KPMG

State of AI in Banking

Artificial intelligence and the banking industry in the Netherlands: research findings

30 April 2026

KPMG. Make the Difference.

Introduction

In 2024 and 2025, we surveyed organizations across the Dutch financial services landscape to assess the state of AI adoption in the industry as a whole. Participants included banks, asset managers, pension funds, and insurers active in the Netherlands.

Building on this broad, cross-sector view, we are taking a deep dive into individual subsectors in 2026. AI is a rapidly progressing and transformative technology, and its impact is unfolding differently across financial verticals. These differences are now sufficiently pronounced to justify focused, sector-specific research, as opposed to a single consolidated study.

Our latest report, *State of AI in Banking*, is therefore part of a broader series of surveys from KPMG the Netherlands, examining the evolution of AI capabilities, maturity, and use cases within financial services. Companion publications include *State of AI in Asset Management* (January 2026) and *State of AI in Insurance* (February 2026), each offering a tailored view of how AI is shaping its respective domain.

About the research

Our research is based on an analysis of structured interviews conducted with nine banks operating in the Netherlands, spanning public institutions, international banks and mortgage specialists, in addition to major Dutch banks. These interviews were carried out to assess the current state of AI adoption, maturity and ambition across the sector. Given the qualitative nature of this research, we have refrained from drawing quantitative conclusions. Instead, our focus has been on extracting directional, qualitative insights from the interviews. These insights have been enriched with our observations from client work, market developments across the Dutch banking landscape and KPMG's global research.

Structure

With this report, our aim is to present the current state of AI in the banking landscape, outline the challenges and dilemmas banks face, and, given that AI is rapidly evolving, provide insights into the way forward.

This report is structured around six research areas:

- AI vision & strategy
- AI use cases
- Responsible AI
- AI platform
- People & AI
- AI roadmap

Contents

Globally, AI is top priority and already reshaping the workforce

KPMG 2025 Banking and Capital Markets CEO Outlook outcomes

AI-powered technological innovation

65% say that GenAI remains a top investment priority.

59% expect agentic AI to have a significant or transformational impact.

88% expect to spend at least 10% of budget on AI in the next 12 months.

69% expect ROI from AI within 1-3 years, up from 13% last year.

56% say that ethical challenges are the biggest obstacle for AI implementations.

AI-ready workforce

83% are prioritizing workforce reskilling to leverage AI.

79% say AI has redefined entry-level skills.

80% say AI is reshaping workforce development.

34% are anticipating workforce reductions due to AI.

AI's potential and importance recognized

The global banking and capital markets industry increasingly views AI as a top priority due to its ability to improve efficiency, decision-making, and customer engagement across the bank. The emerging potential of agentic AI to autonomously execute tasks, orchestrate workflows and operate in more complex environments is broadly recognized as a significant next step in AI-enabled transformation.

AI is budgeted with growing confidence in near-term ROI

CEOs are allocating material budget to AI initiatives. Expectations of near-term ROI are on the rise, showing growing confidence that AI solutions can deliver tangible value in the short run.

AI is making an impact on the workforce

AI is already reshaping the workforce. In response to the AI-driven evolution in role requirements and talent pathways, global banking and capital markets players have started to pursue reskilling and upskilling programs.

State of AI in the Dutch banking market

These findings underline the globally growing importance of AI across the banking and capital markets industries.

We set out to understand how the Dutch banking industry fares in comparison.

This report outlines our findings about where banks operating in the Netherlands currently stand and where they are heading when it comes to the transformative potential of AI.

Source: Banking and Capital Markets: KPMG 2025 Global CEO Outlook, surveying 110 Banking and Capital Markets CEOs

Key takeaways per research area

AI vision & strategy

- All banks recognize AI as strategically important. Larger banks typically have group-level AI strategies closely linked to the broader business strategy, while smaller institutions are still primarily exploring and experimenting with AI.
- Current AI strategies are primarily efficiency-driven, with cost reduction as the main objective.
- AI will reshape the competitive landscape, and banks have valuable assets in data, customer proximity and trust that position them well for a competitive edge.

AI use cases

- AI use case maturity varies by bank size, with larger banks being further advanced and increasingly evaluating AI initiatives as ROI-driven investments rather than simply aiming to keep up with industry adoption.
- Banks generally favor off-the-shelf vendor solutions.
- The current strategic focus remains on efficiency and hyperpersonalization is widely seen as the next step.
- AI use cases remain concentrated in efficiency-led, retail-focused domains, primarily KYC/AML, customer service, software development, and operational processes. GenAI dominates development.
- Agentic AI applications can help banks move beyond rule-based automation and unlock a new level of potential from AI.

Responsible AI

- Responsible AI is well-established at banks, with AI guardrails widely in place and typically built by extending existing privacy, security, and risk frameworks rather than introducing new AI-specific governance models.
- Leading banks involve risk and compliance early, improving the path from experimentation to deployment. However, Responsible AI practices still primarily emphasize control and compliance over actively enabling AI-driven value creation.
- As AI scales and becomes more autonomous, existing governance frameworks may struggle to catch up. Scalable trust will require stronger governance tooling, rethinking the Three Lines of Defense model, and embedding ethical controls by design rather than relying on maturity, principles, or periodic reviews alone.

AI platform

- Many banks require significant modernization of their technology infrastructure before AI can be deployed reliably at scale, with legacy systems, fragmentation, and limited standardization hindering operationalization of AI use cases.
- Leveraging out-of-the-box AI capabilities from established vendors is a key accelerator.
- Data quality remains a persistent constraint and requires investment in foundational capabilities. However, banks cannot afford to wait for perfect data and must adopt pragmatic, risk-aligned approaches to progress while improving data over time.

People & AI

- Uneven understanding of AI use cases and shortages in AI-related skills slow progress, as employees often need foundational upskilling before AI can be applied effectively in their work.
- Cultural readiness across banks ranges from groups that are resistant to change to teams showing strong bottom-up enthusiasm for AI.
- Successfully moving to an AI-driven way of working requires both a transformation of the workforce and a clear adoption strategy from banks.

AI roadmap

- In the short term, banks are moving from AI pilots toward production, with a focus on scaling use cases across the organization. Data, people and technology are the three pillars for a firm-wide AI rollout.
- Banks follow different adoption strategies, but most of them share an ambition to accelerate AI implementation.
- In the longer term, banks recognize the need to redesign processes around an AI-first operating model, including multi-agent solutions.

AI vision & strategy

Why AI vision matters

A clear vision and strategy for AI enable banks to anchor AI initiatives in business objectives and financial priorities. Such direction ensures that use cases are deployed with purpose, and that the organization can mobilize and prepare its workforce for AI adoption at scale.

Scope

This section provides an overview of how banks are shaping their AI vision and strategy. It reflects their stance on AI adoption, the degree of strategic alignment across the organization, and the key challenges encountered as they advance in their AI journey.

Takeaways

- All banks recognize AI as strategically important. Larger banks typically have group-level AI strategies closely linked to the broader business strategy, while smaller institutions are still primarily exploring and experimenting with AI.
- Current AI strategies are primarily efficiency-driven, with cost reduction as the main objective.
- AI will reshape the competitive landscape, and banks have valuable assets in data, customer proximity and trust that position them well for a competitive edge.

Despite broad recognition of AI's importance, the strategic ambition diverges sharply between large banks and their smaller counterparts

The strategic relevance of AI is widely acknowledged among banks

Across all institutions interviewed, the strategic relevance of AI is widely acknowledged. Banks expect AI to play a substantial role in strengthening customer engagement and improving operational efficiency.

Despite this shared recognition, banks differ markedly in how this awareness is currently translated into concrete strategy and execution. We observe that the main driver of the difference is the relative size of the institution.

Smaller banks take a more exploratory approach

Smaller banks tend to adopt a measured, wait-and-see stance toward AI. Rather than setting a comprehensive long-term AI strategy, they prioritize experimentation with a limited set of targeted use cases and seek to build out their strategic approach gradually.

Consequently, AI initiatives in smaller institutions often emerge bottom-up, led by individual teams implementing use cases to address specific operational challenges. This pragmatic approach reflects both organizational and strategic considerations: many smaller banks operate with more specialized business models and serve distinct customer segments, lowering the immediate need for a broad, bank-wide AI strategy. Furthermore, some of the smaller banks have less room to allocate budget and resources to AI initiatives. A more pragmatic, bottom-up approach comes at the cost of scalability. As the smaller banks rely less on well-established IT platforms for their AI developments, a firm-wide rollout is technically more challenging.

Larger banks are more mature and ambitious, but face more organizational challenges

Larger banks have established more ambitious, enterprise-level AI strategies. At board level, they define priority domains where AI is expected to generate value, supported by explicit financial targets and more budget for technology investments and recruiting the necessary IT talent. Internationally, banks are appointing Chief AI Officers at board level, yet Dutch banks lag in this respect, despite the strong central governance frameworks in place at the larger banks.

However, this top-down approach introduces its own challenges. While management articulates clear strategic ambitions, not all employees fully align with or embrace these initiatives. Resistance to change, uncertainty about new ways of working and concerns about lay-offs can lead to friction between the AI strategy of the board and employees. For this reason, successful implementation of the AI strategy in larger banks depends on management's ability to drive broad workforce adoption of the AI strategy and its implications.

Leadership has a good understanding of AI and its strategic implications

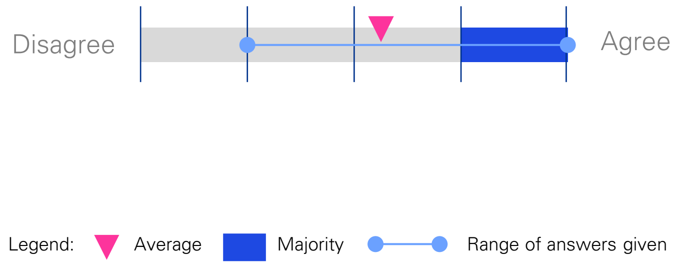

Efficiency is the core objective of current AI strategies

Primary goal of current AI efforts is increased efficiency
For most banks surveyed, the primary driver behind AI adoption is its potential to increase operational efficiency and reduce costs. Furthermore, they also acknowledge that AI can help to improve the quality of their product offering and customer interaction. Both efficiency and quality are seen as necessary to remain competitive in the current landscape. However, efficiency gains are mostly seen in back and middle office functions of the bank, while productivity and quality objectives are more relevant for front-office roles like business development and relationship managers.

This focus of Dutch players is consistent with findings from KPMG's 2025 Global Intelligent Banking report, which highlights that efficiency improvement remains the dominant strategic rationale for AI investment across leading financial institutions worldwide. Concerns about losing a competitive edge to those embracing AI is also a highly pronounced concern among global banking players.

In our view, AI has a clear potential to deliver cost savings and efficiency improvement, but these gains are more visible at a macro level than in individual use cases. In the midterm future, banks that successfully implement AI are likely to benefit from lower cost bases and higher productivity. At the micro level, however, cost savings are difficult to pinpoint, as AI rarely replaces isolated tasks. Realizing these benefits requires a fundamental rethink of the operating model and ways of working. As a result, AI-driven cost savings are harder to achieve in practice than they may initially appear.

Safety considerations induce reluctance to deploy customer-facing AI
Banks remain acutely aware of the need to safeguard operational integrity and emphasize the need to finely balance innovation with operational safety. Operating within a highly regulated and trust-driven environment leaves no margin for error. AI-related failures such as data breaches, biased models, or flawed decision-making can have significant reputational and regulatory consequences.

This risk aversion is especially evident in customer-facing applications. Banks remain highly cautious when implementing customer-facing AI solutions. Most of the AI use cases at this point in time are deliberately designed to operate in the background and support internal processes. In KPMG's 2025 Global Intelligent Banking report, 42% of the banks aim to enhance customer experience with AI. In this regard, banks will need to do so securely when they start implementing AI in customer-facing solutions.

Global banking & capital markets reference
KPMG 2025 Global Intelligent Banking report

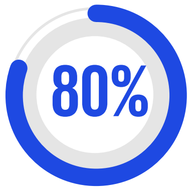

believe that banks that embrace AI will develop a competitive edge over those who do not.

are seeking to reduce costs.

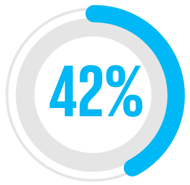

are seeking to enhance the customer experience.

Banks are uniquely positioned to capture AI value

Banks struggle to move from experimentation to at-scale adoption

Despite recognizing AI's strategic importance, many banks remain in an experimental phase due to concerns around regulatory, operational and reputation risk. AI initiatives are often limited to proofs of concept or isolated use cases. One notable exception is KYC and AML, where larger banks are already deploying AI at scale and even reporting workforce impacts driven by AI transformation. However, a recent opinion piece of KPMG colleagues tempers expectations of short-term gains in this domain.

The KYC and AML domain is a logical field in which AI solutions can be rolled out, as there are many repetitive tasks that depend on a large amount of manual effort. AI can support these processes and deliver tangible results. In other domains, such as finance and HR, the underlying technology platforms limit the rollout of AI solutions. Large upgrades and investments in the platform will be required to obtain AI functionality.

AI will reshape competitive landscapes, but banks have valuable assets that, if leveraged adequately, position them well for a competitive edge

As AI adoption accelerates across the financial sector, 'traditional' banks will need to reassess their position relative to fintech's and digital-first challengers. These newer, often digital-native, competitors benefit from technological agility, in part thanks to a lack of legacy systems and platforms traditional banks often need to contend with.

However, incumbent banks also possess structural advantages that are difficult to replicate: large, rich datasets, extensive customer relationships, and long-established trust with customers as well as regulators. A key objective for banks is therefore to leverage these advantages with a clear AI strategy and mature guidelines on responsible AI.

These assets position banks uniquely well to deploy AI at scale and unlock differentiated value. With deeper customer insight, broader operational footprints and more robust risk management capabilities, banks have the potential to translate AI into significant competitive advantage going forward. Particularly, banks should focus on the following:

1. **Shift towards an AI-first mindset**
   Moving beyond viewing AI as efficiency tools, banks should adopt an AI-first mindset and define AI's role in their long-term strategy. This includes clarifying how AI initiatives should be developed and how stakeholders should be involved. This also includes the bank's ambition with AI and its efficiency targets.

2. **Fine-tune their risk appetite for AI**
   Banks' reluctance to scale AI in all domains is partially caused by risk aversion. Banks should articulate their risk appetite for AI, including acceptable levels of model risk, explainability, and regulatory uncertainty.

3. **Move beyond AI use cases and towards partnerships**
   A clear and broad view on AI use cases allows banks to put their AI strategy into practice and focus on the most relevant domains. However, as outlined in the next section on AI use cases, partnerships with technology vendors are an essential next step to enable scalable AI adoption. Vendors are able to rapidly deploy customer-tested, platform-integrated AI solutions at a pace that banks cannot match through in-house development. Banks must therefore rethink their IT architecture and invest to capture the benefits of AI.

AI use cases

Why AI use cases matter

A clear view of AI use cases, and a structured approach to identifying and assessing them, enables banks to focus on initiatives that deliver tangible business value.

Scope

This section outlines the AI use cases currently identified by banks, along with the key considerations guiding their development.

Takeaways

- AI use case maturity varies by bank size, with larger banks being further advanced and increasingly evaluating AI initiatives as ROI-driven investments rather than simply aiming to keep up with industry adoption.
- Banks generally favor off-the-shelf vendor solutions over bespoke development.
- The current strategic focus remains on efficiency and hyperpersonalization is widely seen as the next step.
- AI use cases remain concentrated in efficiency-led, retail-focused domains, primarily KYC/AML, customer service, software development, and operational processes. GenAI dominates development.
- Agentic AI applications can help banks move beyond rule-based automation and unlock a new level of potential from AI.

Larger banks are further ahead in AI adoption and strategy than smaller peers, but all are actively leveraging vendor solutions and looking to move beyond efficiency use cases

Larger banks are at more advanced stages of AI use case development than smaller banks

Banks in the Dutch market are at varying stages of AI adoption. A clear distinction emerges between larger and smaller institutions: larger banks are generally further ahead in both identifying and implementing AI use cases and have made more progress toward scaling them. Smaller banks, by contrast, remain primarily in the experimentation phase.

Larger banks are also making the shift from seeing AI investments as a strategic necessity to ROI drivers in their own right

Among larger banks, the strategic framing of AI investments is evolving. Early use cases were often justified as a necessary investment to keep pace with technological developments and competitive pressure. Some banks also invested in order to start building capabilities and be prepared for further AI acceleration. Recently, larger institutions have begun assessing AI use cases through explicit return-on-investment (ROI) criteria. They are setting clear expectations for cost savings or revenue impact and view AI increasingly as a source of competitive advantage. This shift is consistent with KPMG's global view, where banks this year are prioritizing short-term ROI and scalable solutions, moving away from last year's focus on exploring a larger volume of use cases.

Most banks prefer to buy solutions off-the-shelf if available

Banks continue to make strategic choices between developing AI solutions internally and adopting those offered by technology vendors and outsourcing partners.

Several banks shared that they initially started out building their bespoke AI use cases, only to find out a few months later that the same use cases became native functionalities in vendor products. Developing AI capabilites also involves costly upgrades on an ongoing basis, which can be more quickly and reliably delivered by vendors that operate at scale. As such, the prevailing approach is to experiment broadly, and only to consider significant investment into in-house development when no suitable vendor solution is expected to reach production in the near term. To manage this, banks are increasingly more engaged with their suppliers about which AI capabilities are in the pipeline and what they would like to see being developed.

Hyperpersonalization is being seen as the next step

While current AI efforts focus mainly on efficiency, several banks are already looking further ahead. They see AI as a potential catalyst for reshaping customer engagement and product design. Hyperpersonalization was frequently cited as a next step in which AI enables tailored products, advice, and interactions for individual customers.

We have a clear view on which AI use cases deliver business value

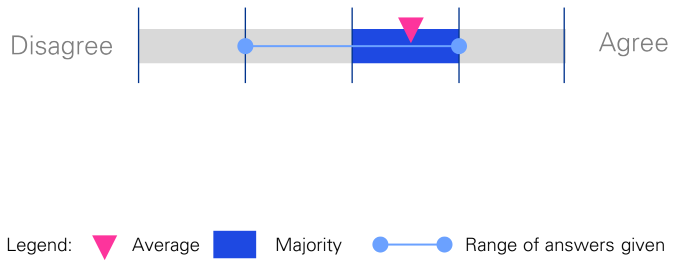

Dutch banks prioritize retail banking and GenAI use cases

Use cases are concentrated in retail banking

Most of the customer-oriented use cases reported during our interviews were from the retail domain. Banks indicate that the product and delivery models of retail banking make it the logical starting point for customer-facing use cases. Data volumes in retail are larger and customer interactions in this area are more homogeneous and standardizable, providing more opportunities to train AI agents better.

GenAI dominates use case development, but machine learning is still relevant

Traditional machine learning focuses on prediction and classification based on structured data, while generative AI enables the creation of new content and interactions, often using unstructured data and natural language. Most AI use cases today are built on GenAI models, reflecting recent advances and the opportunities these models create. At the same time, banks emphasize that traditional machine learning models remain highly relevant, particularly for use cases involving structured or tabular data, where such models can be more effective and efficient, such as in credit risk models.

Use cases in practice today generally fall into one of four categories

Most of the current AI use cases mentioned by the banks fall into one of four groups:

01

Know Your Customer (KYC) and Anti-Money Laundering (AML)

Given its effectiveness in dealing with large, diverse data, AI is highly relevant for KYC and AML domains. Currently, AI performs well in assigning customer risk profiles and automating routine checks. AI is also effectively leading to the (earlier) identification of unusual transactions.

02

Customer service

Virtually all banks now have AI to create summaries for human call center agents, and for internal tools such as knowledge banks. Some banks also have live AI chatbots directly interacting with clients.

03

Software development

Banks are using AI for programming and software development, often leveraging general purpose AI such as Copilot.

04

Operational processes

In addition to the three categories mentioned above, banks are increasingly looking at use cases in other operational processes, such as credit risk assessment for loans, primarily with an ambition to drive lower C/I ratios.

The efficiency focus of the Dutch banking industry is in line with other financial services; personalization and intelligent work emerge as next opportunities

AI use case development in the Dutch banking industry largely mirrors the efficiency focus of the Dutch insurance and asset management industries

The Dutch banking industry largely follows the same AI adoption pattern as the insurance and asset management industries, with efficiency-driven use cases acting as the primary entry point. In banking, AI is first applied to high-volume, repetitive processes where productivity gains are tangible and risk is manageable, such as core banking operations, customer servicing, compliance activities, reporting, and data aggregation. This mirrors insurance, in which claims handling and administrative processes dominate early AI efforts, and asset management, in which compliance and reporting benefit from automation. In all three sectors, AI is initially positioned to perform existing work faster and at a lower cost.

In addition to this shared starting point of the Dutch financial services industry, there are also similarities among large banks and large insurers in how AI value is assessed and governed. Both place increasing emphasis on explicit ROI criteria and business value measurement, particularly as organizations seek to scale AI beyond isolated use cases. While challenges remain in translating early benefits into durable, enterprise-wide impact, value articulation by large banking and insurance players is becoming more structured, with clearer links to productivity, cost reduction, and operational performance.

Banks currently focus on cost reduction use cases. Looking ahead, we see **two immediate opportunities** in the following use case domains:

Personalized customer engagement and product advisory

AI enables more granular client and segment profiling, personalized offers in retail banking, and tailored deal structuring in corporate banking, supported by the automated routing of next-best offers to the most relevant sales and service channels.

Particularly in retail banking, beyond efficiency and conversion uplift, these use cases allow banks to reintroduce a more personal, human dimension to customer interactions. This can be a strategic opportunity to restore the personal touch in a bank’s branding at a time when physical branch presence has largely disappeared.

Intelligent work with agentic AI

Banks can move beyond rule-based automation by deploying intelligent AI agents that work alongside employees to orchestrate end-to-end workflows, apply dynamic reasoning, and handle complex exceptions. In practice, this could include finance agents that support bookkeeping processes, resolve data quality issues, and assist with reconciliations, as well as HR agents that can query internal HR policies, validate expense declarations, and support case handling. In product and process development, agentic AI can act as a sounding board for employees by continuously checking compliance with regulatory and internal standards. In this fashion, AI can augment day-to-day work, shifting from isolated task automation toward integrated, intelligent ways of working.

Responsible AI

Why responsible AI matters

Responsible AI enables banks to safeguard security, privacy, fairness and ethical use, as AI becomes embedded in core processes. With European AI regulation still evolving, banks continue to rely primarily on existing governance frameworks to manage AI-related risks.

Scope

This section outlines how banks currently govern for responsible AI and its impact on the development, testing and deployment of AI.

Takeaways

- Responsible AI is well-established at banks, with AI guardrails widely in place and typically built by extending existing privacy, security, and risk frameworks rather than introducing new AI-specific governance models.
- Leading banks involve risk and compliance early, improving the path from experimentation to deployment. However, Responsible AI practices still primarily emphasize control and compliance over actively enabling AI-driven value creation.
- As AI scales and becomes more autonomous, existing governance frameworks may struggle to catch up. Scalable trust will require stronger governance tooling, rethinking the Three Lines of Defense model, and embedding ethical controls by design rather than relying on maturity, principles, or periodic reviews alone.

Responsible AI governance is maturing, but not yet enabling informed risk-taking and value creation

Banks established strong AI guardrails by extending existing risk frameworks

Banks report that robust controls are in place across security, privacy, and ethics to guide AI implementation. These guardrails have been established by extending existing risk frameworks rather than introducing entirely new AI-specific ones. Common approaches include centralizing controls on AI platforms and reviewing AI use cases in multidisciplinary risk forums. Maturity per bank still varies, with some still defining what responsible AI means in their context, operating within conservative governance structures and legacy systems.

Leading banks incorporate the second line of defense early in the AI lifecycle

The survey also shows that more advanced banks have started involving risk and compliance functions early in the AI lifecycle instead of waiting for fully developed solutions for their review. This early involvement is intended to reduce friction later in development and is widely seen as a positive practice.

Current risk management practices may be overly cautious, holding banks back from translating their structural advantages to lasting value

Overall, we see that the current AI risk management approach at Dutch banks focus primarily on avoiding incidents, rather than enabling informed risk-taking and value creation with AI. While current guardrails are effective at preventing harm, they do not on their own help banks optimize risk appetite or embed risk management as a true business partner in AI innovation. Managing AI risk purely through guardrails can therefore lead to an overly conservative posture, which by itself becomes a strategic risk if it slows AI adoption and limits competitive progress.

Earlier in this report, we mentioned that banks benefit from unique structural advantages such as large, rich datasets, extensive customer relationships, and long-established trust with customers and regulators. An overly cautious risk appetite may be their Achilles heel. While existing guardrails are effective at preventing harm, they do not sufficiently help banks to translate their structural advantages into scalable AI-driven outcomes. Early second-line involvement and strong guardrails represent important progress, but must be complemented by evolving risk frameworks, skills, and toolkits that are effectively embedded across the organization to ensure responsible AI would support both control and value creation.

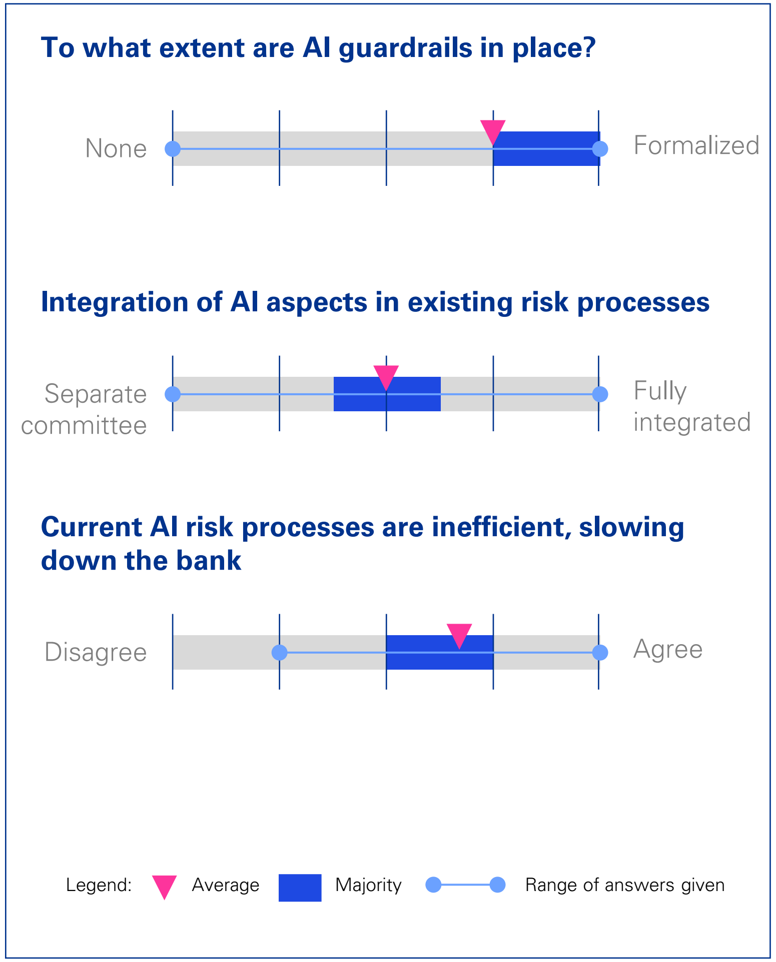

AI's growing autonomy requires that banks rethink governance to enable trust at scale

Why governance maturity is no longer enough

Many banks have reached a high level of maturity in AI risk and governance frameworks. However, as AI adoption accelerates, these frameworks increasingly struggle to scale. Fragmented processes, manual controls, and siloed oversight make it difficult to keep pace with faster model iteration, broader deployment, and rising regulatory and reputational scrutiny.

What worked for a limited set of AI use cases does not necessarily work at enterprise scale. Banks indicate that existing AI risk processes can slow down AI adoption rather than enable it. These limitations become even more pronounced as AI systems evolve from stand-alone models toward more dynamic and agent-based solutions.

Tooling as the enabler of scalable governance

To move from mature governance to scalable trust, banks must complement frameworks with robust AI governance tooling. Agent-based AI solutions fundamentally change processes by autonomously executing tasks and orchestrating workflows, but doing so requires significantly stronger oversight. Their value depends on the tools, permissions, and integrations they are given, while those same capabilities introduce new vulnerabilities. Effective tooling enables continuous monitoring, logging, guardrails, and control of autonomous behavior. This shift is structural and enduring: AI is not static, and as development, adoption, and autonomy increase, governance must be designed to scale alongside it.

AI may become a turning point for the Three Lines of Defense model

As AI-enabled use cases scale and decision-making becomes faster and more complex, internal audit functions are being asked to assess increasingly sophisticated risks that cut across technology, data, and business judgment. In many organizations, risk capabilities and experience are therefore still maturing alongside rapid innovation. Risk frameworks and tooling also continue to evolve, but not always at the same pace as AI adoption, which can at times lead to friction or inefficiencies across lines. AI may therefore represent a turning point for the Three Lines of Defense model, forcing banks to rethink whether risk is best managed through stronger, more embedded controls in the first line or through an expanded role for the second and third line.

Ethical AI: trust by design, not by exception

As AI systems gain autonomy and are delegated tasks, decisions, and tools, ethical risk shifts from intent to execution. In agent-based systems, errors, bias, or unintended behavior are no longer limited to single model outputs but can propagate across processes, systems, and customers. Ethical AI therefore cannot rely on principles or periodic reviews alone. It must be operationalized through tooling that enforces boundaries on what agents may do, records why actions were taken, and enables intervention when behavior deviates from expectations. In an agentic environment, trust is not created by good intentions, but by enforceable limits, traceability, and continuous control at scale.

AI governance principles from KPMG and INSEAD

KPMG set out five principles to govern AI responsibly:

- Strategic oversight for long-term value creation: aligning AI to the organization's mission, creating sustainable value across key business models, investments and ecosystem partnerships.
- Active technology and security oversight: actively overseeing strategic AI technology choices including speed, cost, sovereignty, privacy, safety and security.
- Workforce transformation and human accountability: ensuring human judgment remains central as AI reshapes roles, skills, and culture.
- Building trustworthy AI: ensuring AI reflects organizational values and regulatory obligations, including fairness, transparency, and accountability in ethical deployment.
- The work of the board of directors: evolving board structures, processes and capabilities to remain effective in the AI era.

Source: KPMG and INSEAD launch global AI Governance Principles for Boards

AI platform

Why a strong data and technology foundation matters

Scaling AI effectively requires a strong data and technology foundation. Centralized AI platforms, consistent data standards, and strong governance enable secure, reusable and scalable AI development across domains. For banks, establishing this foundation is particularly challenging due to legacy core systems, siloed data ownership across products/geographies, data fragmentation, the persistence of inconsistent quality alongside the enormity of unstructured data volumes.

Scope

This section examines how technology infrastructure and data foundations together shape banks' ability to scale AI.

Takeaways

- Many banks require significant modernization of their technology infrastructure before AI can be deployed reliably at scale, with legacy systems, fragmentation, and limited standardization hindering operationalization of AI use cases.
- Leveraging out-of-the-box AI capabilities from established vendors is a key accelerator. Although upgrading to modern, cloud-based platforms requires significant investments, it can unlock a substantial share of AI use cases and significantly reduce manual effort.
- Data quality remains a persistent constraint and requires investment in foundational capabilities (such as scanning, labeling, vectorization capabilities for providing unstructered 'AI-ready' data). However, banks cannot afford to wait for perfect data and must adopt pragmatic, risk-aligned approaches to progress while improving data over time. Unstructured data in particular require new quality and data quality remediation approaches.

Scaling AI requires decisive investments in modern technology and new ways of working with unstructured data

Modernization of the current technology infrastructure is essential for scaling AI

Most banks report that their technological infrastructure requires significant modernization before AI can be implemented and scaled reliably. Combined with legacy systems, fragmented technology landscapes, and limited standardization, banks currently find it challenging to translate successful AI prototypes into operational processes at scale. While banks are making progress in building more centralized technology environments and in making unstructured data available, the pace of development varies widely, and dependencies on global IT functions or external vendors often slow down implementation and rollout.

Leveraging out-of-the-box AI capabilities of vendors can be the main accelerator for scalable AI

From our perspective, the main accelerator for scalable AI lies in leveraging out-of-the-box AI capabilities from established technology vendors. Banks that upgrade to the latest, cloud-based versions of core systems can increasingly activate prebuilt AI agents with relatively limited implementation effort. Although such platform modernization typically requires substantial upfront investments and often involves full reimplementations rather than simple upgrades, these initiatives can be underpinned by a clear business case. Vendor-delivered AI functionality often covers a vast majority of AI use cases and hence can significantly reduce manual effort in core processes, automate routine tasks at scale, and lower ongoing development and maintenance costs.

Data quality is considered a major bottleneck for AI scaling in Dutch banks

Unstructured data is repeatedly cited as a bottleneck, also by banks who are more advanced in their AI adoption journey. All banks note that improvements are required before more advanced AI techniques can be applied effectively. In addition, data is often scattered across legacy systems or different business units, making it difficult to develop firm-wide AI solutions and reuse data across the organization.
Progress in addressing the issue is varied across institutions. Some have made progress in structured data management but still face data gaps that limit GenAI effectiveness. Others remain at an early stage of data transformation and have not yet invested sufficiently in the foundational capabilities required.

Banks cannot afford to wait for perfect data

Data quality and fragmentation should not be treated as prerequisites that must be fully resolved before AI adoption can begin. Banks cannot afford to wait for perfect data; instead, they need clear policies that define scope, tolerance levels, and acceptable degrees of imperfection on a use-case-by-use-case basis, aligned with their risk appetite and responsible AI frameworks. Responsible AI and data governance are therefore closely intertwined, determining where and how AI can be safely applied despite data limitations.
Moreover, AI itself can play a role in addressing these challenges: by helping in structuring data, identifying gaps or inconsistencies, proposing mappings, and supporting human validation, AI can turn data quality from a blocking issue into a gradual improvement opportunity over time.

Technology and platform challenges for a firm-wide AI rollout

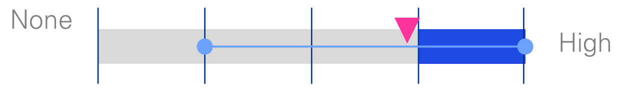

Data challenges for a firm-wide AI rollout

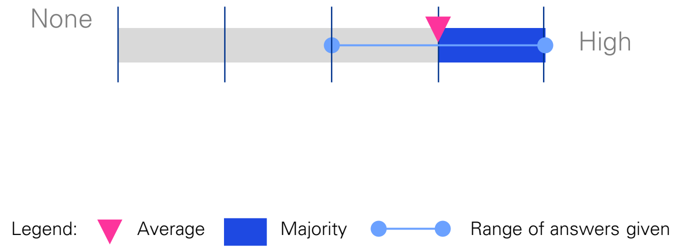

People & AI

Why people and AI matters

People play a decisive role in the success of AI initiatives. Management buy-in, as well as levels of enthusiasm and understanding of the AI agenda across the organization, are critical success factors. As a result, adoption strategy should be an integral component of any organization-wide AI strategy.

Scope

This section examines how employees within Dutch banks experience the organizational and workforce changes introduced by AI, and how these experiences influence the pace and success of AI transformation.

Takeaways

- Uneven understanding of AI use cases and shortages in AI-related skills slow progress, as employees often need foundational upskilling before AI can be applied effectively in their work.
- Cultural readiness across banks ranges from groups that are resistant to change to teams showing strong bottom-up enthusiasm for AI.
- Successfully moving to an AI-driven way of working requires both a transformation of the workforce and a clear adoption strategy from banks.

People and culture play a critical role in AI velocity

AI literacy and skills are uneven and often insufficient

Limited AI literacy and skill gaps remain a recurring challenge. Several organizations explicitly identify shortages in AI-related skills as their largest barrier, requiring employees to first build a foundational understanding of AI as a technology and its capabilities.

Even among employees who have a foundational grasp of AI, adoption can be slow as employees lack a clear understanding of AI use cases or how AI applies to their day-to-day work.

Organizational attitudes around change and its governance can also slow progress

Cultural readiness varies widely across banks. On the one hand, some organizations struggle with a conservative stance among their employees toward innovation. Reluctance to adopt new skills or rethink existing ways of working can slow transformation progress.

On the other hand, some banks need to temper the AI expectations among employees, as groups show enthusiasm to experiment with AI and suggest use cases. This eagerness shows a bottom-up desire to innovate and an interest to use AI to enhance their tasks. While banks certainly do not want to discourage AI experimentation, they must ensure it happens within clear governance and guardrails, as outlined in our section on Responsible AI. Boundaries that are perceived as restrictive or unclear by enthusiastic employees can lead to frustration and fatigue among teams who have been initially excited to leverage the technology. Furthermore, banks should not solely rely on the eagerness of early AI adopters within their workforce for AI innovation. Scalable solutions will require a firm-wide AI adoption stance.

AI-driven change raises concerns about roles and job security

AI introduces significant change for employees, creating uncertainty about roles, ways of working, and the future of work. Concerns about job loss, increased pressure from rapid technological change, and the need to fundamentally redesign processes contribute to hesitation. Where leadership communication and direction are lacking, these concerns can intensify and further hinder adoption.

In the KYC and AML domain, AI is explicitly mentioned by banks as a reason for FTE reductions, showing AI efficiency gains. However, in practice it is hard to pinpoint which FTEs are directly impacted by the layoffs, as many roles are supported by AI rather than fully replaced. We therefore believe the current impact on the workforce is limited and the layoffs are a clear indication that banks have caught up with earlier negligence and shortcomings in these checks.

The people aspect is a challenge in our AI adoption

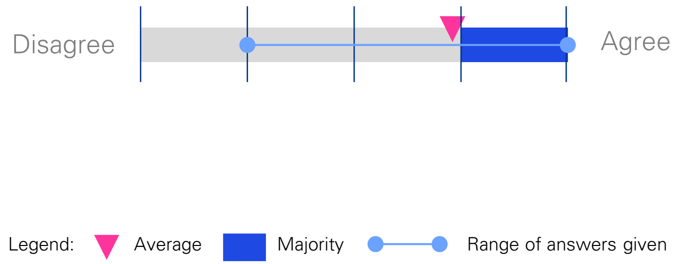

To fully realize the potential of AI, banks must go beyond AI adoption and rethink how their workforce operates

AI workforce

Utilizing the value of AI requires a new approach that is built on the identification and improvements of those tasks that are difficult to automate. A goal-oriented approach and clear success criteria are needed to effectively transform and improve an organization's workforce. KPMG has designed an approach consisting of three important steps:

- **Identify opportunities and make a value assessment**
  Map the departments, roles, tasks, activities, and skills that are impacted by AI and where the greatest value can be achieved through AI adoption. Prioritize the opportunities and assess the available technologies. Ensure they are value-driven and future-focused.

- **Augment the workforce**
  Look at new roles and capabilities, support employees in their development, and explore how AI adoption can be promoted. At the same time, it is important to make these improvements visible. In doing so, we not only look at metrics, but also at the value added to service delivery, customer experience, and so on.

- **Reshape the workforce**
  To fully realize the potential of AI, organizations must go beyond adoption and rethink how their workforce operates. This phase focuses on redesigning roles, optimizing talent, and aligning the entire organization with a future-ready operating model.

AI adoption strategy for a smooth transition

Given the disruptive power of AI on the workforce, a smooth adoption is needed. Four key focus points are:

- Build understanding and trust to enable AI adoption.
- Make AI leadership-led, not just employee-driven.
- Redesign work for teams combining people and AI agents.
- Adopt AI through a coordinated, organization-wide approach.

KPMG's AI workforce approach

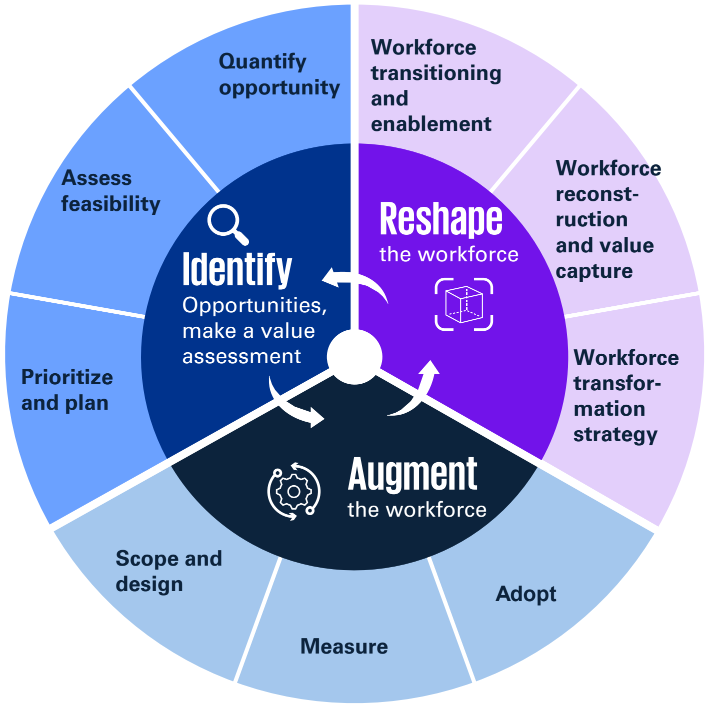

Source: KPMG AI Workforce

AI roadmap

Why an AI roadmap matters

Banks need a clear AI roadmap to steer both short- and long-term AI ambitions, grounded in a well-defined AI vision and strategy. An effective roadmap aligns people, processes and technology towards a shared future state. However, given the rapid pace of AI developments, such roadmaps must remain agile and adaptable to evolving technologies, regulatory expectations and emerging opportunities.

Scope

This section outlines banks' short- and longer-term plans for AI adoption. It highlights the key challenges they face, the opportunities ahead, and the strategic focus areas shaping their AI journey.

Takeaways

- In the short term, banks are moving from AI pilots toward production, with a focus on scaling use cases across the organization. Data, people and technology are the three pillars for a firm-wide AI rollout.
- Banks follow different adoption strategies, but most of them share an ambition to accelerate AI implementation.
- In the longer term, banks recognize the need to redesign processes around an AI-first operating model, including multi-agent solutions.

In the short term, Dutch banks are shifting from AI experimentation toward organization-wide production and scale

Banks are shifting from AI experimentation toward production and scale

Dutch banks are increasingly moving beyond isolated AI pilots toward bringing use cases into production and scaling them across the organization. Short-term AI roadmaps are therefore becoming more rollout-oriented, with greater emphasis on embedding AI into core processes, thus clarifying ownership, and enabling firm-wide deployment rather than continuing stand-alone experimentation.

Banks pursue different adoption strategies, yet share an ambition to accelerate

Banks have made deliberate choices about their position on the AI adoption curve. Some aim to act as front-runners, accepting higher uncertainty in exchange for early-mover advantages. Others follow a more cautious path, prioritizing regulatory certainty, stability and proven solutions, often also influenced by resource and capability constraints. While these approaches result in different pacing and sequencing, most banks indicate a shared ambition to accelerate AI adoption in the near term.

Short-term AI roadmaps focus on three enablers for organization-wide rollout

In the short term, banks consistently highlight three focus areas as critical to scaling AI:

- Data: improving data quality, reducing fragmentation, and addressing missing metadata to ensure AI outputs are reliable and actionable.
- People: addressing shortages in technical AI skills, strengthening AI literacy across the workforce, and proactively managing workforce impact to support adoption and avoid resistance.
- Technology: investing in centralized platforms that support secure, governed and scalable development and deployment of AI use cases. Furthermore, banks should leverage the AI solutions of platform providers. Prioritizing platform upgrades and vendor capabilities allows banks to adopt proven AI solutions quickly and at scale.

Banks need to shift from applying AI to existing processes toward an ‘AI-first’ operating model

A recurring theme from the interviews is that successful banks do not simply layer AI onto existing processes. Instead, they redesign end-to-end journeys with an ‘AI-first’ perspective, structuring processes, decisions and interactions around what AI can enable. While narrowly defined use cases remain a practical starting point, banks increasingly recognize that an AI-first approach is required to capture AI’s broader strategic potential.

An AI-first operating model is emerging as the end state, but most banks are still building foundational maturity

Multi-agent AI solutions will reinforce the need for an AI-first perspective

Within the 'AI-first' operating model that banks will have to implement, multi-agent AI solutions are expected to become more prominent over time. Rather than single models performing isolated tasks, these solutions involve multiple AI agents collaborating across processes and functions. This evolution requires banks to plan for orchestration and interaction between agents, reinforcing the need to move beyond siloed use cases toward AI-first process design.

KPMG's AI transformation framework reveals that the majority of interviewed banks are still at early stages of AI maturity

KPMG's AI transformation framework distinguishes three layers that banks should focus on:

- The Enterprise layer orchestrates transformational change of the whole enterprise, starting with how AI can adjust strategy, business models and key objectives for the enterprise.
- The Functions layer drives AI-enabled transformation across business functions, prioritizing customer-facing value streams and enabling end-to-end processes and workflows which enhance the flow of value.
- The Foundations layer establishes the AI-first technology stack, including infrastructure, platforms and cloud.

The framework puts a bank in one of three phases:

- In the Enable phase, banks focus on establishing foundational capabilities, including leadership, strategy, use case identification, AI literacy, regulatory alignment and guardrails.
- In the Embed phase, AI is integrated into workflows and operating models, supported by platforms and firm-wide data, with a strong emphasis on safety.
- In the Evolve phase, business models transition toward an AI-first approach, enabling firm-wide solutions and sustained value creation.

Based on the interviews, AI front-runners are entering the Embed phase by actively integrating AI into core workflows. While data and platform constraints remain, these banks are willing to make investments to overcome them. Most other banks remain in the Enable phase, focusing on individual AI practices rather than end-to-end transformation.

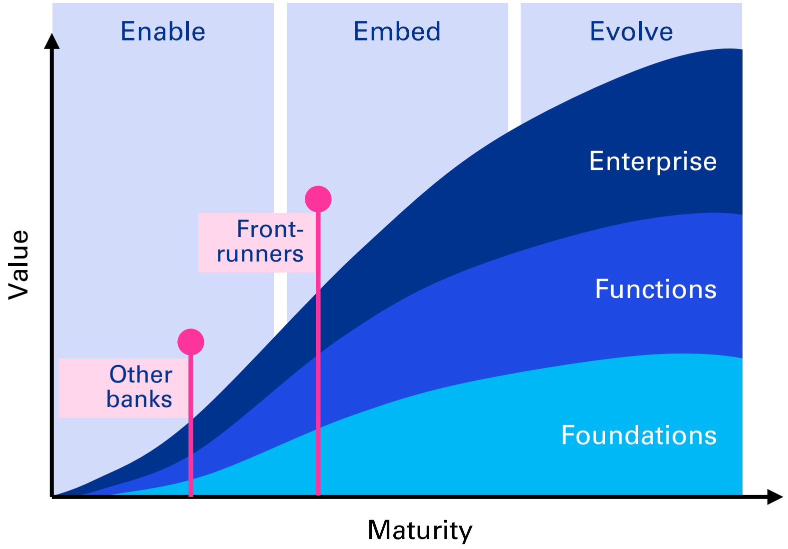

Source: Intelligent banking: A blueprint for creating value through AI-driven transformation

Unlocking AI Value with KPMG

There are four main challenges on the AI journey

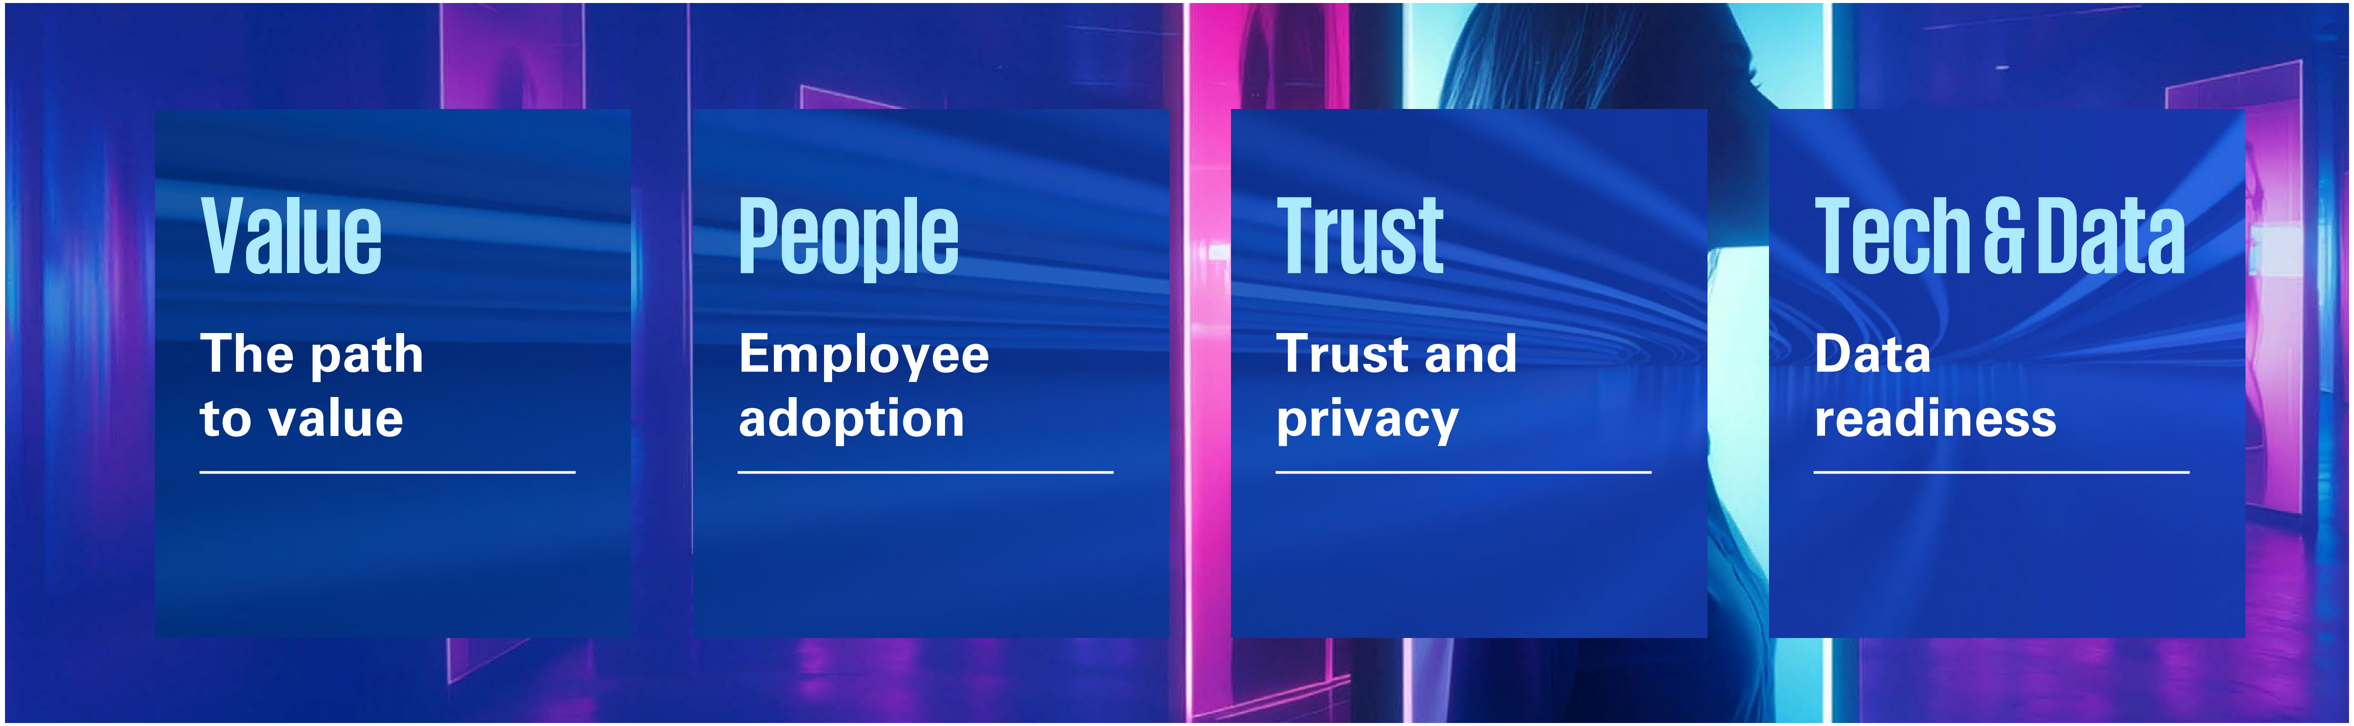

Value
The path to value

People
Employee adoption

Trust
Trust and privacy

Tech & Data
Data readiness

Seeking a path to value?

KPMG research into AI scaling across almost 2,110 senior executives reveals that AI is delivering meaningful business value at 64% of the organizations, yet only 8% report established ROIs.

Navigating your path to value requires addressing at least five strategic questions over time, depending on where you are and how well you are progressing versus peers in your AI transformation.

Value | People | Trust | Tech & Data

Value at stake?

"What is the potential value of adopting AI for my company and where are the highest opportunity areas?"

- Total addressable opportunity in euros.
- Highest value areas within the business?
- Realistic targets by function/ leader initiative?

Current state?

"How strong is my starting position vs. peers and how do I get 'unstuck' and start scaling faster"?

- How are we doing versus our peers?
- How do we get 'unstuck' from pilots to scaling?
- What is my readiness for talent and technology/data?
- Are we ready to leverage both generative and agentic AI?

Change architecture?

"How to design my change architecture, portfolio of initiatives, and governance to successfully transform?"

- How to set up the AI control tower/PMO?
- Value capture by function, value stream, or product?
- Do I need to create an AI Center of Excellence?
- Do I need to set up my own Agent factory or buy/partner?

Strategic choices?

"What strategic choices do I need to make now to attain my future operating and business model and succeed?"

- How to articulate our AI strategy to the Board?
- What is my future agentic operating model and ecosystem partners?
- How do I disrupt others and avoid being disrupted?

Business case?

"How to convince Management and Board to invest with realistic expectations on returns?"

- What is the cost of acquiring and operating all this AI?
- Where and when will we see measurable P&L benefit?
- How much do we plan to invest now and over time?
- What is the possible ROI on our investments?

Source: Global AI Pulse: Q1 2026

Navigating employee adoption?

Executives expect generative AI to have an enormous impact on business but are often unprepared for immediate and ongoing adoption.

Tech is often not the main barrier; human factors and business processes are. AI requires a lot of change to reinvent the organization.

Based on our experience with clients, we have identified four key considerations for successful employee adoption:

Value | People | Trust | Tech & Data

People

A new mindset to employee adoption

GenAI technology at its core is different from previous digital technologies and requires a new way of thinking. For organizations to gain the most value from GenAI tools and ensure sustainability, it is essential to adopt new approaches and mindsets tailored to yield better business results.

Need for continuous upskilling and training

While they might receive initial training on how to use their new GenAI technology, workers need more than that to do their job differently today. They require continuous training, upskilling and knowledge sharing due to the ever-evolving GenAI technology capabilities.

Use of AI is optional and up to the individual

Workers choose to do a specific task with GenAI while completing a broader activity because there is nothing forcing them to use GenAI technology. This contrasts with traditional digital solutions where workers are required to use the system to complete the process end-to-end.

Tailoring to individual roles and needs

Adoption needs to be tailored to individual roles, as each role will likely use GenAI technology very differently – completing different tasks, using different systems, and requiring different data.

Concerns about trust, security and privacy?

Managing the risks associated with your AI transformation **must not be seen as a compliance or risk mitigation alone.** It is a critical driver of the speed of change, as your AI transformation cannot progress faster than 'the speed of trust'.

This starts with ensuring you have appropriate **Trusted AI policies, frameworks, and governance** in place. It is also crucial to immediately raise your cyber defenses and safeguard privacy during your data transformation.

Value | People | **Trust** | Tech & Data

Trust

Concerns about AI fairness and reliability can create apprehension among the workforce.

- Bias and discrimination can be amplified
- Dependence on AI and automated insights
- A lack of explanation towards decisions made

Security

Fears about AI system vulnerabilities to cyberattacks and unauthorized access.

- Loss of control of AI systems
- Security risks and algorithm vulnerabilities
- Misalignment of AI goals

Privacy

The handling of vast amounts of personal data by AI can lead to privacy concerns.

- Loss of privacy
- Rapid changes to AI environment
- Keeping up with evolving AI regulation

Facing data readiness challenges?

To operationalize AI at scale, organizations need a reliable Digital Foundation that integrates AI with the cloud and all data.

This involves migrating data to modern cloud platforms and ensuring connectivity between data and AI tools.

The GenAI era accelerates this need, as automation depends on secure, high-quality data. Starting now is crucial due to long lead times.

Value | People | Trust | **Tech & Data**

75%

"Data security, privacy and risk are barriers to scaling AI"

According to the KPMG AI Quarterly Pulse Survey Q1 2026, 75% of respondents are somewhat or greatly concerned about data security, privacy and risk as barriers to scaling AI

Poor data quality	**Technology sprawl**	**Incorrect integration**
Data verification	**Real-time data needs**	**Issues in scaling**

Source: Global AI Pulse: Q1 2026

How KPMG can help

Wherever you are in your AI journey, KPMG can help you with the following:

Shift towards an AI mindset

Define your AI goals, identify opportunities and risks, and create a tailored strategy and execution plan. Build a business case with clear metrics to secure investments and ensure measurable success by scaling AI for enterprise-wide impact and building lasting capabilities

Empower your workforce with AI

KPMG's AI-enabled workforce solutions deliver personalized adoption and upskilling experiences, helping your team to embrace generative AI and infuse it into everyday work.

Ensure AI trust and compliance

Scaling AI introduces complexities and risks. KPMG Trusted AI teams can help ensure your AI solutions are ethical, secure and compliant. Our Trusted AI Framework, built on ten ethical pillars, empowers organizations to boldly deploy AI responsibly, transparently and with confidence.

Build a sustainable AI technology infrastructure

Leverage KPMG professionals' experience to integrate AI frameworks, platforms and accelerators, while ensuring that your technology infrastructure is ready to scale AI initiatives.

Key contacts

Tom Sprong
Partner
Strategy & Operations
+31 (0)6 - 4676 1952
sprong.tom@kpmg.nl

Bart Veenman
Lead
AI Services
+31 (0)6 - 2663 8924
veenman.bart@kpmg.nl

Jaap van Reeuwijk
Partner
Financial Risk Management
+31 (0)6 - 5253 2679
vanreeuwijk.jaap@kpmg.nl

Paul Rothwell
Partner
Finance Transformation
+31 (0)6 - 8208 3708
rothwell.paul@kpmg.nl

Key contributors

Chris Hoffman
Partner
AI & Data

Paul van Dommelen
Partner
Technology & Digital Advisory

Frank van Praat
Partner
Responsible AI

Dan Marsh
Global Director
KPMG Digital Banking Hub

Caroline Tervoort
Partner
People & Culture

Paul Kromhout
Partner
IT Advisory & Assurance

Dennis Schneider
Senior Manager
KPMG Digital Banking Hub

Aycan Nur Sağır
Senior Manager
Strategy & Operations

Jonas van Nifterik
Consultant
Financial Risk Management

Douwe Nobels
Consultant
Financial Risk Management

KPMG

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.