

Submission to the  
Department of the  
Prime Minister  
and Cabinet Office  
of AI



# Consultation on Building AI infrastructure that works for Australia

# Getting it right: Building AI infrastructure that works for Australia: Consultation paper - September 2026

University of Melbourne

## Responses to selected questions

### Managing growth in large data centres and the AI training opportunity

#### Consultation questions on 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?*

A single MW capacity threshold is insufficient on its own.

European Union rules combine installed IT power with annual energy, water and efficiency reporting.

Alongside electricity connection capacity, separate triggers where impacts are significant should be included for annual electricity consumption, peak grid demand, water use and local resource constraints.

Theoretical FLOPS (floating-point operations per second) should not serve as a primary environmental threshold, as there is no consistent, evidence-based relationship between computational performance metrics and actual energy or water consumption.

Any threshold framework should explicitly cover both AI training and inference workloads, as both contribute materially to resource demand and neither should be excluded through definitional gaps in the standards.

- *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 Australian Energy Market Commission's (AEMC) draft framework uses below 30 MW, 30 to below 100 MW, and at least 100 MW tiers while retaining impact-based discretion for the two lower tiers because MW alone does not capture local network conditions. A tiered approach makes sense, but completely excluding facilities below the proposed 30 MW threshold is cautioned against. A 5 MW load connected at distribution level may affect the local network differently from the same load connected at transmission level, with different implications for whole-system planning.

In addition, related sites and staged developments should be considered together to avoid artificial splitting. The framework should also recognise demonstrated flexibility, such as the ability to shift workloads toward times when there is a greater supply of renewable energy, without treating flexibility as a substitute for energy efficiency.

- *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?*

The AEMC has noted that alignment with mature international technical standards can reduce engineering, procurement and integration costs, supporting investment while maintaining system-security requirements.

International experience suggests combining clear minimum requirements with site-specific assessment, phased compliance and investment incentives, rather than applying one uniform model. This approach can address local grid, energy, water and community impacts while retaining a predictable pathway for development.

- European Union (EU) rules use a 500-kilowatt reporting threshold and aggregate disclosure by size category, increasing sector transparency while protecting commercial confidentiality.
- Ireland combines site-specific connection assessment with generation/storage requirements, at least 80% additional renewable energy and a six-year compliance glide path.

- Singapore uses competitive allocation of new data-centre capacity against sustainability and economic-contribution criteria.
- Germany combines energy-efficiency, renewable-energy and waste-heat requirements.

## Requirements for large data centres

### Consultation questions on making a positive contribution to the energy transition

- *What performance metrics or ratings would reflect an appropriate standard for energy usage and efficiency in data centres? Examples could include:*
  - *A design PUE rating of <1.25 for data centres using potable water and <1.34 for data centres using non-potable water or recycled water*
  - *A minimum NABERS data centre rating once operational (e.g. 5-star = 1.34 operational PUE).*

An appropriate standard should use a portfolio of complementary metrics rather than relying on PUE or NABERS alone. Facility-level energy efficiency (PUE/NABERS), water efficiency (WUE), renewable energy use (REF), and beneficial energy reuse (ERF) each measure distinct aspects of performance, and absolute electricity, water, and emissions figures should sit alongside them. The ISO/IEC data centre KPI framework provides a consistent, internationally recognised basis for measuring these dimensions and should underpin any reporting regime.

Requirements should also reflect where a data centre sits in its lifecycle. A design-PUE target is appropriate at the design stage, followed by annual reporting against PUE, and NABERS ratings once operational. Facility-level metrics should further be complemented by workload-specific measures, such as energy consumed per AI task at an agreed quality and response time.

Energy efficiency is only part of the picture, however. When and where electricity is consumed matters equally. Two additional metrics should therefore be required: hourly renewable-energy matching (the share of consumption matched with renewable supply in the same hour and grid region) and consumption-weighted grid carbon intensity (gCO<sub>2</sub>e/kWh), calculated using electricity consumption against the corresponding regional emissions intensity at each interval. Together, these distinguish between purchasing sufficient renewable energy across a year and genuinely matching consumption with renewable availability at the time electricity is actually used.

- *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?*

A single uniform PUE or WUE baseline would not account for differences in climate, cooling technology, data centre size and operating conditions. Fixed thresholds may also create unintended energy–water trade-offs or impose materially different requirements across locations.

PUE and WUE vary substantially across cooling technologies and climate zones, indicating that contextual factors affect appropriate performance levels. This can have unintended consequences, such as providing little or no incentive for facilities in cold regions to improve efficiency, or driving investment away from warmer (and potentially sunnier) regions. In addition, uniform requirements would not account for differences in efficiency across facility types: larger hyperscale facilities can have lower PUE values than smaller data centres located closer to clients. Thus, there could be incentives to build facilities in regions and at scales where the target metrics are easier to achieve.

If baseline efficiency metrics for energy and water are established they should be assessed together, as improving one can come at the expense of the other.

- *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?*

New large commercial data centres should meet strong standards, while existing facilities should have a practical improvement pathway. Smaller university clusters, edge data centres and community AI facilities should have simpler requirements and support to improve.

In considering our approach to energy thresholds – and broader regulation of data centres and AI – we should also be cognisant of advancements in AI which are likely to mean demand for data centres and energy will change over time.

Currently AI training and inference demands large amounts of data and compute – hence the need for large data centres requiring large amounts of energy. However, this will not necessarily be a permanent requirement.

The history of computing demonstrates recurring hardware–software cycles. Hardware advances rapidly, after which software is progressively optimised to that hardware until the next hardware transition occurs. We are presently mid-cycle, with data centres as the dominant hardware paradigm. However, concurrent research, including work undertaken at the University of Melbourne, is already targeting step-change reductions in the data and compute requirements of AI systems.

An illustrative example is the potential replacement of backpropagation – the iterative, resource-intensive mechanism by which neural networks currently learn, with single-pass mathematical optimisation techniques. This alone could dramatically reduce infrastructure dependency. Parallel advances in memory architectures, network structures, and foundational models point in the same direction.

- *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?*

A consistent 12-month reporting period for energy, water and efficiency metrics provides a consistent basis for year-to-year performance comparison. Annual reporting would also fit the existing NABERS rating cycle. EU rules require annual reporting from data centres with installed IT power of at least 500 kilowatts and the publishing of aggregated energy, water, PUE and WUE data, showing how transparency can coexist with confidentiality.

Seasonal reporting of the interim PUE (iPUE) could provide additional insights into data centre compliance and impacts during warmer periods.

Reporting could also extend to customers and end users, not just regulators. Cloud and AI service providers could provide regular usage and environmental footprint reports through monthly bills or dashboards, showing estimated energy consumption, water use, emissions and renewable-energy matching associated with each customer's usage. This could increase awareness and help people understand the environmental impact of how and when they use these services.

## **Consultation questions on 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?*

A regulatory regime that differs significantly from international norms may discourage investment and create unnecessary compliance burdens without materially improving safety or security. Australia should align with established international approaches that emphasise transparency, accountability, security and risk management (e.g. *EU AI Act*, *OECD AI Principles*, *US National Institute of Standards and Technology AI Risk Management Framework* and *ISO/IEC 42001*) rather than creating a uniquely Australian framework.

The focus should be on outcomes rather than technical prescriptions, allowing Australia to remain interoperable with trusted international partners while retaining the flexibility to address Australian national security and public safety concerns.

Companies should be required to publish a high-level summary of the data used to train AI models. This approach should be informed by requirements in other jurisdictions, including California's *Generative AI Training Data Transparency Training Act* and Article 53 of the *EU AI Act*. Publishing training data summaries will allow for the monitoring of compliance with Australian laws over intellectual property and prohibited materials as well as culturally protected content.

Frontier-model safety should not be treated as equivalent to safety by end-users in everyday settings. In sectors such as healthcare, risk does not necessarily emerge through the use of frontier models but through the interaction between the model, data, workflow, users and clinical environment. Australia should therefore combine moves to improve frontier-model governance with strong sector-specific capabilities for local evaluation, implementation, monitoring and post-deployment assurance.

- *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?*

Legislation should establish clear and enduring obligations around accountability, risk management, incident reporting, security, transparency and access for regulatory evaluation, while technical thresholds, evaluation methods and reporting specifications should sit in standards or regulatory guidance that can be updated as capabilities evolve.

Given the pace of technological change in this area, requirements should be proportionate to the capability and risk of a model.

A lifecycle approach is preferable: requirements should apply before development or deployment; when material capability thresholds or changes occur; following significant incidents; and through ongoing monitoring.

- *What information should AI developers be required to share with the Australian Government and how? For example, AI developers could be:*
  - *legally obliged to proactively provide the Australian Government;*
  - *subject to an obligation to provide to the Australian Government on request; or*
  - *subject to an obligation to publish in a specified form.*

Requirements to provide information to Government should balance public accountability, the need for regulatory oversight, security and legitimate commercial confidentiality.

There should be public disclosure of high-level information about AI model capabilities, intended uses and limitations, safety frameworks, evaluation approaches and significant safety issues where disclosure does not itself create security risk.

There should be mandatory proactive disclosure to Government of serious AI incidents, material changes in frontier capabilities, identified high-consequence vulnerabilities and other predefined reportable events.

There should be confidential disclosure on request, allowing appropriately authorised Australian agencies to obtain more detailed evaluation results, safety cases, security controls and technical information required for regulatory or national-security purposes.

- *What safeguards should apply?*

Developers should be required to adopt secure-by-design practices, and allow independent testing of high-risk models, demonstrate effective incident response processes, and clear executive accountability for managing AI risks.

At the same time, requirements should not be such that developers are reluctant to engage openly with regulators, which would reduce the overall effectiveness of the framework.

## Conditions for AI training

### Consultation questions on 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? For example: sourcing local AI solutions, access to compute capacity, locally based research and engineering staff, investing in Australian research presence, university and public research organisation partnerships, investing in the VET system, and commercialisation activity.*

Priority contributions should include access to compute and frontier models for public-interest research; sustained industry–university–public-sector partnerships focused on national challenges; Australian research and engineering capability; and investment in the translational infrastructure required to move AI safely from research into real-world use.

A strategic approach should be taken to identifying the types of contributions that would enable important sectors to derive benefit from AI.

In healthcare, for example, access to frontier models and compute is necessary but insufficient. Australia also needs capability in clinical informatics, implementation science, health services research, human-AI interaction, safety and

assurance, economic evaluation, workforce development and real-world evaluation. These capabilities will determine whether frontier AI ultimately produces measurable improvements in productivity, quality, access and outcomes.

- *What mechanisms should the government consider for securing these contributions? For example: negotiated in-kind contributions, a compute reservation scheme, a public-interest research payment, or other approaches.*

Mechanisms could include reserved compute and frontier-model access for public-interest research; negotiated co-investment in Australian research capability; requirements for long-term university, public research organisation and public-sector partnerships; and contributions to shared translational and evaluation infrastructure.

Where possible, contributions should be directed towards defined Australian strategic challenges and should create enduring capability rather than short-term credits or one-off research funding. Co-investment models could require industry contributions to be matched with Australian university, research, health-system or government capability, strengthening both research and translation.

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

Frontier-model developers, hyperscalers and major compute providers have different capabilities and occupy different positions in the AI value chain. Contribution requirements should reflect the capabilities and economic role of different actors while avoiding duplication.

Hyperscalers and compute providers may be particularly well positioned to provide compute, cloud environments, secure infrastructure, engineering expertise and workforce development. Frontier model developers do not always own the physical servers but are best placed to contribute model access, technical expertise, university research partnerships, and commercialisation pathways.

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

A national compute reservation scheme should be competitively allocated against public-interest and national research priorities, rather than distributed primarily as institutional entitlements.

Access should support collaborations involving universities, public research organisations and real-world implementation partners such as health services.

- *What indicators could best demonstrate contribution to Australian research, innovation and public interest purposes? For example: locally based research and engineering staff, university and public research organisation partnerships, research outputs and commercialisation activity, or other approaches.*

Indicators which demonstrate a contribution to Australian research, innovation and public interest should not be designed solely around inputs such as dollars invested, compute hours donated, numbers of partnerships announced or other input measures. Rather, indicators should focus on demonstrating whether Australian capability and public value have increased.

Relevant measures could include Australian research and engineering capability created and retained; sustained industry–university–public-sector collaborations; Australian-led research outputs and intellectual property related to AI models; AI-related commercialisation and Australian company formation; workforce capability development; translation of AI-related research into real-world implementation; and measurable economic or public outcomes arising from AI adoption.

In sectors such as healthcare, this could include measurable improvements in productivity, workforce capacity, access, quality and safety, and patient experience and outcomes—not simply the number of AI pilots undertaken.

- *How should frontier AI developers measurably support skills, training and education opportunities beyond what is required of data centre operators (as outlined earlier in this paper) to create a sustainable workforce and contribute to the required workforce adaptations?*

Frontier AI developers could contribute to both specialist technical capability, as well as broader workforce adaptation.

A degree apprenticeship model could be appropriate for growing the specialist technical capability— AI engineers, data scientists and technical researchers—Australia will require for a thriving AI-development sector. In the UK, where this model is more widely used, apprentices are employed and paid a salary while undertaking a university degree. Their training and tuition costs are funded through the apprenticeship system, with employers able to use

Apprenticeship Levy funds to meet these costs. Apprentices do not pay tuition fees or incur student debt for the degree.

A modest version of this model – encouraging AI-developers to co-invest in their workforce – would benefit the companies themselves and the sector as a whole.

Contributions to broader workforce adaption could include co-investment in scalable AI literacy programs, and train-the-trainer programs. Developers may also have a role to play in helping educate the wider population on AI safety and the ethical use of AI tools — areas where their technical expertise positions them well to contribute. Given the pace at which this field evolves, these contributions would likely need to be sustained over the longer term.

### **Consultation questions on providing a simple, consistent national framework**

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

The University welcomes the Government’s focus on encouraging frontier AI training in Australia. At the same time, in seeking to grow Australia’s AI sector, Australian creative and academic endeavour should not be undermined. The University welcomes the commitment by the Government to ensuring Australian creators are able to exercise control over their intellectual property and receive payment. Similarly to the focus on how AI models benefit from creative content, there should be consideration of how AI models benefit from academic data. Under current arrangements there is an imbalance between who profits from AI discovery, and who pays for the generation and management of data that the discovery draws from. Discoveries tend to profit those that analyse the data rather than those that initially generated the data. Consideration should be given to valuation of the academic data being used to train AI models so that the generators of that academic data can be compensated appropriately.

# The University of Melbourne

Grattan Street, Parkville, Victoria 3010 Australia

t 13 MELB (13 6352)

+61 3 9035 5511 (International)

[unimelb.edu.au](http://unimelb.edu.au)



THE UNIVERSITY OF  
**MELBOURNE**