

10 June 2026

Ms Nonkqubela Jordan-Dyani
Director-General
Department of Communications and Digital Technologies
Attention: Mr. M. Mashologu, Deputy Director-General: Digital Society and Economy
Block A, iParioli Office Park, 1166 Park Street, Hatfield, Pretoria
Private Bag X860, Pretoria, 0001

Re: Comments of ACT on the Draft South Africa National Artificial Intelligence (AI) Policy

The Association for Competitive Technology (ACT) appreciates the opportunity to comment on the Draft South Africa National Artificial Intelligence (AI) Policy (the Draft Policy) published by the Department of Communications and Digital Technologies (DCDT) on 10 April 2026. We commend DCDT for advancing a national policy that is expressly technology-neutral, risk-proportionate, and grounded in extensive stakeholder consultation, and for recognising micro, small, and medium enterprises (MSMEs) and startups as central beneficiaries of South Africa's AI ecosystem.

I. Introduction and Statement of Interest

ACT is a global trade association representing the small business technology developer community both in South Africa and around the globe. Our members are entrepreneurs, independent developers, and small firms that build AI-enabled tools across nearly every sector of the economy, including healthcare, education, agriculture, and financial services, which are among the priority sectors the Draft Policy identifies. The small developers we represent stand to gain a great deal from a clear and enabling national framework, and they are also the participants in the value chain least able to absorb compliance burdens that are not matched to demonstrated risks.

II. ACT's General AI Policy Recommendations

Initially, we recommend that DCDT align its Policy with consensus AI policy recommendations that ACT has developed in consultation with its members both in and outside of South Africa.¹ A successful South African government approach to AI should reflect the following:

1. **Understand the technology.** ACT members have been using AI for decades and even with the more advanced technologies in use today, humans are always responsible for its output. Although there is no consensus definition that encompasses the various types of AI in use today, basic taxonomies are nonetheless helpful for anybody seeking an understanding of the technology, and we urge definitional alignment with leading international standards such as ISO/IEC 22989.²

¹ ACT, AI Policy Principles (2023), available at <https://actonline.org/wp-content/uploads/2023-11-16-ACT-AI-Policy-Principles-FINAL.pdf>.

² International Organization for Standardization and International Electrotechnical Commission, ISO/IEC 22989:2022, Information technology, Artificial intelligence, Artificial intelligence concepts and terminology.

2. **Understand how current law applies.** A wide range of South African laws already prohibit harmful conduct regardless of whether AI is involved, including consumer protection, data protection, competition, and anti-discrimination law. Policymakers should first establish how existing frameworks apply to AI activity and avoid creating sweeping new authorities or agencies that overlap awkwardly or inconsistently with current law.
3. **Quality assurance and oversight.** Frameworks should use risk-based approaches calibrated to recognised standards of safety, efficacy, and equity. Risk and liability should be distributed across the value chain so that those best positioned to minimise a given risk, based on their knowledge and ability, have the appropriate incentives to do so.
4. **Thoughtful design.** Frameworks should encourage AI systems informed by real-world workflows, human-centred design, usability, and end-user needs, developed through collaboration among users, developers, and other stakeholders.
5. **Access and affordability.** Frameworks should keep AI-enabled products and services accessible and affordable, support the resources needed to scale systems, and avoid policies that limit accessibility options for developers.
6. **Research and transparency.** Frameworks should fund and facilitate research and development and should treat the costs and benefits of transparency as a subject for study, so that transparency obligations are targeted to the circumstances in which they actually address risk.
7. **Modernised privacy and security frameworks.** Frameworks should protect personal data while allowing the responsible flow of information, applying reasonable data minimisation, consent, and consumer-rights rules rather than undue barriers to data processing.
8. **Bias.** Stakeholders across the chain, including developers, deployers, and users, should identify, address, and mitigate the risk of harmful bias, and agencies should examine data provenance so that bias in datasets does not result in unlawful discrimination or other harm.
9. **Ethics.** Ethical norms should apply across the AI lifecycle, remain consistent with international human-rights conventions, prioritise inclusivity, and protect the sensitive information AI systems may reveal.
10. **Education.** Frameworks should support AI education and workforce development for users, consumers, and students, including instruction on how intellectual property law applies to AI.
11. **Intellectual property.** Protection of IP rights is critical to the evolution of AI. Compliance measures should not undercut IP or trade-secret protections, and IP analysis should focus on the degree of human involvement rather than on the mere involvement of AI.

ACT also encourages DCDT to ensure that its approach to AI reflects the interdependencies in AI value chains and the shared responsibility that all have to ensure safety and efficacy, whether they are AI developers, deploying organisations, end users, standard-setting organisations, certification bodies and test beds, accrediting and licensing bodies, academic institutions. We urge DCDT to

align with ACT's AI Roles and Interdependencies Framework³, which provides baseline recommendations for supporting AI transparency and safety, centred on the concept that responsibility for safety should rest with the actor best able to manage a known or expected use or risk.

We also urge DCDT to align the Policy with leading international standards and frameworks. These widely adopted, voluntary, consensus instruments give South African developers an interoperable vocabulary and lower the cost of cross-border participation, and they include the OECD Recommendation on Artificial Intelligence,⁴ the United States National Institute of Standards and Technology AI Risk Management Framework,⁵ ISO/IEC 42001 on artificial intelligence management systems,⁶ and ISO/IEC 23894 on artificial intelligence risk management⁷ alongside the foundational terminology standard ISO/IEC 22989 referenced above. Building on these consensus frameworks rather than divergent national requirements will help South African developers compete globally and allow international partners to participate readily in South Africa's AI ecosystem.

III. ACT's Specific Input on the Draft Policy

a. Apply existing law first, and consolidate rather than multiply new institutions

The Draft Policy contemplates an unusually large set of new bodies, including a National AI Commission or Office, an AI Ethics Board, an AI Regulatory Authority, an AI Ombudsperson Office, an AI Insurance Superfund, a National AI Safety Institute, and an Integrated AI-Powered Monitoring Centre, alongside a repositioned mandate for ICASA and a new National AI Regulatory Forum (sections 4.6 and 4.7). South Africa already has a substantial body of law that applies to AI-related conduct regardless of whether AI is involved, including the Protection of Personal Information Act, the Competition Act, consumer protection law, and the intellectual property statutes the Draft Policy itself lists in section 1.3.

We encourage DCDT to establish how these existing frameworks apply to AI before standing up new authorities, and to fill only genuine gaps. This approach gives developers a predictable and stable environment in which to invest and build, and it avoids the barriers to entry that a fragmented set of overlapping requirements would otherwise create. A proliferation of overlapping bodies risks duplicative or inconsistent obligations, jurisdictional uncertainty, and higher compliance costs that fall hardest on MSMEs, which must interpret and answer to each new regulator. Where new coordination is needed, we encourage DCDT to favour a single, clearly mandated coordinating function over multiple standalone entities, and to define precisely how any new body relates to existing regulators so that a small developer faces one coherent set of expectations rather than several.

³ ACT, AI Roles and Interdependencies Framework, published as Appendix A to ACT, Comments on the Development of an Artificial Intelligence (AI) Action Plan (Mar. 14, 2025), available at <https://actonline.org/wp-content/uploads/ACT-Comment-re-AI-Action-Plan-14-Mar-2025-w-appendix.pdf>.

⁴ Organisation for Economic Co-operation and Development, Recommendation of the Council on Artificial Intelligence, OECD/LEGAL/0449 (2019, amended 2024).

⁵ National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1 (2023).

⁶ ISO/IEC 42001:2023, Information technology, Artificial intelligence, Management system.

⁷ ISO/IEC 23894:2023, Information technology, Artificial intelligence, Guidance on risk management.

b. Keep the risk-based approach proportionate and tied to demonstrated harms

We welcome the Draft Policy's commitment to a risk-based, risk-proportionate approach (sections 3.3 and 7.3). We respectfully caution against importing the structure of the European Union AI Act, which the Draft Policy cites as an inspiration in section 7.3. In our engagement with policymakers we have found that prescriptive, ex-ante regimes of that kind impose fixed compliance steps that do not scale to the resources of small firms and that often capture benign, low-risk uses. The relevant question is not whether an AI system is present in a given product, but whether a specific use creates a meaningful risk of harm.

Accordingly, we encourage DCDT to define high-risk use cases narrowly and by reference to concrete, demonstrated harms, to confirm that obligations such as algorithmic impact assessments, human-rights and gender impact assessments, mandatory testing, audits, and certification (sections 7 and 8) attach only to genuinely high-risk uses, and to right-size each obligation to the size and role of the actor. A flexible, iterative posture, including the regulatory sandboxes the Draft Policy proposes in section 5.2, is well suited to this end and we support it.

c. Allocate responsibility to the actors best able to manage each risk

The Draft Policy's governance provisions require traceable lines of responsibility and a named accountable official or entity (section 7.1). We support clear accountability, and we encourage DCDT to allocate obligations using the shared-responsibility logic of our Roles and Interdependencies Framework⁸. Responsibility for a given risk should rest with the actor that has the knowledge and the practical ability to mitigate it. A small deploying organisation or downstream developer frequently cannot inspect or alter the behaviour of an upstream foundation model, and it should not bear obligations it has no means to satisfy. Tying duties to control, and ensuring that upstream actors furnish the information downstream participants need, protects consumers more effectively than placing undifferentiated obligations on every actor in the chain.

d. Refrain from establishing the AI Insurance Superfund or any new AI-specific liability regime

The Draft Policy proposes an AI Insurance Superfund, modelled on the Road Accident Fund, to compensate those harmed in ambiguous liability scenarios (sections 4.6 and 7.3). We share the goal of ensuring that those harmed have meaningful redress. We are, however, firmly of the view that a standing AI Insurance Superfund is the wrong instrument for achieving it, and we respectfully urge DCDT not to establish one. A fund financed by levies on developers and deployers would impose recurring, fixed costs that are immaterial to large firms but significant for small ones, and a regime that pools and socialises liability can weaken the incentive for the party best able to prevent a harm to do so, which runs counter to the distribution-of-responsibility approach described above. Rather than create a standing fund now, we encourage DCDT to first assess how existing liability, contract, consumer-protection, and insurance law allocate responsibility for AI-related harms, to identify the specific gaps that remain, and to address any residual gap with the most targeted and least burdensome measure available, calibrated to the size and role of each participant.

⁸ ACT AI Roles and Interdependencies Framework, supra note 3.

In practical terms, a superfund financed by levies on developers and deployers would function as a recurring tax on AI innovation. That cost would be trivial for the largest firms but material for the small developers and startups the Draft Policy rightly seeks to support, and it would arrive before many of them have earned any revenue from the products at issue, discouraging entry and investment at precisely the stage South Africa most wants to encourage. A pooled, no-fault fund also dampens the incentive to prevent harm in the first place, because the actor best able to design a safer system no longer internalises the cost of failing to do so. The Road Accident Fund comparison, on reflection, illustrates the difficulty rather than resolving it. That fund addresses a narrow and well-understood category of physical-injury claims and has still faced persistent solvency and administration challenges, whereas AI-related harms are heterogeneous, context-specific, and hard to price, which would make a broad standing fund especially prone to cost overruns, to cross-subsidisation of higher-risk activities by lower-risk ones, and to unpredictable liabilities that small firms cannot plan around. For these reasons we respectfully encourage DCDT to rely on existing law and on targeted, proportionate measures rather than on a new AI-specific fund.

e. Target transparency obligations and protect IP and trade secrets

We support meaningful transparency, including plain-language notice to individuals affected by automated decisions and avenues to contest them (section 7.2). Consistent with our sixth and eleventh principles, we encourage DCDT to treat the costs and benefits of transparency as a matter for study so that obligations are targeted to the circumstances in which they address real risk, and to confirm that disclosure, audit-publication, and explainability requirements do not compel the release of proprietary information such as source code, model weights, or training data. Disclosing the data on which a model was trained does not establish that the model is fit for a given use. The sounder course is to assess performance in the intended environment using representative data. Protecting trade secrets in this way is fully compatible with the explainability the Draft Policy seeks.

f. Pursue international interoperability and avoid data-localisation mandates

We welcome the Draft Policy's intention to align with international instruments, including the OECD AI Principles⁹ and the NIST AI RMF¹⁰ (sections 4.5 and 5.4). Interoperability with widely adopted, voluntary, consensus frameworks lowers the cost of cross-border participation for small developers and strengthens South Africa's position as a destination for AI investment. We encourage DCDT to favour these interoperable instruments over more prescriptive models, and to ensure that the concept of data sovereignty referenced in the Foreword does not translate into data-localisation requirements. Forced localisation raises costs, fragments the cloud infrastructure on which small firms depend, and tends to disadvantage the smallest participants most, without delivering commensurate protection.

⁹ OECD Recommendation on Artificial Intelligence, supra note 4.

¹⁰ NIST AI Risk Management Framework, supra note 5.

g. Favour voluntary, industry-led standards and guard against standard-essential patent hold-up

The Draft Policy envisions technical standards across several sectors and a National AI Safety Institute working with international counterparts (sections 4.6 and 4.7). We support standards development that is industry-led, voluntary, and based on international consensus, which provides the flexibility to adapt as the technology evolves. As South Africa develops AI standards, we encourage DCDT to take care that AI standardisation is not undermined by standard-essential patent hold-up of the kind that has affected cellular and Wi-Fi standards, and to support licensing of standard-essential patents on fair, reasonable, and non-discriminatory terms.

Standard-essential patents are patents that a user of standards cannot avoid if it builds to a technical standard. When holders of such patents depart from their commitments to license on fair, reasonable, and non-discriminatory (FRAND) terms, they can seek royalties untethered from the value of the patented contribution, condition licences on non-disclosure terms that make it difficult to confirm that offers are in fact non-discriminatory, and use the threat of injunctions to extract payments beyond that value. Small firms are the most exposed, because they lack the large patent portfolios and litigation budgets that larger implementers rely on, and the resulting hold-up raises costs and deters participation in standardised technologies. ACT has examined these dynamics and proposed concrete reforms in its guidance on standard-essential patent licensing,¹¹ including its recent comments to the United Kingdom Intellectual Property Office on fair licensing,¹² and it maintains ongoing resources on this subject. Because AI systems increasingly depend on standardised connectivity and interoperability components, these risks will extend to AI, and we encourage DCDT to anticipate them as it develops AI standards.

h. Build on the Draft Policy's strong support for MSMEs, skills, and infrastructure

We commend several features of the Draft Policy and encourage DCDT to preserve and strengthen them. The express identification of MSMEs and startups as beneficiaries (section 4.3), the support for startups through sandboxes and accelerators (section 4.2), and the emphasis on skills, STEAM education, and reskilling (Strategic Pillar 1) all align closely with ACT's principles on access, affordability, and education. The investments in connectivity, devices, data infrastructure, and compute (section 4.4) address the access barriers that determine whether small developers can participate at all. We encourage DCDT to ensure that the compliance obligations elsewhere in the Draft Policy are calibrated so that they do not offset these supports by imposing costs that only larger firms can bear.

¹¹ ACT, Decoding SEP Licensing: A Comprehensive Guide for the IoT and AI Revolutions (May 9, 2024), available at <https://actonline.org/2024/05/09/decoding-sep-licensing-a-comprehensive-guide-for-the-iot-and-ai-revolutions/>.

¹² ACT, Five Reforms Proposed in Comments to UKIPO Consultation on Fair SEP Licensing (Oct. 8, 2025), available at <https://actonline.org/2025/10/08/act-the-app-association-proposes-five-reforms-in-comments-to-ukipos-consultation-on-fair-sep-licensing/>.

IV. Conclusion

ACT supports South Africa's effort to build an enabling, inclusive, and risk-proportionate national AI framework, and we share DCDT's goal of ensuring that the benefits of AI reach MSMEs, startups, and underserved communities. We encourage DCDT to apply existing law before creating new authorities, to consolidate rather than multiply institutions, to keep risk-based obligations narrow and proportionate, to allocate responsibility to the actors best able to manage each risk, to study carefully any new liability mechanism, to target transparency while protecting IP, and to pursue international interoperability without data-localisation mandates. We would welcome the opportunity to share our Policy Principles and Roles and Interdependencies Framework with DCDT and to serve as a resource as the Draft Policy is finalised.

Sincerely,



Brian Scarpelli
Senior Global Policy Counsel

Chapin Gregor
Policy Counsel

Association for Competitive Technology