IP BRIEFS®

MADIP
GUSTO DE DIFFERENTIA

South Africa's Draft National AI Policy: An IP practitioner's and AI-using business perspective

Citation: Government Gazette No. 54477, 10 April 2026 (Draft South Africa National Artificial Intelligence (AI) Policy, dated March 2026).

Disclaimer: This article is general information and not legal advice. Obtain advice for your specific facts.

Executive summary

South Africa's Draft National Artificial Intelligence (AI) Policy is ambitious, values-driven, and explicitly anchored in the Constitution and Bill of Rights. It adopts a risk-based governance posture, proposes a suite of new institutions (AI Commission/Office, AI Ethics Board, AI Regulatory Authority, AI Ombudsperson, AI Safety Institute, and an "AI Insurance Superfund"), and stresses "sufficient transparency" and "sufficient explainability" in high-risk settings.

From an intellectual property (IP) and commercialisation standpoint, the draft is directionally helpful, particularly in its recognition of copyright and related rights, data sovereignty, and the need to support local innovation ecosystems. However, it remains high-level and, in several places, blends IP concepts with governance tools in a way that could create uncertainty for creators, rightsholders, and businesses.

For businesses that develop or deploy AI, the document sends a clear message: *"expect more governance, more auditing, and more accountability—especially in public-sector or high-risk use cases".*

Yet it does not always specify who must do what, by when, under what standard, or with what safe harbors.

Why this policy matters for IP and business

The draft policy positions AI as a general-purpose technology and frames the policy imperative as inclusive growth, job creation, and ethical deployment, with special sensitivity to South Africa's historical inequities and the digital divide.

For IP owners and IP advisers, AI affects:

- Copyright and related rights (training inputs; generated outputs; infringement risk; exceptions/limitations).
- Trade secrets and confidential information (model prompts, fine-tuning data, weights, evaluation, and safety artefacts).
- Patents (AI-enabled inventions; inventorship issues; patentability of AI-implemented claims).
- Brands and reputation (deepfakes, synthetic endorsements, disinformation).
- Data rights and database value (data access, licensing, and sovereignty in cross-border settings).

For businesses that use AI, the policy affects:

- Compliance (POPIA, automated decision-making, impact assessments, audit readiness).
- Contracting (vendor risk allocation, audits, data processing, IP warranties, indemnities).
- Product design (explainability, human oversight, bias testing).
- Market access (public-sector procurement conditions, certification/licensing of "high-risk" systems).

IP specific relevant highlights from the draft policy

The policy situates AI within a broader legal ecosystem that includes the Constitution, POPIA, ECTA, the Cybercrimes Act, and notably copyright and performers'/artists' rights (Copyright Act 1978 and the Copyright Amendment Bill / Performers' Protection Amendment Bill, as referenced).

The draft also flags the need to distinguish text and data mining (TDM) from AI web scraping; emphasises POPIA-aligned data governance, including "data protection by design and default," minimisation, purpose limitation and storage limitation; adopts "sufficient explainability" and "sufficient transparency," particularly in high-risk contexts; contemplates algorithmic audits, bias testing, and impact assessments and proposes institutions that could issue certifications and oversee compliance.

Critical assessment from an IP practitioner's perspective

IP is recognised—but the policy needs sharper doctrinal separation.

It is positive that the policy references copyright and related rights and treats IP as a major policy interface. However, it often uses "IP" as a broad synonym for "innovation" without distinguishing:

- what is owned (data vs content vs models vs outputs).
- which acts are regulated (copying, adaptation, communication to the public, extraction).
- what permissions exist (licenses, exceptions/limitations).
- which enforcement pathway applies (civil IP enforcement vs administrative AI governance).

This may lead to uncertainty; and when the rules are not clear, deals slow down and become more expensive—especially fundraising, M&A, procurement, and cross-border partnerships where IP due diligence is a key focus.

"Mandatory watermarking" not clearly specified and risks becoming an unworkable compliance lever.

The draft contemplates "mandatory water-markings" (notably in the Large Language Model (LLM) context) as an IP protection measure. But watermarking is a technical and operational mechanism, and the draft does not clarify:

- watermarking of outputs vs models vs training materials.
- whether the obligation is universal or risk tiered.
- the standard for robustness, detection, and interoperability.
- how it applies to open-source models and private fine-tunes.

A suggested improvement would be to adopt a risk-based approach. Where AI content could cause actual harm, require clear labelling and technical markers (like watermarking) where practical. For lower risk uses, focus on governance measures such as contracts, internal policies, and keeping records that can be audited.

Training data

The acknowledgement of TDM as a legitimate academic practice is useful. But the draft does not set out what a compliant training pathway looks like for commercial actors.

What remains absent is a clear policy position on whether South Africa should introduce a statutory TDM exception—and, if so, whether it should be confined to research or extend more broadly—together with workable opt-out mechanisms for rightsholders. The draft does not address how TDM would operate in practice alongside real-world access controls and restrictions, including paywalls, contractual terms, technical protection measures, and robots.txt; nor does it set out a practical standard

for data origin and traceability documentation—that is, maintaining auditable records of where training data is sourced and the terms on which it may be used.

In the absence of clear rules, businesses should treat model training as a high-risk IP activity and maintain disciplined data origin and traceability records.

Ownership of AI outputs and authorship

The draft is governance-heavy and does not address the ownership and protectability of AI-generated outputs. Businesses need clarity on how South African copyright doctrine will treat human contribution, tool use, and authorship thresholds.

The policy should at least flag this as a priority workstream for guidance (including procurement rules on deliverables and output rights).

Transparency and explainability must be designed to protect trade secrets.

The policy endorses transparency and explainability—appropriately so for rights protection and contestability. But without guardrails, transparency obligations can force disclosure of confidential business information (CBI) and trade secrets.

A suggested improvement is to adopt confidentiality-safe audit mechanisms: accredited assessors, regulator confidentiality protections, and public summaries that are meaningful but non-proprietary.

New institutions may create overlapping mandates unless dispute routing is clear.

With an AI Ethics Board, AI Regulatory Authority, Ombudsperson, and multi-regulator coordination (Information Regulator, ICASA, Competition Commission, financial regulators), the draft risks forum complexity.

This means the policy should include a simple, public facing "roadmap" that explains which regulator or body is responsible for which kind of AI problem, and where a business or affected person must take a dispute first. Because the draft proposes several new AI institutions (e.g., an AI Regulatory Authority, an Ethics Board, an Ombudsperson) alongside existing regulators (like the Information Regulator under POPIA, ICASA, the Competition Commission, and sector regulators), it is easy for complaints and enforcement to become fragmented or duplicated. Clear dispute-routing rules would also explain how an AI regulator's audit, certification decision, or non-compliance finding would be treated if the matter later ends up in court as an IP dispute (for example, copyright infringement or misuse of trade secrets)—i.e., whether the regulator's findings can be relied on as evidence, whether they carry any legal weight, and how court enforcement and regulatory enforcement fit together.

Critical assessment from the perspective of businesses that use AI.

Strong intent; limited operational detail

Businesses will welcome the policy direction, but will need implementation instruments (standards, codes, templates) that are proportionate for SMEs.

What businesses need next is clearer operational guidance: agreed definitions of what counts as "high-risk" and what regulators will accept as "sufficient explainability" and "sufficient transparency," as well as a set of baseline controls that organisations can apply to low- and medium-risk AI systems without over-engineering compliance. They will also need phased implementation timelines and practical compliance toolkits (templates, checklists, and standards) so that organisations—especially SMEs—can implement the requirements consistently and cost-effectively.

Procurement and vendor management will shape the market.

The policy's emphasis on data sovereignty and third-party risk—particularly in the context of public procurement—is likely to influence contracting norms across the private sector as well.

AI contracting terms can be expected to become more exacting, with increased emphasis on audit rights, clearer allocation of POPIA roles and responsibilities (including distinctions between operators and responsible parties), tighter controls on cross-border data transfers, and more detailed assurances regarding data origin and traceability for both datasets and models. In parallel, IP protections are likely to be strengthened, with more robust warranties and indemnities—particularly in relation to training data and outputs—alongside more prescriptive incident response obligations addressing security breaches, model failures, and defined notification timelines.

Impact assessments are likely to become standard operating practice.

The draft contemplates Human Rights Impact Assessments and Regulatory Impact Assessments, plus bias and gender impact assessments.

The practical implication is that these will become board-level governance artefacts, similar to privacy impact assessments. Companies should standardise internal processes and retain auditable records.

Recommendations

In my view, the next iteration of the Draft National AI Policy should:

- ✓ define AI-related IP artefacts with greater precision—covering training datasets, fine-tuning sets, prompts, model weights, synthetic data, evaluation outputs, and generated content—so that regulation targets the appropriate layer of the AI stack.
- ✓ replace the current blanket watermarking language with tiered data origin and

traceability requirements, calibrated to risk, technical feasibility, and sector.

- ✓ articulate a lawful training pathway by setting out a clear policy position on TDM, including rightsholder opt-outs, the potential role of licensing registries (where appropriate), and practical expectations for data origin and traceability documentation.
- ✓ ensure that trade secrets are adequately protected within audit and transparency frameworks, for example through the use of accredited audits and redacted public reporting.
- ✓ clarify institutional mandates and dispute-routing mechanisms to avoid duplication, forum shopping, and regulatory uncertainty; and
- ✓ publish model procurement and contracting templates—particularly for government and SMEs—to support compliant adoption and reduce transaction costs.

What should businesses do now?

In the absence of settled regulatory guidance, organisations should take a structured, risk-based approach to managing AI-related IP and data exposure. The checklist below provides a practical, defensible baseline for governance—focusing on training data, data origin and traceability, vendor accountability, and oversight of high-impact use cases. It is intended as a minimum standard that can be implemented now and scaled as policy and regulatory expectations evolve.

Inventory & Classification of AI Use Cases

☐ Identify all AI systems and tools in use (internal and vendor-based)

☐ Map each use case to business function (e.g., legal, marketing, product)

☐ Classify use cases by risk/impact level (low, medium, high)

☐ Flag high-stakes uses (e.g., automated decision-making, profiling, sensitive data)

☐ Document purpose, inputs, outputs, and decision influence

POPIA-Ready Data Governance

☐ Apply data minimisation (only necessary data collected/processed)

☐ Define data retention periods and deletion protocols

☐ Conduct cross-border data transfer assessments (POPIA + adequacy safeguards)

☐ Identify lawful basis for processing personal information

☐ Implement appropriate security safeguards (technical and organisational)

☐ Align with internal privacy policies and PAIA/manual requirements

Dataset Provenance Register

☐ Record all data sources (internal, third-party, public datasets)
☐ Document licensing terms and usage rights
☐ Capture restrictions (territory, purpose, duration)
☐ Record opt-outs, consent conditions, and withdrawal mechanisms
☐ Track data transformations (anonymisation, pseudonymisation)
☐ Maintain version control and audit trail for datasets

Human Oversight & Governance Controls

☐ Define human-in-the-loop or human-on-the-loop controls
☐ Establish escalation pathways for high-risk decisions
☐ Assign accountability (roles and responsibilities)
☐ Document override and intervention protocols
☐ Train staff on AI risk awareness and escalation triggers
☐ Ensure governance oversight (e.g., risk committee / legal review)

Vendor & Contractual Accountability

☐ Conduct vendor due diligence (security, compliance, reputation)
☐ Include data protection obligations aligned with POPIA
☐ Define audit rights and compliance verification mechanisms
☐ Include IP protections (ownership, licensing, indemnities)
☐ Require incident response and breach notification obligations
☐ Clarify data use limitations (no secondary use without consent)

Audit Readiness & Impact Assessments

☐ Maintain comprehensive documentation of AI systems and decisions
☐ Prepare Data Protection Impact Assessments (DPIAs) for high-risk use cases
☐ Keep logs of data processing, model outputs, and decisions
☐ Document risk assessments and mitigation measures
☐ Ensure traceability and explainability where required
☐ Establish internal audit and review schedule

Key takeaways for IP practitioners and AI-using businesses

- The policy strongly signals future regulation and potentially sector legislation; businesses should begin compliance and governance readiness now.
- IP and data are treated as strategic national assets, but the draft needs sharper distinctions between ownership, permitted uses (including TDM), licensing models, and enforcement.

- "Watermarking" for LLMs is flagged as mandatory, but the draft does not define scope, standard, or enforceability—creating risk of fragmented requirements.
- Procurement and third-party vendor governance are central: data sovereignty, POPIA compliance, and contestability will likely become deal terms.

Conclusion

The Draft National AI Policy provides a strong ethical and constitutional foundation. For IP practitioners and AI-using businesses, its principal value lies in signalling a shift toward risk-based governance and enhanced accountability.

To unlock investment and reduce regulatory friction, the next iteration should strengthen operational clarity around training data and IP, adopt a workable approach to data origin and traceability, and design transparency obligations that support contestability while preserving trade secrets.

Why action now matters.

The Draft National Artificial Intelligence (AI) Policy is not just a statement of principles—it is a clear signal of the regulatory and commercial direction. It anticipates a risk-based governance model, new oversight institutions, increased audit and impact-assessment expectations, and stronger requirements around privacy, transparency, and accountability. For businesses building or deploying AI, and for creators and rightsholders whose works and data may be used in AI systems, these choices will shape compliance costs, procurement requirements, product design constraints, and competitive advantage.

The consultation window (open until 10 June 2026) presents a timely opportunity to reduce legal uncertainty before it crystallises into standards, procurement requirements, and enforcement practices. This is particularly critical for IP- and data-driven sectors, where unresolved issues—such as lawful pathways for training data use, the scope and practicality of data origin and traceability requirements (including watermarking), and the interaction between transparency obligations and trade secret protection—can materially influence innovation, investment decisions, and cross-border collaboration. Early engagement enables stakeholders to shape workable, risk-calibrated rules that safeguard rights and public interests without inadvertently constraining legitimate research, commercialisation, or the growth of local AI ecosystems.

About the Editor/Author

Dr. Madelein M. Kleyn is a South African registered patent attorney and admitted attorney of the High Court, with over 30 years' experience in intellectual property strategy, licensing, technology transfer, and data governance across industry and university innovation ecosystems. She is the CEO and Founder of Mad K IP Consulting and recently served as Chief Legal & IP Officer at Omnisient, a privacy-enhancing data collaboration platform. Dr. Kleyn is a recognised international IP strategist, a Vice President of the Licensing Executives Society South Africa, and a regular speaker and author on IP, data, and commercialisation.

Dr Kleyn will be a workshop speaker at the LES International Annual Conference 2026 in Dublin (26–29 April 2026) on "AI vs. the IP Professional: Adapting Expertise and Credibility in the Age of Change," moderated by Anji Miller of LifeArc.