




Insights | 6 October 2026

INSIDE CBK'S AI PLAYBOOK

Where the draft nails it and where it falls short.

SAVARIC AFRICA | savaric.africa



Inside CBK's AI Playbook: Where the Draft Nails It and Where It Falls Short

Kenyan banks are already using AI solutions across their operations, including to help decide who gets a loan. A Central Bank of Kenya (CBK) survey in 2025 found that half of the institutions surveyed had adopted AI, including 66 per cent of commercial banks. Credit risk assessment was the most common use, at 65 per cent. Among institutions that had adopted AI, 68 per cent reported using generative AI tools and frameworks. Another 35 per cent used third party service providers for AI and machine learning services. CBK has now released a draft Guidance Note on Artificial Intelligence in the Banking Sector, open for public comment in 2026. It applies to banks, microfinance institutions, credit reference bureaus, money remittance providers, non deposit taking credit providers and holding companies. It gets a lot right, and it also leaves a few questions open.

Draft or not, the expectations set out in the text are already clear. The draft requires institutions to seek CBK's approval before deploying an AI system. The application must include a data governance and AI checklist. A register of every AI system is due as soon as the note takes effect and then every year. The draft also requires notice to CBK within 24 hours of a material AI incident. Every other requirement should be read with this in mind.

CBK has done well to link its approach to Kenya's National AI Strategy 2025 - 2030, which was launched in March 2025. That link shows that the draft expects AI adoption in institutions to support the country's broader priorities. Boards are also expected to oversee an AI strategy, which sets out directions on how the institution will adopt and derive value from AI. The board must set the business objective and justification for each use of AI and keep it aligned with the institution's strategy. This matters because only 30 per cent of surveyed institutions had a formal AI strategy.

Every stakeholder has a part to play in AI integration in the banking sector. AI creates legal, reputational and operational risks that need oversight beyond the tech team. The board carries primary responsibility. Beneath it, CBK asks for a data governance committee, an AI governance committee, and a model risk committee that stays independent of the people who build the models. The data and AI committees draw on ICT, legal and the business units. The data committee includes the Data Protection Officer. Internal audit reviews the whole framework every year. The use of AI systems is a collective responsibility and should not be left to the ICT teams alone.

The approach also leans on risk practice that banks already know. AI risk sits within enterprise risk management, and banks are asked to scale the process to their size, complexity and risk profile. Each system is assessed by materiality, criticality and

consequences. CBK says it drew on practice from leading and comparable jurisdictions and adapted it to the supervision framework Kenyan banks are used to. Its strictest label is for agentic AI, which it classifies as high risk. This is a sensible approach, rather than simply importing something like the EU risk tiering without context.

On oversight, the draft is firm. Agentic AI must remain open to human oversight, with the ability to review, override, suspend or terminate it, and it must keep audit logs. For critical decisions such as credit denial, fraud flagging, account freezing and pricing changes, customers can request a human review. Reviewers must use independent judgment. CBK warns them against leaning too heavily on the AI's recommendation. This suggests a human who simply approves the output is not enough. Customers must also be told when they are dealing with AI, and they can ask for the factors behind a decision. Taken together, these requirements embed a human in the loop for high impact AI decisions in the banking sector.

Vetting before launch is built into the approval step. Before a system goes live, banks should train, test and validate it. That testing should include security testing to protect sensitive training data. It should also prevent data manipulation, contamination and model hijacking, which covers what is often called model poisoning. Banks must also run a bias impact assessment. Training data should represent the setting where the model will run. The draft warns against proxies such as postal code that can stand in for race or ethnicity.

However, AI models built from scratch in Kenya are likely to be minimal, so most of the models in use are probably foreign built. The CBK survey notes that 35 per cent of respondents that had implemented AI used third party service providers for AI and machine learning services. Some partnerships included AI and fintech start ups, to deliver capabilities such as credit scoring and financial crime monitoring. This matters because the register records where a service is performed and where data is hosted, not where the model originated. A bank that licenses a model may not hold its training data. A key concern is inherited bias. Fine tuning a foreign model on Kenyan data can reduce it, and it should not be assumed that this removes it. Banks should ask vendors about provenance before they sign.

Smaller players may find the rules hard to absorb, although some duties do scale to size. The risk management process and its controls should be commensurate with an institution's size, complexity and risk profile. The data and AI governance committees may be combined or kept separate according to size. The approval requirement carries no such scaling. It covers AI systems in general, with a checklist asking for a data strategy, AI strategy, policy, risk framework and committees. CBK's own survey named limited skills and high costs as barriers. The note lists responsible innovation among its aims, yet in the current draft there is no sandbox, no approval timeline and no transition period. Without a sandbox or scaled approval for small and new entrants, critical systems

and agentic AI face full review while low risk ones face the same hurdle. Institutions already running systems have no clear timeline or transition path, even though consumer protections would still apply in full.

On AI regulation, sector specific rules make sense for Kenya. The draft is a banking specific note written by the banking regulator, and that fits how risks show up in practice. A health regulator will judge health risks better than a general AI body. A central bank understands what a credit decision does to a household. The risk is gaps between regulators. A shared baseline can close them, and the draft already leans on the Data Protection Act for automated decisions and on consumer protection law for fairness.

Data is where the draft is strongest. Banks are expected to keep data lineage, document and authorise their AI data sources, and keep a record of the original datasets used to train their models. Banks stay accountable for third party AI. They are asked to consider how they keep control over a system's parameters and outputs. The draft reaches downstream providers in the vendor's supply chain as well as the vendor itself.

For critical systems, smaller models whose data and behaviour a bank can inspect, hosted where the bank can control them, are preferable. Smaller is not automatically safer, and the point is inspectability. Such models are easier to validate, and easier to keep within the control over parameters and outputs that the draft asks banks to consider. Large language models are like black boxes, since they are harder to validate, harder to explain to regulators and customers, and riskier when they drive high impact decisions such as credit, pricing or account restrictions.

In sum, the draft Guidance Note sets a strong foundation for safe and accountable AI in Kenya's banking sector. It gets the big things right: clear governance, human oversight, data discipline and alignment with the national AI strategy. The main gaps lie in flexibility for smaller players and clarity on model provenance.

Sources: Central Bank of Kenya, Survey on Artificial Intelligence in the Banking Sector (2025); Central Bank of Kenya, Draft Guidance Note on Artificial Intelligence in the Banking Sector (consultation draft, 2026).