



REGULATORY

Speak now or forever hold your peace. The draft AI policy has been published and parties have 60 days to comment

10 April 2026 6 min read

by Ahmore Burger-Smidt, Director and Head of Regulatory

On 10 April 2026, South Africa's Department of Communications and Digital Technologies published its Draft National Artificial Intelligence Policy and opened a sixty-day public comment window.

At eighty-six pages, the document covers an extraordinary amount of ground, everything from supercomputing infrastructure to the digitisation of indigenous languages.

If your company develops, deploys, or procures AI systems with any connection to South Africa, you need to be reading this document carefully. And you need to be paying attention to what it doesn't say as much as what it does.

The policy's headline vision, "*AI for inclusive economic growth, job creation, cost reduction, and a developing Africa*", is hard to argue with. Education, healthcare, agriculture, and public administration are flagged as priority sectors, and the policy sets out six objectives covering skills development, public-service modernisation, ethical governance, and cultural preservation.

Where things get really ambitious, perhaps overly so, is in the institutional design. The draft proposes:

- 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, designed to compensate people harmed by AI-driven decisions.

The risk-based regulatory approach borrows openly from the EU AI Act, with stricter rules for high-risk applications and lighter treatment elsewhere, plus provision for regulatory sandboxes. Its six principles of:

- responsible AI,
- fairness,
- reliability and safety,
- privacy and security,
- inclusiveness, transparency, and
- accountability,

will feel familiar to anyone who has spent time considering the OECD AI Principles. None of this is controversial. But the real question is whether the detail behind these commitments is adequate, and in relation to privacy, it is doubted.

Credit where it is due: the draft makes the right noises on data protection. It commits to harmonising AI privacy controls with the Protection of Personal Information Act (**POPIA**), enforcing its eight conditions for lawful processing, and embedding data protection by design and default, data minimisation, purpose limitation, and storage limitation into AI governance. It calls for Privacy Impact Assessments when sensitive information is at stake and points to POPIA's Section 71 on automated decision-making as a transparency safeguard.

The problem is that the policy never seems to get beneath the surface. Given the purpose limitation, in machine learning, training data is routinely repurposed across models and applications in ways that bear little resemblance to the original reason it was collected. The policy says nothing meaningful about how to handle that. Or consider data minimisation. The draft simultaneously champions minimisation and calls for a "*sustained national effort to curate large, diverse datasets in AI-ready formats*," treating non-private data as a "*public good*". You cannot have it both ways without explaining how you intend to square the circle, and the draft does not try.

Then there is Section 71 of POPIA. The policy rightly identifies it as relevant, but stops there. Section 71 gives individuals a right not to be subject to decisions based solely on automated processing, but it is a narrow provision. How does it interact with the broader rights of data subjects, the right to object, or the right to have personal information corrected? The policy does not explore this. When the economy considers rolling AI out across healthcare diagnostics, credit scoring, law enforcement, and public administration, that is a gap with real consequences for real people.

Working across the UK and EU data protection regimes, it does not take much effort to spot the policy's influences and, unfortunately, its shortcomings.

The EU AI Act provides a legally binding, granular risk classification system backed by conformity assessments, post-market surveillance, and meaningful penalties. South Africa's draft uses the same vocabulary of risk categorisation, but punts the substance, what counts as "high-risk," "medium-risk," or "low-risk", to future regulations and sector strategies. That leaves organisations in limbo, uncertain of what they actually need to do.

The rights gap is just as concerning. Under the UK GDPR and the Data Protection Act 2018, individuals have the right to meaningful information about the logic behind automated decisions, the right to human intervention, and the right to challenge outcomes. POPIA offers less, and the draft policy's language around "*sufficient explainability*" and "*sufficient transparency*" risks entrenching a lower standard than what many multinationals already meet under UK or EU law. The word "*sufficient*" introduces flexibility, but it also invites interpretation by those who may not share the same commitment to individual rights.

Cross-border data flows deserve a mention, too. The policy invokes the National Policy on Data and Cloud and frames data sovereignty partly as a guard against "*perpetuation of colonial-era data extraction practices*". That language resonates politically, but it needs to translate into functioning legal mechanisms. The adequacy frameworks under the UK GDPR and the EU's standard contractual clauses are well-established

tools; South Africa's own regime for cross-border transfers under POPIA Section 72 remains comparatively undeveloped, and this policy does not move the needle.

More worrisome is the institutional design. Seven new bodies, on top of existing regulators like the Information Regulator, ICASA, and the Competition Commission, is a recipe for overlap, turf disputes, and diluted accountability. The policy acknowledges the need for a National AI Regulatory Forum to coordinate these bodies, but the governance lines remain vague. South Africa is not a country with limitless public resources. The danger is that the country will end up with impressive-sounding institutions that lack the funding, people, and political independence to do anything meaningful.

But what should organisations do now? Three things.

- First, do not wait. Start mapping AI systems against the risk categories and ethical principles in the draft, and benchmark data protection practices against POPIA, the UK GDPR, and the EU AI Act together, apply the highest common standard.
- Second, respond to the consultation. The government has framed this policy explicitly as a "*point of departure*" and a "*work-in-progress*". Submissions that push for sharper data protection obligations, clearer risk definitions, and stronger individual rights around automated decision-making would make a genuine difference.
- Third, keep a close eye on who ends up doing what. Whether the Information Regulator, the proposed AI Regulatory Authority, or the AI Ethics Board takes the lead on privacy enforcement will shape the entire character of South Africa's AI governance regime.

This is a creditable piece of policy work and it reflects a serious engagement with international AI governance thinking. But good intentions are not the same as good regulation.

On the issues that matter most to individuals, privacy, data protection, the right to understand and challenge decisions made about you by a machine, the draft stays at the level of aspiration. Turning those aspirations into enforceable, practical obligations is the hard part, and it is the part that still lies ahead.

The comment window is open.

Use it.



Ahmore Burger-Smidt

Head of Regulatory

[View Profile](#)