POLICY COMMENTARY / OCTOBER 2026

An AI Answer Is Not a Policy Decision: The Case for Auditable AI in African Public Institutions

The missing design principles for responsible language-model adoption in public-sector decision support

Eric John Nimba | Full-Stack Software Engineer & AI Officer | Nairobi, Kenya

Imagine a county officer asking an AI assistant: "Which climate projects missed their reporting deadlines, and why?" The assistant responds in seconds, names three projects and offers a persuasive explanation. The problem is that one explanation came from an outdated report, another confused planned expenditure with actual spending, and the third had no supporting source. Speed has produced the appearance of knowledge, not a defensible public decision.

This scenario is hypothetical. But it points to the central risk of deploying generative AI in government: a fluent answer can conceal weak evidence. Having worked on Python-based document-retrieval and monitoring applications, I believe the question for African institutions is no longer simply whether an AI assistant can summarise policy. It is whether an official can verify, contest and safely act on its output.

AI strategy must become operational governance

The African Union endorsed a Continental Artificial Intelligence Strategy in 2024, advancing a development-focused, ethical approach to AI. Kenya's AI Strategy 2025–2030 similarly identifies data and AI governance as a central pillar. Those ambitions matter, but responsible practice is decided inside systems: what evidence systems use and who can override their outputs.

The distinction is particularly important for climate governance, where policy documents change, reporting indicators evolve and agencies hold overlapping responsibilities. A retrieval-augmented generation system can search approved documents before drafting a response. Yet retrieval alone does not guarantee truth. A model can still misread a table, merge two versions of a policy or overstate what an indicator establishes.

Six safeguards before deployment

Start with traceable sources. Every material claim should link to the document, section, date and version used to generate it. An answer without a retrievable source should be presented as unverified, not authoritative.

Second, require meaningful uncertainty. If the available records cannot establish whether an expenditure was approved or actually paid, the assistant must say so. Refusing to answer is sometimes better public service than inventing certainty.

Third, define access by role. A public-facing climate information assistant and an internal programme officer should not see the same underlying records. Access controls, retention rules and logs are essential, especially where personal information or sensitive procurement records are involved. Kenya's Data Protection Act provides the relevant domestic legal baseline for personal-data processing.

Fourth, test beyond impressive demonstrations. Evaluation should include outdated documents, contradictory files, local terminology, Kiswahili queries, poorly formatted tables and deliberately misleading prompts. Institutions should measure unsupported statements, citation accuracy, fairness and the rate at which the system appropriately declines to answer.

Fifth, assign responsibility. Human reviewers should approve consequential outputs and retain authority over funding, compliance and policy decisions. An AI-generated recommendation must not quietly become the official institutional record.

Sixth, create a route for correction. Users should be able to flag errors, see when an answer has been superseded and understand how the system is improved. This is not just an engineering feature; it is a practical form of administrative accountability.

Before procurement, an agency should publish a limited use-case definition: which decisions the assistant may support, which data it may process, who supervises it and which decisions are strictly excluded. A controlled pilot can compare AI-supported responses against answers produced by trained staff using the same reference material. Testing should include ordinary questions as well as ambiguous or adversarial ones. Evaluate results across user groups and pause deployment when harmful or unsupported outputs appear. Responsible AI needs a stop button with an accountable owner.

The social-science questions are technical requirements

A technically accurate response may still be inequitable. Who chose the documents in the knowledge base? Which languages are missing? Can a community organisation challenge a classification? Do the people most affected by a decision understand why the system recommended it? These are questions of institutional power, representation and trust—not optional additions after deployment.

UNESCO's Recommendation on the Ethics of Artificial Intelligence emphasises human rights, fairness and oversight. The US National Institute of Standards and Technology's Generative AI Profile likewise offers a framework for evaluating AI-specific risks. African public institutions can adapt these principles to local law, capacity and context rather than treating imported technical benchmarks as complete solutions.

A more useful definition of innovation

The success of a public-sector AI system should not be measured by the number of questions it answers. Better measures include the share of answers backed by valid sources, the number of errors detected before decisions, the time needed to correct a misleading result and whether staff and citizens retain meaningful avenues for scrutiny.

African governments have an opportunity to make AI-assisted administration more transparent than the manual processes it supplements. But that opportunity depends on deliberate design. The best policy assistant is not the one that sounds most confident. It is the one that makes evidence visible, limitations legible and human accountability impossible to bypass.

References and further reading

1. African Union — Continental Artificial Intelligence Strategy (2024)
2. Kenya Ministry of ICT — AI Strategy 2025–2030
3. Kenya Office of the Data Protection Commissioner — Data Protection Act
4. UNESCO — Recommendation on the Ethics of Artificial Intelligence
5. NIST — Generative AI Profile (2024)

AUTHOR NOTE Draft policy commentary prepared with AI assistance for the author's independent review, fact-checking, substantial personal revision and outlet-specific disclosure. It reports no original field study or institutional finding. Views are personal, not official AGNES positions.