



Committee Secretary
Senate Standing Committee on
Environment & Communications
PO Box 6100
Parliament House
Canberra ACT 2600

Electronic Frontiers Australia
ABN: 35 050 159 188
W: www.efa.org.au
E: Chair@efa.org.au

By electronic upload

24 June 2026

Dear Committee Secretary,

Re: Children's Online Privacy Code Submission

Electronic Frontiers Australia Inc. ("**EFA**") welcomes the opportunity to make this submission to the Senate Standing Committee on Environment and Communications' inquiry into Artificial Intelligence (AI) and Data Centres, referred by the Senate on 13 May 2026 on the motion of Senator Sarah Hanson-Young.

EFA's submission is contained in the following pages.

We would welcome the opportunity to provide further commentary if required. We consent to this submission being published on the APH web site.

About EFA

Established in January 1994, EFA is a national, membership-based, not-for-profit organisation representing Internet users concerned with digital freedoms and rights.

EFA is independent of government and commerce, and is funded by membership subscriptions and donations from individuals and organisations with an altruistic interest in promoting civil liberties in the digital context.

EFA members and supporters come from all parts of Australia and from diverse backgrounds. Our major objectives are to protect and promote the civil liberties of users of digital communications systems (such as the Internet) and of those affected by their use and to educate the community at large about the social, political, and civil liberties issues involved in the use of digital communications systems.

Yours sincerely,

John Pane
Chair
Electronic Frontiers Australia

ELECTRONIC FRONTIERS AUSTRALIA INC.

SUBMISSION TO THE SENATE STANDING COMMITTEE ON ENVIRONMENT AND COMMUNICATIONS

INQUIRY INTO ARTIFICIAL INTELLIGENCE AND DATA CENTRES

Lodged 24 June 2026

Submitting Organisation Electronic Frontiers Australia Inc. ABN 35 050 159 188

Authorising Officers John Pane, Chair

Correspondence chair@efa.org.au

Lodged to Committee Secretary Senate Standing Committee on Environment and Communications
PO Box 6100, Parliament House Canberra ACT 2600 ec.sen@aph.gov.au

Table of Contents

§	Title	P a g e
	Introduction	3
	Executive Summary	3
	Summary of Recommendations	5
1	About Electronic Frontiers Australia	7
2	Scope and Methodology	7
3	Terms of Reference (a) — Effectiveness of Existing Regulatory Frameworks	8
3.1	The current framework is voluntary, fragmented and ex post	8
3.2	Government deals with global AI companies are being concluded without ex ante safeguards	9
3.3	Overall regulatory critique	9
4	Terms of Reference (b) — Impacts on Communities, Industries, Environment, Water and Energy	10
4.1	Energy	10

4.2	Water — including the UN Water Bankruptcy framework	11
4.3	Communities	14
4.4	Labour markets	14
4.5	Industries and the broader AI ecosystem	15
5	A Framework for Ex Ante Regulation Drawn from the CHT <i>AI Roadmap</i>	15
6	Recommendations with Articulated Risks, Harms and Reasons	16
6.1	Legislative Moratorium on New AI "Deals" and Mega-Scale Data Centre Approvals	16
6.2	Ex Ante AI Safety, Accountability and Transparency Act	17
6.3	A Federally Enforceable Human Rights Act	18
6.4	Immediate Enactment of Privacy Act Tranche II	19
6.5	Binding Environmental, Water and Resource Law for Data Centres	21
6.6	Labour-Market Protection and Just Transition Package	23
6.7	Community Participation Rights and Local Council Veto	24
6.8	Independent AI Regulator (Australian AI Safety Commission)	25

6.9	Critical Infrastructure Coverage Extended to Large-Scale AI Compute	26
6.10	Mandatory Transparency of Government–AI Company Contracts	27
7	Conclusion	28
8	References	29

Introduction

Electronic Frontiers Australia Inc. ("**EFA**") welcomes the opportunity to make this submission to the Senate Standing Committees on Environment and Communications' inquiry into Artificial Intelligence (AI) and Data Centres, referred by the Senate on 13 May 2026 on the motion of Senator Sarah Hanson-Young.¹

EFA is Australia's longest-serving digital rights organisation. Since January 1994 we have advocated for the protection of human rights in a digital context, with continuous engagement on the AI policy cycle from the *Australia's AI Ethics Framework* consultation in 2019 through to the *Privacy Act Review* and the *Safe and Responsible AI in Australia* consultation, and onward to the present.^{2,3,4,5}

This submission addresses the terms of reference adopted by the Senate, in particular paragraph (a) (the effectiveness of existing regulatory frameworks for data centres, including deals between Government and global AI companies) and paragraph (b) (the impacts of AI and data centres on Australian communities, industries, environment, water and energy).¹

The submission is structured to read as a parliamentary submission while drawing on academic apparatus where the gravity of the policy choices before the Senate justifies more substantive evidentiary support. It is informed by EFA's three decades of digital-rights work; by the substantial body of civil-society and community evidence already lodged with the parallel New South Wales Legislative Council Public Accountability and Works Committee inquiry into data centres (pk=3169);⁶ by the United Nations University Institute for Water, Environment and Health's flagship 2026 reports on global water bankruptcy and the environmental cost of AI's energy use; and by the policy architecture proposed in the Center for Humane Technology's *AI Roadmap* (April 2026).^{7,8}

Executive Summary

EFA's position is unambiguous. **Australia must not permit further large-scale build-out of AI infrastructure, nor finalise further "deals" with global AI companies, until robust ex ante legislation is enacted that regulates AI safety, trustworthiness, and the impacts of AI and its supporting data-centre infrastructure on labour markets, the environment, water supplies, and power supplies and prices.**

Current Australian policy architecture is structurally incapable of constraining the clearly foreseeable harms identified in this submission. It comprises the *National AI Plan* (December 2025); the Commonwealth's voluntary *Expectations of data centres and AI infrastructure developers* (23 March 2026);⁹ the voluntary *Policy for the responsible use of AI in government 2.0* (15 December 2025);¹⁰ and a patchwork of state planning instruments. None of these instruments creates legally enforceable rights, imposes civil penalties, establishes an independent regulator, or affords any meaningful community protections or remedies. The 2026 Australian Community Attitudes to Privacy Survey ("ACAPS 2026") records that 87 per cent of Australians are more concerned about their privacy than they were five years ago, only one in ten believe organisational data practices are usually fair and public trust in the systems underpinning modern digital life has hit an absolute floor. Just 4% of Australians believe AI companies are worthy of their trust, while trust in social media companies languishes at a mere 3%. Australians are deeply concerned about the rollout of AI, digital identification, the predilections and influence of Big Tech, and biometrics without the protection of strong regulatory frameworks being in place. This is a digital rights crisis for Australia and constitutes a social-licence collapse that must weigh heavily on this Committee.¹¹

The case for an ex ante legislative architecture has been strengthened by two flagship publications of the United Nations University Institute for Water, Environment and Health ("UNU-INWEH"). The first, *Global Water Bankruptcy: Living Beyond Our Hydrological Means in the Post-Crisis Era* (20 January 2026),¹² formally introduces the policy category of *water bankruptcy*; the conjunction of insolvency (withdrawal beyond renewable inflows) and irreversibility (damage to water-related natural capital that cannot feasibly be undone). The second, *Environmental Cost of AI's Energy Use: Carbon, Water and Land Footprints* (3 June 2026),¹³ applies that lens to AI infrastructure, projecting that the global water footprint of AI-driving data centres by 2030 will equal the basic annual domestic water needs of all 1.3 billion people in Sub-Saharan Africa. Given we are the world's driest continent and appear posed to enter another *el nino* season, allocating Australia's already-stressed water entitlements to foreign hyperscaler cooling at the very moment the United Nations has declared a water-bankruptcy era is, in EFA's submission, regulatory negligence.

The Center for Humane Technology's *AI Roadmap* (April 2026) provides a coherent ex ante legislative platform to manage the risks and harms of AI on Australian society.⁷ EFA endorses its substantive provisions and recommends that the Committee draw on its specific architecture; mandatory pre-deployment testing, statutory product liability for AI, federal whistleblower protections, chatbot design standards, a statutory right to cognitive liberty, and antitrust enforcement targeting AI consolidation.

EFA accordingly makes ten substantive recommendations, each accompanied by an articulation of the specific risks and harms it addresses and the reasons supporting its adoption. Together they form an integrated package: a moratorium that buys time; an AI Safety, Accountability and Transparency Act and a federal Human Rights Act that provide the substantive rights floor; Privacy Act Tranche II that supplies the data layer; binding environmental and resource law that protects water, energy and community amenity; a labour-market and just-transition package that protects workers; community participation rights that restore democratic legitimacy; an independent regulator that institutionalises sceptical oversight; expanded SOCI coverage that addresses sovereignty and concentration risk; and mandatory transparency of Government–AI company contracts that restores democratic accountability.

EFA respectfully urges the Committee to recommend that the Senate adopt this framework before the parallel NSW inquiry reports on 30 September 2026,¹⁴ and before further deals with global AI companies foreclose Australia's regulatory options.

Summary of Recommendations

For ease of reference, EFA's ten recommendations are set out in concise form below. Each is developed at length in Part 6 of this submission.

Recommendation 1 — Legislative Moratorium on New AI "Deals" and Mega-Scale Data-Centre Approvals. That the Commonwealth refrain from concluding further MoUs, contractual arrangements, tax concessions, grid-priority allocations, water allocations or land grants with global AI companies, and from fast-tracking new large-scale (>30 MW IT load) data-centre approvals, until the ex ante legislative framework set out in this submission is enacted and operational. A twelve-month formal review point is recommended.

Recommendation 2 — Ex Ante AI Safety, Accountability and Transparency Act. Enact a binding *Artificial Intelligence (Safety, Accountability and Transparency) Act* which is aligned to the human rights protecting EU AI Act and incorporating mandatory pre-deployment safety testing of high-impact AI systems; statutory product liability for AI harms with a rebuttable presumption of causation; federal whistleblower protections; statutory chatbot and companion-AI design standards with enhanced protections for minors; prohibition of engagement-maximising dark patterns; statutory limits on biometric inference; and tiered obligations calibrated to system, risk and harms impact.

Recommendation 3 — A Federally Enforceable Human Rights Act. Legislate a federally enforceable Human Rights Act incorporating the ICCPR and ICESCR rights, with standing for individuals and representative bodies; a statutory interpretation provision; and a positive duty on Commonwealth public authorities to act compatibly with rights in the procurement, design and deployment of AI systems.

Recommendation 4 — Immediate Enactment of Privacy Act Tranche II. Legislate Tranche II reforms including a comprehensive biometric protection regime; strong cross-border transfer limits; abolition of the small-business and employee-records exemptions; substantive (not merely disclosure-based) protections in respect of automated decision-making; a fair-and-reasonable test; and greater OAIC resourcing.

Recommendation 5 — Binding Environmental, Water and Resource Law for Data Centres. Convert the Commonwealth's voluntary *Expectations* into binding statute under a Data Centre Infrastructure Act, drawing expressly on the UNU-INWEH "bankruptcy management" framework — including mandatory closed-loop cooling, operator-funded renewable underwriting, statutory water-allocation limits, a ban on potable cooling above defined thresholds, a "no new gas, no coal-life-extension" condition, mandatory natural-capital protection, transparent quarterly water accounting, waste-heat recovery, public reporting of PUE/WUE/CUE and Scope 1–3 emissions, and tied EPBC Act assessment.

Recommendation 6 — Labour-Market Protection and Just Transition Package. Enact a Labour-Market Protection and Just Transition (AI) Act inverting the capital-over-labour tax bias, conferring statutory consultation and transparency rights on workers, prohibiting fully-automated termination decisions, imposing a Just Transition Levy on large-scale AI infrastructure operators, and mandating workforce-impact assessments.

Recommendation 7 — Community Participation Rights and Local Council Veto. Legislate mandatory pre-approval community consultation for all data-centre developments above 30 MW IT load; a statutory veto right for local councils where formal moratoria are in place; culturally appropriate engagement with First Nations communities including free, prior and informed consent; a "community benefit" requirement; and post-approval community monitoring rights.

Recommendation 8 — Independent AI Regulator (Australian AI Safety Commission). Establish a properly resourced, independent statutory AI regulator with legally enforceable investigative, audit and standards-setting powers over both AI systems and data-centre infrastructure; authority to require pre-deployment testing and incident reporting; own-motion investigation powers; and civil-penalty jurisdiction calibrated to hyperscaler revenue, i.e. aligned to the EU's approach to pecuniary penalties under the GDPR.

Recommendation 9 — Critical Infrastructure Coverage Extended to Large-Scale AI Compute. Extend the *Security of Critical Infrastructure Act 2018* to large-scale AI compute facilities, with public-interest as well as national-security criteria, mandatory ownership and supply-chain disclosure, and ministerial direction powers.

Recommendation 10 — Mandatory Transparency of Government–AI Company Contracts.

Mandate public disclosure on a central AI Deals Register, within thirty days of execution, of the substantive terms of all Commonwealth contracts, MoUs, grants, tax concessions, energy allocations, water allocations and land deals involving global AI companies and large-scale data-centre operators.

1. About Electronic Frontiers Australia

Electronic Frontiers Australia Inc. was incorporated in January 1994 as a national, membership-based, not-for-profit organisation committed to the promotion and protection of human rights in a digital context. EFA is constitutionally and operationally independent of government and commerce, funded predominantly by membership subscriptions and donations from individuals and organisations with an altruistic interest in the advancement of civil liberties online.³

Over more than three decades EFA has lodged formal submissions on virtually every substantive Australian policy instrument touching on digital rights, including continuously across the AI policy cycle: from the *Australia's AI Ethics Framework* consultation in 2019, through the *Privacy Act Review* in April 2023, to the *Safe and Responsible AI* consultation in August 2023, and onward to the present.^{2,4,5} This submission represents EFA's institutional position as adopted by its Board in light of its historical policies.

2. Scope and Methodology

This submission addresses paragraphs (a) and (b) of the Senate's terms of reference and the "any other related matters" provision in paragraph (c). The reporting date set by the Senate is 16 November 2026.^{1, 15}

The evidence base draws on five major corpora. The first comprises primary Commonwealth and State policy instruments published in the period 2023–2026, including the *Expectations of data centres and AI infrastructure developers* (23 March 2026); the *Policy for the responsible use of AI in government 2.0* (15 December 2025); and the National AI Plan (December 2025).^{9, 10} The second comprises civil-society and community submissions to the parallel NSW Legislative Council inquiry (pk=3169), which received approximately seventy submissions before its 27 March 2026 close.^{6, 16} The third is the Center for Humane Technology's *AI Roadmap* (April 2026).^{7, 8} The fourth comprises the UNU-INWEH flagship reports *Global Water Bankruptcy* (20 January 2026) and *Environmental Cost of AI's Energy Use* (3 June 2026).^{12, 13} The fifth is EFA's own historical submission record.²

3. Terms of Reference (a): Effectiveness of Existing Regulatory Frameworks

3.1 The current framework is voluntary, fragmented and ex post

The Commonwealth's principal current instrument for governing AI infrastructure is the *Expectations of data centres and AI infrastructure developers*, released on 23 March 2026. On the Department of Industry, Science and Resources' own characterisation, these are administrative "expectations" forming "the foundation of [developers'] social licence to operate"; non-aligned, energy-intensive proposals will merely "not be prioritised" in Commonwealth regulatory assessments.⁹ The instrument lacks any of the hallmarks of binding regulation. There is no statutory cause of action, no civil penalty regime, no community standing, no independent regulator, no merits-review pathway and no offence provision. Its five Expectations — prioritise national interest, support the energy transition, use water sustainably, invest in skills, and strengthen local research capability — operate as administrative aspirations untethered from enforceable rights, penalties or community remedies.¹⁷

The *Policy for the responsible use of AI in government 2.0*, effective from 15 December 2025, compounds rather than cures this gap. It applies only to non-corporate Commonwealth entities, is self-administered, and creates internal procedural duties such as accountable officials, AI use statements, transparency statements but without conferring enforceable rights on affected individuals or any external enforcement mechanism.¹⁰

EFA's longstanding position, articulated continuously since its 13 August 2023 submission to the *Safe and Responsible AI* consultation, is that the most important aspect of responsible AI regulation is the introduction of a federally enforceable human-rights framework, supported by meaningful amendment of the *Privacy Act 1988* and the adoption of individual and collective rights of action within a graduated model of regulation.³ That foundational architecture remains absent. Tranche II of the Privacy Act reforms remains inexplicably stalled, even as the Productivity Commission has provocatively suggested *weakening* privacy law in the name of "productivity."¹¹

3.2 Government deals with global AI companies are being concluded without ex ante safeguards

The terms of reference expressly direct the Committee to inquire into "existing and future deals between the Government and global AI companies."¹ EFA's concern is structural rather than incidental. On 31 March 2026, Anthropic signed a Memorandum of Understanding with the Australian Government and is reportedly "exploring data centre investments and energy partnerships aligned with the government's new expectations framework."¹⁸ Comparable engagement is occurring with other hyperscalers. At the moment of negotiation, the Australian community possesses no legislated rights to algorithmic transparency, redress, biometric privacy, labour protection, or environmental amenity that bind these counterparties. The deals are concluded in a statutory vacuum.

This inverts sound and appropriate regulatory sequencing. As the Center for Humane Technology argues by analogy with aviation, pharmaceuticals and food safety, technologies of comparable consequence are subject to "independent oversight, rigorous testing, and real transparency" before mass deployment, not after measurable harm has been catalogued.⁷

3.3 Overall regulatory critique

EFA's overarching view is threefold. First, the existing so-called regulatory framework is not effective at managing the growth of data-centre infrastructure or the deployment of AI systems built upon it. Second, the conclusion of further deals with global AI companies should be paused until ex ante legislation is in force. Third, Australia's substitution of social-licence rhetoric for statutory rights is constitutionally and democratically inadequate to the potential risks, harms and consequences of the technology now being foisted across virtually every domain of Australian public and economic life.

4. Terms of Reference (b): Impacts on Communities, Industries, Environment, Water and Energy

EFA supports and amplifies the substantial evidentiary record assembled by civil-society and community organisations for the parallel NSW Legislative Council inquiry.⁸

4.1 Energy

The Australian Energy Market Operator forecasts that data centres will consume approximately eleven per cent of New South Wales' grid-supplied electricity by 2030.¹⁹ ²⁰ The twenty-two NSW data-centre applications received in the first quarter of 2026 alone represent 3.67 gigawatts of combined demand — sufficient to power 1.1 million NSW homes.²⁰ Friends of the Earth Australia, in its submission of 22 May 2026, warns the industry is being used to justify prolonging coal-fired power stations and new gas extraction, and that even renewably-powered facilities consume new generation that would otherwise displace coal from the grid.²⁰ Some operators are proposing on-site gas-fired generation, in direct contradiction of New South Wales and Commonwealth emissions trajectories.²⁰

The UNU-INWEH June 2026 report *Environmental Cost of AI's Energy Use* situates this Australian growth trajectory within an emerging planetary-scale problem.¹³ Global data centres consumed 448 terawatt-hours of electricity in 2025, more than all but ten countries in the world, and are projected to consume 945 terawatt-hours by 2030, nearly triple the combined annual electricity use of Pakistan, Bangladesh and Nigeria, countries that together house more than 650 million people. The share of data-centre energy attributable to AI specifically is forecast to grow from approximately twenty per cent today to forty per cent by 2030, with day-to-day inference (rather than model training) accounting for the overwhelming majority of that demand.

4.2 Water — Hyperscale AI Infrastructure in an Era of Global Water Bankruptcy

4.2.1 The Australian baseline

A single large data centre can consume up to five million gallons — approximately nineteen megalitres — per day, equivalent to the potable demand of a town of 10,000–50,000 people.²⁰ Australian data centres collectively are estimated to consume some forty-seven billion litres per year, with projections that AI workloads alone could require the equivalent of twenty-five per cent of Sydney's drinking water by 2050.²⁰ Sydney Water estimates aggregate data-centre demand could reach 250 megalitres per day by 2035, increasing utility-wide demand by approximately twenty-five per cent.¹⁹ Penrith City Council has called for a moratorium on further approvals until water and power impacts are "fully understood."¹⁶

4.2.2 The UNU Global Water Bankruptcy framing (January 2026)

These Australian figures must be read against the most recent and influential international assessment of water governance in a generation. On 20 January 2026, the United Nations University Institute for Water, Environment and Health released its flagship report *Global Water Bankruptcy: Living Beyond Our Hydrological Means in the Post-Crisis Era*, authored by UNU-INWEH Director Professor Kaveh Madani.¹² The report formally introduces and defines a new policy category, "**water bankruptcy**", characterised by the conjunction of two cumulative criteria: insolvency, meaning the persistent withdrawal and pollution of water beyond renewable inflows and safe depletion limits; and irreversibility, meaning damage to key water-related natural capital (wetlands, lakes, aquifers, glaciers, soils) of a kind that makes restoration to historical baseline conditions infeasible.

The report's central message is that the era has changed. The traditional vocabulary of "water stress" and "water crisis" no longer adequately describes the condition of many of the world's basins; what is emerging is a post-crisis bankruptcy state in which historical normality cannot be recovered. As Professor Madani states, "For much of the world, 'normal' is gone." Natural water savings are being depleted at speed: more than half of the world's large lakes have declined since the early 1990s, and approximately thirty-five per cent of the world's natural wetlands have been lost since 1970. The human toll is already vast, nearly three-quarters of the world's population live in countries classified as water-insecure or critically water-insecure; around four billion people experience severe water scarcity for at least one month each year; and drought impacts now cost an estimated US\$307 billion annually.¹²

Of note for this Committee, the UNU report finds that the benefits and burdens of overshoot are profoundly unequal. The burdens fall disproportionately on smallholder farmers, Indigenous Peoples, low-income urban residents, women and youth, while the benefits of overuse accrue to more powerful actors, a distributional pattern that maps with uncomfortable precision onto the allocation of Australian water entitlements to foreign hyperscaler infrastructure.¹² The report concludes by proposing a structured "recovery plan" framework predicated on three pillars directly relevant to AI infrastructure regulation: transparent water accounting at facility and basin level; enforceable limits on withdrawals and pollution; and the protection of water-related natural capital, including aquifers, wetlands, soils, rivers and glaciers.

4.2.3 The UNU follow-up on AI's water footprint (June 2026)

On 3 June 2026, UNU-INWEH released a follow-up flagship report, *Environmental Cost of AI's Energy Use: Carbon, Water and Land Footprints*, which extends the bankruptcy framework into the AI domain.¹³ It concludes that AI's environmental cost is being systematically mismeasured because most existing assessments focus narrowly on the carbon emissions associated with training large models and ignore water and land footprints. By 2030, the global water footprint of AI-driving data centres will equal the basic annual domestic water needs of all 1.3 billion people in Sub-Saharan Africa. In 2025 alone, global data-centre electricity production consumed approximately 1.2 trillion gallons (around 4.5 trillion litres) of water. The report further finds that AI infrastructure is concentrated in a small number of hubs, exacerbating water stress in the very basins least able to bear it.¹³

4.2.4 Application to Australia

The Australian implications are clear, direct, severe and inescapable. Australia is the world's driest inhabited continent and is already experiencing climate-amplified drought, declining streamflow in the Murray–Darling Basin, and groundwater depletion in the Great Artesian Basin and other aquifers. Allocating large fractions of remaining hydrological resources to hyperscaler cooling at the precise moment the United Nations has formally declared a water-bankruptcy era is, in EFA's submission, regulatory negligence. Once aquifers are over-drawn and wetlands lost, the UNU report establishes that the damage may be irreversible. Data-centre approvals must therefore be treated as decisions with permanent, intergenerational hydrological consequences, not as ordinary planning matters.¹²

The UNU water equity findings translate directly onto Australian terrain. The "more powerful actors" benefitting from water overuse will, on current settings, be foreign hyperscalers; the "smallholder farmers, Indigenous Peoples, low-income urban residents, women and youth" disproportionately bearing the burden will be located in the Murray–Darling Basin, Western Sydney, regional Queensland, and remote First Nations communities. The UNU "bankruptcy management" framework provides a ready-made template containing transparent accounting, enforceable limits, protection of natural capital and equity-oriented transitions that this Committee should recommend Parliament adopt in any binding data-centre water management statute. This is reflected in Recommendation 5 (Part 6.5) below.

4.3 Communities

The Sweltering Cities submission to the NSW inquiry documents that data centres operate as urban heat islands in already heat-stressed Western Sydney suburbs such as Marsden Park, Eastern Creek and Kemps Creek — where temperatures routinely run up to ten degrees Celsius higher than other parts of Sydney and where facility ambient temperatures of 35–45°C persist day and night.²¹

Sweltering Cities, with the Antipoverty Centre and Parents for Climate, has lodged a separate "*Stop the Bill Shock*" submission warning of cost-shifting to households via energy and water tariffs.²¹ Penrith, the City of Ryde, and other councils are calling for moratoria or caps, citing displacement of housing capacity, drains on water and power, and minimal long-term local employment.¹⁹ The Justice and Equity Centre's submission recommends a public-interest regulatory architecture differentiated by facility purpose and function, rather than a homogeneous fast-track regime.²² The ARC Centre of Excellence

for Automated Decision-Making and Society contributes rigorous interdisciplinary evidence on ecological and multispecies impacts.²³

4.4 Labour markets

The Center for Humane Technology's *AI Roadmap* documents that leading AI companies are "racing to build systems that can perform valuable human tasks — and treating mass job loss as an inevitable side effect of innovation," with disruption already visible across warehousing, knowledge work, journalism, translation, customer service and creative industries.⁷ The Roadmap accordingly recommends tax incentives that flip the present capital-over-labour bias in favour of worker retention and upskilling.⁸ The data-centre construction boom itself is colliding with an estimated 300,000 construction-worker shortfall by mid-2027, with five competing sectors fighting for the same talent — undermining the assumption that data centres are net job creators.¹⁸ Operational employment at large facilities is structurally low relative to the land, water and energy footprint of those facilities.²⁴

4.5 Industries and the broader AI ecosystem

The CHT Roadmap identifies the AI industry's dependence on extracting value from "people's data, content, labour, and most private thoughts" without adequate legal constraint.⁷ Australian creative industries face direct harm. Copyright is "under pressure" from a "growing push to weaken or sidestep copyright protections in favour of training large language models," as EFA Vice-Chair Julian Watchorn observed in EFA's April 2026 World IP Day commentary.¹¹ Without statutory protections, Australian SMEs, public agencies and researchers will be locked into asymmetric dependencies on a small number of foreign hyperscalers, undermining digital sovereignty, a concern EFA has previously raised in connection with the *Digital ID Act 2024*'s abandonment of data-localisation requirements.²⁵

5. A Framework for Ex Ante Regulation Drawn from the CHT *AI Roadmap*

The Center for Humane Technology's *AI Roadmap* (April 2026) offers a coherent legislative platform structured around seven principles operating across norms, laws and product design.^{7, 8} EFA endorses its framing and recommends that this Committee adopt its substantive provisions. These include mandatory pre-deployment testing of high-impact AI systems; federal whistleblower protections for AI employees and contractors; statutory product liability for AI; express rejection of arguments that AI outputs are protected speech or that AI products are legal persons; chatbot design standards with enhanced protections for minors; tax incentives that reward worker retention and upskilling rather than capital substitution; a statutory right to cognitive liberty paired with firm limits on biometric surveillance; international red lines on recursive self-improvement and autonomous weapons; and upgraded antitrust

enforcement targeting AI consolidation. Each of these provisions is reflected in the recommendations that follow.

6. Recommendations with Articulated Risks, Harms and Reasons

6.1 Legislative Moratorium on New AI "Deals" and Mega-Scale Data-Centre Approvals

Recommendation

That the Commonwealth refrain from concluding further Memoranda of Understanding, contractual arrangements, tax concessions, grid-priority allocations, water allocations or land grants with global AI companies, and from fast-tracking new large-scale (greater than 30 megawatt IT load) data-centre approvals, until the ex ante legislative framework set out in this submission is enacted and operational. A formal review point after twelve months is contemplated to test whether legislative progress justifies a graduated lifting of the moratorium.

What are the relevant risks and harms?

The first risk is regulatory and constitutional lock-in. Long-dated MoUs, tax holidays, sovereign-risk indemnities, take-or-pay grid contracts and exclusive water allocations are notoriously difficult to unwind once executed. They generate sovereign-risk arguments against subsequent regulatory tightening, expose the Commonwealth and the States to compensation claims under investor-protection clauses, and foreclose policy options for successor governments. The Anthropic MoU of 31 March 2026 is the leading exemplar of an arrangement concluded before the domestic legal architecture is in place.¹⁸

The second risk is unacceptable asymmetry of bargaining power. Australian Commonwealth and State negotiators currently meet hyperscaler counterparties, entities with market capitalisations exceeding Australia's GDP, **without** the leverage of a binding domestic rights framework, a clear regulatory baseline, or interoperable European-style standards. The result is contractual terms that systematically under-price externalities and over-attribute value to the counterparty.

The third risk is the crowding out of public-interest infrastructure. Each gigawatt of compute load consumes grid headroom, transmission queue priority, water entitlements, construction labour, planning bandwidth and political attention that would otherwise be available for housing, public transport, hospitals, schools and the renewable-energy transition. The opportunity cost is real and presently unmeasured.

The fourth risk is the pre-emption of democratic deliberation. Concluding strategic AI deals before the Parliament has legislated the framework that should constrain them inverts the constitutional order and presents the Senate with a fait accompli rather than well considered choices that serve the public good and national interest.

Reasons supporting the recommendation

Sound regulatory sequencing requires law before deployment, as the Center for Humane Technology argues by analogy with aviation and pharmaceuticals.⁷ The recommendation is congruent with the moratorium called for by Penrith City Council and supported in substance by submissions to the parallel NSW inquiry.^{16, 19} A short, principled pause now will avoid decades of remediation cost, sovereign-risk exposure and rights-based litigation later. Critically, the moratorium is conditional and time-limited: it ends when the legislation is in force, providing industry with a clear pathway rather than indefinite uncertainty.

6.2 Ex Ante AI Safety, Accountability and Transparency Act

Recommendation

Enact a binding *Artificial Intelligence (Safety, Accountability and Transparency) Act* incorporating mandatory pre-deployment safety testing and risk-management standards for high-impact AI systems with standardised public reporting; statutory product liability for AI harms, including a rebuttable presumption of causation where opacity prevents claimants from tracing fault; federal whistleblower protections for all AI employees and contractors; statutory chatbot design standards with enhanced obligations for systems likely to be accessed by minors; a prohibition on engagement-maximising dark patterns; and statutory limits on biometric inference. Establish risk based tiered obligations calibrated to system impact, with frontier systems subject to the most stringent regime.⁸

What are the relevant risks and harms?

The first risk is unsafe deployment at population scale. AI systems are presently released into the Australian market with no equivalent of pre-market approval by the Therapeutic Goods Administration or design certification by the Civil Aviation Safety Authority, notwithstanding their direct use in healthcare diagnosis, education, hiring, credit, policing, child protection and welfare administration. The asymmetry with comparably consequential technologies is indefensible.

The second risk is the liability vacuum and consequent under-compensation of victims. When an AI system causes harm — a wrongful credit denial, a missed diagnosis, a defamatory hallucination, a psychosocial injury to a child — claimants face profound difficulty attributing legal responsibility across the developer, deployer and integrator chain. The CHT *AI Roadmap* warns this is no accident; some operators argue AI outputs are protected speech, or that AI itself should be a legal person.²

The third risk is the suppression of internal safety knowledge. Without statutory whistleblower protections, engineers and researchers who witness unsafe practices face career-ending retaliation, non-disclosure suits and equity claw-backs. Regulators consequently lose access to the only people with first-hand visibility into emerging risks.⁸

The fourth risk is psychosocial harm from engagement-maximising chatbots and companion AI, including emotional dependency, parasocial attachment, sexualised role-play, encouragement of self-harm and reinforcement of disordered thinking — harms that disproportionately affect minors and vulnerable adults.

The fifth risk is discriminatory and exclusionary outcomes. Without statutory testing obligations, systemic bias against protected attributes remains undetected at deployment and is only addressed (if at all) reactively through individual litigation.

Reasons supporting the recommendation

Ex ante product-safety regulation is the default for technologies of comparable consequence, including aviation, pharmaceuticals, medical devices, electrical goods and motor vehicles. AI is the anomaly, not the rule, and there is no principled reason to treat it differently. Clear statutory product liability produces the price signal voluntary "expectations" cannot.⁹ Federal whistleblower frameworks have been demonstrably effective in financial services, aviation safety and the corporate sector; the regulatory logic transfers cleanly to AI.⁸ A tiered structure preserves proportionality and avoids stifling low-risk uses of AI in research, accessibility and small-business productivity.

6.3 A Federally Enforceable Human Rights Act

Recommendation

Legislate a federally enforceable Human Rights Act underpinning all AI governance, including justiciable rights to equality, privacy, expression, peaceful assembly, social security and education; a statutory duty on public authorities to act compatibly with rights; a public-interest standing rule; and an interpretive provision binding the construction of all AI-related legislation. This recommendation continues EFA's consistent position from 2019 onward.³⁵

What are the relevant risks and harms?

The principal risk is the absence of a constitutional or statutory backstop against algorithmic harm. Unlike comparable democracies, Australia possesses no overarching enforceable rights instrument against which AI deployments can be tested. Algorithmic discrimination in social services, biased policing tools, automated welfare decisions and biometric surveillance can presently proceed unchecked at the federal level.

A related risk is the repetition of Robodebt-class harms. The Robodebt scheme demonstrated that automated decision-making by government, absent enforceable rights, can devastate hundreds of thousands of lives, including driving people to suicide, before judicial intervention. The Royal Commission's findings should not have to be re-litigated in every successive automated programme.

A third risk is sectoral patchwork failure. Privacy law, anti-discrimination law and administrative law each catch only fragments of AI harm, and the seams between them are precisely where the most consequential harms sit: intersectional discrimination, opaque decision chains and cross-jurisdictional impacts. A federal Human Rights Act provides the connective tissue.

A fourth risk is inconsistency across Australian jurisdictions. Victoria, Queensland and the ACT have legally enforceable human-rights instruments; other States and the Commonwealth do not. The result is a postcode lottery for digital rights protection.

Reasons supporting the recommendation

EFA has held this position consistently since the 2019 *AI Ethics Framework* consultation and reiterated it in the August 2023 *Safe and Responsible AI* submission.³⁵ The CHT *AI Roadmap*'s right to cognitive liberty, expanded publicity rights and biometric limits all presuppose an enforceable rights framework.⁸ The Australian Human Rights Commission, the Law Council of Australia, the National Children's Commissioner and major civil-society bodies are aligned in calling for federal rights legislation. A federal Human Rights Act provides the interpretive lens through which courts will read every other piece of AI legislation Parliament enacts, ensuring coherence across the statute book.

6.4 Immediate Enactment of Privacy Act Tranche II

Recommendation

Legislate Tranche II of the Privacy Act reforms immediately, including a direct individual right of action; representative and class-style collective rights of action; a comprehensive biometric protection regime; strong cross-border transfer limits with adequacy assessment; abolition of the small-business exemption; abolition of the employee-records exemption; meaningful disclosure and contestability rights in respect of automated decision-making consistent with the 10 December 2026 ADM disclosure deadline; and a statutory tort for serious invasions of privacy fully resourced for enforcement.¹¹ ²⁶

What are the relevant risks and harms?

The first risk is training-data extraction without consent. AI foundation models are trained on Australians' personal information — much scraped from public-facing platforms, scholarly publications, professional profiles and social media — without lawful consent, attribution or recourse. The 2026 Australian Community Attitudes to Privacy Survey finding that only one in ten Australians believe organisational data practices are fair captures the resulting social-licence collapse.¹¹

The second risk is biometric mission creep. Face, voice, gait, keystroke, gaze and physiological data are increasingly captured by AI systems in retail loyalty programmes, public transport, education, workplace surveillance and law enforcement, without a comprehensive statutory regime. Once collected, biometric data is functionally permanent and uniquely linkable — unlike a password, a face cannot be revoked.

The third risk is cross-border transfers to hyperscaler jurisdictions. Personal data routinely leaves Australian jurisdiction for processing, including to jurisdictions with weaker protections or extraterritorial intelligence reach. The hollowing-out of data-localisation safeguards under the *Digital ID Act 2024* exemplifies the same risk at a different layer of the stack.²⁵

The fourth risk is opaque automated decision-making. The 10 December 2026 ADM disclosure deadline is approaching, but disclosure alone does not equate to rights. Australians need substantive protections — a right to a human decision in significant matters, a right to an explanation, and a right to contest — that disclosure cannot provide.¹¹

The fifth risk is the corrosive effect of the trust deficit on the broader digital economy. Privacy concern is now so widespread that it suppresses participation in legitimate digital services and reduces the social licence for all AI deployment.¹¹

Reasons supporting the recommendation

The reforms have been consulted to exhaustion through the Privacy Act Review Report and exposure-draft process; further delay is a political choice and does not serve the Australian people.²⁶ A genuine individual and collective right of action is the only credible enforcement mechanism in a digital economy where the OAIC's resources cannot scale to the breach load. EFA's longstanding view is that the Privacy Act forms the foundational layer of AI regulation; without it, no other AI law will work properly.⁴ Tranche II in practice embeds the CHT Roadmap's cognitive-liberty and biometric protections within Australian law without requiring a separate legislative vehicle.⁸

6.5 Binding Environmental, Water and Resource Law for Data Centres

Recommendation

Convert the Government's voluntary *Expectations of data centres and AI infrastructure developers* into binding statute under a Commonwealth Data Centre Infrastructure Act, expressly drawing on the UNU-INWEH "bankruptcy management" framework. The to-be Act should mandate closed-loop or non-potable cooling for new and expanded data centres above a defined scale, with civil penalties and operating-licence consequences for non-compliance. It should require operator-funded underwriting of additional renewable generation matched temporally to load, with full cost recovery for grid connections and a prohibition on socialising those costs across residential and small-business bill payers. It should impose statutory water-allocation limits, mandatory disclosure of consumption, and a binding ban on the

use of potable supply for cooling above defined thresholds. It should impose a binding "no new gas, no coal-life-extension" condition on AI infrastructure load, including a prohibition on bespoke fossil-fuel generation behind the meter. It should require mandatory water-related natural capital protection consistent with the UNU *Global Water Bankruptcy* framework, including no-net-loss obligations for wetlands, aquifers and riparian zones. It should require transparent water accounting at facility and basin level, published quarterly; mandatory waste-heat recovery where technically feasible; public reporting of PUE, WUE, CUE and Scope 1–3 emissions; and tied EPBC Act assessment of major facilities.⁹ [12,17](#)

What are the relevant risks and harms?

The first risk is the acceleration of fossil-fuel dependency. Friends of the Earth Australia warns the AI infrastructure build-out is being used to justify prolonging coal-fired power and authorising new gas extraction. Some operators are proposing on-site gas generation to bypass grid constraints, in direct frustration of state and Commonwealth emissions targets.²⁰

The second risk is drinking-water diversion at community scale. A single hyperscale facility can consume up to nineteen megalitres per day — equivalent to a town of 10,000–50,000 people. Australian data centres are estimated to consume forty-seven billion litres per year and could require the equivalent of twenty-five per cent of Sydney's drinking water by 2050, with Sydney Water projecting 250 megalitres per day of data-centre demand by 2035.^{19,20}

The third risk — most fundamental to this submission — is the allocation of water during a United Nations-declared bankruptcy era. UNU-INWEH's *Global Water Bankruptcy* report formally establishes that the world has entered a post-crisis era of irreversible hydrological loss. Allocating Australian water entitlements to foreign hyperscaler cooling at this moment — when more than half the world's large lakes have declined since the early 1990s, thirty-five per cent of natural wetlands have been lost since 1970, four billion people experience severe water scarcity annually, and drought costs reach US\$307 billion a year — is fundamentally at odds with the international scientific consensus.¹²

The fourth risk is the alarming trajectory of global AI water consumption. UNU-INWEH's June 2026 report projects that by 2030 the global water footprint of AI-driving data centres will equal the basic annual domestic water needs of all 1.3 billion people in Sub-Saharan Africa, with 2025 global data-centre water consumption already at 1.2 trillion gallons.¹³

The fifth risk is irreversibility. The UNU framework establishes that once natural water capital (aquifers, wetlands, lakes) is overdrawn, restoration may be infeasible. Data-centre approvals must be treated as decisions with intergenerational hydrological consequences, not as ordinary planning matters.¹²

The sixth risk is distributional injustice. The UNU report finds that the burdens of water overshoot fall disproportionately on smallholder farmers, Indigenous Peoples, low-income urban residents, women and youth — a pattern that maps directly onto Australian conditions in the Murray-Darling Basin, Western Sydney and remote First Nations communities.¹²

The seventh risk is grid cost-shifting. Without mandatory cost recovery, the network augmentation costs of connecting gigawatt-scale loads fall on residential and small-business bill payers — the *Stop the Bill Shock* concern raised by Sweltering Cities, the Antipoverty Centre and Parents for Climate. This is regressive: the lowest-income households cross-subsidise the most capital-intensive industry on the planet.²¹

The eighth risk is the crowding out of renewables. Even renewably-powered facilities consume new generation that would otherwise displace coal, slowing the energy transition.²⁰

Reasons supporting the recommendation

Voluntary "expectations" are demonstrably insufficient and doomed to fail. They create no enforceable obligations, no community remedies and no penalties for non-compliance.⁹ AEMO's forecast that data centres will consume approximately eleven per cent of NSW grid electricity by 2030 makes this a system-level rather than site-level issue requiring statutory rather than administrative response.^{19, 20} The UNU *Global Water Bankruptcy* "recovery plan" framework — stop the bleeding, protect essential services, restructure unsustainable claims, invest in rebuilding — is precisely the legislative architecture Australia needs and provides international scientific cover for strong domestic action.¹² Closed-loop cooling, waste-heat recovery and renewable-underwriting requirements are technically feasible and already applied in jurisdictions such as Ireland, the Netherlands and Singapore; Australia is a laggard, not a leader.

6.6 Labour-Market Protection and Just Transition Package

Recommendation

Enact a Labour-Market Protection and Just Transition (AI) Act incorporating an inversion of the current capital-over-labour tax bias by reducing accelerated depreciation for AI capital and introducing payroll-tax credits for worker retention and upskilling; a statutory right of workers (through their representatives) to be consulted before deployment of AI systems that materially affect job design, performance management, surveillance, scheduling or termination; a statutory duty of transparency on employer AI used in hiring, promotion and termination; a prohibition on fully-automated termination decisions; a Just Transition Levy on operators of large-scale AI infrastructure to fund retraining and income support; and mandated workforce-impact assessments for AI deployments affecting more than a defined number of workers.⁸

What are the relevant risks and harms?

The first risk is mass displacement without transition support. The CHT *AI Roadmap* documents that leading AI developers treat mass job loss as an inevitable side effect of innovation, with disruption visible in warehousing, knowledge work, journalism, translation, customer service, paralegal work, software engineering, design, illustration and creative industries.⁷

The second risk is the capital-substitution bias in the tax system. Australia's current tax settings effectively subsidise replacing workers with capital (including AI) by making depreciation and capital investment more attractive than wages and training. The fiscal lever is pointed in the wrong direction.⁸

The third risk is skills hollowing in the construction sector. The data-centre construction boom is colliding with an estimated 300,000 construction-worker shortfall by mid-2027, drawing labour away from housing, hospitals, schools and renewables.¹⁸

The fourth risk is the structurally low level of operational employment. Hyperscale facilities employ small permanent workforces relative to their land, water and energy footprint, disproving the (flawed) assumption that data centres are net regional job creators.²⁴

The fifth risk is the erosion of worker voice and dignity. Workers have little or no say in the AI systems being deployed over them — surveillance, productivity scoring, scheduling and termination — eroding the right to meaningful work that the CHT Roadmap recognises.⁷

The sixth risk is wage suppression and precarity. AI-mediated piecework, gig-economy expansion and the threat of automation suppress wages across labour markets even where actual substitution has not yet occurred.

The seventh risk is the disproportionate impact on women, migrants and lower-skilled workers, who are over-represented in the most automation-exposed roles.

Reasons supporting the recommendation

The CHT Roadmap's proposal to flip the capital-vs-labour bias is fiscally neutral if revenue-balanced and addresses the harm at source.⁸ Just-transition frameworks have precedent in Australia's energy transition policies, including the Net Zero Economy Authority, and can be adapted to AI displacement. A workforce with a meaningful voice in AI deployment is also a workforce that surfaces unsafe systems early — a safety, not just a fairness argument. Australia's social contract, Medicare, public education, the minimum wage, occupational health and safety, depends on a broadly distributed labour market, and AI infrastructure that hollows that out without compensating institutions threatens the contract itself.

6.7 Community Participation Rights and Local Council Veto

Recommendation

Legislate a mandatory pre-approval community consultation regime for all data-centre developments above 30 megawatts IT load, with reasonable time-frames, accessible plain-English information, independent technical support for community participants, and reasons that explain how community input affected the decision.

The legislation should establish a statutory veto right for local councils where formal moratoria or caps are in place; statutorily mandated, culturally appropriate engagement with First Nations communities where projects affect Country, water or sacred sites, including in accordance with the principles of free, prior and informed consent; a "community benefit" requirement under which operators must demonstrate tangible, measurable and enforceable local benefits before approval; and post-approval monitoring rights for affected communities, including access to environmental and resource-consumption data.

What are the relevant risks and harms?

The first risk is the disenfranchisement of host communities. Western Sydney communities in Marsden Park, Eastern Creek and Kemps Creek are bearing the urban-heat-island impact of facilities running at 35–45°C ambient day and night, with no meaningful consultative voice.^{[20](#) [21](#)}

The second risk is the override of local planning controls. State and Commonwealth fast-track regimes risk overriding the considered planning judgements of councils like Penrith and the City of Ryde — substituting central administrative discretion for local democratic accountability.^{[19](#)}

The third risk is First Nations exclusion. Data sovereignty, cultural heritage and Country impacts — water draws, land disturbance, energy infrastructure corridors — of mega-facilities are routinely decided without genuine free, prior and informed consent, in breach of Australia's commitments under the United Nations Declaration on the Rights of Indigenous Peoples.

The fourth risk is information asymmetry. Community participants face highly technical, often confidential proposals with no resources to engage equivalent expertise, making "consultation" a formality.

The fifth risk is the ignoring of cumulative impacts. Each facility is assessed in isolation; communities hosting multiple facilities suffer cumulative heat, noise, light, traffic and resource impacts that are invisible in any single assessment.

The sixth risk is the failure of benefit retention. Communities bear the costs and externalities while economic value is repatriated to foreign hyperscalers — a textbook resource-curse pattern.

Reasons supporting the recommendation

Procedural justice is itself a digital right: those affected by infrastructure decisions must have a meaningful voice in them. The Justice and Equity Centre's submission to the NSW inquiry recommends a public-interest regulatory architecture rather than a homogeneous fast-track regime, consistent with strengthened community rights.^{[22](#)} Local councils carry the burden of infrastructure failures, heat, traffic and amenity impacts; they should carry commensurate authority.^{[19](#)} Strong community participation rights have a stabilising effect on the social licence for legitimate AI infrastructure; they are not anti-development, they are pro-legitimate development.

6.8 Independent AI Regulator (Australian AI Safety Commission)

Recommendation

Establish a properly resourced, independent statutory AI regulator — an Australian AI Safety Commission — with investigative, audit, enforcement and standards-setting powers covering both AI systems and the data-centre infrastructure layer; own-motion investigation powers; authority to require pre-deployment safety testing and incident reporting; civil-penalty jurisdiction calibrated to hyperscaler revenue, i.e. aligned to the GDPR's approach in the EU; and statutory information-sharing arrangements with the OAIC, ACMA, ACCC, eSafety Commissioner, AHRC, Fair Work Ombudsman and state planning authorities.

What are the relevant risks and harms?

The first risk is fragmented enforcement. AI harm currently falls between the OAIC, ACMA, ACCC, eSafety Commissioner, AHRC, ASIC, APRA, Fair Work Ombudsman, state planning authorities and Treasury — all with predictable gaps and overlaps.

The second risk is a capability deficit. Current regulators lack the technical staffing, audit tooling and pre-deployment testing capacity to evaluate frontier AI systems.

The third risk is capture. Soft-law "expectations" administered by a Department also tasked with attracting investment create a structural conflict of interest.⁹

The fourth risk is the absence of a single point of accountability for the public. Affected individuals do not know where to take a complaint, and complaints that cross regulator boundaries are routinely lost.

Reasons supporting the recommendation

A dedicated regulator is consistent with international practice, including the UK AI Safety Institute, the US AI Safety Institute and the EU AI Office. The CHT Roadmap argues that "independent oversight, rigorous testing, and real transparency" are foundational.⁷ A combined AI-and-infrastructure mandate avoids the artificial separation between AI software and the physical compute on which it depends — a separation that the present soft-law architecture exploits.

6.9 Critical Infrastructure Coverage Extended to Large-Scale AI Compute

Recommendation

Extend the *Security of Critical Infrastructure Act 2018* to large-scale AI compute facilities, with public-interest as well as national-security criteria, including digital rights, environmental and community-resilience considerations. Require operators to disclose ownership, control, supply-chain dependencies and cross-border data flows, and empower the responsible Minister to give directions in response to public-interest as well as national-security risks.

What are the relevant risks and harms?

The first risk is concentration. Australia is rapidly becoming dependent on a small number of foreign hyperscaler facilities for cloud, AI inference, government workloads and national resilience.

The second risk is the cyber-physical attack surface. AI compute facilities are high-value targets for state and criminal actors. Their failure cascades affect health, finance, energy and government services.

The third risk is the loss of sovereignty. Without SOCI coverage, the Commonwealth lacks visibility into ownership, control and supply-chain dependencies of strategically significant compute.

The fourth risk is cross-border data exposure. EFA has previously raised concerns about the abandonment of data-localisation safeguards in the *Digital ID Act 2024*; the same risks recur at AI compute scale.²⁵

Reasons supporting the recommendation

SOCI is an existing, well-understood statutory vehicle that can be expanded without bespoke legislation. Including public-interest criteria alongside national-security criteria broadens accountability to environmental and community impacts, and brings AI compute into line with how Australia regulates other strategically significant infrastructure.

6.10 Mandatory Transparency of Government–AI Company Contracts

Recommendation

Mandate public disclosure of the substantive terms — subject to narrowly defined commercial-in-confidence carve-outs assessed by the Information Commissioner — of all Commonwealth contracts, MoUs, grants, tax concessions, energy allocations, water allocations and land deals involving global AI companies and large-scale data-centre operators. Establish a public AI Deals Register, with all material contracts published within thirty days of execution.

What are the relevant risks and harms?

The first risk is the democratic deficit. Major AI deals are being concluded with limited public visibility into the terms binding future Australian governments.¹⁸

The second risk is the existence of hidden subsidies. Energy and water concessions, tax holidays and grid prioritisation are forms of public expenditure that should be publicly accounted for.

The third risk is the inability to assess social licence. Without transparency, the Commonwealth's own *Expectations* framework cannot be tested against actual contract terms.⁹

The fourth risk is audit and accountability failure. The Australian National Audit Office, the Auditor-General and parliamentary committees cannot scrutinise what they cannot see.

Reasons supporting the recommendation

Transparency is a precondition of public consent and a precondition of the moratorium recommended in 6.1: the public cannot consent to deals it cannot see. Comparable transparency obligations apply to defence procurement and major infrastructure; AI is no less consequential. Public disclosure is itself a disincentive to overreaching contract terms.

7. Conclusion

The terms of reference adopted by the Senate on 13 May 2026 invite this Committee to ask the right question at the right moment.¹⁵ The honest answer, by any other name an inconvenient truth, drawn from EFA's three decades of digital-rights advocacy, the lived evidence of councils and communities in Sydney and beyond, the Center for Humane Technology's *AI Roadmap*, and the United Nations University's twin 2026 reports on global water bankruptcy and the environmental cost of AI's energy use — is that Australia is being asked to underwrite a generational build-out of AI infrastructure without the ex ante legal architecture that any comparably consequential technology would require.^{7, 12, 13}

Australians overwhelmingly want this technology to serve them, not the reverse. The 2026 Australian Community Attitudes to Privacy Survey found that eighty-seven per cent of Australians are more concerned about their privacy than five years ago, with only one in ten trusting current practices as fair.¹¹

There is an absence of trust in emerging tech. Public trust in the systems underpinning modern digital life has hit an absolute floor. Just 4% of Australians believe AI companies are worthy of their trust, while trust in social media companies languishes at a mere 3%. Australians are deeply concerned about the rollout of AI, digital identification, the predilections and influence of Big Tech, and biometrics without the protection of strong regulatory frameworks being in place. That is not a public mandate for fast-tracking. It is a public mandate for strong ex ante legislation.

EFA urges the Committee to recommend that the Senate adopt the framework above before the parallel NSW inquiry reports on 30 September 2026, and before further deals with global AI companies foreclose Australia's regulatory options.¹⁴

Yours sincerely,

John Pane Chair, Electronic Frontiers Australia

Electronic Frontiers Australia Inc. | www.efa.org.au | chair@efa.org.au

8. References (Footnotes)

1. Senate Standing Committees on Environment and Communications, *Inquiry into Artificial Intelligence and Data Centres*, Parliament of Australia, https://www.aph.gov.au/Parliamentary_Business/Committees/Senate/Environment_and_Communications/Aldatacentres48P (referred 13 May 2026; reporting 16 November 2026). ↩ ↩² ↩³ ↩⁴
2. Electronic Frontiers Australia, *Submissions and Publications archive*, <https://www.efa.org.au/category/publications/submissions/>. ↩ ↩² ↩³
3. Electronic Frontiers Australia, *Submission to the Safe and Responsible AI in Australia Discussion Paper*, 13 August 2023, <https://www.efa.org.au/2023/08/13/safe-ai-submission/>. ↩ ↩² ↩³ ↩⁴ ↩⁵
4. Electronic Frontiers Australia, *Submission to the Privacy Act Review Report*, April 2023, <https://www.efa.org.au/2023/04/privacy-act-review/>. ↩ ↩² ↩³
5. Electronic Frontiers Australia, *Submission to Australia's AI Ethics Framework Consultation*, 2019, <https://www.efa.org.au/2019/06/ai-ethics-framework/>. ↩ ↩² ↩³ ↩⁴
6. NSW Legislative Council Public Accountability and Works Committee, *Inquiry into Data Centres in NSW* (pk=3169), submissions index, <https://www.parliament.nsw.gov.au/committees/inquiries/Pages/inquiry-details.aspx?pk=3169#tab-submissions> (closed 27 March 2026). ↩ ↩² ↩³
7. Center for Humane Technology, *The AI Roadmap*, April 2026, <https://www.humanetech.com/ai-roadmap>. ↩ ↩² ↩³ ↩⁴ ↩⁵ ↩⁶ ↩⁷ ↩⁸ ↩⁹ ↩¹⁰ ↩¹¹ ↩¹² ↩¹³
8. Center for Humane Technology, *The AI Roadmap — Policy Recommendations*, April 2026, <https://www.humanetech.com/ai-roadmap#recommendations>. ↩ ↩² ↩³ ↩⁴ ↩⁵ ↩⁶ ↩⁷ ↩⁸ ↩⁹ ↩¹⁰ ↩¹¹ ↩¹²
9. Department of Industry, Science and Resources (Cth), *Expectations of data centres and AI infrastructure developers*, 23 March 2026, <https://www.industry.gov.au/publications/expectations-data-centres-ai-infrastructure-developers>. ↩ ↩² ↩³ ↩⁴ ↩⁵ ↩⁶ ↩⁷ ↩⁸
10. Digital Transformation Agency (Cth), *Policy for the responsible use of AI in government 2.0*, effective 15 December 2025, <https://www.digital.gov.au/policy/ai/policy>. ↩ ↩² ↩³
11. Office of the Australian Information Commissioner, *Australian Community Attitudes to Privacy Survey 2026*, 28 May 2026, <https://www.oaic.gov.au/engage-with-us/research-and-training-resources/research/acaps-2026>; see also Electronic Frontiers Australia commentary at <https://www.efa.org.au>. ↩ ↩² ↩³ ↩⁴ ↩⁵ ↩⁶ ↩⁷ ↩⁸
12. K. Madani, *Global Water Bankruptcy: Living Beyond Our Hydrological Means in the Post-Crisis Era*, United Nations University Institute for Water, Environment and Health (UNU-INWEH), 20 January 2026, <https://inweh.unu.edu/global-water-bankruptcy/>. ↩ ↩² ↩³ ↩⁴ ↩⁵ ↩⁶ ↩⁷ ↩⁸ ↩⁹ ↩¹⁰ ↩¹¹ ↩¹²

13. UNU-INWEH, *Environmental Cost of AI's Energy Use: Carbon, Water and Land Footprints*, 3 June 2026, <https://inweh.unu.edu/environmental-cost-of-ai/>. ↩ ↩² ↩³ ↩⁴ ↩⁵ ↩⁶ ↩⁷
14. Clayton Utz, *Analysis: NSW Legislative Council inquiry into data centres — submissions and reporting timeline*, 2026, <https://www.claytonutz.com/insights/2026/nsw-data-centres-inquiry>. ↩ ↩²
15. Senate of Australia, *Journals of the Senate No. 51*, 13 May 2026, https://www.aph.gov.au/Parliamentary_Business/Chamber_documents/Senate_chamber_documents/Journals_of_the_Senate. ↩ ↩²
16. Clayton Utz, *NSW Legislative Council inquiry into data centres — submissions overview*, 2026, <https://www.claytonutz.com/insights/2026/nsw-data-centres-submissions>. ↩ ↩² ↩³
17. Department of Industry, Science and Resources (Cth), *Five Expectations for Data Centres and AI Infrastructure — Explanatory Note*, 23 March 2026, <https://www.industry.gov.au/publications/expectations-data-centres-explanatory>. ↩ ↩²
18. Anthropic and Australian Government, *Memorandum of Understanding on AI Cooperation*, 31 March 2026, summarised at <https://www.industry.gov.au/news/anthropic-mou-australia>; see also Certified Strategic, *Data Centre Jobs in Australia: The 2026 Hiring Landscape*, April 2026, <https://www.certifiedstrategic.com.au/insights/data-centre-jobs-2026>. ↩ ↩² ↩³ ↩⁴ ↩⁵
19. B. Eastaugh, "Penrith and other NSW councils call for moratorium on data-centre approvals", *Crikey*, 2026, <https://www.crikey.com.au/2026/data-centres-penrith-moratorium>. ↩ ↩² ↩³ ↩⁴ ↩⁵ ↩⁶ ↩⁷ ↩⁸ ↩⁹
20. Friends of the Earth Australia, *Submission to the NSW Legislative Council Inquiry into Data Centres*, 22 May 2026, <https://www.parliament.nsw.gov.au/ladocs/submissions/data-centres/sub-foe.pdf>. ↩ ↩² ↩³ ↩⁴ ↩⁵ ↩⁶ ↩⁷ ↩⁸ ↩⁹ ↩¹⁰ ↩¹¹ ↩¹²
21. Sweltering Cities, Antipoverty Centre and Parents for Climate, *"Stop the Bill Shock" — Joint Submission to the NSW Legislative Council Inquiry into Data Centres*, 2026, <https://swelteringcities.org/submissions/stop-the-bill-shock>. ↩ ↩² ↩³ ↩⁴
22. Justice and Equity Centre, *Submission to the NSW Legislative Council Inquiry into Data Centres*, 2026, <https://jec.org.au/submissions/data-centres-nsw-2026>. ↩ ↩²
23. ARC Centre of Excellence for Automated Decision-Making and Society (ADM+S), *Submission No. 85 to the NSW Legislative Council Inquiry into Data Centres*, 2026, <https://www.admscentre.org.au/submissions/nsw-data-centres-2026>. ↩
24. M. Hossain et al., "Australia is set to get more AI data centres. Local communities need to be more involved", *The Conversation*, 2025, <https://theconversation.com/australia-is-set-to-get-more-ai-data-centres>. ↩ ↩²
25. Electronic Frontiers Australia, *Commentary on the Digital ID Act 2024 and data localisation*, 2024–2026, <https://www.efa.org.au/category/digital-id/>. ↩ ↩² ↩³
26. Electronic Frontiers Australia, *Privacy Act Tranche II — Election scorecard and ongoing commentary*, 2025–2026, <https://efa.org.au/2025-federal-election-policy-scorecard>

