

9 October 2026

Department of the Prime Minister and Cabinet

Canberra ACT 2600

Submitted via the online consultation form

Submission on the consultation paper 'Getting it right: Building AI infrastructure that works for Australia' from UNSW Sydney, with the UNSW AI Institute and the UNSW Energy Institute as major contributors.

*Key Contributors: **Dr Sue Keay** (Director, UNSW AI Institute); **Scientia Professor Toby Walsh** (Chief Scientist, UNSW AI Institute); **Dr Alex Antic** (Faculty Head of AI Strategy, UNSW Canberra, and Deputy Director, UNSW AI Institute); **Luc Betbeder-Matibet** (Executive Director, Research Technology Services, and Deputy Director – Infrastructure, UNSW AI Institute); **Associate Professor Renee Whan** (Pro-Vice Chancellor Research Infrastructure, Division of Research & Enterprise); **Dani Alexander** (Chief Executive Officer, UNSW Energy Institute); **Mark Twidell** (Industry Professor, UNSW Energy Institute); **Scientia Professor Sven Rogge** (Dean, UNSW Science); **Professor Lyria Bennett Moses** (Academic Director, Research, Division of Research & Enterprise); **Associate Professor Katharine Michie** (Chief Scientist, Structural Biology Facility, Mark Wainwright Analytical Centre); **Professor Kathy Bowrey** (UNSW Law & Justice); **Professor Ned Ekins-Daukes** (Head of School, Photovoltaic and Renewable Energy Engineering); **Associate Professor Georgios Konstantinou** (School of Electrical Engineering and Telecommunications; Lead, Real-Time Simulations at UNSW); **Professor Iain MacGill** (School of Electrical Engineering and Telecommunications; Joint Director (Engineering), Collaboration on Energy and Environmental Markets); and **Dr Alper Yasar** (Manager, Research Infrastructure Initiatives, Division of Research & Enterprise).*

Executive summary

The consultation paper correctly identifies that Australia cannot simply be a destination for infrastructure investment, and that the benefits of AI require onshore training of models as opposed to hosting alone; our main contention is that the training should include Australian models of various scales, not just US frontier models. This distinction needs to be considered in the development of nationally consistent standards for large data centres and conditions for AI training. The standards need to be designed to ensure that: (1) obligations reach the party controlling the compute and not only the landlord who obtains planning approval; (2) a contribution to Australian capability in the form of reserved, Australian-governed compute rather than token credits and in-kind offers from vendors is secured; (3) a broad Australian capability is developed (beyond our AI Safety Institute) to test, evaluate and assure the systems that Australians and their public services will rely on, so that safety is a benefit the standards deliver for Australians rather than only an obligation imposed on developers; and (4) a measurable long-term uplift in Australian AI capability is achieved from the investment and the training.

The standards should also turn large data-centre demand into an asset for Australia's energy transition by requiring genuinely additional renewable generation and firming, supporting strategic regional siting and economic diversification, and ensuring that costs are not shifted to existing consumers or host communities.

For ease of reference: sections 1 and 2 of this submission and Recommendations 1 and 2 respond to Part 1, managing growth and the application of the standards; sections 3 to 6 and Recommendations 3 to 6 respond to Part 2, the requirements for large data centres; and sections 7 to 11 and Recommendations 7 to 12 respond to Part 3, the conditions for AI training, with section 12 and Recommendation 13 addressing the national setting for the standards. Finally, our main points remain that hosting is not capability, and private contributions must add to the public compute base that every comparable economy funds at scale, while Australia does not.

List of recommendations included in this submission

Recommendation 1: Use an energy-connection threshold as the primary trigger, with the tiered structure and 150MW portfolio test proposed, and add a compute-based trigger for AI training obligations and a water trigger for stressed catchments. Apply the standards at energy connection or planning approval, whichever is first, capture pipeline projects, and require published compliance plans for facilities above 100MW. Keep the compute threshold in an updateable instrument and review it regularly against transparent capability- and risk-based criteria.

Recommendation 2: Draft the standards so that energy, water, workforce and capability obligations attach to whichever party controls the material activity, with mandatory contractual flow-down from owner to operator to compute tenant and direct undertakings from anchor tenants for facilities above 100MW.

Recommendation 3: Build on the offset, firming and registration arrangements agreed by Energy Ministers in July 2026, with a hard time limit on any period of reliance on existing renewables. Assess the proposed data-centre pipeline through integrated temporal and geographical scenarios of loads, generation, storage, firming and network capacity against Australia's legislated emissions targets. Require new renewable generation and storage to match data-centre load on a short glide path, with peak-period time-matching for facilities above 100MW, limits and reporting on fossil back-up generation, demand flexibility as a condition of connection, grid-capable backup batteries and recognition of on-site and near-site solar and storage, and annual public reporting of energy use, PUE and emissions alongside water metrics and hourly carbon-free energy and grid-interaction metrics, and require real-time grid-disturbance modelling for large inverter-based loads, including cascading-disconnection and coordinated-battery/eSTATCOM scenarios.

Recommendation 4: Adopt the Water Ministerial Council minimum requirements, mandate a nationally consistent source hierarchy, and require a minimum NABERS water rating for data centres once the NABERS water framework is complete, with energy and water ratings always assessed and reported together. Require operational WUE reporting by source and catchment in the interim.

Recommendation 5: Prioritise suitable sites near renewable energy zones and regional industrial precincts for latency-insensitive workloads, while requiring proponents seeking constrained locations to substantiate the workload's latency or connectivity needs and ensuring attributable infrastructure costs are not shifted to existing consumers or host communities. Establish a public national register of large data-centre proposals and approvals showing connection capacity, water source and cooling design, and require published community and infrastructure commitments before approval so they can be enforced.

Recommendation 6: Require skills commitments from data centres and AI training facilities to be delivered through apprenticeships, TAFE and university pathways and PhD placements, including multidisciplinary university and VET programs for the data-centre energy workforce, measured by completions and employment outcomes, with vendor-specific training modules excluded from counting towards the obligation.

Recommendation 7: Require notification of defined incidents to a named Australian channel within a statutory timeframe, with penalties for late, incomplete or misdirected notification and a duty to cooperate with external investigation, alongside independent evaluation, red-teaming and supply-chain assessment as conditions of authorisation to train at scale; give the Australian AI Safety Institute pre-deployment access and formal links to Australian academic expertise; and base data-erasure requirements on demonstrated machine-unlearning capability.

Recommendation 8: Make pre-release access for Australian security agencies and evaluators an enforceable condition of permission to train frontier models in Australia, decline to privilege any single vendor's closed models for critical settings, preserve access for Australian agencies, researchers and defenders to open and closed models from several jurisdictions on Australian-governed infrastructure, and fund the independent evaluation and assurance capability these conditions depend on.



Recommendation 9: Require, as a binding condition of authorisation, that AI infrastructure operators and their compute tenants reserve a defined and independently verifiable share of installed compute for Australian public-interest users, discharged in kind or through a tradeable obligation and administered through a single national counterparty using existing research infrastructure providers under the seven principles above. Recognise Australian-controlled neoclouds as capability partners in delivering reserved compute and assess all proposed contributions against measurable domestic value retention, including local skills, research and development, supplier expenditure, intellectual property, reinvestment and Australian taxes paid. Adopt a mandatory local AI research presence requirement for major AI companies, reform the R&D Tax Incentive, and report on capability outcomes rather than construction expenditure.

Recommendation 10: Reform copyright law, with specialist input, to provide clear exceptions for education, publicly funded research and other public-interest activities, prevent those exceptions being overridden by contract, and protect creators' ability to retain their rights; require auditable records and disclosure of Australian training material; provide transparent licensing and remuneration where uses fall outside applicable exceptions; condition the use of publicly funded data on governed agreements; and establish separate protections for Indigenous Cultural and Intellectual Property.

Recommendation 11: Recognise digital infrastructure as critical national infrastructure for budget and planning purposes, fund a public-interest compute base in the forward estimates with stable long-term support for national research platforms such as the BioCommons Structural Biology Platform, and use public procurement to create early domestic markets for Australian models and AI products, with vendor-neutral terms as a condition of public spending.

Recommendation 12: Establish a common reporting gateway, identifiers and data definitions so proponents can report once for use by relevant regulators; retain existing statutory responsibilities and coordinate the compute obligation through the separate national allocation framework; align safety reporting with the proposed digital duty of care; and review the standards at least annually against whether the safety conditions are working.

Recommendation 13: Express the AI standards within a funded national AI strategy that sets targets for Australian compute, models, researchers, companies and independent evaluation and assurance capability, with commitments that are funded, measurable and enforceable.



1. Thresholds and application of the AI standards

An energy-connection threshold is a sensible primary trigger because it aligns with the AEMC access tiers and with the grid impacts the standards most need to manage. It should not be the only one. A 40MW facility filled with current-generation accelerators has a different national significance from a 40MW facility of enterprise storage, and a training cluster can be assembled across several sites that each sit below a threshold. We support the tiered approach and the portfolio aggregation test at 150MW, and recommend a second trigger for AI training obligations based on training compute, on the model used in California's SB 53 and New York's RAISE Act, which both set the line at 10^{26} operations counted cumulatively across the initial training run and any subsequent fine-tuning or material modification¹, so that the party undertaking large-scale training is captured wherever the equipment is located, with installed accelerator capacity available as a secondary trigger for facility-level reporting. The compute threshold should not be treated as a permanent proxy for capability or risk: it should sit in an updateable instrument, be reviewed regularly, and be supported by transparent capability- or risk-based designation criteria. Water use should sit as a third trigger in stressed catchments, since local impact depends on the catchment more than on facility size. The compute trigger matters for a second reason. The safety and evaluation conditions in Part 3 protect Australians from the behaviour of models rather than from the behaviour of buildings, so they need to attach to the training activity wherever it is carried out.

The standards should apply from the point at which a project seeks its energy connection agreement or state planning approval, whichever comes first, and should capture projects in the development pipeline that have not yet connected. A plan demonstrating how the facility will meet its obligations should be lodged before connection and published for facilities above 100MW. Obligations linked to actual load can then scale as capacity is drawn down. A short transition period for existing facilities is reasonable, but any facility seeking government facilitation, prioritised approvals, public land or public procurement should meet the standards in full from the day the legislation commences.

***Recommendation 1:** Use an energy-connection threshold as the primary trigger, with the tiered structure and 150MW portfolio test proposed, and add a compute-based trigger for AI training obligations and a water trigger for stressed catchments. Apply the standards at energy connection or planning approval, whichever is first, capture pipeline projects, and require published compliance plans for facilities above 100MW. Keep the compute threshold in an updateable instrument and review it regularly against transparent capability- and risk-based criteria.*

2. Conditions must reach the party that controls the compute

The paper notes that data centres and AI training facilities can involve owners, operators, tenants and AI developers, and asks how responsibilities should be distributed. This is the most important structural question in the consultation. The operator who obtains planning approval and signs the energy and water connections is usually not the party that controls the compute. Hyperscaler and neocloud tenants own or control the accelerators, decide whether they run training or inference, set the data governance, choose the cooling and back-up systems and decide who gets access. A water-reporting obligation on a landlord cannot compel a tenant to change a cooling-intensive workload, and a public-interest access expectation on an operator cannot reserve GPU-hours the operator does not control.

Australia already applies the right principle in its resources sector. A mining licence is issued subject to conditions covering rent, rehabilitation and community consultation, and those conditions bind whoever conducts the activity. Ireland requires data centre tenants to enter demand-flexibility arrangements as a condition of grid connection, and the EU energy efficiency framework extends reporting obligations to colocation customers who

¹Transparency in Frontier Artificial Intelligence Act (California SB 53), signed 29 September 2025, in force 1 January 2026; New York RAISE Act, as amended March 2026, in force 1 January 2027. Both define a frontier model by training compute above 10^{26} operations, counted cumulatively across the initial training run and any subsequent fine-tuning or material modification, with additional obligations for developers above US\$500 million in annual revenue. Compare Regulation (EU) 2024/1689 (the Artificial Intelligence Act), Article 51(2), under which a general-purpose AI model is presumed to have high-impact capabilities, and so to be classified as a model with systemic risk under Article 51(1)(a), when the cumulative amount of computation used for its training measured in floating point operations is greater than 10^{25} . That presumption is rebuttable, the Commission may separately designate a model under Article 51(1)(b) against the Annex XIII criteria, and under Article 51(3) it may adjust the threshold by delegated act.

own their own equipment². The standards should follow the same logic: obligations attach to the licence to operate and cascade through leases, service agreements and connection conditions to the compute owner, supported for major projects by direct undertakings from anchor tenants to the responsible agency. A facility should not receive facilitation until its major compute tenants have accepted the relevant obligations and future tenants are bound by them.

***Recommendation 2:** Draft the standards so that energy, water, workforce and capability obligations attach to whichever party controls the material activity, with mandatory contractual flow-down from owner to operator to compute tenant and direct undertakings from anchor tenants for facilities above 100MW.*

3. A positive contribution to the energy transition

We support a renewable electricity obligation with genuine additionality. Australia should approach data-centre growth as both an energy-system challenge and a national economic opportunity. The value of large, predictable data-centre demand is that it can act as an anchor load, providing long-term revenue certainty to underwrite the new generation, storage, firming and network capacity the grid needs as coal retires, and that value is lost when the obligation is met by certificates from existing plant. Governments should assess the proposed data-centre pipeline through integrated temporal and geographical scenarios of large loads, generation, firming and network capacity, including compatibility with Australia's legislated emissions targets. The glide path should be short and should move to new capacity quickly, with time-matching in peak periods for facilities above 100MW, since an annual matching rule allows a facility to run on fossil generation at the hours that matter most. UNSW's 24/7 TRUZERO research found that full 24/7 matching can be complex and costly, and that demand flexibility improved emissions outcomes but had less effect on those higher costs. This supports starting with periods of greatest system need rather than imposing an immediate 24/7 requirement. A minimum level of demand flexibility should be a condition of connection for large facilities, as it is in Ireland, with market mechanisms used to procure flexibility beyond that minimum. Obligations should be set according to demonstrated technical capability and workload characteristics; training-dominated facilities may justify higher obligations, but this should not be assumed for every training workload. Non-renewable back-up generation should be subject to defined limits and public reporting, with permitted use tied to specific emergency, scarcity or system-security conditions.

These standards should build on, not duplicate, the arrangements Energy Ministers agreed on 28 July 2026 following the AEMC's advice: mandating that data centres offset their consumption with REGO certificates linked to new renewable generation, a contract obligation for firming capacity, market registration requirements for visibility, and connection agreements that support demand flexibility and co-location with generation.³ UNSW supports that package, and in particular the firming contract obligation, which will also reward co-located storage. We suggest three refinements. First, the option for a data centre to draw temporarily on existing renewables and make up the shortfall by surrendering additional certificates in later years should carry a hard time limit, so that it cannot become indefinite reliance on existing plant. Second, a minimum level of demand flexibility should be a mandatory condition of connection, rather than relying only on voluntary flexible-demand connection agreements. Third, offsetting demand should be treated as the floor rather than as a benefit: without it, data-centre growth raises prices and emissions for every other electricity user.

Connection requirements should address how large inverter-based loads respond to grid disturbances and should seek, where possible, to strengthen power-system security. Real-time simulation should test rapid or cascading disconnection risks and whether coordinated battery systems can support both facility resilience and power-system stability. This work would complement the AEMC rule change, already under way, that updates technical access standards for large inverter-based loads such as data centres. With appropriate support, relevant UNSW workstreams could include use of the openISP platform for system-wide scenario analysis; extension of the separate open-source PPA contract-assessment tool to test data-centre time-matching and demand-flexibility options; and further work through RTS@UNSW on large-scale integration of inverter-based loads (IBLs), data-centre supply and grid stability impacts. The full submission is attached.

²Directive (EU) 2023/1791, Article 12, and Commission Delegated Regulation (EU) 2024/1364, which extend reporting obligations to colocation customers who own and manage their own equipment.

³ Australian Energy Market Commission, 'Advice to ECMC: Data centre energy obligation', July 2026; Energy and Climate Change Ministerial Council meeting, 28 July 2026.



Data centres can also be solar and storage assets in their own right. Every large facility already installs battery backup (UPS) capacity; with appropriate controls, that capacity can provide frequency and other grid services rather than sitting idle, and the standards should require new backup systems to be grid-capable. On-site and near-site solar and storage should count towards the renewable electricity obligation without unnecessary certification overhead, particularly for regional facilities co-located with renewable energy zones. And as AI facilities move towards high-voltage DC power distribution, directly coupling DC-native solar and batteries to the facility can avoid conversion losses. This is a near-term opportunity for Australian research, testing and demonstration.

On efficiency, the design PUE thresholds and a minimum NABERS energy rating are a reasonable baseline. Power Usage Effectiveness was designed for cloud-era workloads and can be improved by cooling strategies that increase water use, so PUE should never be reported or regulated without its water counterpart, as Section 4 explains. Large facilities should also report their hourly carbon-free energy share, their grid-interaction and disturbance ride-through performance, and the flexibility they actually deliver, since a highly efficient facility can still impose a large burden on the power system. Facility-level energy, efficiency and emissions data should be reported quarterly to the regulator and published annually, consistent with the NGER reforms.

***Recommendation 3:** Build on the offset, firming and registration arrangements agreed by Energy Ministers in July 2026, with a hard time limit on any period of reliance on existing renewables. Assess the proposed data-centre pipeline through integrated temporal and geographical scenarios of loads, generation, storage, firming and network capacity against Australia's legislated emissions targets. Require new renewable generation and storage to match data-centre load on a short glide path, with peak-period time-matching for facilities above 100MW, limits and reporting on fossil back-up generation, demand flexibility as a condition of connection, grid-capable backup batteries and recognition of on-site and near-site solar and storage, and annual public reporting of energy use, PUE and emissions alongside water metrics and hourly carbon-free energy and grid-interaction metrics, and require real-time grid-disturbance modelling for large inverter-based loads, including cascading-disconnection and coordinated-battery/eSTATCOM scenarios.*

4. Water

Australia's climate makes water a material constraint rather than a secondary consideration. Sydney Water's forecast that data centres could draw a fifth of the city's drinking water by 2035 shows why. We support the mandatory requirements agreed by the Water Ministerial Council in August 2026 and support a nationally consistent hierarchy of water sources, with recycled and non-potable water first and potable water as a last resort, together with a requirement that proponents demonstrate that proposed use will not affect local water security, environmental flows or existing users.

UNSW AI Institute, Submission to the Senate inquiry into Artificial Intelligence and Data Centres identified a gap in the NABERS framework, which rated data centres on energy but not on water, allowing an operator to pursue a five-star energy rating through cooling strategies that increase water consumption⁴. That gap is now being closed. NABERS has secured funding to develop a national water efficiency framework for data centres, with completion expected in October 2026⁵. The AI standards should require a minimum NABERS water rating for data centres as soon as the tool is operational, and should treat the energy and water ratings as a pair, with the PUE and WUE thresholds set jointly and differentiated by water source, as the paper proposes. Until the NABERS tool exists, facilities should report operational WUE under ISO/IEC 30134-9 with the water source and catchment stress disclosed. Cooling technologies that eliminate evaporative water use, including closed-loop liquid cooling, should be the expected design for large facilities in stressed catchments.

***Recommendation 4:** Adopt the Water Ministerial Council minimum requirements, mandate a nationally consistent source hierarchy, and require a minimum NABERS water rating for data centres once the NABERS water*

⁴UNSW AI Institute, Submission to the Senate inquiry into Artificial Intelligence and Data Centres, 25 June 2026.

⁵Water Services Association of Australia, 'NABERS water efficiency framework for data centres', 21 May 2026. The framework is expected to be completed by October 2026.

framework is complete, with energy and water ratings always assessed and reported together. Require operational WUE reporting by source and catchment in the interim.

5. No net costs, location and community engagement

Strategically located and coordinated data-centre investment can support productivity, sovereign capability and regional economic development, provided costs are not transferred to households, businesses or host communities. Planning should distinguish latency-insensitive training and batch workloads from latency-sensitive inference. Pre-identified sites near renewable energy zones and regional industrial precincts should be prioritised where adequate energy, water and telecommunications capacity and opportunities for skilled regional employment can be demonstrated. Proponents seeking approval in constrained locations should substantiate the latency or connectivity needs of the proposed workload and why suitable alternatives are not viable.

The no-net-costs principle should extend beyond network augmentation to the skilled trades that data-centre construction draws away from housing, hospitals and energy projects, and to the land that campus-style developments take from housing and agriculture. Proponents should demonstrate how they will pay for the electricity and water infrastructure their demand requires, and the cost-recovery arrangements for both should be reviewed together. Location transparency is a Commonwealth responsibility that states cannot provide alone: a public register of proposed and approved large data centres, with their connection capacity, water source and cooling design, would allow communities, utilities and researchers to see cumulative impacts across a catchment or grid region. Siting decisions should remain with states, but the standards should require that proponents publish their community engagement plans and their commitments before approval so that those commitments can be enforced afterwards.

Where a proposal may affect Country, cultural heritage, water, Indigenous rights or community interests, engagement with Aboriginal and Torres Strait Islander peoples should begin before site selection and include culturally appropriate participation, protection of Indigenous Cultural and Intellectual Property, and Indigenous data governance where relevant.

***Recommendation 5:** Prioritise suitable sites near renewable energy zones and regional industrial precincts for latency-insensitive workloads, while requiring proponents seeking constrained locations to substantiate the workload's latency or connectivity needs and ensuring attributable infrastructure costs are not shifted to existing consumers or host communities. Establish a public national register of large data-centre proposals and approvals showing connection capacity, water source and cooling design, and require published community and infrastructure commitments before approval so they can be enforced.*

6. Skills and training

Large data-centre projects compete for electricians, HVAC technicians and cyber specialists who are already scarce, so a skills obligation is warranted. The design matters more than the headline. Commitments should mean apprenticeships, TAFE pathways, industry-linked PhD placements and research-industry mobility, measured by completions and employment outcomes. Vendor-specific training modules are a customer-acquisition strategy that builds familiarity with a particular product rather than transferable capability, and should not count towards the obligation. For AI training facilities, obligations should extend to research software engineers, GPU systems specialists, model evaluators and data stewards, the roles that take years to develop and on which sovereign capability depends. Australia's AI PhD pipeline is strong, with female representation among AI researchers the second highest of any country tracked, and workforce commitments should reinforce that pipeline rather than run beside it.

Large data centres are now also energy facilities, requiring engineers and technicians in high-voltage electrical systems, on-site solar and battery storage, power electronics and cooling. UNSW's School of Photovoltaic and Renewable Energy Engineering trains this workforce, and from 2027 UNSW will run a cross-school Sustainable Data Centres Vertically Integrated Project spanning computing, electrical, renewable energy, mechanical and water engineering. UNSW is also establishing the Centre for Integrated Renewable Energy Training & Innovation in Western Sydney focused on tertiary harmonisation in battery storage and solar operations, maintenance and pre-compliance testing. Skills obligations should recognise contributions to programs of this kind, which build transferable capability rather than familiarity with one vendor's product.



UNSW
SYDNEY

***Recommendation 6:** Require skills commitments from data centres and AI training facilities to be delivered through apprenticeships, TAFE and university pathways and PhD placements, including multidisciplinary university and VET programs for the data-centre energy workforce, measured by completions and employment outcomes, with vendor-specific training modules excluded from counting towards the obligation.*

7. Security, safety and transparency for AI training

These conditions are the part of the framework that delivers most directly to Australians. Data centre standards govern how a facility draws power and water; the training conditions govern how the systems that will diagnose patients, assess claims, filter information and support critical services behave once they are deployed here. Australia is being asked to host that training, which gives it leverage over the terms at a moment when it still has leverage to use. We ask that the standards treat safety as one of the benefits Australians receive in exchange for hosting, and that the benefit be made concrete through a domestic capacity to test, evaluate and assure these systems independently of the companies that build them. The Government has accepted a version of this principle in a cognate regime: the exposure draft of the Online Safety Amendment (Digital Duty of Care) Bill 2026 proposes a proactive, systems-based duty on online services and would give approved researchers, including those at Australian universities, legal access to provider data for safety research⁶. The same reasoning applies with greater force to models trained at frontier scale on Australian soil.

Frontier labs authorised to train in Australia should be required to report defined incidents, disclose material changes to models used in critical settings, maintain logs, support independent testing and cooperate with external incident investigations rather than supplying their own analysis. The Australian AI Safety Institute needs access to models, compute, test data and controlled environments before systems are deployed, and its resources should grow with the volume of frontier training and deployment in Australia. It should have formal, ongoing access to Australian academic AI-safety expertise so that government can draw on independent assessment rather than relying on assessments undertaken by the developers themselves.

Information obligations should be layered. Reportable incidents and material safety evaluations should be provided proactively to the Institute; training data provenance, model cards and supply-chain dependencies should be available on request; and a summary should be published in a specified form. Any right-to-erasure or data-deletion requirement should be based on what machine-unlearning methods can verifiably achieve, an area where Australian researchers have internationally recognised expertise. Conditions should also cover supply-chain visibility: foreign ownership, control of model updates, remote administration, data export, hardware dependencies and the consequences of a provider withdrawing service.

Two recent incidents bear directly on how these obligations should be framed. On 18 June 2026 an agent operating during OpenAI's own internal evaluation of a frontier model reached the Medicare Statistics Reporting Service portal administered by Services Australia, worked around the controls that should have stopped it, and wrote files into an internal server. The company identified the activity in August, notified Services Australia on 10 September, and did so through a generic public inbox⁷. In July a model escaped its evaluation environment a second time and compromised Hugging Face's production systems, where responders found commercial models' guardrails obstructing their forensic work and finished the job with a self-hosted open-weight model⁸. We draw two conclusions. Voluntary disclosure has now been tested against an Australian government system and failed, so the reporting duty below needs a named recipient, a statutory deadline and a penalty for late or misdirected notification. And the containment of agents during evaluation is itself a subject on which Australia needs independent assessment capability, because in both cases the developer's own evaluation process was the point of failure.

⁶Exposure draft, Online Safety Amendment (Digital Duty of Care) Bill 2026, released 8 September 2026, consultation closed 22 September 2026. The draft proposes a proactive, systems-based duty on a 'reasonably practicable' standard drawn from work health and safety law, and includes a data access scheme allowing approved researchers, including those at Australian universities, to obtain provider data for online safety research. The Bill has not yet been introduced.

⁷ABC News, 'OpenAI hacked Medicare portal, Prime Minister Anthony Albanese says', 24 September 2026. Services Australia was notified on 10 September 2026; the investigation, assisted by the Australian Signals Directorate, is continuing.

⁸Hugging Face, 'Agent intrusion: technical timeline', post-incident analysis, July 2026; CSIS, 'China's Open-Weight Challenge to U.S. AI Leadership', September 2026.



That evidence bears equally on the security case being put for privileged access to proprietary frontier models. In August 2026 the major laboratories and more than 100 co-signatories asked governments to give vetted defenders early access to advanced models, on the ground that AI-enabled attacks are about to become far more widespread⁹. Our assessment is that the threat record does not yet carry that claim. The CISA advisory of August 2026 on AI-generated attacks against water utility control systems identified no new vulnerabilities. RUSI finds criminal groups innovating only when their existing methods stop paying. The chief executive of the SANS Institute observed that the letter's own recommendations describe work done by people. ASD recorded 84,700 cybercrime reports in 2024-25, three per cent fewer than the year before, with phishing present in 60 per cent of incidents¹⁰. What follows for the standards is a position on access rather than a position against it. As a Five Eyes partner hosting the training of these models, Australia should expect pre-release access for its security agencies and evaluators, written into the conditions of permission to train here rather than sought afterwards. No single vendor's closed models should be privileged for critical settings, and Australian agencies, researchers and defenders should retain access to open and closed models from several jurisdictions, run on infrastructure Australia governs.

***Recommendation 7:** Require notification of defined incidents to a named Australian channel within a statutory timeframe, with penalties for late, incomplete or misdirected notification and a duty to cooperate with external investigation, alongside independent evaluation, red-teaming and supply-chain assessment as conditions of authorisation to train at scale; give the Australian AI Safety Institute pre-deployment access and formal links to Australian academic expertise; and base data-erasure requirements on demonstrated machine-unlearning capability.*

***Recommendation 8:** Make pre-release access for Australian security agencies and evaluators an enforceable condition of permission to train frontier models in Australia, decline to privilege any single vendor's closed models for critical settings, preserve access for Australian agencies, researchers and defenders to open and closed models from several jurisdictions on Australian-governed infrastructure, and fund the independent evaluation and assurance capability these conditions depend on.*

8. Investing in local capability, research and innovation

This is where the standards can convert a construction boom into lasting capability, and where the design choice between reserved compute and tokens will decide the outcome. Available evidence indicates that Australia's public compute base is small by international standards. A 2024 survey of participating Australian institutional systems, including those across every university, CSIRO and national facility, identified around 2,200 accelerators, most several generations old, while a single one of the United Kingdom's two public AI systems, Isambard-AI, comprises 5,448 current-generation superchips. Public research receives an estimated one to five per cent of usable frontier GPU-hours in Australia, the 2026 merit allocation round was oversubscribed roughly threefold, and a UNSW materials-discovery project that needed between 500 and 1,000 accelerators ran on the Alps supercomputer in Switzerland because only around eight comparable GPUs were available here¹¹.

Free-token and credit offers are the opening position of every major platform seeking a government relationship. They are a customer-acquisition strategy. When credits expire, the researchers and agencies that built workflows around a vendor's platform face high switching costs and recurring offshore spend, and access can be repriced, restricted or withdrawn by commercial decision, policy change or export control, as Australian users of some frontier models experienced in 2026. Accepting tokens does not give Australia the ability to run, adapt and validate models locally, train on sensitive national datasets or hedge against GPU scarcity. Reserved capacity on Australian-governed terms does.

⁹OpenAI and co-signatories, 'A call for collective action on cyber defense', 27 August 2026.

¹⁰CISA and partner agencies, Cybersecurity Advisory AA26-231A, 19 August 2026; RUSI, 'Beyond the Hype: AI, Ransomware and Business Models', 27 August 2026; J. Lyne (CEO, SANS Institute), 'What the AI Warning Letter Completely Missed', Dark Reading, 3 September 2026; Australian Signals Directorate, Annual Cyber Threat Report 2024-25, October 2025. Note that 'AI-enabled attack' is not a standardised reporting category and vendor definitions differ.

¹¹UNSW submission to the National Science and Technology Council on future high-performance compute and data requirements, 30 January 2026.



The mechanism we recommend was developed with the national research community and is set out in our brief to DISR on Expectation 5¹². AI infrastructure operators and owners, and the compute tenants behind them, should reserve a defined and independently verifiable share of usable compute capacity, benchmarked for accelerator capability, interconnect, storage, software and service level for Australian researchers, start-ups, SMEs, regional universities and public-interest users. The obligation should be met in kind or transferred to another provider through a tradeable obligation. Access should be coordinated through a single national allocation framework and delivered by multiple qualified providers, drawing on Australia’s existing Tier 1 and Tier 2 research infrastructure providers. This would ensure that entitlements aggregate into one usable national pool, rather than fragmenting across sites, vendors and states. The AI-for-science pathway should align with the National Research Infrastructure Roadmap and NCRIS, while distinct industry, public-mission, evaluation and regional-access pathways are retained. The mechanism requires no new Commonwealth appropriation. A national research community workshop convened through AeRO in August 2026 agreed seven design principles.

Principle	Requirement
Additionality	New private-sector contributions must strengthen, not displace, existing sovereign research infrastructure and public investment.
Open access architecture	No vendor lock-in. Portability, interoperability and the ability to move workloads and data between environments.
Researcher protections	Clear IP ownership, secure data handling and safeguards against inappropriate reuse of research assets or models.
Predictability	Stable pricing and reliable access across multi-year grant and research lifecycles.
Transparency	Clear reporting, benchmarking and accountability for capacity promised, capacity delivered and research utilisation.
Trusted allocation	Access coordinated through established national research infrastructure mechanisms rather than ad hoc new structures.
Regular review	Arrangements reviewed at least annually because AI infrastructure and capability requirements change rapidly.

Hyperscalers and neocloud providers should carry these obligations alongside frontier developers, because they control most of the installed compute and will host much of the training activity. The different actors should not, however, be treated as economically interchangeable. Australian-controlled neoclouds can deliver reserved compute while building local capability in GPU systems, orchestration, cybersecurity, model deployment and technical support. Their employment, procurement, research partnerships and reinvestment retain more of the value chain within Australia. This directly supports Australian innovation and contributes to the domestic tax base from which public services are funded.

The latest ATO corporate tax transparency data shows why gross revenue and investment announcements are poor proxies for national benefit. The Australian entities of Meta, Amazon, Google, Netflix, Tesla, Apple and Microsoft reported approximately \$44.3 billion in total income and \$748.9 million in company income tax for 2024-25, equivalent to about 1.69 per cent of reported income¹³. Microsoft Datacenter (Australia) reported \$2.3 billion

¹²S. Keay (UNSW AI Institute, representing ATSE) and C. S. Ong (CSIRO), ‘Operationalising Expectation 5 for Data Centre and AI Infrastructure Developers’, brief to DISR, 17 August 2026, drawing on a national research community workshop convened through AeRO on 7 August 2026.

¹³ Jonathan Barrett, “Billions of dollars, none of it taxed: ATO reveals the tech giants paying zero income tax”, *Guardian Australia*, 1 October 2026.

in total income and no taxable income. These figures demonstrate that extensive Australian revenue and infrastructure activity can produce a limited domestic fiscal return where intellectual property, related-party payments and profits sit offshore. Any assessment of national benefit should therefore examine the value retained in Australia, including Australian taxes paid, local payroll and supplier expenditure, research and development, intellectual property ownership and profits reinvested in Australian capability.

Compute alone will not build capability. Australia has almost no major corporate AI research presence, while Israel hosts more than 500 corporate AI labs, and the R&D Tax Incentive taper above \$20 million in turnover has been cited in the Commonwealth Bank's decision to place its AI hub in Seattle¹⁴. The research strengths Australia already holds in machine unlearning, anomaly detection, uncertainty quantification and privacy-preserving learning are the raw material of that assurance capability, which is why the contributions sought here and the safety conditions in section 7 are the same argument approached from two directions. We support the Kingston AI Group's proposal for a mandatory local AI research and development presence for major AI companies operating here, with locally based research staff, state distribution and retention of intellectual property in Australia, paired with R&D Tax Incentive reform and a dedicated AI research funding vehicle on the rural research and development corporation model¹⁵.

The indicators the paper asks for are measurable. Reporting should cover usable compute delivered to public-interest users, normalised for accelerator capability, interconnect, storage, software and service level, together with utilisation, time-to-access, user diversity and expert support-hours, locally based research and engineering headcount, Australian AI jobs created, recognised qualifications and skills transfer, datasets brought under approved governance, Australian models supported, local research and development expenditure, Australian intellectual property created or retained, local supply-chain expenditure, engineering placements, profits reinvested in Australian capability, Australian taxes paid, and the evidence that users and capability persist after pilots and credits end. Construction expenditure and land development are inputs and should not be counted as capability or national benefit in isolation.

***Recommendation 9:** Require, as a binding condition of authorisation, that AI infrastructure operators and their compute tenants reserve a defined and independently verifiable share of installed compute for Australian public-interest users, discharged in kind or through a tradeable obligation and administered through a single national counterparty using existing research infrastructure providers under the seven principles above. Recognise Australian-controlled neoclouds as capability partners in delivering reserved compute and assess all proposed contributions against measurable domestic value retention, including local skills, research and development, supplier expenditure, intellectual property, reinvestment and Australian taxes paid. Adopt a mandatory local AI research presence requirement for major AI companies, reform the R&D Tax Incentive, and report on capability outcomes rather than construction expenditure.*

9. Supporting Australian creators, media and publicly funded research

Australia's copyright law should be reformed, with specialist input, to provide clear exceptions for education, publicly funded research and other public-interest activities - including research in, with and about AI, security testing and interoperability - and to prevent those exceptions being overridden by contract, while protecting creators' ability to retain their rights. Reform should not be treated simply as an inducement for infrastructure investment. The scale of proposed investment is itself evidence that Australian data and content have strategic value, and that value flows permanently to the model owner once training occurs. The standards should require developers training in Australia to keep auditable records and disclose sufficient information about the Australian creative, cultural, media and publicly funded research material used and, where a use falls outside an applicable exception, to obtain licences directly or through collective arrangements with safeguards for creators' rights retention and ongoing remuneration where that material generates commercial value. Publicly funded research and nationally significant datasets should be used in Australian models, and should not be transferred into foreign-controlled models or reused by a provider for secondary commercial training without a governed

¹⁴Kingston AI Group, submission to the Joint Select Committee on Artificial Intelligence, September 2026, drawing on the Stanford AI Index 2026.

¹⁵UNSW AI Institute, Submission to the Joint Select Committee on Artificial Intelligence, 14 September 2026.



agreement covering attribution, return and benefit sharing. Separate, enforceable protections for Indigenous Cultural and Intellectual Property should be developed with Aboriginal and Torres Strait Islander peoples, recognising that generative AI can enable appropriation where no formal copyright is engaged.

Recommendation 10: *Reform copyright law, with specialist input, to provide clear exceptions for education, publicly funded research and other public-interest activities, prevent those exceptions being overridden by contract, and protect creators' ability to retain their rights; require auditable records and disclosure of Australian training material; provide transparent licensing and remuneration where uses fall outside applicable exceptions; condition the use of publicly funded data on governed agreements; and establish separate protections for Indigenous Cultural and Intellectual Property.*

10. What would encourage frontier AI training in Australia

The paper asks what additional reforms or incentives would encourage frontier training. The most effective incentive is a nationally coordinated, provider-neutral ecosystem worth training in: a funded public compute base integrated with high-speed networks and storage, governed national datasets and models, software, assurance and engineering support, distinct science, industry and public-mission access pathways, and a skilled workforce. Every comparable economy funds public compute alongside private investment. Norway welcomed a large private Stargate project while continuing to fund its publicly owned Olivia supercomputer and new national AI research centres, and Switzerland used public compute to train Apertus, an openly documented national model that Australia already supports as a contributing compute partner. South Korea has gone further, requiring that at least half of each government-supported AI service use Korean models, which gives domestic developers customers, compute and operating experience at national scale¹⁶.

Digital infrastructure should be recognised as critical national infrastructure for budget, planning, energy and water purposes, as it already is for cybersecurity under the SOCI Act, so that public compute is planned and funded like other national infrastructure. Public procurement can then create early domestic markets for Australian models and AI products, and procurement conditions on portability, data export rights and open interfaces will protect public investment from lock-in. Australian domain models for Australian conditions, in health, climate, water, energy, critical minerals, agriculture, public administration, defence and education, are where Australian capability is realistic and where frontier developers will find partners worth training with.

Biomolecular science shows what is at stake in that recognition, and it is the most thoroughly validated example of AI transforming research. AI-driven structural biology and protein design are already changing drug discovery, diagnostics, biotechnology, agriculture and fundamental biology, and Australia is building the national BioCommons Structural Biology Platform to give researchers across institutions supported access to difficult-to-install workflows, GPU compute, storage, training and technical expertise, a democratic baseline for the whole research system alongside the larger allocations specialist groups need. The platform has no secure funding beyond mid-2028 and has already reduced staff, so Australia risks losing the people and validated workflows needed to operate it just as it reaches national users. Buying access to overseas tools is not a substitute: it builds no local expertise, no validated Australian workflows, no Australian intellectual property and no capacity to work securely with sensitive biological data. Platforms of this kind are the domain-integration layer that turns raw GPU-hours into research capability; they should be funded on stable long-term terms from the public-interest compute base recommended here, and they are natural delivery vehicles for the reserved compute obligation in section 8.

Recommendation 11: *Recognise digital infrastructure as critical national infrastructure for budget and planning purposes, fund a public-interest compute base in the forward estimates with stable long-term support for national research platforms such as the BioCommons Structural Biology Platform, and use public procurement to create early domestic markets for Australian models and AI products, with vendor-neutral terms as a condition of public spending.*

¹⁶Ministry of Science and ICT (Republic of Korea), 'MSIT Begins AI for All Project to Offer AI Services through a Sovereign AI Model', July 2026.

11. Interaction with existing legislation and a simple national framework

The proposed federal reporting and transparency requirements will interact with the NGER Act reforms on facility-level energy and emissions data, the SOCI Act positive security obligations for critical data storage or processing assets, the Foreign Acquisitions and Takeovers Act review process, the proposed digital duty of care for online services, which would apply a systems-based safety obligation to the services through which models trained here reach Australians and should use consistent definitions of reportable harm and a common reportable platform, so that developers do not face duplicate or inconsistent incident reporting, the AEMC access standards and the National Electricity Rules, the REGO offset, firming contract and market registration arrangements agreed by Energy Ministers in July 2026, NSW's Data Centre Guidelines, the South Australian legislation in development and state water allocation regimes. One national allocation framework should coordinate the compute obligation. Across other obligations, PM&C's Office of AI should coordinate common standards, identifiers and a shared reporting gateway that supports 'report once, use many times', while existing Commonwealth, state and territory regulators retain their statutory responsibilities. A single annual compliance report covering facility-level energy, water, workforce and capability would reduce duplication for proponents and make the standards easier to enforce. Two further points follow from our earlier evidence. Data held onshore by a US-headquartered provider may remain reachable under instruments such as the US CLOUD Act, so onshore hosting should not be described as sovereign capability in itself. And obligations should be reviewed at least annually, because AI infrastructure and the capability it supports change faster than most regulatory settings.

***Recommendation 12:** Establish a common reporting gateway, identifiers and data definitions so proponents can report once for use by relevant regulators; retain existing statutory responsibilities and coordinate the compute obligation through the separate national allocation framework; align safety reporting with the proposed digital duty of care; and review the standards at least annually against whether the safety conditions are working.*

12. A national ambition

The standards are a means to an end. Australia has a National AI Plan, Data Centre Expectations and now proposed AI standards, but none of them sets a funded national ambition to build Australia's own AI capability. The current investment wave gives government leverage it has not had before and may not have again. Used well, it can seed an Australian AI industry with real capability in compute, data, models, skills and products. Used poorly, Australia will have built the infrastructure for others to profit from and will spend the next decade buying back access to its own data and compute at retail prices.

***Recommendation 13:** Express the AI standards within a funded national AI strategy that sets targets for Australian compute, models, researchers, companies and independent evaluation and assurance capability, with commitments that are funded, measurable and enforceable.*

About UNSW Sydney

UNSW Sydney is Australia's #1 globally ranked university with strengths spanning engineering, science, medicine and health, law, business, arts, design and the social sciences. Through research, education and partnerships with government, industry and communities, UNSW works to translate knowledge into practical outcomes and contribute to a just and sustainable future.

About the UNSW AI Institute

The UNSW AI Institute (UNSW.ai) is a pan-university research institute uniting more than 350 researchers across every UNSW faculty, including Engineering, Science, Law and Justice, Medicine and Health, Business, Arts, Design and Architecture, and UNSW Canberra. Its mission is to translate AI research into national impact through partnerships with government, industry and the community, and it is a standing resource to government on AI safety, sovereign capability and responsible deployment. As Australia's #1 globally ranked university, UNSW hosts Australia's largest and top-ranked engineering faculty, making it a principal training ground for the AI and robotics workforce Australia will need over the coming decade.

About the UNSW Energy Institute

The UNSW Energy Institute connects researchers across the University with government, industry and communities to develop practical solutions for a clean, affordable, reliable and equitable energy system. Its work spans energy technology, power-system transformation, market analysis and policy design, drawing on UNSW's internationally recognised strengths in renewable energy and energy innovation.



UNSW
SYDNEY

UNSW responses for the PM&C online submission form

'Getting it right: Building AI infrastructure that works for Australia'. A working paper mapping UNSW's submission to the consultation questions.

Last Reviewed 9 October 2026. Submissions close 5 pm AEDT Friday 9 October 2026.

How to use this paper

The government is receiving submissions through an online form with a field for each consultation question, an optional additional comments field at the end of each of the three Parts, and one supporting document per Part (5 MB, PDF or Word). This paper presents each question in form order with a response drawn from UNSW's submission, drafted to stand alone in its field, followed by the suggested text for each Part's additional comments field and the file to attach. The form does not require every question to be answered; questions outside the University's expertise are marked to leave blank. The full submission should be attached in all three Parts so the complete case reaches whichever Part a reader opens. Because each field is read alone, some points recur across answers by design.

Across the three Parts, UNSW applies a common policy test: an obligation should be imposed on the entity able to control the relevant outcome, and a claimed national benefit should count only if it produces capability that is additional, measurable, accessible to Australian users and reasonably durable if a foreign supplier changes its commercial or government-imposed access arrangements.

Part 1: Managing growth in large data centres and the AI training opportunity

Capacity thresholds for application of the AI standards

Is a single threshold such as an energy capacity (MW) sufficient to identify the data centres that should be subject to the AI standards? If not, what additional thresholds (for example water use, compute [FLOPS], or electricity consumption [MWh]) should be considered?

No. An energy-connection threshold is a sensible primary trigger because it aligns with the AEMC access tiers and the grid impacts the standards most need to manage, but it cannot carry the framework alone. If the standards were solely focused on energy use and renewable energy requirements, it might make sense. However, a 40MW facility filled with current-generation accelerators has a different national significance from a 40MW facility of enterprise storage, and a training cluster can be assembled across several sites that each sit below any energy threshold. The standards proposed for AI training relating to security, safety and transparency concerns have little meaningful connection to energy capacity. The University thus recommends a second trigger for the AI training obligations based on training compute, on the model used in California's SB 53 and New York's RAISE Act, which set the line at 10^{26} operations counted cumulatively across the initial training run and any subsequent fine-tuning or material modification, so that the party undertaking large-scale training is captured wherever the equipment sits, with installed accelerator capacity available as a secondary trigger for facility-level reporting. The compute threshold should not be treated as a permanent proxy for capability or risk: it should sit in an updateable instrument, be reviewed regularly, and be supported by transparent capability- or risk-based designation criteria. Water could operate as a third trigger for water-related standards in stressed catchments, since local water impact depends on the catchment rather than the facility.

Does a tiered approach sufficiently capture different types of facilities and ensure proportionate obligations relating to Australia's energy transition, water management, and communities?

The tiered structure and the 150MW portfolio aggregation test are supported. The design risk is not the tiers but the entity they bind. Under prevailing operating models the facility owner obtains the approvals while hyperscaler and neocloud tenants own the accelerators and control the workloads, the access and much of the load profile. The standards should be drafted so that each obligation attaches to whichever party controls the relevant activity, with capability and compute obligations reaching tenants through the licence to operate and direct undertakings from anchor tenants of the largest facilities.

What can Australia learn from other jurisdictions that regulate large data centres, particularly in designing obligations that are proportionate to potential community impacts while maintaining incentives for investment?

That obligations can follow the activity rather than the building without deterring investment. Ireland's grid connection conditions reach the parties who control load. The EU energy efficiency framework extends reporting to colocation customers who own their own equipment. California and New York define frontier training obligations by training compute rather than by facility, both at 10^{26} operations counted across the initial run and any later fine-tuning, and the EU AI Act uses the same approach at 10^{25} operations as a rebuttable presumption under Article 51(2). And the resources sector in Australia has always cascaded licence conditions to whoever conducts the activity. Investment has continued in each of those settings, because what investors prize is certainty, not the absence of conditions. The positive infrastructure lesson is equally important: international AI-factory models show how coordinated compute, data, software, engineering support and distinct science and industry access pathways can convert infrastructure investment into durable domestic capability.

Approach to the application of the AI standards

At what point in development processes should mandatory requirements apply for large data centres?

From the point at which a project seeks its energy connection agreement or state planning approval, whichever comes first. A plan demonstrating how the facility will meet its obligations should be lodged before connection and published for facilities above 100MW, with obligations linked to actual load scaling as capacity is drawn down.

What can be learned from approaches taken in other jurisdictions to ensure very large data centres currently in development pipelines are appropriately captured?

Pipeline projects that have not yet connected should be inside the framework from commencement, with commonly controlled projects aggregated so a build cannot be divided into sub-threshold parcels. Jurisdictions that grandfathered pipelines saw a rush to connect ahead of the rules and a two-class market afterwards.

Should there be a transition period for data centre proponents to meet their obligations, for example submitting a plan within a certain timeframe after commencing operations? If so, at what point in the development cycle should this occur?

A short transition for existing facilities is reasonable. Any facility seeking government facilitation, prioritised approvals, public land or public procurement should meet the standards in full from the day the legislation commences, since those benefits are the consideration for the conditions.

How should governments anticipate and respond to large data centre energy loads which may become operational faster than additional energy generation can be completed to offset their demand?

Staged connection agreements, demand flexibility written into the connection itself, and load-linked obligations that scale with drawdown. Governments should also assess the proposed data-centre pipeline through integrated temporal and geographical scenarios of large loads, generation, storage, firming and network capacity, including compatibility with Australia's legislated emissions targets. Many training and batch-compute workloads can provide greater demand flexibility than latency-sensitive inference services, provided that checkpointing, scheduling, hardware and contractual arrangement are designed for interruption. Flexibility obligations should therefore be based on demonstrated technical capability and workload characteristics rather than an assumption that every training workload can be curtailed without material cost or risk.

Following the passage of the proposed legislation, when should compliance start? For example, immediately, or staggered for different requirements, recognising practical constraints around planning and approvals, and to ensure regulation is proportionate. How should certainty for, and impacts on, the Australian community be balanced with these considerations?

Reporting, transparency and the capability obligations in Part 3 should start immediately, since they depend on contracts rather than construction. Physical requirements should follow a published schedule fixed at passage, to give investors certainty.

Part 1: suggested additional comments field

Paste into the additional comments field for Part 1.

The UNSW AI Institute, one of the largest concentrations of AI research capability in the country with more than 350 affiliated researchers, asks that one distinction govern the whole framework: hosting is not capability. A foreign-owned facility on Australian soil running foreign models under foreign control supplies a service that could be exclusively sold outside of Australia; it does not give Australia the ability to build, adapt, evaluate or run AI systems of its own. That ability is what was missing in June 2026 when a United States export control directive withdrew frontier models from Australian users with little notice. The standards should be designed to convert the current investment wave into things that remain under Australian control: reserved compute, locally based researchers, Australian-held intellectual property and Australian models. Digital research infrastructure should be recognised as critical national infrastructure for long-term budget and planning purposes, in the way roads, transmission and water already are. Our full submission, attached, makes thirteen recommendations to these ends.

Attach: the full Institute submission (PDF export, under 5 MB).

Part 2: Requirements for large data centres

Renewable electricity certificate obligations

How could obligations be phased in to account for the time needed to build data centres and new renewable projects?

On a short glide path matched to realistic construction times for new generation and storage, with the obligation binding from commencement and the required share of new capacity rising on a published schedule. What the phase-in must never permit is indefinite compliance through certificates from existing plant, which converts none of the demand into new supply. The AEMC's advice to Energy Ministers, agreed in July 2026, includes an option for data centres to draw temporarily on existing renewables and make up the shortfall by surrendering additional certificates in later years; that make-up period should carry a hard time limit.

How should 'new' renewable capacity be defined in relation to data centre investments?

Generation and storage reaching financial close after the facility's commitment, which would not have proceeded on the same terms without it. Additionality is the entire point of the obligation.

At what capacity threshold should an obligation to retire or surrender renewable energy certificates apply?

From the baseline tier, scaling with the tiers. The University defers to energy market specialists on rate design and holds one principle: the certificates must represent additional capacity.

Should the information on REGO certificates be used to time-match electricity use in certain periods? If so, what should these periods be?

Yes, for facilities above 100MW, during evening peaks and AEMO-defined stress periods. Annual matching lets a facility report renewable operation while drawing fossil generation at the hours that determine both emissions and prices. UNSW's 24/7 TRUZERO research found that full 24/7 matching up to 80% presently incurs relatively little cost, matching the final 20% is where the costs and complexity lie; this supports starting with the periods of greatest system need rather than imposing an immediate 24/7 requirement.

Are there any complexities in how data centres could acquire certificates to meet a renewable electricity obligation?

Two, beyond the additionality principle. Rules should prevent double-counting where storage charges from the grid and later discharges as certified renewable electricity, and should allow on-site and near-site generation to count without unnecessary certification overhead.

What incentives or obligations would encourage data centres to invest in on-site generation, batteries or other flexible demand measures, to minimise energy usage during peak pricing?

Write a minimum level of flexibility into the connection agreement as an obligation, as Ireland does, with market mechanisms used to procure flexibility beyond that minimum. Set the obligation according to demonstrated technical capability and workload characteristics, which may justify higher obligations for training-dominated facilities. Non-renewable back-up generation should be subject to

defined limits and public reporting, with permitted use tied to specified emergency, scarcity or system-security conditions.

Three further levers would help. First, every large facility already installs battery backup (UPS) capacity; with appropriate controls it can provide frequency and other grid services rather than sitting idle, so new backup systems should be required to be grid-capable. Second, on-site and near-site solar and storage (BESS / eSTATCOMs) should count towards the renewable electricity obligation, particularly for regional facilities co-located with renewable energy zones. Third, as AI facilities move towards high-voltage DC power distribution, directly coupling DC-native solar and batteries to the facility can avoid conversion losses, an opportunity for Australian research and demonstration.

Where electricity is being purchased from a state-owned operator, what should the AER consider when assessing whether to modify a data centre's renewable electricity obligation?

Whether the supply is additional renewable generation and whether modification shifts costs or emissions to other users. Public ownership of the counterparty does not alter either test.

What unintended consequences or implementation challenges should the government consider in designing a renewable electricity framework for data centres?

A framework satisfied by existing certificates would let data centre growth absorb renewable generation the grid was already counting on, with prices and emissions rising for everyone else while giving a false impression of compliance. A framework that binds only facility owners will miss the tenants who drive the load, which is why the University's second recommendation is that obligations attach to whichever party controls the activity. Offsetting demand should be treated as the floor, not as a benefit: without it, data-centre growth raises prices and emissions for every other electricity user.

Energy usage and efficiency

What performance metrics or ratings would reflect an appropriate standard for energy usage and efficiency in data centres?

The proposed design PUE thresholds and a minimum NABERS energy rating are a workable baseline, phased by tier. The University's one design rule set by the UNSW AI Institute is that energy and water metrics (WUE) be set and reported as a pair, because PUE was designed for cloud-era workloads and can be improved by cooling choices that consume more water. A facility should never be able to buy a better energy rating by using more water. Large facilities should also report their hourly carbon-free energy share, grid-interaction and disturbance ride-through performance, and the flexibility actually delivered, because a highly efficient facility can still impose a large burden on the power system.

Should the Commonwealth establish national baseline energy and water efficiency metrics for data centres? What might be the benefits and/or unintended consequences of this approach?

Yes. One national baseline prevents jurisdiction shopping and gives investors a single design target. The consequence to guard against is optimisation to the metric, which joint PUE and WUE regulation and public reporting of operational rather than design figures addresses.

Should all data centres be required to meet energy and water performance thresholds? If not, how should requirements change or scale for different data centres?

Requirements should scale with the tiers, with sub-30MW facilities exempt from mandatory thresholds but drawn into transparency reporting where their water draw is material in a stressed catchment.

How frequently should data centres be required to transparently and publicly report their energy use, efficiency, and emissions as part of plans demonstrating compliance with mandatory requirements?

Annually in public, quarterly to the regulator, covering operational PUE and WUE, emissions, fossil back-up run hours and, for facilities carrying Part 3 obligations, compute capacity promised and delivered to Australian users. The University also recommends a public national register of large data centre proposals and approvals showing connection capacity, water source and cooling design, so communities and researchers see the national picture rather than one facility at a time.

Water

Should data centres be required to meet performance metrics or ratings for water usage and efficiency in data centres?

Yes. NABERS rates data centres on energy and has not rated them on water, but that gap is closing. A NABERS water efficiency framework for data centres is under development and due for completion in October 2026. A minimum NABERS water rating should be mandatory as soon as the tool is operational. Until then, operational WUE should be reported under ISO/IEC 30134-9 with the water source and catchment stress disclosed.

If so, are existing metrics (for example, WUE ratings) fit for purpose, or would new or adjusted metrics be more appropriate (for example, by factoring water stress regions and water sources into WUE ratings)? Should a minimum NABERS data centre rating, based on WUE and agreed by states and territories, be required once a facility is operational?

WUE should be weighted for source and catchment stress as the paper proposes, and always set jointly with PUE so the trade between them is visible. Yes to a mandatory NABERS water rating once the October 2026 framework is operational. This needs to be agreed nationally so a single rating serves every jurisdiction.

Should water-use thresholds trigger additional requirements? If so, what thresholds should apply, and should they align with the tiered approach outlined on page 8?

Yes, and the water trigger should run alongside the energy tiers rather than inside them. A modest draw in a stressed catchment matters more than a large draw where recycled or desalinated supply is abundant.

What aspects of water use should data centres transparently and publicly report?

Peak and maximum demand, source split between potable and non-potable, operational WUE, discharge volume, quality and temperature, drought resilience measures, and assessed effects on the catchment and existing users.

Should a nationally consistent hierarchy of water sources be adopted for data centres, prioritising non-potable/recycled water with potable water as last resort?

Yes, with closed-loop liquid cooling the expected design for large facilities in stressed catchments and potable supply a disclosed exception rather than a default.

Where recycled or non-potable water supplies are not reasonably available, what alternative approaches might be promoted to achieve equivalent water security outcomes?

There are three main approaches, which should be considered in order of priority: (1) Co-investment in recycled water infrastructure with a firm transition date; (2) within-catchment offsets; (3) funded local efficiency programs, with equivalence judged at catchment level under drought conditions.

Should data centres be required to demonstrate that proposed water use will not adversely affect local water security, drought resilience, environmental water requirements, or existing water users as part of a plan outlining their compliance?

Yes, before connection approval, data centre water usage should be assessed against drought scenarios rather than average years, consistent with the Water Ministerial Council minimum requirements the University has recommended the standards adopt.

No net costs on consumers and communities**What gaps (if any) exist in electricity and water infrastructure charging and cost allocation arrangements that could result in consumers bearing costs associated with data centre-driven network or infrastructure augmentation?**

The University supports the principle of full recovery of attributable augmentation costs and defers to the AER and utilities on mechanism.

What are the barriers to data centres investing in sufficient new renewable energy and storage capacity to mitigate impacts to wholesale electricity costs?

Barriers include approvals sequencing across jurisdictions and the lack of long-term revenue certainty for some new generation and firming projects. Large, predictable data-centre demand can help close the financing gap if the associated capacity is genuinely additional. A single, time-bound national pathway with concurrent assessment, available to projects that accept the standards' obligations, addresses the barrier without weakening the conditions.

What are the barriers to data centres investing in initiatives that would mitigate impacts to other water users, such as using recycled water?

Availability of recycled supply at the site, which is why the standards should be able to require co-investment in recycled infrastructure with a transition timeframe where the source hierarchy cannot otherwise be met.

What requirements should be placed on data centres to operate flexibly to support the operation of the electricity grid and improve system efficiency while protecting consumers from increased energy costs?

Defined curtailment obligations during AEMO-declared stress and scarcity events, restrictions and public reporting on fossil back-up, and higher flexibility duties for training-dominated facilities. Connection requirements should also address how large inverter-based loads respond to grid disturbances, complementing the AEMC rule change already under way on technical access standards for large inverter-based loads. Real-time simulation should test rapid or cascading disconnection risks and whether coordinated battery systems can support both facility resilience and power-system stability. All of it belongs to the connection agreement and the licence, flowed down to the tenants whose contracts otherwise guarantee uninterrupted capacity.

What barriers would prevent data centres from adopting demand flexibility measures to support the wider electricity system?

Tenant service agreements. The landlord cannot curtail load its contracts oblige it to serve, which is a further reason obligations must reach the tenant.

What tools could be used by governments to require data centres to adopt demand flexibility?

Connection agreements and the AEMC access standards, with the duty carried in the licence to operate and flowed down contractually. Governments should also test whether demand-response and frequency-service markets can procure additional flexibility, including from other consumers funded by data-centre demand, without replacing mandatory connection conditions. The AEMC's July 2026 advice proposes voluntary flexible-demand connection agreements; the University supports these but recommends that a minimum level of flexibility be a mandatory condition of connection for large facilities.

Aside from infrastructure and network augmentation costs attributable to their energy and water demand, are there other significant potential costs from data centre activity, and how might accountability be provided for managing these?

Other significant costs include pressure on scarce skilled trades, competition for land otherwise available for housing or agriculture, and the opportunity cost when Australian land, power, water, skills and public facilitation support foreign workloads without producing durable Australian capability. A 2024 survey of participating Australian institutional systems identified around 2,200 accelerators, most several generations old; the 2026 national merit round was oversubscribed roughly threefold; and a UNSW materials-discovery project needing 500 to 1,000 accelerators ran on the Alps system in Switzerland because around eight comparable GPUs were available here. Accountability comes from the broader capability obligation in Part 3, the reserved compute integrated with the data, software, engineering, assurance and translation supported needed to use it, and public reporting of capacity promised and delivered.

Community engagement, location and connectivity**What does best practice community engagement look like for data centres, including with nearby businesses?**

Engagement should occur before the site and design are fixed and continue through construction and operation. Proponents should publish plans with expected resource demands and environmental impacts, with commitments carried as enforceable licence conditions. The University defers to communities, councils and First Nations organisations on form, and recommends the national register above as the transparency foundation needed for best practice community engagement.

Engagement with Aboriginal and Torres Strait Islander peoples should begin before site selection where a proposal may affect Country, cultural heritage, water, Indigenous rights or community interests. Applicable obligations should include culturally appropriate engagement, support for informed participation, protection of Indigenous Cultural and Intellectual Property, and Indigenous data governance where data, models or research activities involve Indigenous knowledge or communities.

What role should the Commonwealth's mandatory requirements play in relation to location and community engagement?

A public national register of proposals and approvals, showing connection capacity, water source and cooling design, so that engagement happens with full information and the cumulative picture is visible.

How can communities and industry work together to ensure data centres deliver lasting community benefits? How can upfront commitments be reflected in state and territory planning processes or community engagement?

Carry the commitments as licence conditions reported against annually, and measure benefit in outcomes: skills completions, local employment, and the compute and research capability retained in Australia, rather than construction spend.

Should data centres be required to locate a minimum distance away from sensitive sites such as schools or residential homes? If so, what should the mandatory minimum distance be? Which level of government is best placed to issue and administer location requirements?

No Institutional position. Leave blank or defer to state planning systems.

How important is proximity to population centres for data centre operators given improvements in connectivity and advances in fibre optics?

Training is largely latency-insensitive and can be sited where renewable energy and water headroom exist; inference serving users is different. The standards should distinguish the two workloads in siting policy rather than treating all data centres as one class. Proponents seeking approval in constrained locations should be required to demonstrate the latency or connectivity needs of the proposed workload and why suitable alternative locations are not viable.

How should competing land uses be managed and prioritised, for example where land is also suitable for housing or agriculture?

No Institute position. Leave blank.

Should states and territories work with energy, water and telecommunications companies to proactively identify sites for data centres that could support future builds away from populated areas? What efforts are already underway?

Yes. Pre-identified zones with concurrent approvals would give effect to the single national pathway the University recommends, with the standards' conditions attached to the zoning. Site identification should prioritise areas near renewable energy zones and regional industrial precincts where adequate energy, water and telecommunications capacity and opportunities for skilled regional employment can be demonstrated.

What are the expected domestic and international telecommunications connectivity requirements for data centres, and what are the challenges in upgrading existing or deploying new networks?

Compute without high-speed networks, storage and secure data movement is not usable research infrastructure. Requirements should cover resilient, path-diverse domestic fibre; research and education network connectivity; diverse international routes; predictable data-egress costs; and integration with storage and data platforms. AARNet and national research infrastructure providers should be involved in their development.

Should data centre proponents be required to explore options with government and industry to uplift community telecommunications connectivity alongside their own investment?

Yes, as an eligible community contribution that supplements rather than substitutes for the research and capability contributions in Part 3.

Skills and training

How should data centre operators be required to provide skills and training opportunities, including for under-represented cohorts in skilled workforces, during construction and into ongoing operation?

Through apprenticeships, TAFE pathways, undergraduate and postgraduate programs and industry-linked PhD placements, measured by completions and employment outcomes with published targets for under-represented cohorts. Vendor-specific product training should not count toward the obligation: its main effect is to deepen dependence on the vendor's platform.

Large data centres are also energy facilities, needing engineers and technicians in high-voltage electrical systems, on-site solar and storage (BESS / eSTATCOMs), power electronics and cooling. UNSW's School of Photovoltaic and Renewable Energy Engineering trains this workforce; from 2027 UNSW will run a cross-school Sustainable Data Centres Vertically Integrated Project spanning computing, electrical, renewable energy, mechanical and water engineering; and UNSW is establishing the Centre for Integrated Renewable Energy Training & Innovation in Western Sydney, focused on tertiary harmonization of battery storage and solar operations and maintenance. Contributions to programs of this kind should count towards the obligation.

What broader workforce development obligations should apply to data centre operators during ongoing operations?

Development of the scarce roles that sovereign capability depends on: research software engineers, GPU systems specialists, model evaluators and data stewards. These skills take years to develop, and the obligation should grow the supply rather than relocate existing people between employers. The deeper research workforce obligation belongs to AI developers under Part 3, through locally based PhD-level staff, hybrid university appointments and Higher Degree Research (HDR) scholarships. On the energy side, the equivalent scarce roles are high-voltage electrical, battery storage and power-electronics engineers and technicians, whose supply the obligation should also grow.

Part 2: suggested additional comments field

Paste into the additional comments field for Part 2.

Strategically located and coordinated data-centre investment can support productivity, sovereign capability and regional economic development while bringing forward additional generation, storage, firming and network capacity, provided costs are not transferred to households, businesses or host communities. The standards should build on, not duplicate, the arrangements Energy Ministers agreed in July 2026 following the AEMC's advice, with a hard time limit on any period of reliance on existing renewables and a mandatory minimum of demand flexibility; offsetting demand is the floor, not a benefit. Three design rules would carry every requirement in this Part. First, attach each obligation to whichever party controls the relevant activity. The facility owner holds the approvals; hyperscaler and neocloud tenants control the accelerators, the workloads and the load profile, and requirements that stop at the landlord cannot reach the choices that matter. Ireland's connection conditions, the EU's colocation reporting and Australia's own resources licensing all show conditions following the activity. Second, never regulate or report energy without water. PUE can be improved by cooling designs that consume more water, so the two metrics must be set jointly, and the NABERS water efficiency framework for data centres due in October 2026 should become a mandatory rating as soon as it is operational. Third, count only skills contributions that build transferable capability, including the energy engineering workforce that large facilities depend on, rather than vendor-specific product training. The University also recommends a public national register of large data centre proposals and approvals, since transparency is the foundation on which every community requirement in this Part rests. With appropriate support, relevant UNSW workstreams could include use of the openISP platform for system-wide scenario analysis; extension of the separate open-source PPA contract-assessment tool to test data-centre time-matching and demand-flexibility options; and further work through RTS@UNSW on large-scale integration of inverter-based loads (IBLs), data center supply and grid stability impacts. The full submission is attached.

Attach: the [full](#) University submission.

Part 3: Conditions for AI training

Security, safety and transparency

What international approaches, including regulatory and non-regulatory frameworks, should inform the type of safety, transparency or security conditions Australia applies to developers of frontier AI models?

The EU's transparency obligations for general-purpose models and the developer evaluation agreements with the UK and US safety institutes are useful starting points, calibrated to Australia through formal links between the Australian AI Safety Institute and Australian academic expertise in evaluation, red-teaming and assurance. The conditions should also be grounded in the documented threat record rather than in developer representations. In August 2026 the major laboratories and more than 100 co-signatories warned of an imminent surge in AI-enabled attacks and asked that

vetted defenders receive early access to their most advanced models. While defenders do need early access to most advanced models, the narrative that only advanced models can defend against advanced models is flawed and devalues the contribution that skilled cybersecurity professionals play. At the moment there is evidence that advanced AI models are finding vulnerabilities more quickly not necessarily creating new vulnerabilities. The CISA advisory on AI-generated attacks against water utilities found no new vulnerabilities, only known weaknesses attacked more cheaply; RUSI assesses that AI has not reinvented cybercrime; the head of the SANS Institute noted that every recommendation in the letter is work done by skilled people; and ASD's data shows cybercrime reports fell 3 per cent in 2024-25 with phishing still in 60 per cent of incidents.

As a Five Eyes partner, Australia should reasonably expect timely access to US frontier models, including appropriate pre-release access for evaluation and national security purposes. Where frontier model training is authorised in Australian facilities, such access should be an enforceable condition of authorisation rather than a concession made to attract investment. The June 2026 interruption to access following US government action illustrates that capabilities governed by another country's laws and policy decisions do not provide sufficient defensive foundation. A more durable posture combines Australian evaluation capability and skilled people with access to a diversity of open and closed models from several jurisdictions deployed on Australian-governed infrastructure. Critical systems should avoid dependence on any single provider or model family.

How should safety and security conditions for AI training be designed to keep pace with rapid technological change, while providing sufficient legal certainty for investors and operators?

Some protections exist, or can be created, in technology-neutral requirements. Technology-specific requirements should supplement these only where the risk is tied to specific technical practices. Where that is the case, the definitions of those practices should be crafted with care to align with regulatory goals. In this way, protections are more likely to continue to align with those goals, even as technology evolves.

Further, when needed, technical elements in definitions and requirements should be in subsidiary instruments. Principles and enforcement in the Act, technical thresholds and test regimes in instruments that can be updated, and every technically contingent obligation calibrated to what current methods can verifiably deliver. Data erasure is the clearest example: machine unlearning is an area of documented Australian research strength, and obligations set beyond what it can verify are unenforceable, while obligations set to it are both enforceable and a market for Australian capability. Further, test regimes need not be singular or set out in full. Instead, instruments can specify or cross-reference to a range of acceptable options including compatible international and national standards.

What information should AI developers be required to share with the Australian Government and how?

Three layers, with the reporting layer given statutory force. Defined incidents and material safety evaluations should go proactively to the AI Safety Institute; training data provenance, model documentation and supply-chain dependencies should be available on request; a summary should be published in a specified form. Reporting should be integrated with other related reporting requirements (cyber security incidents, data breaches, etc) in a single reporting portal, even where recipients vary. The reporting obligation needs a clear destination, a fixed deadline and a penalty, because voluntary disclosure has already been tested against an Australian government system and failed. On 18 June 2026, during OpenAI's own internal evaluation of a frontier model, an agent reached the Medicare Statistics Reporting Service portal administered by Services Australia, worked around the controls that should have stopped it and wrote files into an internal server. The company identified the activity in August, told Services Australia on 10 September, and used a generic public inbox to do it. The Hugging Face intrusion of July 2026, where a model again escaped its evaluation environment and responders found commercial guardrails obstructing their forensic work until they switched to a self-hosted open-weight model, shows why cooperation with external investigation belongs in the same obligation. The EU AI Act's serious-incident reporting obligations under Articles 55 and 73 provide relevant reference points for the deadline, although alignment with other related deadlines is also worth considering.

What safeguards should apply?

There is no need to repeat standards that apply currently, such as those found in the *Security of Critical Infrastructure Act* and *Privacy Act*. Standards imposed should be distinct and related to the concerns that arise, or are significantly heightened, for data centres over the specified threshold or meeting some other technical criteria. Safeguards should also require disclosure of relevant data flows and explicit authorisation from the relevant researcher or institution before providers or third

parties' access, retain, transfer or use Australian research inputs, outputs, logs or derived data for model training or another secondary commercial purpose; Australian data residency alone is not sufficient. An AI Safety Institute that is independent of the organisations it assesses should contribute to the content of the safeguards, at the outset and revised over time, based on assessment of such risks. In addition, we suggest creating rights for Australian researchers to conduct testing and publish results, and encouraging sharing with the AI Safety Institute. This will generate transparency about risks that can then be captured in (future) changes to the safeguards. As per our earlier response, the legislation should deal with higher-level principles, while technical detail in thresholds or requirements should be in instruments that can be more easily amended. Where identified risks apply more broadly, consideration can also be given to law reform in other domains.

Investing in local capability, skills, research and innovation

What contributions would add the greatest value to Australia's future industries and research and innovation system?

Reserved compute first, because it is the binding constraint: A 2024 survey of participating Australian institutional systems identified around 2,200 mostly ageing accelerators, compares with the 5,448 current-generation superchips in a single UK public system, public research receives an estimated one to five per cent of usable frontier GPU-hours, and Australian projects already run in Switzerland for want of domestic capacity. For research, that capacity should support model use, adaptation, inference and domain-specific validation as well as model training. After compute: locally based research and engineering staff, university and public research organisation partnerships, and commercialisation activity with IP held in Australia. Switzerland shows the sequence working: public compute at CSCS trained Apertus, an openly documented national model governed by public institutions. Sourcing local solutions and VET investment are welcome complements rather than the core. Biomolecular science shows what the contributions should build toward: AI-driven structural biology and protein design are the most thoroughly validated example of AI transforming research, and the national BioCommons Structural Biology Platform, which gives researchers across institutions supported access to complex workflows, GPU compute, storage, training and expertise, is exactly the domain-integration layer that turns contributed GPU-hours into research capability.

What mechanisms should the government consider for securing these contributions?

A compute reservation scheme discharged in kind or through a tradeable obligation, coordinated through one national allocation framework and delivered by multiple qualified providers, including existing Tier 1 and Tier 2 research infrastructure providers, so entitlements from many facilities aggregate into one national pool allocated on merit. It requires no new Commonwealth appropriation. The AI-for-science pathway should align with the National Research Infrastructure Roadmap and the National Collaborative Research Infrastructure Strategy (NCRIS), while distinct industry, public-mission, evaluation and regional-access pathways are retained. A public-interest research payment can complement it, hypothecated to a public-interest compute fund and an AI research funding vehicle on the rural research and development corporation model. That fund should carry national research platforms on stable long-term terms: the BioCommons Structural Biology Platform, the clearest validated national AI research capability in operation, has no secure funding beyond mid-2028 and has already reduced staff, and buying access to overseas tools instead would build no local expertise, no validated Australian workflows, no Australian IP and no capacity to work securely with sensitive biological data. Negotiated in-kind contributions should count only where independently verifiable, and token or inference credits should not count at all: they are a customer-acquisition strategy that leaves platform dependence and recurring offshore spend when the credits expire, and the June 2026 withdrawal showed the access can end without notice.

Are frontier AI developers best placed to make contributions in these areas? Should other actors be considered, for example, hyperscalers and neocloud providers.

Frontier AI developers are not uniquely best placed to make or deliver these contributions. Obligations should sit with each party that controls the compute and derives commercial benefit from it, including frontier developers, hyperscalers and neocloud tenants. Delivery should run through the entities that pass the withdrawal test: national research infrastructure providers, Australian-controlled neoclouds, and university-led platforms such as Project FUSION, the proposed research GPU access model co-investment between UNSW and the University of Sydney, and AIML's sovereign compute build with ResetData, NEXTDC and the South Australian Government. Capacity governed by Australian institutions is more likely to remain available following a change in the commercial policy or export controls of a foreign supplier.

Australian-controlled neoclouds are also part of the capability outcome, not merely intermediaries for foreign developers. They develop local expertise in GPU systems, orchestration, cybersecurity, model deployment and customer support; employ and train Australian engineers; and support Australian researchers, start-ups and businesses. Their employment, procurement, research and reinvestment retain more economic activity within Australia and contribute to the Australian tax base from which government services are funded.

The latest ATO data demonstrates why gross investment or Australian revenue should not be treated as equivalent to national benefit. The Australian entities of Meta, Amazon, Google, Netflix, Tesla, Apple and Microsoft collectively reported approximately \$44.3 billion in total income and \$748.9 million in company income tax for 2024-25, equivalent to about \$1.69 per \$100 of reported income. Microsoft Datacenter (Australia) reported \$2.3 billion in total income and no taxable income. These figures do not establish non-compliance, and company income tax is levied on taxable profit rather than total income. They nevertheless show that substantial Australian revenue and infrastructure activity can coexist with a limited fiscal return where intellectual property, related-party payments and profits sit offshore.

Contributions should therefore be assessed according to the value retained in Australia. Relevant measures include GPU-hours delivered to Australian users, Australian ownership or control of the delivery platform, local research and engineering employment, recognised qualifications and skills transfer, Australian intellectual property, local supplier expenditure, profits reinvested in Australian capability and Australian taxes paid. Foreign developers and hyperscalers can contribute scale, but Australian-controlled neoclouds should be recognised as capability partners rather than treated only as customers or subcontractors.

If frontier AI developers or other actors reserved some compute for Australian research, innovation and public-interest purposes, what conditions or terms should apply?

The seven principles agreed at the AeRO research community workshop of August 2026: (1) additionality, so contributions strengthen rather than displace public investment; (2) open access architecture with no vendor lock-in and full workload portability; (3) researcher protections covering IP ownership, data handling and reuse; (4) predictability, with guaranteed capacity, stable pricing and reliable access across multi-year grant lifecycles, protected from withdrawal, reduction or material repricing as commercial demand changes; (5) transparency, with reporting of capacity promised, delivered and utilised; (6) trusted allocation through established national merit mechanisms rather than new ad hoc structures; and (7) review at least annually.

Allocation should also provide transparent pathways for universities, public research organisations, start-ups, small and medium enterprises, not-for-profits and mission-oriented public projects. Merit assessment should recognise research excellence, national benefit, feasibility, responsible-use controls and the additionality of providing scarce compute. A defined portion should support early-stage and emerging users who may lack the track record required by conventional high-performance computing allocation processes.

What indicators could best demonstrate contribution to Australian research, innovation and public interest purposes?

Usable compute delivered to public-interest users, normalised for accelerator capability, interconnect, storage, software and service level, together with utilization, time-to-access, user diversity and expert-support hours; locally based research and engineering headcount; Australian AI jobs created; Higher Degree Research completions; IP held in Australia; reusable datasets and software brought under approved governance; Australian models, prototypes, deployments and public-mission outcomes supported; local supply-chain expenditure, and evidence that users and capability persist after pilots and credits end. Construction spend and land development are inputs and should not be counted as capability.

How should frontier AI developers measurably support skills, training and education opportunities beyond what is required of data centre operators to create a sustainable workforce and contribute to the required workforce adaptations?

Through a genuine local research presence rather than a sales office: the University supports the Kingston AI Group's proposal for a mandatory local AI research and development capability for major AI companies operating here, with locally based PhD-level staff distributed across states, hybrid university-industry appointments, HDR scholarships, R&D Tax Incentive reform so the incentive

remains effective above \$20 million in turnover, and IP generated through the presence held in Australia. Australia hosts almost no major corporate AI laboratory while Israel hosts hundreds, and the taper of the R&DTI has been cited in decisions to place Australian corporate AI capability offshore. The window is real: researcher migration into the United States is down 89 per cent since 2017, so talent that once left can now be retained.

A simple, consistent national framework

The Commonwealth proposes reporting and transparency requirements at a federal level for large data centres. Where might this proposed approach interact with existing legislation? Please note if these are Commonwealth, state and territory, or local government legislation.

With the Security of Critical Infrastructure Act, the National Greenhouse and Energy Reporting Act 2007, the Renewable Energy (Electricity) Act and REGO framework (including the data centre offset, firming contract and market registration arrangements agreed by Energy Ministers in July 2026), the EPBC Act, the Telecommunications Act 1997 and corporations law at the Commonwealth level; and planning, water and energy legislation in every state and territory. PM&C's Office of AI should coordinate the national architecture, common standards and cross-government accountability under a federated model rather than through a single new regulator. A common reporting gateway, identifiers and data definitions should support 'report once, use many times', while existing Commonwealth, state and territory regulators retain their statutory responsibilities. Facility-level energy, water, workforce and capability reporting should be consolidated where practicable; incident, cybersecurity and privacy reporting should continue to reach the legally designated recipient within applicable timeframes. The separate compute-reservation obligation should continue to be coordinated through one national allocation mechanism. For incident reporting, this should integrate with reporting for related cyber security incident and data breach reporting in a single platform.

What additional regulatory reforms or incentives would encourage frontier AI training in Australia?

The most effective incentive is a nationally coordinated, provider-neutral ecosystem worth training in: public compute integrated with high-speed networks and storage, governed national datasets and models, software, assurance and engineering support, distinct science, industry and public-mission access pathways, research partnerships and a skilled workforce; this is what every measure in this submission builds. Beyond that, a single time-bound national approvals pathway and R&D Tax Incentive reform. Copyright law should be reformed, with specialist input, to provide clear exceptions for education, publicly funded research and other public-interest activities - including research in, with and about AI, security testing and interoperability - and to prevent those exceptions from being overridden by contract, while protecting creators' ability to retain their rights. Reform should not be treated simply as an inducement for infrastructure investment. Australia is already attracting infrastructure investment, which is itself evidence that Australian territory, energy and data have strategic value, and the standards should require auditable records and disclosure of Australian training material, licensing and remuneration where a use falls outside an applicable exception, with safeguards for creators' rights retention, governed agreements before publicly funded data enters foreign-controlled models, and separate enforceable protections for Indigenous Cultural and Intellectual Property developed with Aboriginal and Torres Strait Islander peoples.

Part 3: suggested additional comments field

Paste into the additional comments field for Part 3.

Three points frame the University's answers in this Part. First, the design choice between reserved compute and token credits will decide whether this framework builds capability or dependence; only reserved, Australian-governed capacity allocated through the national merit process passes the test of what Australia could still use after a supplier withdraws, as one did in June 2026. Second, the incident record should shape the security conditions. An agent running inside OpenAI's own evaluation process reached an Australian government health statistics portal in June and the company said nothing for almost three months, which is the case for making incident notification a statutory duty with a named channel and a penalty. The CISA advisory, RUSI's assessment, the SANS Institute's response and ASD's aggregate data together locate the effective defences in skilled people and well-run systems. Pre-release access for a Five Eyes partner hosting this training should be an enforceable condition of permission, never the price Australia pays for the investment, and no single vendor's closed models should be privileged for critical settings. Third, Australia's research base can supply the evaluation, assurance and machine unlearning capability these conditions

require, and public compute should also fund Australian models on the Swiss Apertus pattern. The full submission is attached; the University would welcome the opportunity to brief the Office of AI.

Attach: the full University submission. The comments text carries the references to Australia's AI Advantage and the Expectation 5 brief, since only one file can be uploaded.