



Australian Priorities for Navigating the Crisis in AI policy

Submission to the Joint Select Committee on Artificial Intelligence

14 September 2026

[Dr Greg Austin](#)

Co-founder and Research
Director
Social Cyber Institute

[Dr Gary Waters](#)

Distinguished Fellow
Social Cyber Institute

[Prof Glenn Withers](#)

Co-founder and Director
Social Cyber Institute

This submission does not represent an agreed view of all researchers associated with the Social Cyber Institute. It follows our tradition where we invite individual researchers to contribute to a submission or joint research paper. This submission reflects work by several Institute-affiliated researchers working with their peers in the UK, the US, Malaysia, India, Bangladesh and East Africa. Important background research and analysis by us on these issues is summarised under our project “AI Consequences for Australia” on our website, especially the page :

<https://www.socialcyber.co/social-cyber-institute/submissions-and-notes>.

Australian Priorities for Navigating the Crisis in AI Policy

AI is already having radical impacts. This includes impacts on Australia's national security, its counter-intelligence and the human security of its citizens. Focusing on only a few cases of AI use from one or several sectors risks losing sight of how massive the challenges from AI are, how broad-ranging they are, and what depth of knowledge and expertise is required for a country like Australia to address them coherently and consistently. With rapid change in this revolutionary technology affecting the whole of society, a focus only on selected sub-sets of policy is insufficient. So this submission addresses general strategic settings without which Australia will not be as secure from AI or as competitive with AI as it needs to be. This focus sits squarely under item (n) ("any other related matters") in the Terms of Reference for the Inquiry at the same time as addressing almost all of the other items (a) to (m) in some way.

To meet the Committee's request for submissions to be short, we have based our analysis here on a series of background notes and other references available on the website of the Social Cyber Institute on the project page: "[AI Consequences for Australia: Submissions and Notes](#)".

Recommendation 1: Australian stakeholders need to act urgently and explicitly recognise the AI ecosystem as presenting a national and international crisis.

Policymaking for AI should be treated as a crisis of the highest order rather than as a series of isolated technical or regulatory problems. AI presents existential risks, rapidly accelerating change, profound uncertainty and intensified competition among major powers. Although governments and corporations have issued serious warnings about these risks, they have not established coherent or enforceable safeguards. The crisis is further compounded by fragmented responsibilities, weak accountability and the growing influence of major AI companies with clear political agendas.

Defining the AI challenges as constituting a crisis is important because we need to understand that the dynamics of decision-making for a crisis can be different from those for what might be called normal times. Above all, crises can erode accountability for the power-holders and undermine the democratic political influence of the stakeholders, both citizens and business.

We can identify five clusters of urgent action all of which are inter-linked and need to be addressed comprehensively, simultaneously and holistically by national authorities:

- geopolitical and security risks
- power and sovereignty, who controls AI, and on whose terms
- democracy and rights
- socioeconomic and environmental costs
- governance failure and institutional lag.

To reflect this urgency, one immediate concrete policy is for every substantive actor with a stake in AI policy to be called upon to adopt a rolling 90-day AI action plan, with the cycle scaled to the actor: government, peak body or individual organisation. They should report publicly against it at the end of each cycle. Doing nothing remains an option, but it is probably not a socially responsible one. The value of the 90-day discipline is not that three months is sufficient time to resolve anything. It is that a short,

repeating cycle produces visible movement and a public record, which is precisely what the pattern of talking about AI risk while declining any binding constraint on it avoids.

A workable template, proposed in discussion and offered here for debate rather than as a settled model, runs as follows: a top-level directive; capability and asset mapping; an interim supervising team or helpline; progression to a dedicated AI policy working group; interim guidelines pending fuller discovery; AI literacy training; and, in later cycles, development of a governance framework and safety evaluation capability. Not every element will suit every actor, and the sequence is open to argument.

Recommendation 2: Australia must design a far more specific AI plan that is markedly more comprehensive and much more detailed in its policy ambitions than the current National AI Plan.

This plan must be based on experimentation with the mechanisms of government, the country's business environment, its universities and civil society to meet the circumstances of crisis in AI policy.

More radical and high-risk measures must be more prominent. Education at all levels needs to be transformed. A small number of policy breakthroughs each year (the under-16 social media ban, new federal guidelines for data centres,¹ and digital duty of care confined to one or two parameters) should not be seen as meeting the scope and gravity of the AI challenge.

Australia could follow Singapore's model of a good AI Plan, directed by a national AI council, chaired by the Prime Minister, sitting above a policy, infrastructure and assurance stack.² That architecture could take on the EU's more explicit security spine, writing AI-enabled threats and agentic systems explicitly into critical-infrastructure obligations. And Australia could adopt the UK's discipline of a published delivery dashboard, so the national AI roadmap is accountable against dated commitments rather than aspiration and generalities.

Recommendation 3: Significantly greater diversification of sources of AI technologies and expertise for Australia away from the US must be one of the highest priorities for national policy, with self-reliance in the most important operational sectors as a primary reference point all buttressed by new international partnerships.

This recommendation means that Australia must build sovereign AI capability much more deliberately. For this it must adopt a layered national AI strategy that routes work across frontier, open-weight and specialised models by cost, capability and risk; establish a public-option capability through public-interest laboratories, lawful high-quality data and reserved domestic compute; and negotiate dependable frontier access with the United States, subject to review for geopolitical and commercial concentration risk. Define sovereignty operationally across model, infrastructure, knowledge, continuity and human-authority layers rather than as technological independence.

A key further step for Australia should be to look to lead international partnerships. This can include Five Eyes members but also other like-minded countries. The Social Cyber Institute, for

¹ On 15 July 2026, the Australian Prime Minister Anthony Albanese announced a major policy initiative on how new data centre projects should be managed. One element of this speech was the commitment by the Federal Government to "work with State and Territories to ensure large data centres are built in the most appropriate locations, and with input from local communities." See "AI in Australia's interests", Media release, Wednesday 15 July 2026, <https://www.pm.gov.au/media/ai-australias-interests>.

² On 15 July 2026, in the same announcement referenced above, the Prime Minister announced the creation of an AI office in his own department but it looks set to be a pale reflection of the Singapore governance model in its totality.

instance, has proposed that it lead Australian membership of the International Cyber Education Collaborative (ICEC) alongside UK and US partners. Singapore and Canada representatives are also negotiating entry. But funding in support from the Australian government is yet to be forthcoming. This violates the urgency principle.

The strong general point is to proceed by establishing actions on multiple fronts, while reining in an occasional Australian disposition to go it alone in order to claim a world first.

No single lever (legal, institutional or diplomatic) is sufficient on its own. The responses commonly proposed for controlling existential-risk technologies (international conventions, coalitions of the willing, democratic oversight, enforceable safety thresholds) each depend on the others to function. Australia must therefore act on a multi-front basis, engaging many different parts of society, internationally and domestically.

Promising precedents have not been converted into durable regimes: the agreement between former President Biden and President Xi Jinping that humans, not machines, must remain in the decision loop for the use of nuclear weapons has not been institutionalised. Something is holding governments back, and the pattern is visible in the record. Coordination is the defining failure of the international response. Not until July 2026 did the world's governments convene as a whole on AI, through the Global Dialogue on AI Governance mandated by General Assembly resolution A/RES/79/325 and held in Geneva. Everything before it (Bletchley, Seoul, Paris) was invitation-based and excluded most of the world from decisions affecting all of it. In the interim, governments leaned on AI Safety Institutes, which evaluate but do not bind, allowing measurement to substitute for regulation.

The UN has been working the military slice of AI for over a decade. The GGE on Lethal Autonomous Weapons Systems has been convened under the CCW since 2014, and the General Assembly has run a parallel track: resolutions 78/241 (2023) and 79/62 (2024), followed by a third consecutive resolution adopted by the First Committee on 6 November 2025 and by the plenary as Resolution 80/57 on 1 December 2025. The UN has been engaged for over a decade but its AI machinery is split between a mature but consensus-bound weapons track and a new, universal, but non-binding governance track.

What follows from this is that Australia cannot wait for international regimes to stabilise. Operational action at the micro and meso levels will be essential before macro-level regimes settle, and that action is available now through partnerships. The European Union, the United Kingdom, France, Germany, Canada and Singapore should become first-rank AI partners for Australia, alongside engagement with the US and China. Smaller and middle powers making the best pacts they can with one another is not a retreat from multilateralism; it is the layer that can move while the machinery is built.

Recommendation 4: Australia should adopt a layered AI strategy.

Expensive frontier models should be reserved for complex research, sophisticated agents and genuinely high-stakes tasks. Routine work should be routed to cheaper open-weight or specialised models according to cost, capability and risk. This approach can reduce expenditure, diversify supply and give Australian organisations greater control, while recognising that no model is neutral and that open-weight systems can carry hidden flaws, unwanted data flows or embedded assumptions.

Comparative advantage is most likely to come from the layers beneath frontier development: small language models trained on organisational data; safe AI applications; sovereign or trusted

compute; high-quality Australian data; evaluation and assurance capability; and a skilled workforce able to deploy AI responsibly. A public-option AI stack supported by public-interest laboratories with lawful access to quality data, compute and open-weight models would give government, researchers, start-ups and smaller businesses an affordable foundation for innovation.

Recommendation 5: Establish a trust and resilience architecture for autonomous systems through a comprehensive package.

- create an open standard for agent identity and accountability covering verifiable identity, attributable actions, revocable permissions and auditable delegation
- mandate secure-by-design controls including least privilege, sandboxing, offline testing, runtime monitoring and automatic termination of unauthorised behaviour;
- accelerate machine-speed defence alongside specialist human capability;
- recognise AI as an emerging critical-infrastructure sector;
- regulate data-centre infrastructure, hardware, software administration and data as separate objects with fit-for-purpose controls.

Recommendation 6: Establish an independent technology impact assessment capability for AI, sited outside the agencies that deploy AI and reporting to the Parliament, with its first assessment commissioned in 2027.

Affected communities are not the only stakeholders relevant to the decision about the value and risks of using new technologies. International practice recognises an important role for scientific experts and civil society organisations in environmental consultations of regional or national significance. Their expertise can inform issues such as national security, resources and climate impacts, economic viability, social impacts and future trends. This broader approach to environmental impact assessments has been established in jurisdictions for decades. It is also reflected in the Australian Government-funded research on technology impact assessment (TIA) for peace and stability and led in Australia by the Social Cyber Institute.³ That work promoted the following benchmarks for effective TIA:

- the scope and granularity of the technology being assessed
- the depth and detail of specialist input
- recognition of the significant role of non-technical social, political, legal and economic impacts
- the breadth and depth of stakeholder input
- the comprehensiveness of the analysis, which should include international and alternative perspectives
- the timeliness of the assessment
- a need for clarity about the ethical frameworks to be applied.

This two year study identified three key weaknesses in Australian approaches:

- the technology assessment work was largely left to parliamentary committees
- even this mechanism was rarely used

³ Social Cyber Institute, 'Technology Impact Assessment for Peace and Stability: A Comparative Study on Australia and India', June 2025, p. ii, https://www.socialcyber.co/files/ugd/15144d_c5acc66a4a014035a939a1b534f06822.pdf.

- parliamentary examination tended to be relatively superficial, focussing more on AI as an aggregated set of technologies rather than on the distinct implications of the many sub-fields of AI.

The habits of advanced TIA have so far not been evidenced consistently for analysis by governments, scholars or businesses in their investigation of AI impacts on Australia. Occasional inquiries by parliamentary committees, while important, are no substitute for mature and comprehensive technology impact assessment.

The frameworks that might have been applied (anticipatory, technology-wide, publicly consultative, and independent of the bodies deploying the technology) remain available but mostly unused for the case of AI adoption. That gap has been filled by an instrument of a different kind. The Policy for the responsible use of AI in government, version 2.0, effective 15 December 2025, requires agencies to conduct an AI use case impact assessment for in-scope use cases, designate an accountable owner for each, maintain an internal register shared with the Digital Transformation Agency every six months, and review deployed use cases annually. The AI Plan for the Australian Public Service 2025 requires every agency to appoint a Chief AI Officer.

Each requirement is defensible on its own terms. Together they have supplanted technology impact assessment by occupying its vocabulary. A use case assessment examines one system, in one agency, for one purpose, against the AI Ethics Principles. It cannot ask what a class of technology is doing to Australian society, what the cumulative effect of many deployments will be, or whether a deployment should proceed at all. It excludes incidental and lower-risk uses, which is where diffusion effects are found, and it does not reach private sector deployment, which is where most impact will land.

The assessor is also the adopter. Agencies describe the Chief AI Officer as championing adoption and propelling innovation; in some, that officer is the accountable official to whom risk assessments report. We apply to industry the test that an undertaking enforceable only by the party that gave it is a statement of intent. The same test should apply to government.

Recommendation 7: Establish a National AI Education Mission under the National AI Plan, jointly governed by the Commonwealth and the states, with scheduled milestone dates and published capability targets.

Deliver this plan through nationally recognised units of competency, shared subjects and micro-credentials that can be stood up in months, rather than through degree cycles that will not produce graduates until 2030.

The gap is not in technology courses; it is in everything adjacent to them. The evidence assembled by the Social Cyber Institute shows expansion concentrated in specialist AI degrees and AI specialisations within computing awards. Policy-facing provision is comparatively thin when measured against the volume of AI governance activity in Australian government. The demand side has already moved; the supply side has not. The APS AI Plan 2025 mandates foundational AI literacy for all APS staff and requires agencies to appoint Chief AI Officers. The people who will draft, administer and audit Australian AI regulation are, on present evidence, served in terms of formal AI policy education by a single graduate certificate.

Pace is the binding constraint, and the sequencing is wrong. Announcement dates are not commencement dates. The Monash Bachelor of CyberAI was announced in July 2026, opened

applications in August, takes its first cohort in Semester 1 2027 and graduates that cohort in 2030: four years from announcement to capability. That is an appropriate timeline for a profession but less relevant to short term needs in a policy crisis. Escalation means inverting the order of reform: fund the fast layer first and let the degrees follow.

Ambitions should be stated as numbers with dates and published. Begin with a national register of AI provision: we cannot reform what we cannot see. Possible actions could include:

- Every Chief AI Officer and SES officer should be newly credentialed by end-2027
- a mandatory secure-AI and adversarial machine learning unit should be included in every cyber security award by 2028
- policy, law and governance provision for AI analysis should be increased tenfold
- a public delivery dashboard, based on the UK model, could report against each commitment.

Even on the accelerated path proposed here, meaningful policy effects would be three to five years away: credentials issued in 2027 may change decisions in 2029. The broader transition will be protracted, running well beyond the life of any single plan or parliament. That is an argument for starting the fast layer immediately and for building AI governance expertise that survives ministerial and electoral cycles, and not for treating the degree pipeline as sufficient.

Australia cannot claim leadership in AI policy through symbolic gestures without massive investment in AI-related education and research, including about the technologies and the social implications. The initial reform budget would need to approach at least \$1 billion over the first five years, with a commitment to education (teaching) at least as great as that allocated to research.

Recommendation 8: Treat the political power and intent of the AI corporations as major policy variables in their own right and design Australian regulation so that it does not depend on the goodwill, self-restraint or stated principles of any firm.

The central dilemmas in AI policy are already highly politicised through the direct interventions and dispositions of a technology oligarchy loosely associating around anti-democratic dispositions. These executives now constitute a third set of powerful actors alongside governments and publics, conducting their own contest over AI policy often outside democratic process.

The intensity of the political onslaught mounted by technology company leaders against the political system in the United States over the past year exceeded what was anticipated even eighteen months ago. Taking the money and the political power that accompanies it together, there is reason for considerable scepticism about the prospects for good outcomes from US politics. Australia should not design its own arrangements on the assumption that the American political system will discipline these actors on Australia's behalf.

Australia should therefore legislate the expectations rather than publish them, require disclosure of political and lobbying expenditure by major AI providers operating here, fund evaluation capability that is independent of the firms being evaluated, and remove self-attestation from procurement.

Recommendation 9: Policy should privilege binding legal obligations over voluntary pledges and hold leading actors to standards of enforceability rather than intent.

The Institute's consistent position is that binding legal obligations should be preferred to voluntary industry commitments. Voluntary regimes allow the most capable private actors to choose their regulatory venue, delay implementation, or shape less than rigorous standards through lobbying. Safety institutes evaluate; they do not bind. Australia's own AI Safety Institute appears to allow the government to demonstrate seriousness without accepting any binding constraint on the firms it is meant to be governing.

Two international benchmarks are instructive. Illinois enacted the Artificial Intelligence Safety Measures Act (SB 315) on 6 July 2026, making it the third US state after California and New York to impose transparency, safety and reporting obligations on large frontier developers and the first US state to require an annual independent third-party audit of compliance rather than accept self-certification. Auditors must be free of financial conflicts of interest, critical safety incidents must be reported to the state within 72 hours, and whistleblower protections are built in. Separately, the UN Secretary-General has called on every major technology firm to power all data centres with 100 per cent renewable energy by 2030, and in June 2026 extended that call to full public disclosure of environmental footprint, including emissions, water and land use. Jurisdictions that have converted such positions into binding law (Illinois, California, New York, Minnesota, the European Union) are the appropriate comparators for Australia, and they should be studied immediately.

The preference for enforceability applies at each level, with a different obligation at each. Governments should legislate rather than publish expectations. Where Australia has accepted an argument in an international forum, as it did with the two-tier approach to autonomous weapons in 2023, it should accept the procedural step that follows. Peak bodies should stop treating voluntary codes as the endpoint of self-regulation and instead draft model clauses that are capable of being legislated, publish member compliance rather than member intent, and decline to offer governments a voluntary alternative to law as a means of deferring it.

Individual organisations should convert their own AI commitments into contractual and auditable form: warranties and audit rights in supplier agreements rather than vendor assurances, board-level accountability with named officers, and independent evaluation of AI systems by parties without a commercial interest in the result. The common test is simple, and each actor can apply it to itself. If an undertaking cannot be enforced by someone other than the party that gave it, it is a statement of intent and should be described as such rather than counted as a safeguard.

Recommendation 10: Build a movement: mobilise effective actors at home and internationally, and secure the political commitments needed to get them into the room.

Mobilisation of affected populations is a recurring prescription in the Institute's analysis of AI consequences for Australia. It can act as a countervailing force against both corporate and governmental inertia. Getting affected populations into the room for these conversations is fundamentally important, because the control dimension of the crisis is currently the weakest, and because the actors outside the AI corporate giants (citizens, workers, small business, civil society) are the ones whose influence a crisis most readily erodes.

The precedents are instructive on both the potential and the effort involved. Climate change required sustained work over decades to get the right people in the room, and the same is true of the major international disarmament movements. Neither succeeded through public concern alone; each

required institutional entry points, technical capacity within the movement, and political decisions to admit those voices to processes that had previously excluded them.

In the US, community mobilisation has succeeded in defeating a number of proposed data centre developments, and there is now a substantial body of local organising directed at AI infrastructure. The caution is that this is at a very early stage, and that local success against a specific development is not the same as influence over national policy.

The requirement applies at all three levels, with different work at each. Governments hold the decisive lever, and it is the practical bottleneck identified here: a commitment from political leaders to ensure the right people are in the room is both part of the challenge and part of the answer. That commitment should be concrete through measures such as standing representation for affected populations on AI advisory and governance bodies, funded participation rather than unpaid consultation, and the same access to officials that industry currently enjoys. Peak bodies are the natural aggregators, capable of converting dispersed concern into technically credible submissions, and of representing workforces and sectors that cannot individually sustain engagement. They should be resourced to do so.

Individual organisations, such as universities, unions, professional associations, community groups, and firms exposed to AI-driven change can supply the evidence base and the affected constituencies, and should be prepared to give evidence, publish their own exposure, and support staff and members to participate.

The point is not that mobilisation substitutes for regulation. It is that, absent a countervailing constituency, the control that democratic institutions have effectively outsourced to a small number of corporations will not be recovered.

Conclusion

Taken together, our account of the character of the AI policy crisis rests on a specific diagnosis: an existential, urgent, and deeply uncertain problem being mismanaged by institutions still operating in non-crisis mode, worsened by corporate capture, an unreliable and internally unstable US, and governments (Australia prominent among them) that talk about risk while quietly accommodating the industry causing it. Our prescription follows the same logic point for point: adopt genuine crisis-mode decision-making; use technology impact assessment as the methodological foundation for consultation and action; mobilise the populations actually affected; prefer binding law over voluntary pledges; act on a multi-front basis involving government, corporations, universities, community groups and even the military; and, in the absence of functioning multilateral machinery, be prepared for middle powers to build their own coalitions rather than wait for one to arrive.

In pursuing economic advantage, Australia's AI policy should treat sovereignty, national security, cyber security, and critical infrastructure resilience as inseparable priorities. AI may not create wholly new attack methods, but it will accelerate exploitation (especially at the juncture of associated frameworks), expose weaknesses faster, and enable autonomous systems to act at machine speed.

The central policy task is to build national AI capability without creating unmanaged dependence on foreign models, infrastructure or providers. This means focusing on the main control layers that will enhance sovereignty: data residency, model behaviour, inference, and auditability.

Australia's highly skilled workforce and strong R&D sector can be leveraged and enhanced still further to gain comparative advantage in generative and agentic AI, especially through upskilling the workforce, preparing enterprises, and accelerating technology capability.

Annex A

Background papers supporting this submission

- #1 Navigating the AI Policy Crisis
- #2 Maximising Self-Reliance in Cyber Security and Critical Infrastructure
- #3 Technology Impact Assessment
- #4 Impacts of AI on Cyber Workforce Development in the US, UK and Australia

Available at: <https://www.socialcyber.co/social-cyber-institute/submissions-and-notes>