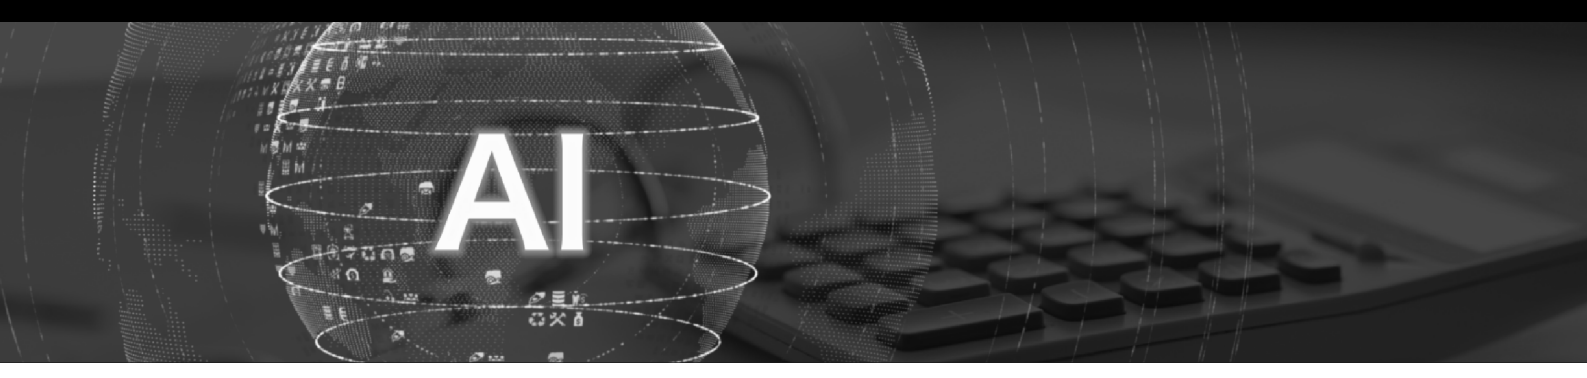

REGULATORY

The AI Governance Stack and South Africa's Draft National AI Policy: An Operational Gap in Search of a Framework

14 April 2026 *9 min read*

Author's Note

I am presently reading Noah M Kenney's *Governing Intelligence: Law, Privacy, Security, and Compliance*,[1] and it has given me genuine cause to reflect, which I suspect was precisely the author's intention. The book lands at a time when South Africa has published its own Draft National Artificial Intelligence Policy,[2] opening a public comment period on what is intended to become the foundational instrument for AI governance in this jurisdiction.

The timing is fortunate. Kenney's central thesis, that AI governance must be understood and implemented as a structured, layered, interdependent system, throws into sharp focus both the ambitions and the shortcomings of South Africa's Draft Policy.

What follows is an attempt to read the Draft Policy through the organising framework at the heart of Kenney's text, the AI Governance Stack, and to consider whether South Africa's proposed approach is structurally adequate for the task it sets itself.

The AI Governance Stack as an Organising Discipline

Kenney's AI Governance Stack is a five-layer operational model, drawn from a decade of practical implementation work across regulated industries.[3] Built from the base upward, it comprises:

- **Layer 1** (Data Governance) constituting data inventory, quality management, bias assessment, provenance tracking, and consent mechanisms;
- **Layer 2** (Model Governance) focusing on architecture review, fairness testing, robustness evaluation, interpretability, and model documentation;
- **Layer 3** (System Integration Governance), considering integration architecture, pipeline security, cascading failure analysis, human-AI interaction design, and boundary condition testing;
- **Layer 4** (Control and Monitoring Governance) addressing access controls, performance monitoring, anomaly detection, incident response, and deployment governance; and
- **Layer 5** (Audit and Evidence Governance) calling for documentation standards, evidence preservation, audit mechanisms, regulatory reporting, and stakeholder communication.[4]

The framework's real force lies in its insistence on cascading dependency.

Each layer of the Kennedy AI Governance Stack creates the foundation for the one above it, and points out that a governance failure at Layer N cannot be fully remediated at Layer N+1.[5] This is not simply an architectural preference. It is a testable operational claim: piecemeal governance, attending to audit whilst neglecting data quality, or monitoring without model documentation, will produce governance that is structurally unsound, no matter how many resources are thrown at it.[6] This is in fact true for any form of regulatory compliance.

The practical upshot, which Kenney demonstrates through a detailed walkthrough of an AI credit decision system, is that each layer must have exactly one primary organisational owner and that governance must be sequenced to follow the dependency chain.[7] Done properly, the Stack transforms governance from a set of aspirational commitments into something closer to an executable specification, with defined requirements, thresholds, decision rules, and verification criteria.[8]

South Africa's Draft Policy: Ambition Without Architecture

There is much to celebrate in the Draft Policy. It is rightly anchored in the Constitution of the Republic of South Africa, 1996, and expressly provides that AI must not be used to violate the rights enshrined in sections 9 (equality), 10 (human dignity), 14 (privacy), 16 (freedom of expression), and 33 (just administrative action), amongst others.[9] It identifies the Protection of Personal Information Act 4 of 2013 (POPIA), the Cybercrimes Act 19 of 2020, and the Promotion of Access to Information Act 2 of 2000 as part of the legislative architecture within which AI governance must operate.[10] It goes further still, proposing the establishment of a National AI Commission, an AI Ethics Board, an AI Regulatory Authority, an AI Ombudsperson Office, a National AI Safety Institute, and an AI Insurance Superfund modelled on the Road Accident Fund.[11]

These are serious institutional commitments that should not be dismissed. But when one measures them against the operational specificity of the Governance Stack, a conspicuous gap opens up. The Draft Policy proceeds largely at the level of principles and institutional mandates. It sets out six key principles of responsible AI, fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability, and proposes embedding these across the AI lifecycle.[12] It calls for "sufficient explainability" and "sufficient transparency" in high-risk systems.[13] It contemplates risk-based classification, drawing some inspiration from the European Union AI Act.[14]

What it does not do is specify the operational infrastructure through which any of these principles can be enforced. One looks in vain for anything equivalent to the Stack's requirement that organisations maintain data catalogues with provenance records documenting origin, transformations, and lineage, or its mandatory quality thresholds, completeness at 95 per cent, accuracy at 98 per cent for labelled data, cross-source consistency at 90 per cent, below which data must not be used for model training without documented exception approval.[15] Cascading failure analysis, circuit breaker requirements for systems with downstream dependencies, boundary condition testing protocols, none of these features.[16] The Policy's reference to "AI-specific data governance frameworks that ensure provenance, quality control, and interoperability of datasets" reads as aspiration, not specification.[17]

Key Tensions and Risks

Three tensions in the Draft Policy deserve close scrutiny.

The first concerns accountability. The Draft Policy's treatment of it is structurally incomplete. It provides that "organisations must take responsibility for the outcomes of their AI systems" and that "accountability must ultimately point to an attributable official or entity."[18] That is necessary, but it is not enough. Kenney's point is that diffuse accountability is the primary organisational failure mode in AI governance. The remedy is to assign determinate accountability at each Stack layer: data stewards at Layer 1, ML engineering leads at Layer 2, platform and infrastructure teams at Layer 3, security and operations teams at Layer 4, and compliance and legal teams at Layer 5.[19] Without that degree of granularity, the Draft Policy's accountability requirement risks becoming what Kenney aptly terms a "compliance fiction", formally satisfied but operationally hollow.

The second tension arises from the Draft Policy's reliance on POPIA as the primary data governance instrument for AI, which is, at best, partial. POPIA's conditions for lawful processing, including purpose limitation (section 13), minimality (section 10), and security safeguards (section 19), were simply not designed with the demands of AI training data in mind. The friction between data minimisation and the data-hungry requirements of machine learning model training, which Kenney identifies as a fundamental governance challenge under the analogous provisions of the GDPR,[20] is not acknowledged in the Draft Policy. Nor does the Draft Policy grapple with how section 71 of POPIA, which governs automated decision-making, will interact with the proposed AI Ombudsperson's jurisdiction or the AI Regulatory Authority's audit mandate.[21]

The third tension is regulatory fragmentation. The Draft Policy proposes an elaborate institutional architecture involving the DCDT, ICASA, the Information Regulator, the Competition Commission, the South African Reserve Bank, and the Financial Sector Contingency Forum, among others.[22] Kenney's argument on this point is direct: the Governance Stack provides a unified architecture through which organisations can satisfy the requirements of multiple regulators by means of a single layered governance system, rather than maintaining separate compliance programmes for each.[23] Without a unifying operational framework, the Draft Policy's multi-regulator model risks imposing precisely the kind of compliance fragmentation that the Stack was designed to resolve.

A Considered View

What, then, should organisations operating in or entering the South African market actually be doing?

It is submitted that the Draft Policy should be treated as a signal of regulatory direction, not as a governance blueprint. Its principles are sound and its institutional ambitions are genuine. But the operational gap between principle and implementation is wide, and organisations that wait for the regulatory apparatus to mature before building their own governance frameworks will find themselves badly exposed. If the EU AI Act teaches us anything, and Kenney documents this in considerable detail, it is that compliance costs compound rapidly when governance is retrofitted rather than designed in from the outset.[24]

Organisations would be well advised, now, to map their AI systems against the five layers of the Governance Stack, assign primary ownership at each layer, and begin building the documentation, testing, and monitoring infrastructure that any competent regulator will eventually demand.[25] They should ensure that their data governance practices satisfy POPIA's existing requirements whilst also anticipating the more demanding standards that the Draft Policy foreshadows.[26] And they should engage meaningfully with the public comment process, not merely to protect commercial interests, but to press for the kind of operational specificity that separates effective governance from well-intentioned aspiration.

If Kenney's book can be reduced to a single proposition, it is that governance must be engineered, not merely declared.[27]

South Africa's Draft Policy has declared its intentions. The engineering remains to be done.

---

[1] Kenney NM *Governing Intelligence: Law, Privacy, Security, and Compliance* (Digital 520 2026).

[2] Draft South Africa National Artificial Intelligence (AI) Policy (March 2026) published in GG 54477 of 10 April 2026.

[3] Kenney (n 1) page 22.

[4] Kenney (n 1) page 22-25.

[5] Kenney (n 1) page 18-19; see also page 25 ("Failure at any layer cascades upward; governance cannot be implemented piecemeal").

[6] Kenney (n 1) page 30.

[7] Kenney (n 1) page 26-28.

[8] Kenney (n 1) page 30.

[9] Constitution of the Republic of South Africa, 1996, section 9, 10, 14, 16, 33; Draft AI Policy (n 2) page 8.

[10] Protection of Personal Information Act 4 of 2013; Cybercrimes Act 19 of 2020; Promotion of Access to Information Act 2 of 2000; Draft AI Policy (n 2) page 7.

[11] Draft AI Policy (n 2) page 26–27.

[12] Draft AI Policy (n 2) page 62.

[13] Draft AI Policy (n 2) page 35–36.

[14] Draft AI Policy (n 2) page 36; cf Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (EU AI Act).

[15] Kenney (n 1) page 23, 27.

[16] Kenney (n 1) pages 24, 32–33.

[17] Draft AI Policy (n 2) page 53.

[18] Draft AI Policy (n 2) page 58.

[19] Kenney (n 1) page 20, 26.

[20] Kenney (n 1) page 258; see also Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data (GDPR) art 5(1)(c).

[21] POPIA section 71; Draft AI Policy (n 2) page 72, 26.

[22] Draft AI Policy (n 2) page 28, 61.

[23] Kenney (n 1) page 123, 154.

[24] Kenney (n 1) page 26; see also Regulation (EU) 2024/1689 (EU AI Act).

[25] Kenney (n 1) page 22, 26–28.

[26] Draft AI Policy (n 2) page 55–56; POPIA section 10, 13, 19.

[27] Kenney (n 1) page 30.

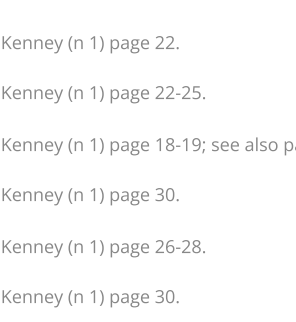

Ahmore Burger-Smidt

Head of Regulatory

View Profile