AUSTRALIA'S POSITION

IN THE AGE OF AI

AUGUST
2026

Good Ancestors

Table of Contents

Executive Summary

AI development is rapid, accelerating, and on track to upend geopolitics and the global economy.

Australia needs to find valuable positions in a world with highly capable AI. We then need to use these positions to secure our national interests, whatever other countries and companies do and however the technology unfolds.

A position can be valuable if it secures Australia a place in the economy and supply chains that will underpin future prosperity; gives Australia something to offer our allies and regional partners that they cannot easily obtain elsewhere; or reduces Australia's dependence on decisions made offshore.

We have identified and analysed ten possible positions that Australia could pursue to establish its place in an AI future:

1. Attract frontier AI training
Australia could attract one or more frontier AI companies to make their future AI models in Australia.

2. Build inference capacity for Australia and the region
Australia could have enough domestic AI inference compute to meet domestic demand and be a net exporter.

3. Secure industries across the AI value chain
Australia could secure multiple valuable places across the AI supply chain, beyond hosting AI training and inference compute.

4. Fund narrow AI models in areas of national strength
Australia could support the development of truly sovereign small or narrow AI models in areas of national strength.

5. Develop a sovereign fast chaser
Australia could support the development of an Australian "fast-chaser" AI model that aims to have most of the performance of frontier models at a fraction of the cost.

6. Coordinate with middle powers on AI
Australia could use its influence and trusted relationships to coordinate middle powers on reducing AI risks and mitigating power concentration.

7. Fund a world-leading AI Safety Institute
Australia could have a world-leading AISI that makes a meaningful contribution in one or more critical areas of technical AI safety, in ways legible to our allies and frontier AI companies.

8. Drive AI adoption across industry
Australia could accelerate the safe and responsible adoption of AI across domestic industry.

9. Invest in AI-proof industries
Australia could invest in industries unrelated to AI but that are likely to become much more valuable in a world with cheap intelligence available at scale.

10. Host the new industries that AI will create
Australia could prepare for a large expansion in energy, land use, and industrial capacity so that the new industries AI creates are built here.

Assessing the positions

Good Ancestors assessed each of the ten positions on two measures (see below figure). **Value** is how much Australia gains if we attain the position, through a place in the future economy, value to allies, or reduced dependence on decisions made offshore. **Feasibility** is how likely Australia is to attain the position, given legal, financial, competitive, and political hurdles.

Two positions emerged as most promising:

- 'Funding a world-leading AI Safety Institute' (position 7) rated highest overall. The AISI has recently been established, and sharing world-leading research has significant value. The main hurdles to overcome are addressing the funding shortfall and driving a research agenda that interests our allies and leading AI companies.
- 'Attracting frontier AI training' (position 1) rated highest on value. Australia is an attractive destination to train frontier AI models, and AI companies are showing interest. Feasibility is constrained by copyright law rather than Australia's suitability on other dimensions.

Of the ten positions evaluated, developing a sovereign fast chaser rated lowest overall. An Australian sovereign model may not outperform comparable models on key attributes, like price. The required investment – while lower than a frontier model – is still substantial, and licensing training data under Australian copyright law would raise costs relative to Chinese and US competitors.

Copyright law is a hurdle across several of the positions, including hosting frontier AI training, building inference capacity, and developing sovereign fast chasers or narrow AI models.

Value and feasibility of each position

01 Attract frontier AI training

Australia could attract one or more frontier AI companies to make their future AI models in Australia.

What's the case for it?

Frontier AI companies are on track to secure significant economic and geopolitical power. The creation of frontier models, rather than just their operation, is a critical place in the AI supply chain. Frontier models will continue to drive the industry, even if small and narrow AI models find niches.

Training models in Australia would be economically valuable, strengthen our alliances, and allow us to negotiate for access to other important capabilities.

Frontier AI training, unlike inference, will only happen in a small number of countries, making it particularly desirable. In supply chain terms, inference is like a petrol station, and training is like a fuel refinery. It's rarer and more strategically valuable.

What are the hurdles?

Frontier AI training has energy and community impacts that need managing. Training happens in data centres, meaning it faces the same hurdles as data centres generally. This is primarily high energy consumption, but also local community impacts if planning is inadequate. Ongoing standards development is important to minimise impacts, provide companies with certainty about what is required, and secure social licence.

Copyright law blocks frontier AI training. Australia's current copyright framework has no viable pathway to AI training:

- Collective licensing covers only some categories of copyright material in Australia and does not reach the long tail of blogs, forums, and open-source code that developers want to train on.
- Australia's narrow fair dealing exceptions prevent deals made in the US from being replicated here, such as News Corp selling its archive to OpenAI. Australian news can include third-party material, such as photographs and quotes, under the fair dealing exception for reporting the news. That exception does not allow the news company to license that third-party material on to someone else. An AI company using the archive would remain exposed to infringement claims. Fair use covers that gap in the US, but Australia has no equivalent. This means the archive cannot be practically licensed for AI training here.
- Government should develop a copyright pathway that enables AI training in Australia while supporting creators. Good Ancestors' consultation paper sets out one potential model, but other uniquely Australian approaches may also be possible.

The benefits of hosting AI training depend on the terms Australia negotiates. Pre-release model access, ongoing access to frontier models, and public-compute reserves will not come automatically from data centres being on Australian soil. Australia will need to negotiate with governments, AI companies, and providers. Maintaining robust relationships with frontier AI companies may become harder as they become more influential.

02 Build inference capacity for Australia and the region

Australia could have enough domestic AI inference compute to meet domestic demand and be a net exporter.

What's the case for it?

The share of global compute dedicated to running AI models, rather than building them, is rising fast. Before OpenAI o1 in September 2024, capability came almost entirely from training. 'Test-time reasoning' — the process of a model 'thinking' before answering a question — caused a jump in inference demand. Since then, inference has grown from a third of all AI compute in 2023 to half in 2025 and roughly two-thirds in 2026.[1]

While building AI models sits further up the value chain, running AI models and offering them regionally supports Australia's economy and security. As AI is adopted across the economy, including in critical infrastructure, failure to meet domestic inference demands introduces a new foreign dependency and exposes Australia to additional AI security risks. Countries that control our inference compute could drive up costs, remove access, or exploit their access to create vulnerabilities. Offshore inference is exposed to attacks like the substitution of the served model for a backdoored variant[2] and the extraction[3] or leakage[4] of prompts, outputs, and weights.

The best global estimate is that inference compute supply is growing at 3x per year, but demand is growing at 10x per year — pointing to a 'compute crunch'.[5]

- Today, Australia only has enough inference compute supply to meet less than half of domestic demand.
- Future estimates have uncertainty about supply, demand, and how they interact during a compute crunch. While at least 19 data centres possibly intended for AI are under construction, only two have confirmed tenants with detailed amounts of AI compute. Based on limited data, we assess that Australia will meet less than a quarter of its inference demand by the end of 2027. **(Attachment A)**

To safeguard AI and data sovereignty, Australia should be able to meet its domestic inference demand – including as that demand grows – with a significant surplus. Additional compute would position Australia to serve regional partners for whom hosting compute at national scale is not viable. This would reduce our regional partners' dependence on alternatives, such as Chinese-supplied AI infrastructure. Pursuing this position requires new work to measure Australian supply and demand and set a policy target to ensure we are a net compute exporter.

What are the hurdles?

AI inference has energy and community impacts that need managing. Inference happens in data centres, meaning it faces the same hurdles as data centres generally. This is primarily high energy consumption, but also local community impacts if planning is inadequate.

Australia has no target for how much inference capacity is enough. Australia is building AI inference compute, but is not catching up with demand and lacks a strategic target. An AI-inference compute target that exceeds domestic demand forecasts would provide both a justification and a boundary for an inference data-centre build-out, which may in turn support data centre social licence.

Legal uncertainties around copyright and inference are growing. AI inference relies on exceptions in the Copyright Act for copies that are "temporary, necessary and incidental". While this was defensible before OpenAI o1, copies created during AI inference may not be any of those things. If so, this could stall AI adoption and AI inference data centre builds in Australia.

---

[1] Stewart, D. et al. (2025). *Why AI's next phase will likely demand more computational power, not less*, Deloitte.
[2] Cai, W. et al. (2025). *Are you getting what you pay for? Auditing model substitution in LLM APIs*. arXiv:2504.04715.
[3] Anthropic & Irregular. (2025). *Confidential inference systems: Design principles and security risks*.
[4] Gu, Z. et al. (2026). *Blueprint, bootstrap, and bridge: A security look at NVIDIA GPU confidential computing*. Proceedings of Machine Learning and Systems.
[5] Emberson, L. & Sevilla, J. (2026). *Is a compute crunch coming?*. Epoch AI.

03 Secure industries across the AI value chain

Australia could secure multiple valuable places across the AI supply chain, beyond hosting AI training and inference compute.

What's the case for it?

Occupying positions in the AI value chain can generate export revenue, develop domestic expertise, give Australia something partners need in exchange for access to things we can't build ourselves (e.g., AI chips, access to frontier AI), and reduce dependence on offshore supply chains.

Critical minerals. AI chips, infrastructure, and robotics all need critical minerals, including copper, gallium, lithium, and rare earths. Australia is a major producer of several of these, although most are refined offshore. Australia should continue partnering with allies and investing in AI-relevant minerals. The Alcoa-Sojitz gallium project is one example, where the Government is partnering with the US and Japan to build out the non-Chinese gallium supply chain. Australia's position can bring us to the right tables, like with Pax Silica. Australia must use that position to secure access to capabilities we lack, including AI chips.

Clean energy. AI data centres are heavy energy consumers. Almost any amount of supply Australia can bring online would be met by AI demand if other regulatory settings are right. Selling sunlight to AI companies at a scale no other country can achieve would be highly valuable. AI demand could fund Australia's clean energy build-out, while also creating a position Australia can sell into.

AI assurance technology (AIAT). AIAT refers to the software, hardware, and services that allow organisations to understand and mitigate AI risks. The UK Government values its AI assurance market at £1.0 billion in 2024 and projects £18.8 billion by 2035.[6] One forecast puts the global AI assurance technology market at USD 276 billion by 2030.[7] Australia's success in safety-critical industries such as mining, aviation, and food biosecurity shows Australia can succeed in such an industry. Further, buyers may prefer assurance products independent of the AI companies they oversee. This is an upstream economic opportunity and would position Australia as a direct contributor to AI products consumed globally.

These are only a few examples, and supply chain analysis may reveal other areas where Australia can perform valuable AI-related activities.

What are the hurdles?

Australia needs supply chain analysis to understand where the biggest opportunities are. The Australian Government should commission a segment-by-segment capability assessment of the AI supply chain, benchmarking Australia's position at each stage against projected global demand and competitor cost curves.

Holding a valuable position in the supply chain is not the same as capturing value from it, particularly when the supply chain is complicated. Most critical mineral refining happens offshore, so turning Australia's strength in mining into either export value or leverage requires deliberate negotiation.

Capturing these positions requires investment. Government, for example, could seed an AI assurance technology venture fund with $90 million in cornerstone funding. This would match the inflation-adjusted investment that established CSIRO's Main Sequence Ventures, which attracted over $1 billion in total funds and created companies now valued at $6.8 billion.

[6] UK Department for Science, Innovation and Technology. (2025). *Trusted third-party AI assurance roadmap.*
[7] AIAT Report. (2024). *Risk & Reward: AI Assurance Technology Market Report 2024.*

04 Fund narrow AI models in areas of national strength

Australia could support the development of truly sovereign small or narrow AI models in areas of national strength.

What's the case for it?

Australia could make strategic investments in sovereign AI models in areas of national expertise where data availability and expert feedback are uniquely available, and where these models are likely to stay ahead of the general-purpose frontier. These could include computer vision and models for remote sensing, mining, agriculture, geospatial analysis, seismology, and health.

Current frontier general-purpose models access other tools to achieve their generality — a process known as 'tool calling'. For instance, if you ask a frontier model whether you need an umbrella, it will access an existing weather forecast tool, rather than perform its own weather forecast. Even if the frontier of general-purpose models progresses quickly, they will likely continue to access other models to complete tasks in narrow domains where specific training data and design are required for advanced performance.

What are the hurdles?

The general-purpose frontier has previously overtaken narrow or targeted AI models, causing AI start-ups to fail. However, others have captured significant value. Australia will have to accept the risk that some narrow domestic AI models will fail.

Making informed decisions will require evaluating areas of national strength. This analysis would seek areas of rich data and human experts who can participate in reinforcement learning. Rather than picking winners directly, Government could seed a venture fund based on the Main Sequence Ventures model, or run a dedicated stream through the National Reconstruction Fund's critical technologies priority.

Australia will face challenges unlocking data sets for AI training.

- Australian datasets often include third-party material and material subject to restrictions on subsequent use. Copyright is the most obvious example. Some material may be included in a dataset under a fair dealing exception (e.g., for research or reporting the news), but that exception would not extend to the dataset holder training a commercial AI model or licensing the data to someone who would. This is less of a challenge for narrow models in countries with fair use or text-and-data-mining exceptions, but it is a challenge in Australia.
- The same problem arises when data is collected under contracts, privacy settings, or sector-specific regulation that did not contemplate AI training. For instance, health imaging datasets can include radiologists' reports, geospatial data can include commercial base maps, mining data can include licensed seismic surveys, and agricultural data can be subject to ag-tech platform contracts.
- 'Unlocking' a dataset for AI training in Australia requires identifying the restricted material, then either removing it, which can degrade the dataset, or obtaining consent or a licence. The latter may be impractical or impossible depending on the scale and age of the material, whether rights holders or individuals can be found, and whether a deal can be reached. This is not a universal problem — some narrow AI models have been trained in Australia where the problem doesn't arise in their specific dataset, or where the dataset can be realistically unlocked. But in many cases, the time and cost of this process may remove the edge that made the model worth building in the first place.

05 Develop a sovereign fast chaser

Australia could support the development of an Australian "fast-chaser" AI model that aims to have most of the performance of frontier models at a fraction of the cost.

What's the case for it?

Fast-chaser models, often released with open weights, implement innovations discovered by frontier labs but without the research spend required to discover them. Currently, many of these models, like DeepSeek, Kimi, and Qwen, are Chinese. Relying on Chinese models does not give Australia a valuable position and exposes Australia to dependencies and security risks.

Fast-chaser models are not intended to compete economically in fields where users demand the best model they can access. However, for many use cases, users will want an AI model that can perform the relevant task cheaply, rather than the AI model that is strictly the best.

In a future with highly capable AI models, fast-chaser AI models with 90% of frontier performance may still be able to perform many economically valuable tasks. If Australia developed our own fast chaser, it could capture that economic value and reduce our dependence on Chinese-supplied models and continued access to the frontier.

What are the hurdles?

Fast chasers still require significant resourcing and claims that they are cheap to train are likely overstated. The Swiss AI Initiative was started in December 2023 with 10 million GPU hours and 20m CHF (A$35m) in federal funding through to 2028.[8] The resulting Apertus AI model performed well on certain multilingual benchmarks, but was far behind equivalent DeepSeek or Qwen models. Mistral, the French AI model, has $3b in funding, much of it from the private sector, but is still behind the Chinese models. There was widespread media reporting that DeepSeek V3 cost US$5.6M to train, but this likely reflected only the compute cost of the final training run and not pre-training compute, staffing, and other costs.

If Chinese models are built using illicit techniques, their performance may be impossible to replicate. US AI companies argue that Chinese companies use 'model distillation attacks' to extract capability from US models.[9] Chinese espionage may also target the AI industry.[10] If this is the case, it might be that an Australian fast chaser is not able to match Chinese performance levels unless there is unprecedented cooperation from frontier US labs.

The future market share of fast chasers is uncertain. Previously, Chinese fast chasers had a very small share of the Western market, but they have experienced rapid growth in the last few months.[11] This may suggest that these models have become 'good enough' to do a range of cost-sensitive tasks inside automated workflows, but it is not clear how much of the market will be satisfied with good enough rather than best available.

Australian copyright laws are a hurdle to a sovereign fast chaser. AI capability is a function of compute, data, and talent. Funding-limited fast chasers already have limited compute and talent. Australian copyright law would also limit their access to data and increase their overheads. For instance, computer code is one of the most valuable types of AI training data. In Australia, there's no collective licence for computer code. The Australian fast chaser would have to negotiate and license computer code on a case-by-case basis.

Australians adopting inferior models could be a national disadvantage. If an Australian fast chaser performs poorly, but Australian businesses are incentivised or obligated to use it regardless, they may underperform international competitors.

---

[8] ETH Zurich. (2025). *Apertus: A fully open, transparent, multilingual language model.*
[9] Anthropic. (2026). *Detecting and preventing distillation attacks*.
[10] United States Department of Justice. (2026). *Former Google engineer found guilty of economic espionage and theft of confidential AI technology*.
[11] Steinschaden, J. (2026). *Chinese AI models overtake US rivals in global token consumption*, Trending Topics.

06

Coordinate with middle powers on AI

Australia could use its influence and trusted relationships to coordinate middle powers on reducing AI risks and mitigating power concentration.

What's the case for it?

AI is on track to centralise a large amount of wealth and power into a small number of countries and companies. That degree of power concentration makes other countries significantly weaker economically and geopolitically. Smaller countries already struggle to protect their interests in relationships with large companies, and AI is on track to intensify that risk.

Middle powers have a lot to lose from AI, and a lot to gain by working together. Australia has a record of building coalitions and should use our influence and expertise to safeguard the interests of Australia, our allies, and like-minded nations through partnership with our peers.

Specific actions could include:

- **Host the 2029 AI Summit**. The UAE is hosting the 2028 Summit. No host for 2029 has been chosen. An Australian summit could focus on how all nations secure their interests in a world of increasingly capable AI, or on how red lines around the capabilities and uses of AI are agreed and enforced,
- **Coordinate with middle powers on supply terms across the AI value chain**, similar to the Minerals Security Partnership or the US–Japan–Netherlands lithography arrangement. Middle powers hold valuable positions in memory, lithography, semiconductor materials, minerals, and energy. Together, they are more likely to secure durable terms on issues such as access to frontier models and supply-chain security.
- **Pool resources with partners to develop fast-chaser AI models**. If AI model development is too expensive for Australia, a shared program with a small group of partners could help. More ambitiously, a large group of like-minded nations could jointly run a non-profit frontier AI company committed to sovereign access.[12] CERN, Airbus, or the Square Kilometre Array are templates.
- **Coordinate risk reduction among middle powers**. The Australia Group aligns export controls on chemical and biological weapons inputs across more than 40 countries. The same model could extend to AI risks, starting with AI-enabled biosecurity risks. Synthetic DNA screening standards, controls on biodesign tools, and evaluations of frontier model uplift would mitigate one of the most serious frontier AI risks and could scale as new risks emerge.
- **Contribute to chip controls and compute governance**. Compute is the most governable AI input. Chips are detectable, quantifiable, and produced through a concentrated supply chain.[13] Australia can contribute by preventing diversion of chips that are imported here, implementing know-your-customer standards for large-scale compute access, and verifying what data centres on Australian soil are running. This could also help build expertise for the AISI.
- **Lead AI standards development**. When a consumer product fails, standards help determine who in the supply chain is responsible. This allows courts to apply existing law to new technologies and contributes to a safety culture. Standards also deliver many of the benefits of AI regulation at a global scale without the downsides of unilateral domestic legislation.
- **Build agreement on AI red lines**. Australia endorsed the Hiroshima AI Process commitments. Australia could build on this by coordinating middle powers to urge global powers toward enforceable limits. Frontier AI development is already pressing against existing commitments, including the Hiroshima principles and the limits proposed in the Global Call for AI Red Lines.
- **Support a ceiling or slowdown as a contingency if red lines are crossed**. If frontier development moves toward systems that evade human control, create catastrophic risks, or cross other red lines, Australia should support coordinated limits. Australia supported the Hiroshima Guiding Principles but is failing to act as those lines are crossed, jeopardising our credibility on AI and other issues.

[12] Tertium AI. (n.d.). *Sovereign access to frontier AI*.
[13] Sastry, G. et al. (2024). *Computing power and the governance of artificial intelligence*.,

What are the hurdles?

Several of these actions are no-regrets and face few hurdles beyond resourcing — standards development, extending Australia Group-style coordination to AI biosecurity, and building compute governance expertise can begin now.

AI is moving faster than diplomacy. Without strong political leadership, coordination among governments can be slow. Today's AI issues sounded far away and speculative to many leaders only a few months ago. Proposals that hit the mark in the future may sound ambitious by today's standards but have to start before they're needed.

Working with partners can create friction. Shared programs spread costs but can move slower or be less focused. A pooled AI company would need governance designed for speed as much as for fairness. AI is already contributing to AI development, so catching up becomes harder each day.

Supplier coalitions are vulnerable to defection. Any individual country can be offered a better bilateral deal to leave the group. Coordination is more likely to hold in small, high-trust groupings around specific inputs than in broad declarations.

Coordination among middle powers will sometimes mean negotiating harder with allies, including the US. Allies negotiate terms constantly, and ideally the US should be in the tent, but Australia should be clear-eyed that better deals for middle powers are not costless to those currently setting terms.

07 Fund a world-leading AI Safety Institute

Australia could have a world-leading AISI that makes a meaningful contribution in one or more critical areas of technical AI safety, in ways legible to our allies and frontier AI companies.

What's the case for it?

Countries need the sovereign capability to evaluate, advise on, and mitigate frontier AI risks. Countries should contribute to technical research on how to make models safe — including preventing misuse and ensuring models are aligned with human values and under human control. Understanding and mitigating risks is relevant to industry adoption, effective governance, and public trust. In part because of the shortcomings of Australia's AISI, Australia has not been able to obtain pre-release access to frontier AI models to evaluate their risks and harden critical infrastructure.

The United Kingdom's AI Security Institute (UK AISI), by contrast, has successfully obtained pre-release access. Because of its technical expertise, it secured agreements with OpenAI and Anthropic to share frontier model access for evaluation before deployment. These collaborations allow them to understand the state of frontier technology, advance innovation and security in their AI industry, and prompt improvements in the most widely-used AI systems.[14]

Australia's AISI should develop world-leading capability in one or more areas of AI risk. This capability and contribution must be of the kind recognised and valued by key countries and companies. That contribution gives others a reason to share information with Australia in their areas of expertise.

What are the hurdles?

Australia's AISI is underfunded. The AISI receives an average of $7.5 million per year, with funding set to decline over time. This should increase to at least $50 million per year plus compute access, roughly matching the UK's world-leading AISI when adjusting for population or GDP.

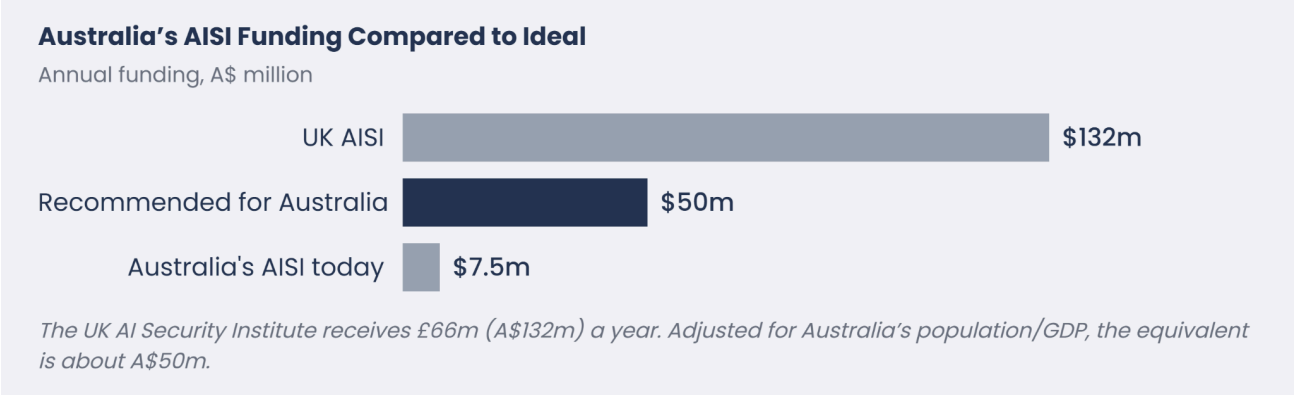

Australia's AISI hasn't developed expertise in a strategically relevant niche. AI risk is too broad for even a well-funded AISI to cover. Australia's AISI should deliver a work program that the US, UK, and frontier AI companies value.

Compute verification is a possible area of specialisation. This would let Australia see what foreign-operated data centres on Australian soil are running, and leading AI companies (e.g., Anthropic, OpenAI) are calling for tools that let them make verifiable commitments about when and how they train new models. Experts called for compute verification to be a focus area of the Australian AISI, alongside independent model evaluations and red-teaming.[15]

[14] AI Security Institute. (2025). *How we're working with frontier AI developers to improve model security*.
[15] Good Ancestors. (2026). *Designing Australia's AI Safety Institute: Expert survey report*.

08 Drive AI adoption across industry

Australia could accelerate the safe and responsible adoption of AI across domestic industry.

What's the case for it?

Highly capable AI is on track to upend the global economy. This period of change is likely to see many businesses fail, while others grow rapidly. If Australian businesses are relatively slow to adopt AI and relatively AI-illiterate, they're more likely to be losers than winners during this uncertainty.

By staying at the forefront of frontier AI adoption, Australian businesses will be better placed to see opportunities in their spheres of expertise and more confident in seizing them. Disruption will remain significant, but AI adoption will improve Australia's odds.

What are the hurdles?

Australia's current efforts at driving AI adoption, led by the National AI Centre (NAIC), are going in the right direction but are far too small in scale.

The NAIC's focus and scale are not proportionate to the AI future we are tracking towards. The NAIC's current work imagines AI as a significant, but largely non-revolutionary, technology. It does not imagine that it's preparing Australia for a transformation akin to the Industrial Revolution. The NAIC should anticipate transformative AI and deliver products suited to that future.

Australia's trust in AI is one of the lowest in the world.[16] Trust is a key hurdle to adoption. Broader work on Data Centre Standards, AI assurance and safety work, and a positive vision for Australia's AI future, is required to build trust and thus drive AI adoption across Australian industry.

Australia lacks committed access to the world's best AI models, and Australia has no capable domestic models. If Australia loses access to foreign models, Australian companies that have integrated AI into their operations may be left stranded. Australia needs to secure frontier access and to backstop that access with viable alternatives.

Currently, Australian copyright law puts AI use on an unclear legal footing. AI inference relies on exceptions in the *Copyright Act 1968* that may not hold for current AI models. Specifically, Section 43B of the *Copyright Act* only allows the reproduction of copyright works in technical processes (like AI inference) if the reproduction is 'temporary', 'necessary', 'incidental', and complies with other requirements, such as prohibiting subsequent use of the temporary reproduction. This provision was designed to cover temporary copies of text, images or videos in a computer's RAM, but there are clear legal arguments that it no longer fits sophisticated AI inference. If a court or influential legal advice finds that to be the case, it would undercut AI adoption in Australia and AI inference data centre builds.

Separately, the Prime Minister has said that using Australian books, music, art or news to train AI without the artist's control is "theft".[17] This is how models are currently trained overseas, and those are the models Australian businesses are being encouraged to adopt. Mixed messages — that AI models are akin to stolen property, but that Australian businesses should adopt them anyway — don't help AI adoption.

[16] Gillespie, N. et al. (2025). *Trust, attitudes and use of artificial intelligence: A global study 2025*. The University of Melbourne and KPMG.
[17] Albanese, A. (2026). *AI in Australia's interests*.

09

Invest in AI-proof industries

Australia could invest in industries unrelated to AI but that are likely to become much more valuable in a world with cheap intelligence.

What's the case for it?

Professional, financial, and administrative services generate a large share of Australia's income outside of resources. AI performing these tasks could be extractive, shifting jobs and profits offshore. Cheap intelligence could also make many goods and services abundant, with goods that stay scarce capturing a rising share of total spending. These scarce goods would be those that AI and robotics cannot increase the supply of because of physical scarcity, location requirements, or the need for a human to make or provide them.

Australia could anticipate which industries are "AI-proof industries of national significance" and build policy frameworks that support them and emphasise their value propositions.

In the 1970s, quartz technology made cheaper and more accurate watches than mechanical alternatives. Within 15 years, Swiss watchmaking employment fell by two-thirds.[18] *The industry recovered by pivoting to hand assembly, heritage, and Swiss origin. Mechanical watches now sell at prices unrelated to their function, and watchmaking is again a nationally significant Swiss industry. Policies like legally enforced origin claims supported the recovery.*

What are the hurdles?

Australia currently lacks a framework to help identify "AI-proof industries of national significance" and develop policy settings to foster them. Such a framework likely incorporates at least four factors:

1. **Physical scarcity / restricted reproducibility** — supply cannot be expanded by intelligence or automation. The underlying asset is finite or output-limited.
2. **Value in authenticity** — the buyer is paying for Australian provenance.
3. **Physical co-location** — the good or service must be consumed where it's produced.
4. **Australian significance** — Australia holds the asset exclusively or with a decisive advantage.

Examples might include premium wine from popular regions; reef and outback tourism; art; premium Australian food; and major sporting and cultural events. New industries are discussed separately (position 10).

While some industries will require transition support due to the negative impacts of AI, others should receive it because of positive impacts. Policy interventions for "AI-proof industries of national significance" could include:

- **Establishing and promoting a certified human-made mark** — a legally protected stamp, like Woolmark or "Swiss made", that producers can apply to goods and services made by people, so human work stands out as AI-made output becomes the default.
- **Extending the legal protection of regional names**, which currently applies only to wine, to a wider range of food, craft and tourism experiences, so that producers in more regions can trade on their region's name and reputation.
- **Establishing matched industry investment schemes for tourism and the creative industries**, modelled on agriculture's levy system, where sectors pay a levy on sales, the Commonwealth matches it, and industry directs the pool.
- **Funding additional protection of natural assets** like the reef and the outback through access charges on the visitors who use them.
- **Requiring that content meet set standards of being human-made to qualify for certain existing subsidies**, such as film and television tax offsets.

[18] Roulet, C. (2023). *A view of the Swiss watch industry*. Fondation de la Haute Horlogerie.

10 Host the new industries that AI will create

Australia could prepare for a large expansion in energy, land use, and industrial capacity so that the new industries AI creates are built here.

What's the case for it?

The Productivity Commission expects AI to lift labour productivity by around 4.3% over the next decade,[19] and Jobs and Skills Australia has noted that only 4% of Australian jobs have high exposure to generative AI.[20] The Commission's view reflects historical examples of general-purpose technology replacing some jobs but creating new and better jobs. Others expect the AI transformation to be larger and faster. One survey of economists found that, conditional on rapid progress in AI, annualised GDP will grow at around 4% a year, and labour force participation will fall to 55% by 2050.[21] Researchers at Stanford found that early-career workers in the US, in AI-exposed occupations, have already experienced a 19% decline in employment. Those at the forefront of AI developments suggest more dramatic changes may be possible. Dario Amodei said that AI could eliminate half of all entry-level white-collar jobs and push unemployment to 10–20% within one to five years, while the same systems drive economic growth of 10% a year.[22]

Australia has to think about where new jobs might be created. Previous technologies automated physical and routine work, and jobs shifted to cognitive work. AI automates cognitive work, so new jobs are more likely to appear in new industries that AI facilitates. AI-directed manufacturing is one example. New industries are likely energy- and land-hungry.

New industries will be built somewhere, but won't arrive in Australia by default. Data centres are the first wave of AI-driven physical demand. The second wave is industries built on the highly capable AI that is running in those data centres. Growth will concentrate in places that can supply the inputs the new industries need, like energy, land, and industrial capacity, and permit their use at speed. This could be an opportunity for Australia if our allies have capital and technology, but we have energy and space.

What are the hurdles?

The nature and speed of the jobs and industries created by AI is unknown. What we can do is prepare energy supply and grid capacity beyond forecast domestic demand, create industrial precincts with power, water, and transport, and fast approval pathways. These preparations pay off across a wide range of futures, including moderate ones — surplus clean energy and fast approvals are valuable even if AI does not cause explosive growth.

Hosting new industries is not the same as capturing value from them, particularly when the facilities are foreign-owned and automated. A facility of that kind could generate little domestic value beyond land rents, so converting Australia's land and energy into economic benefit and human wellbeing requires negotiating them as one package and will not happen by default.

New, large industries might face community resistance, as with data centres. Zoned remote precincts, genuine benefit-sharing with regions and traditional owners, and enforceable standards on water and land use are the difference between social licence and a decade of contested approvals.

---

[19] Productivity Commission. (2025). *Harnessing data and digital technology*, Australian Government.
[20] Jobs and Skills Australia. (2025). *Our gen AI transition: Implications for work and skills*, Australian Government.
[21] Karger, E. et al. (2026). Forecasting the economic effects of AI. FRB of Chicago Working Paper
[22] Vandehei, J., Allen, M. (2025). *Behind the Curtain: A white-collar bloodbath*, *Axios*; Brynjolfsson, E., Chandar, B. & Chen, R. (2025). *Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence*. Stanford Digital Economy Lab.

Estimate of Australian AI inference capacity and demand

We estimate that Australia currently hosts enough AI compute to serve around 39% of our own inference demand. This accounts for around 0.4% of global AI compute capacity (see below figure).[23] This attachment outlines how we reached that figure by estimating: 1) the amount of AI compute Australia currently hosts and 2) the amount of compute Australia currently demands.

Our estimate is based on the best information available, but that information has gaps, and our calculations rely on several assumptions. The difficulty of making this estimate suggests there is inadequate information about Australia's current and future compute supply and demand. Government should obtain this information.

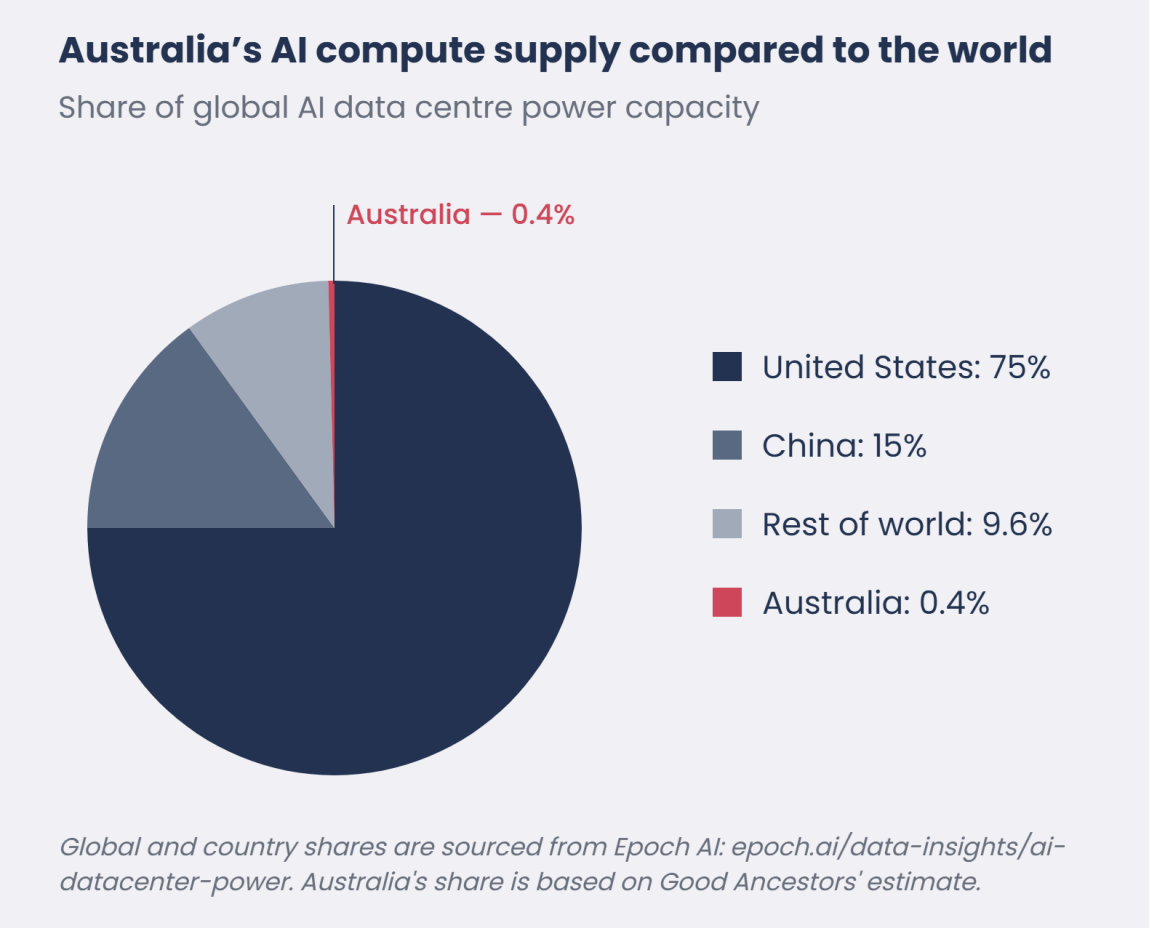

How much AI compute is in Australia?

We estimate that Australia hosts approximately 124 MW of nameplate capacity (range 90-160), which delivers around 87 MW of average output, rising to around 430 MW in nameplate capacity by late-2027.

Nameplate capacity is the power facilities would draw if every chip ran at full load. This is the standard way of measuring data centre capacity. We count all AI chips physically located in Australia, regardless of who owns them or who has contracted to use them.

Other than for hyperscalers, we only include facilities with details on how much capacity would supply AI accelerators. Around 1850 MW of data centre capacity is under construction in Australia, but many of those projects have not announced whether they will house AI accelerators.[24] If a substantial amount of that capacity does, our forecast of the self-supply ratio would be higher.

---

[23] Country shares are from Epoch AI's GPU cluster database (epoch.ai/data/gpu-clusters), which covers an estimated 10–20% of global compute. Epoch considers the distribution broadly representative. We apply those shares to Epoch's estimate of approximately 30 GW of total global AI data centre power capacity in late 2025.

[24] ITK Research. (2026). *Australia's ten largest data centres under construction*, citing Renewmap.com.au.

The table below outlines the projects which make up the current and predicted AI compute capacity, with ranges where operators have not published their chip counts:

Capacity	Status	Nameplate MW
AWS — Sydney and Melbourne regions; A$20bn program early in rollout — extrapolated from investment figures[25]	Operational	20 (10–30)
Microsoft — Sydney, Melbourne and Canberra regions; A$25bn program largely undeployed — extrapolated from investment figures[26]	Operational	17.5 (9–26)
Google — Sydney and Melbourne regions — estimate[27]	Operational	6.5 (3–10)
Oracle — Sydney, Melbourne and Canberra government regions, leased — estimate [28]	Operational	5 (2–8)
Groq — Equinix SY5 Sydney, dedicated LPU site[29]	Operational	4.5
Sharon AI — NEXTDC M3, Equinix SYD3/SYD5; ~2,500 GPUs incl. B200 and Blackwell Ultra[30,31]	Operational	4 (3–6)
Research HPC — Pawsey, NCI, CSIRO, universities; ~2,500–3,500 accelerators, mostly pre-Hopper — estimate	Operational	6 (4–8)
Enterprise on-prem and small neoclouds — ResetData, DUG, Macquarie/Dell, SambaNova partners — estimate	Operational	10 (5–15)
Firmus Southgate Melbourne — 18,500 GB300 contracted[32]	Scheduled Apr 2026; energisation unconfirmed	50
Total (Current)		**~90–160 (Mid 124)**
Firmus Tasmania — ~36,800 GB300[33]	Scheduled late 2026	~95
Sharon AI expansion — >55,000 GPUs total; 132 MW leased[34]	Scheduled mid-2027	~100–130
Ongoing hyperscaler rollout — AWS A$20bn and Microsoft A$25bn	Progressive to end-2027	~ 70-120
Total (2027)		~355–505 (Mid ~430)

[25] Amazon. (2025). *Amazon investing AU$20 billion to expand data center infrastructure in Australia and strengthen the nation's AI future*.
[26] Microsoft. (2026). *Microsoft deepens commitment to Australia with A$25 billion investment in AI infrastructure, security, and skills*; Bloomberg. (2023). *Microsoft to invest $3.2 billion to boost AI and tech in Australia*.
[27] Data Center Dynamics. (2026). *Google warns Aussie gov't that high taxes could prevent AU$20bn investment*; Google Cloud. (2026). *GPU availability by regions and zones*.
[28] Data Centre Magazine. (2025). *Oracle's data centre strategy: cloud, AI and more capacity*.
[29] Equinix. (2025). *Equinix supports Groq in launching low-latency AI inference in Australia.*
[30] Data Center Dynamics. (2026). Sharon AI signs $950m cloud deal with unnamed global tech company; Cisco. (2026). Sharon AI & Cisco launch Australia's first Cisco Secure AI Factory with NVIDIA.
[31] Cisco. (2026). *Sharon AI & Cisco launch Australia's first Cisco Secure AI Factory with NVIDIA*..
[32] Firmus. (2025). *Firmus expands Project Southgate across Australia*; Data Center Dynamics. (2025). *AI cloud firm Firmus raises AU$500m for data center roll out*; Firmus. (2026). *Firmus AI Cloud*.
[33] Firmus. (2026). *Firmus signs multi-year agreement with global hyperscale customer at Project Southgate*.
[34] SharonAI Holdings Inc. (2026). *Sharon AI announces six year strategic compute collaboration with NVIDIA*

How these figures were calculated

The megawatt figures above have mostly been calculated from announced GPU counts, as that is what companies typically report. Converting a GPU count into megawatts requires two estimates:

1. We estimated each GPU draws about 1.9 kW. An NVIDIA GB300 rack holds 72 GPUs and draws around 135–142 kW,[35] so dividing by 72 gives roughly 1.9 kW per chip.
2. We estimated each GPU accounts for about 2.4–2.6 kW, once facility power is included. A data centre draws more power than its chips alone, because it also runs cooling and electrical systems. These add roughly 25–35% on top of what the chips themselves use, taking each GPU from 1.9 kW to 2.4–2.6 kW.

To test these estimates, we applied them to Firmus, which has announced GPU counts for both its Australian sites. At 2.6 kW per GPU, its 18,500 GB300s in Melbourne come to approximately 50 MW and its 36,800 in Tasmania to approximately 95 MW. That totals 145 MW, close to Firmus's own figure of 150 MW for the two first stages combined.

Ranges reflect uncertainty about how many chips each operator holds and, therefore, how much power they draw. Firmus, Groq and Sharon AI have announced their deployments, so those figures are firmer.

The hyperscalers have not published country-level accelerator counts. We based estimates on announced capital programs and time horizons. Each provider has disclosed its Australian investment, but not what fraction of the overall investment is AI accelerators. We estimate the capital deployed so far and an accelerator share of 10–20% for older investment and 25–35% for newer investment. We estimate A$45–50 million per megawatt of all-in facility capacity. These figures imply Australia hosts well under 1% of each provider's global accelerator fleet, against a ~2% share of revenue — consistent with fleets concentrating in large US training campuses.[36] The forward row applies the same method to the remainder of the AWS and Microsoft programs through late 2027, minus the capacity already counted.

How much of that capacity is actually running

Nameplate capacity is a theoretical maximum, but no data centre operates every chip at full load continuously. At approximately 70% fleet utilisation, Australia's operating capacity delivers an average output of around 87 MW. This is about 15% of national data centre load, which AEMO estimated was 600 MW across the National Electricity Market in Q1 2026.[37]

[35] NVIDIA. (2026). System hardware and components, NVIDIA NVL72 AI Factory Enterprise Reference Architecture; Lenovo. (2026). Lenovo NVIDIA GB300 NVL72 rack scale AI product guide.
[36] Epoch AI. (2026). *Data on AI chip owners*.
[37] Westerman, D. (2026). AEMO CEO speech at Australian Energy Week 2026. Australian Energy Market Operator.

What's Australia's demand for AI compute?

We estimate Australia's AI inference demand at about 225 MW of average power draw (range 130-400). We used two methods, their central estimates are 220 MW and 230 MW, and adopted the average.

Demand here means the average power a facility would draw to serve inference for users in Australia, running at typical efficiency. We estimated this figure as it has not been published.

Method A — Australia's share of global inference

The first method starts with how much compute the world uses for inference, then estimates Australia's share of it. This requires four estimates:

- **Total AI compute installed worldwide: 30 GW**.[38] This includes all AI chips running globally, both for training and inference.
- **Utilisation of 65–75%.** Data centres do not run at full capacity continuously.
- **Inference accounts for about 65% of AI compute.**[39]
- **Australia accounts for 1.2–2.0% of global tokens.**[40] Tokens are the units of text AI models read and generate, so Australia's token share approximates the share of global inference compute serving Australian users.

Multiplying these gives a range of 150–290 MW. Using the middle of each range indicates about 220 MW.

Method B — Australia's share of global AI spending

The second method starts with global revenue from AI models, takes Australia's share, and works backwards to the compute that revenue must be paying for. This requires five estimates:

- **Global model-layer revenue: US $110–140 billion** a year.[41] This is what AI companies earn from selling access to their models.
- **Australia accounts for around 2% of that spending**, or US $2.2–2.8 billion a year.[42]
- **Inference compute accounts for 35–50% of an AI company's costs.** Applying this to Australian spending gives US $0.8–1.4 billion a year going to compute that serves Australian users.
- **Running one megawatt of accelerators costs US $7–9 million a year**,[43] including the chips, power, and facility. Dividing into the Australian share gives 86–200 MW of compute paid for by Australians.
- **Free users add 50–100% on top.** Most people using AI pay nothing, so paid demand understates total demand.

That gives a range of 130–400 MW, with a middle estimate of **230 MW.**

Methods A and B produced similar estimates, averaging **225 MW**.

---

[38] Edelman, Y., You, J., Somala, V., & Emberson, L. (2026). *Global AI power capacity is now comparable to peak power usage of New York State*. Epoch AI.
[39] Stewart, D. et al. (2025). *Why AI's next phase will likely demand more computational power, not less*, Deloitte.
[40] This is our estimate. No figure is published for Australia's share of global tokens. We start from Anthropic's finding that Australia accounted for 2.65% of global Claude.ai and Cowork usage in May 2026, ranking first of 121 countries per capita. We reduced this for two reasons. Australian usage skews away from coding and mathematics, which consume more tokens per conversation than other tasks. And Australia over-indexes on Claude relative to other providers, using it at 6.4 times the rate its population would predict, so its share of Claude usage may overstate its share of all AI usage. Anthropic. (2026). *Anthropic Economic Index report: Cadences*; McCrory, P. (2026). *How Australia uses Claude: Findings from the Anthropic Economic Index*. Anthropic.
[41] Assumption based on Epoch AI revenue tracking: Epoch AI. (2026). *Data on AI companies*.
[42] Assumption based on Australia's GDP, Australia's share of token usage where disclosed, and hyperscaler revenue reports.
[43] SemiAnalysis (2026) H100 Daily & Contract Pricing Index, SemiAnalysis.

How much of our own demand can Australia serve?

Today, compute located in Australia could meet ~39% of our own demand.

This is our best estimate, as calculated from 87 MW of supply / 225 MW of demand.

Announced builds will not keep pace with demand. Scheduled projects with confirmed AI accelerator installations would take supply to approximately 355–505 MW of nameplate capacity by late 2027, or around 250–355 MW of deliverable output. But global token demand grew roughly tenfold over the past year.[44] If Australian demand followed that trajectory, it would reach around 4,800 MW by late 2027,[45] and Australia's self-supply ratio would fall to between 5% and 7%.[46]

This is a simplified estimate. In practice, demand will fall as token prices rise in response to constrained supply. It also assumes a 1:1 relationship between token demand and electricity consumption, which likely overstates the physical capacity required, since efficiency gains may increase the amount of tokens servable per MW.

Australia has about 1850 MW of data-centre projects under construction that have not been expressly announced as containing specific AI accelerators.[47] Even on a generous assumption that 60% of that capacity supplied AI accelerators, **deliverable Australian supply would still only meet 21% to 24% of demand by late 2027**.[48]

Existing research

There's limited other analysis of self-supply today or into the future. One industry forecast estimates approximately 10.5 GW of total Australian IT demand by 2035, of which approximately 3.8 GW would be attributable to AI inference.[49] Accounting for power usage effectiveness, this corresponds to approximately 12.6–13.7 GW of total facility power demand, including 4.5–4.9 GW attributable to inference.[50]

Separately, an analysis prepared by Baringa for the Clean Energy Finance Corporation finds that Australia's rated data-centre capacity could grow by 4.7–7.4 GW by 2035.[51] On the conservative interpretation that this capacity is additional to the existing 1.35 GW baseline, total rated capacity would reach approximately 6.1–8.8 GW.

Overall, the existing research confirms that supply falls short of demand both currently and in the medium term. Governments may have further data to produce a more accurate assessment of self-supply.

---

[44] Emberson, L. & Sevilla, J. (2026). *Is a compute crunch coming?*, Epoch AI.
[45] 225 MW growing at 10x per year for the 16 months from August 2026 to December 2027 gives a multiplier of about 21.5 — tenfold in the first year, then a further 2.15x over four months — meaning it reaches approximately 4,800 MW.
[46] Nameplate capacity of 355–505 MW in late 2027 delivers 250–355 MW of average output at 70% utilisation. Against demand of 4,800 MW, that gives 250 ÷ 4,800 ≈ 5% and 355 ÷ 4,800 ≈ 7%.
[47] ITK Research. (2026). *Australia's ten largest data centres under construction*.
[48] 60% of 1,850 MW is 1,110 MW nameplate, giving 1,465–1,615 MW with announced builds, or 1,026–1,131 MW deliverable at 70% utilisation. Against demand of 4,800 MW, that is 21% to 24%.
[49] Idem Est Research & Advisory. (2026). Idem Est Research launches AI Demand & Connectivity Atlas, forecasting Australia's data centre IT load to reach 10.5 GW by 2035.
[50] At a power usage effectiveness of approximate 1.2, which aligns with the numbers given in Baringa. (2025). Getting the balance right: Data centre growth and the energy transition. Clean Energy Finance Corporation.
[51] Baringa. (2025). Getting the balance right: Data centre growth and the energy transition. Clean Energy Finance Corporation.

Published
20 August 2026

Authors
Greg Sadler, Emily Grundy, Daniel Marns

About Good Ancestors
Good Ancestors is an Australian charity dedicated to improving the long-term future of humanity by providing rigorous, evidence-based, and practical policy recommendations for Australia's biggest challenges. We have been deeply engaged in the AI policy conversation since our creation, working with experts around the world and helping to organise Australians for AI Safety.

Contact
If you would like to discuss the report or propose further research, please let us know at contact@goodancestors.org.au.