

#AIMATURITY

The AI value gap

Why AI deployment is outpacing value realization – and how organizations can close the gap



Roland
Berger

The AI value gap

Why AI deployment is outpacing value realization –
and how organizations can close the gap

Companies across industries are facing a widening AI value gap: They are investing heavily in the technology, but financial returns are not keeping pace. This dynamic resembles a form of "profitless prosperity" in AI. Almost 90 percent of companies report that returns fall short of investment levels – not because the technology is failing, but because organizations struggle to manage the transformation. AI is filtering down into production faster than companies can break even financially, causing operating costs to rise long before meaningful value materializes.

To understand what lies behind this value gap, Roland Berger's AI Lab surveyed more than 200 top executives in December 2025 across five major industries and five geographies, focusing on the adoption of generative AI applications and agentic AI systems. The results show that companies fall into four distinct groups, and only a small elite – around ten percent – consistently captures significant value. These leaders are not simply spending more or moving faster: They are operating differently, treating AI as an industrial capability to be engineered rather than a series of isolated initiatives. This enables them to align speed-to-production with speed-to-value, generating more than USD 10 million in annual AI impact far more often than the lowest-performing group.

In this study, we diagnose the root causes of today's profitless prosperity and show what organizations can do to close the AI value gap. Performance varies by geography and industry, with Japan leading globally, the DACH region (Germany, Austria and Switzerland) trailing somewhat behind and financial services struggling across markets. But whatever the starting point, the conclusion is clear: Companies cannot buy their way out of this transformation – they have to engineer their way through it.

Contents

P 4	1/ The state of AI adoption Widespread deployment, rising budgets, limited financial returns
P 8	2/ Uneven outcomes Four AI performance archetypes defined by strategy and execution
P 15	3/ Root causes of the value gap How measurement failure and weak operating models erode returns
P 21	4/ The Industrializer code Five protocols that convert AI into durable value
P 24	5/ Moving forward A mandate for the C-suite
P 29	6/ Conclusion
P 30	Methodology

The state of AI adoption

Widespread deployment, rising budgets, limited financial returns



In just a few short years, organizations across industries have moved from exploring the possibilities of AI to embedding it in the core of how they operate. What was once an innovation project at the edge of the enterprise is increasingly becoming a standard operating condition, with companies scaling both generative AI (GenAI) and agentic AI into real business environments. Yet this surge in activity masks a growing divergence between successful and less successful firms: While AI momentum is now widespread, performance and value capture are not.

This study examines both generative AI and agentic AI adoption. Organizations increasingly tackle both as part of a single AI transformation agenda. Yet agentic systems, which autonomously plan, act and adapt, represent the operational frontier where traditional processes meet intelligent automation. We therefore address the broader AI landscape while also providing dedicated deep-dives into the agentic dimension.

To understand how far these developments have progressed, we surveyed 203 top executives across five major industries and five geographies in December 2025.¹ The results show that executive engagement with AI is now nearly universal: 99 percent of firms in our survey report formal leadership involvement in AI initiatives, while 98 percent have already defined a strategic investment philosophy. This commitment extends to the newest frontier as well: All the companies we surveyed have established a strategic approach toward AI agents, and 90 percent have moved beyond planning into action, with agentic AI either in proof of concept or already in

production. The shift from chat-based generation to action-oriented execution is underway, and more than half of organizations have now deployed autonomous agents in a live production environment.

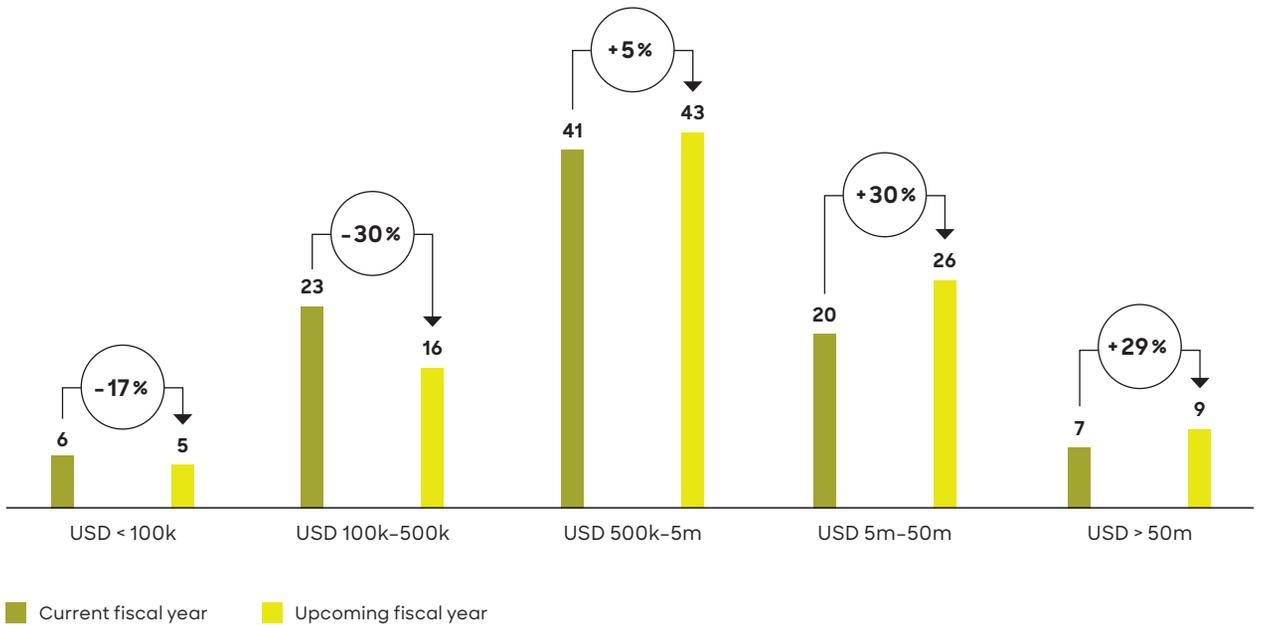
Adoption, however, is only part of the story. This strategic momentum is also being matched by a surge in investment. Budgets are rising quickly, with over a quarter of firms investing more than USD 5 million in GenAI and agentic AI in fiscal 2025. That share increases to 34 percent in companies' plans for fiscal year 2026, reflecting a sharp rise in the number of high-spend organizations. Even the elite tier is growing: The proportion of firms investing more than USD 50 million annually is set to rise from 6.8 to 8.8 percent. AI investment is no longer experimental: It is becoming a material line item in enterprise strategy. ▶ **A**

A common narrative is that the market is engaged in a simple "build or buy" debate. In reality, the dominant implementation model is neither pure in-house development nor full outsourcing, but what we might call "hybrid orchestration". Chosen by almost 40 percent of respondents, off-the-shelf solutions remain the most common approach, closely followed by hybrid models that combine internal and external capabilities. Only six percent of firms build purely in-house, while 18 percent rely primarily on external partners. Crucially, however, "buying" can mean purchasing either finished solutions or capabilities. Firms focused on the former risk slipping into dependency without building internal competence, while the real differentiator in fact lies in buying capabilities that compound over time.

¹ See Methodology for details of the scope of this study and definitions of key terms.

A Spending on AI is surging

AI budget allocation trajectory by investment tier
(2025 vs. 2026, % of all companies)



Source: Roland Berger

FAILURE TO DELIVER

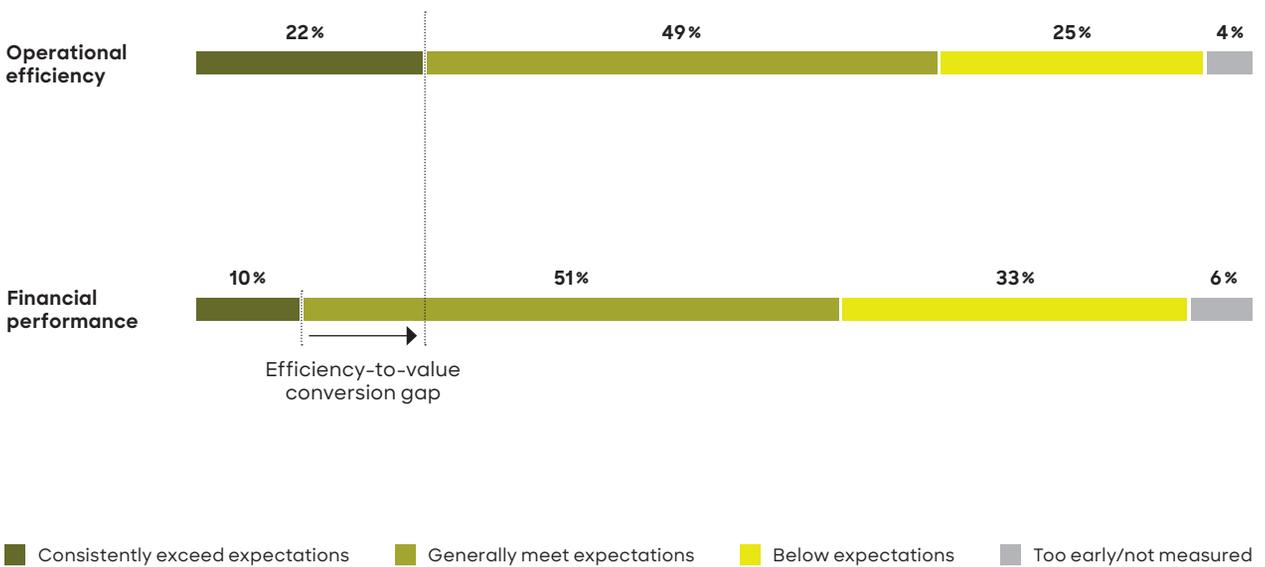
Despite this decisive shift in the market, companies are often disappointed with the results. The survey data reveals what we call a "velocity gap": Companies are bringing AI into production faster than they are capturing value from it. Thus, while 37 percent of organizations reach production within nine months, only 31 percent achieve financial breakeven in the same timeframe. This disconnect means that firms systematically incur production costs while waiting

significantly longer for returns. Over time, this velocity gap compounds into a breakeven trap. Nearly three-quarters of the surveyed firms report that their AI projects face delays or extended timelines, while just 14 percent consistently reach breakeven on time. For most organizations, getting AI live is no longer the hardest part: Making it pay is.

The pattern becomes even clearer when firms compare outcomes against their initial projections. What emerges

B The efficiency-to-value conversion gap

Operational efficiency vs. financial returns (% of all companies)



Source: Roland Berger

is an efficiency-to-value conversion gap: Organizations are far better at achieving operational improvements than converting those gains into financial impact. The dominant, buy-heavy approach can deliver process optimization, but it is often insufficient for sustained value creation. Cutting costs is easier than generating new revenue. ► B

why, we must examine how the market has splintered into distinct realities, with some players trapped in this pattern and others already breaking through. This is the subject of the following chapter.

In short, companies have learned to make AI work, but they have not yet learned how to make it pay off. To understand

Uneven outcomes

Four AI performance archetypes defined by
strategy and execution

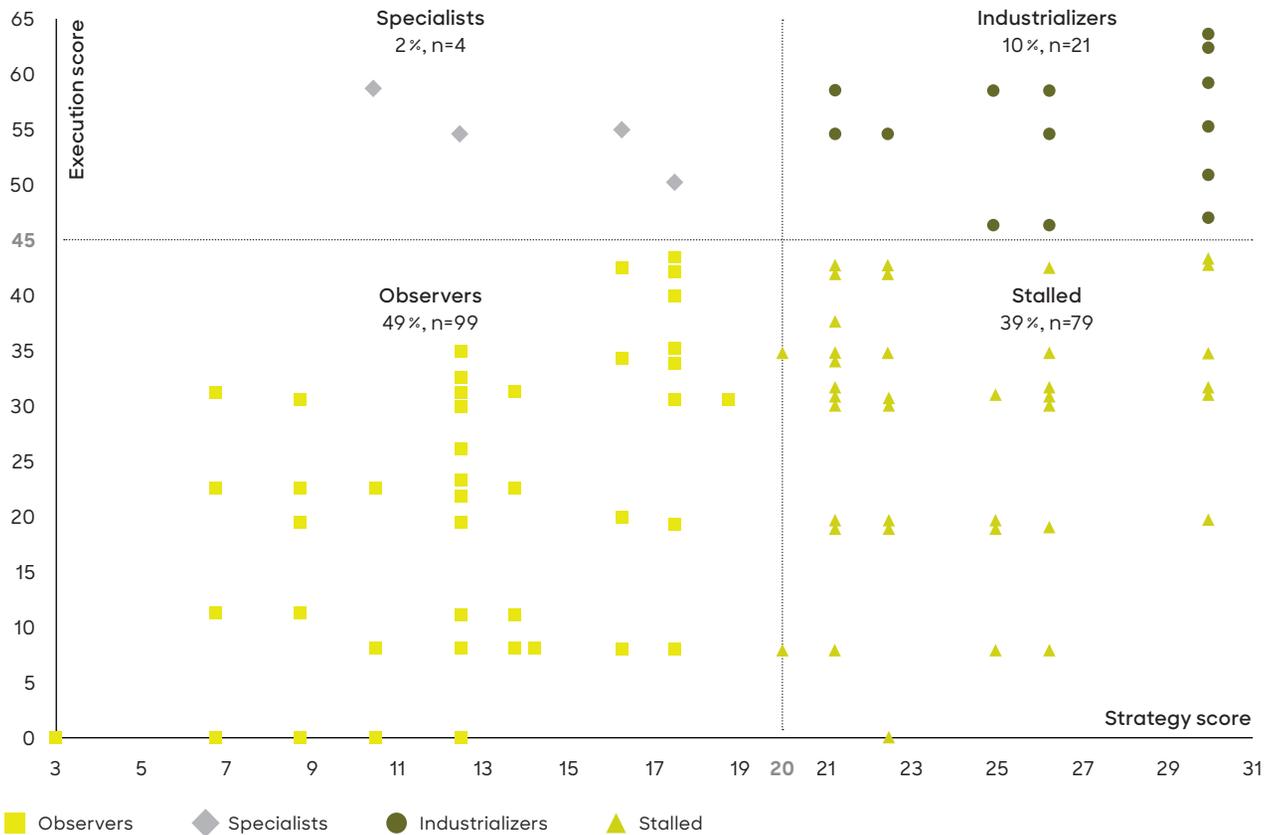
A large, bold, black number '2' is positioned to the left of a large, white number '1'. The '1' is partially cut off by the right edge of the page. The numbers are set against a solid olive green background.

When it comes to AI, high budgets and rapid deployment are no longer reliable predictors of success. Although AI commitment is now widespread, outcomes are increasingly uneven. To make sense of this performance split, we grouped the 203 firms in our survey

into four distinct segments based on their strategic intent and execution capability. Each segment reflects a different relationship between investment and returns – and each, as we will see, requires a different response. ▶ **C**

C High investment does not guarantee high returns

AI performance matrix



Source: Roland Berger

At the top right are the **Industrializers**, representing around ten percent of the sample. They set the performance benchmark for large-scale AI deployment, having successfully synchronized speed-to-production with speed-to-value and created a compounding advantage that justifies their high investment levels. In other words, when they bring AI into production, they monetize it effectively.

The Industrializers contrast with the larger share of firms that we describe as having **Stalled**. This group, which makes up 39 percent of our sample, matches the Industrializers in ambition and budget allocation but sees only laggard-level returns. They are running fast, but they are running on the spot. Understanding where they are going wrong is a key focus of the following chapter.

The largest group are the **Observers**, which account for 49 percent of the market. These firms are piloting and experimenting, but they are not scaling. As we will see later, such caution at the corporate level often masks unsanctioned activity within individual departments.

Finally, a small group of **Specialists**, around two percent of the total, achieves strong results with limited budgets, often through narrow or siloed approaches. They demonstrate that massive spend alone is not a prerequisite for impact. As they included just four companies in our survey and are thus statistically of very minor importance, we have excluded them from the analysis that follows.

Importantly, these uneven outcomes are **not distributed evenly across geographies**. Japan leads globally, with nearly three times the average share of Industrializers, reflecting a disciplined execution culture combined with a willingness to commit resources. The DACH region presents something of a puzzle: Despite its engineering heritage and strong industrial base, it shows the lowest share of Industrializers and the highest share of Observers, perhaps a reflection of

strong compliance requirements or cultural risk aversion. The United Kingdom, meanwhile, has the highest Stalled rate at 52 percent, suggesting systematic barriers to execution rather than a lack of ambition on the part of the firms themselves. ▶ **D**

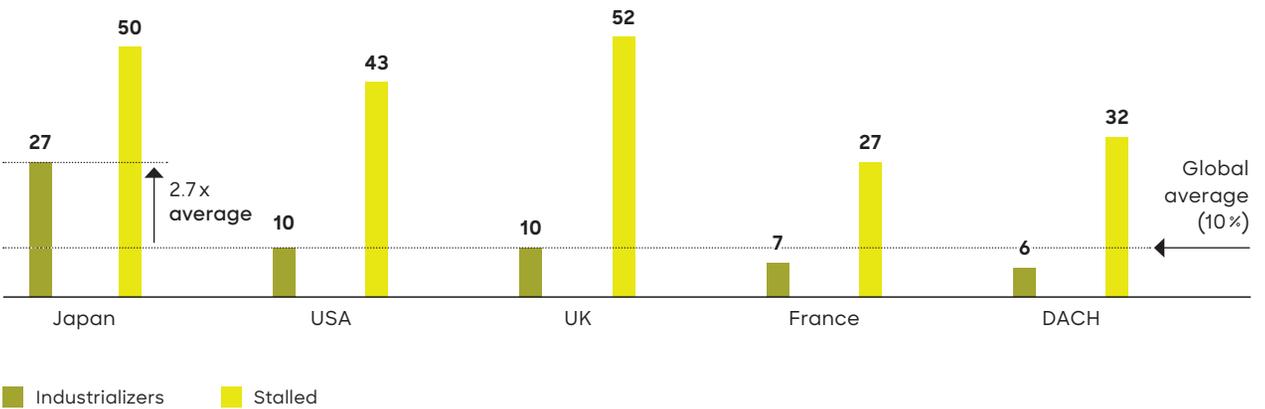
A similar pattern emerges across industries. Financial services players struggle in every geography, with only five percent of firms in the Industrializers group. In both DACH and the United States, we find no Industrializers in financial services at all. Their absence underscores how regulation, legacy infrastructure and risk constraints can exacerbate each other as inhibiting factors. Technology, media and telecommunications deliver a counterintuitive outcome: Despite presumed technical sophistication, 54 percent of firms fall into the Stalled category – the highest share of any industry. Health likewise exhibits a major gap in value realization. By contrast, industrials and automotive firms lead on performance, likely reflecting greater operational discipline and experience with managing complex, regulated processes. ▶ **E**

● ● "We identified 4 AI archetypes: Industrializers turn deployment into profit, Stalled firms waste millions despite big budgets, Observers remain stuck in pilots, Specialists excel in niches."

Maria Mikhaylenko
Senior Partner

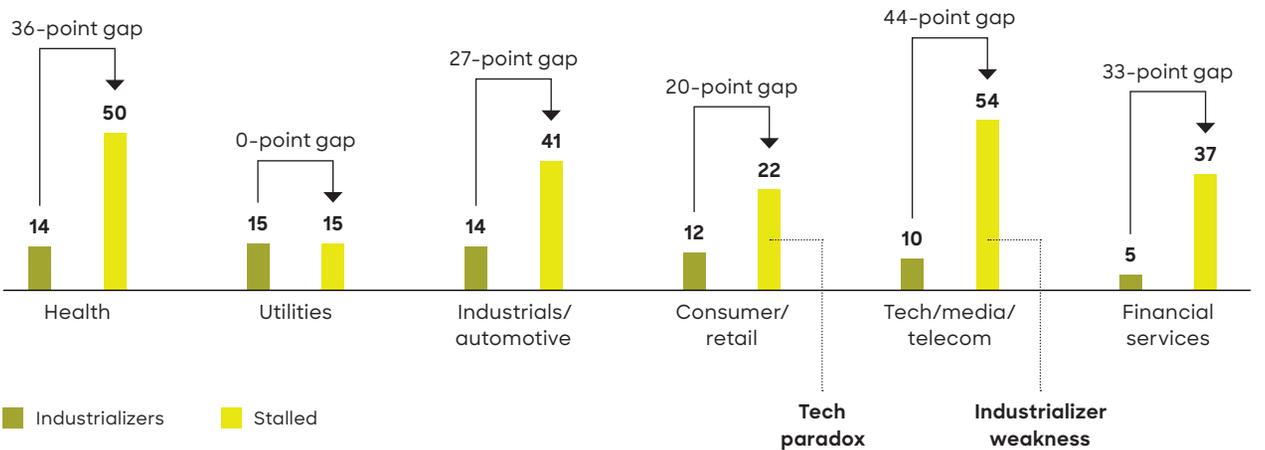
D Geography matters

Share of Industrializers and Stalled companies by country/region



E Surprising industry patterns

Share of Industrializers and Stalled companies by industry



Source: Roland Berger

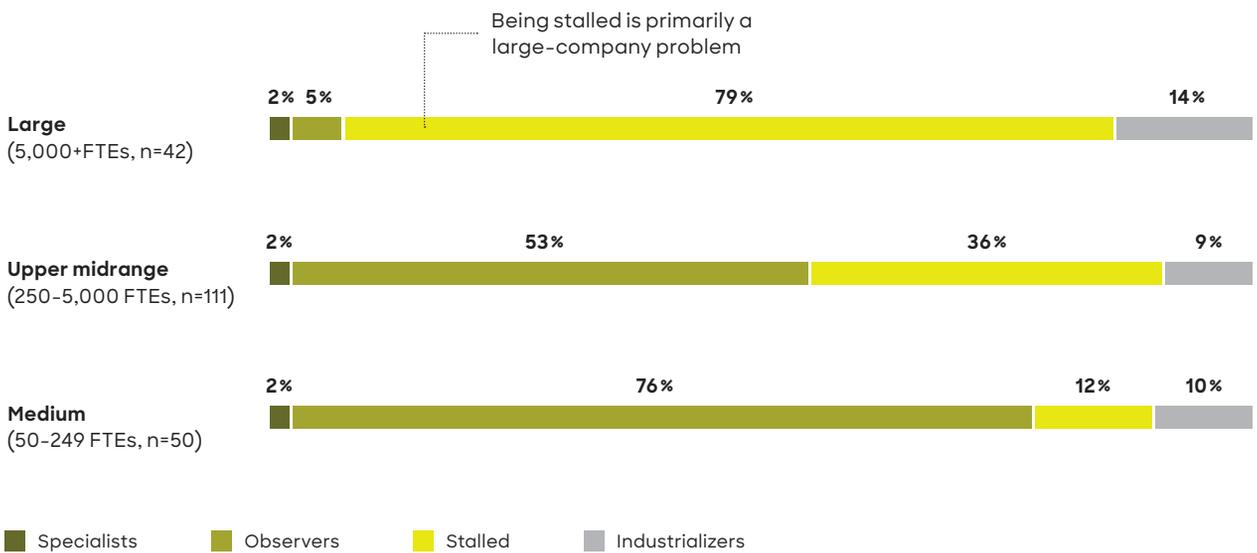
Company size adds another layer to this performance divide. One might assume that larger firms would enjoy structural advantages, from bigger budgets to deeper talent pools. However, the data suggests otherwise. Among organizations with more than 5,000 employees, almost 80 percent are in the Stalled group. They have committed to AI, but struggle to execute at scale. Medium-sized firms, by contrast, are often constrained by resources, with just over three-quarters remaining Observers. The strongest performance appears in the USD 1-2 billion revenue bracket, while the USD 100-500 million band emerges as a dead zone, too large to remain agile but too small to invest

at scale. Being stalled is therefore less a dysfunction than a complexity problem: The firms in the latter revenue bracket have the budget, but the challenge of integration overwhelms them. We examine the operating failures behind these patterns in Chapter 3 and the specific interventions required to overcome them in Chapter 4. ► **F**

This becomes even clearer when we examine the velocity gap at segment level. Somewhat counterintuitively, Stalled firms are not slow in the traditional sense. In fact, they get AI into production just as quickly as Industrializers: 52 percent do so within nine months, identical to the top

F Almost 80% of large companies have stalled

Share of different segments by size (% of all companies)



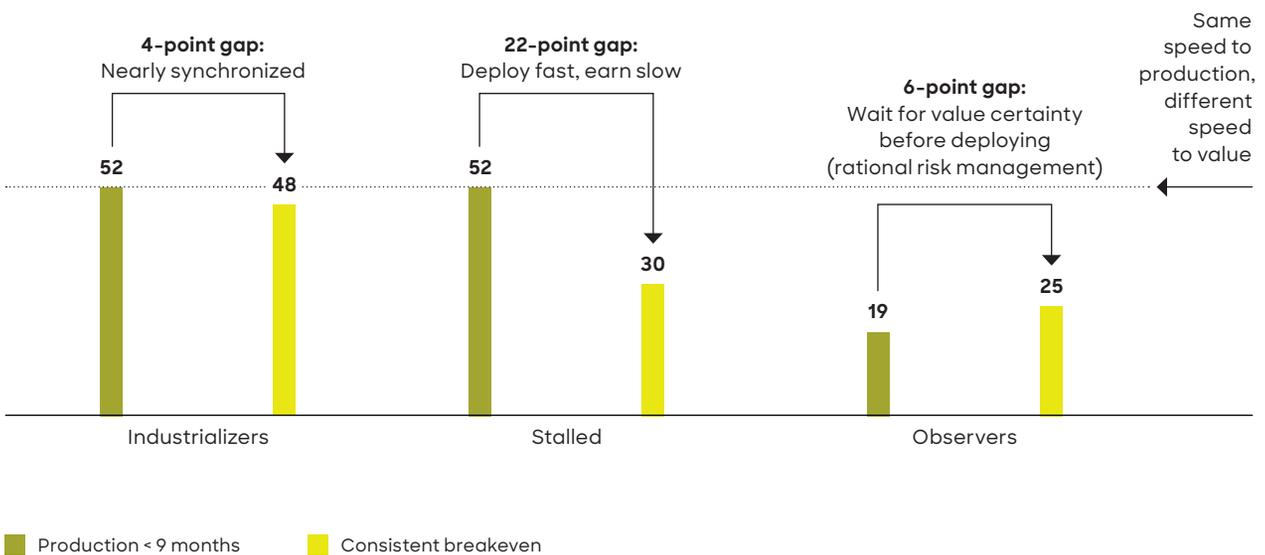
Source: Roland Berger

segment. The difference emerges after go-live: Industrializers nearly synchronize production speed and breakeven speed, while Stalled firms show a massive 22-point gap, carrying the full cost of production deployment for months or even years before seeing returns. Observers, by contrast, often wait for value certainty before scaling, reflecting risk management rather than failure. Among firms that move fastest into production, the paradox intensifies: 68 percent find themselves Stalled. Our conclusion? Speed is necessary, but profoundly insufficient in itself. ▶ G

Beyond speed, the split is also one of **predictability**. Industrializers move to production only after establishing strong evidence of value, be it through pilot validation, clear business case modeling or deliberate risk acceptance within defined guardrails. As a result, 62 percent consistently hit their breakeven targets on time. For Stalled firms, launching is a gamble, not a guarantee: Only nine percent meet breakeven on schedule, while nearly a third report that initiatives take much longer than expected to pay back. In this sense, Industrializers have largely taken surprise out of the equation. They have built systems that make success not just possible, but predictable.

G The velocity gap

Time-to-production vs. time-to-breakeven by segment



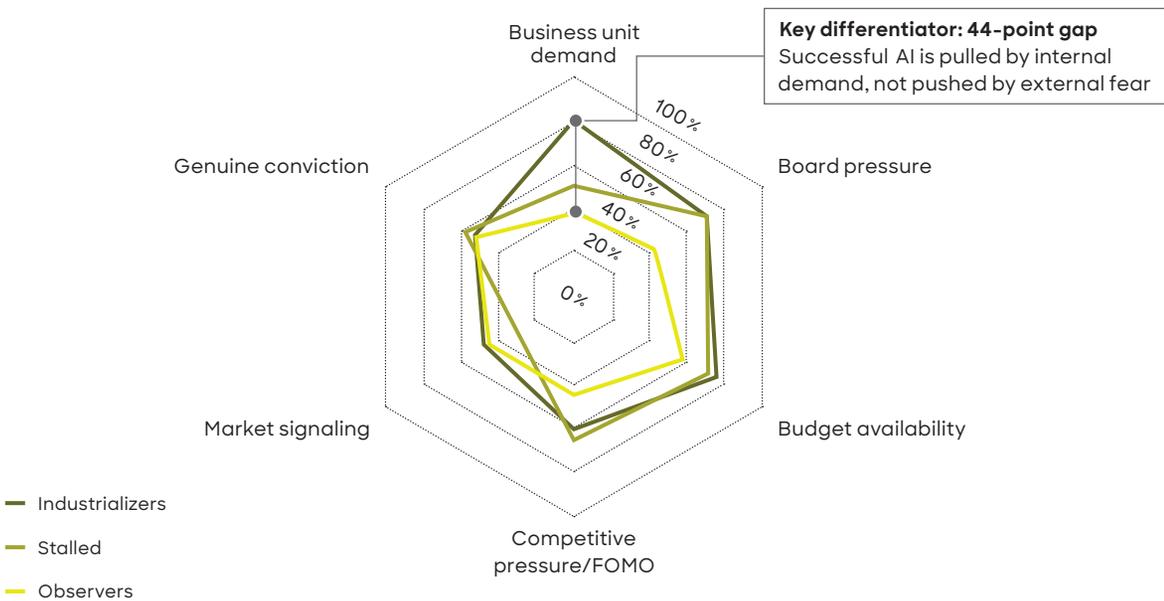
Source: Roland Berger

Investment motivations further reinforce this distinction. While one might expect Stalled firms to invest out of fear and Industrializers to invest out of conviction, the data complicates this narrative. Industrializers experience equal or greater board pressure and competitive urgency, but the true differentiator is internal demand. Business units drive AI investment at 81 percent of Industrializers, compared with 51 percent of Stalled firms and 37 percent of Observers. In other words, successful AI transformation is pulled by internal demand, not pushed by external fear. ► **H**

Lastly, the "**efficiency illusion**" helps explain why so many firms seem to remain trapped. Stalled firms often achieve operational improvements: Just under a quarter exceed efficiency targets and over half (57 percent) meet expectations. Yet they fail to convert these gains into financial outcomes. Only around six percent consistently exceed revenue or value targets, compared with 43 percent of Industrializers. For Stalled firms, AI activity looks productive and dashboards show progress, but they confuse busyness with business. Crucially, their bottom line fails to move.

H Internal demand is the true differentiator

Investment drivers by segment (% citing factor as "decisive" or "very influential")



Source: Roland Berger

Root causes of the value gap

Measurement failure and weak operating
models erode returns

3/

How can companies invest millions in AI and still remain trapped in uneven outcomes? The answer lies in two root causes: Many organizations are still flying blind on measurement and steering, while their operating model and technology architecture remain hollow. Together, these failures explain why speed-to-production so often fails to translate into speed-to-value.

#1: MEASUREMENT FAILURE - "FLYING BLIND"

The velocity trap identified in Chapter 2 is, at its core, a measurement failure that manifests as a steering problem. Stalled companies are moving quickly, but they are watching the wrong dashboard. Across industries, measurement maturity remains low: 63 percent of companies still rely on one-time pre-post analysis or informal gut feelings to judge success, while only 25 percent have implemented continuous ROI (return on investment) modeling through automated dashboards. Within this baseline blindness, however, the gap between segments is wide: Industrializers are nearly four times as likely as Observers to use continuous ROI steering, and they embrace A/B testing at five times the rate. Even so, almost half of Industrializers also partly rely on intuition – the difference being that Industrializers back it up with rigorous data. ▶ |

This lack of discipline is compounded by **KPI sprawl**. Stalled companies do not fail because they measure nothing, but because they try to measure everything at once. Without clear priorities, every metric becomes urgent and teams cannot make tradeoffs. While both Stalled companies and Industrializers typically track four high-priority KPIs, the difference lies in focus and leadership alignment. Just over half of Industrializers report strong alignment on the implementation approach, compared with 30 percent of Stalled companies and just 13 percent of Observers. When leaders disagree on what matters most, teams are forced to optimize for speed, cost, accuracy and autonomy simultaneously, which often results in paralysis.

One tradeoff - **autonomy versus safety** - captures this tension particularly well. To generate value, an AI agent must be granted autonomy to act. Yet organizations also seek to minimize risk. Stalled companies often engineer an impossible contradiction by treating autonomy and accuracy as equal priorities. The KPI data confirms this pattern: Stalled firms assign near-identical weight to autonomous task completion and decision accuracy, while Industrializers prioritize autonomy 24 percentage points more heavily than accuracy. The result - for the Stalled firms - is paralysis.

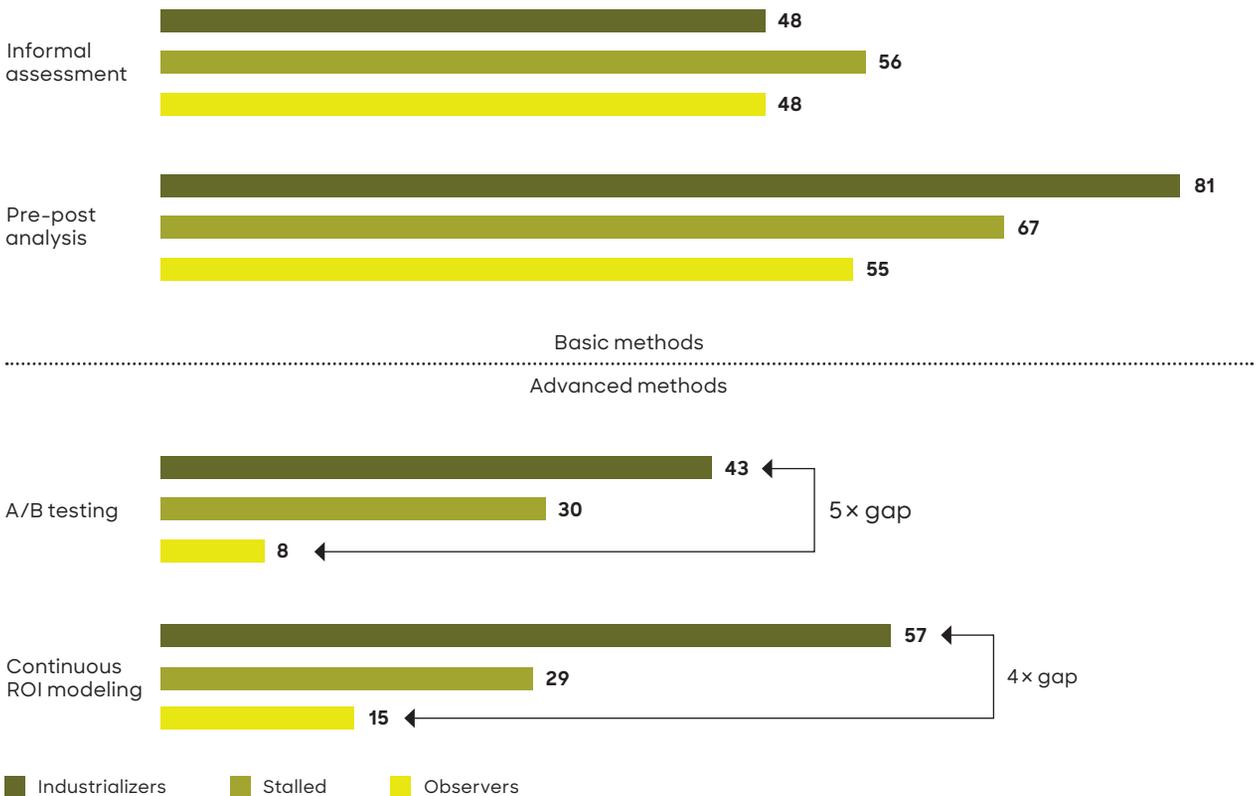
Industrializers take a different approach, allowing autonomy to lead while embedding strong monitoring, governance and feedback systems that safeguard performance. This is not recklessness: As we show in the next section, Industrializers have resolved governance and monitoring at the platform level, freeing them to measure output rather than police errors. Rather than preventing action upfront, they manage risk systematically and improve accuracy over time. Observers opt for safety first, prioritizing accuracy over autonomy, which fits their cautious posture but limits capturable value.

● "The velocity gap opens up
● when firms optimize for deployment milestones instead of value milestones. This creates the illusion of progress while financial returns lag behind."

Patrick Ruthardt
Principal

From gut feelings to real-time dashboards

Measurement approach by segment (% citing approach as "standard practice" or "frequently used")



Source: Roland Berger

When we asked respondents to describe the single most pressing challenge in moving AI from proof of concept to production, their answers confirmed these patterns. Observers say they are blocked by the "why," struggling to justify investment and prove value. Stalled companies are hindered by the "how," citing data fragmentation,

integration complexity and resource allocation. Industrializers, meanwhile, focus on optimization at scale: adoption, trust and consistent quality. Each segment faces challenges appropriate to its maturity, which is why fundamentally different interventions are required.

#2: WEAK OPERATING MODELS - "HOLLOW TECH ARCHITECTURE"

Strategic confusion is physically embodied in the **technology stack**. Many Stalled organizations have in effect bought the Ferrari of modern AI infrastructure – multi-cloud platforms, best-of-breed tooling, premium vendor ecosystems – but are running it with the go-kart engine of shallow integration. They have confused buying modern software with building modern architecture. Beyond cloud infrastructure, agentic AI introduces an additional technology layer: orchestration frameworks, tool-use protocols such as Model Context Protocol (MCP), multi-agent coordination and state management. As with the cloud infrastructure, Stalled firms often acquire these components without integrating them into a coherent orchestration stack. The result? Agents that can reason but cannot reliably act within enterprise systems.

This fragmentation is reinforced by a persistent **talent gap**: Around two-thirds of Stalled companies cite technical skills as a significant barrier, compared with only one-third pointing to business acumen. Development approaches reflect this divide, with 38 percent of Industrializers relying primarily on high-code custom development versus 14 percent of Stalled companies and just four percent of Observers.

At the core lies the question of **integration depth**, or "wrappers versus wiring". Stalled companies are five times more likely to deploy agents in departmental silos than Industrializers, building disconnected tools rather than an interconnected ecosystem. The underlying issue is not purely technical but organizational: Scaling AI agents beyond departments requires end-to-end process transformation across functions and business units. Data integration, by definition, is cross-organizational. Yet most Stalled firms pursue it within departmental boundaries, effectively fragmenting both the data layer and the process logic on which agents depend.

While the surface architecture of Stalled companies may appear modern, Industrializers wire agents directly into the data layer: 91 percent use data-centric approaches such as RAG (retrieval-augmented generation) pipelines that connect to source systems. Stalled companies, by contrast, rely more heavily on platform-centric connectors, wrapping agents around legacy environments rather than rewiring the core. Without end-to-end process transformation, these wrappers preserve fragmentation, such that agents operate within silos rather than across the value chain. This explains how Stalled firms can reach production quickly yet struggle after going live: Wrappers enable speed upfront but cannot carry the weight of enterprise operations. ▶ J

Legacy systems remain a universal constraint – 57 percent of Industrializers see them as a significant challenge – but the difference lies in the response. Seventy-one percent of Industrializers explicitly pursue an AI-adapted brownfield strategy, transforming existing systems rather than building around them, while Stalled companies allow legacy complexity to stand in the way of deep integration.

● "What looks like technical debt
● is often organizational debt:
The cumulative cost of choosing wrappers and connectivity over rewiring and integration leads to systems that demonstrate capability but fail to deliver at scale."

Daniel Rohrhirsch
Senior Partner

Governance further separates the four segments. Industrializers have largely solved security and compliance, holding a 34-point advantage over Observers. More strikingly, only 19 percent of Industrializers see governance as a barrier, compared with 56 percent of Stalled companies. The difference is not the presence of governance, but the way it is designed. While Stalled companies often rely on structured committees and approval chains, Industrializers embed governance in platforms through documented policies, automated enforcement and integrated controls – governance that answers the "how" question. ▶ **K**

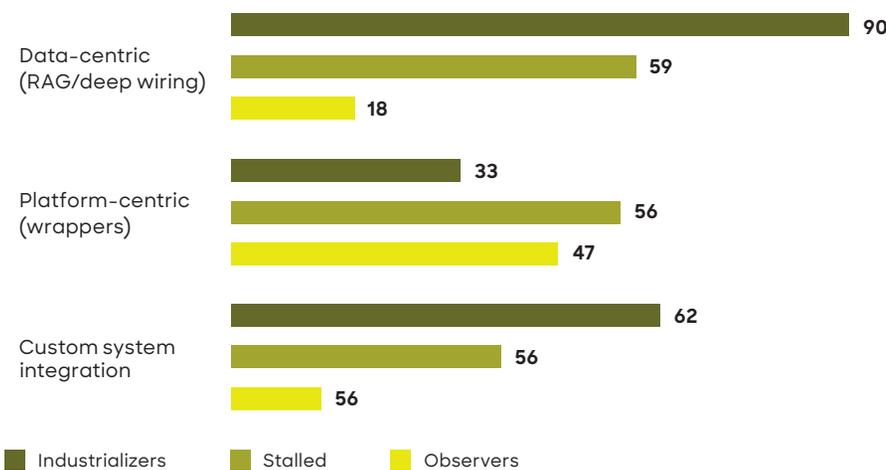
A different risk emerges among Observers, namely the rise of "shadow AI" or unsanctioned tools and agentic applications built outside of central IT oversight. Sixty-seven percent of Observers rely on low-code or no-code

platforms and 51 percent operate in siloed deployment mode, while 64 percent cite governance as a barrier. This combination suggests that when the corporate center cannot scale AI formally, activity shifts to the edges of the organization. The result is invisible tech debt and growing security exposure, as disconnected solutions proliferate without enterprise guardrails.

To summarize: Stalled companies have modernized their shopping list but not their wiring. They have invested in advanced tools but remain organizationally siloed and technically shallow. The result is the friction, latency and value gap discussed in Chapter 2 – and the reason why closing that gap requires not more spending, but a fundamentally different operating model.

J Industrializers' wiring goes deep into the data

Integration approach by segment (% selecting approach)



Source: Roland Berger

Net integration score
(data-centric minus platform-centric)

+57 points

Industrializers
▶ Deep wiring

+4 points

Stalled
▶ No clear direction

-29 points

Observers
▶ Surface wrappers

K Industrializers have mastered governance

Severity of challenges by segment (% citing "significant" or "critical barrier")

	Industrializers	Stalled	Observers	Gap between Stalled and Industrializers
Governance & monitoring	19%	56%	64%	37 points ¹
Safety & control	33%	58%	66%	25 points
Integration & security	29%	53%	60%	24 points
Reliability & consistency	29%	53%	55%	24 points
Accountability & liability	14%	28%	41%	14 points
Cost & processing speed	14%	34%	39%	20 points

■ Solved
 ■ Moderate (manageable)
 ■ Significant (active blocker)
 ■ Critical (transformation-blocking)

¹ **Largest gap** - Industrializers have mastered governance

Source: Roland Berger

The Industrializer Code

Five protocols that convert AI into
durable value

A large, bold, black number '4' is positioned in the bottom right corner of the page. The number is stylized with a thick stroke. The rightmost vertical stroke of the '4' is white, creating a sharp contrast with the olive green background.

Our diagnosis is now complete: The AI value gap persists not because companies lack ambition, capital or access to technology, but because most continue to treat AI as a series of projects to be bought rather than a capability to be engineered. The Industrializer Code describes how a small group of companies have broken out of profitless prosperity. They have not done so by spending more or moving faster, but by operating differently. Across the data, five consistent protocols emerge that govern how Industrializers approach AI transformation and succeed in converting activity into durable value. While these protocols apply to AI transformation in a very broad sense, they are particularly critical for agentic AI, where autonomous, multi-step execution amplifies both value potential and the risks of weak operating models. ► **L**

PROTOCOL 1

ORCHESTRATE HYBRID APPROACHES - DO NOT OUTSOURCE OWNERSHIP

Many Stalled companies and Observers assume that competitive advantage lies in selecting the right LLM (large language model) or purchasing the perfect off-the-shelf solution. The data reveals this to be a category error: The advantage does not derive from model selection or product choice, but from how the organization orchestrates capabilities around those tools. Industrializers illustrate the alternative by prioritizing speed (62 percent) and governance (71 percent) over cost (14 percent). And in doing so, they invert conventional procurement logic. While most work with external experts, they themselves retain ownership of the core. Partners handle the plumbing – the infrastructure, vector databases and RAG pipelines. But internal teams own the brain – the workflow logic, business rules and prompt design. This approach resolves the apparent contradictions observed earlier. Industrializers spend heavily on vendors, but they buy builders, not black boxes. By purchasing capabilities rather than finished

products, they gain vendor leverage instead of vendor lock-in. They see AI as a raw material to be shaped, not a solution to be installed.

PROTOCOL 2:

DON'T WRAP, REWIRE - RISE TO THE BROWNFIELD CHALLENGE

Stalled firms rely on wrappers to float AI on top of legacy systems, but they avoid integration at the core. Industrializers take the opposite approach: They transform existing systems rather than bypassing them with isolated greenfield solutions. Most pursue an AI-adapted brownfield strategy, going deep into the stack while ensuring orchestration. This enables 91 percent data-centricity, allowing systems not only to reason about data but to act on it. The work is harder upfront, but it compounds over time. Each deep integration iteration reduces friction for the next one, while wrappers merely accumulate as technical debt.

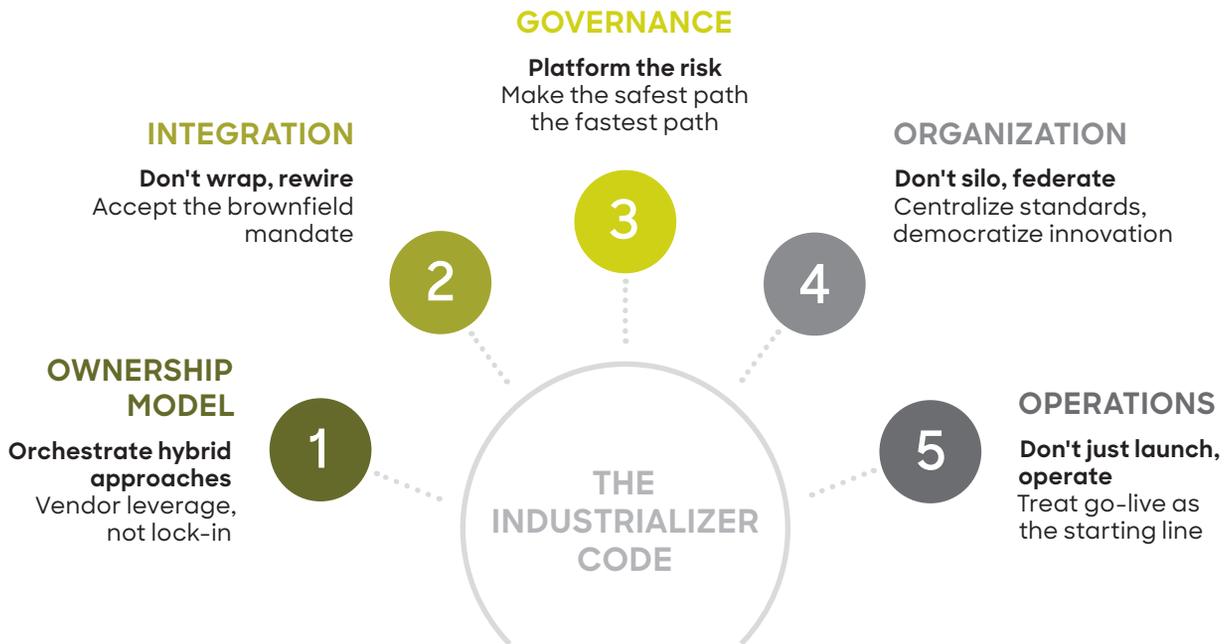
PROTOCOL 3

PLATFORM THE RISK - MAKE THE SAFEST PATH THE FASTEST PATH

Stalled firms experience governance as a series of traffic stops – lengthy approval processes that delay project initiation, and bureaucratic checks during deployment that add friction and tempt departments to find workarounds. Industrializers treat governance as path-building: They adopt a governance-by-design approach. By automating compliance as an integral part of platforms, they build controls directly into the operating environment. Governance checks run continuously rather than at the stage gate, removing the tradeoff between speed and safety. It is no coincidence that Industrializers are the only segment where speed and governance rank equally as top drivers. As one Industrializer put it: "Broader usage requires discipline; unified data and rollout guardrails are necessary to prevent uneven performance."

L The Industrializer Code

Five protocols for AI transformation



Source: Roland Berger

PROTOCOL 4

DON'T SILO, FEDERATE - CENTRALIZE STANDARDS AND DEMOCRATIZE INNOVATION

The market as a whole vacillates between two failure modes: siloed chaos, where departments deploy tools independently, and centralized bottlenecks, where IT blocks persist in the name of control. Industrializers choose a third path: 71 percent of them (versus 11 percent of Observers) operate under an enterprise-wide mandate for broad augmentation, while minimizing siloed deployment. A central hub provides the wiring – standardized interfaces, identity management and architectural guardrails – while distributed teams build use cases within that framework. This federated model is not a compromise between centralization and decentralization: It is a synthesis that captures the benefits of both.

PROTOCOL 5

DON'T JUST LAUNCH, OPERATE - TREAT GO-LIVE AS THE STARTING LINE

For many Stalled firms, success is declared at the moment of deployment. For Industrializers, that is where the real work begins. Industrializers focus on "Day 2 operations": continuous maintenance, monitoring and optimization. They are the only group to consistently cite operational

maintenance and troubleshooting as critical success factors, and they are far more likely to run complex, multi-agent systems in production. Their advantage does not come from launching more initiatives, but from running them better.

Together, these five protocols form the Industrializer Code – the key to how companies convert AI activity into scalable, repeatable value. They show that the path forward is not about spending more on AI, but about operating differently: engineering capabilities rather than assembling projects, building platforms rather than launching pilots and running systems rather than celebrating launches. The million-dollar question, of course, is how to translate these protocols into concrete leadership actions – as we discuss in the final chapter below.

● "The Industrializer Code
● inverts the question from 'What can AI do for us?' to 'What can we engineer with AI at scale?'"

Michael Schich

Vice President Tech & AI Innovation

Moving forward

A mandate for the C-suite



As we have seen, the decisive barrier to AI success has shifted. Technology readiness is no longer the key constraint, and access to capital is rarely the limiting factor. Instead, many organizations struggle to convert AI activity into sustained business impact they lack discipline in governing AI and operating it at scale. The evidence is consistent: Success no longer depends on which tools companies adopt, but on whether they adapt their operating model to run AI as a system rather than manage it as a collection of projects.

This shift has important implications for leadership. The market will not wait for Stalled organizations to catch up. As Industrializers continue to build on their advantage, the gap between them and the rest will widen into a lasting competitive divide. Crossing that divide requires a deliberate change in focus, moving away from exploration toward operations and from launching pilots toward running systems in production.

THE C-SUITE MANDATE

The Industrializer Code demands a fundamental shift in C-level responsibilities. AI industrialization cannot be delegated to a single function, nor can it be driven through isolated initiatives. Each role in the C-suite has a distinct mandate, and value emerges only when these mandates reinforce one another:

- **For the CEO: Don't innovate - orchestrate.** The CEO's role is not to champion individual AI initiatives, but to create the conditions for systemic value creation. This requires a clear in incentivization. Innovation theater, measured by the number of pilots or experiments, must give way to systemic value, measured by impact that multiplies across the organization. Crucially, the CEO must mandate a federated operating model: AI cannot be scaled through departmental silos. A target operating model is needed that orchestrates initiatives across the

enterprise by centralizing standards while democratizing innovation.

- **For the CTO and CIO: Build the control tower.** The tech leader's mandate is to build infrastructure that makes doing the right thing easy and doing the wrong thing hard. This starts with stopping the purchase of wrappers. Pre-built connectors may speed up deployment, but they create the illusion of integration while perpetuating disconnection from core systems. Instead, organizations need an AI control tower. A digital workforce of agents cannot be managed with spreadsheets: What is needed is a centralized infrastructure that monitors KPIs across production environments. The tech leader must also establish enterprise-wide data governance, ensuring a unified data layer that enables reuse, comparability and coordinated decision-making across AI initiatives. Without this foundation, even the most sophisticated agent architectures will fail to deliver at scale. In addition, governance itself must be automated: Compliance needs to move from manual stage gates to platform-level controls, making the safest path the fastest path.

- **For the CFO: Fund platforms, not projects.** The finance chief's task is to allocate capital in ways that enable scale rather than perpetuate fragmentation. Project-based funding creates silos: When every initiative has its own budget, every team reinvents the wheel. Industrializers therefore shift capital toward shared infrastructure. Efficient AI requires a "build-once, deploy-many" architecture, and funding a central platform is not seen as a cost center: It becomes the efficiency engine that eliminates redundant work across business units. This also changes how success is measured. Individual initiatives may fail even while overall capability is growing. CFOs must therefore assess portfolio value rather than project-level return.

This C-suite agenda is not theoretical: It directly reflects what Industrializers told us when asked about their most critical success factors. Three themes stand out. First, Industrializers consistently emphasize discipline and guardrails. They explicitly define rollout guardrails to prevent uneven performance across the organization. Governance, in this context, is not a constraint: It enables scale by ensuring consistency and reducing friction as AI systems expand.

Second, they stress the importance of **unified data**. They reject fragmented experimentation in favor of a single source of truth. Data management is not an optimization choice, but a prerequisite for enterprise-wide AI, enabling reuse, comparability and coordinated decision-making.

Finally, Industrializers highlight **operational maintenance**. They understand that the real work begins after going live. Sustainable value comes from continuous development and disciplined portfolio management across both AI use cases and core capabilities.

CASE STUDY

Transforming AI operations at a global industrial technology company

Challenge

We put this theory into practice with one of our clients, a leading global industrial technology company that had reached a tipping point. While the company had successfully launched multiple AI pilots, these initiatives remained confined to isolated silos. Leadership recognized that moving from experimentation and observation to industrialization would require more than ambition or additional pilots. The organization faced a governance vacuum, with scattered audit trails, fragmented budget tracking and AI agents that could not connect to core systems. Data governance also varied widely across regions, offering limited visibility into regional differences and resulting in inconsistent standards across the enterprise. The transformation reflected the patterns we describe as the Industrializer

Code: orchestrating rather than outsourcing, rewiring rather than wrapping and treating governance as infrastructure rather than bureaucracy.

Solution

Instead of launching further pilots, the company consolidated its fragmented efforts into a unified transformation program. In partnership with Roland Berger, it implemented a comprehensive initiative grounded in our Perform.AI orchestration framework, structured around three integrated pillars: Business & Portfolio, People & Enablement, and Technology & Architecture. Execution was organized across five coordinated workstreams: ► M

- **Use case and value steering:** A centralized approach was established to evaluate, prioritize and track AI initiatives across the enterprise, ensuring that investment decisions were guided by business impact rather than local experimentation
- **AI governance and roadmap:** Enterprise-wide governance policies were designed through the implementation of an AI control tower, while data governance was consolidated across regions. This approach harmonized standards without losing transparency into regional variation. It also replaced manual policy checks with automated enforcement embedded in the platform
- **IT technology enablement:** A shared technical foundation was deployed, including a central Model Context Protocol (MCP) hub that provided standardized connectivity and enabled AI agents to connect directly to core systems
- **Target operating model:** Roles, responsibilities and decision rights were clarified across the organization, creating a clear framework for coordinated execution between central functions and business units
- **Speedboats:** A fast-track mechanism was introduced for high-potential use cases that required rapid validation outside standard timelines without undermining enterprise-wide governance

Result

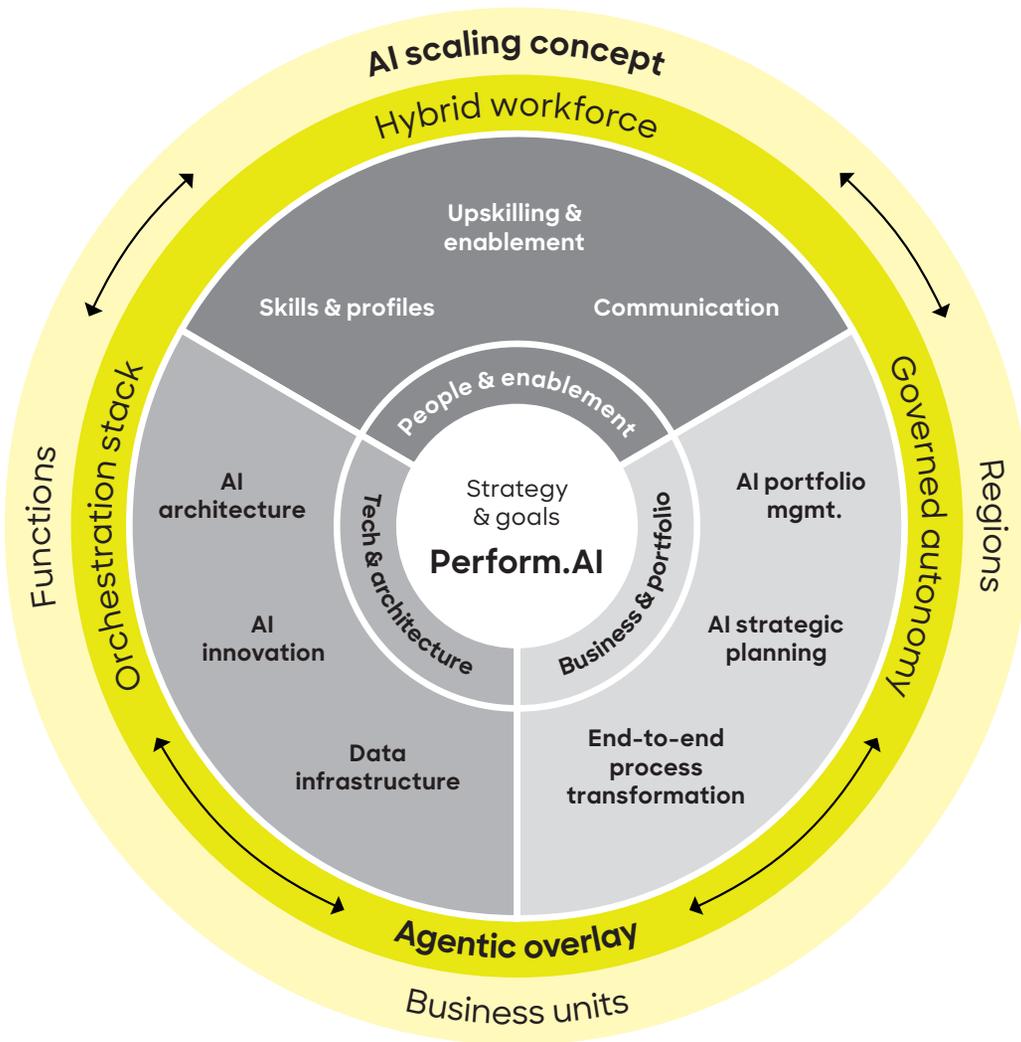
The company is now transitioning to the Industrializer group. By standardizing infrastructure through the AI control tower and aligning its operating model with the Perform.AI framework, it has established a scalable foundation capable of supporting advanced and autonomous workflows across the enterprise.

●● "The C-suite mandate isn't to accelerate adoption, it's to synchronize speed-to-production with speed-to-value. By securing governance at platform levels and making continuous operations, monitoring and maintenance a modus operandi, each deployment becomes the foundation for the next."

Maria Mikhaylenko
Senior Partner,
Global Managing Director

M Perform.AI

Unlock the potential of AI for superior business performance



Broad reach

Enable every employee to make use of AI in their daily work routine

Deep impact

Identify, evaluate and scale high-value AI use cases across functions, divisions and roles

6/ Conclusion

We began this study with a simple question: Why are almost 90 percent of firms underperforming despite record levels of AI investment? The data points to an uncomfortable answer: The market has not stalled for lack of ambition, budget or access to technology. It has stalled because many organizations have misunderstood the task. While AI deployment has accelerated, most companies continue to approach it as a sequence of projects to be launched and optimized in isolation. A small minority, by contrast, has adapted its operating model to treat AI as a standard industrial capability to be engineered, governed and run at scale. This difference explains today's phenomenon of profitless prosperity – and the widening AI value gap.

What separates these groups is not the models they buy, but the mechanisms they build around those models. Industrializers do not win by selecting better tools: They win by wiring AI into the enterprise, aligning speed-to-production with speed-to-value and building operating systems rather than celebrating launches.

As 2026 unfolds, this distinction will matter more, not less. The AI value gap is already widening, and it will morph into a lasting competitive disadvantage for organizations that fail to adapt. The lesson from the data is twofold: First, you cannot buy your way out of this transformation – you have to engineer your way through it. And second, as intelligence itself becomes a commodity, increasingly accessible to everyone, the only competitive advantage left is how well you industrialize it.

WHERE DO YOU STAND?

The client in our case study is on the path to success because it first diagnosed its specific blockage. Closing the AI value gap starts in exactly the same way: by understanding where you stand today.

Roland Berger's AI Lab diagnostic supports this process in two structured steps:

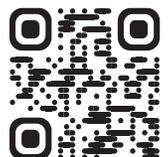
#1: Identify your architectural archetype

We assess how AI is currently built and operated in your organization. Most firms fall into one of four patterns: platform sprinters (speed seekers relying primarily on ecosystems); sovereign engineers (builders focused on intellectual property control); hybrid orchestrators (balancing best-of-breed tools with internal ownership); or capacity buyers (outsourcers of execution to scale quickly)

#2: Benchmark performance and value gaps

We benchmark your organization against the four performance segments identified in this study to reveal where value is leaking, how you compare with peers in your industry and which interventions will have the greatest impact

Want to know which archetype best fits your organization and how you perform relative to your market? Use the QR code to identify your position and define the most effective path forward.



This study examines the adoption of both generative AI and agentic AI. Many organizations increasingly address both as part of a single AI transformation agenda. However, agentic systems – which autonomously plan, act and adapt within defined parameters – represent the operational frontier where traditional processes intersect with intelligent automation. We therefore analyze the broader AI landscape while also providing dedicated deep-dives into the agentic dimension.

The study draws input from a survey of 203 executives at C-level and C-2 level with technology-related mandates and direct oversight of, or involvement in, AI initiatives. The sample spans Europe, the United States and Japan, with geographic representation as follows: DACH region (Germany, Austria, Switzerland) at 35 percent, France at 15 percent, the United Kingdom at 15 percent, the United States at 20 percent and Japan at 15 percent. Respondents covered all the major industries served by Roland Berger, including consumer, retail and logistics; health and utilities; financial and professional services; industrials and automotive; and technology, media and telecommunications.

To ensure conceptual consistency across responses, all participants were provided with the following definitions:

Generative AI (GenAI): Models that learn statistical patterns in data and can generate new, coherent content such as text, images, audio, video or code

Agentic AI: Systems in which AI models, often LLMs (large language models), repeatedly look at the current context, plan what to do next, choose actions such as calling tools, make adjustments based on feedback, and use this step-by-step loop to work toward goals over multiple actions

Workflow: Predefined, code-driven sequences of steps that deterministically control process logic, branching and context handoffs

Credits

AUTHORS

Maria Mikhaylenko

Global Managing Director
+39 389 762 0882
maria.mikhaylenko@rolandberger.com

Michael Schich

Vice President Tech & AI Innovation
+49 160 744 7298
michael.schich@rolandberger.com

Daniel Rohrhirsch

Senior Partner
+49 160 744 8014
daniel.rohrhirsch@rolandberger.com

Prince Punnackal

Principal
+49 160 744 6210
prince.punnackal@rolandberger.com

Patrick Ruthardt

Principal
+49 160 744 8014
patrick.ruthardt@rolandberger.com

We sincerely thank the following colleagues for their valuable contribution:

Friedrich-Nicolas Puhl and Hannes Gerstel

03.2026

ROLANDBERGER.COM

CONTACTS

AMERICAS

Brandon Boyle | brandon.boyle@rolandberger.com

Philipp Leutiger | philipp.leutiger@rolandberger.com

UK

Michael Knott | michael.knott@rolandberger.com

FRANCE

Gilles Oubuih | gilles.oubuih@rolandberger.com

Mathieu Jamot | mathieu.jamot@rolandberger.com

JAPAN

Yuzuru (Joe) Ohashi | yuzuru.ohashi@rolandberger.com

FURTHER READING



Toward data-centric AI

Why good data management is the key to unlock value in generative AI (2025)



This publication has been prepared for general guidance only. The reader should not act according to any information provided in this publication without receiving specific professional advice. Roland Berger GmbH shall not be liable for any damages resulting from any use of the information contained in the publication.

© 2026 ROLAND BERGER GMBH. ALL RIGHTS RESERVED.

ROLAND BERGER is one of the world's leading strategy consultancies with a wide-ranging service portfolio for all relevant industries and business functions. Founded in 1967, Roland Berger is headquartered in Munich. Renowned for its expertise in transformation, innovation across all industries and performance improvement, the consultancy has set itself the goal of embedding sustainability in all its projects. Roland Berger generated revenues of around 1 billion euros in 2024.



Publisher

Roland Berger GmbH

Sederanger 1
80538 Munich
Germany
+49 89 9230-0