

AI Training and Infrastructure



Submission

Date: Oct 9, 2026

Attention: Department of the Prime Minister and Cabinet

Beyond Zero Emissions welcomes the opportunity to provide a submission on the future of AI training and infrastructure in Australia.

[Beyond Zero Emissions](#) (BZE) is an independent think tank creating solutions for a prosperous zero-emissions Australia. We work with industry and government to develop and drive zero-emissions economic opportunities that deliver community benefits and long-term prosperity for Australia.

This submission draws upon our extensive body of research, including the development of our [National Action Plan](#) (NAP), a place-based, multi-sector approach to accelerate and track the development of Renewable Energy Industrial Precincts (REIPs). The NAP provides a snapshot of the readiness of industrial regions to host decarbonised industry across five themes: infrastructure, industrial capability, policy, social acceptance and finance. We have completed NAP assessments in the [Hunter](#), [Gladstone](#), [Kwinana and Port Hedland](#) – all of which underscore the importance of sequencing and coordinating major infrastructure alongside targeted forward planning – insights that apply equally to the development of data centres in Australia. Currently our work continues in Far North Queensland, focusing on the high potential areas of Townsville and Mount Isa.

BZE supports the federal government's goal of enforceable, uniform standards for data centre development across workforce, renewable energy, community outcomes and water use.

We regard the proposed Renewable Electricity Guarantee of Origin (REGO) liability as an effective way to drive investment in additional renewable energy, but only if surrender is restricted to certificates from new-build projects, so the market is not diluted. Allowing certificates from existing renewable generation would enable data centres to meet the obligation without necessarily causing any new renewable capacity to be built. Certificates alone, however, provide no assurance of the contracted revenue renewable developers need to reach financial close. The REGO liability should therefore be matched by long-term power purchase agreements (PPAs) with new generators.

We also encourage the government to treat data centre construction over the next decade as an opportunity to invest in Australia's emerging green economy, including the green aluminium, steel and low-carbon cement value chains. Programs such as the Hydrogen Production Tax Incentive and the Green Aluminium Production Credit provide supply-side support for these industries, and the Future Made in Australia Innovation Fund is helping to overcome technical barriers. What is missing is demand certainty and foundational offtake agreements. Hyperscalers such as [Microsoft](#) and [Google](#) have

already shown a willingness to support green commodity markets in Europe. Australia should harness developer interest to do the same here.

More detailed responses to relevant consultation questions are outlined below.

Yours Sincerely,

Matt McKee

Chief Researcher

Beyond Zero Emissions

Questions

Consultation questions on capacity thresholds for application of the AI standards

Q1: 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?

Q2: Does a tiered approach sufficiently capture different types of facilities and ensure proportionate obligations relating to Australia's energy transition, water management, and communities?

Q1-Q2: A tiered approach to data centre development is essential to their orderly development. Capacity is a useful overarching proxy for facility size, and should form the base of the tier system. However, given the goal of these standards is to maximise benefits while mitigating impact, several other factors will be relevant for an effective tiered approach. These include

- **Network connection type:** Operating data centres have considerably smaller capacity than the ones currently being proposed, and are connected to the grid at the distribution level. The AI-oriented facilities now progressing through the development pipeline are considerably larger, and as a result, are increasingly planning to connect directly to the transmission network.

Though both distribution and transmission network costs are passed on to consumers, distribution networks have a much larger impact on household bills, contributing about 33% of the total bill – about the same as wholesale costs. Transmission, on the other hand, only contributes 6% to this total ([AEMC, 2025](#)).

This means that smaller, distribution-connected data centres may have a greater impact on household electricity bills than larger data centres connected directly to the transmission network. The tiered system should therefore consider where a data centre connects to the grid, rather than relying on size alone.

- **Facility type/intended workload:** Data centres come in all shapes and sizes. Facilities handling AI training and inference workloads should have different standards than those hosting critical digital services often invoked by the data centre industry.

These workloads have a range of tolerances for flexibility and latency, which should be reflected in the standards. Beyond this, we note that much of the community opposition to data centres has arisen because they are seen as the physical manifestation of artificial intelligence and its perceived impact on jobs and safety. This comes with a higher burden to demonstrate tangible benefits compared with data centres which enabled valuable social services.

Q3: 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?

BZE has not responded to this question as it falls outside the scope of our expertise.

Consultation questions on the approach to the application of the AI standards

Q4: At what point in development processes should mandatory requirements apply for large data centres?

The mandatory requirements for data centres cover a range of different aspects of facility development and have different natural implementation points. However, as all require thorough preparation, meeting these standards should be embedded in the early stages of project development.

For example, surrendering REGOs from ongoing electricity consumption will only ramp up once a facility becomes operational. However, if – as we recommend – this is reinforced with a mandatory requirement for signing PPAs with new renewable energy generators, this would have to be done quickly after a data centre reaches final investment decision to not exacerbate the inevitable lag between new demand and generation.

On the other hand, the use of closed-loop cooling systems requires the engineering design of a facility to be completed with reference to mandatory requirements, and skills and training commitments would apply from construction all the way through operation. This means there should not be a single point at which all mandatory requirements begin. Instead, each requirement should apply at the earliest stage of development needed to ensure it can be effectively delivered.

Q5: What can be learned from approaches taken in other jurisdictions to ensure very large data centres currently in development pipelines are appropriately captured?

BZE has not responded to this question as it falls outside the scope of our expertise.

Q6: 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?

The overarching goal of these standards should be to maximise benefits while minimising impacts. The highest level of compliance would see these impacts mitigated throughout construction and from day one of operation. If, as in the case of new generation lagging demand, this cannot be achieved, the expectation should be for data centres to address this by overcompensating and surpassing actual demand. For example, new data centres will temporarily be unable to meet the REGO liability while additional generation is constructed. To compensate for this, their liability should subsequently surpass their actual energy consumption.

Q7: 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 core tension in requiring data centres to be powered by additional renewable energy comes from the mismatch in their construction times. Industry estimates indicate data centre halls can be built in roughly 18 months, whereas solar and wind farms, even when shovel-ready, take approximately two and three years respectively to complete ([AEMO 2025](#), [Net Zero Australia Topic Report #2 2025](#)).

This creates a sequencing and coordination problem. Concerns that data centres will drive up grid emissions are, in the short term, well founded. Until new renewable generation comes online, new data centre demand will largely be met by the existing generation fleet, including coal and gas. This will temporarily raise electricity prices and slow grid decarbonisation. Our National Action Plan assessments routinely identify similar sequencing issues in industrial precincts, where manufacturers seeking to decarbonise cannot access the renewable energy they need because of limitations in existing infrastructure and market structures. Governments must take a stronger role in aligning energy system planning with industrial strategy.

For data centres specifically, there are two broad options for approaching this issue. Governments can delay data centre construction so that completion aligns with corresponding renewable energy projects. Alternatively, they can allow data centre development to proceed and define a mechanism that ensures data centres compensate for their initial lack of additionality.

Delaying data centres. Making data centre operation contingent on the construction of the underlying renewable energy infrastructure would avoid any additionality problem. Such an arrangement is already a necessity for off-grid data centres. However, industry analysis suggests that the most important investment factor for data centres is speed-to-power: the time it takes to connect the initial data centre hall. We acknowledge that a temporary pause on construction would likely reduce companies' willingness to invest in Australia. This tension could be eased by allowing data centres to proceed where they can demonstrate that they will reduce curtailment of existing renewable generators, as that demand would be met by renewable energy that would otherwise be wasted.

The alternative is to initially offer data centre developers speed-to-power in exchange for catalysing our energy transition and investing in Australia's emerging green economy. Where this happens, strong safeguards are required to prevent this from becoming a repeating cycle every time a new facility is built. We discuss options for this below.

Renewable Electricity Guarantee of Origin Overshoot. The Australian Energy Market Commission (AEMC) has proposed that data centres meet the expectation to bring their own renewable energy by surrendering Renewable Electricity Guarantee of Origin (REGO) certificates. A REGO represents 1 MWh of renewable electricity generation, tagged with the generating facility, its location and the hour it was generated. Under the AEMC's proposal, every MWh a data centre consumes must be matched by a surrendered REGO in the corresponding time bucket (for example, peak or off-peak). Additionality would be secured by limiting qualifying REGOs to those from new renewable generators.

In the lag before new generation can be built, data centres will be unable to meet this liability. This consultation paper therefore proposes that data centres temporarily meet their obligation using REGOs from existing facilities, with the required proportion from new facilities increasing over time until existing facilities are phased out. Allowing this could help establish the market, but the number of qualifying facilities must be strictly limited, or a price floor set, for the scheme to offer a meaningful price signal. Even then, it does not address the more fundamental issue: the cycle of delay is likely to repeat with every new data centre.

This could be addressed through a system of REGO overshoot. This would see data centres compensate for their initial period of unmatched demand by having their REGO liability build up and surpass their actual energy use. **Signalling this inflated future demand now, tied to data centres already under construction, gives the market a more secure basis to invest proactively in renewable generation beyond what data centres themselves need, genuinely turning them from net-users into net-generators.**

Furthermore, data centre tenants should be required to sign power purchase agreements or [equivalent contracts](#) with renewable energy generators to offset their own power use and ensure that additive capacity actually has the long-term revenue stream required to secure their own finance. This requirement should apply promptly after the data centre reaches final investment decision to avoid exacerbating the lag between demand and supply.

Special Investment Vehicles. Alternatively, the cycle of generation lagging demand could be avoided by coordinated and proactive investment in renewable energy. For example, data centre operators could be required to pay into special investment vehicles – e.g. the Net Zero Fund or a fund overseen by the Clean Energy Finance Corporation – which can be used to proactively provide finance for renewable energy projects their demand would eventually need.

Q8: 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?

The direction of travel for these requirements was made clear when the federal government released its [expectations of data centres and AI infrastructure developers](#) on 23 March 2026. Requirements should apply to all data centres which reached final investment decision after this date. Acknowledging the complexity of amending engineering plans late in project development, any hard infrastructure requirements (e.g. closed-loop cooling) should not apply to projects that submitted applications for development approval before this date – however, suitable workarounds should be implemented to ensure they still comply with the overall standards.

Consultation questions on renewable electricity certificate obligations for data centres

Q9: How could obligations be phased in to account for the time needed to build data centres and new renewable projects?

See Q7.

Q10: How should ‘new’ renewable capacity be defined in relation to data centre investments?

For the purpose of the REGO liability, ‘new’ renewable capacity should constitute renewable energy generation projects not yet at final investment decision by the date the liability is enacted.

However, we do not consider projects with successful Capacity Investment Scheme (CIS) tenders as ‘additional’ or ‘new’ for the purposes of direct PPAs. The CIS provides government-backed revenue certainty for renewable energy projects, reducing investment risk and enabling developers to offer PPAs at more competitive prices. At present, though CIS wind projects have not yet been reaching financial close at the desired rate, we are hopeful that the REGO liability combined with the forthcoming Electricity Network Entry Mechanism contracts are sufficient to get these projects over the line. To date, CIS underwriting has been used by developers to offer more competitive PPAs. We do not consider public underwriting to be appropriate for data centre electricity consumption.

Given the significant purchasing power of large data centre operators, we do not consider it appropriate for their electricity demand to benefit from public underwriting intended to bring forward new renewable generation. Instead, data centres should use their own purchasing power to support additional renewable projects through direct PPAs.

Q11: At what capacity threshold should an obligation to retire or surrender renewable energy certificates apply?

The REGO liability should apply to data centres above 10 MW. Beyond this, it is essential that the size of an organisation’s data centre portfolio can also cumulatively trigger requirements for all capacity above this threshold. Data centres in the 10–50 MW range are precisely those which will connect to distribution networks and can play an important role in decreasing the network components of bills, if their implementation is managed well.

Q12: Should the information on REGO certificates be used to time-match electricity use in certain periods? If so, what should these periods be?

BZE supports the use of REGO time-stamps to match electricity use in certain periods. If no time restrictions are enforced, data centres will simply purchase REGOs from solar facilities during the middle of the day, even though their energy consumption is relatively uniform. At the very least, data

centres should have peak, off-peak and overnight liabilities which must be independently satisfied. Crucially, this creates a stronger price signal for renewable energy generation when supply dwindles outside the solar window. This is especially useful for wind and storage.

Q13: Are there any complexities in how data centres could acquire certificates to meet a renewable electricity obligation?

See Q7.

Q14: 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?

Project tracking database RenewMap estimates the total pipeline of data centres in Australia at over 32 GW – far in excess of what will actually be able to connect to the grid. As a result, connection agreements should be prioritised for developers and tenants that are willing to curtail when directed by grid operators. AI data centres have a range of options available to them to meet these requirements:

- **Battery storage:** Utility scale batteries are increasingly being used in data centres as a buffer between their IT equipment – which can ramp up and down very quickly – and the grid, which has stringent ramp-rate and voltage ride-through requirements. Embedding this as a technical requirement for large data centres would ensure that temporarily switching to battery storage during peak times is available as an option. Switching to backup diesel should not be viewed as a comparable approach to flexibility, as it would drive up emissions and increase our already world-leading reliance on a fuel we do not produce domestically.
- **Geographic flexibility:** Data centre tenants, who have facilities at different locations, can reroute jobs in high-demand locations to facilities located elsewhere.
- **Temporal flexibility:** Flexibility can also be achieved by delaying non-essential AI workloads. It has regularly been established that training workloads are straightforward to delay and do not have any immediate time-sensitivity. This can be achieved by pausing jobs outright or changing GPU resource allocation.

In addition, both [OpenAI](#) and [Anthropic](#) already offer flex/batch API rates for lower-priority tasks, in exchange for slower response times or service unavailability. Although neither are marketed specifically for grid stability at this stage, they should be extended to meet this obligation. AI agents are increasingly being given longer problem-solving tasks, with fewer human interactions. A priority system attached to this – or slower agent speeds during peak periods – will also help manage this large emerging workload.

- **Power capping:** US start-up Emerald AI has experimented with capping the power draw of NVIDIA GPUs using a system of [dynamic voltage frequency scaling](#). The physics of GPUs means that, for many tasks, lowering the voltage at which the chip is operating reduces its power draw more than performance is affected. In a study performed at an Oracle cloud data centre in Phoenix (Arizona), a 25% reduction in GPU power consumption was found to reduce chip throughput by less than 10% for fine-tuning (model adjustment) and inference workloads, while a 50% power reduction only sacrificed 20% of performance.

Emerald AI has developed software and predictive simulations to coordinate power capping and temporal flexibility in response to instructions from grid operators. So far, this has been tested in [Phoenix](#) and in the United Kingdom with [National Grid](#). The federal government should establish pilot projects targeting this with organisations like Emerald, so their compatibility with AEMO's workflow and standards can be established.

Q15: 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?

In this circumstance, it is incumbent on the state-owned operator to provide transparent and detailed modelling demonstrating how the removal of a renewable energy obligation would make energy more affordable for its customers, or that the data centre's siting close to renewable energy generation unlocks significant curtailed wind and solar. In addition, any exception should only be for the requirement to sign PPAs with generators – with the REGO liability remaining intact. Where thermal sources of energy are used to power data centres while PPA-linked renewables projects are in construction, the water consumption of that energy should be added to the facility's water obligation.

Q16: What unintended consequences or implementation challenges should the government consider in designing a renewable electricity framework for data centres?

See Q7.

Consultation questions on making a positive contribution to the energy transition

Q17: 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).

BZE has not responded to this question as it falls outside the scope of our expertise.

Q18: 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?

Water and energy efficiency are important aspects of sustainable data centre development – and can act as a baseline for compliance – however, in many cases will be the result of appropriately applied standards, rather than the standards themselves. For example, requiring data centres to use closed-loop cooling by default will guarantee water efficiency, and comprehensive flexibility arrangements can do more good for grid management and renewable energy matching than efficient but constant electrical load.

Q19: 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?

When designing energy and water performance thresholds, it is important to distinguish data centres by scale and purpose. Data centre impositions upon a host community are cumulative rather than isolated. As a result, efficiency standards matter more for large facilities, or in regions with a large number of facilities congregating, and this context must be captured.

The social acceptance issues currently stalling data centre development globally has arisen because of rapid buildout for AI training and inference. These facilities – bigger and more powerful than their predecessors – impose more upon the utility systems of their host regions, all in the name of advancing a technology which many fear will contribute to mass job loss, or even societal collapse. This is a very different prospect to building data centres to run gaming servers, streaming platforms and host enterprise data, and the standards should reflect this.

Q20: 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?

BZE supports quarterly disclosure of data centre energy use, efficiency and emissions.

Consultation questions on using innovative, efficient, and sustainable solutions to minimise water use

Q21: Should data centres be required to meet performance metrics or ratings for water usage and efficiency in data centres?

Q22: 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)?

Q23: Should a minimum NABERS data centre rating, based on WUE and agreed by states and territories, be required once a facility is operational?

Q24: 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?

Q25: What aspects of water use should data centres transparently and publicly report? E.g. maximum and peak demand, potable and non-potable water sources, efficiency measures, drought resilience measures and impacts on local water systems.

Q26: Should a nationally consistent hierarchy of water sources be adopted for data centres, prioritising non-potable/recycled water with potable water as last resort?

Q27: Where recycled or non-potable water supplies are not reasonably available, what alternative approaches might be promoted to achieve equivalent water security outcomes? For example, water offsets within the same catchment, investment in recycled water infrastructure with an associated transition timeframe, desalinated water supply, and local water efficiency programs.

Q28: 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?

Q21–28: It is essential that data centres meet both energy and water efficiency standards, acknowledging the interrelated nature of these two metrics. A clear hierarchy of solutions is necessary to simplify design trade-offs and give communities comfort concerning water availability. BZE does not have specialist expertise in data centre system design. However, industry examples suggest that minimising water use for cooling is now largely a solved problem.

The baseline expectation should be to use **closed-loop** systems. Increasingly, GPUs are cooled via heat exchange with sealed loops of water. Closed-loop systems reject heat within this internal loop using dry cooling or chillers. Such systems draw down upon water when filled for the first time, but

unlike evaporative cooling, do not require ongoing water consumption.

Major operators already treat closed-loop cooling as standard:

- Since August 2024, [Microsoft](#) has used a new data centre design that consumes zero water for cooling, using chip-level cooling without water evaporation.
- [Meta](#) says the majority of its newest AI-optimised data centres use closed-loop liquid cooling.
- [Oracle](#) says its data centres use closed-loop, non-evaporative cooling, in response to community concerns about local water supplies.

Domestically, [Canberra Data Centres](#), which operates data centres totalling 615 MW of capacity, with a further 1,351 MW in construction and 1,828 MW in development, already uses closed-loop systems to cool all of its existing facilities.

Where closed-loop systems are not possible, or result in an untenable electricity efficiency penalty, alternative designs can be used as long as they demonstrate minimal impact on water availability in the region. Examples could include using wastewater or seawater to efficiently cool the GPUs. Such a system is already in use at the [Sines Start Campus](#) in Portugal, which uses sea-water cooling to reject heat from the inner closed loop, achieving a PUE of 1.1 and WUE of 0.

Above these baseline requirements, there will also be opportunities for data centres to invest in local water resilience. Examples of this are already appearing domestically in other industries. As part of their work on the Uungala Wind Farm, Squadron Energy – along with the Dubbo Regional Council – has invested in an [advanced wastewater treatment facility](#) as a legacy benefit for the region. Similarly, to offset water use in Chile, Microsoft has partnered with agriculture technology company [Kilimo](#) to deploy their precision irrigation systems in affected communities. However, where these indirect offsets or investments are used, it is essential that they invest in the same local water environment they are drawing from.

Finally, facilities with a modified renewable energy obligation must have the water use associated with electricity generation included as part of their cumulative impacts. Sites seeking an exemption will have to show not only improved cost outcomes for consumers, but also verify that this does not come at the expense of local water security.

Consultation questions on imposing no net costs on consumers and communities

Q29: 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?

Like all large electrical loads, data centres pay for the costs of their own grid connection infrastructure – including some upgrades to common-user infrastructure (e.g. substations) required for them to connect, but which have benefits to the local community. However, when a large number of data centres are proposed in a specific region, it can accelerate the need for large-scale investment in the transmission network.

This scenario is already playing out in Western Sydney. Within 10km of Transgrid's West Sydney bulk supply point, over 1 GW of data centre capacity is already operating or under construction, with another 3.5 GW in development. For reference, this amount is larger than Australia's total aluminium smelting demand, all concentrated in a single corner of Sydney. Demand projections in this region have already brought forward the need for Transgrid's Sydney Ring South project from 2037-38 to 2033 in the 2026 Integrated System Plan. This is being addressed in a rule change request to the Australian Energy Market Commission.

The key issue is ensuring that the costs of network augmentation driven by concentrated data centre demand are not passed on to other electricity consumers. Where data centre development brings forward or increases the scale of transmission investment, charging arrangements should ensure those costs are appropriately attributed to the developments creating the additional demand, rather than being socialised across the broader customer base.

Q30: What are the barriers to data centres investing in sufficient new renewable energy and storage capacity to mitigate impacts to wholesale electricity costs?

See Q7.

Q31: What are the barriers to data centres investing in initiatives that would mitigate impacts to other water users, such as using recycled water?

BZE has not responded to this question as it falls outside the scope of our expertise.

Q32: 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? This could include restrictions on the use of non-renewable back-up generation, and scenarios under which flexibility should be required, for example during specified system stress events and scarcity events?

See Q14.

Q33: What barriers would prevent data centres from adopting demand flexibility measures to support the wider electricity system?

Some data centre workloads – both AI and more traditional cloud compute – have legitimate reliability requirements which cannot be as reliably flexed (though batteries can still be used). However, it is worth noting that the industry, in some circumstances, has been resistant to offer flexibility, not due to any technical barrier, but instead because remaining online is an economic preference. Data centre operators will be more willing to sign flexibility agreements if it is understood as a requirement for obtaining grid connection. This is already evident in industry support for flexibility frameworks like that of [Emerald AI](#).

Q34: What tools could be used by governments to require data centres to adopt demand flexibility?

Willingness to offer demand flexibility should be built into the connection agreements of a facility, with an indication of the level of flexibility required. This would provide hard constraints on the number of inflexible workloads able to be offered by any one facility.

Q35: 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?

BZE has not responded to this question as it falls outside the scope of our expertise.

Consultation questions on meaningful engagement with communities and councils

Q36: What does best practice community engagement look like for data centres, including with nearby businesses?

Best practice engagement for data centres – particularly larger training and inference facilities – should match the standards now expected of renewable energy developers in those regions. Applying the principles from our National Action Plan community engagement assessments, we consider best practice for data centre development in these regions to involve:

- establishing a genuine local presence, such as a staffed local office, rather than engaging remotely;
- holding regular town halls and other accessible ways for residents to raise concerns;
- actively listening, and asking communities what matters most to them;
- being open about where design choices remain flexible, such as cooling technology and water sourcing, so outcomes can be shaped around local priorities;
- being upfront about proposed sites early, and not using non-disclosure agreements where they aren't strictly necessary, which has regularly been the case in, e.g., [Virginia](#).

Engagement must also be honest about both benefits and risks. That means realistic accounts of ongoing local benefit, including operational rather than construction-phase employment, alongside a clear account of the risks and how they will be mitigated. A data centre is a long-lived asset expected to operate in a community for decades, and the proponent is investing in that community and its people, not only in the land.

Q37: What role should the Commonwealth's mandatory requirements play in relation to location and community engagement? For example, the Commonwealth could promote transparency around where new data centres are being developed or plans for community engagement.

The Commonwealth's mandatory requirements should do three things: promote transparency about where data centres are being developed, set clear standards for best-practice community engagement, and build the capacity of local communities to engage effectively.

Communities frequently learn of large proposals late, sometimes only once sites are secured or approvals are well advanced. The Commonwealth should maintain a national public register of proposed data centres, requiring proponents to disclose location, contracted and ultimate capacity, phasing, and connection status. This would give communities and councils early visibility of what is coming. It would also give governments and network planners a clear view of cumulative demand, and reduce duplicate connection applications.

Second, the Commonwealth should set clear national standards for best-practice community engagement, and offer its support only to projects willing to meet them. With a pipeline of around 32 GW of proposed capacity competing for finite grid capacity (according to the RenewMap database), the

Commonwealth is in a competitive position to only offer support to developers willing to engage in best-practice via vehicles such as the Investor Front Door, or Foreign Investment Framework.

Third, the Commonwealth should support local councils and communities to engage effectively. As we state in Q38 below, converting engagement into lasting benefit is a two-way street. Councils and communities need the capacity to understand what is proposed, to articulate what they want from it, and to negotiate effectively. The Commonwealth can build this capacity through funding, education programs, and by sharing case studies of what has worked in renewable energy host communities.

Q38: 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?

Delivering lasting community benefit is a two-way street. Proponents must be willing to build relationships, listen and find out what a community needs. But communities will only see the best results when they are coordinated, having done the work to agree at a broad level on what they want major investment to deliver. That might be jobs, training, housing, or on-the-ground infrastructure such as water recycling.

Renewable Energy Zone development in New South Wales has already given us examples of this working well. For example, Hay Shire council, working with RE-Alliance, ran a series of workshops which unearthed the key issues affecting the town ([Australian Government 2025](#)). The council has since adopted [Renewable Energy Position Statements](#) in November 2025, supported by a Community Benefit Framework, designed to inform conditions of consent and guide negotiations with proponents. Similarly, Squadron Energy, through their Uungala Wind Farm, has partnered with Dubbo Regional Council to construct an advanced wastewater treatment facility in the region.

Councils cannot do this unsupported. Regional councils in particular are often left to absorb the workload of major project engagement, strategic planning and benefit negotiation without the funding or staff to do it well. State and Commonwealth governments should directly fund regional council capacity to prepare for large-scale investment, including data centres.

Q39: 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?

BZE does not have specialist expertise in acoustics or site design, and does not propose a specific minimum distance. However, noise and emissions from backup power are significant local impacts that must be actively managed. Data centres emit a constant, low-frequency hum that rarely breaches conventional noise limits but can affect nearby residents' wellbeing ([JLARC 2024](#)). Minimum setbacks should therefore work alongside good site design and noise standards that account for low-frequency sound, assessed from the earliest stage of development.

Enforcement will largely fall to state planning agencies through zoning and conditions of approval, as Victoria's recent 150m buffer shows. The Commonwealth's role should be to set consistent national minimum standards that states can strengthen but not weaken, preventing a race to the bottom between jurisdictions.

Q40: How important is proximity to population centres for data centre operators given improvements in connectivity and advances in fibre optics?

Although proximity to population centres has been an important consideration for previous generations of data centres, these restrictions are not as strict for many emerging AI workloads. The clearest case here is AI training, which is highly tolerant of latency to end users. Instead, it needs very low latency within the cluster (between GPUs) which drives demand for very large single facilities. Large sites, abundant power and proximity to generation make regional locations the natural home for training.

Inference is more nuanced, but latency tolerance is far more workload-specific than a simple training/inference split suggests. A narrow class of applications does depend on minimising latency, such as algorithmic trading, where proximity is itself the competitive advantage. By contrast, most chatbot and frontier model use can tolerate even up to a few seconds before degrading user experience ([Maslych et al. 2025](#)). Given that light in fibre travels approximately 200,000 km per second, this allows considerable geographic flexibility in where these workloads are located. This is a contributing factor for our recommendation that intended workload form part of the tiering system for data centre requirements.

The relaxation of proximity requirements is no more evident than in Anthropic's choice of the Western Downs Digital Park as an [inference hub in Australia](#), despite being 250 km from Brisbane, and being sized to service markets much larger than Brisbane itself. Regional siting of inference is therefore readily achievable, provided proponents fund the infrastructure it requires. Indeed, extending this connectivity to surrounding communities offers a ready-made benefit for regional communities which currently lack high quality internet access.

Q41: How should competing land uses be managed and prioritised, for example where land is also suitable for housing or agriculture?

Competing land uses should be managed using the approaches now emerging for renewable energy siting. Many states and territories already have publicly available geographic information system (GIS) databases that map land by its value for different uses. These should factor into planning decisions and proactive studies identifying land suitable for industry, high-value agricultural land, areas of high biodiversity value, and land of cultural significance, where this information can appropriately be shared.

Data centres should not consume the most productive agricultural land. The picture is more nuanced for land of moderate agricultural value. The renewable generation needed to supply data centres can support farm incomes – wind turbines occupy a small fraction of the land they sit on, and solar farms coexist well with sheep grazing. For farmers facing growing climate pressures, hosting this infrastructure is a well-established way to diversify income. Realising this benefit requires involving landholders and

the wider community early in consultation, reinforcing the engagement and benefit-sharing principles outlined above.

Housing requires the same strategic approach. Governments should identify priority housing locations, particularly areas well served by public transport, such as land around regional and suburban train stations. Large data centres should not be built on land better suited to housing where demand is foreseeable. They should instead be directed to industrially zoned sites poorly suited to residential use. This issue is most acute in Western Sydney and Melbourne's west, where data centre development is concentrating. AI training facilities, which need very large sites and have no latency-based reason to locate in dense urban areas, should not be sited there. As discussed at Q40, much inference can also be located elsewhere.

Q42: 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. A central recommendation of BZE's National Action Plan is that infrastructure investment must be sequenced and a region's needs appraised holistically. The grid is being built out for the renewable transition, and data centres are seeking to secure whatever network capacity they can find. Because they can be built far faster than green industrial facilities and renewable energy generation, they risk pre-empting capacity that those industries will need and entrenching a cycle where generation cannot keep up.

States and territories should work with network, water and telecommunications businesses to strategically identify regions where data centres can be developed and managed well. Responsible organisations should take the significance of a region for emerging green industries into account when making these decisions. Some locations close to generation or existing grid infrastructure may be well suited, provided host regions are supported through the process, the benefits are clearly explained, and communities do not feel the development has simply been imposed on them.

Q43: 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?

BZE has not responded to this question as it falls outside the scope of our expertise.

Q44: Should data centre proponents be required to explore options with government and industry to uplift community telecommunications connectivity alongside their own investment?

Telecommunications uplift is one of the clearest levers available to data centre proponents for delivering lasting community benefit. The bandwidth needed to substantially improve internet services for a region such as the Western Downs, the Latrobe Valley or the Riverina is minuscule compared with the data flowing to and from a multi-gigawatt facility, particularly where new fibre is already being built. For

communities, the benefit is substantial and highly visible. Improving regional and rural connectivity has proven difficult and expensive for governments everywhere. In the United States, the federal government has devoted years of effort and tens of billions of dollars to rural broadband programs.

Energy infrastructure is already being used this way in Australia. On HumeLink, Transgrid, in partnership with Lumea, is [installing telecommunications boosters](#) on transmission towers to improve mobile and internet coverage along the 365km corridor. Data centre proponents building regional connectivity could deliver comparable benefits should they fit the priorities identified through community engagement.

Consultation questions on building skills and training

Q45: 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?

Data centres draw on the same pool of construction workers, tradespeople and engineers needed across the economy, from housing construction to the energy transition. That pool is already under strain, and Australia is not training enough people to meet emerging demand ([Powering Skills Organisation 2025](#), [Master Builders Australia 2026](#)). Workforce obligations must take a cumulative view. Using a local workforce is valuable, and proponents should prioritise it, but drawing too heavily on a region's existing tradespeople leaves too few to meet other local needs, from the energy transition to household renovations. The central obligation should be to add to the skilled workforce, not simply draw it down.

BZE recommends a compulsory apprenticeship requirement for data centre construction, at a minimum matching the Commonwealth's [Australian Skills Guarantee](#). For major construction projects, the Skills Guarantee requires at least 10 per cent of labour hours to be undertaken by apprentices and trainees, with targets for women rising annually to 12 per cent of apprentice and trainee hours by 2030. Equivalent targets should apply for other under-represented cohorts, including First Nations workers.

Obligations should also prioritise transition pathways for workers leaving sunset industries, where appropriate. Several proposed data centre sites are former or closing coal-fired power stations, whose workforces have skills in high-voltage systems, cooling and plant operation that transfer well.

Training opportunities alone will not grow the workforce if apprenticeships remain unaffordable to take on and complete. According to Master Builders Australia, more than half of construction apprentices do not finish their training, with financial barriers cited as a significant barrier. Alongside operator obligations, BZE supports the financial reforms proposed in Master Builders Australia's [Future Workforce Blueprint](#), in particular a Commonwealth payroll tax rebate on apprentice wages in critical occupations and a Construction Training Bond to improve completion rates. These measures address the cost barriers that prevent employers from taking on apprentices and apprentices from finishing. The public benefit comes from the trained worker, who will support housing, infrastructure and the energy transition well beyond any single project.

Q46: What broader workforce development obligations should apply to data centre operators during ongoing operations?

Workforce obligations should continue beyond construction. Our own investigation of public data centre proposals puts ongoing employment at around 0.4–0.5 FTE per MW, reinforced by similar international studies ([Verdant 2026](#)). While planning documents and economic impact assessments do not disaggregate these by occupation, industry reports identify key jobs, including ([Mandala 2026](#)):

- Building and engineering technicians
- Electricians
- Air conditioning and refrigeration mechanics
- Air conditioning and mechanical services plumbers
- Electronic equipment trades workers
- Industrial, mechanical and production engineers
- Electrical engineers
- Software engineers
- Computer network professionals
- Database and system administrators, and ICT security specialists

In light of this, data centre operators should extend skills and training programs established in construction through to the operational phase, with a focus on upskilling local workers. This could include local scholarships and internships for technical study and occupations. Sponsored education of this kind has been used by major resources companies in regional Australia, and it builds a genuine, lasting connection between an industry and the communities it operates in.

Consultation questions on security, safety and transparency

Q47: 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?

BZE has not responded to this question as it falls outside the scope of our expertise.

Q48: 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?

BZE has not responded to this question as it falls outside the scope of our expertise.

Q49: 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.

BZE has not responded to this question as it falls outside the scope of our expertise.

Q50: What safeguards should apply?

BZE has not responded to this question as it falls outside the scope of our expertise.

Consultation questions on investing in local capability, skills, research and Innovation

Q51: 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.

All of the contributions listed in the consultation paper have value. BZE anticipates the greatest value will come from those that build lasting Australian capability, anchored in host communities and directed towards the industries Australia needs to build.

Investing in people. Proponents should fund scholarships, cadetships and sponsored study for people from the regions where they operate, creating local pathways into AI and engineering. Research partnerships should span the full technology stack, from hardware through to models and applications, so Australian researchers build capability across the whole system rather than only as its users.

Targeted access to compute. A workable compute reservation scheme would be highly valuable. GPU access is a real constraint for Australian researchers, yet the compute needed for most research and industrial applications is small relative to what large-scale training and inference consume. Reserved compute should be directed to defined priorities, including green manufacturing and industrial productivity, from primary processing through to advanced manufacturing. Hysata's [AI-optimised electrolyser production line](#) at Port Kembla, designed to produce up to 15% more product from the same physical line, is an example of the kinds of improvements possible in this space.

Green supply chains. Though AI itself will undoubtedly offer economic benefits, the exact scale and nature of these remain highly uncertain. Supply chain investment offers a concrete, immediate contribution which can help build demand certainty for green manufacturing in Australia. Data centres are highly material-intensive, and proponents should commit to supporting Australian production of green commodities, such as green iron and steel, low-carbon cement and green aluminium. These products currently carry a green premium while the industry scales. Data centre developers, with vast capital and a demonstrated willingness to pay for speed, are well placed to cover it through long-term commitments.

Hyperscalers have already shown a willingness to support green commodities. Microsoft has agreed to use near-zero emission steel from Swedish producer Stegra in its data centres ([ref](#)), while Google has agreed to buy environmental attribute certificates for up to 91,000 tonnes of Stegra's first-year production ([ref](#)).

Grid transformation. The data centre boom is proof that a "build it and they will come" approach to transmission is effective. Developers are scouring the world for spare grid capacity and directing capital wherever they find it. Similarly, large pockets of grid capacity will also be needed to grow the industries

of a decarbonised economy, from green iron to advanced manufacturing. BZE's [Safeguarding our Future](#) report with ACIL Allen shows the cheapest way to build that grid is through fewer, larger projects, which can be up to 75% cheaper than a string of incremental upgrades delivering the same capacity.

Building at that scale runs into the investment tests that decide what gets built. National rules [exclude any benefit](#) that cannot be measured within the electricity market, so the industries a line makes possible never enter the calculation, and lines are sized to the demand we can already see.

Data centres can help break this cycle. NSW has proposed a [Major Network Upgrade Fee](#), charged per megawatt and paid into an infrastructure fund outside network revenue. If payments like this were earmarked and extended nationally, they could underwrite capacity on build-ahead lines, much as the [US Transmission Facilitation Program](#) buys anchor capacity on new transmission and resells it as customers arrive.

The only way to build a grid fit for the future economy is to let investment tests weigh economic opportunity and the cost of underbuilding, while putting systems in place that shield consumers from the downside.

Q52: 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.

BZE favours a standards-based approach over case-by-case negotiation. A clear code is more transparent, sets consistent expectations across proponents, and avoids the uneven outcomes and opacity of bespoke deals.

On compute reservation, BZE supports the principle and suggests it be set as a percentage of capacity, tiered by facility size. The detailed design of any scheme should be developed through wide consultation with those who would use it, including researchers, manufacturers and host communities, and with technical experts on what is workable.

Green supply chain commitments should also be built into the baseline standard. Australia's support for green commodities to date has focused almost entirely on the supply side. The Hydrogen Production Tax Incentive, a refundable tax offset of \$2 per kilogram of eligible low-emissions hydrogen, and the \$2 billion Green Aluminium Production Credit both reduce the cost of producing green products. But they do not fully close the cost gap, and nothing yet creates a market willing to pay the remaining premium.

BZE proposes a certificate-based liability for green commodities, running parallel to the REGO liability in the data centre energy obligation. Proponents would accrue an obligation proportional to their size, discharged by purchasing certificates from Australian green commodity production, with the required share rising over time. Because many of these industries are not yet producing at scale, the obligation should be able to be met through forward commitments. This follows the same logic BZE proposes for the energy obligation: a temporary shortfall is permitted while supply matures, and obligations step up once it does. Alternatively, proponents could meet the obligation through long-term firm offtake

contracts, which give emerging producers the revenue certainty and risk mitigation they need to reach final investment decision.

Q53: Are frontier AI developers best placed to make contributions in these areas? Should other actors be considered – for example, hyperscalers and neocloud providers.

Industrial contributions should be expected from all developers and major tenants, including frontier AI developers, hyperscalers and neocloud providers. For example, given data centres consume construction materials irrespective of their operator or tenant, all facilities should carry a green supply chain obligation.

Beyond this, however, different actors are suited to different contributions. Frontier AI developers are best placed to offer research partnerships at the frontier of model development, or in AI integration in manufacturing. Hyperscalers such as Google, Microsoft and Amazon can offer contributions aligned with the services they provide, from cloud and tooling to industry-specific technology partnerships. What matters in every case is genuine investment in economic partnerships with Australian research, industry and communities.

Multi-tenant facilities are harder to orchestrate, as no single tenant may hold enough capacity to make a meaningful contribution alone. In these cases, obligations may be better placed on the data centre developer, allowing contributions to be pooled across tenants, where they can deliver more than a series of small individual commitments.

Q54: 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 portion of reserved compute should be directed to defined priority areas in the national economic interest, set by government rather than left to providers. Examples include:

- **Orchestrating distributed energy.** Australia has among the highest per-capita rooftop solar uptake in the world, with home batteries following, and EV uptake accelerating amid fuel price increases sparked by the conflict in the Middle East. Coordinating millions of these devices to lower household bills and support the grid is a demanding computational task. Work is already underway to tackle this. For example, UNSW is actively developing AI-driven platforms for forecasting, demand-side management and digital-twin modelling to coordinate distributed energy and virtual power plants, and Wollongong's Gridsight has raised capital for an AI platform helping network operators integrate rooftop solar and batteries. Reserved compute could accelerate this work and help ensure its benefits flow to consumers.
- **Managing a renewable grid.** A least-cost grid will increasingly be dominated by inverter-based resources, with grid-forming inverters taking over stability services traditionally provided by spinning machines. AI can support forecasting, real-time optimisation and fault response. Fortescue's Pilbara network is an early demonstration of this phenomenon. When bushfires removed all thermal generation in October 2025, its solar and battery systems responded instantly and kept supplying around 85MW of load, and Fortescue is deploying AI systems for real-time weather forecasting and grid responsiveness.

- **Green manufacturing.** As outlined at Q51, AI can lift productivity across green manufacturing, from primary processing to advanced manufacturing, with Hysata's AI-optimised electrolyser line at Port Kembla a defining example.

Q55: 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.

For green supply chain commitments, the indicator should be the share of a facility's construction materials that meets a recognised low-carbon or near-zero emissions standard, rising over time.

Work has been underway on global guarantee of origin systems for some time. The IEA's near-zero emissions definitions for steel and cement provide a robust starting point, and industry and multi-stakeholder standards, including ResponsibleSteel and the Global Cement and Concrete Association's concrete ratings, are converging around them. Australia also has its own working definitions: the Hydrogen Production Tax Incentive applies an emissions intensity threshold of 0.6kg of CO₂ per kilogram of hydrogen, verified through the Guarantee of Origin scheme, and the Commonwealth is funding work on green aluminium standards alongside the Green Aluminium Production Credit. Data centre supply chain obligations should align with these definitions as they mature.

More broadly, indicators should be treated as iterative. Governments should not legislate a fixed set once and leave it unchanged. They should build the capability to measure contribution well, and sharpen indicators as better metrics become available. At the same time, indicators should not be so prescriptive that they lock proponents into supporting Australian industry and research in only a few predefined ways. The framework should leave room for proponents to be creative in how they contribute.

Q56: 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?

Much of the skills and training contribution should follow from the economic investment outlined at Q51 and Q52. Research partnerships, compute access for Australian industry, and green supply chain commitments all build demand for skilled Australian workers and create pathways into those roles. Skills development is best secured as an outcome of genuine economic partnership, not as a separate compliance exercise.

Frontier AI developers should also invest directly in education for people in the communities that host their infrastructure, as a core part of their community benefit commitments. This could include scholarships, sponsored study, cadetships and internships that give local people, particularly young people in regional host communities, a pathway into AI, engineering and related fields.

BZE recognises that this question also reflects broader concerns about the long-term effect of AI on employment, including the risk of large-scale workforce displacement. These concerns are legitimate and warrant dedicated policy attention. They are, however, beyond the scope of BZE's expertise and this submission.

Consultation questions on providing a simple, consistent national framework

Q57: The Commonwealth proposes reporting and transparency requirements at a federal level for large data centres, so that impacts and benefits to communities are transparent and trusted. Where might this proposed approach interact with existing legislation? Please note if these are Commonwealth, state and territory, or local government legislation.

BZE has not responded to this question as it falls outside the scope of our expertise.

Q58: What additional regulatory reforms or incentives would encourage frontier AI training in Australia?

BZE has not responded to this question as it falls outside the scope of our expertise.