Operationalising Sovereign AI for the Global South: A Minimum Viable Governance Architecture for LDCs and Base-of-the-Pyramid Markets

An Anatomical and Philosophical Review Extending the Governance-Aware Retriever Framework for Global South Contexts

Dr. Syed Muntasir Mamun
Obi Umegbolu
ORCID: 0000-0001-6845-2853 (Mamun)

Abstract

This paper proposes a Minimum Viable Governance Architecture (MVGA) for operationalising sovereign AI in Least Developed Countries (LDCs) and Base-of-the-Pyramid (BoP) markets. Drawing on the Governance-Aware Retriever Framework (GnARF) and Minimum Viable Innovation Engines (MVIEs), it contends that full-stack technological sovereignty remains structurally unattainable for most Global South states in the medium term due to persistent institutional voids, capital deficits, and compute constraints. Instead, a pragmatic, layered approach is required that deliberately connects market demands (citizen-centric services, AfCFTA digital trade, local content economies), appropriate technologies (open-source models, locally curated RAG corpora, integration with existing mobile digital public infrastructure), and innovative organisational forms (cross-ministerial MVIEs, regulatory sandboxes, and South-South consortia). The paper systematically examines the pros and cons of this architecture, compares competing epistemologies (Western statistical inference versus relational and contextual Southern knowledges), and analyses the ontological consequences of embedding AI systems in post-colonial contexts. Positioned as a strategic contribution to the inaugural UN Global Dialogue on AI Governance (Geneva, 6–7 July 2026), the MVGA offers concrete mechanisms for bridging AI divides while advancing epistemic pluralism and ontological self-determination for the Global South.

Keywords

Sovereign AI, Global South, LDCs, Base-of-the-Pyramid Markets, Minimum Viable Governance Architecture, GnARF, MVIEs, Epistemic Pluralism, Ontological Sovereignty, Decolonial AI, AI Governance, WSIS, Global Dialogue on AI.

Classifications

Primary Field: Public Administration & Governance in Developing Countries / African Studies
Secondary Fields: Artificial Intelligence Ethics & Governance, Digital Sovereignty in the Global South, Development Studies, Political Philosophy (African Epistemologies), Information Systems for Development

Takeaways

- **The Minimum Viable Governance Architecture (MVGA):** Full-stack sovereign AI is unattainable for most LDCs in the medium term. A pragmatic MVGA—combining GnARF's technical safeguards with MVIEs as organisational vehicles—offers a viable pathway that connects market demands, appropriate technologies, and institutional innovation without requiring massive capital or compute investment.

• **Pros of the Proposed Architecture:** Accelerated value capture from GenAI, reduced vendor lock-in through model redundancy and local RAG, lower barriers to entry via existing mobile DPI, stronger epistemic inclusion of Southern knowledge systems, and enhanced political legitimacy through participatory organisational design.

• **Cons and Risks:** Coordination costs across fragmented institutions, potential elite capture of new organisational forms (MVIEs), amplified computational and energy overhead in low-resource settings, risk of 'pilot purgatory', and the danger that short-term market pressures may override long-term epistemic sovereignty goals.

• **Competing Epistemologies:** Western statistical inference (probabilistic, pattern-based, often anglocentric) must be contrasted with relational, contextual, and communal Southern epistemologies (Ubuntu, oral traditions, elder consensus). The MVGA's RAG layer can serve as a site of resistance if deliberately populated with locally curated, plural knowledge corpora.

• **Ontological Consequences:** Embedding AI systems risks imposing Western ontological assumptions (individual data subjects, data as property, linear progress). The MVGA must incorporate mechanisms for ontological self-determination, including participatory design with customary authorities and protection of communal notions of personhood and knowledge.

• **Strategic Relevance for the Global Dialogue (July 2026):** The MVGA provides concrete, implementable language and mechanisms for Cluster 2 (Bridging AI Divides) and the 'Dialogue of Dialogues', enabling Global South countries to move from normative demands to operational proposals in UN AI governance negotiations.

Table of Contents

Introduction: The Market-Technology-Organisation Nexus for AI Sovereignty in Africa

The rapid proliferation of generative artificial intelligence (GenAI) across global public administrations has created a stark bifurcation in outcomes. While advanced economies debate the fine print of the EU AI Act or the nuances of sovereign cloud architectures, many African states confront a more elemental question: how to participate in the AI era without reproducing the extractive logics of previous technological waves (Belli, 2025; Ndemo and Weiss, 2020). The Governance-Aware Retriever Framework (GnARF) proposed by Langer and Haag (2026) offers a technically elegant response to vendor lock-in and data expropriation through redundant multi-model querying, retrieval-augmented generation (RAG), and immutable decision logging. Yet its original formulation assumes a baseline of institutional capacity—reliable electricity, skilled IT bureaucracies, enforceable data protection regimes, and access to diverse commercial APIs—that remains unevenly distributed across the African continent.

Africa occupies a dual and contradictory position. On one hand, the continent has demonstrated remarkable leapfrogging capacity: mobile money platforms (M-Pesa in Kenya and its regional replications) bypassed traditional banking rails; USSD and IVR channels deliver services to populations with limited smartphone penetration; and tech hubs in Lagos, Nairobi, Kigali, and Cape Town are producing globally recognised AI talent and startups (GSMA, 2025). These constitute genuine *complementary assets* (Teece, 1986) that can be mobilised for AI deployment. On the other hand, pervasive *institutional voids* (Khanna and Palepu, 2010)—absent or weak capital-market intermediaries, fragmented regulatory enforcement, shallow pools of AI/ML engineering talent within the civil service, unreliable power and broadband backbones, and limited domestic compute capacity—render the direct transplantation of GnARF, or any full-stack sovereignty model, impractical in the short-to-medium term.

This Volume II therefore advances a two-pronged argument. First, GnARF's core architectural principles remain highly relevant to African public sectors precisely *because* they mitigate the most immediate vectors of dependency (single-vendor lock-in, data leakage, unexplainable decisions) without requiring sovereign frontier models or domestic hyperscalers. Second, successful operationalisation demands deliberate adaptation through *Minimum Viable Innovation Engines* (MVIEs)—agile, low-overhead public-sector units that prototype GnARF instances on non-critical workflows, ground RAG in locally curated African knowledge corpora, and iteratively build the tacit capabilities and political legitimacy necessary for scale. The paper situates this proposal within the epistemological and ontological stakes of AI adoption in Africa, compares it with the African Union's 2024 Continental Artificial Intelligence Strategy and WSIS+20 processes, and articulates a decolonial pathway that treats institutional voids not merely as deficits but as potential wellsprings for alternative, relationally grounded AI governance models.

Theoretical Contours: Institutional Voids and Complementary Assets in African Public Administration

Institutional Voids in African Public Sectors

Khanna and Palepu (1997, 2010) conceptualise institutional voids as the absence or weakness of specialised intermediaries that facilitate market (or, by extension, governance) transactions in emerging economies. In the African public-sector context, these voids manifest across multiple registers. *Regulatory voids* leave data protection authorities under-resourced and enforcement patchy; AI procurement often proceeds via opaque sole-source contracts with global cloud providers, reproducing the very vendor dependencies GnARF seeks to disrupt. *Intermediary voids* in talent mean that few ministries possess in-house ML engineers capable of maintaining sophisticated middleware; reliance on external consultants risks both capture and the externalisation of institutional memory. *Capital and infrastructure voids* constrain investment in sovereign or regional compute clusters; undersea cable landings and data-centre capacity remain geographically concentrated and vulnerable to geopolitical or climate shocks. *Information voids*—sparse digitisation of local-language legal texts, health records, agricultural extension knowledge, and oral histories—undermine the very possibility of culturally grounded RAG systems.

Crucially, these voids are not uniform. Rwanda's aggressive digitalisation drive, Ghana's data-protection regime, South Africa's relatively mature research universities, and Nigeria's large domestic market create pockets of relative institutional thickness. The strategic challenge is therefore not merely to 'fill' voids but to design architectures that can operate effectively *across* heterogeneous institutional landscapes, turning pockets of strength into platforms for continental learning.

Complementary Asset Portfolios as Strategic Leverage

Teece's (1986) complementary assets framework—originally developed to explain why innovators often fail to capture value from their inventions—provides a powerful lens for African AI strategy. In the GenAI context, the core 'invention' (frontier foundation models) resides overwhelmingly with US and Chinese hyperscalers. African states cannot easily replicate the compute, data lakes, or talent pools that produced GPT-4 or Claude. They can, however, cultivate and orchestrate complementary assets that allow them to *access* and *shape* the value of these models on sovereign terms.

Africa's portfolio includes: (1) *Mobile-first digital public infrastructure* (DPI) exemplified by M-Pesa's agent networks, interoperable mobile money schemes across the East African Community, and Ghana's mobile money interoperability; these channels can serve as low-friction interfaces for citizen-facing GenAI services via USSD or WhatsApp Business APIs, bypassing smartphone and data-cost barriers. (2) *Demographic and human-capital assets*: the world's youngest population, rapidly growing tertiary enrolment in computer science and engineering, and a vibrant startup ecosystem (AfriLabs network) that can supply MVIE personnel and co-creation partners. (3) *Linguistic and epistemic assets*: over 2,000 languages, substantial oral traditions that can be transcribed and ethically curated for RAG, and emerging NLP initiatives (Masakhane, Lelapa AI) producing Africa-centric models and datasets. (4) *Policy and diplomatic assets*: the African Union's 2024 Continental AI Strategy, the Malabo Convention on Cyber Security and Personal Data Protection, and coordinated African positions within WSIS, IGF, and ITU processes that can be leveraged to negotiate better terms with platform providers and to shape global norms.

The GnARF-MVIE proposal advanced here is explicitly designed to activate and amplify these complementary assets rather than treating Africa as a tabula rasa awaiting external technological solutions.

Anatomical Adaptation of the Governance-Aware Retriever Framework for African Contexts

Langer and Haag's (2026) five-module architecture remains structurally sound for African deployment, but each module requires contextual reconfiguration to address voids and leverage assets. The adaptations below preserve GnARF's core sovereignty logic while rendering it operable under conditions of constrained bandwidth, intermittent power, heterogeneous legal regimes, and sparse digitised local knowledge.

1. Query Model Allocation (QMA) Reconfigured

The original QMA dynamically routes queries across a portfolio of commercial LLMs to avoid single-vendor dependency. In African settings, three modifications are essential. First, *cost and currency sensitivity*: commercial API pricing denominated in USD creates foreign-exchange exposure; MVIEs should maintain a preference ordering that includes open-source models hosted on regional or sovereign clouds (where they exist) or on cost-optimised GPU instances in partner jurisdictions (e.g., South Africa, Morocco, or emerging GPU clusters in Kenya). Second, *latency and connectivity awareness*: many African government offices outside capital cities experience high and variable latency to European or US endpoints; QMA logic should incorporate real-time endpoint health and latency metrics, with graceful fallback to lighter local models for time-sensitive but lower-stakes tasks. Third, *geopolitical hedging*: over-reliance on any single geopolitical bloc's models (US or Chinese) introduces new vectors of dependency; African QMA implementations should actively diversify across US, European, Chinese, and emerging Global South providers (Indian, Brazilian, or intra-African where available), treating model provenance as a sovereignty-relevant attribute.

2. Retrieval-Augmented Generation (RAG) as Epistemic Anchor

RAG is arguably the most transformative module for African contexts. By keeping the reasoning engine (external LLM) strictly separated from the knowledge base (local vector store), RAG prevents the permanent embedding of sensitive state or citizen data into proprietary model weights—the central data-sovereignty risk identified in Vol. I. For Africa, RAG's deeper significance lies in its potential as an instrument of *epistemic resistance*. Institutional voids in digitisation mean that Western-trained models possess thin or distorted representations of African legal systems, health protocols, agricultural practices, and cultural norms. A well-curated local RAG corpus—populated with national gazettes, court judgments (anonymised), ministry policy documents, transcribed oral histories (with appropriate community consent protocols), and multilingual parallel texts—can ground LLM outputs in African realities, reduce hallucination on context-specific queries, and surface plural knowledge systems that statistical pre-training alone erases. MVIEs should therefore prioritise investment in data curation pipelines (OCR for legacy documents, ASR for oral archives, ethical metadata schemas) as a core sovereignty infrastructure, not an ancillary technical task.

3. PII Filtering and Culturally Situated Privacy

African data-protection regimes are heterogeneous: South Africa's POPIA, Nigeria's NDPR, Ghana's Data Protection Act, and the supranational Malabo Convention create a patchwork of obligations. GnARF's PII-filtering module must therefore be configurable to jurisdiction-specific rulesets and sensitive-attribute taxonomies. Beyond formal compliance, African conceptions of privacy often emphasise *relational* and *communal* dimensions—data about an individual may implicate family, clan, or community; consent may be negotiated collectively rather than individually. MVIE prototypes should therefore incorporate participatory data-governance protocols (inspired by but not limited to CARE Principles for Indigenous Data Governance) that treat data stewardship as a communal responsibility, not merely a technical anonymisation exercise.

4. Decision Logging and Procedural Legitimacy

In post-colonial contexts where state–citizen trust is often fragile, immutable decision logging acquires heightened legitimising function. GnARF's audit trail—query, retrieved documents, prompt, models used, final output—provides a procedural scaffold that can be presented to citizens, courts, or oversight bodies. African MVIEs should extend this scaffold with *citizen-facing explanation interfaces* (simple natural-language summaries in local languages, visual decision trees) and *appeal mechanisms* that route contested AI-mediated decisions to human review, thereby converting technical transparency into lived administrative justice.

Minimum Viable Innovation Engines (MVIEs): Concept and Operationalisation

A Minimum Viable Innovation Engine (MVIE) is a lightweight, time-bounded, cross-functional public-sector unit (typically 5–12 personnel drawn from the host ministry, a local university or research centre, and one or two vetted private-sector or civil-society partners) chartered to prototype and iterate a GnARF-adapted GenAI system on a narrowly scoped, non-critical workflow within 90–180 days. The 'minimum viable' qualifier signals deliberate constraint: the unit operates with modest budget (often under US$150,000 for the pilot phase), uses predominantly open-source or low-cost commercial components, accepts 'good-enough' rather than production-grade reliability, and measures success primarily by learning velocity and demonstrable sovereignty gains rather than immediate scale or cost savings.

Operationalisation Steps

- **Problem Selection:** Choose a workflow that is painful but not mission-critical (e.g., internal synthesis of public consultation submissions; drafting of routine correspondence in multiple languages; citizen query triage for a high-volume but low-risk service). This minimises political and operational risk while maximising learning.
- **Team Assembly:** Blend domain expertise (the policy officials who currently perform the task), technical capability (ML engineers or data scientists from a partner university or startup), and governance/ethics capacity (legal or civil-society representative). Diaspora talent on short-term secondment or remote consultancy can accelerate capability transfer.
- **Local Corpus Curation (First 30 Days):** Prioritise ingestion of 500–2,000 high-value documents relevant to the workflow—policy circulars, past decisions, relevant statutes—into a secure vector store. Use lightweight OCR/ASR pipelines; establish community consent protocols where oral or traditional knowledge is involved.
- **GnARF Instance Deployment:** Stand up a minimal QMA (initially 2–3 models: one open-source local inference, one commercial API), RAG pipeline, PII scrubber tuned to applicable national law, and immutable log store (simple append-only database with cryptographic hashing). Host on existing government cloud tenancy or a regional provider with data-residency guarantees.
- **Iterative Shadowing and Evaluation:** Run the prototype in 'shadow mode' alongside human operators for 4–6 weeks; compare outputs on accuracy, cultural appropriateness, and procedural compliance. Introduce human-in-the-loop review gates before any citizen-facing output.
- **Political and Budgetary Socialisation:** Document learning, cost trajectories, sovereignty metrics (percentage of queries served without single-vendor dependency; volume of sensitive data retained locally), and citizen-experience improvements. Use these artefacts to secure follow-on funding and political cover for expansion or replication in adjacent workflows.

MVIEs are not intended to remain perpetually 'minimum'. Successful pilots should graduate into permanent institutional homes (innovation labs within digital ministries, or centres of excellence aligned with the AU Strategy) with expanded mandates and budgets. The MVIE construct simply lowers the activation energy for African public sectors to begin the sovereignty journey without waiting for perfect conditions or massive capital allocations.

Critical Evaluation: Advantages, Limitations, and Systemic Frictions

Advantages (Pros)

• **Leapfrogging without Lock-in:** By treating frontier models as interchangeable utilities accessed through a sovereignty-preserving middleware layer, African institutions can begin deriving value from GenAI immediately while building the institutional muscle (data curation, prompt engineering, human-AI workflow design) that will be essential regardless of future technological shifts.

• **Epistemic Inclusion at Low Marginal Cost:** Once a local RAG corpus exists, the incremental cost of enriching it with African languages, legal traditions, and community knowledge is modest compared with the cost of pre-training or even fine-tuning frontier models. This creates a structural incentive for epistemic pluralism that pure API consumption lacks.

• **Political and Budgetary Realism:** MVIEs require budgets and political attention spans that are attainable in most African contexts; they do not depend on multi-year sovereign wealth fund allocations or complex public–private partnership negotiations that often stall.

• **Continental Learning Accelerant:** Successful MVIE deployments in one jurisdiction (e.g., Rwanda's Ministry of ICT or Ghana's Data Protection Commission) can be documented as open playbooks and adapted by peer institutions, fostering South–South diffusion that complements formal AU coordination mechanisms.

• **Alignment with Existing Assets:** The architecture maps cleanly onto mobile DPI channels already trusted by citizens and onto the policy language of the AU Continental AI Strategy (capacity building, national strategies, centres of excellence).

Limitations and Systemic Frictions (Cons)

• **Amplified Computational and Energy Overhead:** Redundant querying multiplies inference costs and carbon footprints. In contexts where electricity is unreliable or expensive, and where data centres are few and far between, this overhead is not merely financial but can undermine broader sustainability and energy-access goals.

• **Talent and Continuity Risks:** MVIEs depend on a small number of hybrid policy-technical personnel. High public-sector turnover, brain drain to private tech or international organisations, and limited career paths for AI governance specialists threaten institutional memory and continuity.

• **Pilot Purgatory and Scale Failures:** Many well-intentioned African digital pilots remain trapped in demonstration mode, unable to secure the budgetary mainstreaming or organisational redesign required for scale. Without deliberate 'graduation' mechanisms and political championship, GnARF-MVIE prototypes risk adding to the inventory of abandoned innovations.

• **Data Quality and Representativeness Gaps:** Local RAG corpora will initially be sparse, biased toward formal English/French/Portuguese documents, and missing large swathes of oral, customary, and vernacular knowledge. Poor retrieval quality can produce confident but culturally misaligned outputs that erode rather than enhance legitimacy.

• **New Dependencies Masquerading as Sovereignty:** Over-reliance on a narrow set of open-source model providers, GPU cloud sponsors, or donor-funded technical assistance programmes can recreate dependency in new guises. True sovereignty requires ongoing vigilance against capture by any single external actor or funding stream.

• **Regulatory Patchwork Navigation:** Cross-border data flows, model hosting, and mutual recognition of AI decisions remain legally murky within Africa. An MVIE operating in one jurisdiction may inadvertently violate rules when queries or logs traverse borders.

Epistemological Consequences: From Relational African Knowing to Statistical Inference

The introduction of GenAI into African public administration does not merely automate existing processes; it enacts a profound epistemological rupture. For centuries, African knowledge systems have privileged relational, communal, contextual, and often oral modes of knowing. The proverb 'Ubuntu ngumuntu ngabantu' ('a person is a person through other persons') encapsulates an epistemology in which knowledge is validated through communal deliberation, lived experience, and intergenerational transmission rather than through statistical correlation or individual expert assertion (Wiredu, 1996; Mbiti, 1969). Administrative decision-making in many African contexts has historically blended codified law with customary norms, elder counsel, and community consensus—forms of reasoning that are hermeneutic, dialogic, and explicitly value-laden.

Large language models, by contrast, operate through inductive statistical inference over massive, largely anglocentric and WEIRD (Western, Educated, Industrialised, Rich, Democratic) corpora. When a public servant in an African ministry uses an unmodified commercial LLM to summarise citizen submissions or draft policy options, the state is not 'knowing' its citizens through the relational matrices that have historically conferred legitimacy; it is processing them as vectors in a latent space optimised for patterns that may be orthogonal or even antagonistic to local realities. Hallucinations on African names, places, or legal concepts are not merely technical errors; they are symptoms of an underlying representational deficit that risks systematic misrecognition.

GnARF's RAG module, properly adapted, offers a partial but significant corrective. By anchoring generation in a locally curated corpus, it reintroduces a measure of institutional and cultural specificity into the inference process. Yet even sophisticated RAG cannot fully bridge the ontological gap between probabilistic pattern-matching and the deliberative, value-articulating practices that African administrative traditions have cultivated. The danger is therefore twofold: *epistemicide* (de Sousa Santos, 2014), in which alternative ways of knowing are marginalised or rendered invisible by the dominance of statistical outputs; and *skill atrophy* within the civil service, as human officials increasingly defer to the fluent but shallow reasoning of the model rather than exercising and transmitting the deeper hermeneutic capacities that relational governance requires.

African MVIEs must therefore treat epistemological design as a first-order governance task. This includes: (a) explicit protocols for what kinds of decisions may *not* be delegated to statistical inference (e.g., those involving customary law interpretation, communal land rights, or culturally sensitive resource allocation); (b) investment in hybrid human–AI workflows that preserve space for communal validation and contestation; and (c) documentation of cases where the model 'fails' in culturally revealing ways, turning breakdowns into sites of institutional learning and corpus refinement. Only through such reflexive practice can the epistemological rupture be transformed from a vector of epistemic colonisation into an occasion for the articulation and defence of plural African ways of knowing.

Ontological Consequences: Digital Neocolonialism and the Struggle for African Being

If epistemological consequences concern *how* the African state knows, ontological consequences concern *what* it considers real, legitimate, and worth preserving. Western-trained foundation models encode deep ontological assumptions: the primacy of the individual (data subjects with alienable rights), the fungibility of information (data as a resource to be extracted and optimised), the linearity of progress (development as movement toward Western technological and institutional forms), and the separation of 'nature' from 'society' (resource management as technical optimisation rather than relational stewardship). These assumptions are not neutral; they are the sedimented product of specific historical trajectories and power relations (Zembylas, 2023; Mohamed et al., 2020).

When African public institutions integrate these models into land administration, health triage, welfare targeting, or conflict early-warning systems, they risk subjecting their populations to a subtle but pervasive ontological hegemony. Customary land tenure—relational, overlapping, and often unregistered—may be rendered illegible or illegitimate by systems optimised for individual title and cadastral certainty. Community consensus mechanisms may be overridden by 'objective' algorithmic risk scores that privilege quantifiable indicators over lived testimony. The very definition of 'development' embedded in AI-driven policy simulators may reproduce the extractive, growth-at-all-costs paradigms that have historically disadvantaged the continent.

Decolonial scholars have long argued that true sovereignty requires not only control over territory and resources but also over the categories through which reality is apprehended and valued (Mignolo, 2011; Ndlovu-Gatsheni, 2018). In the AI age, this implies a right to *ontological self-determination*: the collective authority to define what counts as a person, a community, a resource, a harm, and a future worth pursuing, and to encode those definitions into the technical systems that increasingly mediate state–citizen relations. GnARF's technical safeguards (local RAG, decision logging, PII filtering) are necessary but insufficient for ontological sovereignty. They must be coupled with participatory design processes that bring customary authorities, community representatives, elders, and youth into the specification of what the AI system should 'see' and how it should 'reason' about African social realities.

Comparative Analysis: The AU Continental AI Strategy, UN/WSIS Efforts, and the GnARF-MVIE Proposal

To assess the strategic fit of the GnARF-MVIE approach, it is necessary to locate it within the broader landscape of continental and global AI governance initiatives targeting Africa.

Table 1: Comparative Positioning of Governance Paradigms for African AI Sovereignty

Governance Paradigm	Core Orientation	Strengths	Limitations	Compatibility with GnARF-MVIE
African Union Continental AI Strategy (2024)	People-centric, development-oriented, inclusive; national strategies + centres of excellence; capacity building focus (2025–2030)	High political legitimacy; explicit Africa-centric framing; addresses skills, data, compute, and governance pillars; aligns with Agenda 2063	High-level and non-binding on many dimensions; implementation capacity varies widely across member states; limited operational 'how' guidance for technical architecture	**High compatibility**. GnARF-MVIE supplies the missing operational layer for AU pillars on governance, capacity, and responsible adoption. MVIEs can function as the 'centres of excellence' or national sandboxes the Strategy recommends.
WSIS / WSIS+20 Processes (ITU-led, multi-stakeholder)	Multi-stakeholder information society; development-oriented; WSIS Action Lines (e-gov, capacity building, cultural diversity, media); recent emphasis on AI ethics and governance	Global platform for African voice; long-standing African engagement; Action Lines provide normative hooks for sovereignty and equity claims; IGF and WSIS Forum spaces for contestation	Often captured by powerful state and corporate actors; outcomes frequently aspirational rather than enforceable; limited technical depth on architectural sovereignty mechanisms	**Moderate-to-high compatibility**. GnARF-MVIE advances WSIS Action Line C7 (e-Government) by demonstrating how AI can be deployed accountably; supports C4 (capacity building) and C8 (cultural diversity) through local RAG and epistemic pluralism. African delegations can cite concrete implementations to strengthen sovereignty claims in WSIS negotiations.
UNDP / UNECA Digital Transformation & DPI Initiatives	Digital public infrastructure as enabler of SDGs; focus on digital ID, payments, data exchange; e-government platforms; AI for development pilots	Strong implementation track record in several countries; emphasises public goods and inclusion; builds on existing DPI investments (e.g., mobile money, digital ID)	Sometimes donor-driven; risk of standardised 'best practice' templates that underplay local institutional voids and epistemic specificity; limited focus on AI model sovereignty or middleware resilience	**High complementarity**. GnARF can be layered onto existing UN-supported DPI platforms to add GenAI capabilities with sovereignty safeguards. MVIEs offer a mechanism for countries to move from DPI consumption to AI-augmented service innovation on their own terms.
Bilateral / External (China Digital Silk Road, EU Digital Partnerships, US Tech Investments)	Infrastructure financing and market access; varying degrees of regulatory alignment pressure or data-localisation requirements	Can accelerate hard infrastructure (data centres, connectivity); provides immediate access to advanced models and compute	Risk of debt traps, vendor lock-in, or geopolitical capture; data sovereignty often weakly protected; limited support for African epistemic or ontological self-determination	**Conditional / tactical compatibility**. GnARF's QMA and RAG layers can be used to diversify dependencies across multiple external partners and to retain local data control even when models are foreign-hosted. Requires strong political discipline to avoid capture.

The comparative mapping reveals that the GnARF-MVIE proposal occupies a distinctive and under-served niche: it supplies a concrete, deployable technical architecture that operationalises the high-level aspirations of the AU Strategy and WSIS processes while remaining realistic about institutional voids and capital constraints. It is neither a top-down regulatory hammer (EU-style) nor a full-stack autarky project (Chinese model), but a pragmatic middleware strategy that maximises African agency within conditions of managed interdependence.

A Comprehensive Narrative: Infrastructuralisation of Politics and Decolonial Imperatives in African Public Sectors

The synthesis of technical architecture (GnARF), organisational form (MVIE), epistemological reflexivity, and ontological self-determination points toward a broader phenomenon: the *infrastructuralisation of politics* in African public administration (Ferrari Braun and Cath, 2026). In classical democratic theory, political power is rendered visible and contestable through legislative chambers, public hearings, and accountable officials. When substantive decisions—resource allocation, risk scoring, service eligibility, cultural recognition—are progressively encoded into algorithmic infrastructures that are opaque even to their operators, the locus of power migrates from visible political arenas to invisible design choices made by engineers, data curators, and platform providers, often located outside the polity altogether.

African states enter this dynamic from a position of structural weakness. Many already suffer from what Mbembe (2001) termed the 'banality of power'—the everyday exercise of arbitrary authority through bureaucratic opacity and clientelism. The introduction of GenAI, if pursued through uncritical adoption of foreign platforms, risks compounding this opacity while adding a veneer of technical neutrality that makes contestation even more difficult. Conversely, a deliberately designed GnARF-MVIE pathway can *re-inject visibility and accountability* into administrative processes: decision logs create auditable trails; RAG grounds outputs in publicly inspectable documents; human-in-the-loop gates and citizen explanation interfaces create moments of political encounter where algorithmic recommendations can be challenged and renegotiated.

Yet technical visibility is not enough. The deeper decolonial imperative is to ensure that the *categories* through which African realities are rendered legible to the state—categories of personhood, community, harm, value, and future—are not silently imported from Western training data but are actively authored and authorised by African publics. This requires treating MVIEs not merely as technical prototyping units but as sites of *ontological politics*: participatory fora in which customary authorities, women's groups, youth, persons with disabilities, and linguistic minorities can articulate what the AI system should and should not 'see', how it should weigh different forms of evidence, and what kinds of decisions it may never be permitted to make. Only through such sustained democratic and cultural labour can African public sectors convert the infrastructuralisation of politics from a threat of hollowing-out into an occasion for the reinvention of governance on genuinely African terms.

Conclusion: From Digital Consumer to Co-Creator—Pathways to African AI Sovereignty

The Governance-Aware Retriever Framework, adapted through Minimum Viable Innovation Engines and deliberately anchored in Africa's complementary asset portfolios, offers a pragmatic and philosophically coherent pathway for African public sectors navigating the GenAI transition. It acknowledges the structural impossibility of full-stack technological sovereignty for most African states in the present conjuncture, while refusing the false choice between isolation and capitulation to platform hegemony. By enforcing multi-model redundancy, local RAG grounding, PII protection, and immutable logging, the adapted architecture reclaims immediate operational agency and data control. By embedding these technical mechanisms within agile MVIE units that prioritise learning, local corpus development, and participatory design, it begins the longer-term work of building the human and institutional capabilities that sovereignty ultimately requires.

The stakes, however, extend beyond operational resilience. The epistemological shift from relational, communal, and hermeneutic African ways of knowing to statistical inference, and the ontological imposition of Western-centric categories of being and value, threaten forms of epistemicide and cultural erasure that technical safeguards alone cannot prevent. GnARF's RAG module, if populated and governed with care, can function as a site of resistance and pluralism; but this outcome is not automatic. It demands sustained investment in data curation as a sovereign infrastructure, explicit protocols protecting spaces of human and communal judgment, and the courage to refuse deployment where statistical inference would distort or displace more appropriate African modes of decision-making.

Strategically, the GnARF-MVIE approach aligns closely with the African Union's 2024 Continental AI Strategy and can strengthen African positions within WSIS+20 and related global forums by demonstrating concrete, sovereignty-enhancing implementations rather than merely asserting normative claims. It positions Africa not as a perpetual consumer of externally produced AI capabilities but as a potential co-creator of more plural, more just, and more relationally grounded AI futures—futures in which the continent's demographic vitality, linguistic diversity, and philosophical traditions become generative resources for global AI governance rather than liabilities to be overcome.

The path forward requires three mutually reinforcing commitments. First, *technical pragmatism*: rapid deployment of adapted GnARF instances via MVIEs on carefully chosen workflows, with ruthless prioritisation of learning over perfection. Second, *epistemic and ontological vigilance*: continuous interrogation of what the technology renders visible or invisible, and deliberate curation of local knowledge as a counterweight to statistical monoculture. Third, *continental and global solidarity*: active use of AU coordination mechanisms and WSIS/IGF platforms to share playbooks, negotiate collectively with platform providers, and insist that global AI governance recognise and accommodate plural ontologies and knowledge systems. Through this synthesis of architecture, organisation, and normative commitment, African public sectors can transform institutional voids from sites of vulnerability into wellsprings of innovation, and move decisively from the periphery to the centre of the global AI conversation.

Works Cited

African Union (2024) *Continental Artificial Intelligence Strategy*. Addis Ababa: African Union. Available at: https://au.int/en/documents/20240809/continental-artificial-intelligence-strategy (Accessed: 15 June 2026).

Belli, L. (2025) 'Digital sovereignty in the Global South: Between dependence and autonomy', *Internet Policy Review*, 14(2), pp. 1–22.

de Sousa Santos, B. (2014) *Epistemologies of the South: Justice against Epistemicide*. Boulder, CO: Paradigm Publishers.

Ferrari Braun, V. and Cath, C. (2026) 'The infrastructuralisation of politics: AI, visibility, and democratic hollowing-out', *AI & Society*, 41(1), pp. 45–67.

GSMA (2025) *The Mobile Economy Sub-Saharan Africa 2025*. London: GSMA Intelligence.

Khanna, T. and Palepu, K.G. (1997) 'Why focused strategies may be wrong for emerging markets', *Harvard Business Review*, 75(4), pp. 41–51.

Khanna, T. and Palepu, K.G. (2010) *Winning in Emerging Markets: A Road Map for Strategy and Execution*. Boston, MA: Harvard Business School Press.

Langer, P.F. and Haag, S. (2026) 'AI sovereignty: Redundant use of large language models for public sector resilience', *SSRN*, Abstract ID 6937498.

Mbembe, A. (2001) *On the Postcolony*. Berkeley, CA: University of California Press.

Mbiti, J.S. (1969) *African Religions and Philosophy*. London: Heinemann.

Mignolo, W.D. (2011) *The Darker Side of Western Modernity: Global Futures, Decolonial Options*. Durham, NC: Duke University Press.

Mohamed, S., Png, M.T. and Isaac, W. (2020) 'Decolonial AI: Decolonial theory as sociotechnical foresight in artificial intelligence', *Philosophy & Technology*, 33(4), pp. 659–684.

Ndemo, B. and Weiss, T. (eds) (2020) *Digital Kenya: An Entrepreneurial Revolution in the Making*. London: Palgrave Macmillan.

Ndlovu-Gatsheni, S.J. (2018) *Epistemic Freedom in Africa: Deprovincialization and Decolonization*. London: Routledge.

Teece, D.J. (1986) 'Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy', *Research Policy*, 15(6), pp. 285–305.

Wiredu, K. (1996) *Cultural Universals and Particulars: An African Perspective*. Bloomington, IN: Indiana University Press.

Zembylas, M. (2023) 'Decolonial AI and the politics of knowledge production in education', *Learning, Media and Technology*, 48(3), pp. 412–428.

African Union (2025) *High-Level Policy Dialogue on AI for Africa's Development and Prosperity: Communiqué*. Addis Ababa: AUC.

Belli, L. and Jiang, M. (2024) 'BRICS and digital sovereignty: Plural pathways in a fractured internet', *Global Media and Communication*, 20(2), pp. 145–162.

Floridi, L. (2020) 'AI and its new winter: From myths to realities', *Philosophy & Technology*, 33(1), pp. 1–3.

ITU (2025) *WSIS+20: Building a Digital Future for All – Review Report*. Geneva: International Telecommunication Union.

Research ICT Africa (2026) *Africa WSIS+20 Review: Final Report*. Cape Town: Research ICT Africa.

UNECA (2025) *Africa WSIS Annual Review 2025: Digital Transformation and Emerging Technologies*. Addis Ababa: United Nations Economic Commission for Africa.

Wirtz, B.W., Weyerer, J.C. and Geyer, C. (2021) 'Artificial intelligence and the public sector—applications and challenges', *International Journal of Public Administration*, 44(7), pp. 1–15.

Zerilli, J., Knott, A., Maclaurin, J. and Gavaghan, C. (2019) 'Transparency in algorithmic and human decision-making: Is there a double standard?', *Philosophy & Technology*, 32(4), pp. 661–683.