
Buying Safety: A Model AI Procurement Standard for African Public Sectors¹

Ayodeji Adesegun
Independent Researcher

With
Apart Research

Abstract

African governments are deploying AI in public services, welfare targeting, tax administration, health triage, identity systems, faster than they are governing it. The African Union's Continental AI Strategy (2024) named AI procurement standards as a regulatory gap to be filled, yet that gap remains open: no continental standard exists, and national procurement laws predate AI. This paper argues that procurement is an underused but immediately available governance lever, as the largest buyer of AI in most African economies, the state can set a de facto safety floor through what it requires before signing, without waiting for comprehensive AI legislation. We assess the procurement frameworks of four African countries (South Africa, Nigeria, Kenya, Rwanda) against nine AI-safety safeguards and find that none impose any AI-specific requirements. We then propose a model AI procurement standard built around four Africa-specific constraints, namely foreign-trained models, data dependency, capacity gaps, and local-language failure, that states can adopt directly, and the AU can harmonise continentally.

¹ Research conducted at the [The Global South AI Safety Hackathon](#), June 2026

1. Introduction

Across Africa, AI systems are entering public service delivery: welfare and subsidy targeting, tax and revenue administration, healthcare triage and diagnostics, biometric identity, and citizen-facing service platforms. These systems increasingly make or shape consequential decisions about citizens.

The significance for AI safety is substantial: global efforts have concentrated on regulating frontier development in the US and China, leaving deployment-side risks in Africa largely unaddressed. High-stakes public AI systems are procured through processes originally designed for physical assets, not for statistical systems with opaque behaviour, performance that relies on unseen data, and errors that disproportionately affect vulnerable populations. Systems that deny welfare claims, misinterpret clinical cases, or misidentify citizens, without mandated impact assessment, bias testing, transparency, or appeal mechanisms, constitute a safety failure. Acquisition represents the final opportunity for the state to impose safety conditions before harm occurs; currently, most African jurisdictions impose none.

The African Union's Continental AI Strategy (2024) explicitly identifies AI procurement standards as a regulatory gap. This recommendation remains largely unimplemented: no continental standard exists, and national procurement statutes predate the technology. As the government is typically the largest purchaser of AI in emerging economies, its pre-contractual requirements establish a de facto market standard. Procurement thus serves as a uniquely powerful and immediately available governance lever, obviating the need to wait for comprehensive AI legislation.

The main contributions of this paper are as follows:

1. A comparative procurement gap analysis of four African countries' public procurement frameworks (South Africa, Nigeria, Kenya, Rwanda) against nine AI-safety safeguards, demonstrating that none imposes any AI-specific requirement (28 of 36 cells marked "None," with the remaining 8 only indirect).
2. A model AI procurement standard comprising a risk-tiered set of adoptable clauses, including four safeguards specifically adapted to African deployment constraints: local-context performance validation, anti-lock-in/data sovereignty, capacity transfer, and local-language failure.
3. An implementation pathway for near-term adoption by national procurement authorities and for continental harmonisation by the AU and Regional Economic Communities, as envisaged by the Continental AI Strategy.

2. Related Work

This work builds on three strands of prior work. First, the African Union Continental AI Strategy (2024) identified AI procurement standards as a regulatory gap but did not specify what such standards should contain. Second, the literature on pragmatic African AI governance argues that oversight is best advanced by strengthening existing regulatory regimes (data protection, consumer protection, sectoral rules) rather than enacting new AI-specific laws (OECD, 2026); procurement is precisely such a regime. Third, existing AI procurement guidance from non-African contexts (UK Government, World Economic Forum) and the academic treatment of AI governance through markets (see [arXiv:2501.17755](https://arxiv.org/abs/2501.17755)), which establishes the concept but is not adapted to African constraints.

The paper also situates itself within the emerging African AI safety agenda, in particular Segun et al. (2025), "[Toward an African Agenda for AI Safety](#)", as a concrete deployment-side instrument that

operationalises that agenda. It draws on the human-rights-based AI lifecycle framing of the African AI and Equality Toolbox (Turing Institute & OHCHR).

This work addresses a gap in existing AI-procurement guidance, which is largely non-African and generic. While African policy literature identifies procurement as a regulatory gap, it does not provide a practical instrument to address it. This paper offers (a) a comparative evidence base detailing what African procurement frameworks require and omit, and (b) a standard tailored to the four constraints defining African deployment risk. The proposed standard is intended for contexts where deployment is in Africa, with foreign-trained models, data and infrastructure dependencies, capacity scarcity, and low-resource-language failure as central design considerations.

3. Methods

This is a comparative legal-policy desk study producing a normative instrument. It has two phases: a gap assessment and a standard derivation.

3.1 Country Samples

We selected 4 countries, namely South Africa, Nigeria, Kenya, and Rwanda, chosen for contrast in national AI strategy maturity and for the public availability of their procurement legislation, not for statistical representativeness (see Limitations). For each, we identified the primary public procurement statute(s) and any AI- or digital-specific procurement guidance, using public legal sources, the Africa AI Policy Tool, and the OECD.AI Policy Navigator to locate instruments.

3.2 Safeguard Taxonomy

We assessed each framework against nine safeguards that a procurement process could impose on an AI system. The taxonomy was derived by combining (i) the gaps the AU Strategy names (procurement standards, bias/discrimination protection, healthcare approvals), (ii) elements common to existing UK/WEF AI procurement guidance, and (iii) the human-rights AI-lifecycle stages in the African AI and Equality Toolbox. The nine safeguards are: (1) AI risk classification/impact assessment pre-tender; (2) bias and fairness testing on local populations; (3) performance validation on local languages and contexts; (4) transparency/explainability documentation; (5) human oversight and citizen contestability; (6) data governance and data-protection compliance; (7) vendor liability and audit rights; (8) anti-lock-in / data portability; (9) post-deployment monitoring.

3.3 Coding

For each country–safeguard pair, we coded the framework's requirement as Yes (explicitly required), Partial (indirectly or partially addressed), or None (silent), citing the governing provision. We assessed what frameworks require on paper as of June 2026, not enforcement in practice.

3.4 Standard Derivation

Each safeguard found absent across the sample was translated into an adoptable procurement clause. Four clauses were specifically adapted to the African deployment constraints: foreign-trained models, data-colonial dependency, compute/capacity scarcity, and low-resource-language failure. A risk-tiering gateway was added so requirements scale to the consequences of error, keeping the standard proportionate and adoptable rather than a blanket burden.

Reproducibility: The study can be reproduced by retrieving each country's procurement statute (sources listed in References), applying the nine-safeguard coding scheme in 3.2–3.3, and deriving clauses as in 3.4. No code or proprietary data is required.

4. Results

4.1 The procurement gap

#	Safeguard	South Africa	Nigeria	Kenya	Rwanda
1	Risk classification/impact assessment pre-tender	None	None	None	None
2	Bias & fairness testing on local populations	None	None	None	None
3	Local-language / local-context performance validation	None	None	None	None
4	Transparency/explainability documentation	None*	None*	None*	None*
5	Human oversight & citizen contestability	None†	None†	None†	None†
6	Data governance & data-protection compliance	Partial‡	Partial‡	Partial‡	Partial‡
7	Vendor liability & audit rights	Partial§	Partial§	Partial§	Partial§
8	Anti-lock-in /data portability	None	None	None¶	None
9	Post-deployment monitoring	None	None	None	None

Table 1. *AI-safety safeguards required by national procurement frameworks. SA = Public Procurement Act 28 of 2024; NG = Public Procurement Act 2007; KE = Public Procurement and Asset Disposal Act 2015 (rev., 2022); RW = Law N°62/2018.*

Coding notes:

- *** Transparency (4):** All four Acts mandate strong process transparency (open tendering, public scrutiny, e-procurement) but not model transparency or explainability of AI decisions. None as defined.
- **† Contestability (5):** All provide procurement dispute/review for losing bidders (SA's debarment and review; KE's Review Board and judicial review; RW's Independent Review Panel), but nothing for a citizen contesting a decision a deployed AI makes about them. None as defined.

- **‡ Data governance (6):** The procurement Acts are silent; coded Partial only because separate national data-protection law applies to deployed systems (SA: POPIA; NG: NDPA 2023; KE: Data Protection Act 2019; RW: Law N°058/2021) - indirect, not via procurement rules.
- **§ Vendor liability/audit (7):** All have general supplier-liability and oversight/audit provisions (NG: Bureau of Public Procurement; KE: Public Procurement Regulatory Authority investigation/audit; RW: RPPA monitoring + anti-corruption/obstructive-practices rules; SA: audit and debarment), but none AI-specific (no liability for harmful model outputs, no model-audit right) - Partial, general-only.
- **¶ Anti-lock-in (8), Kenya:** Still None as defined (no data portability/exit requirement), but note Kenya's regulations do require a tender plan for technology and knowledge transfer to Kenyan citizens, an adjacent precedent for Clause 9 (capacity transfer).

Finding: Across all four frameworks, of 36 cells (9 safeguards × 4 countries), 28 are None, 8 are Partial, and none is a full Yes. Every one of the eight Partials is indirectly arising from general procurement/contract provisions (vendor liability, audit) or from a separate national data-protection law, never from an AI-specific procurement rule. The result is uniform across countries of very different AI-strategy maturity: no sampled procurement framework imposes a single AI-specific requirement on the AI systems governments acquire. The safety floor for public-sector AI procurement in the sample is, on paper, zero.

Per-country notes.

South Africa: Public procurement is set to be governed by the Public Procurement Act 28 of 2024 (signed July 2024), which consolidates the previous framework (Constitution §217; PFMA 1999; PPPFA 2000) into a single Act. While the Act is currently undergoing a phased implementation, awaiting the finalisation of Section 63 regulations, its text reveals a distinct gap in technology governance. Its references to technology concern using technology to run procurement efficiently and combat corruption, and it promotes innovation. However, it contains no provision governing the AI systems being procured: no impact assessment, bias testing, explainability, contestability, or anti-lock-in requirement. Its transparency provisions govern the procurement process.

Nigeria: The Public Procurement Act 2007 remains the active statutory framework, establishing the Bureau of Public Procurement and the National Council on Public Procurement, and focusing on transparency, competition, and accountability. Its only technology that empowers the Council to "approve changes in the procurement process to adapt to improvements in modern technology", modernising the process, not governing acquired AI. No AI-specific safeguard appears. This sits alongside the fact that Nigeria has no dedicated AI statute and instead relies on the NDPA 2023, the Cybercrimes Act, and the 2025 National AI Strategy.

Kenya: The Public Procurement and Asset Disposal Act 2015 (revised in 2016 and 2022), administered by the Public Procurement Regulatory Authority, requires entities to adopt ICT in the procurement process and establishes a Review Board and judicial review for bidders. None of this governs acquired AI. Notably, Kenya's procurement *regulations* already require tender documents to include a plan for technology and knowledge transfer to Kenyan citizens, not an AI-specific one. However, a concrete precedent shows that African procurement law can carry capacity-building conditions of the kind proposed in this paper's Clause 9. Kenya's relatively advanced national AI strategy is not reflected in its procurement law: strong strategy, silent procurement.

Rwanda: Public procurement is governed by Law N°62/2018, administered by the Rwanda Public Procurement Authority through the UMUCYO e-procurement system, with an Independent Review Panel

for bidder challenges. The law includes local-preference provisions (Articles 79–82), and Rwanda approved a Sustainable Public Procurement Policy Framework in October 2024, demonstrating that its procurement regime already pursues non-price policy goals (e.g., local content and sustainability). It pursues none of the following for AI safety: impact assessment, bias mitigation, explainability, contestability, or anti-lock-in requirements governing acquired AI.

4.2 The model AI procurement standard

Our second contribution is a model standard for insertion into existing national procurement processes as conditions on the acquisition of AI systems, scaled to risk. Clauses marked [**Africa-specific**] address the four named deployment constraints and are the core adaptation.

Clause 1 - Risk classification (gateway). Before tender, the procuring entity classifies the system by risk tier: *Tier 1 (Low)* - internal tools, no direct citizen impact (minimal requirements); *Tier 2 (Moderate)* - citizen-facing but non-determinative (Clauses 3, 4, 6, 7, 9); *Tier 3 (High)* - systems that determine access to rights, benefits, or services, or affect health or liberty (all clauses). All later requirements scale to the assigned tier, keeping the standard proportionate.

Clause 2 - Algorithmic impact assessment. For Tier 3, a documented assessment of intended use, affected populations, foreseeable harms, failure modes, and mitigations is a published condition of award.

Clause 3 - Local-context performance validation [Africa-specific**].** Vendors demonstrate performance on data representative of the actual deployment population, including validated performance in relevant official and local languages. Vendor benchmarks obtained from foreign data are not accepted as evidence of fitness. *Rationale: foreign-trained, English-centric systems degrade in African languages, names, and contexts; procurement that accepts foreign benchmarks imports that fail in public services.*

Clause 4 - Transparency and explainability. Vendors provide documentation sufficient for oversight (at a minimum, a model card that includes training data provenance, known limitations, and tested performance). For Tier 3, the capacity to give an intelligible explanation of an individual decision on request.

Clause 5 - Human oversight and contestability. For Tier 3, a qualified human retains override authority, and the entity establishes a mechanism for citizens to contest an automated or AI-assisted decision and obtain human review. No citizen is subject to a solely automated determination of a right or benefit without appeal.

Clause 6 - Data governance. Compliance with the national data-protection regime [name per country], documented provenance, lawful basis, and security controls; equivalent contractual obligations where no adequate regime applies.

Clause 7 - Vendor accountability and audit rights. Warranties of fitness for the stated purpose, allocation of liability for harms from error or misrepresentation, and a reserved right to audit performance and compliance through the contract term.

Clause 8 - Anti-lock-in and data sovereignty [Africa-specific**].** Guaranteed ownership of, or unrestricted access to, the entity's own input/output data in portable form, with specified exit and transition terms to prevent single-foreign-vendor dependency. *Rationale: services that cannot leave a foreign vendor without losing their own data are not sovereign, directly addressing data-colonial dependency.*

Clause 9 - Capacity transfer [Africa-specific**].** For Tier 2–3, knowledge-transfer obligations (documentation, staff training, appropriate local technical participation) so the entity builds rather than

erodes internal oversight capacity. *Rationale: addresses the skills-gap constraint central to African AI policy.*

Clause 10 - Ongoing monitoring. Post-deployment monitoring against the metrics set at procurement, with periodic reassessment for Tier 3 and a trigger for re-evaluation in response to material performance or context changes. Compliance is a lifecycle obligation, not a point-of-sale event.

5. Discussion and Limitations

5.1 Implications

The matrix shows that the safety floor for public-sector AI in the sampled states is, on paper, near zero - yet the lever to raise it already exists and can be operated by a single procurement authority without new primary legislation. This reframes a problem usually posed as "Africa needs AI laws" into one that is actionable now. The standard's two adoption routes are complementary: a national route, where a procurement authority adopts the clauses as guidance or regulation under existing powers (the fastest path); and a continental route, where the AU/RECs harmonise the standard as the Strategy already calls for, giving smaller states a shared baseline and vendors a single predictable standard rather than 54 divergent ones. Risk-tiering (Clause 1) is what makes either route adoptable: light-touch, where the stakes are low, or stringent, where harm is real, addressing the fear that safety requirements stifle innovation or exclude local vendors.

5.2 Limitations

This study assessed what frameworks require on paper, not enforcement, and African procurement systems face documented challenges of capacity, corruption, and political interference that a standard alone cannot solve (OECD, 2026). The country sample is small and selected for documentation and contrast rather than representativeness; findings should not be generalised to all 54 states without further work. Procurement is one lever with blind spots: it governs acquisition, not systems built in-house or accessed informally, and not private-sector or consumer harms. The standard is a template that requires legal adaptation in each jurisdiction. As a weekend desk study, it has not been validated with procurement officials or vendors. **Key assumption:** that procurement conditions translate into deployed-system behaviour via vendor compliance - if enforcement is absent, the standard yields paper compliance without real safety, which is the central risk below.

5.3 Dual-Use Considerations.

A standard built to raise the safety floor can be turned against its purpose.

(i) A gap map is also a roadmap, documenting exactly which safeguards each state does not require, hands a bad-faith vendor or official, a guide to what they can get away with; we judge the net effect positive but pair the analysis with advocacy to close the gaps rather than merely enumerate them (see Code and Data re: note).

(ii) Risk-tier gaming: The gateway incentivises mis-classifying high-risk systems as low-risk; mitigation is independent, reviewable classification defaulting to the higher tier when contested.

(iii) Safety-washing: Superficial adoption cited for legitimacy without enforcement produces false assurance; without audit and monitoring (Clauses 7, 10), the standard becomes reputational cover.

(iv) Protectionism/incumbency: Compliance costs could exclude smaller or local vendors or entrench large foreign incumbents; the proportionality design limits but does not eliminate this.

5.4 Future Work.

I intend to do the following:

- Validate the clauses with procurement practitioners and through at least one pilot tender; expand the country sample to achieve regional representativeness using the Africa AI Policy Tool.
- Build standardised assessment templates and a model algorithmic-impact-assessment form to lower the capacity burden on procurement officers.
- Specify enforcement and audit mechanisms in detail, and test whether local-context validation (Clause 3) can be operationalised against existing African-language evaluation benchmarks (e.g. IrokoBench, AfroBench, UbuntuGuard).

6. Conclusion

Public-sector AI is being deployed across Africa faster than it is being governed, and the dominant policy response, comprehensive AI legislation, is too slow for the deployment-side risks already materialising. This paper showed, across four national frameworks, that procurement law currently imposes no AI-safety requirements at all - 28 of 36 safeguard cells were absent, and the remainder only indirect, leaving the state's largest channel for AI acquisition ungoverned. It then offered a concrete, adoptable remedy: a risk-tiered model procurement standard, adapted to the four constraints that define African deployment risk, that a national authority can adopt now, and the AU can harmonise continentally.

The broader point for AI safety is that governance leverage in lower-resource settings often lies not in new institutions but in repurposing existing ones. Procurement is one such lever, and using it turns AI safety from something Africa must wait for into something it can buy.

6.1 Code and Data

This is a policy analysis; it involves no code, models, or datasets. The primary sources are public legal and policy documents (listed in References).

Note: The gap analysis identifies, by design, where procurement frameworks lack AI-safety requirements; this is disclosed to motivate reform rather than to enable evasion, and the Dual-Use Considerations (5) address the residual risk.

- Code repository: N/A
- Data/Datasets: Public legal sources (see References)

References

African Union. (2024). *Continental artificial intelligence strategy*.

https://au.int/sites/default/files/documents/44004-doc-EN-Continental_AI_Strategy_July_2024.pdf

Alan Turing Institute & OHCHR. African AI and equality toolbox.
https://africa.aiequalitytoolbox.com/wp-content/uploads/2025/09/AI-EQ-Africa-Toolbox_Sept-2025-2.pdf

Data Protection Act, 2019. (2019). Office of the Data Protection Commissioner (Kenya).
https://www.odpc.go.ke/wp-content/uploads/2024/02/TheDataProtectionAct_No24of2019.pdf

Future of Privacy Forum. (2024, November 18). *The African Union's continental AI strategy: Data protection and governance laws set to play a key role in AI regulation*.
<https://fpf.org/blog/global/the-african-unions-continental-ai-strategy-data-protection-and-governance-laws-set-to-play-a-key-role-in-ai-regulation/>

Law No. 058/2021 of 13/10/2021 relating to the protection of personal data and privacy. (2021). National Cyber Security Authority (Rwanda).
https://www.risa.gov.rw/fileadmin/user_upload/RISA/Publications/2.Laws/Law_relati on of personal data and privacy.pdf

Law No. 62/2018 of 25/08/2018 governing public procurement. (2018). Rwanda Public Procurement Authority. <https://rwandalii.org/akn/rw/act/law/2018/62/eng@2018-09-07>

Mhlanga, D. (2026). Using artificial intelligence to improve governance and public services in Africa. *Frontiers*. <https://doi.org/10.3389/fdata.2026.1835663>

Michalsons. (n.d.). *AI procurement: Guidance for the public sector*.
<https://www.michalsons.com/blog/ai-procurement-guidance-for-the-public-sector/70297>

Naija AI. (2026). *AI in Nigeria: The complete 2026 guide*. <https://naijaai.com.ng/ai-in-nigeria-guide/>

Nigeria Data Protection Act. (2023). Nigeria Data Protection Commission.
https://ndpc.gov.ng/resources/#flipbook-df_2442/3/

OECD. (2026). *AI governance in Africa (2026)*.
https://www.oecd.org/content/dam/oecd/en/publications/reports/2026/04/oecd-artificial-intelligence-case-studies_71322177/strengthening-ai-governance-in-africa_812bb149/1ff55135-en.pdf

Protection of Personal Information Act 4 of 2013. (2013). Department of Justice and Constitutional Development (South Africa). <https://www.justice.gov.za/legislation/acts/2013-004.pdf>

Public Procurement Act 28 of 2024. (2024). Government of South Africa.
<https://www.gov.za/documents/acts/public-procurement-act-28-2024-english-setswana-23-jul-2024>

Public Procurement Act No. 14 of 2007. (2007). Bureau of Public Procurement (Nigeria).
<https://www.bpp.gov.ng/wp-content/uploads/2019/01/Public-Procurement-Act-2007pdf.pdf>

Public Procurement and Asset Disposal Act No. 33 of 2015. (2015). Public Procurement Regulatory Authority (Kenya). <https://new.kenyalaw.org/akn/ke/act/2015/33/eng@2022-12-31>

Segun, S. T., Adams, R., Florido, A., Timcke, S., Shock, J., Junck, L., Adeleke, F., Grossman, N., Alayande, A., Kponyo, J. J., Smith, M., Fosu, D. M., Tetteh, P. D., Arthur, J., Kasaon, S., Ayodele, O., Badolo, L., Plantinga, P., Gastrow, M., ... Bissyandé, T. F. (2025). *Toward an African agenda for AI safety*. arXiv. <https://doi.org/10.48550/arxiv.2508.13179>

Tomei, P. M., Jain, R., & Franklin, M. (2025). *AI governance through markets*. arXiv. <https://doi.org/10.48550/arxiv.2501.17755>

LLM Usage Statement

Large language model assistance (specifically Claude) was utilised for project conceptualisation, report structuring, and preliminary drafting of several sections, including the framework, methodology, model standard, and discussion. The comparative gap analysis detailed in Section 4.1 was independently executed and subsequently verified by the author against primary legal authorities. The author has independently verified all factual assertions, citations, and research findings, and the final manuscript has been revised to reflect the author's personal academic style.