



**Actuaries
Institute.**

9 October 2026

Office of AI
Department of the Prime Minister and Cabinet

Via: [Online Submission Form](#)

Getting it right: Building AI infrastructure that works for Australia – Submission to Consultation Paper

The Actuaries Institute ('the Institute') welcomes the opportunity to make a submission on this consultation.

The Institute is the peak professional body for actuaries in Australia. Our members work in a wide range of fields including insurance, superannuation, investments and retirement incomes, banking, enterprise risk management, data science and AI, climate change impacts and government services. The Institute has a longstanding commitment to contribute to public policy discussions where our members have relevant expertise. The comments made in this submission are guided by the Institute's '[Public Policy Principles](#)' that any policy measures or changes should promote public wellbeing, consider potential impacts on equity, be evidence-based, and support effectively regulated systems.

This submission also draws on the Institute's recent thought leadership publications and submissions to consultations, particularly:

- [AI Risk Management in the Financial Services Sector](#), a practical guidance resource produced jointly with the UTS Human Technology Institute, recognising the critical role that financial services play in the economy which can create new and amplified risks of AI (September 2026);
- [Where Are Australia's Data Science Leaders? The Case for Technical Career Pathways](#) (October 2025);
- [Building Tomorrow: Preparing Australia for the Age of AI](#) (May 2025);
- [Response to Department of Industry, Science and Resources Consultation on Mandatory Guardrails for Safe and Responsible AI](#) (October 2024); and
- [Response to Department of Industry, Science and Resources Consultation on Safe and Responsible AI in Australia](#) (July 2023).

Overall comments

We strongly support the Australian Government's recognition that AI brings opportunities but must be used safely and responsibly. Our members, often referred to as the 'original data scientists', have been leveraging data, analytics and increasingly AI to drive productivity and innovation across the economy in a safe and sustainable manner.

The Institute supports the Government's aim of attracting investment in AI infrastructure, data centres and training on terms that deliver lasting benefits for Australians. The Institute considers

that development of AI infrastructure should not materially reduce the affordability, reliability or availability of electricity and water services for existing households, businesses and essential services. Assessment processes should explicitly consider whether costs, service impacts or resource constraints are being transferred to existing users, and proponents should demonstrate how these impacts will be mitigated.

Data centres and the supporting infrastructure are long-lived assets. Decisions made now will shape costs, risks and benefits for decades. Therefore they should be assessed over their expected lifetimes, allowing for uncertainty in demand, technology, climate conditions and infrastructure delivery.

The attachment sets out our responses to the consultation paper questions. Our main comments include:

- We strongly encourage a focus on regulating outcomes, but recognise that certain technologies may require mandatory processes to demonstrate that those outcomes are likely to be achieved.
- We encourage the Australian Government to clarify how existing liability rules apply to actions conducted by AI developers' AI agents and determine which agencies should be involved in enforcement. Further, we strongly support timely disclosure of incidents, for example within 72 hours of discovery by AI developers, consistent with APRA requirements on financial institutions for managing cyber security incidents.
- Access to local compute and education are two areas that are important to focus on enabling, in addition to those set out for feedback in the consultation paper, to support Australian businesses and people being able to benefit strongly from AI. We strongly encourage the Australian Government to consider requirements for frontier AI developers and data centre operators to reserve a minimum proportion of data centre computing capacity in Australia for inference to ensure Australian organisations and individuals can access leading AI models at internationally competitive prices.
- Compliance with individual requirements, such as renewable certificate obligations or payment of connection costs, should not be treated as sufficient evidence that overall consumer and environmental impacts have been addressed. Overall impacts on system outcomes in a range of scenarios and stress conditions should be explicitly considered as part of connection approvals.

We would welcome the opportunity to discuss any aspect of this submission in detail. If you would like to do so, please contact us via public_policy@actuaries.asn.au or (02) 9239 6100.

Yours sincerely,

(Signed) Elayne Grace
Chief Executive Officer

Attachment: Responses to consultation paper questions

Approach to the application of AI standards

Question: **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?**

The Institute supports the consultation's consideration of staged connections and transitional arrangements. We do not support either a blanket prohibition on early operation or automatic permission based on a single outcome, such as a Final Investment Decision.

Planning, approvals, workforce, supply chain and construction constraints can delay supporting generation, storage and network infrastructure. Early operation should be permitted where assessment demonstrates that the initial load can be accommodated while meeting defined consumer protection, reliability, system security and emissions requirements.

Assessment for a staged connection should consider financing, approvals and construction milestones, and credible contingencies for delay.

Connection arrangements should establish:

- the initial permitted load and conditions for increasing it;
- delivery milestones for supporting infrastructure;
- enforceable responses to delay, including demand limits where necessary; and
- responsibility for additional costs and risks during the transition.

Where these conditions cannot be met, demand should be staged or connection deferred. Governments should coordinate approvals and provide clear decision timeframes, while preventing speculative capacity reservations and repeated extensions that undermine the requirements.

Requirements for large data centres

Additional comments

The Institute supports the proposed direction of requiring large data centres to contribute additional renewable supply, provide demand flexibility and bear their fair share of network costs.

Australia's renewable resources provide an opportunity to attract data centre investment that also supports the energy transition. Long-term electricity procurement can underpin new generation and storage, improve financing certainty, reduce costs, and make better use of infrastructure.

That said, these benefits should be demonstrated across the system, rather than assumed.

The Institute recommends:

- clear additionality requirements that support new or expanded renewable generation, with appropriate firming and network capacity, and prevent double counting;
- flexibility to meet obligations through retailers, pooled procurement and shared infrastructure;
- demand flexibility requirements that reflect operational capabilities and service obligations, with commitments that can be measured and relied upon during system stress;
- coordinated planning of generation, storage, networks and major new loads, including the cumulative demands on skilled labour, equipment and construction capacity; and
- location decisions that consider delivered energy costs, network capacity, water security, telecommunications, climate exposure and community impacts together.

The Institute does not support a blanket requirement for each data centre to match its consumption with new renewable supply at every point in time. Such a requirement could overlook the benefits of an interconnected system and impose unnecessary costs. Equally, matching annual electricity consumption with renewable certificates is insufficient to establish that reliability, affordability and emissions impacts have been addressed.

The appropriate test is whether additional supply, firming and flexible demand adequately address the development's effects on the wider system, including during scarcity periods. Any time-matching requirements should be targeted to demonstrated system needs.

The Institute would not support treating either payment of connection costs or projected national economic benefits as sufficient evidence that affected consumers and communities are protected. The assessment method, assumptions, uncertainty, and treatment of residual risks should be transparent.

How data centres can use innovative, efficient and sustainable solutions to minimise water use

Additional comments

The Institute supports the proposed requirements to minimise water use, improve efficiency, report performance transparently and build resilience to drought and climate change.

Water requirements should reflect local conditions and cumulative demand within the relevant catchment or supply system. National averages can conceal material local pressures.

Recycled or other suitable non-potable water should be the default for water-intensive cooling. Any proposed use of potable water should require justification to the relevant authority, demonstrating that alternative sources, cooling designs and site options have been assessed. Interim potable use should be subject to agreed limits, drought arrangements and a credible transition plan.

Water availability and infrastructure requirements should inform site selection from the outset. Proponents should assess proximity to wastewater treatment plants and recycled-water networks alongside electricity, telecommunications and other location requirements. Developers should fund dedicated water infrastructure and their attributable share of shared upgrades, without transferring those costs to existing customers.

Energy and water efficiency should be assessed together. Cooling choices can trade lower water consumption against higher electricity use, or vice versa. The Institute would not support applying either measure in isolation where this produces a worse overall outcome.

Climate assessment should extend across infrastructure lifetimes and include extreme heat, drought, flood and bushfire. It should test concurrent stresses, particularly when cooling demand rises while electricity and water supplies are constrained. Operating plans should explain how these conditions will be managed without transferring unacceptable burdens to other users.

The Institute supports transparent reporting using consistent definitions and existing frameworks wherever practicable.

Reporting should cover actual energy and water consumption, sources, efficiency, peak demand, energy sourcing and emissions, together with progress against infrastructure and flexibility commitments. Renewable procurement claims should be distinguished from physical emissions impacts, and reporting should include backup and interim fossil-fuel generation.

Efficiency improvements should be reported alongside absolute consumption. Better efficiency does not necessarily mean lower total resource use or emissions.

Requirements should identify the accountable entity, provide proportionate verification and specify consequences for material non-compliance. This would allow governments and communities to assess whether promised benefits and protections are being delivered.

Conditions for AI training

Security, safety and transparency

Question: 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?

We recommend the Australian Government clarify how existing liability rules apply to actions conducted by AI developers' AI agents. We acknowledge this is an untested area of law and it may be difficult to quantify or qualify appropriate redress. Regardless, enforcing liability is necessary to incentivise AI developers to implement controls that minimise the likelihood and severity of incidents.

Law enforcement needs sufficient support to address the fallout from incidents, including access to experts. As the volume of incidents grows, it may be difficult to manage these purely at a Federal level. In the short term we anticipate these events should be managed by Federal agencies in collaboration with State level agencies (e.g. the AFP's Joint Policing Cybercrime Coordination Centre). The collaboration is necessary so that knowledge of how to manage these incidents can diffuse into the States, enabling them to manage incidents independently.

Frameworks for agentic risk are nascent. The 'Model AI Governance Framework for Agentic AI' published by Singapore's IMDA (<https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/mgf-for-agentic-ai.pdf>) and the vendor-operated AIUC-1 Standard (aiuc-1.com) are both focused on AI deployers rather than AI developers. However, these examples illustrate the direction that future frameworks may take, including broad considerations that may be relevant for AI developers to address.

Question: 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?

We strongly encourage a focus on regulating outcomes, but recognise that certain technologies may require mandatory processes to demonstrate that those outcomes are likely to be achieved. Refer further comments later in this submission, including with regard to mandatory notification requirements for incidents, and the Australian AI Safety Institute and safeguards.

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

During 2026 several AI developers disclosed cyber incidents conducted by their AI agents. When AI developers discover these incidents (or similar harms) have occurred, they should be legally obliged to advise both the Australian Government and the affected organisation(s) within 72 hours of discovery, consistent with APRA requirements on financial institutions for managing cyber risk. The rapid turnaround is necessary because containment effort may be required. We recommend the Australian AI Safety Institute document a taxonomy of harms that require disclosure.

We also recommend that AI developers grant the Australian AI Safety Institute access to new models ahead of public release so it can conduct independent evaluations, as well as unredacted findings from other independent evaluators.

Question: What safeguards should apply?

We recommend AI developers grant the Australian AI Safety Institute access to independently audit the safety controls applied and activities executed by AI developers, especially in the context of research and testing that may have a higher likelihood of resulting in adverse outcomes. Rather than prescribing specific safeguards, the Government should clarify the categories of harm that are considered relevant and material to Australians. This would help focus mandatory independent evaluations and auditing activities on the risks that matter most, while creating clear incentives for AI developers to design appropriate controls. Relevant areas for assessment could include cyber capabilities, biosecurity capabilities, deceptive or manipulative behaviour, autonomous operation, and capabilities that may materially accelerate AI research and development. Assessments should consider both the likelihood and potential severity of harm, including its scale, reversibility, persistence and distribution, as well as the potential for cascading or systemic risks. The specific safeguards applied should then be proportionate to the demonstrated level of risk

Investment in local capability, skills, research and innovation

Question: 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.

We agree that attracting frontier AI training to Australia can deliver substantial benefits through local R&D capability, skills development, research partnerships, innovation ecosystems and commercialisation activity. In addition, we consider broad access to frontier AI inference capability to be an important public benefit in its own right. While only a relatively small number of organisations will train frontier models, access to inference enables the benefits of AI to diffuse across the broader economy, including SMEs, start-ups, universities, not-for-profits and government agencies. This is particularly important where data sovereignty, security or regulatory considerations limit the use of offshore infrastructure.

As signposted in the Dialogue Paper [Building Tomorrow: Preparing Australia for the Age of AI](#), access to local compute, and education, are two areas that will yield high value for Australia.

Locally hosted inference is a significant bottleneck to the widespread diffusion of AI across Australia, particularly for organisations with desires or obligations related to data sovereignty (as noted above). The limited capacity restricts the choice of models available, their serving mode (e.g. pay-as-you-go vs provisioned compute), and their geographic availability. These restrictions may be prohibitive for small and medium enterprises (SMEs).

However, capability must also be developed so Australians can make use of the compute. This could include investment in local talent via the establishment of offices in Australia, training programs like internships and apprenticeships as one strategy to reduce talent shortages, and free or subsidised AI subscriptions for university and school students to develop AI literacy from a young age.

Question: **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.**

We recommend the Australian Government require frontier AI developers and data centre operators to reserve a minimum proportion of data centre computing capacity in Australia to host current frontier models for inference, with pricing aligned to global market rates, so that they can be used by Australian individuals and organisations. As noted in response to the question above, access to AI models for inference is the mechanism that allows AI benefits to diffuse across Australian society.

Question: **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 the leaders in models and product capability. However, they have competing compute commitments between training, safety, internal uses, and serving inference to customers. Therefore, the application of inference guarantees is most relevant for AI developers to ensure Australian customers have access.

Both AI developers and hyperscalers are well placed to support education and training. Hyperscalers are likely to have broader geographic reach and talent capacity which enables greater scale.

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

Access should not be limited to research institutions. Compute should be made available at sufficient scale so that it can diffuse across broad segments of the Australian economy, rather than a small set of projects. That includes (but is not limited to) enterprises both large and small, start-ups, universities, not-for-profits, and government.

Question: **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.**

We support the consultation paper's emphasis on local R&D capability, skills development, research partnerships and commercialisation activity. In addition, broad access to frontier AI inference capability can help these benefits diffuse across the economy by enabling Australian businesses, start-ups, researchers, not-for-profits and government agencies to make practical use of frontier AI systems, particularly where data sovereignty or regulatory considerations limit the use of offshore infrastructure.

We therefore suggest Frontier AI developers should build locally based research hubs, not just sales functions and forward-deployed engineers. Local R&D will foster an ecosystem of talent, including training and research collaborations with other organisations. Demonstrated outputs include publications, technical reports, safety research, and novel products or services.

Question: 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 form) to create a sustainable workforce and contribute to the required workforce adaptations?

Employing staff in activities such as model training, safety (and broader) research, and product development will create incentives for top local talent to seek employment in Australia rather than going overseas for these opportunities. Over time as individuals move into different roles and organisations, their knowledge will diffuse into the broader Australian economy.

At the same time, this requires organisations to develop culture and practices for technical leadership to thrive, as noted in the Dialogue Paper [Where Are Australia's Data Science Leaders?](#) Frontier AI developers can contribute to diffusing this culture and mindset.

Creating a simple and consistent national framework

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

Although the Institute does not have specific expertise in copyright law, equity is a core principle of actuarial work particularly when competing stakeholder interests need to be balanced (such as for insurance or insurance like schemes). We consider it critically important that the Australian government resolve a fair and workable system for protecting and compensating creators when Australian content is used to train AI models.