88% of organisations use AI in at least one function, but fewer than 10% have genuinely scaled it** in even one of them; agent deployment remains “single-digit” across almost every function. On the agent side, **KPMG** (*Global AI Pulse*,

Q1 2026) confirms it: **32% of organisations deploy and scale AI agents, but only 11% are “AI leaders”** capturing clear business value. Adopting has become easy; industrialising remains rare.

To industrialise is to **turn private, fragile, invisible use into collective, measured and governed use.**

Go deeper: The five fatal mistakes that doom AI transformations in SMEs.

4. The seven frictions of the last mile, seen from the field

The [HBR article](#) identifies **seven structural frictions**. Here is their on-the-ground translation for SMEs.

“Last mile” friction (HBR 2026)	On-the-ground reality in SMEs	Resolution lever
Pilot proliferation	A dozen PoCs, zero industrialisation	Tie every initiative to a business priority
Productivity gap	“AI changes nothing in our numbers”	Adoption <i>and</i> impact KPIs before deployment
Process debt	Automating an already-broken process	“Blank-page” redesign
Tribal-knowledge identity	The long-standing expert resists	Reposition them as an “architect”
Agentic governance	Agents with no control framework	AgentOps & a control plane
Architectural complexity	Stacked tools, scattered data	Multi-model orchestration, sovereignty
The efficiency trap	AI framed as cost-cutting	Reframe around value creation

Three frictions, up close

Pilot proliferation. In a board meeting I supported in late 2025, we counted together: eleven active PoCs, carried by seven different people, zero in production. Each had its demo; none had a mandate. The triage took one meeting: every initiative had to attach itself to one of the year’s three business priorities, or stop. Seven stopped. The four that remained finally had oxygen: a sponsor, a budget, an indicator.

Process debt. Automating a broken process means producing disorder faster. The healthy reflex is the “blank-page” redesign: redesign the process *for* AI, then automate. It is the exact mirror of the “instant-ROI mirage” described in [the first white paper](#): the gain does not come from a tool dropped onto the status quo; it comes from the status quo rethought.

Tribal-knowledge identity. The best quoting expert in an industrial SME has twenty years with the company. His knowledge *is* his status. Telling him that “AI will do the quotes” is an identity threat, and the resistance that follows is rational. Repositioning him as **architect of the reference base** (he is the one who validates what the agent is allowed to say) turns the keeper of knowledge into the guarantor of quality. Same person, same knowledge, new role.

The five human resistances, one by one

Beneath these frictions, I address five recurring human resistances:

1. **Fear of replacement.** It is lifted by a clear frame, not by slogans: state explicitly what AI will do, what it will not do, and what happens to the hours freed up. The [BCG AI at Work 2026](#) study (3 June 2026) measures that **72% of workers already see their role's skills transformed by AI**, and **61% think agents could do at least half their work within three years**: the answer is a training plan, not a memo.
2. **Distrust of outputs.** It is treated with **native verification**: an agent that proofreads itself and keeps a log (ch. 7) is an agent you can audit, and therefore trust.
3. **The skills deficit** (**70.9% of European non-adopters**, Eurostat, Dec. 2025). That is the purpose of the three-tier training plan (ch. 9).
4. **The absence of ownership.** A PoC with no business owner dies at the first friction. Business champions (ch. 8, lever 2) create that ownership.
5. **Regulatory uncertainty** (52.5% of non-adopters, same source). It dissolves on contact with the facts: the current European framework is an accessible floor (ch. 6), not a wall.

Go deeper: Junyr Agents™: delegating to AI without losing control.

5. The proof: what randomised controlled trials tell us

The argument that “change management creates value” is now backed by the highest level of evidence: the **randomised controlled trial**. The study by Tinelli, Ashley-Timms, Ashley-Timms and Phillips (2023) in the *Journal of Work-Applied Management*, *Impacts of adopting a new management practice: Operational Coaching™*, randomly assigned managers in English SMEs (40 trained SMEs versus 22 controls), under the ethical oversight of the *London School of Economics* and backed by the public *Business Basics* fund. The statistically significant result ($p < 0.01$): training leaders to adopt a **questioning, facilitative** stance **doubles** the time they spend coaching their teams (+13.8 percentage points), with positive trends in **productivity**.

In other words: the effect of a change in managerial posture can now be measured to the standard of medical research. The SMEs' long-standing productivity gap is not closed by buying a tool, but by changing the way you lead. Large firms (250+ staff) produce on average **close to twice the output per hour worked** of small firms, according to the [OECD Compendium 2025](#).

6. The European framework as a lever, not a brake

The [AI Act](#) still applies from **2 August 2026** for transparency. But the political agreement on the **Digital Omnibus** has pushed back the heavier deadlines: according to [Gibson Dunn](#), **high-risk Annex III is deferred to 2 December 2027**, and Annex I **to 2 August 2028**. Most SME use cases do **not** fall under high risk. Article 4 on [AI literacy](#) (in force since 2 February 2025) has been softened: to **support** rather than **guarantee** skills.

Compliance is an **accessible floor**, and the obligation to build AI literacy **gives leaders a legitimate mandate** to invest in training. The question “*where does our data go?*” turns the hosting choice into a **lever for adoption**: that is the sovereignty angle detailed in **the first white paper**.

And that sovereignty is now *practicable*. In late 2025 a French lab, **Mistral**, released a frontier-class model (Mistral Large 3) under an **Apache 2.0** licence: you can run it on your own infrastructure, with no data ever leaving it. And Mistral is building a **sovereign data centre in France** (target of 1 GW by 2030, an \$830M debt raise led by Bpifrance). The balance of power remains skewed (US hyperscalers control **more than 85% of the European cloud market**), but that is precisely why choosing a European building block is a matter of **risk management**, not activism: for an SME, sovereignty is achieved through data governance and reversibility, not through a slogan.

Go deeper: The “all-cloud” era is over · Cryptography 2026.

7. AgentOps: industrialising agents, not just deploying them

In 2026, we no longer deploy mere assistants but **agents** that perceive, reason and act. The **BCG AI at Work 2026** study (3 June 2026, 11,749 workers across 14 markets) finds that **47% of workers already spend more time directing and supervising AI than doing the work themselves**: the agent has become a synthetic colleague to manage. Yet you do not industrialise a colleague the way you industrialise a piece of software.

A colleague that is improving fast, but not yet reliable on its own. **METR** measures the length of the tasks an agent completes on its own at 50% reliability: since 2023, this horizon has been **doubling every ~130 days (~4.3 months)**, faster than any annual planning cycle. But the same work sets the guardrail: the horizon at **80% reliability (the level production demands) remains far lower and grows much more slowly** (**METR, January 2026**). An agent that succeeds at a two-hour task one time in two does not succeed one time in five. It is exactly this gap (capable *and* not yet reliable without a frame) that AgentOps fills.

The risk is not theoretical. A **Sinch** survey (13 May 2026, 2,527 decision-makers across 10 countries) reveals that **74% of enterprises have already rolled back a live customer-facing AI agent**, and that figure **rises to 81% among those with a mature governance model**, not because they fail more often, but because they *see* the drift that others miss. Gartner, for its part, expects in its *Hype Cycle for Agentic AI* (April 2026) that **more than 40% of agentic AI projects will be cancelled by 2027**, for want of cost and risk control. Governance does not remove the risk: it makes it visible in time.

From DevOps to AgentOps: what actually changes

AgentOps is the discipline that extends DevOps/MLOps/LLMOps to **observe and govern the lifecycle of autonomous agents in production**. **Red Hat** describes it as a framework built on five pillars: **observability, evaluation, governance, security and resilience**. Translated into SME-sized gestures, those pillars come down to five questions:

- **Observability**: can I re-read what the agent did yesterday, step by step?
- **Evaluation**: do I have a reference set of business cases to test every change against?
- **Governance**: who approved what the agent is allowed to do, and where is it written?
- **Security**: which data and tools can the agent reach, and which can it *not*?
- **Resilience**: if the agent drifts or fails, who takes over, and how?

A conventional programme fails noisily; an agent can fail **silently and confidently**. That is the whole difference, and the whole discipline.

The Junyr triptych: Plan → Execute → Verify

A concrete best practice: three LLM calls per step. Where a naïve agent chains its actions in a single call, **Junyr Agents™** breaks each step into **Plan → Execute → Verify/Test**. This triptych, close to the “generator/critic” pattern, makes verification a **native step**: each action is gated by a test before the next, which curbs drift and makes the trajectory auditable. This is AgentOps made flesh, and the direct answer to the distrust of outputs (ch. 4): you trust an agent that proofreads itself.

The cost (three calls instead of one) is a deliberate FinOps trade-off, managed through **model routing**: a fast, economical model to plan and verify, a more capable model to execute. The cost-control machinery (LLM gateway, caps, monthly review) is detailed in [the FinOps chapter of the first white paper](#).

The minimum control plane for an SME

You do not need an industrial group’s observability platform to start. The minimum control plane I install on engagements comes down to five guardrails, aligned with what the [2026 agentic-observability tooling](#) documents:

1. **Session interruption**: a human can stop the agent at any time, and the agent stops itself outside its perimeter.
2. **Human validation** on any irreversible action (external sending, accounting entries, contractual commitments).
3. **Auditable log**: every Plan → Execute → Verify step is traced and re-readable.
4. **Capped budget**: a consumption ceiling per agent per month, with an alert before cut-off.
5. **Weekly review**: thirty minutes to re-read the week’s trajectories, adjust instructions, record refusals.

*Go deeper: **Junyr Agents™**: delegating to AI without losing control.*

8. Change management in five levers (the MATIA Method™)

The move from **Artisan to Orchestre** is the move from PoC to industrialisation. Five levers, in the order I pull them on engagements.

Lever 1: Mandate and narrative

Everything starts with the leader, not as a soloist deciding everything, but as a **“leader-as-teacher”** who carries the meaning: why AI, why now, what it changes for each person, what it will not change. The narrative fits on one page, is repeated at every meeting, and names where the freed-up hours go: *“to the customer, not to a redundancy plan.”* An AI project without a narrative is a technical project; a technical project in an SME is a dead project.

Lever 2: Business champions

One credible champion per process creates **ownership**. The right profile is not the most tech-savvy: it is the operational person **respected by their peers**, the one whose adoption tips the others. Two to four champions are enough for an SME of 50 to 500 staff. Their role: first-line help day to day, collecting objections, and **monthly reporting to the executive committee**. They are how ground truth travels upwards unfiltered.

Lever 3: Three-tier training

The antidote to the **skills wall** (70.9% of non-adopters, Eurostat). Three tiers, three audiences, three goals, detailed in chapter 9.

Lever 4: Minimum viable governance + AgentOps

A two-page usage charter approved by the executive committee, a data-classification policy, and the agentic control plane of chapter 7. “Minimum viable” means: sufficient to be applied, light enough not to be bypassed. Governance that lives in a binder protects no one.

Lever 5: Adoption measurement

What is not measured degrades in silence. The indicators, defined *before* deployment, are the subject of chapter 10.

The golden rule throughout: **no level is skipped**. An Artisan-level SME that sinks €200,000 into an agent platform fails far more often than not, for want of foundations. Level-by-level transition durations and budgets are documented in **the first white paper**.

Go deeper: *The MATIA Method™ Scale.*

9. The three-tier training plan & the augmented leader

Tier 1: Acculturation, dispelling fear

For **all** staff, without exception: what AI does, what it does not do, what the company allows itself and forbids itself. This is the building block required by [Article 4 of the AI Act](#), and the direct answer to the [72% of workers whose role already sees its skills transformed by AI](#) (BCG, *AI at Work* 2026). A well-run half-day is enough to turn diffuse fear into precise questions. SMEs have lived through this shift before: [the computerisation of the 1980s](#) followed exactly the same path: refusal, curiosity, normality.

Tier 2: Practitioners, commanding AI, not talking to it

For the teams who use AI inside their process: learning to *command* (structured templates, roles, task chains, acceptance criteria) rather than “chatting with a bot”. It is the shift from artisanal prompting to **context engineering**, and it is what separates a reproducible practice from a daily lottery.

Tier 3: Architects & the augmented leader

For champions and the executive committee: deciding and governing (choosing use cases, reading an ROI, running AgentOps, holding the compliance line). To this is added the work on the leader: turning the cognitive overload tied to [burnout](#) into the posture of an **“augmented leader”** (AI as a second brain for preparation, never as the decision-maker). Drawing on [France Num](#) and Bpifrance’s [“Osez l’IA” \(Dare AI\) plan](#) (an initial envelope of €200 million, within a €10 billion drive).

Go deeper: [The end of prompt engineering](#).

10. Measuring adoption: KPIs and the agentic control plane

The three families of KPI, in practice

Three families of indicators, defined *before* deployment:

- **Adoption:** weekly active users, usage rate per team, frequency per use case. This is the leading indicator: impact follows adoption with a two-to-three-month lag.
- **Impact:** hours saved, output quality, ROI. External benchmark: the [AI & ROI Barometer for French SMEs 2022-2025](#) (Denis Atlan, over 200 B2B deployments analysed) measures a **median ROI of 159.8% over 24 months** for well-framed projects, with a 73% success rate, while the [Stanford AI Index 2026](#) reminds us that **fewer than 10% of organisations have genuinely scaled AI** in a single function.
- **Control** (FinOps & AgentOps): inference cost per use case, human-validation rate, logged incidents and refusals, compliance score.

When to stop a use case

The discipline of stopping is part of measurement. My engagement rule: a use case that stays **below its adoption or impact threshold for two consecutive monthly reviews** is either redesigned or stopped, and the why is documented. Stopping a use case cleanly costs one meeting; letting it die in silence costs the credibility of the whole programme. It is the operational twin of the deployment discipline described in [the five-phase methodology of the first white paper](#).

Go deeper: [Token budgets and AI APIs: the FinOps guide for SMEs in 2026](#).

11. The deflation of expertise: why value is shifting

Karim Lakhani's work (HBS) puts it plainly: AI **lowers the marginal cost of expertise**. What once made a specialist SME scarce becomes abundant. The phenomenon is accelerating: in twelve months, the best open-weight model has gone from ~22 to **54 on the Artificial Analysis Intelligence Index**, **six points behind the best closed model** (April 2026), and the inference cost at constant performance **falls 5-10× per year** (study revised March 2026). Raw capability has become a commodity you download or rent for a few cents. The **right response** is therefore not to endure it but to **redeploy** towards what AI does not commoditise: judgement, empathy, non-standard problems, storytelling, trust, and towards what no competitor can download: your own data and your own processes.

Go deeper: [Where is AI's value heading? Three signals of commoditisation](#).

12. INDUSTEC, from the inside: the human mechanics of a 182% ROI

INDUSTEC (an industrial SME with 78 staff) posts **a 182% ROI documented at nine months into the engagement** (consolidated over 12 months), **275 hours saved per month**, **+18% revenue** on the commercial scope concerned, and **1.7 FTE redeployed**. What made the difference was not the technology: **a narrative carried by the leader, three business champions** trained before deployment, **adoption indicators tracked every week**, and **light but genuine governance**. The full figures are in [the first white paper](#).

What the three champions actually did, week after week

The ritual took thirty minutes a week: re-reading the quotes produced with the assistant, recording what had been rejected and why, raising two or three concrete irritants. The three champions (a senior salesperson, a sales-administration assistant, a design-office technician) had no technical background. Their strength lay elsewhere: when the most sceptical of the three showed his own quote done in twenty minutes instead of two hours at a team meeting, he converted more people than every project presentation combined. Nine months on, the teams were using the tool every day, with confidence, and the monthly board report fitted on one page (data: INDUSTEC engagement reporting, validated at the monthly executive committee).

What nearly sank the engagement

Two critical moments, rarely told. First, the **product reference base**: incomplete and contradictory, it forced three unplanned weeks of clean-up before anything went to production (the classic “data blindness”, the fifth fatal mistake documented in the first white paper). Second, the **temptation to scale too fast**: at the first results, the leader wanted to open three new use cases simultaneously. Holding the line (one use case industrialised end-to-end before extending) was a governance decision, not a technology one. That is exactly what the monthly committee made it possible to arbitrate.

*Go deeper: **INDUSTEC**: 182% ROI in nine months.*

13. A 12-month roadmap & European funding

Reference budget (SMEs and mid-caps of 50-500 staff): **€30,000 to €80,000** over 12 months.

Months 0-1: Framing and narrative

Maturity diagnostic, selection of three to five use cases tied to business priorities, the leader writes the narrative. **Deliverable**: a framing note approved by the board. **Completion indicator**: every use case has a sponsor, a success indicator and a named business owner.

Months 2-3: Human foundations

Champions appointed and trained, acculturation session for everyone (tier 1), usage charter signed; audit and clean-up of the first use case’s data (observed order of magnitude: **€3,000-8,000**). **Indicator**: at least 80% of staff have attended the acculturation session; the data for use case no. 1 is qualified.

Months 4-7: Industrialising the first use case

“Blank-page” redesign of the process, joining the **30% who redesign their core processes** (Deloitte 2026), then governed production roll-out, tier-2 training for the practitioners concerned. **Indicator**: use case no. 1 holds its adoption and impact thresholds for two consecutive months.

Months 8-10: Extending and orchestrating

Second and third use cases, introduction of agents **under supervision, with AgentOps** (the control plane of ch. 7). **Indicator**: three use cases in production, the agentic log active, no untraced incidents.

Months 11-12: Measuring and consolidating

Twelve-month review, gains consolidated, quarterly AI committee constituted, year-two roadmap. **Indicator**: ROI documented with its methodology, durable governance in place.

Funding the journey

Funding in Europe: the “*Osez l’IA*” **plan** was **reinforced in June 2026** (615 AI ambassadors, an “Accélérez avec l’IA” platform, 75+ documented use cases), a sign that the public window is widening rather than closing. Within that framework, Bpifrance’s **Diag Data IA** has covered, since 1 January 2026, **25% of a fixed-price €10,000 (excl. VAT) diagnostic** (€7,500 remaining payable; mid-caps at full rate); the **IA Booster France 2030** funds support work at **up to 80%** (a dedicated €25M budget, open until 31 December 2026, firms of 10 to 2,000 staff). The **European Digital Innovation Hubs** (Digital Europe programme) complement this, with 50% co-funding and a ceiling that can reach **€200,000 per SME over three years** (GBER framework), often with no direct cost to the company. Rates and ceilings evolve: check the conditions in force at the time of application.

14. Conclusion: to industrialise is to transmit

The wall between experimentation and everyday use is not made of silicon. It is the one where **only 46% of PoCs reach production** (IDC-Lenovo, *CIO Playbook 2026*) and **43% of large AI initiatives are judged doomed to fail** (HCLTech, May 2026). Technology accounts for 10%. The rest is change-management work, now **validated by randomised controlled trials**. To industrialise AI is to **transmit a capability** to an entire organisation, humans and agents alike.

And the frontier is advancing fast: the horizon of the tasks agents complete on their own **doubles every ~4 months** (METR), while raw capability becomes a commodity you download or rent. It is precisely because the technology is becoming commonplace that the **only durable advantage for an SME is organisational**: knowing how to frame, verify and govern. BCG puts a number on it: a **clear AI strategy** brings **+25 percentage points of impact**, against +5 points for better tools alone (*AI at Work 2026*). The wall is not made of silicon; nor is the key.

The next rung is already actionable: agent interoperability. *Transactional* two-sided agent-to-agent business, generalised inter-agent commerce and payments, remains a frontier (first vertical deployments, **2028 horizon**). But its **interoperability layer** already exists today: the **A2A** protocol (v1.0 “production-ready”, stewarded by the *Linux Foundation*, April 2026) and the **Model Context Protocol (MCP)** let you cleanly expose your processes to agents, your own and those of your customers and suppliers. For an SME industrialising in 2026, making its IS *agent-ready* is therefore a **2026-2027 objective** that costs little when the APIs already exist, and the precondition for dealing, tomorrow, with the agents of its ecosystem.

Take action

AI Express Audit & Roadmap: 60 minutes, no commitment. [Book your audit](#) · [The blog](#).

*Paul-Antoine TUAL, AI Transformation Leader · Engagements delivered through the SAS **Croissance et Transitions**.*