AKILI AI

Strathmore University
@iLabAfrica Centre

think
TECH INNOVATORS NETWORK

Banking and AI in Kenya

GenAI's Role in the New Wave of Banking and Financial Services

Contents

Key Takeaways

- Generative AI marks a step change in technology every bit as profound as the Internet
- Like the Internet, it will usher in major changes in operational efficiency, customer experience and service excellence
- While individual solutions can be rolled out rapidly and at relatively low cost, the AI revolution will profoundly change banking processes
- AI will touch every part of the organisation: from product, to compliance, to HR
- Organisations must embrace the huge opportunities, and challenges, that AI represents

The transformative potential of AI

Generative AI (the technology that underpins tools such as ChatGPT) has the potential to change society. For the first time, these "Large Language Models" are able to undertake tasks that traditional software has struggled to conquer but humans find easy – tasks like natural language, nuance, content generation and handing incomplete or conflicting data.

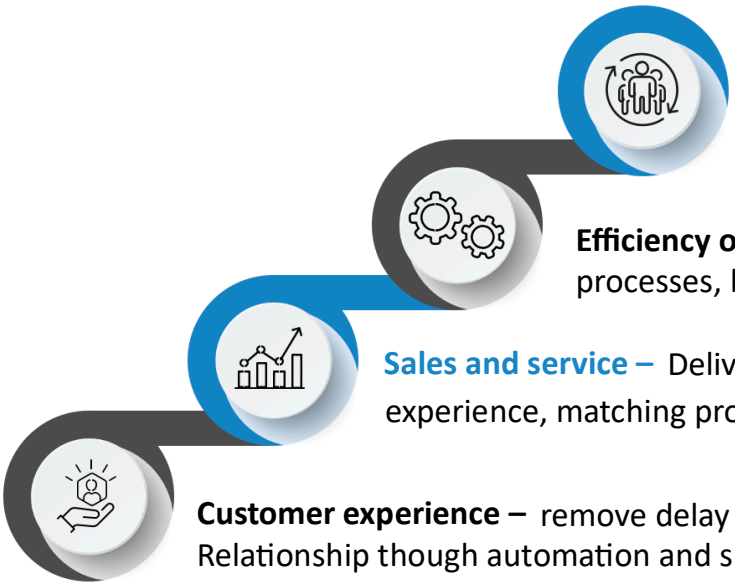

Workforce and workplaces – boost productivity and staff engagement through smart AI tools

Efficiency optimization – transform legacy processes, back-office operations and compliance

Sales and service – Deliver a fully personalised customer experience, matching products and services to customer needs

Customer experience – remove delay and frustration in the customer Relationship though automation and smart service delivery

In the banking context, this means that AI can underpin call centres, summarise complex documents, automatically "understand" paper documents, cope with imprecise and unstructured data and create personalised content that engages customers as individuals.

It can be thought of as having an army of low cost assistants that are able to undertake a wide range of often mundane and repetitive tasks, freeing staff to focus on more important activities such as customer engagement. This transition represents a profound shift in the way of doing business that will transform the industry in the coming years.

How we Got Here

On March 27, 2025, the Ministry of ICT and Digital Economy formally published the Kenya AI strategy. The document underlines the importance of AI to the future of the Kenyan economy and seeks to make Kenya a leading provider of AI in Africa, with a projected budget of some KSH.156Bn investment over the next 5 years.

"The majority of Bank CEOs believe that AI will transform their businesses within the next 5 years but only a minority of Kenyan banks have yet developed a comprehensive AI strategy"

Artificial intelligence is not new – it has been a topic of research since the 1950s and has been deployed in financial services since the 1980s. But conventional AI (so called, "Predictive AI"), while powerful, is expensive and complex to deploy. As a result, it has been used in a limited number of areas in the financial services sector, such as real-time trading or estimation of loan default risk.

In 2022, a new generation of AI came to public view in the form of ChatGPT. Unlike the earlier predictive AI, these "Generative AI" systems can be used directly to solve problems, without training or the need for data scientists. Generative AI can be thought of as a general tool that can be applied in a huge number of ways and is likely to be as impactful to finance as the Internet or the Spreadsheet.

Most importantly, the effort needed to use generative AI is much lower than predictive AI. Applications can be deployed in weeks rather than years. The first high-profile applications have been typically in call centers or for complex document management. But these AI systems are becoming increasingly ubiquitous and can be used to replace or augment a wide array of mundane tasks, much as the spreadsheet did for paper ledgers.

AI is more than a technological change. It will require adjustment across the bank, from HR to compliance and from customer care to risk management. The applications are widespread and diverse. As a result, while the core technologies will be available from product suppliers, customization, integration and ongoing management will be required by the bank.

For example, Microsoft has its Copilot range of AI services that are widely available. But while these, and others, are comprehensive platforms, they provide only a foundation layer. Copilot itself is not a product as such but rather an umbrella term for a collection of very different technology offerings. While users can easily benefit from an assistant that will help write an email or a powerpoint (Microsoft Office Copilot), it is a very different proposition to create a tool in the Copilot AI Stack that will mine social media in order to create qualified leads for youth bank accounts.

While this may at first sound daunting, the relatively low barrier to adoption of generative AI means that early adopters can take the first steps on the ladder and demonstrate business value without large scale investment programs. In addition, a new generation of AI-as-a-Service providers are emerging to remove much of the complexity of AI adoption.

"Challenger Banks that have embraced AI and digital transformation have transaction overheads 8-10x lower than traditional banks"

Generative AI represents both an opportunity and a threat to established banks. The opportunity is improved margins, improved products (for example, much faster customer onboarding and more responsive and targeted marketing) and reduced operating costs. The threat is digital banks and fintechs that have already invested heavily in digital technology, don't have the complex legacy of disparate in-house systems and so are able to transact with a fraction of the operating costs of traditional banks.

For those that have yet to do so, banks need to put in place an organization-wide AI strategy for the next 3-5 years. In a field that is evolving as fast as AI, it is not possible to predict with confidence the precise state of the market in 5 years time. However, it is possible to provide a direction of travel, to embrace AI for specific solutions that add business value today and to position the bank in the best possible place to take advantage of the new AI technologies as they emerge.

The state of the market

McKinsey estimates that generative AI will fundamentally change the nature of banking operations by 2030 and be responsible for a 3-5% total growth in banking revenue globally. AI has been enthusiastically embraced by the financial sector in the US, Europe and Asia but is at a formative stage in Africa (with the notable exception of South Africa). Similarly, major equipment suppliers such as Microsoft are present in African markets but adoption has so far been much slower than elsewhere in the world.

"AI has the potential to augment virtually every job"
- Jamie Dimon, CEO, JP Morgan

While Kenya remains one of the leading technology centers in Africa, adoption of AI in banking is lagging. Relatively few banks have a clearly defined AI strategy and deployment is well behind international players. This represents an opportunity for the bank to capture early advantage with the strategic and carefully targeted adoption of AI.

Steps in building a company-wide approach to AI

Top-down:

- Align the capabilities of AI with overall strategy and shorter term KPIs
- Select the initial application areas for AI in the bank
- Build the governance, ethics, reporting and management structures needed to ensure successful rollout
- Build a change-management plan for the evolution of the bank into the AI age
- Share the AI mission of the bank to staff to ensure widespread engagement

Bottom-up:

- Socialize AI widely through the bank to identify key pain points that AI can address
- Create a small number of rapid proof-of-concept solutions to test the benefit of AI in business use
- Communicate the risks associated with AI, even today (for example, putting PII into ChatGPT at work) and ensure compliance with best practice

Adoption of generative AI has been rapid and widespread in markets such as the US, Europe and Asia. A report by Emerj highlights that major U.S. banks are investing heavily in AI, particularly in areas like fraud detection, compliance, and customer service. For instance, JPMorgan Chase has over 2,000 AI and machine learning experts and has identified around 400 AI use cases across its operations

"AI is not software as we know it. Rather, it is a profound change in the way we do business that we are only beginning to understand"

Examples of successfully deployed generative AI solutions in banks include:

JPMorgan's COIN platform, which reads and interprets complex legal documents like loan agreements. It automates the extraction of key clauses and data, significantly reducing the manual workload of legal teams. What once took hundreds of thousands of hours annually is now handled in seconds with improved accuracy.

Wells Fargo's mobile app uses generative AI to create natural, context-aware responses for customers asking about balances, spending, and financial advice. The assistant learns from past queries and tailors answers, offering a more human-like interaction and helping reduce reliance on call centers for routine questions.

ING uses generative AI to auto-generate risk model documentation by pulling data from model outputs, audit trails, and regulatory requirements. ING's compliance and risk teams save time by not having to manually compile lengthy technical documents, while ensuring consistency and completeness across reports.

Citi uses generative AI to create highly personalized marketing messages tailored to individual spending habits and preferences. The AI generates custom email and app content that aligns with a customer's financial behavior, driving higher engagement rates and improving return on investment for marketing initiatives.

BNP Paribas uses generative AI to draft large sections of regulatory reports. It summarizes transaction data, compliance notes, and internal reviews into coherent narratives. This reduces the burden on compliance teams, ensures uniform language, and helps meet strict deadlines for cross-border reporting.

The uses for AI that are changing banking

AI is a step-change in the banking sector. While it may at first appear to be "just software", it has very different capabilities that allow it to perform tasks that humans find easy but traditional software finds very difficult.

"Large AI applications get the headlines. But the real value is in the hundreds of simpler and more mundane solutions that are transforming operational performance"

For conventional software, the programmer has to imagine every possibility. For example, in the case of reconciliations, the software needs to know every possible date format or field layout. In the real world, where data is in multiple locations, in different formats and sometimes incomplete or inaccurate, human intervention is often needed to undertake what are otherwise mundane and boring tasks.

Generative AI on the other hand can much more easily determine a date and distinguish it from an account balance on an invoice, in much the same way that a human can easily do. This opens up a wide range of new applications for automation. This may be automation of existing processes such as simplifying customer onboarding. Or it may be undertaking processes that humans can readily do but it would simply be uneconomic for a bank to undertake.

For example, this ability to "do what I mean, not just what I say" means that AI can be used for entirely new purposes that range from smart chatbots in a call centre to creating highly personalized marketing content for customers based on transaction history (much as social media platforms currently do). This capability can then be linked to existing banking systems including CRM, account history and transaction ledgers.

AI is not a "product" but rather a technology that can be used in a wide variety of different ways. It can be thought of like a spreadsheet or a database – the core application is provided by a supplier but then it is customized at low cost to meet the needs of the bank.

This ability to deploy AI-based tools at low cost and at high speed means investment cycles for AI are generally much faster than conventional enterprise software. But AI also requires the bank to gain new skills. It is not sufficient for organisations to simply buy in products. Questions of governance, ethics, policy, privacy, reporting, supplier management and organisational culture all need to be addressed. Getting in early with small but valuable AI solutions allows the bank to adjust to the new reality and embrace the new technologies as they arrive.

Solution	Department	What it does	Size
Smart Knowledge Base	Company-wide	Tool that allows large numbers of documents to be queried in natural language. Typically for compliance, loan officers, customer service, front-line staff	S-M
AI Chatbot	Customer Care HR Sales / Marketing	Interactive tool that automates 70%+ of customer enquiries (essentially level 1 support). First application in online chatbots in customer care but also application to automated sales campaigns and internal knowledge base (e.g Human Resources)	Varies, depending on Solution
Customer segmentation / profiling	Sales / marketing	Analysis of customer transaction data to micro-segment customers into different usage groups (similar to how social media categorizes users). Allows hyper-personalized marketing	M
Hyper-personalised marketing	Sales / marketing	Use segmentation scorecard to match product offers to customers. Use AI to create compelling, targeted outreach. Use AI chatbot to field responses to allow massive scale up and timely responses	M
Social Media Engagement	Sales / marketing	AI agents that monitor and interact with social media for targeted marketing, sentiment analysis and customer engagement. Reduce the response time to social media from weeks to minutes.	M
Intelligent data handling	Operations, customer onboarding, reconciliations	AI is able to handle unstructured and inconsistent data. AI Agents can undertake tedious, human-intensive tasks including suspense account reconciliation, kyc validation and smart OCR	M-L
NPLs	Operations	AI can analyze NPLs in normal conditions but don't work well following a shock (for instance a tax rise or inflation). Generative AI can monitor customer accounts and identify customers at future risk to enable loan officers to intervene proactively	M-L
Fraud management	Compliance	Generative AI is able to combine multiple sources of unstructured data to identify patterns in behavior that indicate possible fraudulent activity, freeing up compliance officers to focus on the most promising cases	M-L

Regulation, Compliance and Legislation

As generative AI becomes increasingly embedded in the financial services landscape, banks are exploring how to harness its potential responsibly. However, due to the regulated nature of banking, adoption must be carefully aligned with regulatory requirements, compliance frameworks, legislative obligations, and robust internal operations. Unlike larger multinational banks with deep legal and risk teams, many African banks must be especially strategic in how they address these issues in order to balance resource and impact.

> *"AI compliance is like tax returns. It is not especially difficult but it requires dedicated attention and is not something to be ignored"*

Before deploying generative AI tools at scale, banks must understand the relevant legal and regulatory frameworks. This includes data protection laws (such as GDPR, Kenya's Data Protection Act and similar laws), financial sector regulations (e.g., central bank guidelines), consumer protection standards, and sector-specific AI guidance. Banks should also take into consideration the Kenya AI Strategy 2025 and ensure that operations are aligned with the overall direction of their recommendations.

Generative AI is a very new technology, with the first commercial implementation, ChatGPT, only arriving in November 2022. As such, the regulatory environment for AI is still in its infancy. There are a number of well-defined AI policy frameworks including NIST (USA), the EU AI Act and Singapore's Model AI Governance Framework. All the frameworks have similar underlying principles and approaches.

However, it is clear that banks will be held to account, not on the regulations that exist today but those that will exist in the future. As a result, banks should consider defining standards and policies internally that are likely to be futureproof and demonstrate the bank's use of industry best practise as the regulatory environment evolves.

Banks should create an internal AI governance policy that clearly defines he responsible roles for overseeing AI projects (e.g., a compliance officer, risk officer, or dedicated AI lead), the approval processes for AI tools and reporting lines and escalation procedures for potential AI-related issues. It is also recommend that the bank constitute a high level, cross functional AI Steering Committee as a subset of the Board that can provide oversight across business units and help maintain transparency and alignment with compliance expectations.

Every generative AI application should undergo a formal risk and impact assessment before implementation. This includes Model Risk (evaluating how the AI makes decisions, its potential for bias, and how errors might affect customers or operations) Compliance Risk (whether outputs may breach legal obligations, such as making misleading claims in marketing

"It is not sufficient to blame a supplier for shortcomings relating to AI. Ultimate responsibility for AI compliance rests with the bank"

or generating discriminatory decisions in lending) and a data Privacy Impact Assessment. While this may at first sight appear daunting, for most simple AI applications that are low risk (such as tools to optimise customer onboarding to automate reconciliation) the AI risk is low and, while the process needs to be undertaken, a lightweight approach is quite sufficient. This should therefore not be a major barrier to adoption.

As part of this, generative AI tools should be clearly bounded in their application, especially in customer-facing and decision-making scenarios. Specifically they should:

- Define what AI can and cannot do. For instance, it can assist with drafting internal reports, but not issue credit decisions without human review.
- Implement content filtering, prompt controls, and human-in-the-loop review mechanisms.
- Use AI output disclaimers where appropriate to maintain customer trust and transparency.

This is critical to avoid "AI hallucinations" (fabricated or misleading output) being passed off as authoritative information.

It is also important that employees understand how to work with AI responsibly, particularly in relation to data privacy responsibilities and ethical considerations. In this early (and largely experimental) phase of AI adoption, it is important that all staff have clear training on how to escalate AI-related anomalies or concerns so that any issues can be addressed quickly and effectively.

The responsibility for AI rests with the bank, it cannot be delegated to suppliers. For example, if a bank uses AI in a third-party CV selection tool and bias is discovered in the application of the tool, it is likely that the bank will face liability rather than the tool vendor. As a result, banks must implement robust third-party supplier engagement policies that ensure the suppliers are compliant with both the bank and wider regulations. Contracts should include provisions for audit rights, liability, and service level agreements that reflect regulatory expectations and banks should engage such specialist support as may be needed in ensuring the appropriate provisions are embedded in supply contracts and that these provisions are audited to ensure compliance.

Adopting generative AI can bring powerful capabilities to banks, enabling them to operate more efficiently and serve customers more effectively. However, success hinges on a responsible, well-governed approach that integrates regulatory compliance, internal controls, and ethical AI use. By choosing initial applications of AI that are inherently low risk but still highly valuable, and taking appropriate proactive steps across legislation, reporting, and operations, banks can confidently harness AI's benefits while minimizing risk and maintaining regulatory trust.

Steps to becoming AI ready

The future use of AI is likely to be so widespread that it's sometimes difficult to know where to start. AI will become broadly embedded in commercial products and a wide variety of AI platforms will emerge. But none of these are useful until they are integrated into the bank's systems and data.

The advantage of generative AI over earlier technologies is that it can be applied to a wide range of tasks in a period of weeks rather than years. Therefore banks can take a step-wise approach to AI adoption, selecting individual profitable solutions one at a time and evolving towards a business model where AI plays in integral part in almost all business operations.

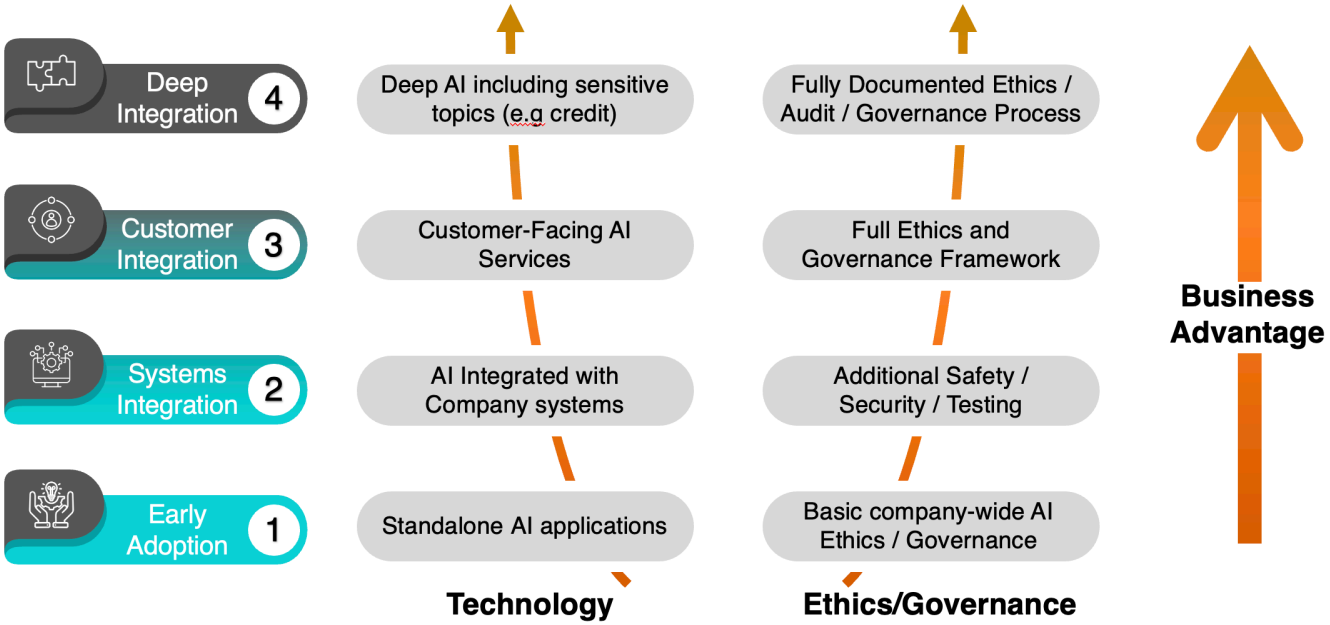

While it would be possible to go all-in on artificial intelligence (and some challenger banks have), for most established banks the adoption of AI is a more stepwise approach. This is in part because many banks with legacy operations have data stored in multiple siloes with no unified means of access. While this needs ultimately to be addressed, it limits the extent to which AI can easily be rolled out across the organisation as a single layer. In addition, AI technology is moving at such a pace, it is prudent for banks not to become too aligned to one specific technology that may be superseded in short order.

The journey toward adopting generative AI should be strategic and grounded in practical steps that align with resources, the competitive landscape and customer needs. The first step is to identify specific business problems or processes where generative AI can add value. Banks should focus on high-impact, low-risk areas such as customer service (e.g., AI chatbots), content generation (e.g., automated customer emails or reports), marketing personalization,

document summarization, and compliance support. Selecting well-defined, narrow use cases helps ensure early wins that build internal confidence and stakeholder support.

Generative AI is a new field that requires both technical and domain knowledge to adopt in the most cost-effective way Specifically, it is key for the bank to create the appropriate foundations of reporting and governance to act as a basis for future development work. Banks can partner with technology providers, fintechs, or AI consultants to access needed expertise without having to hire large in-house teams. At the same time, banks should consider upskilling a small internal team through online courses and vendor training programs to ensure that the management of technology providers and consultants is undertaken in the most effective manner. A designated AI lead or project manager can coordinate these efforts.

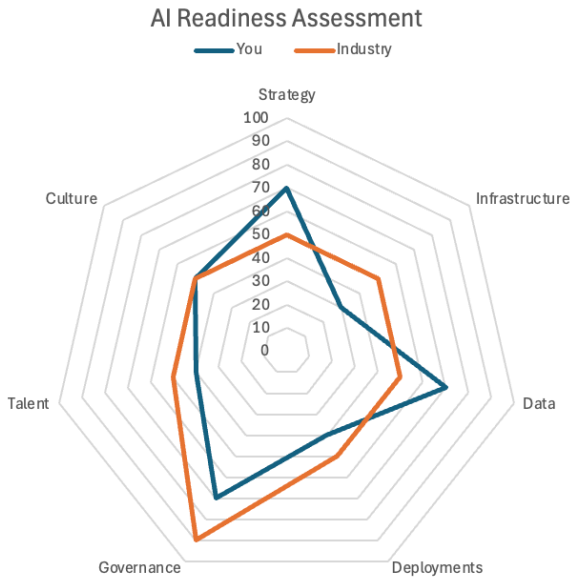

The process can be broken into three phases: The preparation (Engage) phase sets up the organisation with the core capabilities for success without substantial investment. The next phase (Evaluate) involves creating a small number of rapid proof-of-concept applications to demonstrate the value of the technology and to align the organisation to the changes in process, management and governance that will be needed with the new solutions. Only once the foundations are in place and the business alignment has been demonstrated does it make sense to invest in larger-scale rollout of technology.

While this process is comprehensive, it need not be time-consuming or expensive. With proper planning, most organisations will be able to undertake the preparatory Engage and Envision phases within an elapsed period of 6 months.

Most of all, the application of AI should be aligned to a clear business case with well-defined success criteria, such as reduced customer service response times, improved customer satisfaction, or increased marketing campaign conversion rates. Pilots allow organisations to measure impact, learn from implementation, and adjust the AI's role in real-world workflows before scaling.

Incorporating compliance and risk management teams early in the process is crucial. Generative AI outputs can be unpredictable or even biased. Establishing governance frameworks to review AI outputs, monitor ethical concerns, and control access helps avoid regulatory pitfalls and reputational risks.

"AI is moving fast and nobody can predict exactly where it will be in 5 years. But we know it will be everywhere and we ignore it at our peril"

After successful pilots, AI applications can scale across departments. This may include integrating AI into loan origination workflows, internal knowledge management, or fraud detection support. Scaling should be gradual, allowing teams to adapt and systems to mature. To do so will also require consideration of the bank's internal data organisation and cybersecurity to ensure that more expansive applications remain secure and compliant.

It is worth noting that AI continues to evolve at a rapid pace. Banks need to match the rate of investment with the risk that technology being developed could become redundant in the future. To address this, many banks are adopting an AI-as-a Service (similar to Software as a Service) approach to AI in which core data and services are maintained in-house but many AI aspects are outsourced to a commercial supplier. This provides the bank with flexibility to change provider in the future as needed and removes much of the operational overhead of maintaining AI systems.

By taking a step-by-step approach grounded in business needs and operational reality, banks can harness generative AI to stay competitive, even with limited resources.