Deloitte.

HKU Centre for AI, Management and Organization — CAMO

HKU and Deloitte China AI Adoption Index 2026
Enterprise AI Adoption in China:
The Paradox of Promise and Performance

February 2026

CONTENTS

Executive Summary

Across industries worldwide, the race to harness artificial intelligence (AI) is accelerating. Yet, while enthusiasm runs high, execution often lags behind lofty ambitions. Under this backdrop, Deloitte China and the Centre for AI, Management and Organization at the University of Hong Kong (HKU) have jointly established the "Deloitte-HKU Lab for Organizational Transformation" to unveil findings from a survey of more than 100 C-suite executives that there is a disconnect between the transformative promise of AI and the measurable outcomes realised so far.

Most companies have begun integrating AI into customer-facing and operational functions, signalling that the technology has moved firmly beyond experimentation. However, only a small proportion of companies have successfully scaled these initiatives to achieve a meaningful impact on profitability. Nearly half of respondents acknowledge that realized returns fall short of expectations, indicating that the hype surrounding AI's return on investment (ROI) still exceeds the reality.

The **obstacles to success are largely organisational and executional rather than technical**. Legacy structures, talent limitations, and inconsistent implementation strategies remain the main barriers to enterprise-wide adoption. While current AI priorities continue to centre on customer service and process optimisation, emerging investments in research and development mark a slow but significant shift from efficiency towards innovation.

Despite the uneven results, **optimism prevails**. Most executives plan to expand AI budgets over the next three years, confident that the technology will ultimately deliver sustainable growth, competitive advantage, and new sources of enterprise value.

# Respondents Profile (China mainland and Hong Kong)

To understand the corporate AI landscape, the study surveyed more than 100 senior executives across different industries in China mainland and Hong Kong. The leadership composition provides a distinctly strategic perspective: **70% of respondents are CEOs, CTOs, or CIOs** — decision-makers who directly shape organisational priorities and investment frameworks.

The survey also leans toward large-scale enterprises, with 44% representing firms of more than 5,000 employees. As a result, the insights reflect the challenges and successes of companies managing complex, multi-market operations, where AI deployment requires both governance discipline and cultural change.

| Who We Surveyed

44%
Large Enterprises
(with 5,000+ employees)

37%
Medium Enterprises
(with 1,000 – 5,000 employees)

70% CEO/ CTO/ CIO

| Top Three Industries

Consumer/ Retail
27%

17%
Automotive

Top
3

15%
Energy, Resources
and Industrial
Products

Others
41%
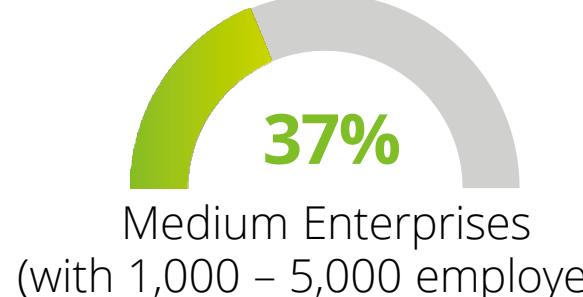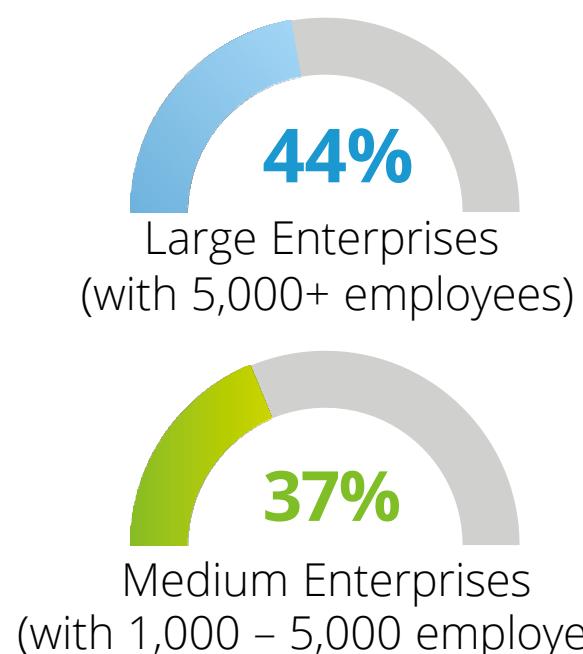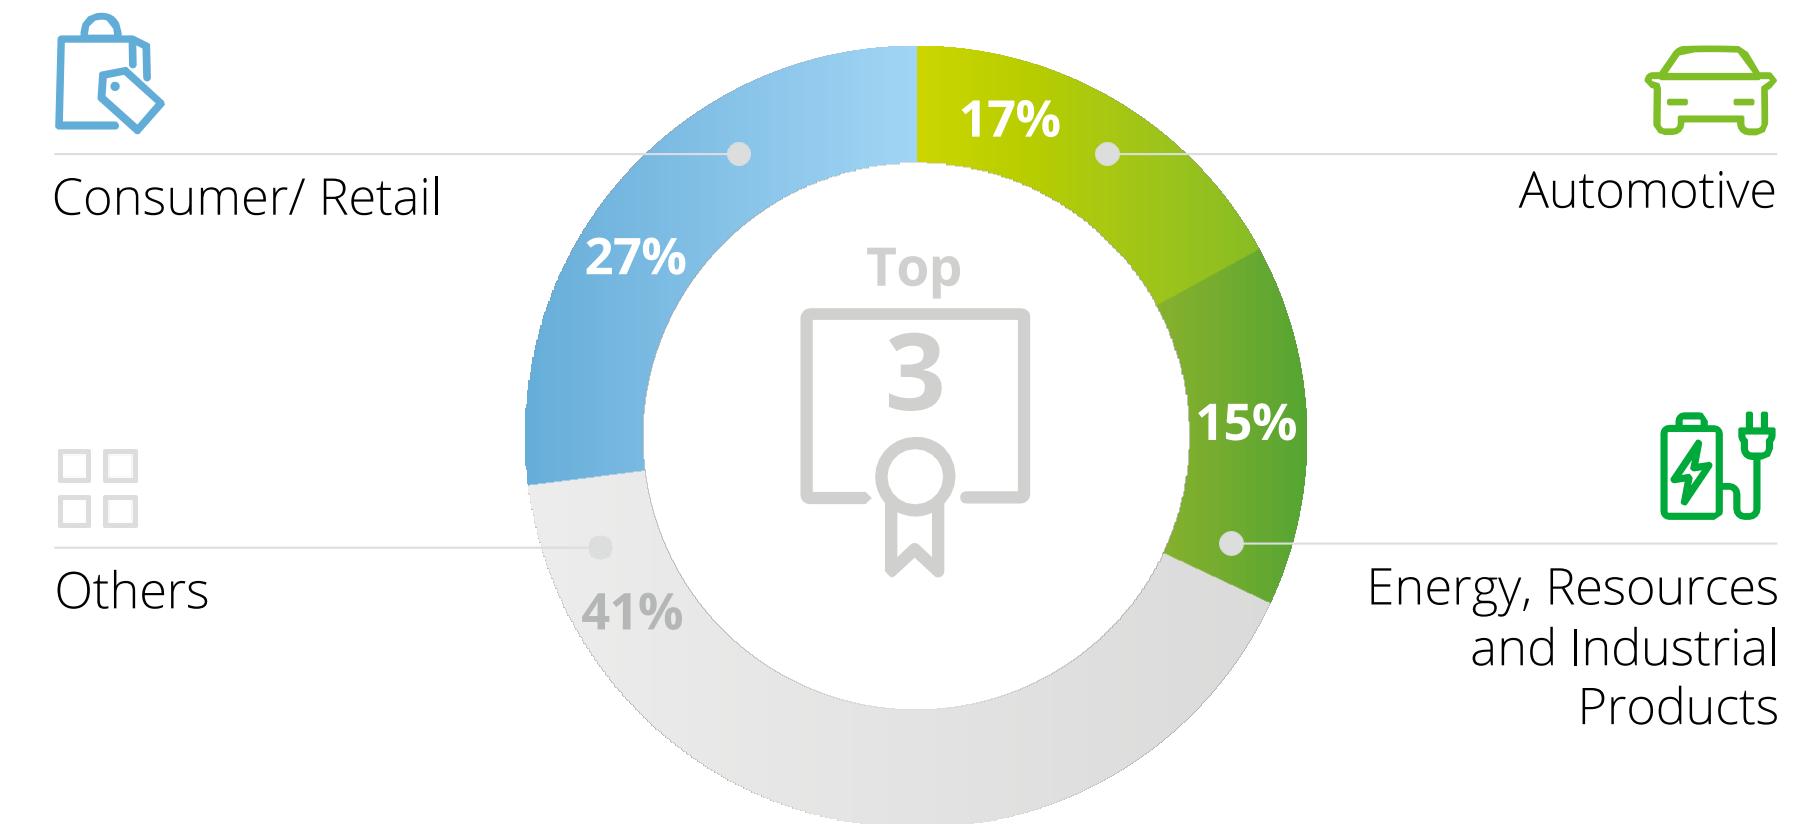
Sector representation is led **by consumer and retail (27 percent), automotive (17 percent), and energy, resources, and industrial products (15 percent)**. These are industries under acute pressure to balance customer-driven transformation with operational efficiency, positioning them at the forefront of AI experimentation and early adoption.

4

Current State of AI Adoption

MAJORITY OF BUSINESSES ARE NOW MOVING FROM THEORY TO PRACTICE, GENERATING TANGIBLE BUT LOCALIZED BENEFITS

AI adoption is now nearly universal among businesses, 69% of organisations are experimenting with or scaling limited pilots, only 23% have achieved operational deployments with measurable financial impact. Just 4% describe themselves as fully transformational, embedding AI across the business

AI Adoption Maturity

Exploring 36%		Implementing 56%		Transforming 8%	
Awareness 8%	Pilot/ Experimental 28%	Early Adoption 41%	Operational 15%	Scaled Operational 4%	Transformational 4%

"Considering" — "Tangible Benefits" — "P&L Impact"
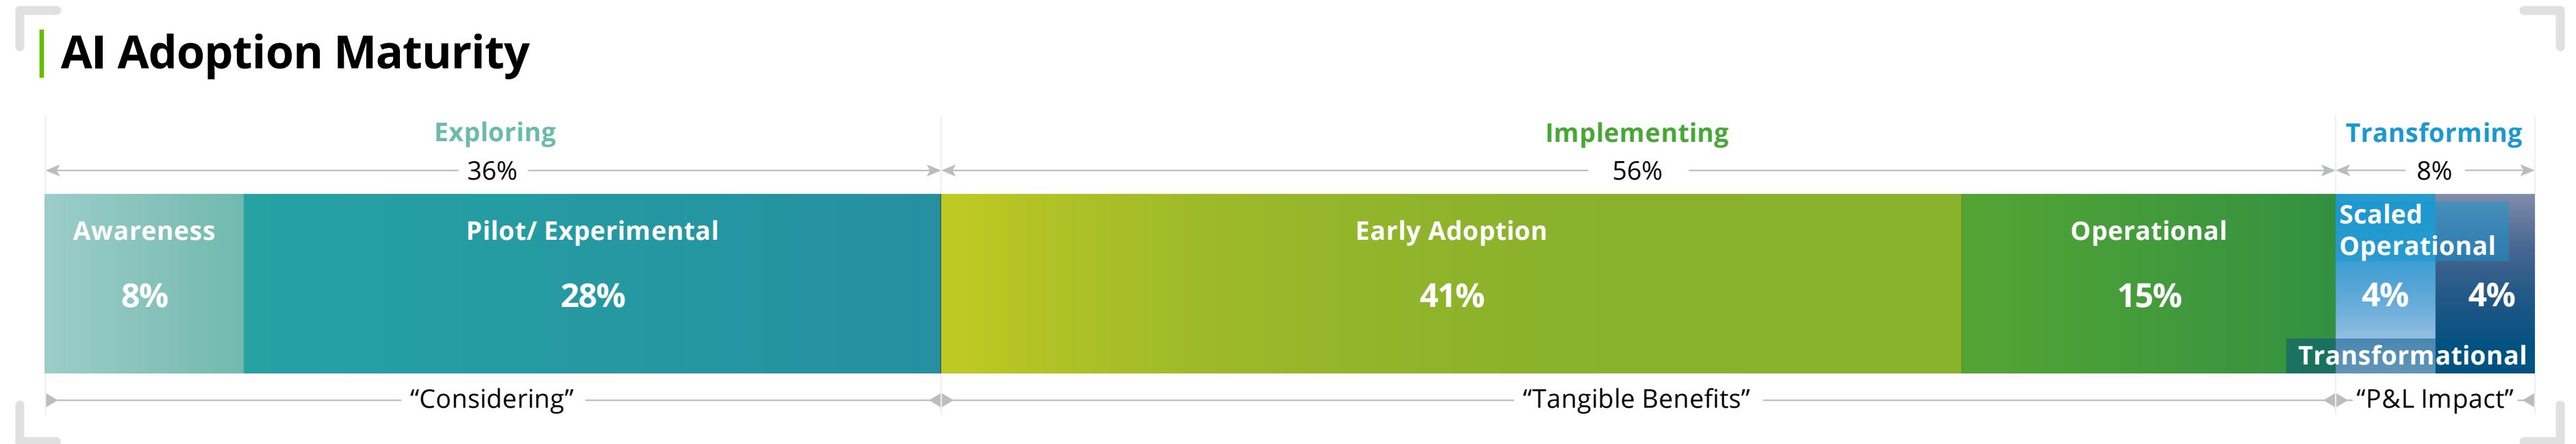
Exploring Phase (36%):
More than one-third of respondents are still in the initial "awareness" phase, focused on understanding AI's potential and testing early concepts. This includes an awareness group (8%) just beginning to identify use cases and build foundational knowledge, and a pilot/ experimenting group (28%) running proof-of-concepts in controlled environments to gauge feasibility and ROI.

Implementing Phase(56%):
The majority of enterprises are now moving from theory to practice, generating tangible but localised benefits. Within this category, 41% fall under early adoption, where successful pilots are being scaled within specific departments. Another 15% have reached operational integration, embedding AI into regular workflows to enhance efficiency and decision-making — though impact remains siloed.

Transforming Phase (8%):
A minority of firms have achieved company-wide maturity, realising significant financial and strategic returns. Among them, 4% have scaled AI-driven efficiencies across the organisation, while a leading 4% have reached the transformational stage, where AI is no longer a support tool but a catalyst for new revenue models, market differentiation, and lasting competitive advantage.

AI Expectations vs. Reality

BUSINESSES SYSTEMATICALLY OVERESTIMATED RETURNS

While executive optimism around AI remains high, the reality of financial returns paints a more mixed picture

Comparing expected versus actual ROI across surveyed projects **reveals a consistent pattern of overestimation**, reflecting both the immaturity of AI governance and the tendency to view emerging technologies through a lens of inflated promise.

Perhaps the most striking finding is the **emergence of negative returns**. None of the surveyed leaders anticipated losses, yet **9%** of AI initiatives ultimately delivered a **return below zero**. These unplanned shortfalls highlight the hidden costs of implementation — from data-integration challenges and change-management hurdles to unforeseen operational disruptions.

Over-optimism also affects the high-yield brackets. For example, **21%** of executives expected returns in the range of 10–25%, but only **13%** reached that level. This overestimation trend persists across nearly every higher-yield category, indicating that even mature enterprises struggle to accurately forecast the economic value of AI.

Equally revealing is the rise in projects classified as "**not relevant**" for direct ROI measurement. Initially, **30%** of leaders considered ROI as secondary, viewing these efforts as strategic bets or learning investments. After implementation, that proportion grew to **36%**, signaling that many projects failed to develop quantifiable success metrics. This lack of measurable impact underscores a broader weakness in AI project governance and evaluation discipline.

## Expected vs. Actual ROI from AI Initiatives

Bar chart comparing Expected ROI (green) and Actual ROI (blue) across ROI brackets:
- <0%: Expected 9%, Actual (none)
- <10%: Expected 34%, Actual 34%
- 10-25%: Expected 21%, Actual 13% (GAP)
- 26-50%: Expected 9%, Actual 6% (GAP)
- 51-100%: Expected 2%, Actual 1% (GAP)
- >100%: Expected 4%, Actual 2% (GAP)
- Not Relevant: Expected 30%, Actual 36%

Legend: Expected ROI, Actual ROI

6
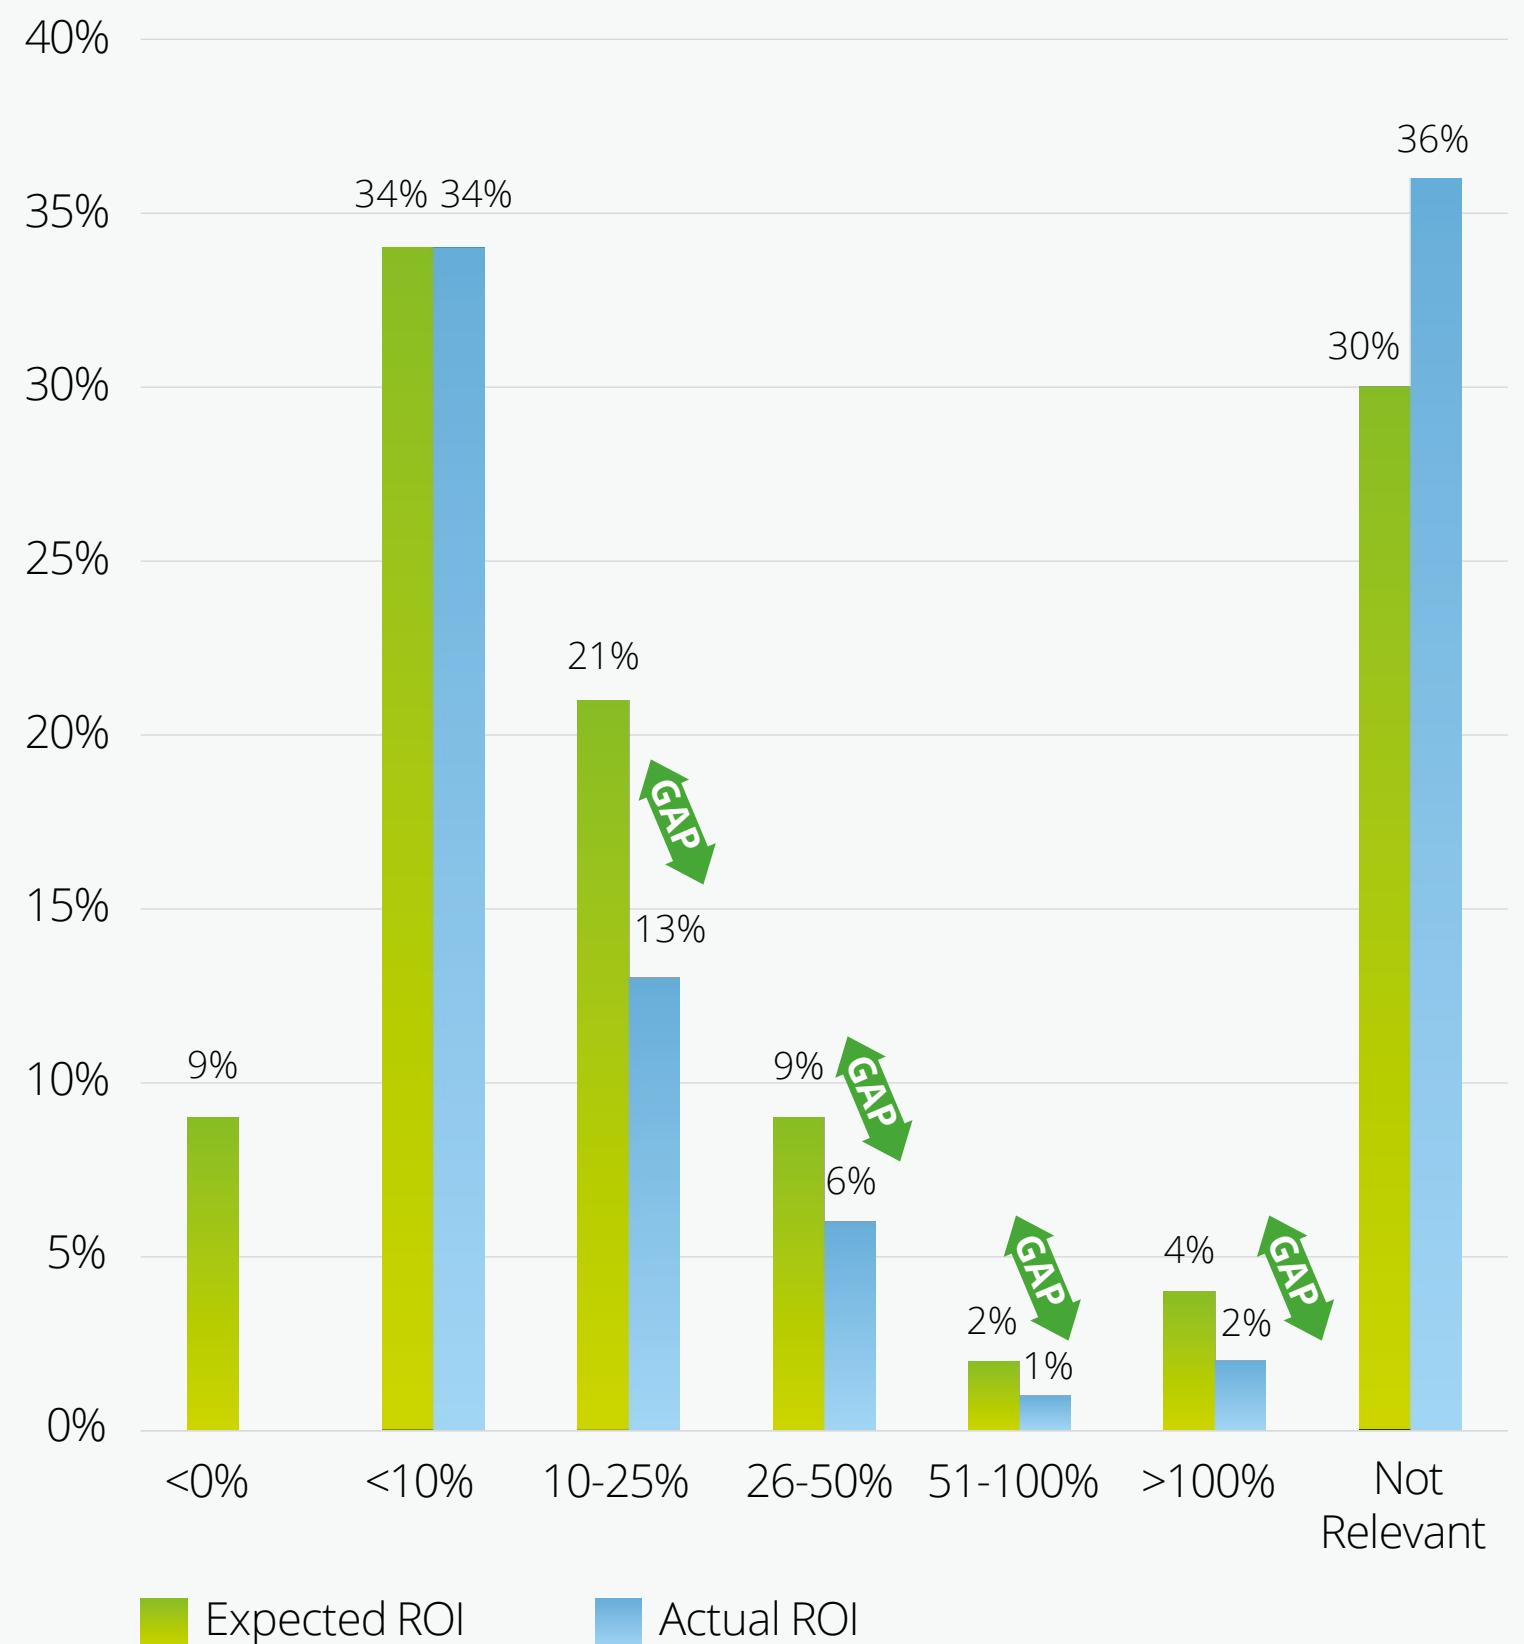
The ROI Expectation Gap

THE REALITY IS THAT MAJORITY OF AI INITIATIVES UNDERDELIVER

Inflated expectations creates unrealistic optimism while key performance indicators are often poorly aligned with strategic goals

ROI Expectation Gap
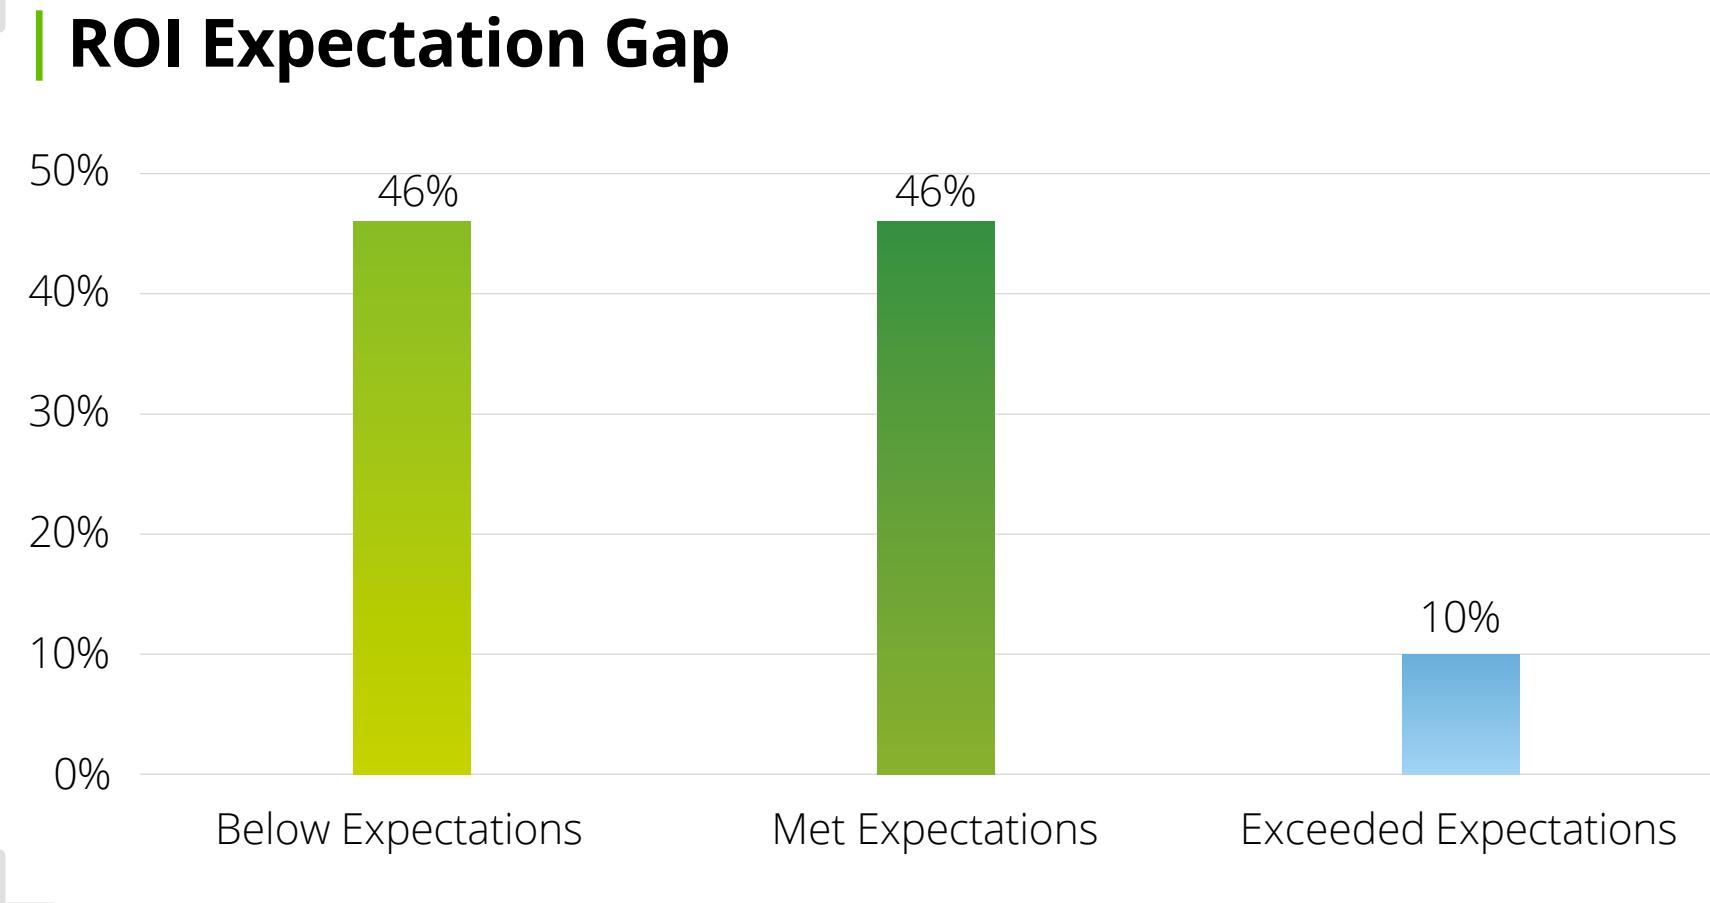
Nearly half of executives **(45%)** reported that their AI initiatives **fell short of expectations**, while only **one in ten** said their projects **exceeded expectations**. The remaining **46%** achieved results that merely **met initial targets**, suggesting progress but not transformation.

Taken together, these results highlight that while enthusiasm for AI remains strong, most businesses are still grappling with how to convert investment and enthusiasm into consistent, measurable business impact.

These results are not random fluctuations. They stem from a fundamental weakness in how AI initiatives are conceived, funded, and managed. Too often, business cases rest on **unrealistic optimism** rather than evidence-based modelling. Furthermore, **key performance indicators** (KPIs) are either undefined or poorly aligned with strategic goals, making it difficult to assess financial impact once projects are deployed.

Closing this gap demands both cultural and operational change. Leaders must move beyond the excitement of emerging technologies and adopt a **rigorous measurement mindset** — one grounded in realistic projections, disciplined KPI design, and continuous financial validation.

The future success of AI adoption by businesses will hinge not only on innovation, but on accountability. Organisations that combine strategic ambition with measurable discipline will be the ones to turn AI's potential into sustainable economic performance.

7

Key Barriers

ORGANISATIONAL AND EXECUTIONAL ARE KEY OBSTACLES TO AI IMPLEMENTATION

AI adoption is not merely a technology rollout, but an organisational transformation across leadership, culture and governance

| The Five Critical Barriers to AI Success

Strategic & Planning

33%

Organisational & Cultural

50%

Technical & Infrastructure

39%

Execution & Results

47%

Risk & Compliance

26%
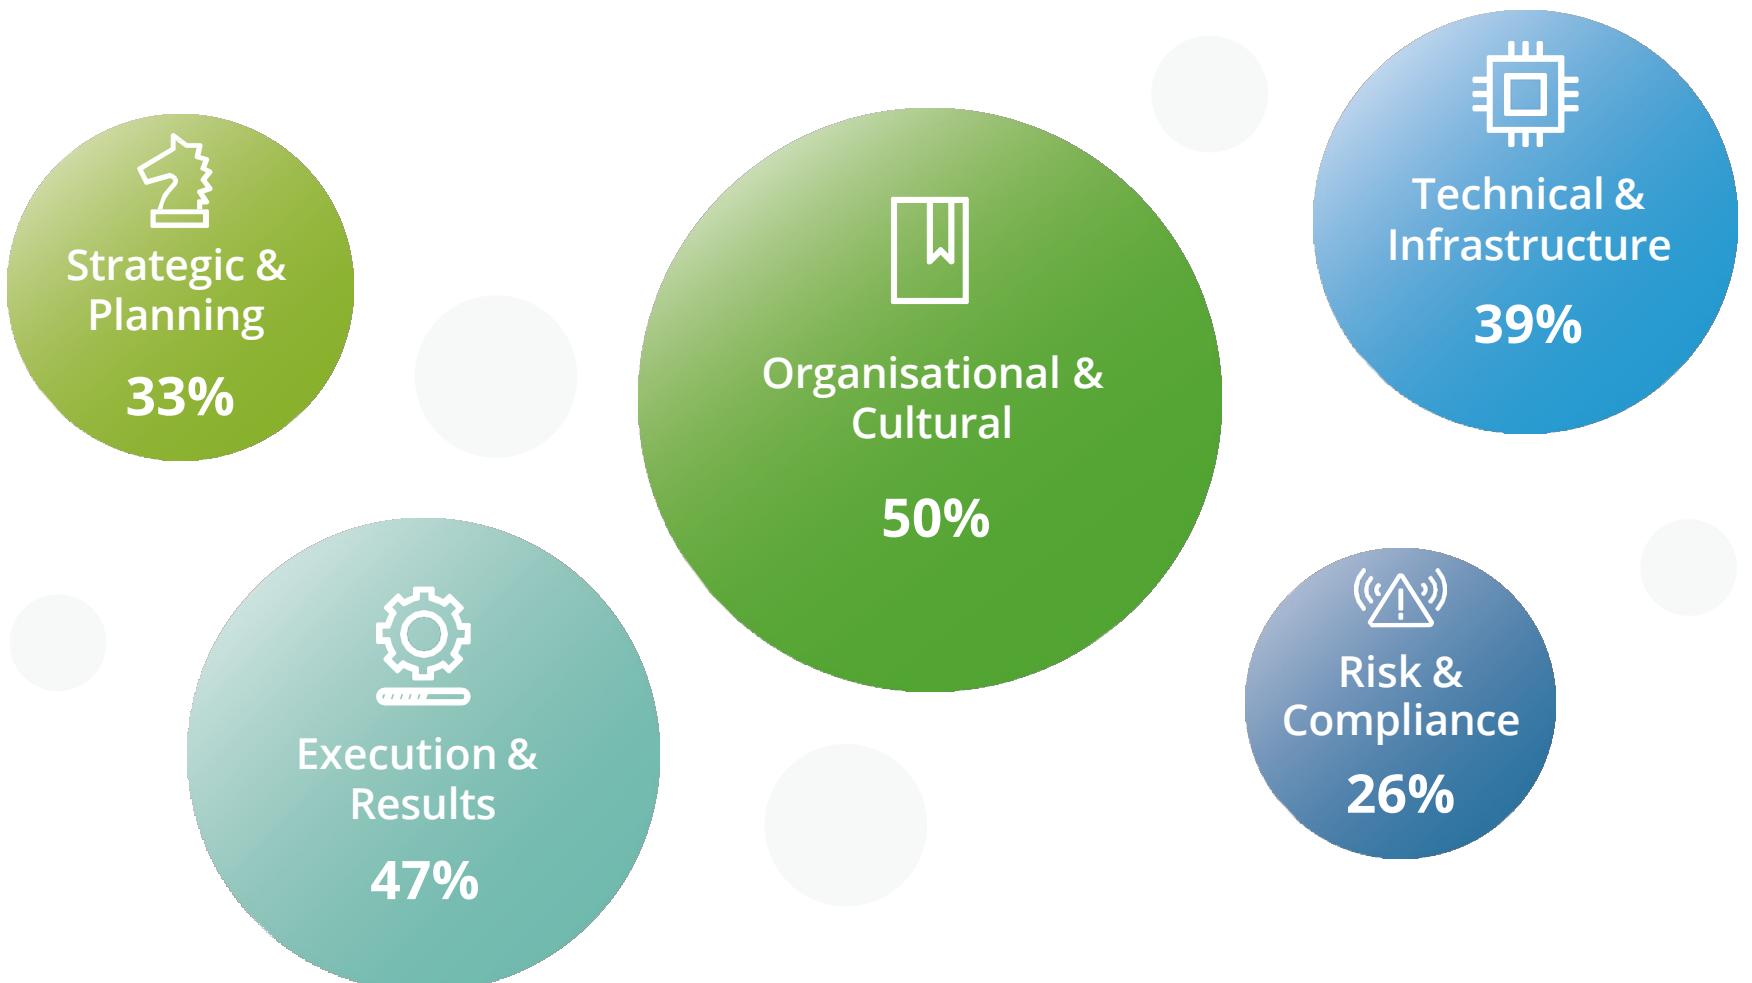
To understand why many AI initiatives fail to deliver the expected financial results, executives were asked to identify barriers both broadly — across organisational, technical, and strategic categories — and within each specific domain.

The findings are clear: **organisational and cultural hurdles (50%)** and **execution challenges (47%)** dominate, far outpacing **technical limitations (39%), strategic misalignment (33%),** and **risk or compliance issues (26%)**. Together, these findings highlight that the **dominant barriers are organisational and executional, not technical**.

These findings underscore a fundamental truth: AI adoption is not merely a technology rollout; it is an organisational transformation.

Without alignment across leadership, culture and governance, even the most advanced models are unlikely to generate sustainable business impact.

Top 10 Reasons AI Underperforms

AI USAGE OFTEN FAILS TO MEET EXPECTATIONS DUE TO SILOS, LACK OF QUICK WINS AND POOR DATA

The three most cited reasons are deeply interconnected and point to systemic organisational dysfunction

1 Siloed departments preventing collaboration **33%**	**2** Lack of immediate/ viable results **32%**	**3** Data quality/ availability issues **31%**	**4** Organisational culture not conducive to AI adoption **28%**	**5** Insufficient understanding of AI capabilities & limiations **27%**
6 Implementation taking longer than expected **24%**	**7** Unclear business case/ ROI **24%**	**8** Lack of skilled AI/data science personnel **24%**	**9** Security & privacy concerns **23%**	**10** Difficulty integrating AI with existing system **22%**

At the heart of AI underperformance lies a systemic challenge: the structure and behavior of organisations themselves. The most frequently cited issues point to fragmentation, resistance and misunderstanding across businesses. The result reveals that the most significant challenges are not purely technical but are deeply rooted in organisational structure, strategy and culture.

9

The Challenge Isn't Technology

THE MOST SIGNIFICANT CHALLENGES ARE DEEPLY ROOTED IN ORGANISATIONAL STRUCTURE, STRATEGY AND CULTURE

| ORGANISATIONAL & CULTURAL BARRIERS

Siloed departments (33%) is the biggest reason for underperformance. AI projects, especially those with transformative potential, require cross-functional data and cooperation. When departments operate as independent fiefdoms, it becomes nearly impossible to access the necessary data, align on objectives, and integrate solutions across the business.

Cultural resistance (28%) — a culture that is risk-averse, lacks a data-driven mindset, or distrusts new technology will actively or passively sabotage AI initiatives. Successful AI requires a culture that embraces experimentation and data-informed decision-making.

Limited understanding of AI's capabilities (27%) — when decision-makers don't understand what AI can realistically achieve, they set unrealistic expectations, approve poorly- conceived projects, and cannot effectively oversee implementation.

| EXECUTION & PROJECT MANAGEMENT FAILURES

Many organisations struggle to sustain momentum once AI projects are underway, leading to a loss of momentum and stakeholder support.

Lack of immediate, visible results (32%) — AI projects can have long development cycles. If there isn't a clear plan to deliver incremental value or "quick wins," stakeholders can lose confidence and withdraw support before the project can deliver on its long-term promise.

Implementation taking longer than expected (24%) reflects the unforeseen complexity of integrating AI into legacy operations. This often results in extended timelines, budget overruns and reduced cross-departmental confidence.

Finally, **unclear business cases or ROI definitions (24%)** — Projects launched without a clear, quantifiable goal are destined to be perceived as failures because there is no agreed-upon definition of success.

Key Takeaways

FOCUS LESS ON TECHNOLOGY AND MORE ON FOSTERING A COLLABORATIVE CULTURE AND ESTABLISHING CLEAR BUSINESS CASES WHILE MANAGING EXPECTATIONS

TECHNICAL & DATA INFRASTRUCTURE CHALLENGES

While technology is not the top obstacle, it remains a critical part of the AI adoption story.

Data quality and availability (31%) — the third most cited challenge — continues to hinder AI scaling. Models are only as effective as the data on which they are trained, and fragmented or inconsistent datasets often result from siloed departments. Without clean, integrated data pipelines, even the most advanced algorithms cannot achieve their full potential.

Difficulty integrating AI with existing systems (22%) further compounds the problem. Legacy IT infrastructures, often rigid and outdated, struggle to support modern AI platforms. The consequences include higher implementation costs, slower deployment, and reduced.

RESOURCE AND GOVERNANCE GAPS

Talent and oversight complete the picture of why AI ambitions frequently stall.

A shortage of skilled AI and data science professionals (24%) continues to constrain businesses' progress in rolling out AI initiatives, as demand for technical expertise far outstrips supply.

Meanwhile, **security and privacy concerns (23%)** increasingly shape the boundaries of AI adoption. As organisations manage increasing volumes of sensitive customer and operational data, compliance with regulations such as GDPR introduces both necessary caution and operational friction. In some cases, these governance requirements can slow or even halt deployment altogether.

KEY TAKEWAY

The overarching conclusion is that **AI adoption is fundamentally a change management challenge**, not just a technology implementation.

The three most cited reasons – silos, lack of quick wins, and poor data quality – are deeply interconnected and point to systemic organisational dysfunction.

To overcome these barriers, leaders must focus less on the algorithms and more on fostering a collaborative culture, establishing clear business cases, managing expectations through incremental wins, and building a solid, unified data foundation.

Current Areas of AI Application

HIGHEST IN CUSTOMER SERVICE, MARKETING AND INFORMATION TECHNOLOGY WHICH DELIVER IMMEDIATE VALUE

Across the business, AI has gained its strongest foothold in functions that directly interact with customers or core business technologies

Current AI Application Usage Area

Area	Value
Customer Service	58
Marketing	54
IT	53
R&D	45
Operation	45
Sales	39
HR	32
Financial	28
Sourcing	20
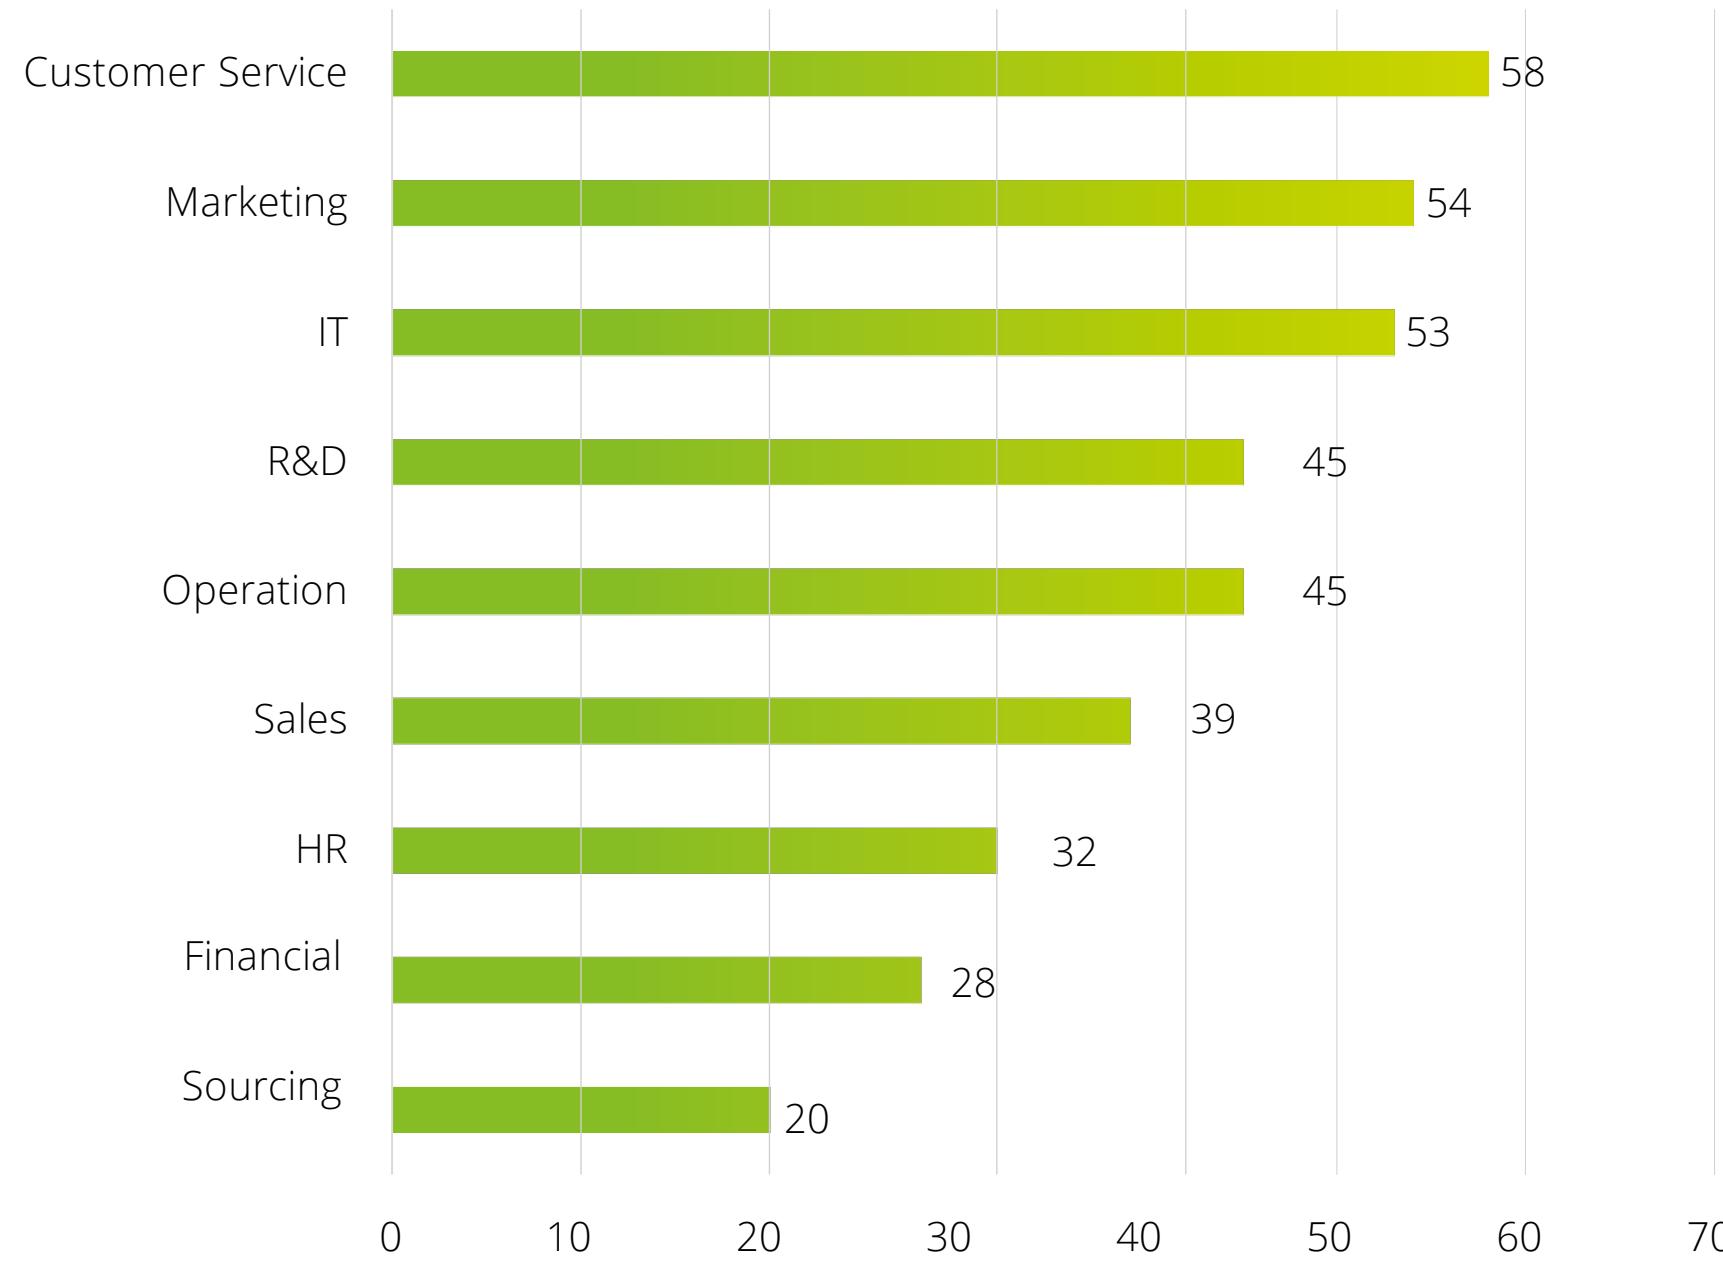
Today, AI adoption is highest in **customer services (58%)**, **marketing (54%)**, and **IT/technology (53%)** — areas where automation, personalisation, and rapid data processing deliver immediate value.

Moderate adoption is taking shape in **research & development (45%)** and **operations (45%)**, highlighting AI's growing role in process optimisation and product innovation. By contrast, **back-office functions** such as **finance (28%)**, **human resources (32%)**, and **sourcing (20%)** remain relatively underdeveloped, still reliant on manual systems and legacy workflows.

While today's adoption map reflects AI's operational roots, the outlook for the next three years shows clear signs of evolution. Executives plan to deepen their investments in proven areas like customer engagement, while expanding AI into innovation-centric functions such as R&D — signalling a strategic shift from efficiency to discovery.

12
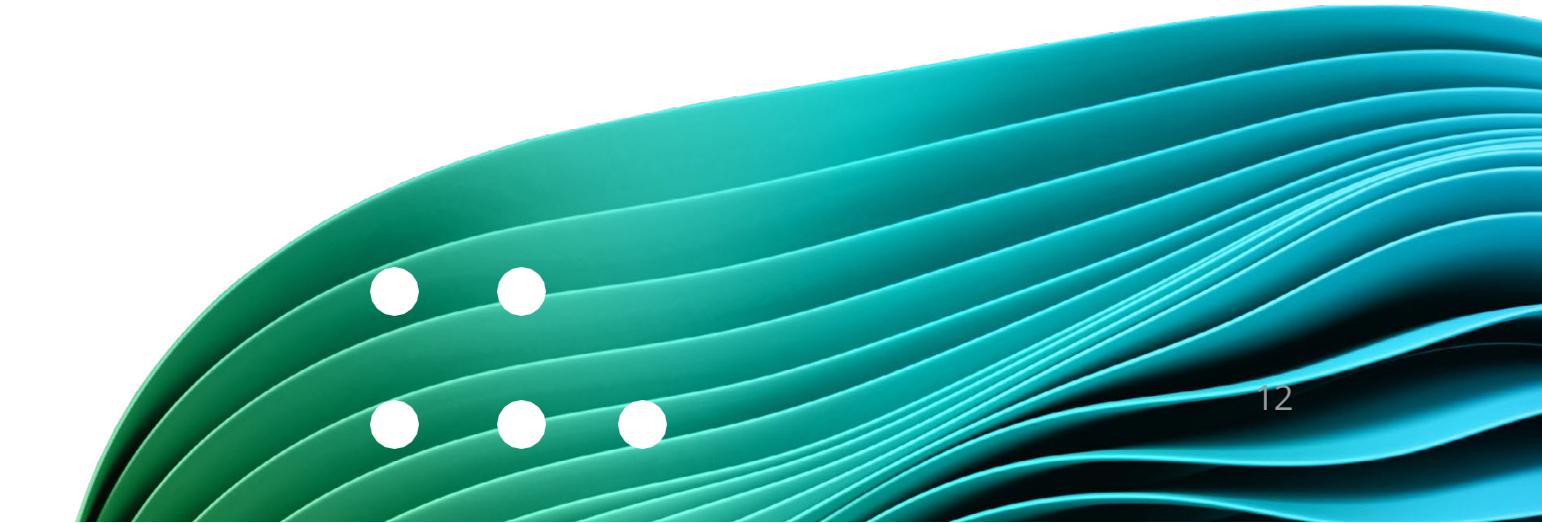
Future Priorities

MOVING FROM CORE EFFICIENCY TO DISCOVERY FUNCTIONS SUCH AS R&D

While future priorities largely mirror the present, the data reveals notable strategic shifts in where executives plan to direct their AI investments

AI Usage vs. Future Priority

Future Priority

80%

70% ● Customer Service

60% Marketing ●
Sales ● ● Operation

50% ● R&D

40% ● IT
● Financial

30%

20% ● HR

10% ● Sourcing

0%
0%   20%   40%   60%   80%
Current Usage
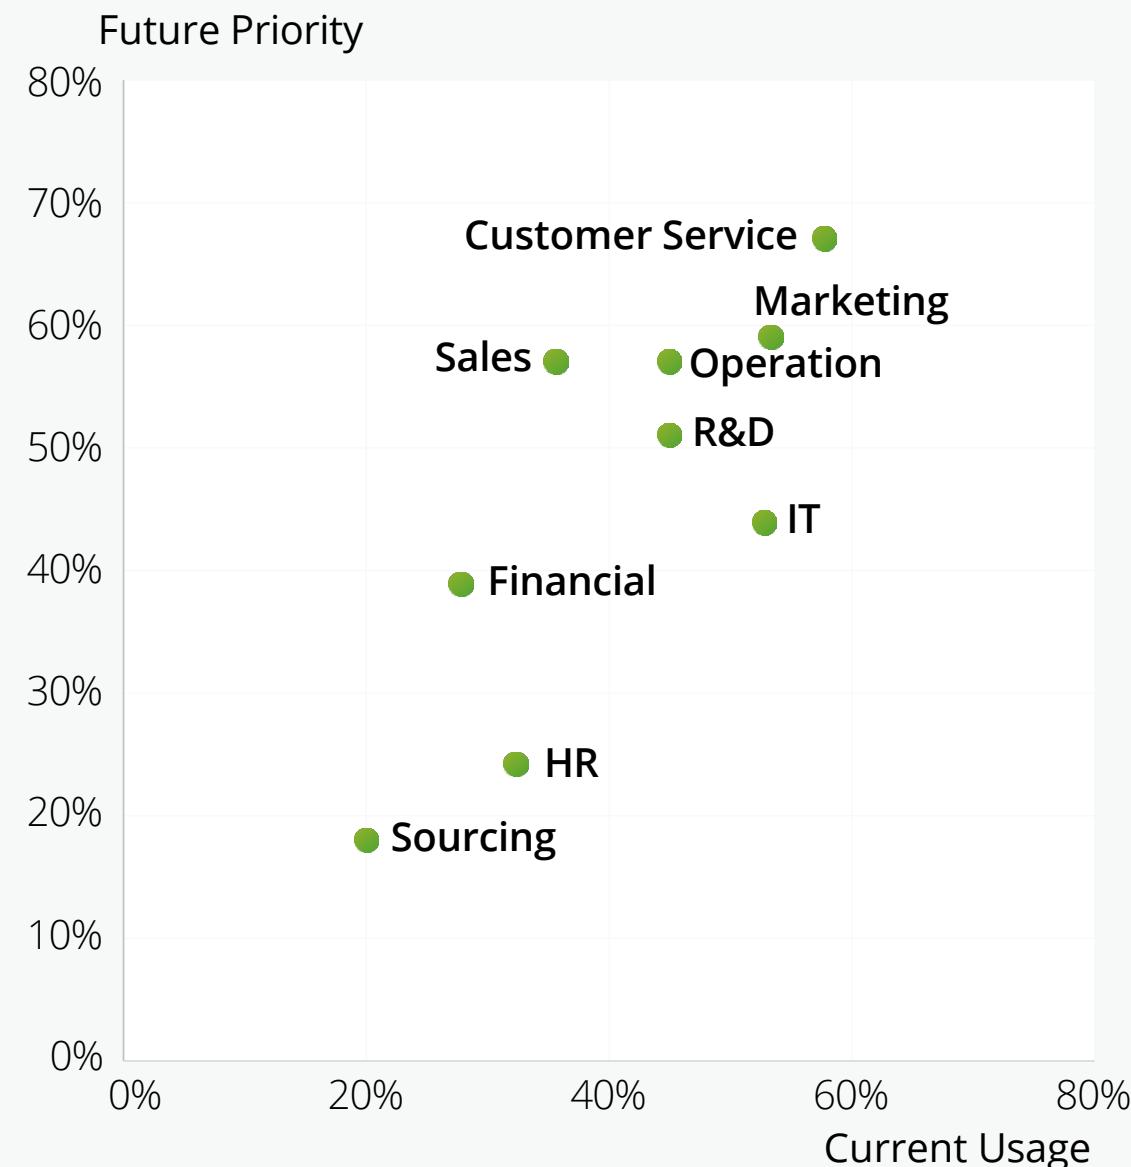
Established Leaders: Deepening Strength in Customer-Facing Functions

Customer service and marketing remain the two pillars of businesses' AI deployment. Survey responses cite adoption rates of **58%** and **54%**. Regardless of the discrepancy, both sources align on one point: these functions are not only mature, but also strategic priorities for future expansion. Customer service, in particular, stands out with a **67% future priority score**, indicating a strategy to deepen investment in this proven area.

The Rise of R&D as a Strategic Frontier

The most revealing shift is occurring in **research and development**. Current adoption sits at **45%**, with future investment priorities expected to reach **51%**, signalling AI's evolution from a cost-efficiency tool to an innovation engine. This transition reflects growing confidence that AI can generate intellectual property, accelerate experimentation and shape entirely new business models. This highlights a clear pivot from using AI for current operational efficiency to leveraging it for future innovation.

Mature but Stabilizing: IT and Technology Functions

IT and technology departments — long the backbone of AI experimentation — are showing signs of maturity. Both textual and graphical data place current adoption at **53%**, yet future priority drops to **44%**, suggesting that foundational investments have largely been made. For many organisations, IT now serves as the enabling infrastructure rather than the next growth horizon for AI.

Lagging but Necessary: The Untapped Potential of Back-Office Functions

Functions such as **finance, human resources** and **sourcing** remain on the periphery of AI transformation, consistently clustered in the lower-left quadrant of the strategic map. Their lower adoption and limited future prioritisation indicate that these areas are still viewed as secondary, despite the significant long-term efficiency gains that intelligent automation could unlock.

This chart maps business functions based on their current level of AI adoption (x-axis) against their stated priority for future investment priority (y-axis).

Future Outlook

SHORT-TERM CONSERVATISM FOLLOWED BY SIGNIFICANT LONG TERM BULLISHNESS

Executives expect AI's impact to grow over time, the gap shows that leaders understand that realizing AI's full potential is a marathon, not a sprint

| Gap Between 5 years and 2 years

Difference in Share
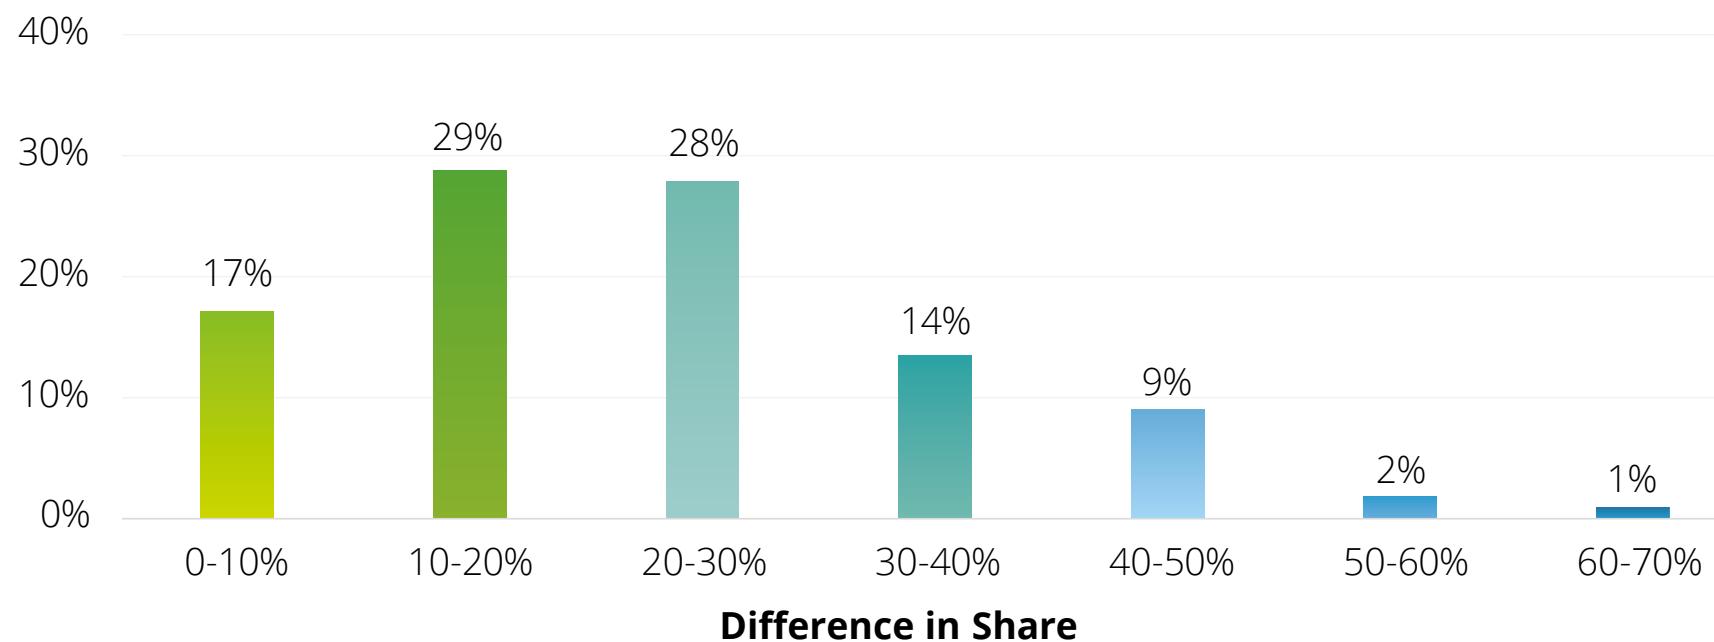
| Expected Business Value Driven by AI

Share of business value driven by AI
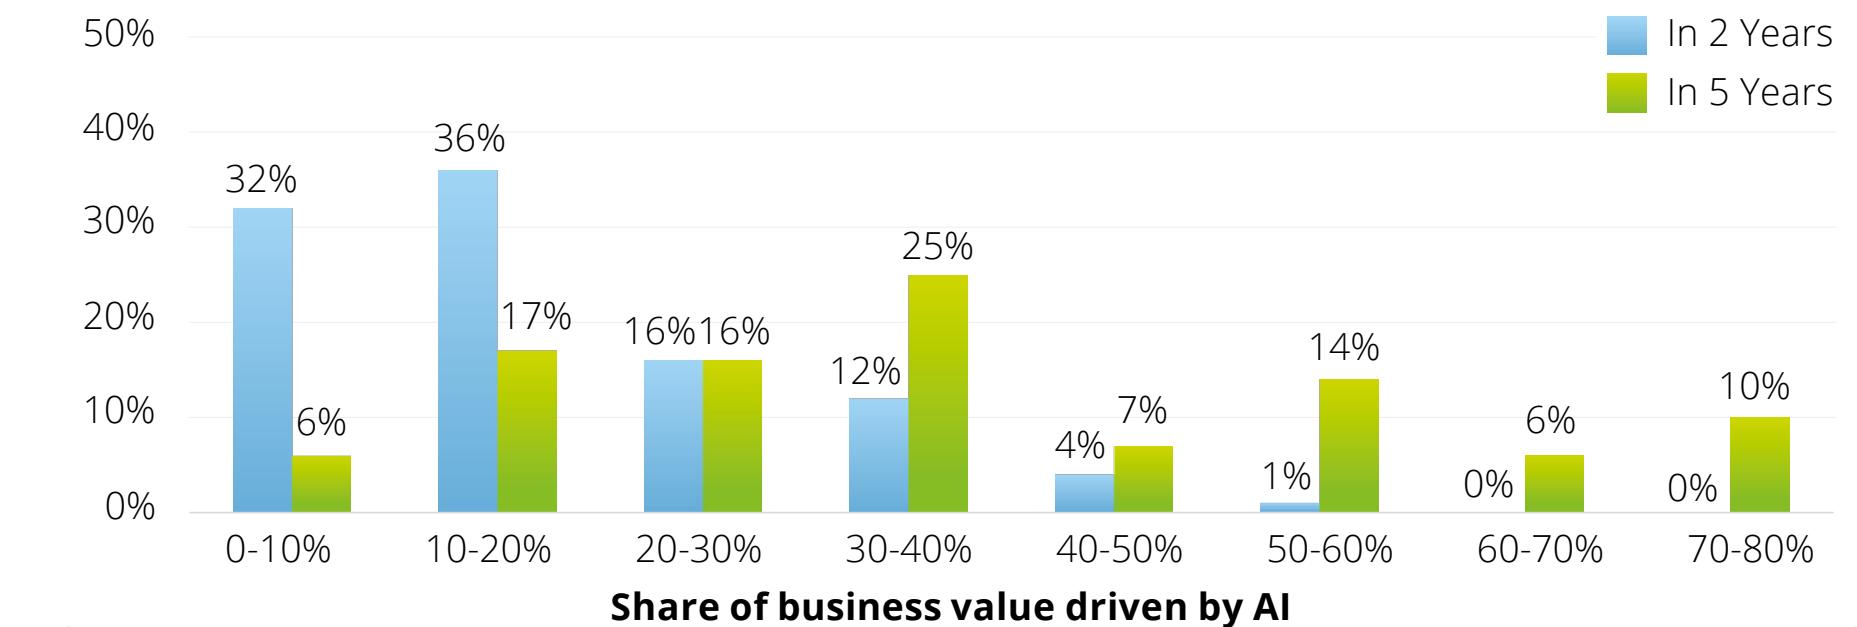
Anticipated Value Acceleration

The distribution clearly shows that leaders expect a significant jump in AI's contribution in the medium term. The largest cohorts of executives anticipate the **value share from AI to grow by an additional 10-20%** (cited by 29% of respondents) **and 20-30%** (cited by 28%) in the three years between the two-year and five- year outlooks.

Quantifying Long-Term Optimism

Nearly 75% of all executives surveyed expect AI's share of business value to grow by at least 10% in this period, with over 45% expecting it to grow by more than 20%. This provides strong quantitative evidence for the conclusion: leaders are patient, expecting modest gains in the immediate future but are highly optimistic about AI's transformative, long- run potential. They are planning and investing now for a major payoff down the road.

Key Implications

As artificial intelligence moves from the experimental fringe to the corporate core, executive leaders are tasked with navigating a significant paradox: while long-term optimism for AI has never been higher, the short-term results are often underwhelming. An extensive survey of C-suite executives reveals a landscape of widespread but shallow adoption, elusive returns, and organisational, rather than technical, barriers. Understanding the following key implications is critical for any leader aiming to convert AI investment into tangible, enterprise-wide value

1 AI Adoption Is Widespread But Shallow

AI has become nearly universal across large businesses, yet its depth of impact remains limited. The data reveals a clear divide between today's reality and tomorrow's ambition.

In the near term, 59% of leaders expect AI to contribute less than 20% of total business value — with the largest single group (32%) predicting an impact below 10%. These figures underscore the early-stage nature of most initiatives: pilot programs, departmental tools, and automation trials that operate in isolation from core business strategy.

Looking ahead five years, the outlook shifts dramatically. A majority (54%) of executives anticipate AI will drive more than 20% of business value creation, with the 20–30% and 30–40% contribution brackets emerging as the new centre of gravity.

The message is clear: while AI is now embedded across most organisations, its integration into the financial and operational core is still incomplete. The next phase of progress will depend on whether leadership can evolve fragmented experimentation into a cohesive, enterprise-wide system of value creation.

2 Returns Remain Elusive

Nearly half of executives acknowledge that **AI outcomes have fallen short of** expectations, exposing a persistent gap between ambition and execution. Yet this underperformance rarely stems from technical limitations; it originates from a combination of two factors rooted in strategy and governance:

The first is **overly ambitious business cases**. In the rush to adopt AI, many projects are launched with inflated expectations and poorly defined goals. The strategic promise of "transformation" often overshadows the practical need for a clear, measurable business case.

The second is the **weak return measurement**. Many early-stage projects are labeled "strategic" and are therefore not subjected to the same rigorous ROI tracking as other investments. This lack of discipline makes it impossible to quantify success, learn from failures, and make data-driven decisions about where to scale investment.

Closing this gap requires more than optimism — it demands **operational discipline**. leaders must instill a culture of disciplined execution, demanding clear KPIs for every AI initiative and balancing the portfolio between long-term exploratory projects and efficiency-focused applications with clear, measurable returns.

3 Barriers Are Organisational More Than Technical

When AI initiatives stumble, the root cause is seldom the algorithm itself. The most persistent and challenging hurdles are usually human and structural. While data quality and system integration are indeed significant tasks, they are often symptoms of deeper organisational issues:

Departmental silos — The most effective AI models require access to broad, cross-functional datasets. Organisational silos create technical data silos, preventing the holistic view needed for high-impact AI.

Cultural resistance — A culture that is risk-averse or not data-literate will naturally resist the changes that AI brings to workflows and decision-making processes. True adoption requires active change management and a commitment from leadership to foster a data-first mindset.

Execution discipline — AI is not a traditional IT project. It requires an agile, iterative approach of "test, learn, and scale." Organisations that stick to rigid, waterfall-style project management will struggle to keep pace with the dynamic nature of AI development.

In essence, AI success depends as much on **organisational transformation** as on technical capabilities. Companies that can align their culture, leadership, and workflows around data-driven decision-making will move faster and extract greater value from their investments.

4 Future Priorities Largely Reinforce The Present

Looking ahead, most businesses plan to deepen their AI investments in the very functions where it already delivers tangible returns: **customer service, marketing** and **operations**. However, two critical shifts are emerging:

In **customer service**, companies are moving beyond basic automation toward predictive and hyper-personalised engagement — using AI not only to reduce costs but to drive loyalty and lifetime value. Chatbots and virtual assistants are giving way to adaptive service ecosystems that anticipate needs before customers articulate them. The goal is to evolve from cost-saving automation to revenue-driving customer loyalty.

At the same time, the most significant strategic pivot is the increasing **prioritization of R&D**. This signals a move from using AI to optimize existing business processes to using it as an engine for innovation. Companies are beginning to leverage AI to accelerate drug discovery, design new materials, and create entirely new products and services, turning it into a true source of competitive advantage.

This signals a gradual but critical evolution in AI strategy—from a tool for efficiency to a catalyst for growth.

5 Optimism Fuels Investment Despite Uncertainty

Perhaps the most revealing survey finding is the persistence of **executive optimism**. Despite elusive short-term ROI and ongoing operational hurdles, confidence in AI's long-term potential remains remarkably strong. This bullish sentiment continues to drive investment and budget increases.

This conviction reflects a deeper strategic view: AI is not simply a new tool but a **foundational technological shift** — one on par with the advent of the internet or the smartphone revolution. For most leaders, the perceived long-term risk of being left behind is far greater than the short-term risk of uncertain project returns. Leaders are willing to invest through the uncertainty because they are confident that mastering AI is not just an option, but an existential necessity for future relevance and success.

As one respondent observed, *"We can't afford to wait for perfect conditions — by then, it will be too late."* Executives are therefore choosing to invest through volatility, betting that proficiency in AI will determine relevance, resilience and leadership in the next decade of global competition.

The AI Transformation Journey

Cultivating People's Readiness for AI

When thinking about employees' readiness for AI, our research identified three problems: uncertainty, fear of replacement, and the self-image problem.

The Uncertainty Problem: "What will this AI actually do?"

Slack's 2024 global survey of more than 17,000 office workers found that 61% of employees had spent less than five hours learning about AI and 30% had received no training at all. In the absence of knowledge, opinions can be polarized. Some employees dismiss AI as mere hype, while others assume it can do everything. Companies' uncertainty about AI extends beyond technical capability. One audit firm, for example, identified AI opportunities across its workflow, but both clients and auditors resisted, citing regulatory risk. In the end, the firm abandoned many of its AI-based approaches. To address these concerns, firms need to embed AI governance into their daily work and make it intuitive for every employee. Effective governance not only safeguards against unintended consequences but also helps demystify AI. As an example, in 2018 DBS Bank introduced the PURE framework—Purposeful, Unsurprising, Respectful, and Explainable—to evaluate every AI use case. Instead of relying on lengthy policy documents, employees are guided by four simple questions: Is the use purposeful and meaningful? Will the results surprise customers? Does it respect customers and their data? Can the outputs be explained? This approach reduces uncertainty while ensuring responsible use. DBS also established a Responsible Data Use Committee to review projects that do not meet PURE requirements. With an easy-to-grasp framework and clear human oversight, the bank empowered employees at all levels to innovate responsibly. By 2023, AI had already generated $274 million in value for DBS.

Fear of Replacement: "Will I keep my job?"

When employees suspect they are training a system that will replace them, they comply minimally. They drag their feet when asked to "label data" or "teach the model." This "training trap" slows down the adoption of AI in service, retail, and manufacturing firms. Companies can counter this by sharing the upside—offering training royalties for data and labeling work, productivity bonuses tied to realized gains, and career guarantees that channel efficiency gains into reskilling rather than layoffs. Because many of these benefits are based on future promises that are easy to break, firms must make promises credible and easy to verify. One e-commerce company, for instance, pledged to increase total labor spending by 1% annually to demonstrate its commitment to investing in employees. This 1% number is easy to check and hard to manipulate, so it helped build trust with workers. It also gave workers formal seats on the AI steering committee and greater influence over personnel decisions. These enhanced power of the workers further reinforce trust. Another form of resistance is fault-finding: holding AI to much higher standards than humans. At a leading insurance company, employees' fault-finding missions led to demand for unrealistically high levels of accuracy from AI systems, which in turn slowed deployment and drove up investment costs. Independent studies and external audits comparing AI and human outputs helped restore realism. Finally, replacement fears ease when AI fuels growth rather than contraction. If technology expands the business, efficiency gains feel like opportunity, not threat.

The Self-Image Problem: "Will I appear competent?"

Fear of status loss can be even more powerful than fear of job loss. We've observed engineers who quietly use AI tools but conceal it to avoid appearing less skilled. Many worry that admitting to using AI could make them seem lazy, incompetent, or even dishonest. Similar image concerns result in radiologists ignoring AI recommendations to protect professional pride. One financial services firm flipped this stigma by launching an "AI Masters" program that fast-tracks employees who demonstrate exceptional AI skills, regardless of their seniority. This celebrates the proficiency with AI as sophistication and forward-thinking, not laziness or incompetence. Equally important is designing AI use for professional dignity. Companies can position AI as a tool that presents facts without judgment, while leaving final conclusions to professionals. This framing reinforces expertise rather than undermining it. Some firms have created private "second-opinion consoles" where employees can consult AI without fear of embarrassment or reputational risk.

The AI Transformation Journey

Redesigning Workflows

AI adoption often falters when organisations treat it as a simple overlay on existing processes. True transformation demands systematic change at three levels: individual workflow (nodes), cross-functional connections (edges), and system-wide coordination (networks).

The Node Level: Transforming How Individuals Work

A consulting firm's legal team initially used AI like a spell-check tool—running it at the very end of traditional reviews. The approach produced negligible benefits, as AI was only 100% accurate for 40% of error types. By restructuring the workflow so that AI conducted the first pass—checking only error types it handled best—lawyers could focus solely on the remaining ones. This redesign demonstrated how rethinking workflows unlocks AI's value. Some firms accelerated the change by setting "mission impossible" goals that force teams to abandon old habits and discover new ways of working. One Boston startup, for instance, faced resistance to using AI in document preparation. To break through, it required that documents—previously completed in a week—be finished within a single day. The extreme time pressure left employees no choice but to integrate AI from the start and redesign their processes around it.

The Edge Level: Redesigning Connections

The edge level focuses on how improved local judgment and data can transform inter-departmental processes and decision-making flows. At a Japanese cosmetics company, beauty advisors in stores once supplied untrusted, anecdotal feedback. Generative AI helped them analyze customer conversations and traffic patterns, giving structured insights. Headquarters, now confident in the data, built a two-way loop: Campaigns could launch faster and be tweaked in real time based on credible field intelligence. The edge between local operations and central planning became a responsive circuit rather than a one-way command.

The Network Level: Orchestrating System-Wide Impact

To generate real business impact from AI, companies must consider the network level—how improvements across multiple nodes and edges interact within the broader system. Without this perspective, AI can simply shift bottlenecks from one part of the network to another, limiting overall performance gains. This phenomenon is common because many organisations concentrate their gen AI efforts in a few high-impact areas—such as marketing, customer service, or software development—while overlooking the interdependence across business units.

A major car manufacturer discovered this when it adopted generative AI to boost productivity in automotive software development (enhancing one set of nodes), enabling faster design iterations, code generation, and feature testing. Yet the overall vehicle production network showed little improvement, as hardware manufacturing became the primary bottleneck. The enhanced software development nodes were now waiting on unchanged hardware nodes, and the edges connecting them couldn't handle the increased pace of software output.

Addressing such network-level challenges requires coordinated action across all nodes and edges. Organisations should begin by mapping the entire network topology—understanding how workflows between teams and identifying potential bottlenecks. AI adoption should be synchronized across interconnected nodes so that capacity improvements are matched throughout the network.

The AI Transformation Journey

Navigating Power and Influence

AI shapes who gains and who loses inside organisations. The resulting politics—over data, hierarchy, and accountability—often prove more formidable than technical issues. Successful AI adoption often requires redesigning governance structures, adjusting incentive mechanisms, and, in some cases, relying on senior leadership to broker agreements and remove barriers. Here are three specific problems we observed in our research:

Resource Hoarding

Organisations quickly discover that AI's hunger for data and knowledge collides head-on with deeply ingrained competitive instincts. At a large Chinese IT firm, researchers uncovered that programmers were 16–18% less likely to recommend AI access to their own teammates, effectively hoarding knowledge to preserve their personal edge. Across business units, larger, more successful divisions that own sophisticated AI models and valuable datasets often see little incentive to share them with smaller units that could benefit most. Sharing can feel like enabling potential internal competitors while diluting their own performance metrics. DBS Bank confronted this resistance by designing an incentive structure that rewarded units for converting proprietary datasets into reusable assets on the central platform. A key metric tracked the percentage of each unit's use-case-specific datasets that had been transformed into shareable resources. This approach breaks down silos by motivating both large and small units to contribute high-quality, accessible data.

Hierarchy Disruption

AI unsettles the traditional hierarchy built on two pillars: experience and headcount. The first weakens when junior employees armed with AI outperform seasoned veterans. In one software firm, programmers with only two years' experience began producing more and cleaner code than colleagues with five years' tenure. Juniors felt that they were doing more for less. Some companies responded by expanding competency models to explicitly include AI mastery and by shortening promotion ladders. When advancement cycles shrink from five years to one or two, and mastery of new tools is rewarded, young employees see immediate payoff for learning. The second pillar, power through headcounts and control over resources, creates even stronger resistance. Managers are the gatekeepers of AI adoption, yet their authority often depends on the size of their teams. When efficiency threatens to shrink those teams, their self-interest can quietly derail otherwise valuable AI initiatives. In one translation department, leaders hesitated to automate because doing so would shrink their headcount, bonuses, and prestige. OPPO, the smartphone maker, tackled this by staging an AI tournament where every employee had equal access to tools, and results were ranked by department. Suddenly, managers had to champion AI adoption or risk public embarrassment if their teams lagged.

Accountability Attribution

AI also disrupts the traditional balance of blame and discretion within organisations. Its precision turns fuzzy responsibility into hard data—and that can create new political friction. At Dingdong Maicai, a Chinese grocery e-commerce company, AI systems began tracing every customer complaint back to the exact department at fault. When a customer received spoiled fruit, algorithms could pinpoint whether procurement bought poorly, storage mishandled goods, or delivery caused the damage. What had once been shared uncertainty became explicit accountability. With this change, departments that had long operated under ambiguity now found themselves publicly exposed. The binary nature of algorithmic judgment—assigning full responsibility to one side—ignored the grey areas of real-world operations. This led to escalated disputes and complaints from department heads. The lesson is that perfect accountability can undermine organisational harmony. Dingdong eventually changed its system of attribution, allowing final judgment to humans. The goal was not to reject transparency but to buffer it with trust. Effective AI adoption requires knowing when precision helps performance—and when it merely fuels internal politics.

The AI Transformation Journey
Pulling Multiple Levers for AI-Driven Transformation

A professional services firm with 2,200 practitioners—primarily software developers and product managers—began piloting GenAI initiatives in mid-2023. Within weeks, individual productivity rose by 30–40%, yet by mid-2024, overall performance—measured by productivity and time-to-delivery—remained flat. Several factors explained the gap. Developers lacked incentives to boost output, fearing efficiency gains might trigger layoffs. The flood of new AI tools created inconsistent practices across teams, disrupting established standards and complicating project management. Meanwhile, junior developers often outperformed senior ones, but work assignments and recognition still followed traditional hierarchies. To address these challenges, the firm pulled levers across people, processes, and politics.

On the people front, it redefined its competency model to explicitly reward AI proficiency, making expertise visible across the organisation and turning mastery into a source of pride. To counter fears of replacement, the compensation structure was overhauled: Base salaries were reduced to 80%, while performance-based incentives of up to 40% were added, directly linking efficiency gains to individual rewards.

The process dimension was overhauled to embed AI throughout the workflow. Developers became data and process stewards, responsible for following standardized data definitions, coding practices, and AI protocols while participating in training to strengthen process consistency. A unified end-to-end framework harmonized AI integration across development stages, with updated SOPs incorporating AI-augmented steps for easier training and compliance. At the organisational level, a centralized governance model with defined checkpoints and Business Process Stewards ensured alignment across data, AI, and workflows.

Political barriers were confronted head-on. Job grades expanded from six to 14, with biannual reviews enabling rapid promotion or demotion. This system rewarded AI adopters with greater responsibility and influence, realigning incentives that once favored tenure over capability.

By mid-2025, these changes began to pay off. Productivity rose by 22%, enabling a 10% price cut that boosted sales by 20%. Labor costs grew by 5% as the firm reinvested in its workforce, reinforcing its commitment to employees. Overall profitability improved by 3%, demonstrating that AI-driven transformation translated into tangible business value. Building on this foundation, the firm expanded into markets that had previously been too price-competitive to enter. AI-supported development also shortened the learning curve for new programming languages, enabling the company to broaden its offerings.

This case shows that true AI transformation goes beyond technology. By aligning incentives, redesigning processes, and reconfiguring organisational power, the firm turned AI adoption into lasting business value.

Next steps (1/3)

Deploy Proven Use Cases to Deliver Early Wins

Deploy proven use cases to deliver early wins

Start with AI initiatives that can show fast, measurable value and then shift from small pilots to projects tied directly to cost savings, productivity gains, or revenue uplift, with clear success metrics and accountable owners.

Design each initiative for deployment from day one, ensuring it can quickly move into real operations and deliver visible business impact

Use a focused roadmap to secure returns

Build a disciplined use-case roadmap that begins with well-established, low-risk applications capable of delivering results within months.

Use early wins to build confidence, release funding, and gradually expand into more advanced and strategic AI use cases.
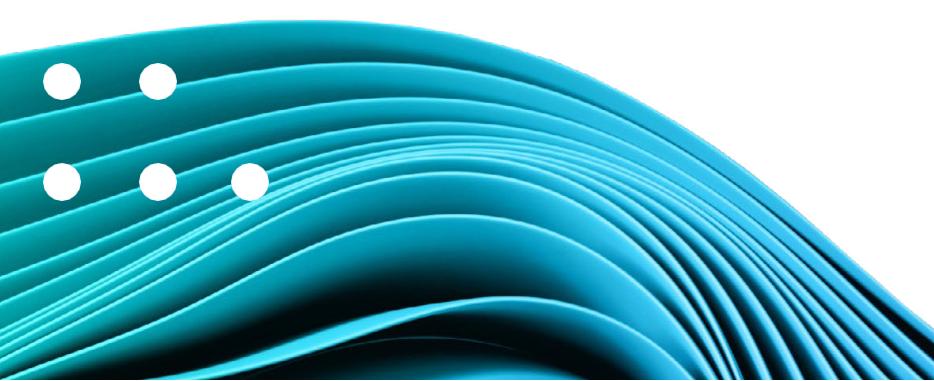
Use Case: GenAI Team Ramp Up for a Multinational Retail Group

BUSINESS CONTEXT

- The Group has defined GenAI strategy, focus on leveraging GenAI to drive operational excellence, and enhance personalization for customer for competitive differentiation

- To drive this GenAI vision, the Group needs to review current technology operating model, in particular, on expanding GenAI capabilities and determining the optimal talent sourcing model to support their strategic direction

KEY ACTIVITIES

- Integrated GenAI capabilities into future technology operating model and capability framework, in line with the Group's GenAI strategy and future vision, and detailed expectations on GenAI competencies across different role levels

- Leveraged GenAI knowledge agent to enrich overall GenAI capability framework / sourcing model

IMPACT

- GenAI Capability Model derived from GenAI strategy and vision

- GenAI Sourcing Model with key considerations across short, medium and long term to build up organisational capability

24

Next steps (2/3)

Make Organisation Transformation and Change Management Top Priority

Set the pace through organisational readiness

AI succeeds when the organisation is ready to adopt it. The pace of progress is set by people, processes, and culture — not the technology. Leading companies treat AI as a business transformation, focus on leadership alignment, build new skills, and establish cross-functional ways of working together.

Break organisational resistance through integrated change and process redesign

Successful adopters overcome resistance by redesigning workflows for AI, defining new roles and decision rights, and supporting employees through structured change and training. They combine process re-engineering, communication, and capability building to help teams confidently use AI in daily work and scale impact.

Use Case: GenAI Change Management for a Leading Chemical Manufacturer

BUSINESS CONTEXT

- As part of the manufacture's digital transformation initiative, GenAI capabilities will be integrated into the new digital operations platform SAP S/4HANA

- The manufacture was looking for a strong change management approach to drive GenAI activation and adoption, and reinforce the digital culture in this transformation

KEY ACTIVITIES

- Developed a comprehensive change management plan to drive GenAI adoption in the organisation, through targeted communication, training and enablement, leveraging digital change champion network etc.

- Leveraged GenAI knowledge agent to develop and enrich communication and training materials

IMPACT

- GenAI Learning and Training Materials to enable upskilling on GenAI capabilities

- Targeted Communications to drive awareness and facilitate adoption of GenAI in the digital platform
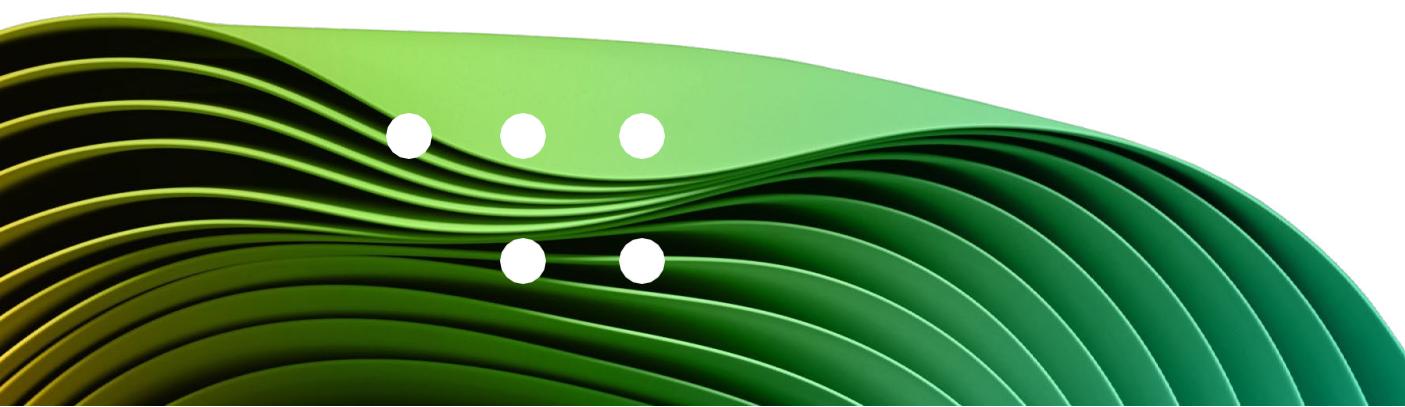
Next steps (3/3)

Shift Mindset from Cost Efficiency to Revenue Generation

Extend focus from cost savings to value creation

AI is increasingly being applied not just to reduce costs or improve operational efficiency, but to generate measurable revenue growth and expand business opportunities. Organisations that take a dual approach (evaluating initiatives for both efficiency and top-line impact) can identify high-value opportunities such as pricing optimization, personalized offers, customer retention, and upselling.

Capture business growth through innovative applications

AI enables the creation of new products, services, and customer experiences. By leveraging data-driven insights, companies can accelerate decision-making, enter new markets, and improve speed-to-market. Balancing efficiency gains with revenue generation allows organisations to capture sustainable competitive advantage in an AI-driven future.

Use Case: Building a Scalable GenAI Foundation

BUSINESS CONTEXT

- A world-leading luxury jewelry and watches brand. Teams across retail, service, and corporate functions face increasing friction in retrieving knowledge, executing tasks, and scaling support. Current tools lack intelligence, consistency, and speed—limiting productivity and undermining the brand's ability to deliver seamless, data-informed experiences.

KEY ACTIVITIES

- Aims to build a scalable and flexible GenAI foundation through two core components: a centralized knowledge base and a modular agent platform.

- The knowledge base is designed to manage and structure the company's unstructured data, enhancing GenAI's ability to parse and understand. The agent platform enables the creation of GenAI powered agents and workflows, with the flexibility to extend across use cases and integrate with business applications.

- Together, they form a future-ready infrastructure that supports rapid deployment, cross-functional utility, and continuous enhancement.

IMPACT

- By establishing a centralized knowledge base and a flexible agent platform, the project enables a strong, scalable GenAI foundation that delivers immediate and long-term value.

- The modular platform architecture allows for rapid PoC development and streamlined deployment of new use cases, significantly shortening the delivery cycle. This accelerates innovation, supports continuous iteration, and ensures GenAI applications evolve efficiently alongside business needs.

Authors

Reference

HKU

Prof. Jin Li
Area Head and Professor
Management and Strategy
Zhang Yonghong Professor in Economics and Strategy

Prof. Jie Gong
Associate Professor
Economics and Management & Strategy

Prof. Matthias Fahn
Associate Professor
Management and Strategy

Deloitte

Pascal Hua
National Managing Partner
Technology & Transformation
Deloitte China

Cheng Zhong
Managing Partner
Technology, Media & Telecommunication Industry
Deloitte China

George Dickinson
Senior Manager
Technology & Transformation
Deloitte China

Jin Li, Feng Zhu, and Pascal Hua, "Overcoming the Organizational Barriers to AI Adoption", Harvard Business Review (2025)

27