RESEARCH REPORT

Bridging the AI Implementation Gap

Findings from primary research with 69 European industrial decision-makers.

Jacek Kieszkowski

Research conducted December 2025 to February 2026 for the Diploma in AI for Business.

Published June 2026.

EXECUTIVE SUMMARY

The market is widely engaged, narrowly successful

Between December 2025 and February 2026, I conducted primary research with 69 European industrial decision-makers across eleven countries. Operations directors, engineers, supply chain leaders, C-suite executives. The objective was operational: how do established industrial B2B organisations operationalise AI adoption, given the data infrastructure constraints they actually face?

The picture the data produced is consistent, pressing, and actionable.

Three-quarters of respondents (75.4%) identify implementation, not strategy, as their primary barrier to AI adoption. Not confusion about what AI can do. Not resistance to change. Not budget. Implementation: the absence of the internal expertise and data infrastructure required to make AI work in the organisation they have.

62.3% are deploying AI in some form. Only 21.7% have reached meaningful scale. A mere 2.9% describe AI as transformative to their operations. This figure converges with BCG's independent finding that just 4% of organisations have built functioning AI value engines. Two datasets, separately collected, arriving at the same conclusion.

Organisational resistance ranked last, cited by just 15.9%. Industrial companies are willing. They lack the means.

These findings expose a structural problem in the AI framework industry. Its primary frameworks were developed against exemplars like e-commerce, platforms, fintech: organisations born digital, operating on unified data infrastructures, with engineering cultures built for continuous experimentation. The prescription is coherent. The applicability to a mid-size manufacturer running fragmented ERPs, with critical knowledge distributed across decades of undocumented practice, is limited.

The AI industry has spent five years focusing on greenfield frameworks whose applicability for brownfield organisations is constrained.

This report documents the data behind that gap, names the three boundary conditions where established AI frameworks systematically misspecify what industrial organisations face, and offers a concrete starting point for executives confronting the implementation barrier directly.

> **Inside this report:** the data, the three boundary conditions, the Strategy-Efficiency Disconnect, what to do Monday, and methodology with references.

THE RESEARCH

Method and sample

The study was designed around three operational hypotheses about industrial AI adoption: that organisations are stuck despite AI availability (the productivity paradox in industrial context), that implementation rather than understanding is the primary barrier, and that generic AI inadequately addresses industrial vertical requirements. A structured survey instrument was developed, piloted, and distributed through professional networks across European industrial markets.

Collection ran for eleven weeks, from 10 December 2025 to 23 February 2026. Sixty-nine respondents completed the full instrument. Geographic distribution covered eleven countries, with concentration in German-speaking markets (Germany, Austria, Switzerland), the Benelux region, and Central European industrial economies.

Quality assurance protocols flagged no critical issues. Distribution variance across responses indicated genuine differentiation rather than 'Yes-saying' bias. The sample is meaningfully representative of the European industrial mid-market while remaining modest in absolute size: a constraint discussed in the methodology notes.

Who responded

BY ROLE

Executive leadership	**54%**
Strategy & digital transf.	**12%**
Operations & manufacturing	**12%**
Engineering & product	**7%**
Supply chain & logistics	**7%**
Other (IT, data)	**8%**

BY INDUSTRY

Construction Materials & Products	**28%**
Industrial & Automation Mfg	**22%**
HVAC & Building Equipment	**16%**
Eng/Constr., distribution, adjacent	**34%**

BY COMPANY SIZE

Under 250	**26%**
250 to 1,000	**12%**
1,000 to 5,000	**38%**
5,000 plus	**25%**

The mid-market concentration is deliberate: this is the segment where AI value capture is most contested and where the implementation gap is most acute.

AI ADOPTION MATURITY

Where are companies now?

EXHIBIT 1

AI adoption maturity, five-stage distribution

Share of respondents at each stage. The dashed line marks the scaling wall.

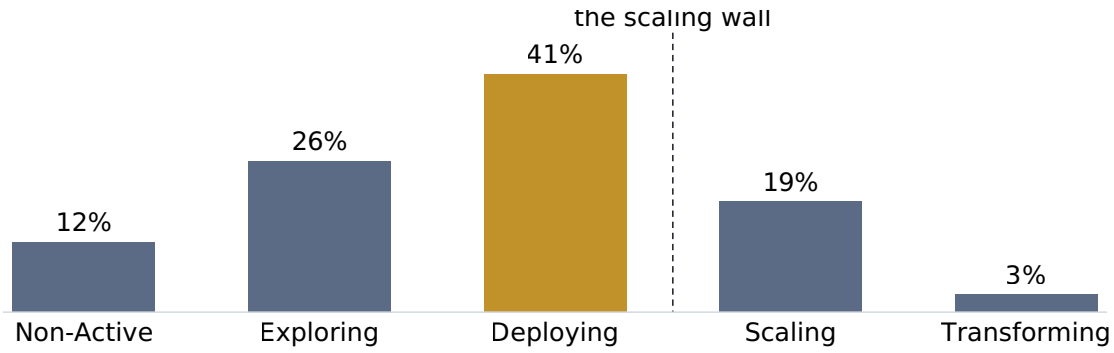

The five-stage maturity distribution reveals the central operational reality of European industrial AI: most organisations are in motion, very few are at scale.

12% are Non-Active. No AI initiatives, no production usage beyond individual ad-hoc tools. **26% are Exploring.** Isolated pilots, proofs-of-concept, no production deployment yet. **41% are Deploying.** One or two AI use cases in production, limited scope. **19% are Scaling.** AI deployed across multiple functions, beginning to standardise. **3% are Transforming.** AI core to strategy and integrated into major workflows.

Two facts the distribution makes plain

62.3% have deployed AI in some form. The exploration phase is largely complete. Industrial organisations have moved from "should we adopt" to "how do we make this work." **Only 21.7% have reached meaningful scale.** Most deployments are stuck at one or two use cases. The scaling wall is the operational reality, not the strategy question. **78.3% remain pre-scale,** including Non-Active, Exploring, and Deploying-in-limited-scope organisations. The market is widely engaged but narrowly successful.

> **The pattern is not an outlier. It is the norm. Just 2.9% describe AI as transformative to their operations. BCG's independent research on AI value engines puts that figure at 4%. Two datasets, separately collected, converging on the same conclusion.**

This is the productivity paradox Brynjolfsson documented in the 1990s, alive and specific in European industrial B2B. The technology is available. The strategic argument is made. The organisational capacity to absorb it remains the constraint.

Where do companies see AI value?

EXHIBIT 2

Where respondents see AI creating the most value

Percentage of respondents selecting each value area. Multiple selections allowed.

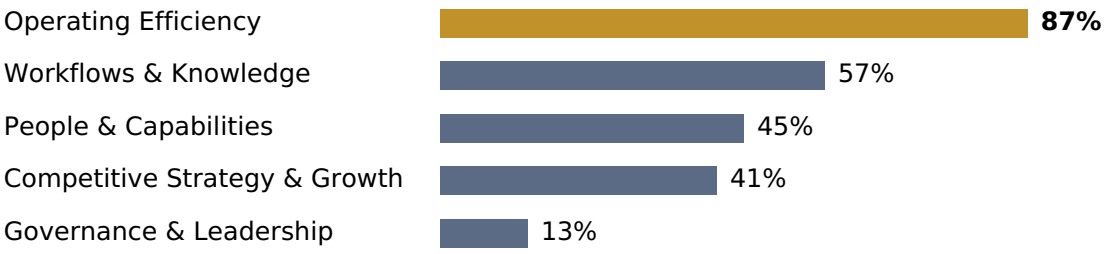

Respondents were asked where they see AI creating the most value in their organisations. The answers split sharply. 87% see AI as an efficiency engine. Only 41% see it as a strategic repositioning tool. Even fewer (13%) see governance as a primary value area.

This matters because the two framings imply different investments, different operating models, different leadership conversations, and different timelines to value capture. Efficiency-led adoption optimises the existing business. Strategy-led adoption rebuilds it. Most European industrial organisations are doing the first, while a growing minority of new entrants are doing the second.

> **The Strategy-Efficiency Disconnect is the single most consequential finding in the dataset. Industrial organisations are not under-investing in AI. They are over-investing in the wrong category of AI: efficiency optimization at the expense of strategic AI repositioning.**

The implication for executive teams is direct. If your AI investment portfolio is 80%+ weighted toward operational efficiency use cases, you are responding to AI the way most of the market responds. That is not, by itself, a strategy. It is the default. The companies that will compound an advantage over the next five years will be the ones treating AI as a strategic repositioning instrument, not an efficiency tool.

PRIMARY BARRIERS

What stands in the way?

EXHIBIT 3

Primary barriers to AI adoption

Percentage of respondents citing each barrier, with respondent counts (n=69).

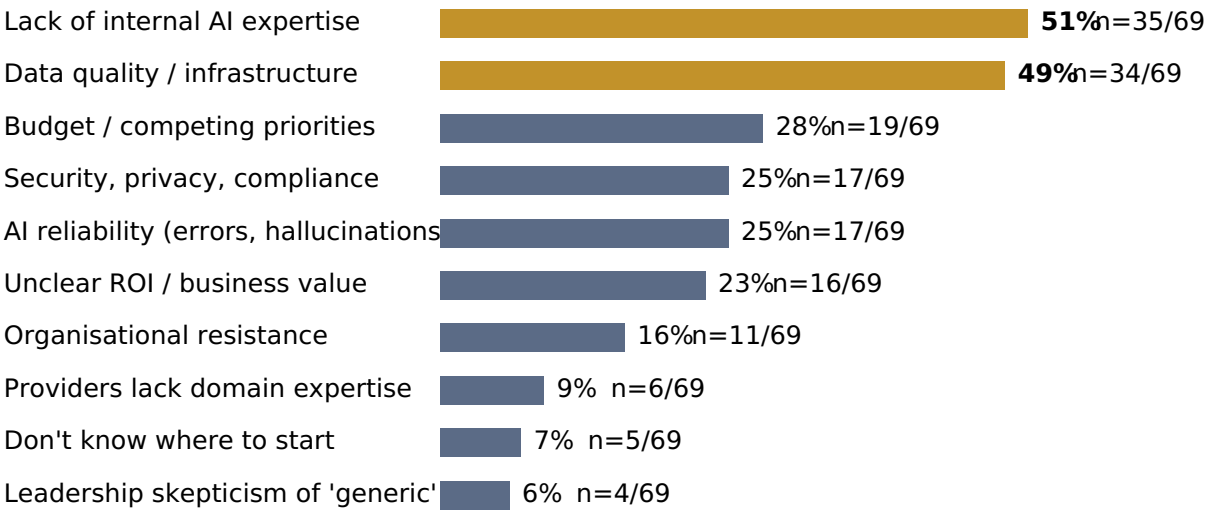

Implementation is the headline

The top two barriers, expertise gaps and data quality, are both implementation problems. They are not strategic problems. They are not understanding problems. They are operational inheritance: the accumulated result of decades of systems added, sometimes not unified; documentation created, not fully maintained; expertise developed, never fully codified.

When these are aggregated with "don't know where to start" into a single Implementation category, **75.4% of respondents cite implementation as the primary barrier to AI adoption.** The next category, Trust and Quality (security plus reliability concerns), reaches 39.1%. The ratio is nearly 2:1.

Organisational resistance ranked last

This finding contradicts AI change management theories which predict resistance is a dominant factor. The N=69 dataset says the opposite. Just 16% of respondents cite organisational resistance as a primary barrier. Industrial leaders are not refusing to adopt AI. They are unable to absorb it without the implementation infrastructure they don't yet have.

> **Industrial companies are willing. They lack the means.**

The size-maturity gradient

Cross-tabulation by company size reveals a clear pattern. No firm with 1,000 or more employees is Non-Active. Scaling rises sharply with size: roughly 6% of firms under 250, 20% of mid-large (1K to 5K), 35% of large (5K plus). Mid-market organisations have started but cannot scale, which validates the implementation-partner positioning rather than the strategy-consultant positioning that the market currently oversupplies.

DEMAND

What solutions do companies want?

SOLUTION TYPE	% OF RESPONDENTS
Automated workflows (multi-step processes, minimal human intervention)	**64%**
Data analysis and insights (patterns, recommendations)	**51%**
Process optimization (identifying and eliminating inefficiencies)	**46%**
Technical knowledge assistants (products, specs, standards)	**41%**
Content generation (documents, reports, proposals, analysis)	**39%**
Predictive analytics and forecasting	**38%**

Zero respondents selected "Don't know enough about AI types to choose." The market has crossed the awareness threshold. Demand is articulated. The question is no longer what AI is, but how to make it work in the organisation that exists.

The solution mix also tells a story about retrieval. Three of the top four solution categories (automated workflows, data analysis, technical knowledge assistants) are fundamentally retrieval problems before they are prediction problems. The systems organisations want most are the ones that find, surface, and route what the organisation already knows.

THREE BOUNDARY CONDITIONS

Where standard frameworks misspecify

The research identified three specific conditions where established AI frameworks systematically misspecify what industrial organisations face. Each represents a place where the standard prescription fails not because it is wrong in general, but because the boundary case has different mechanics.

1 Retrieval, Not Prediction

The most influential framework for AI's economic impact (Agrawal, Gans, and Goldfarb's *Prediction Machines*) argues that AI creates value by making prediction cheap, shifting the economic weight of decisions from prediction to judgment. The framework is largely correct in consumer and digital contexts. In industrial specification contexts, it misses the important bottleneck.

A commercial engineer specifying a complex multi-component industrial system does not face a prediction problem. Given the right inputs (component specifications, performance parameters, integration requirements, supplier data), the calculation is deterministic. The outcome is not uncertain. It is calculable.

The problem is retrieval. Where is the specification document for that device? Is it the 2019 version or the 2023 revision? Does the product catalogue entry reflect current pricing or the one superseded two quarters ago? Which of the three contradictory installation guides is authoritative for this configuration?

The bottleneck is not predicting an unknown future state. It is finding known present facts in systems that have accumulated years of documents, versions, and exceptions that no individual can hold in memory simultaneously.

> **Industrial AI's job is to find what the organisation already knows. Most AI deployments are solving the wrong problem.**

This explains why Retrieval-Augmented Generation (RAG) addresses the actual problem in industrial workflows where predictive machine learning would not. It explains why accuracy on tabular and document data is the design constraint that matters most. And it explains why the verification layer is not a conservative add-on. It is the architecture.

2 The Trust Escalator

Most AI adoption barrier research treats trust concerns as a pre-deployment anxiety that reduces as experience accumulates. The data inverts this.

EXHIBIT 4

Trust and reliability concerns, by deployment status

Share citing trust and reliability as a primary concern.

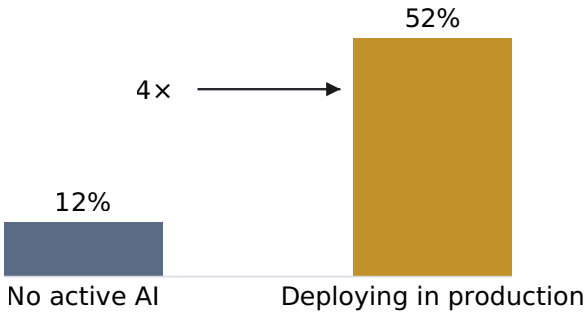

Trust and reliability concerns among organisations with no active AI: 12%. Among organisations actively deploying AI in production: 52%. A fourfold increase.

The logic is industrial. Organisations that have not deployed AI worry about abstract risks. Organisations that are deploying encounter concrete ones: a confident hallucination in a technical specification, a retrieval that looks right but pulls from the wrong product version, an answer that is fluent, cited, and wrong. These are not edge cases in industrial contexts. They are operational realities.

> **In safety-critical environments, trust escalates with exposure. Exposure reveals the failure modes that theory describes but practice makes tangible.**

The design implication is significant. Building for the trust level organisations have *before* deployment produces systems that fail at the moment they matter most: when AI starts performing well enough to be consequential, but the verification infrastructure is not robust enough to catch the errors that still occur. The design question is not how to build AI that works. It is how to build AI governance that remains robust as performance improves and human vigilance naturally weakens.

3 Non-Linear Maturity

Most AI maturity models assume organisations progress through capability stages sequentially and organisation-wide. The data shows this is not how it happens.

A single industrial client typically operates knowledge retrieval at one maturity level and collaborative drafting at another. This reflects different data quality conditions, different risk tolerance, and different organisational mandates across functions. It does not reflect different moments in time. Implementation architectures designed around sequential, organisation-wide maturity assumptions will fail at the function level, which is where implementation actually happens. The same organisation can be Transforming in commercial-engineering knowledge retrieval and Non-Active in customer-service automation. A maturity score is an average that obscures the actual operating reality.

> **Implementation happens at the function. Maturity models that average across functions describe an organisation state that does not exist.**

The Knowledge Access Paradox (developed in the underlying research) is one expression of this. Skills gaps, data quality issues, and integration constraints are partly the same problem: an absence of structured access to what the organisation knows. RAG-based deployment as the entry point simultaneously reduces all three perceived barriers. Waiting for clean data before deploying AI delays indefinitely. Starting AI to surface the structure already present accelerates the data work itself.

WHY THIS MATTERS

The Strategy-Efficiency Disconnect

The AI consulting industry has a structural problem. Its primary frameworks (prediction economics, AI maturity models, AI factory architectures, value engine taxonomies) were developed against highly digitized or digital-native exemplars. Organisations born digital, operating on unified data infrastructures, with engineering cultures built for continuous experimentation. The prescription is internally coherent. The applicability to a mid-size manufacturer running fragmented systems, with critical knowledge distributed across years of not fully documented practice, is limited.

In construction and industrial terminology, brownfield means building on existing reality. Greenfield means starting clean. **The AI industry has spent several years writing greenfield frameworks for brownfield organisations.** The execution gap is the distance between those two worlds.

If 75.4% of industrial decision-makers cite implementation as their primary barrier, the AI transformation industry is structurally mispositioned. It oversupplies strategy frameworks and undersupplies execution infrastructure. Yet we observe, since a few months, the first solutions to bridge this gap, e.g. Claude for Small Business.

What this implies for executive teams

One: adjust benchmarking and ambition levels to industry realities. The frameworks built around digital-natives or global leaders were not designed for the brownfield reality of the European Mittelstand. Benchmarking against them produces strategic recommendations that the organisation cannot operationalise. Benchmark instead against organisations facing comparable data and knowledge infrastructure constraints. That is where the relevant implementation lessons live.

Two: invest in implementation infrastructure before strategy frameworks. The right sequence for brownfield industrials is minimum viable implementation first, strategy refinement second. Strategy emerges from observed reality. Reality is observed by deploying. The conventional sequence (strategy first, implementation after) compounds the constraint of limited data.

Three: treat retrieval architecture as the foundational AI investment. Industrial workflows are retrieval problems before they are prediction problems. The verification layer is not optional. RAG plus structured human-in-the-loop oversight is the load-bearing architecture, not a conservative add-on. Skipping it produces systems that pattern-match consultancy decks but fail at industrial deployment.

WHAT TO DO MONDAY

The first move differs by situation

The data divides European industrial organisations into roughly three operational situations. The first move differs by situation.

If you are Non-Active or Exploring (37.7% of the dataset)

Your scaling wall is downstream. The immediate work is moving from individual ad-hoc AI usage to a single function-level deployment.

1. **Pick one function where retrieval is the primary problem.** Commercial engineering, technical sales support, equipment specification, claims handling, supplier compliance. Pick where the cost of "where is that document" is highest. That is where RAG pays for itself fastest.
2. **Map the knowledge sources.** Before any tool selection, inventory where the relevant documents live: shared drives, ERPs, document management systems, individual mailboxes, the heads of senior employees approaching retirement. The inventory itself surfaces what is fragmented and what is consolidated.
3. **Establish the verification protocol before the first deployment.** Decide upfront what human review process governs AI outputs, who owns it, and what triggers escalation. Doing this after deployment is a forced retrofit. Doing it before is architectural.

If you are Deploying (40.6% of the dataset)

Your scaling wall is the immediate constraint. The work is moving from one or two production use cases to repeatable, function-level deployment.

1. **Diagnose where the second use case stalled.** It almost always stalled at one of: data infrastructure that the first use case worked around, a verification process that did not generalise, or organisational ownership ambiguity. Naming the cause unlocks the response.
2. **Resist the temptation to launch a third pilot.** Pilot proliferation is the most common scaling-wall symptom. The discipline is to deepen one of the existing deployments to second-function adjacency before opening new ground.
3. **Make the verification capacity load-bearing.** Hire, train, or contract for it as deliberately as you would for the AI itself. Verification capacity determines whether trust escalation kills the deployment or strengthens it.

If you are Scaling or Transforming (21.7% of the dataset)

You are already past the scaling wall. The work is institutionalising what got you here so the next function doesn't restart from scratch.

1. **Codify the implementation architecture as reusable internal IP.** The patterns that worked in the first scaled function are the most valuable artefact your AI investment has produced. Most organisations leave this implicit. The cost shows up when the second business unit asks for the same capability.
2. **Re-examine the pricing and contracting structure of AI relationships.** Seat-based and usage-based AI contracts will increasingly mismatch the value delivered as AI capability matures. The Safety-Revenue Paradox is real: success in building capable AI creates commercial pressure on the model that rewards it. Outcome-aligned contracts are the next generation.
3. **Plan the strategic repositioning conversation.** You have the implementation infrastructure most of the market lacks. The next question is no longer how to deploy AI but how to use AI to compound a strategic advantage. This is where the Strategy-Efficiency Disconnect closes for the organisations that have earned the right to address it.

> **Benchmark your position.** A free five-minute AI Readiness Benchmark at **brownfieldbridge.com** returns a diagnosis of where your organisation sits against the N=69 European industrial average, a short analysis of the gaps that matter most for your starting point, and a simple roadmap for what to do first. No registration. No email gate.

ABOUT THE RESEARCH

Methodology and references

Methodology limits and caveats

The study is a primary survey of N=69 European industrial decision-makers, conducted between 10 December 2025 and 23 February 2026 (eleven weeks of collection). The sample is meaningfully representative of the European industrial mid-market that was the focus of inquiry, but is modest in absolute terms. Findings should be read as indicative of pattern direction rather than precise market estimates. This public report uses the full N=69 dataset for LinkedIn-series and benchmark-platform consistency.

Geographic distribution covered eleven European countries with concentration in DACH (Germany, Austria, Switzerland), Benelux, France, Italy and Central European industrial economies.

References

Agrawal, A., Gans, J., Goldfarb, A. (2022) *Prediction Machines: The Simple Economics of Artificial Intelligence* (Updated and Expanded). Harvard Business Review Press.

Brynjolfsson, E. (1993) "The productivity paradox of information technology", *Communications of the ACM*, 36(12), pp. 67-77.

De Bellefonds, N., Charanya, T., Franke, M.R., Apotheker, J., Forth, P., Grebe, M., Luther, A., de Laubier, R., Lukic, V., Martin, M., Nopp, C., Sassine, J. (2024) *Where's the Value in AI?* Boston Consulting Group, 24 October.

Cocoual, C., de Kroon, D., Lefort, F., Chatelet, L. (2025) "From pilots to performance: How COOs can scale AI in manufacturing", *McKinsey & Company*, 15 December.

Leonard-Barton, D. (1992) "Core Capabilities and Core Rigidities: A Paradox in Managing New Product Development", *Strategic Management Journal*, 13, pp. 111-125.

Sukharevsky, A., Krivkovich, A., Gast, A., Storozhev, A., Maor, D., Mahadevan, D., Hämäläinen, L., Durth, S. (2025) "The agentic organization: contours of the next paradigm for the AI era", *McKinsey & Company*, 26 September.

ABOUT THE AUTHOR

Jacek Kieszkowski

Longstanding experience in European industrial B2B: HVAC, construction materials, and building services across Eastern and all European markets.

Academic credentials: Oxford Saïd Business School Diploma in AI for Business (2025–26), IMD Digital Excellence Diploma (2024), SGH Warsaw School of Economics Digital Marketing Diploma (2022–23). PhD in trade strategies (2003). E-MBA in Strategic analysis and programs (1996–98).

linkedin.com/in/jacek-kieszkowski · jacek.kies@outlook.com