A WHITEPAPER FOR MANAGING DIRECTORS OF MID-SIZED COMPANIES

The Decision Before the Investment

Why most AI initiatives fail not because of the technology, but because of the decision that precedes it — and how mid-sized companies navigate confidently between hype and inaction.

KAIRON
The strategic AI decision assistant for executive leadership
2026 Edition · approx. 20-minute read

Contents

1 · Management Summary

> Artificial intelligence has arrived in mid-sized companies — value creation, in many cases, has not. The gap rarely opens up in the code and almost always in the decision that precedes every project: Is this initiative worth it for this business, now, at this cost?

This whitepaper is written for the managing directors of small and medium-sized companies who face mounting pressure to "do something with AI" while sensing that blind activism gets expensive. We show why between 70 and 95 percent of AI initiatives deliver no measurable economic return, which five causes lie behind this, and why none of them is a technical problem. Above all, we provide a practical decision framework: five questions — need, fit, effort and risk, time horizon, and total cost versus return — with which any initiative can be tested before the first euro is invested.

The central thesis: the greatest lever for a mid-sized company is not the next pilot project, but a disciplined **go/no-go decision**. Those who decide early and honestly not only save on failed investments — they gain the very speed that is otherwise lost while others sit for months in pilot gridlock. Specifically, this paper covers the adoption-value paradox in numbers, the five roots of failure, five psychological traps, a five-question framework, an honest TCO methodology, the serious handling of uncertain returns, the build-buy-wait question, the limits of AI, the role of public funding, and a fully worked case study — closing with a checklist you can apply immediately.

2 · The Paradox of the Year: More AI, Barely More Value

The adoption curve is climbing steeply. According to Bitkom, by the summer of 2025 roughly one in three companies in Germany (36 percent) was using AI — almost twice as many as a year earlier; by early 2026 the figure had reached 41 percent. Among mid-sized businesses in the narrower sense, meaning firms with 20 to 500 employees, the share ranges between about a third and 40 percent depending on the survey, with a further 25 percent planning to adopt within twelve months. Nearly half of all companies are actively discussing or planning use.

41%
of companies in Germany use AI (Bitkom, 2026) — up from 20% two years earlier

~39%
of organizations worldwide see any EBIT impact from AI at all (McKinsey, 2025)

95%
of organizations see no measurable return from generative AI (MIT Project NANDA)

Yet behind the adoption curve a value gap yawns. McKinsey's global AI survey from the fall of 2025 found that while 88 percent of organizations use AI in at least one function, only about 39 percent see any effect on operating profit (EBIT) at all — and more than 80 percent report no meaningful impact on enterprise-wide earnings. A widely cited MIT study concludes that 95 percent of organizations derive no measurable return from generative AI. Analysts at Gartner expect at least half — in some forecasts up to 60 percent — of generative AI projects to be abandoned after the proof of concept, and S&P Global observes that on average 46 percent of AI prototypes are scrapped before reaching production; only around half of all initiatives reach production at all.

Particularly treacherous is the **pilot trap**. A pilot is quick to launch and looks impressive in the demo — but the leap into reliable, integrated, day-to-day operation is incomparably more expensive and time-consuming. Many businesses get stuck in an endless loop of ever-new pilots that never go live, because their economic viability was never tested up front. The result is the productivity paradox of our time: lots of activity, lots of tools, lots of licenses — and, at year's end, no visible effect in the profit and loss statement.

These numbers are not a rejection of AI. They are a rejection of the notion that adoption equals value creation. For mid-sized companies the message is especially relevant, because here every misspent euro is felt more sharply than in a corporation: there is no innovation budget to quietly absorb a failed initiative. The good news is the flip side of the same statistic — those who separate the wheat from the chaff before investing belong to the minority that actually sees results. The competition here is less fierce than the hype suggests: a disciplined decision-maker is not competing against 95 percent, but for the five percent who get it right.

There is also a structural advantage of mid-sized companies that gets lost in the hype debate: short paths. While in a corporation an AI decision wanders through committees, IT boards and budget rounds, a managing director of a mid-sized firm can decide within days — provided he has a process he trusts. This very advantage evaporates if the decision is either never made or made only after a lengthy external consultation. Decision speed is itself a competitive factor in the AI era: it is not the fastest implementer who wins, but the one who quickly separates the right initiatives from the wrong ones and channels his limited resources consistently into the former.

3 · Five Roots of Failure — and Why None of Them Is "the Technology"

Ask why AI initiatives fail in mid-sized companies, and the first word that comes up is "data quality." That is correct, but incomplete. A synthesis of current studies reveals five recurring roots — and, strikingly, the actual model or algorithm technology appears in none of them as a primary cause. This is a liberating insight, because it means the decisive levers do not lie in the data center, but on the managing director's desk.

3.1 Missing Strategy and Unclear Business Value

Around 68 percent of surveyed SMEs have no worked-out AI roadmap, and only about 19 percent have even named a responsible person or team. Without that frame, initiatives are launched on a hunch — and abandoned just as quickly once the first costs become visible. Gartner consistently lists "unclear business value" among the top reasons for abandonment. A strategy need not be a 40-page document; an ordered list of initiatives, prioritized by return, effort and risk, is enough. What is missing is not the grand vision, but the small discipline of sequence.

3.2 Choosing the Wrong Use Cases

More than half of SMEs — 54 percent in one survey — do not know which use cases are even relevant to their own business. As a result, they either implement arbitrary, highly visible gimmicks or avoid the value-intensive but uncomfortable cases. Choose the wrong use case and you can implement it technically flawlessly and still create no value. The most valuable use cases rarely lie where AI shines brightest, but where a recurring, well-describable bottleneck costs a lot of time or money — the unremarkable high-volume processes.

3.3 Data That Is Not AI-Ready

About 76 percent of SMEs struggle with inadequate data quality and data silos. Gartner attributes a large share of all failed models to missing or poor data and predicts that by 2026 organizations will abandon around 60 percent of AI projects not supported by AI-ready data. Crucially, this is a question that can be answered *before* the project — not only once the bill arrives. Those who honestly check whether the required data is present, accessible and clean enough recognize early whether the real (and often most expensive) preliminary project is not the AI, but the data work.

3.4 Missing Know-How and Missing Resources

In the Bitkom survey, 53 percent each cite a lack of technical know-how and legal uncertainty as the greatest obstacles, and 51 percent a lack of staff capacity; other studies put the know-how gap as high as 64 percent. For a mid-sized company this means, concretely, that no one has the time to survey and soundly assess a market of hundreds of tools. The consequence is a dependence on the sales pitch of whichever vendor is talking — who is, by nature, no neutral advisor.

3.5 No Change Management

Only about 28 percent of SMEs even have a change-management strategy for AI adoption. Even a technically successful system creates no value if the workforce does not use it. Organizations reporting significant financial returns are twice as likely to have rethought their workflows end-to-end *before*

selecting a model. AI, after all, is not a tool you merely buy, but an intervention in how work is done — and that succeeds only with people, not against them.

THE COMMON PATTERN

Four of the five roots — strategy, use-case selection, data readiness, resources — are questions of decision and preparation, not of implementation. They are asked (or skipped) long before the first line of productive code is written. This is precisely where the underestimated lever lies.

4 · The Blind Spot: the Decision Before the Project

Nearly all the common guides to AI for mid-sized companies start at the same place: implementation. They explain how to pilot, scale, prepare data, train models or integrate tools. That is useful — but it does not answer the most expensive question. Because the greatest loss arises not from a poorly implemented project, but from a well-implemented project that should never have been started in the first place.

An average misinvestment in an ill-fitting AI solution quickly reaches the range of €10,000 to €50,000 for a mid-sized company once licenses, integration, internal effort and opportunity costs are added up — to say nothing of the lost time of leadership. The same sum, invested up front in a clean decision, costs a fraction. The asymmetry is enormous: one hour of thorough thinking before the start beats one month of cleanup afterward. This asymmetry is the economic core of the entire paper.

The reason this decision is skipped anyway is not negligence but overload. The market is opaque, vendors promise everything, the numbers are uncertain and the pressure is high. In this situation managing directors either reach for expensive external consulting — an AI workshop costs €3,000 to €7,000, a strategy roadmap starts at around €5,000, an accompanying retainer runs €2,000 to €8,000 per month — or they make no decision at all and follow the loudest vendor. Both routes are expensive. What is missing is a third: a structured, repeatable decision process that leadership keeps in its own hands.

There is also a rarely named cost type: the cost of non-decision. Whoever, out of uncertainty, makes no decision at all pays not zero, but the price of standstill — forgone efficiency, demotivated employees who have long been waiting for tools, and a creeping disadvantage against competitors who consistently execute their few good initiatives. The task, therefore, is not to approve or reject every AI idea, but to place each one promptly and with justification into one of three drawers: implement now, deliberately later, or discard. A clear "no" is just as valuable as a clear "yes" — what is expensive is only the perpetual "maybe."

KEY STATEMENT

The bottleneck for mid-sized companies is not the ability to *build* AI, but the ability to *decide* in advance, with confidence, whether and which AI should be built. Those who professionalize this step reverse the success rate.

5 · Five Psychological Traps That Sabotage Good Decisions

Before we set up the decision framework, it is worth looking at the forces that undermine it. Investment decisions fail not only for lack of numbers, but because of predictable thought patterns. Knowing them allows you to counter them deliberately.

5.1 The FOMO Trap

"Fear of missing out" is the strongest driver of ill-considered AI spending. When competitors, associations and the media celebrate AI daily, waiting feels like failing. Yet the statistics show: the majority of those in a hurry lose money. The right counter is not courage but measure — the question "What does one year of waiting actually cost us?" can usually be answered calmly.

5.2 The Sunk-Cost Effect

Once a pilot has been paid for, it is hard to stop — "we've come this far, let's see it through." This very reflex keeps dead projects alive. A predefined abort threshold ("If there is no measurable effect by month X, then stop") immunizes against this effect.

5.3 The Shiny-Object Syndrome

The technically fascinating crowds out the economically sensible. The spectacular use case wins attention, the boring high-volume bottleneck is overlooked — even though the value lies there. A fixed evaluation grid forces every initiative through the same sober lens.

5.4 Vendor Dependence

Whoever has no evaluation process of their own adopts the seller's. Vendors estimate returns optimistically and costs selectively. That is not a reproach but their role. Leadership needs its own neutral calculation in order to negotiate on equal terms.

5.5 Activism Instead of Prioritization

"Let's just do something with AI now" is not a strategy but its absence. Five half-hearted initiatives create less value than one carried through consistently. Prioritization — that is, deliberately saying no — is the most uncomfortable but most effective act of leadership in this field.

6 · A Decision Framework in Five Questions

A robust go/no-go decision can be condensed into five questions. They are deliberately kept simple, so that they actually get asked in practice — and yet together they cover the five roots of failure.

Question	What it's about	Which root it closes
a) Need	Where in the business does AI help at all? Which concrete problem, at what volume?	Use-case selection
b) Fit	Does this specific application fit the company's process, data and maturity?	Data readiness, strategy
c) Effort & Risk	What do adoption and operation really cost, and what risks (legal, operational) exist?	Resources, legal
d) Time Horizon	How long until a productive solution — and does that fit the business cycle?	Expectation management
e) TCO vs. ROI	Does the total cost match the expected return? Where is the break-even?	Business value

6.1 Need — the Right Problem First

The first question is the most important and the most frequently skipped. It is not "Which AI do we want?" but "Which problem do we want to solve?" A good need statement is quantified: not "our customer service is slow," but "we handle 400 similar inquiries per week, each taking eight minutes on average." Only this concreteness later permits an estimate of return.

6.2 Fit — Does the Solution Suit This Business?

The same application can be ideal for one company and unsuitable for the next. The fit question checks three things: Is the necessary data present in sufficient quality? Can the solution be connected to the existing systems? And is the process stable enough that automation does not constantly chase special cases?

6.3 Effort and Risk — the Unvarnished Work

Here the adoption and operating effort are realistically quantified and the risks named — legal (GDPR, EU AI Act), operational (downtime, faulty decisions by the system) and organizational (acceptance). A risk you can name is manageable; the unnamed one is dangerous. A word on regulation: with the EU AI Act, documentation, transparency and oversight obligations arise depending on an application's risk class. For the vast majority of mid-market use cases — assistance, text work, internal automation — the obligations are manageable, but they must be known. Legal classification therefore belongs not at the end of the project, but as a brief, honest check at the start of the decision. It rarely decides the whether, but often the effort.

6.4 Time Horizon — When Does the Solution Pay Off?

The time to productive use determines whether an initiative fits the business cycle. A solution that runs in three weeks is something quite different from one that requires nine months of data work. Both can be right — but only if you know which in advance.

6.5 TCO vs. ROI — the Synthesis

The fifth question brings everything together: do the fully captured total costs match the expected return, and from when do cumulative benefits exceed cumulative costs? This synthesis is the heart of the decision — the next two chapters are devoted to it.

FROM GUT FEELING TO COMPARABILITY

Once every initiative runs through the same five questions, initiatives become *comparable*. Instead of "Should we do the chatbot?" the question becomes "Which of our five ideas has the best ratio of return, effort and risk?" The discussion shifts from technical enthusiasm to entrepreneurial prioritization — to where leadership is strongest anyway.

7 · Calculating Total Cost Honestly (TCO)

The most common cause of disappointed expectations is a cost calculation that ends at the license fee. The visible price of an AI tool is regularly the smallest item. As a rule of thumb, you should assume 1.3 to 1.5 times the list price once integration, ongoing operation and governance are added. An honest total-cost-of-ownership view comprises at least seven items.

Cost block	What is often overlooked
License / usage	Per-user and consumption-based models scale with success — cheap in the pilot, expensive in operation.
Data	Cleaning, structuring, connecting sources; for 76% of SMEs the actual bottleneck.
Integration	Connecting to existing systems (ERP, CRM, line-of-business software) — rarely trivial, often outsourced.
Operation & maintenance	Models, interfaces and prompts age; annual upkeep is the rule, not the exception.
Training & change	Without enabling the workforce, the benefit stays on paper.
Governance & legal	Documentation, GDPR, EU AI Act obligations — depending on the application, a serious item.
Internal time	The hours of leadership and the department are real, even if they appear on no invoice.

The point is not that AI is expensive — but that the total cost is *plannable* once you list it in full. An initiative whose TCO you cannot write down in seven lines is not decision-ready. Conversely, a seemingly expensive tool loses its terror when a clearly quantified return stands against it. Another stumbling block is the distribution over time: a large share of the costs falls once at the start (data, integration, training), while the benefit only sets in gradually. Whoever looks only at the monthly price systematically underestimates the initial investment.

8 · Estimating Value Seriously — and Handling Uncertainty

While the cost side is fairly tangible with some discipline, the return side is unavoidably uncertain. This is precisely where many decisions fail — not because the estimate is wrong, but because the uncertainty is handled dishonestly. Optimists pretty up efficiency gains, pessimists talk every assumption to death. Both have a posture, not a method.

A serious approach separates three levels of certainty and makes them visible:

Level	Example	Reliability
Fact	The license price is €X per user per month.	High — verifiable, contractual.
Derivation	With 12 users and our volume, that yields €Y per year.	Medium — follows logically from facts and assumptions.
Estimate	Processing time per case drops by about 30 percent.	Lower — plausible, but to be proven.

The gain of this three-way split is trust. Whoever makes an investment decision must know which number in the calculation is a rock and which a grain of sand. Labeling an efficiency assumption as an "estimate with medium certainty" is not an admission of weakness, but the precondition for making the decision traceable — and later verifiable. From the same logic follows the recommendation to work with **ranges rather than point values**: a pessimistic, a realistic and an optimistic scenario. If the break-even lies within an acceptable period even in the pessimistic case, the decision is robust. If it tips already under slightly less favorable assumptions, caution is warranted. This sensitivity check costs ten minutes and prevents the most expensive errors.

THE BREAK-EVEN TEST

The single most useful metric is not ROI in percent, but the point at which cumulative returns exceed cumulative costs. If a tool prevents even *one* misinvestment of €10,000 to €50,000, its acquisition has, as a rule, already paid for itself at the first well-founded "no."

9 · Build, Buy or Wait?

Once it is established that a need is real and valuable, the next fork follows. It has three directions, and the third is the most frequently forgotten. The choice between them is rarely final: a business can start with a purchased standard tool, gather experience, and only later — once the value is proven and the need is differentiated — consider an in-house solution. This staging substantially lowers risk, because it defers the largest investment to the point at which most uncertainties have already been resolved.

Buy

For most standard problems — text building blocks, document processing, simple classification, customer service — ready-made tools exist. They are quickly available and cheap to start with, but they tie you to a vendor and never fit 100 percent. For mid-sized companies, "buy" is usually the right default assumption; anything else must justify itself.

Build

A custom solution is only worthwhile when the use case concerns a genuine point of differentiation or no suitable product exists. In-house development means full control — and full responsibility for operation, maintenance and further development. In mid-sized companies it is the exception, not the rule, because the running costs are regularly underestimated.

Wait

The underestimated option. In a market where tools become markedly better and cheaper every few months, deliberate waiting can be the most economical decision — provided it is *reasoned* and not merely hesitant. "Wait until our data is ready" or "wait until the market leaders emerge" are legitimate, active decisions. What matters is that waiting has a defined end: a trigger at which it is re-evaluated.

10 · When AI Is *Not* the Answer

An honest decision framework must also know clear no-criteria. AI is powerful, but no universal tool. In at least four constellations the sober answer is "not with AI" or "not yet."

- **The problem is really a process problem.** If a workflow is chaotic, undocumented or constantly changing, AI merely automates the chaos. First order the process, then think about automation.
- **The volume is too small.** If adoption only pays off from thousands of cases but a business has only dozens, the effort permanently exceeds the benefit. A simple template or macro is then the better solution.
- **The cost of errors is too high.** Where a wrong result would have grave consequences and no human control is possible, the risk outweighs the benefit. AI as assistance with control is something different from AI as an unsupervised decision.
- **The data foundation is missing and cannot be created economically.** Without sufficient, clean data, even the best use case remains theory.

These no-criteria are not pessimism but part of the method's strength. A framework that never says "no" is not a decision framework but a justification machine.

11 · Funding as a Lever, Not a Lure

Germany's funding landscape lowers the effective cost of an AI investment considerably — and is nonetheless often forgotten in the decision. Programs such as the Bavarian Digitalbonus fund digitalization and AI initiatives typically with 30 to 50 percent of eligible costs; subsidized advisory via the federal BAFA scheme can, depending on the constellation, cover up to 80 percent of consulting costs. Funding therefore noticeably changes the TCO-versus-ROI calculation.

The sequence matters: funding is a lever on an already good decision, not a reason to make a bad one. An initiative that only pays off with a subsidy is, at its core, uneconomical. An initiative that carries itself even without funding, by contrast, is moved by it from "sensible" to "clearly worthwhile." Eligibility therefore belongs as its own item in every decision dossier — but only after the base calculation stands.

12 · Case Study: a Decision in Practice

An anonymized, illustrative example shows the framework in action. A trading company with 60 employees is considering an AI solution to automatically pre-qualify customer inquiries in its service department.

Question	Answer in the example
a) Need	~500 service inquiries/week, ~7 min. of manual sorting each. A clear, quantified volume.
b) Fit	Inquiries arrive structured in the ticketing system; connection via a standard interface is possible. Fit: good.
c) Effort & Risk	Adoption ~4 weeks; risk low, since a human gives the final answer (assistance, not autonomy).
d) Time Horizon	Productive in ~6 weeks — fits the current fiscal year.
e) TCO vs. ROI	See calculation below.

TCO (year 1, rounded): license €4,800 (fact) + integration €6,000 (estimate, medium) + training €2,000 + internal time €3,000 + operation/maintenance €1,200 = roughly **€17,000**. **Return:** at 500 inquiries/week and an estimated time saving of 3 minutes per inquiry (estimate, medium certainty), that yields around 1,300 hours per year; valued conservatively at €25/hour, that is about **€32,500** (derivation).

Scenarios: in the pessimistic case (only 2 min. saved, higher integration costs) the break-even lies at around 9 months; in the realistic case at around 6 months; in the optimistic case at under 4 months. Even the pessimistic case amortizes within the year. With the Digitalbonus (30–50% on eligible portions), the break-even shifts forward again. **Recommendation: Go** — with the proviso to measure the time saving after three months and to adjust if it falls below the pessimistic assumption.

WHAT THE EXAMPLE SHOWS

The decision was made not out of enthusiasm, but from a traceable calculation with labeled certainty, three scenarios and a built-in abort/adjust threshold. This very structure is what distinguishes an investment from a bet.

13 · From Single Question to Routine: the Decision Dossier

A framework only unfolds its value when it becomes a habit. The practical step is to create, for each AI initiative, a compact, always identically structured **decision dossier**: one to two pages that answer the five questions, list the TCO in seven lines, estimate the return in three scenarios with labeled certainty, state the break-even, and close with a clear recommendation — go, no-go or "later."

Such a dossier has three effects. First, it makes decisions *comparable*, because all initiatives pass through the same grid. Second, it makes them *verifiable*, because after six months you can check which assumption held and which did not — and learn from it. Third, it makes them *delegable and*

communicable: a bank, an advisory board or a co-shareholder understands in five minutes why the investment is being made.

The often-overlooked fourth effect is the learning effect over time. A collection of decision dossiers grows more valuable with each initiative: it shows where the business has regularly estimated its returns too optimistically and its integration costs too low — and calibrates future assumptions. Thus, out of individual decisions, an institutional memory emerges that raises accuracy initiative by initiative. A fixed cadence is therefore sensible: each implemented initiative is briefly checked against its original dossier after three and after twelve months. This closed loop of deciding, measuring and adjusting is the real maturity that the 81 percent of businesses without a roadmap do not reach today.

It is precisely for this routine that **KAIRON** is built. KAIRON guides leadership in a structured way through the five questions, keeps costs and returns separated in the three-level logic of fact, derivation and estimate, sets TCO against ROI, and condenses the result into a decision dossier per initiative — as a tool that keeps sovereignty over your own numbers in-house, instead of handing it to an external consultancy per initiative. The aim is not to replace the decision, but to make it faster, more complete and more honest.

14 · Ten Questions Before Any AI Investment

This checklist can be applied immediately to the next pending initiative.

1. Which concrete problem, at what measurable volume, does this initiative solve?
2. Why now — and what happens if we wait twelve months?
3. Is our data ready for this use case, or would data work be the first, larger item?
4. Does the solution fit our existing processes and systems, or does it force a rebuild?
5. What is the full TCO over three years — all seven cost blocks, not just the license?
6. Which return is fact, which derivation, which estimate — and how certain?
7. Where is the break-even in the pessimistic, realistic and optimistic scenario?
8. Which legal and operational risks exist, and are they bearable?
9. Who in the business uses the system daily — and is that person involved?
10. Is the initiative eligible for funding, and does it carry itself even without it?

> **RULE OF THUMB**
>
> If one of these ten questions cannot be answered, the problem is not the project but its decision-readiness. Then the right decision is not "no," but "not yet — first clarify this question."

15 · Conclusion

Mid-sized companies need no further encouragement to "get started with AI" — that pressure has long been felt everywhere. What is missing is the counterforce: the ability to decide, before every investment, with confidence, completeness and honesty, whether an initiative pays off for this business. The statistic of failed projects is not a story about bad technology, but about skipped decisions.

Whoever asks the five questions — need, fit, effort and risk, time horizon, total cost versus return — with discipline, knows the psychological traps, lists the costs in full, estimates the value in levels of certainty, and records the result in a repeatable decision dossier, shifts from the majority without value creation into the minority with results. The greatest lever lies not in the next pilot project — it lies in the hour before it.

For practice this means an unspectacular but effective resolution: before the next AI initiative, do not ask first about the tool, but about the problem, the costs and the return — and write the answers down before the first budget is released. Whoever makes this a habit turns AI from a source of uncertainty into a manageable entrepreneurial decision like any other. The technology will keep changing rapidly; the discipline of deciding well before the investment remains the constant advantage.

About KAIRON

KAIRON is the strategic AI decision assistant for executive leadership. It guides you through the five questions described here, weighs TCO against ROI, and makes AI investment decisions comparable, verifiable and traceable — with full sovereignty over your own numbers. A suitability assessment for your business is available in a personal conversation.

16 · Sources

Bitkom Research / Bitkom e. V.: Artificial Intelligence in Germany, status quo 2025/2026 — bitkom-research.de; bitkom.org

HKA/KARL study on AI use in mid-sized companies (20–500 employees), cited in bebensee.it

McKinsey: Global AI Survey, November 2025 (EBIT impact of AI adoption)

MIT Project NANDA: study on returns from generative AI (95% without measurable return)

Gartner: forecasts on abandonment rates of generative AI projects after proof of concept (2024–2026); "Lack of AI-Ready Data Puts AI Projects at Risk," 2025

S&P Global Market Intelligence: share of AI prototypes scrapped before production

Informatica: CDO Insights 2025 (obstacles to AI success)

Surveys on AI strategy, use-case and change maturity in mid-sized companies, cited in exponentiellerwandel.de, materna.de, plextec.de, maximal.digital

Costs of AI consulting and funding for mid-sized companies: kimi.consulting, synclaro.de, consulting.de; Digitalbonus and BAFA programs

Data as researched: June 2026. Statistics are drawn from secondary sources and industry surveys; percentages vary by methodology and sample. The case study is illustrative and anonymized. This whitepaper does not replace individual advice.