

Hosting AI Infrastructure Offers Australia More Than Economic Benefits

Submission to the inquiry into Artificial Intelligence and Data Centres

September 2026

Submitted

1 September 2026

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 this submission, please let us know at contact@goodancestors.org.au.



**Good
Ancestors**

Table of Contents

Summary of key points.....	2
AI is not normal technology.....	3
AI is matching or exceeding human performance in many domains.....	3
The technology will continue to improve.....	4
Greater capability brings greater risk.....	5
Data centre benefits go beyond economic value.....	6
Benefits of hosting AI inference compute.....	6
Benefits of hosting AI training compute.....	7
Data centre deals need to get energy and water conditions right.....	9
Sensible regulation can reduce negative energy and emissions impacts.....	9
Water use is contested and requires accurate data.....	10
Realising these benefits requires negotiating a copyright solution.....	11
Existing copyright law does not allow AI to be trained in Australia.....	11
Existing copyright law may not allow inference computing to occur in Australia.....	13
Conclusion.....	14
Attachment A: Building AI while Compensating Creators.....	15
Attachment B: What We've Heard So Far.....	25

Hosting AI Infrastructure Offers Australia More Than Economic Benefits

If done right, hosting AI data centres could offer Australia more than economic investment. Hosting inference compute would reduce our dependence on foreign infrastructure, and could help our regional partners do the same. Hosting training compute would give Australia a say over a technology set to reshape our economy and security. Australia should seize these opportunities to secure our national interests, including by being a net compute exporter, ensuring access to frontier AI models, bringing more AI development under Australian jurisdiction, and having mechanisms to manage the risks posed by AI.

Summary of key points

- AI is not normal technology. AI is poised to have far-reaching consequences for the economy and national security.
- Hosting AI compute is a valuable and achievable contribution to the global AI supply chain that will help Australia mitigate the risks of AI while seizing its benefits.
 - Meeting or exceeding domestic demand for inference compute would make us less dependent on foreign infrastructure and better able to support our regional partners.
 - Hosting AI training compute could secure access to frontier AI models, bring AI development under Australian jurisdiction, and provide a basis for cooperation with our allies.
 - Overall, the better Australia's position on the AI value chain, the more we can secure the wellbeing of Australians as AI disrupts the economy, and our broader national interests as AI reshapes geopolitics.
- AI data centres have real risks and downsides, including energy demand, emissions, community impacts, and water use. Policy should be based on accurate data. Some commonly cited figures, particularly on water use, are overstated or have since been corrected.
 - Data centre standards should be robustly legislated and implemented so risks and downsides are meaningfully addressed.
- Hosting AI training offers a uniquely valuable place in the AI value chain, but current copyright law does not allow frontier AI training to occur in Australia.
 - A negotiated solution that enables training while recognising Australian creators is essential to securing Australia's national interests.
- If Australia does not act, AI infrastructure will be built elsewhere, potentially with fewer conditions and safeguards. Australia will remain exposed to the risks of AI and be left with fewer means to manage risks or capture security and economic benefits.


AI is not normal technology

The case for hosting AI data centres requires understanding the capabilities, trajectory, and risks of the technology.

AI is matching or exceeding human performance in many domains

AI capabilities are improving rapidly. Most people are familiar with AI chatbots—often the free versions—but the frontier of the technology extends far beyond those systems. AI is approaching and surpassing human performance in a wide range of domains, including scientific reasoning, coding, cybersecurity, and persuasion.¹ OECD analysis suggests outcomes by 2030 could range from modest improvements to systems that match or exceed human performance on most or all cognitive tasks.²

Performance of Leading General-Purpose AI Models Over Time across Benchmarks



Data obtained from [EpochAI AI Benchmarking Hub](#).³

In biology, teams of generalist biologists using the Fable 5 AI model (not included in the above graph) outperformed teams of world-leading specialists. Work that would normally take experts 40–95 days was done in 16 hours.⁴

¹ Bengio, Y., et al. (2026). [International AI Safety Report 2026](#). Department for Science, Innovation and Technology.

² Hobbs et al. (2026). [Exploring Possible AI Trajectories Through 2030](#). OECD.

³ Benchmarks are standardised question sets used to track AI capabilities. The four shown are: GPQA Diamond (graduate-level science across biology, chemistry, and physics); FrontierMath (novel research-level mathematics problems, verified by expert mathematicians); MATH Level 5 (the hardest tier of mathematical-olympiad problems); and SWE-Bench Verified (real GitHub software issues a model must diagnose and fix). The most recent models, including Fable 5, Mythos, and GPT-5.5, are not depicted.

⁴ Anthropic. (2026). [Claude Fable 5 and Claude Mythos 5](#).

In cybersecurity, Anthropic's Mythos Preview model (not included in the above graph) found thousands of high-severity flaws across the world's major operating systems and web browsers.⁵ The model was judged too dangerous to release publicly and was restricted to a small group of organisations. Models from other companies have similar cyber capabilities. The UK AI Security Institute found OpenAI's GPT-5.5 has a "similar level of performance" to Anthropic's model on cybersecurity related tasks.⁶

AI cybersecurity capabilities are being used against governments. In July 2026, suspected China-linked attackers ran a near-autonomous AI cyberattack against the Taiwanese government.⁷ The agents mapped 21 Taiwanese government systems, hacked into 85 user accounts, and extracted 2,500 personnel records. The attackers reached a government email system, Taiwan's nuclear safety agency, and at least seven energy companies.⁸ In June 2026, the Five Eyes warned that frontier AI will transform cyber risk in the coming months, and urged leaders to act now.⁹

Models have also begun to act in unintended ways, including breaking into systems they were not authorised to access during testing. In July 2026, OpenAI disclosed that two of its models broke out of a secure testing environment, obtained internet access by exploiting a zero-day vulnerability in third-party software, and hacked into the AI company Hugging Face. Over a thousand individual AI agents collaborated on a message board in an attempt to cheat on the evaluation they were being tested on.¹⁰ The models took roughly 17,600 actions over four days to hack Hugging Face and steal the answers to the test.¹¹ Anthropic, Meta, and the UK AI Security Institute have separately reported that their models, during testing, hacked computer systems.¹²

The technology will continue to improve

Models will get better at executing long-term tasks. The length of the tasks agents can complete unaided has been doubling every seven months.¹³ By 2030, if the trend continues, agents will be able to autonomously complete month-long software tasks.¹⁴

AI is also being used to build itself. Anthropic reported that, as of May 2026, its model wrote more than 80% of the code merged into its systems. This had risen from low single digits sixteen months earlier.¹⁵ The leading AI companies are explicitly working towards recursive self-improvement — AI that's able to autonomously design and develop its own successor.¹⁶ If they succeed, and each model builds a more capable one, progress could accelerate sharply.

⁵ Anthropic. (2026). [Project Glasswing: Securing critical software for the AI era](#).

⁶ UK AI Safety Institute. (2026). [Our evaluation of OpenAI's GPT-5.5 cyber capabilities](#).

⁷ Blanchard, B. & Satter, R. (2026). [AI-driven hacking campaign targets Taiwan government agencies](#), Taipei Times.

⁸ Dream. (2026). [Inside a multi-agent AI framework used to compromise government entities in Asia](#), DREAM Lab Threat Research.

⁹ Australian Signals Directorate. (2026). [Five Eyes cyber security agencies statement](#).

¹⁰ METR & Redwood Research. (2026). [Brief independent investigation of agents' behavior, reasoning and collaboration in the OpenAI / Hugging Face hacking incident](#).

¹¹ Larcher, H. et al. (2026). [Anatomy of a frontier lab agent intrusion: A technical timeline of the July 2026 incident](#). Hugging Face; OpenAI. (2026). [Hugging Face model evaluation security incident](#).

¹² Anthropic. (2026). [Investigating three real-world incidents in our cybersecurity evaluations](#); AI Security Institute. (2026). [Unsanctioned agent behaviour during cyber evaluations](#).


¹³ Kwa et al. (2025). [Measuring AI Ability to Complete Long Tasks](#). METR.

¹⁴ Ibid.

¹⁵ Favaro, M. & Clark, J. (2026). [When AI builds itself](#). Anthropic Institute.

¹⁶ Favaro, M. & Clark, J. (2026). [When AI builds itself](#). Anthropic Institute; Altman, S. & Pachocki, J. (2026). [Built to Benefit Everyone: Our Plan](#). OpenAI.

The scale of investment reflects these expectations. The computing power used to train leading models has increased about fivefold each year, and developers expect the trend to continue.¹⁷ Companies are investing heavily into AI infrastructure, spending over US\$650 billion this year alone.¹⁸ In 2025, Nvidia's revenue from AI chip sales alone was four times larger than Australia's defence budget.



Data obtained from EpochAI's [Data on AI Chip Sales](#), and [SIPRI's Military Expenditure database](#).

We don't know where this scale of investment leads, but the consequences for the economy and national security will be far-reaching. Australia needs a plan for securing our national interests in the face of AI.

Rejecting data centre investment will not protect Australia from the risks and harms of AI. Instead, hosting AI infrastructure can allow us to gain oversight over its development. Good Ancestors recently analysed ten positions Australia could pursue to secure our national interests in a world with highly capable AI, assessing each on value and feasibility. Hosting AI infrastructure underpinned several of them.¹⁹

Greater capability brings greater risk

Advanced AI can pose serious risks. Many of these risks are present today. As outlined by the International AI Safety Report, risks can be placed into three categories:²⁰

- **Being deliberately misused.** For example, being used to find and exploit software cybersecurity vulnerabilities, or assist with the development of biological weapons.²¹
- **Malfunctioning, or behaving in ways their developers and users did not intend.** For example, in a growing number of cases AI chatbots have encouraged and coached users to end their own lives or kill other people.

¹⁷ Bengio, Y., et al. (2026). [International AI Safety Report 2026](#). Department for Science, Innovation and Technology.

¹⁸ Bratton, L. (2026). [Big Tech Set to Spend \\$650 Billion in 2026 as AI Investments Soar](#). Yahoo Finance.

¹⁹ Good Ancestors. (2026). [Australia's position in the age of AI](#).

²⁰ Bengio, Y., et al. (2026). [International AI Safety Report 2026](#). Department for Science, Innovation and Technology.

²¹ Anthropic. (2026). [Claude Fable 5 and Claude Mythos 5](#).

- **Broader harms to society.** For example, labour-market disruption, concentration of power, and eroding trust in information.

The Taiwanese government hacking incident and the Hugging Face incident are examples of the first two categories. As investment continues and capabilities improve, the risks will grow. Unlike some risks and harms Australia manages with geographic distance, AI models being developed and operated offshore offers Australia no protection. Instead, it leaves us more exposed.

Data centre benefits go beyond economic value

Discussion about the benefits of hosting AI infrastructure often focus on headline investment figures. But, significant value comes from how it can give Australia a say over AI development, and opportunities to secure our national interests and protect Australians. Hosting AI training compute and AI inference compute offer different benefits.

- Hosting enough **inference compute** to meet domestic and regional demand would reduce our, and our regional partners', dependence on foreign alternatives.
- Hosting **AI training** in Australia could unlock a degree of scale and influence that inference alone cannot match. In a recently published report, Good Ancestors assessed frontier AI training as the highest value action Australia could take to establish its place in an AI future.²²

Benefits of hosting AI inference compute

Reducing foreign dependence

Our current best estimate is that **Australia's inference compute supply meets less than half of domestic demand**, and Australia only hosts about 0.4% of global compute supply.²³ Future estimates have uncertainty about supply, demand, and how they interact over time. Even considering announcements of new AI data centres, we assess that **Australia will meet less than a quarter of its domestic inference demand by the end of 2027**. This is because infrastructure is not keeping up with rising inference demand. The best global estimate is that inference compute supply is growing at 3x per year, but demand is growing at 10x per year.²⁴

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 economic and security risks. Countries that control our inference compute could drive up costs, remove access, or exploit their access to create vulnerabilities (e.g., substitution of the served model for a backdoored variant, or the extraction and leakage of prompts, outputs, and weights).²⁵ Access could be restricted through export controls, sanctions, or prioritising users in the host country — particularly as compute constraints grow.

Hosting domestic AI inference compute would help mitigate these risks. It would keep sensitive workloads onshore and under Australian jurisdiction, and would help critical services continue running if access to foreign compute is restricted.

²² Good Ancestors. (2026). [Australia's position in the age of AI](#).

²³ Ibid.

²⁴ Emberson, L. & Sevilla, J. (2026). [Is a compute crunch coming?](#) Epoch AI.

²⁵ Cai, W. et al. (2025). [Are you getting what you pay for? Auditing model substitution in LLM APIs](#). arXiv:2504.04715

Exporting inference compute to regional partners

Additional compute would position Australia to serve regional partners, for whom hosting compute at national scale is not viable. Our neighbours in the Pacific and parts of Southeast Asia may lack the power generation, grid capacity, land, or capital to build sovereign AI infrastructure. However, their governments, banks, hospitals, and utilities will still want to adopt AI. That demand has to be served from somewhere.

Australia is politically stable and aligned with regional partners through existing security arrangements. Serving that demand would reduce our partners' dependence on alternative, likely Chinese-supplied, AI infrastructure,²⁶ and it would give Australia standing in how AI is governed across the region.

Benefits of hosting AI training compute

Securing access to frontier models

Few countries will host AI training, because it involves large, concentrated amounts of compute and power in a single location and does not need to be near users. Hosting training in Australia would strengthen our alliances and provide a basis for negotiating terms in our national interest. One potential benefit Australia could negotiate is continued access to frontier AI models, which matters for national security, the economy, and public-interest research.

For national security, effective defence will require AI at least as capable as that available to Australia's adversaries.²⁷ An adversary using more capable AI, including for cyber operations, intelligence, or autonomous systems, would be at an advantage. We cannot realistically build Australian frontier models to rival those of the United States or China, so we will depend on access to models developed elsewhere. Without access, Australia cannot harden systems of national significance, from Medicare to the power grid, before adversaries use the same capabilities against us.

For the economy, those using more capable AI models are likely to outcompete those using less capable AI models. As AI drives more output, the economic gap between “haves” and “have-nots” could grow. If Australia wants to create valuable systems and products built on AI models, as the National AI Plan envisages, we need Australians to have access to capable AI models.

For government and research, the Australian AI Safety Institute needs frontier-model access to evaluate capabilities and risks, and Australian universities and research institutions need access to do high-quality research and build local expertise. Both also need compute, which is becoming scarcer and more expensive as demand grows.

Even as small and narrow AI models find niches, frontier models will continue to be of particular importance in many fields.

Continued Australian access to frontier models is not guaranteed, as shown with Mythos and Fable 5. Anthropic's Mythos cybersecurity model was initially only released to a restricted group of mostly US-based organisations. This did not include the Australian Government. When a public version was released (Fable 5), it was available for only three days before a US Government export directive required Anthropic to suspend access for foreign nationals. As AI capabilities advance, and the demand for compute to train frontier models grows, compute constraints, security concerns, and US strategic interests will all make access harder to rely on. Australia currently has no plan for how to secure such access and few valuable contributions to bring to our alliances.

²⁶ For discussion of this concern, see, e.g., Forbes, K. & Maughan, C. (2026). [ASEAN must strike a balanced AI relationship with China](#).

²⁷ Phillips-Robins, A., Tawil, T. & Winter-Levy, S. (2026). [The compute coalition: How to build the future of AI in the free world](#). Carnegie Endowment for International Peace.

Hosting AI training compute would bring a significant place in the AI supply chain and strengthen Australia's position to negotiate continued access. Companies that have invested billions into building large training clusters here, and that depend on Australian infrastructure and regulatory cooperation, have reason to support access—as would the US government. Australia should seek access at three levels:

- 1) market access for the public, equivalent to the US,
- 2) access to restricted or sensitive models for critical infrastructure operators and other defenders, and
- 3) Government access, equivalent to the US government for national security purposes.

Securing oversight of AI development

Benefits could extend beyond model access. Australia could ask AI companies hosting infrastructure here to allow the AI Safety Institute to test models before deployment, disclose the capabilities and risks of the models they train, and report critical safety incidents. Australia could adopt some of these requirements through regulation independent of hosting, as New York and the EU do.²⁸ But Australia is a smaller market, and companies may be more likely to comply where they have invested a significant amount and rely on continued regulatory cooperation.

Hosting also brings the infrastructure under Australian jurisdiction. Some obligations only attach to infrastructure physically located here, like the security of the facility itself and oversight of training as it happens. The *Security of Critical Infrastructure Act 2018* is one example. The Act requires critical infrastructure assets located in Australia, including data storage and processing assets, to manage their risks and report cyber security incidents.²⁹ The Department of Home Affairs is currently reviewing amendments that would ensure AI data centres are covered by the Act.³⁰

Australia could also build our capability to oversee AI development, rather than relying only on information we obtain from companies. 'Compute verification' refers to a set of techniques for checking whether a facility is training or running models, whether it is running the model its operator claims, and whether it is complying with agreed conditions. Verification would let the Government independently confirm what data centres on our soil are doing, without operators exposing commercially sensitive information. The Australian AI Safety Institute (AISI) is well placed to develop this capability. No allied AISI currently specialises in it, it draws on Australian strengths in cybersecurity and cryptography, and it would be a distinctive contribution to global AI governance—since any future international agreement on frontier AI would depend on these verification tools.

Recommendation 1: Ensure Australia can meet its own inference demand and serve the region.

Australia currently meets less than half of its domestic inference demand, and is on track to meet less than a quarter of demand by the end of 2027. The Australian Government should:

- Measure Australian inference compute supply and demand on an ongoing basis.
- Set a target for Australia to become a net exporter of inference compute, and use the target to guide decisions on energy supply, grid connection timelines, and planning approvals.

²⁸ EU AI Act (Regulation (EU) 2024/1689), article 55; New York State Legislature. (2025). *Responsible AI Safety and Education (RAISE) Act* (S6953-B).

²⁹ Parliament of Australia. (2018). *Security of Critical Infrastructure Act 2018*.

³⁰ Department of Home Affairs. (2026). [Proposed amendments to streamline and modernise the SOCI Act 2018](#).

Recommendation 2: Attract AI training to Australia and apply national interest conditions.

Australia should adopt the policy settings necessary to attract frontier AI training compute to Australia. Leveraging our national strengths, like abundant land and energy, security and stability, Australia could secure commitments including:

- Ongoing access to frontier AI models including public/market access equivalent to the US, access to restricted or sensitive models for critical infrastructure operators and other defenders; and Australian Government access for national security purposes.
- A portion of compute allocated for research and innovation on favourable terms (e.g., to Australian researchers, not-for-profits, and start-ups).
- Oversight over AI development in Australia, including pre-deployment testing of models trained in Australia by the Australian AI Safety Institute, disclosure of model capabilities and risks, and reporting of critical safety incidents to Government.

Other potential terms could include requiring some AI engineers and researchers to be based in Australia, ensuring AI data centres have high security standards, notifying Government of large training runs conducted in Australia, and ensuring fair contributions are made via taxation and other mechanisms.

Recommendation 3: Establish a compute verification workstream in the Australian AI Safety Institute.

The Government should direct the AISI to develop technical expertise in compute verification, including in distinguishing training from inference in Australian facilities and confirming compliance with conditions attached to data centre approvals.

Data centre deals need to get energy and water conditions right

While hosting AI data centres can bring strategic benefits, it also raises concerns about energy use, emissions, and water. Addressing these impacts requires accurate evidence and appropriate conditions.

Sensible regulation can reduce negative energy and emissions impacts

Data centre electricity demand is rising. Australian data centres currently consume about 2.8% of national electricity.³¹ The AEMO estimates this will grow to about 13% by 2036.³² Without additional renewable generation and storage, data centre growth could raise wholesale electricity prices by around 26% in New South Wales and 23% in Victoria by 2035.³³

Legislated data centre standards should address these impacts. The Expectations ask operators to secure new and additional clean energy generation to offset demand, cover grid upgrade costs, and avoid passing demand-driven price increases to consumers.³⁴ There is precedent for AI companies meeting these conditions. In February 2026, Anthropic committed to cover 100% of grid infrastructure upgrade costs and offset consumer price increases.³⁵ Amazon, Google, Meta, Microsoft, OpenAI, Oracle, and xAI have since signed a similar pledge.³⁶

³¹ Australian Energy Market Operator. (2026). [2026 ES00 Data Centre Forecasting Overview](#).

³² Ibid.

³³ Clean Energy Finance Corporation (CEFC). (2025). [Getting the balance right: Data centres and the energy transition](#).

³⁴ Department of Industry, Science and Resources (DISR). (2026). [Expectations of data centres and AI infrastructure developers](#). Australian Government.

³⁵ Anthropic. (2026). [Covering electricity price increases from our data centers](#).

³⁶ White House Office. (2026). [President Trump secures historic commitment to keep electricity costs down amid data center boom](#).

Another option is to build behind-the-meter solutions and microgrids, where operators generate their own power and connect directly to renewable sources rather than drawing from the NEM or other grids.³⁷ This could bypass many grid challenges altogether, including additional infrastructure costs and grid stability pressures.

Building AI infrastructure in Australia could have a positive impact on Australian renewables. Grid operators have said that large, predictable demand backed by long-term renewable contracts can help underwrite new generation rather than strain the grid.³⁸ This could help de-risk clean energy investment and accelerate Australia's energy transition.

At a global scale, data centres in Australia may be *beneficial* for emissions. Currently, global chip supply is constrained,³⁹ meaning the limiting factor for new data centre construction is the chips that data centres contain. That means new data centres built in Australia are likely displacing data centres that would otherwise be built elsewhere. Australia should build clean data centres that displace data centres in jurisdictions less committed to clean energy.

Water use is contested and requires accurate data

The relationship between AI and water use is contested. Analysis can vary widely based on scope, assumptions, and comparisons. Most of a data centre's water footprint is indirect, meaning the water used by power plants to generate its electricity, rather than water used on-site for activities like cooling.⁴⁰ Clean energy requirements, like those proposed for data centre standards, will reduce indirect water consumption. Wind and solar energy, for example, do not use water in generation.⁴¹

Direct water usage in cooling can also be reduced. All new Microsoft data centre designs use cooling systems that recycle water through a closed loop, which use little or no water once filled.⁴² Anthropic has indicated it will use water-efficient cooling in Australian facilities.⁴³ The Expectations of Data Centres and AI Infrastructure Developers require efficient cooling and non-potable water sources where possible.⁴⁴

Public beliefs that AI uses an inordinate amount of water often rest on widely cited figures that are incorrect, sometimes by orders of magnitude. One example is the widespread claim that one AI prompt consumes a bottle of water. This figure originated in a 2024 *Washington Post* graphic and spread rapidly through news coverage.⁴⁵ However, this claim is based on a number of contested assumptions. The authors of the academic work underpinning the figure have since acknowledged that the likely water use is substantially lower – closer to 15 mL per-prompt.⁴⁶ Other analyses suggest it's closer to 2mL (including direct and indirect usage).⁴⁷ For context, a single sheet of paper takes about 2,550 prompts' worth of water to produce.⁴⁸

Another claim, from a best-selling book, was that one proposed data centre would use more than a thousand times a Chilean town's annual water usage. This was found to be overstated by a factor of approximately

³⁷ Fist, T. & Datta, S. (2024). [How to Build the Future of AI in the United States](#). Institute for Progress.

³⁸ Wiggins, J. & Cropp, R. (2026). [Grid giants back data centre boom to bolster renewable build-out](#). Australian Financial Review

³⁹ Sanders, J. et al. (2026) [American AI Companies Can't Get Enough of Chips](#). Center for a New American Security.

⁴⁰ International Energy Agency. (2025). [Energy and AI](#).

⁴¹ Mytton, D. (2021). [Data centre water consumption](#). *npj Clean Water*, 4(11).

⁴² Microsoft. (2024). [Sustainable by design: Next-generation datacenters consume zero water for cooling](#).

⁴³ Anthropic. (2026). [Evidence to the Senate Select Committee](#), Proof Committee Hansard.

⁴⁴ Department of Industry, Science and Resources. (2026). [Expectations of data centres and AI infrastructure developers](#). Australian Government.

⁴⁵ Verma, P. & Tan, S. (2024). [A bottle of water per email: The hidden environmental costs of using AI chatbots](#). The Washington Post.

⁴⁶ Masley, A. (2026). [I think I figured out exactly how the "AI uses a bottle of water per prompt" miscalculation happened](#).

⁴⁷ Elsworth, C. et al. (2025) [Measuring the environmental impact of delivering AI at Google Scale](#), Google.

⁴⁸ Water Footprint Calculator. (2022). [The hidden water in everyday products](#).

1000, leading to a published correction by the author.⁴⁹ These examples show that evidence-based analysis will be needed to assess the water impacts of building AI data centres in Australia.

Data centres require careful local planning—like other large industrial users—but their water use is smaller than often assumed. Further, the water use of data centres is best compared to other industrial uses. By one estimate, a megalitre of water generates several hundred times more economic value in the industries that include data centres, than it does in agriculture.⁵⁰ As with energy, setting effective standards requires accurate data on how and where water is used.

Recommendation 4: Ensure the new data centre standards impose sensible, evidence based requirements on energy and water use.

- Policy on the energy and water impacts of data centres should be based on accurate data about the source, scale, and type of energy and water used.
- Legislated data centre standards should include conditions associated with additional renewable supply, coverage of grid-upgrade costs, efficient cooling, and water sourcing. The Government should not rely on voluntary compliance.

Realising these benefits requires negotiating a copyright solution

Australia must be able to secure its national interests and protect Australians from the harm of increasingly capable AI. Getting the energy and water conditions right helps address the immediate downsides of hosting AI data centres, but it does not manage structural risks from AI and it secures few of the benefits. Hosting AI training and inference in Australia might be the best way for Australia to secure a valuable position in the AI value chain that we can use to protect Australia from the broader risks of AI.

Existing copyright law does not allow AI to be trained in Australia

Under current Australian copyright law, training frontier AI models here is not viable. Australia has no fair use or text-and-data-mining exemption permitting AI models to be trained on copyrighted material, unlike the United States, the European Union, Singapore, Israel, or Japan.⁵¹ As a result, frontier models are trained overseas rather than in Australia. Australian creators receive little or no compensation when their work is used to train AI overseas, and Australia forgoes the national interest benefits that hosting training could bring. This same barrier also prevents Australian AI companies from developing competitive frontier general-purpose AI models.⁵²

Under current Australian copyright law, AI companies and rights holders must strike licensing deals covering the entire global corpus of works used to train AI. A successful licensing path depends on many agreements between any given AI company and a number of rights holder organisations who manage collective licences and reciprocal arrangements as well as agreements for works not covered by rights holder organisations. Rights holders and AI companies do not see eye-to-eye on licensing, meaning this is not on track to happen.

⁴⁹ Hao, K. (2025, December 17). [Empire of AI: Water chapter corrections](#).

⁵⁰ Vardon, M. (2026). [How much water and power will AI data centres use in Australia?](#) The Conversation.

⁵¹ United States: Copyright Act 1976, 17 U.S.C. §107; European Union: Directive (EU) 2019/790, art. 4; Singapore: Copyright Act 2021, s. 244; Japan: Copyright Act 1970, art. 30-4; Israel: Copyright Act 2007, s. 19, 22, 26, and the Ministry of Justice. (2022). [Opinion: Uses of Copyrighted Materials for Machine Learning](#).

⁵² The CEO of Maincode has told a Senate committee that respecting copyright, and implementing data-provenance conditions, means its model Matilda cannot perform as well as other models in certain categories; Lemphers, D. (2025) [Evidence to the Senate Environment and Communications References Committee inquiry into the National Cultural Policy](#).

Reasons AI companies are not willing to license include:

- **Licensing all material seems infeasible to them.** While many economically valuable works (like commercial music, movies and published books) are subject to collective licences, AI companies believe there are many other works (like the vast amount of text, code, and images individuals have shared online) that are part of AI training datasets but not subject to existing collective licences or reciprocal arrangements. For example, there appears to be no collective licensing arrangements for video-game or software code. Without a pathway to lawful AI training on all material, AI companies don't see the value in negotiating to licence some material.
- **Licensing could undermine AI companies' international legal position.** AI companies rely on US case law holding that AI training is fair use.⁵³ Whether a use is 'fair' depends partly on its effect on the market for the original work, including whether a licensing market for that use exists.⁵⁴ AI companies likely worry that if a market for licensing materials for AI training were established in Australia, it would undermine the fair-use argument they rely on in the United States. Rights holders are likely also aware of this, and hope to strike licensing deals precisely to use those licensing deals to rebut precedent in US courts.
- **They are reluctant to pay here for what's free elsewhere.** AI companies may struggle to justify paying for works in Australia that are no or low cost in the US, Singapore, Japan, or elsewhere.
- **AI companies have alternative options.** An AI training data centre buildout does not need to happen in Australia. Companies could double down on the US, expand AI training to Japan where there is already a permissive copyright regime and available nuclear energy, or wait for copyright clarity in other countries, like Canada.
- **Many licensing or data access deals may not work in Australia.** Overseas, AI companies have signed deals with large news organisations like News Corp, covering US-based AI training on Australian mastheads.⁵⁵ Such deals may not be practical in Australia. News articles often contain third-party copyright material, like quotes and images, that the outlet doesn't own. In the US, fair use lets publishers include that material and pass it on in AI deals. In Australia, outlets rely on "fair dealing for the purpose of reporting news" to include it. But that exception likely only permits the use of it in news reporting. It gives the outlet no rights to on-license that content to an AI company.⁵⁶

Reasons rights holders are reluctant to accept any alternative include:

- **Many expect to win in US courts.** Many rights holders believe that they are on track to win ongoing copyright litigation cases in the US. If that happens, they believe they would secure compensation, and potentially damages, for AI trained in the US. This leaves them more inclined to "wait and see" what happens in the US, rather than negotiate to facilitate training here.
- **Many oppose AI outright.** Many rights holders harbour objections towards AI and the companies that develop it, making them less inclined to negotiate. They may value being a roadblock to AI progress more than they value any deal.

This is not to blame rights holders or AI companies for the absence of a deal. Rather, it illustrates that each party operates under incentive structures that make a deal unlikely. This leaves Australia's national interest dependent on a negotiation between two third parties that is unlikely to progress. For that reason, the Australian Government should intervene to facilitate a deal that allows AI training while supporting Australia's national interests, including the interests of Australian creators.

⁵³ *Bartz v. Anthropic* (Alsup, ND Cal, June 2025), *Kadrey v. Meta* (Chhabria, ND Cal, June 2025).

⁵⁴ Copyright Act (1976), 17 U.S.C. §107(4)

⁵⁵ News Corp. (2024). [News Corp and OpenAI sign landmark multi-year global partnership.](#)

⁵⁶ See, eg, the reasoning in *De Garis v Neville Jeffress Pidler Pty Ltd* (1990) 37 FCR 99.

Good Ancestors has proposed one approach in a consultation paper that's currently open to submissions (**Attachment A**).⁵⁷ The Australian AI Training Scheme consists of two separate components: an AI Training Permit and an Australian AI Public Benefit Fund. The AI Training Permit would give companies legal certainty to train in Australia in exchange for a fee paid to Government, an opt-out for Australian rights holders, and compliance with other national interest conditions. An Australian AI Public Benefit Fund, modelled on the *Public Lending Right Act 1985*, would support creators with flexibility in quantum and design.⁵⁸

Keeping the permit and the fund decoupled avoids creating a de facto licensing market, which is what makes the approach workable for AI companies while still delivering creators consent and financial recognition. The same pathway would also benefit Australian AI development. One of the key benefits of the proposal is that it provides Australia with a concrete mechanism to manage the broader risks and harms that Australians are already experiencing from AI and that are on track to intensify.

Creators, rights holders, technology companies, and others are invited to share their views on the proposal – to strengthen it, or to find better alternatives. A summary of initial consultation outcomes is in **Attachment B**.

Polling shows a majority of Australians back some sort of copyright reform over losing AI investment. In April 2026 polling, **61% of Australians said they would prefer to enable AI training under a different copyright arrangement that still supports creators, against 15% who would keep current laws even if it meant AI companies trained overseas**. Requiring AI companies to pay to train in Australia, and having a government-run fund that supports Australian creators, drew 57% support, with 18% opposed.⁵⁹

Existing copyright law may not allow inference computing to occur in Australia

Independent of AI training, copyright reform may be necessary to enable current AI use. Australian AI data centres mostly do inference computing. During inference, copyrighted material gets copied. This can include the AI model accessing works, like news or reports, and storing them in memory. Most assume this does not infringe copyright because section 43B of the Copyright Act allows a 'temporary' reproduction of a copyright work where the copy is a 'necessary' and 'incidental' part of a technical process.

However, making copies during inference may not be temporary, necessary, or incidental. Models can now hold copyrighted material in 'context windows' for extended periods. 'Cross-chat memory' lets models recall material from previous conversations. Longer 'chains of thought' allow models to reason better about copied material. These design choices are made to improve performance, and are perhaps not a necessary or incidental part of AI technology.⁶⁰ Rapid technological developments, like 'continual learning' where a model's weights shift in response to user behaviour, could further blur the boundary between use and training.

The consequences of inference being found to infringe copyright in Australia would be severe. AI use and data centre investment rest on the assumption that the day-to-day use of AI does not infringe copyright. If that assumption were overturned, adoption would stall and ongoing investment would be undercut.

⁵⁷ Good Ancestors. (2026). [Building AI while Compensating Creators: Consultation Paper on an Australian AI Training Scheme](#).

⁵⁸ *Public Lending Right Act 1985* (Cth).

⁵⁹ Good Ancestors. (2026) [New polling shows a majority of Australians back copyright reform over losing AI investment](#).

⁶⁰ Wei, J. et al. (2022). [Chain-of-thought prompting elicits reasoning in large language models](#); Denain, J.-S. & Barry, A. (2026). [Have AI capabilities accelerated?](#) Epoch AI; Epoch AI. (2026). [Trends in artificial intelligence](#).

Recommendation 5: Develop copyright pathways that enables AI training and inference to occur in Australia while supporting creators.

The benefits of hosting AI training depend on AI training being viable here. Existing licensing approaches are not on track to enable AI training.

- Government should progress an Australian approach—not a text-and-data-mining exception—that supports rights holders while allowing training to occur in Australia
- Good Ancestors' consultation paper sets out one potential model, but other creative and uniquely Australian approaches may also be possible.

Recommendation 6: Clarify the position of inference under Australian law and ensure copyright law is fit for purpose.

The Government should seek legal advice on whether AI inference falls within section 43B of the *Copyright Act 1968*, and should clarify its treatment in any copyright reform rather than leaving the question to litigation.

Conclusion

Opting out of hosting data centres would not reduce Australia's exposure to the risks of AI. Instead, it would leave Australia with less oversight of the technology, less ability to secure national interests, and few mechanisms to manage the risks and harms Australians are already experiencing. Australia would remain dependent on foreign infrastructure and decisions made elsewhere.

Hosting AI infrastructure, under the right conditions, is a promising way for Australia to secure benefits and manage risks.

Attachment A: Building AI while Compensating Creators

Consultation Paper on an Australian AI Training Scheme

Executive Summary

Australia faces a decision point on AI training. Australia is an attractive destination for training frontier AI models – we have abundant renewable energy, stable governance, and a skilled workforce. But copyright is a barrier. Unlike some other countries, Australia has no tailored legal regime that makes frontier AI training viable here.

The status quo serves no one. Australian creators receive little or no compensation when their work is used to train AI overseas. AI investment is deterred. Australia loses both the economic benefits and the strategic advantage of hosting AI training infrastructure. Australia's ability to pursue its national interest depends on negotiations between third parties—global AI companies and global rights holders.

The Government has ruled out a “text and data mining (TDM) exception” but has signalled it is open to other approaches. The Copyright and AI Reference Group is considering options, including how to encourage fair, legal avenues for the use of copyright material in AI training.

This paper proposes an Australian AI Training Scheme. The scheme has two components:

An **AI Training Permit** would be available for AI training in Australia. Obligations would include paying a tiered fee into consolidated revenue, granting practical consent and control to rights holders via an opt-out mechanism, complying with the National Data Centre Expectations, and meeting other national-interest conditions. Permit holders would be authorised to train AI with the certainty that they would not be subject to copyright restrictions in respect of that training. Permit holders would also receive support in navigating federal, state and territory requirements to accelerate “time to power” for new data centres.

An **Australian AI Public Benefit Fund** would be established separately by Government, funded from consolidated revenue, to support Australian creators. Its design would draw on the *Public Lending Right Act 1985* – compensation for income loss, and support for Australian culture.

The permit and the fund operate separately. This avoids characterising the fund as a de facto licensing scheme, preserves the Government's ability to adjust each independently, and gives ongoing policy levers as global circumstances evolve.

This paper invites comment from creators, collecting societies and rights management organisations, AI companies, and the broader public (see consultation questions on p.9). Responses are sought within three parameters:

1. Delivering Australian creator expectations for consent, control, and compensation.
2. Consistency with the Government's commitment not to create a TDM exception, while preparing for future copyright challenges emerging from AI.
3. Facilitating economically viable AI training in Australia, in the broader national interest.

Background

AI is becoming increasingly important to the global economy. Investment in AI infrastructure is globally significant – arguably the largest coordinated investment in modern history. Leading US technology companies alone are spending over US\$650 billion on AI infrastructure in 2026.

AI training is a distinct and critical element of the AI value chain. While many countries and regions will host AI "inference computing" – the hardware that runs AI models and serves local, time-sensitive demand – AI training compute is likely to be far more centralised. Training frontier AI models requires concentrations of specialised hardware and energy. Most countries will not have AI training compute. Where training happens matters, because the countries that host it capture greater economic value, strategic leverage, and a practical ability to shape how AI is developed.

Like any large infrastructure project, data centre construction and operation has challenges. These relate primarily to power consumption, but also land use, water, air quality, noise, and other local impacts. The National Data Centre Expectations set out how these issues should be addressed. AI companies are on the record, in Australia and overseas, agreeing to requirements of this kind, including covering grid infrastructure costs and offsetting consumer electricity price increases.

Copyright is the key outstanding barrier. Current Australian copyright law makes it uneconomic, and arguably logistically impossible, for AI companies to train AI models in Australia. Unlike the United States, Israel, the European Union, Singapore and Japan, Australia does not have a fair use or TDM exemption that would permit AI companies to train on publicly available material. Under the status quo, AI companies will continue to train in jurisdictions where training data is available at low or no cost.

Australian creators argue that AI companies should license their materials. AI companies are unlikely to follow this path for three reasons. First, the AI companies are not required to license the material in other jurisdictions. While there is legal uncertainty in some jurisdictions like the US, the law is unequivocal in other jurisdictions, like Japan, Singapore and Israel. Second, Australian copyright law protects all copyrighted material, not just Australian material. Training AI in Australia would require licensing copyrighted works globally – something AI companies believe is unworkable. Third, formally licensing publicly available copyrighted material could undermine favourable court rulings in other jurisdictions, particularly the United States, where AI training is currently treated as fair use.

The status quo serves no one well. Creators receive little to no payment and may not be able to withdraw permission when their material is used to train AI in the US or China, where most AI training occurs today. AI investment is deterred because current settings make training here uneconomic. And Australia's geostrategic position is weakened – without a place in the AI training value chain, Australia has limited leverage to shape the future of AI in the national interest.

Political context

The Australian Government has ruled out a TDM exception to copyright. In October 2025, Attorney-General Michelle Rowland stated:

While the Government is not considering a Text and Data Mining Exception, work is underway to ensure that Australia is prepared for future copyright challenges emerging from AI – so that Australian creators are protected and supported while unlocking new uses of copyright material.

The Government has since tasked the Copyright and AI Reference Group with work across three priority areas, including to "encourage fair, legal avenues for using copyright material in AI." The Attorney-General has framed the challenge as one of balance:

Artificial Intelligence is an exciting technological frontier full of opportunities, and we should be harnessing these opportunities for the benefit of local industries like our creative and media sectors, working with them, rather than at their expense.

This paper works within those parameters. It seeks comment on a scheme that delivers creator protections while creating a viable path for AI training in Australia.

Scope

This paper operates within three parameters:

1. Delivering Australian creator expectations for consent, control, and compensation.
2. Consistency with the Government's commitment not to create a TDM exception, while preparing for future copyright challenges emerging from AI.
3. Facilitating economically viable AI training in Australia, in the broader national interest.

All three parameters must be met simultaneously. This paper does not consider proposals that achieve one or two at the expense of the others.

This paper takes no position on the character or motives of AI companies or rights holders. It focuses on identifying a workable policy framework.

We invite comments that fall within these parameters. The question is whether a practical arrangement exists that satisfies all three.

Proposal

This paper proposes the creation of a two-part Australian AI Training Scheme — a permit framework for AI training in Australia and a separate fund to support Australian creators.

AI Training Permit

Entities conducting AI training in Australia may apply for an AI Training Permit. Training conducted without a permit must comply with existing laws, including copyright law.

The permit would carry the following obligations:

- **Fee.** Permit holders would pay a fee into consolidated revenue. The fee would be set by Government and could be tiered to reflect differences in purpose and scale. Low or no fee permits could be available for training of limited commercial value, education and academic research.
- **Consent and control.** Permit holders would be required to respect an effective, practical, and free opt-out mechanism for Australian rights holders. The design of the opt-out mechanism would be developed in consultation with rights holders, AI companies, and researchers.
- **Data centre expectations.** Permit holders would be required to comply with the National Data Centre Expectations, including commitments on energy, water, emissions, and community impacts.
- **Other national interest conditions.** The permit could include additional conditions where justified, such as giving notice to the Australian Government of training runs, transparency and disclosure requirements around models created in Australia, and a prohibition on circumventing technological protection measures (TPMs) within the meaning of the *Copyright Act 1968*.

Permit holders would be authorised to conduct AI training in Australia with the legal certainty that such training would not be subject to copyright restrictions (the separate prohibition on circumventing TPMs continues to apply). The privilege attaches only to permit holders in compliance with their obligations. It does not apply to entities that do not hold a permit or that fail to meet its conditions. The Department of Industry, Science and Resources would also provide permit holders with direct assistance in understanding and managing regulatory requirements across local, state and territory and federal jurisdictions.

Australian AI Public Benefit Fund

Government would separately establish and administer a fund, resourced from consolidated revenue, to support Australian creators. The *Public Lending Right Act 1985* provides a tested model.

The fund would serve a dual purpose: compensating Australian creators for income loss, and supporting the enrichment of Australian culture. Eligibility would be limited to Australian citizens and permanent residents who are creators, publishers who publish in Australia and their respective rights holders. The Minister would appoint a committee of relevant stakeholders — supported by staff within the responsible department — to administer the fund and develop operational rules, including how eligible material is identified and how payments are calculated.

The fund's detailed design would be developed in consultation with creators to ensure it meets their expectations. Because the fund is decoupled from negotiations with AI companies, the Government has substantial latitude to design it to meet creator and rights holder expectations.

Policy design

Linking obligations to privileges

The proposal places obligations on those engaged in AI training: paying a fee, giving practical consent and control to rights holders, complying with the National Data Centre Expectations, and meeting other national interest conditions. Only entities that meet these obligations receive the privilege of training AI in Australia under the permit.

A key reason a TDM exception is not a promising policy direction is that it grants a valuable privilege without corresponding obligations in the national interest. Once a TDM is in place, Government would have limited levers to require consent, control, and compensation. A TDM also does not allow for additional nuance on national interest matters beyond copyright.

Any policy will require some linkage between the *Copyright Act 1968* and a broader scheme. Under this proposal, the permit is the vehicle for that linkage. Enabling amendments to the *Copyright Act 1968* would grant privileges to permit holders. Those privileges would apply to all content protected by Australian law.

Asymmetry between Australian and global creators

The fund supports Australian creators. It does not serve global creators. In this regard, it mirrors the *Public Lending Right Act 1985*, which only serves Australian authors and publishers.

The asymmetry is practically necessary. If the fund were to serve all creators or rights holders globally, it would likely be uneconomic.

The permit and the fund are decoupled to manage this. The permit would be available for entities wishing to conduct AI training in Australia, and – via enabling amendments to the *Copyright Act 1968* – authorises training on all content protected by copyright under Australian law. The fund is a separate domestic support mechanism, resourced from consolidated revenue, directed at Australian creators. The two operate independently.

Decoupling also creates policy optionality on both sides of the scheme. The permit fee reflects what AI companies will pay for market access, legal certainty, and accelerated time to power. The fund reflects what Government chooses to invest in Australian creators and Australian culture. These are different questions that might need different scopes, quantum and other policy design variables.

Setting the permit fee

The permit fee would be set by the Government.

AI companies make arguments about the value of training data: that copyright material is a small fraction of overall training data, which also includes non-copyright material, synthetic data, and reinforcement materials owned by the companies; that the value of any discrete data point is negligible; that the value comes from aggregate data; or that Australian data is a small fraction of what AI models train on.

These arguments are not material to the fee. The permit fee is not a purchase price for copyrighted material. It is a fee for the privilege of training AI in Australia under the scheme, which includes the authorisation to train without copyright restrictions and the other conditions of the permit.

Through tiering, the fee design can reflect differences in purpose and scale. Some AI training does not produce a commercially valuable product, such as academic research. Some use copyright material to produce highly commercially valuable products. The fee structure can accommodate these differences.

Design and administration of the fund

The *Public Lending Right Act 1985* provides a model that the Australian AI Public Benefit Fund can follow.

Dual purpose. The *Public Lending Right Act 1985* has a dual purpose. First, it makes payments to Australian creators and publishers in recognition of income lost from their books being available in libraries. Second, it supports the enrichment of Australian culture by encouraging Australians to create and publish books in Australia. The AI Public Benefit Fund could adopt a similar dual purpose. This would meet creators' direct ask for compensation and provide a structural basis to support Australian creative industries amid the disruption AI is likely to cause.

Eligibility. As with the *Public Lending Right Act 1985*, eligibility would be limited to Australian citizens and permanent residents who are creators, and to publishers who publish in Australia. The fund could cover creators across the forms of material that AI models are trained on, including text, images, music, code, and video.

Triggering and calculating payments. The *Public Lending Right Act 1985* pays creators based on estimated copies of their books held in Australian public libraries, calculated from an annual survey. The direct equivalent – identifying which works are held in AI training datasets – is not straightforward. Identifying whether a specific work appears in a training dataset can be technically challenging. The fund would need a workable proxy. Options include using existing institutional infrastructure – such as membership of recognised collecting societies like the Copyright Agency or APRA AMCOS, or listing in established catalogues like the ISBN system. Alternatively or additionally, the fund could support the creative industry more broadly, in recognition of the benefits of the creative industry to Australia as a whole. The appropriate mechanism is a matter for further consultation.

Distribution. As with the *Public Lending Right Act 1985*, payments would be calculated from a funded pool, with a tapering formula to ensure the fund supports a broad base of creators rather than concentrating at the top.

Administration. As with the *Public Lending Right Act 1985*, the Minister would appoint a committee comprising relevant stakeholders – creators, publishers, and other appropriate representatives – to administer the fund. The committee would be supported by appropriate staff within the responsible department. Operational rules, including the mechanism for identifying eligible material and calculating payments, would be set by the committee and approved by the Minister via legislative instrument. This arrangement keeps the fund adaptive as AI's impact evolves, without requiring legislative amendment each time operational rules need to change.

Size of the Government contribution. As with the *Public Lending Right Act 1985*, the Government would decide how much to pay into the fund. The amount is a matter for budget processes and ongoing policy judgment.

Designing the consent and control mechanism

Consent and control, alongside compensation, are responsive to the concerns of Australian creators and rights holders. Any solution needs to meet their expectations while being practical to administer.

Permit holders must offer an effective opt-out for Australian rights holders.

The scheme adopts opt-out rather than opt-in. Opt-in would require AI companies to secure permission from every rights holder whose material might be in a training corpus. This is likely unrealistic.

Three design conditions are essential for the mechanism.

Effective. The approach should account for technical mechanisms that lower regulatory burden (such as machine-readable signals like robots.txt), while ensuring Australian rights holders can effectively opt out across all types of works.

Practical for rights holders and AI companies. An opt-out that requires legal knowledge, fees, or complex identification of works will not deliver meaningful control. Rights holders should be able to opt out simply, at the level of individual works, categories of work, or all of their work. Equally, an opt-out that requires manual tracking of an unlimited number of works would not achieve the goal of facilitating economically viable AI training in Australia.

Enforceable. Permit holders must make reasonable efforts to respect opt-outs as a condition of their permit and be able to demonstrate the processes that they use to ensure compliance. Non-compliance should result in fines that escalate to the loss of the permit. Creators should have a practical avenue to raise concerns.

Technological protection measures are a related signal of non-consent. Where a rights holder has applied a TPM to their material, the permit does not authorise circumvention. Permit holders must comply with Part V, Division 2A of the *Copyright Act 1968*, and breach is a ground for enforcement action under the scheme.

The specific design should be developed in consultation with creators, rights holders, AI companies, and other experts.

Enforcement

Enforcement would be administered by the Department of Industry, Science and Resources (the department), which already administers the National Data Centre Expectations, in coordination with the Attorney-General's Department. The department would have the power to monitor compliance, investigate breaches, and take action against permit holders, including imposing financial penalties and suspending or withdrawing the permit. Grounds for action by the department would include failure to pay the permit fee, failure to respect opt-outs, failure to comply with the National Data Centre Expectations, failure to meet other national interest conditions, or failure to comply with other Australian laws.

Entities that train AI in Australia without a valid permit would not receive the copyright authorisation the scheme provides. They would remain exposed to infringement action by rights holders under the Copyright Act 1968.

Consultation questions

On the overall proposal

1. Does the proposed Australian AI Training Scheme — a permit and fund — represent a workable path to meeting all three parameters in the Scope section? If not, what parameter is not met and why?
2. Are there alternative arrangements that meet the three parameters? If so, what is that proposal, how does it meet the three parameters, and how does it compare to the proposed Australian AI Training Scheme?

On the AI Training Permit

3. The scheme proposes opt-out as the primary consent and control mechanism. What approach would best deliver an opt-out that is effective, practical for creators and enforceable by Government? If you prefer opt-in, how could that be achieved within the parameters?
4. Are there national-interest conditions that should be attached to the permit from day 1?

On the Australian AI Public Benefit Fund

5. Who should the fund support, and how should eligibility be defined? Should the fund focus on creators whose materials have been used to train AI, creators who have lost earnings because of AI, the ongoing viability of Australia's creative industry, or other matters? Why?
6. What should the role of established organisations like collecting societies be in designing, distributing and managing remuneration through the fund to creators?
7. How should compensation under the fund be calculated? What approach best balances feasibility, fairness, and administrative cost?

Open question

8. Are there issues raised by this proposal that are not addressed in the paper but are relevant to meeting the three parameters set out in the Scope section?

This paper invites comment from creators, collecting societies and rights management organisations, AI companies, and the public. Responses are sought within the three parameters:

1. Delivering Australian creator expectations for consent, control, and compensation.
2. Consistency with the Government's commitment not to create a TDM exception, while preparing for future copyright challenges emerging from AI.
3. Facilitating economically viable AI training in Australia, in the broader national interest.

All three parameters must be met simultaneously. This paper does not consider proposals that achieve one or two at the expense of the others.

This paper takes no position on the character or motives of AI companies or rights holders. It focuses on identifying a workable policy framework.

Frequently asked questions**Why is the fund funded from consolidated revenue rather than from permit fees?**

The fund and the permit are decoupled. Permit fees flow into consolidated revenue. The fund draws from consolidated revenue. The scheme would be designed so that there is no direct link between the two.

This design has several purposes. It ensures the fund is not characterised as a de facto licensing scheme, where AI companies pay rights holders for the use of their material. It preserves the Government's ability to set the permit fee and the fund independently, based on different considerations. It also mirrors the Public Lending Right Act 1985, where Government funding is not directly tied to any specific revenue source.

Decoupling also avoids constitutional constraints that arise where a Commonwealth charge is tied to payments to private parties. The fee is a charge for the privilege of training. The fund is Commonwealth expenditure.

Why can't AI companies just license the Australian material they need?

Australian copyright law protects all copyright, not just Australian copyright. The corpus of material protected under Australian law is very large, and most of it is available at low or no cost in jurisdictions with fair use or TDM exemptions. Licensing only Australian material would not authorise AI training in Australia under current copyright law. Licensing all material protected under Australian law would be practically unworkable and could undermine AI companies' legal position in other jurisdictions, particularly the United States.

Why not just require AI companies to pay creators directly through a collective licensing scheme?

Collective licensing is one option being considered by the Copyright and AI Reference Group. AI companies would pay licence fees to a collecting society, which would distribute payments.

The proposed scheme differs in three ways.

First, a collective licensing scheme administered under the *Copyright Act 1968* would cover all material protected under Australian law, not just Australian material. AI companies would pay for a licence covering the global corpus, even though the same material is available at low or no cost in jurisdictions with fair use or TDM exemptions. This would reduce the compensation received by Australian creators while leaving the same practical problem as direct licensing – routing it through a collecting society does not reduce the scope of what is being licensed.

Second, a collective licensing scheme treats the payment as a fee for the use of copyrighted material. AI companies have indicated they are unlikely to accept this, as it could undermine court rulings in other jurisdictions where AI training is treated as fair use.

Third, collective licensing addresses copyright only. It does not provide a vehicle for other national interest conditions – such as compliance with the National Data Centre Expectations, notice of training runs, or transparency requirements.

The scheme does not preclude collective licensing arrangements for other uses of copyrighted material, such as outputs of AI models or non-training uses, or for non-permit holders.

The Government ruled out a TDM exception. Isn't this just a TDM?

No. A TDM exception grants a privilege without a corresponding obligation. This proposal grants a privilege only where specific obligations are met – including payment of a permit fee, practical consent and control mechanisms for rights holders, compliance with the National Data Centre Expectations, and other national interest conditions. The privilege attaches to permit holders in compliance with their obligations, not to AI training generally.

Does the scheme infringe Section 301 of the US Trade Act of 1974?

No. The scheme imposes reasonable and justifiable obligations on a specific activity – AI training – occurring within Australia. It is country-agnostic. It does not discriminate between AI companies based on nationality, and it applies equally to domestic and foreign entities. It does not restrict market access. AI companies that do not hold a training permit can continue to offer their AI products in the Australian market or train AI without the permit's privileges and obligations.

What happens to AI models trained overseas but deployed in Australia?

The scheme regulates AI training in Australia only. It does not restrict the deployment of AI models in the Australian market. Models trained in other jurisdictions – including in jurisdictions with fair use or TDM exemptions – can continue to be offered to Australian users.

The scheme is focused on attracting AI training activity to Australia, supporting Australian creators, and securing the benefits of AI. It does not cover AI use, deployment, or adoption.

Does the scheme affect creators' existing rights under copyright?

No. Creators retain their existing rights under the *Copyright Act 1968*. The AI Training Permit only authorises AI training conducted in Australia by permit holders in compliance with the scheme. It does not affect rights with respect to AI training occurring overseas, training not under a permit, AI outputs, or any other use of copyrighted material.

What counts as "training in Australia"?

Training in Australia is based on the location of the AI training compute. It covers AI training occurring wholly or partly using chips physically within Australia's jurisdiction.

Is the permit fee a tax on AI training?

No. The permit fee is a fee for a specific privilege – the authorisation to train AI in Australia under the scheme, including the copyright authorisation and the other conditions of the permit. It is not a general revenue measure. Entities that do not hold a permit do not pay the fee and can still train AI in Australia, just without the scheme's privileges and obligations.

Does all AI training in Australia need a permit?

No. AI training does not require a permit. Permit holders gain legal certainty that using the data for AI training is not subject to Australian copyright restrictions. It would also provide the permit holder with direct engagement and assistance from the department in understanding and managing regulatory requirements at the federal, state and territory, and local levels. AI training can occur without these privileges and obligations.

Who is Good Ancestors and why is it running this consultation?

Good Ancestors is an Australian not-for-profit. We develop evidence-based policy on issues with long-term consequences for Australians, including AI.

We think AI carries generational risks and generational opportunities. Copyright has become the sticking point in Australia's AI policy conversation. We don't think a workable path has to come at the expense of rights holders or Australia's long-term national interest. This paper proposes one way to do both. We're running this consultation to test it.

Attachment B: What We've Heard So Far

Consultation Paper on an Australian AI Training Scheme

The consultation remains open. Below is an initial summary prepared for the Senate Inquiry into AI and Data Centres. Comments are not attributed to protect the anonymity of participants.

A clear preference from rights holders for licensing

All rights holder groups expressed an unambiguous preference for licensing as the path forward. No rights holder group was open to exploring alternatives.

Lack of detail about licensing as a pathway

We heard claims from rights holders that collective licensing could cover most AI training requirements. This included claims like “80% of all music could be licensed under four deals”.

So far, we have found these types of claims hard to interrogate. No group was able to detail the extent of existing collective licences, the extent of possible extended licences, the extent of reciprocal agreements, and the overall extent of overlap with AI training datasets.

We also found the specific claim of “80% of all music could be licensed under four deals” hard to evaluate. We could not evaluate whether “music” refers to “commercially produced music by musicians working with record labels” or “all music, including user-produced music, music on platforms such as YouTube, and music otherwise not linked to record labels”, or something else. We also don't know whether the “80%” refers to market revenue share or the count of works.

The lack of clarity accounts for part of the disagreement between AI companies and rights holders, and underscores that rights holders and AI companies do not see eye-to-eye on the issue or on how licensing could work.

- AI companies focus on the **quantity of publicly available work** to ensure certainty about their entire AI training datasets. The normal commercial value of a work is largely independent of its contribution to AI training. Music is not a highly valuable data type for AI training.
- Rights holder groups focus on the **commercially valuable works created by the artists** they represent and are less likely to represent low-commercial works.

The large number of copyright-protected works with low commercial value are a focus for AI companies but of less interest to rights holder groups.

Lack of clarity about the starting point for licensing negotiations

We heard claims from rights holders that they have not been approached by AI companies to explore licensing as a pathway.

We heard claims from some AI and technology companies that they have approached rights holders.

We heard claims from other AI and technology companies that they don't see value in negotiations for material that can be readily licensed while there's no solution for the remaining works with no plausible pathway.

Lack of incentive to agree on licensing pathways

Rights holder groups expressed lack of trust, or even animosity, towards AI companies. Rights holder groups questioned the consultation's scope, which asked respondents not to attack character or motives. Rights holder groups argued that the historic behaviour of AI companies is relevant to their character, and said they were reluctant to deal with companies that have a track record of breaking the law.

Rights holder groups also argued that they were on track to win in the US courts, after which they would receive compensation and potentially damages for most AI training. That reduces their incentive to reach an agreement in Australia, which would at best cover some models by some companies.

AI and technology companies said they did not consider licensing to be a viable pathway to training frontier models in Australia.

Given that a deal requires agreement between one or more AI companies and many rights holders and a deal has not been reached to date, we assess that a deal is highly unlikely unless a change occurs.

Disagreement about prospects in US courts

We heard from rights holders that they were confident of success in the US courts. Specifically, they believed claims that AI training has been upheld as fair use are overstated, and that US courts are on track to find that AI training is not fair use. Further, rights holder groups believe that – once they succeed in US courts – the Trump administration would not intervene to favour AI companies. This could lead to them receiving compensation and damages from all AI companies training in the US.

We heard from AI and technology companies that they are confident that US court decisions so far uphold that the act of AI training itself can be fair, but concede that in some cases the circumstances of the training (such as certain modes of obtaining data) were not. AI and technology companies think US courts will continue to hold the view that the act of training itself is fair use.

Further, AI and technology companies believe that the Trump administration would likely intervene to favour AI companies over rights holders if courts found that AI training was not fair use.

Other consultation responses tended to support the view that the Trump Administration was likely to side with AI companies on copyright issues, citing specific remarks made by the President.

The disagreement about how US courts are likely to act is important to understanding the path of possible negotiations. There is further discussion of the most likely alternatives to negotiated outcomes below.

Disagreement about plausible alternative countries

AI and technology companies tended to claim that Australia was a highly desirable location for AI training, but that there were viable alternatives with workable copyright regimes and other desirable factors, such as adequate energy availability. Japan was given as an example.

Rights holders tended to claim that alternatives to Australia were overstated.

Low tolerance for uncertainty from tech companies

We heard a low tolerance for legal uncertainty from tech and AI companies. Companies worried that if they followed a licensing path and secured relevant rights for works controlled by rights holders, there would be substantial residual litigation risk from creators and rights holders not covered by those deals. Their view was that extensive investment required a high degree of legal certainty.

Concerns about the certainty of the fund

Rights holder groups expressed concern about the lack of certainty surrounding the Australian AI Public Benefit Fund. Specifically, they worried that a current or future government may not contribute meaningfully to the fund even if the Permit raised significant funds, or that the fund may fail to distribute money in a way that meets the needs of creators and rights holders.

Rights holder groups also expressed concern that if the scheme fails to attract AI training to Australia, AI companies would never buy a permit, and the fund may never operate, meaning Australia's copyright laws would be weakened for no benefit.

Concerns about Australia's treaty obligations

Some rights holder groups expressed concern that extending opt-out only to Australians may violate treaty or other international obligations.

The consultation paper currently addresses this risk in two ways:

- First, opt-out is a feature of the permit, not of the fund. Australian copyright law would not have differential treatment of Australian works. Instead, any differential treatment would stem from the technical practices of AI companies that hold a permit.
- Second, the consultation paper has optionality to extend the opt-out beyond Australian creators and rights holders.

What this means for the proposal

The consultation surfaced specific changes that would improve the model. For instance:

- Legislation to create the fund should mandate a committee like the [PLR Committee](#).
 - The new Committee should include representatives from creator and rights holder groups to increase confidence that the fund would be well administered from their perspective.
- Legislation to create the fund could mandate that distributions be made in part through rights holder groups, declared collecting societies or organisations that administer statutory licences.
 - Funds should also be available to creators not represented by rights holders groups.
- Legislation to create the fund could include mechanisms to increase confidence that future governments would contribute. For instance, there could be a specific endowment and an investment mandate, like the [Disaster Ready Fund](#).
 - The DRF is a \$1 billion fund that distributes \$200 million per year.
- The Permit and the Fund could sunset after a given period (e.g., five years) in the event that no permits are ever issued. This would address rights holder concerns about the legacy of the scheme if it is never used.
 - If permits are issued, the legislation should not sunset to ensure certainty for companies that invest in Australia on the basis of the permit scheme.
- Legal advice is required on whether the opt-out mechanism should be global or can be Australian only.

The consultation surfaced opportunities for clarification. For instance:

- Some commentary suggested that the fund would be a once-off payment.
 - The permit is intended to be annual and the fund ongoing.
- Some commentary suggested that the AI training scheme would undercut licensing / data access deals that have already been made.
 - The scheme's intent is to preserve valuable data access opportunities. A permit would not authorise the circumvention of TPMs, such as paywalls. This means commercial deals between AI companies and companies like HarperCollins and the Associated Press involving content behind TPMs would remain on the table.

Overall, the proposal assumes agreement that the status quo is bad for all parties. Specifically, Australian works are currently being used to train AI overseas, with Australian creators and rights holders receiving little or no compensation. Finding a way to attract AI training *and* support Australian creators would be better overall. The proposal also assumed that parties would seek a deal better than the status quo. Under these assumptions, there is a large zone of possible agreement between stakeholders, and the proposal is better than the status quo for all stakeholders.

The main point the consultation surfaced is that **there is little common ground between stakeholders and little appetite to seek mutually beneficial deals**. One explanation for this is different views about the most likely alternative to a negotiated agreement (MLATNA).

Specifically, the paper assumed that the zone of possible agreement was large because parties agreed that the **most likely alternative was the status quo**—that AI training continues overseas where Australian creators and rights holders get little compensation. However, the consultation suggested that different stakeholders have different views on the most likely alternative:

- Rights holders believe that the most likely alternative is victory in the US courts, leading to payments from all AI companies and potentially damages, as well as future licensing details.
- AI and technology companies believe that the most likely alternative is that they win in the US courts, or that the Trump administration intervenes, or that they train in another jurisdiction with fair use or a TDM exemption.

The below table sets out five possible future scenarios. One is a directly negotiated solution (like licensing). One is a government-negotiated solution (like the AI training scheme). The remaining are alternatives to negotiated solutions based on different outcomes in the US courts and different actions by AI companies. The table illustrates that any negotiated solution is better for Australia's national interest, and that alternative solutions are only better for creators if they prevail in US courts.

Possible future scenarios

The below table sets out five possible future scenarios. A negotiated solution (Scenarios 1–2) is the only path that somewhat secures Australia's national interest. For creators, the alternatives beat a negotiated outcome in just one case (Scenario 3) which requires a US court win and no intervention by the White House.

Negotiated solutions

Scenario	Creators/rights holders	Australia's national interests
1 Directly negotiated solution: AI companies and rights holders agree licences and reciprocal arrangements that allow AI training in Australia	 Australian creators get opt-in and compensation when AI is trained in Australia.	 AI training happens in Australia, but further mechanisms are required to ensure that training happens in the national interest
2 Government-negotiated solution: Australia implements the Australia AI Training Scheme	 Australian creators get opt-out, and compensation via the fund. But exchange a property right for a government fund.	 AI training happens in Australia with a specific mechanism to secure the broader national interest

Alternatives to a negotiated solution

Scenario	Creators/rights holders	Australia's national interests
3 Alternative to a negotiated solution: AI companies double-down on US investment Creators win key court cases in the US, and the Trump administration does not intervene.	 Australian creators get opt-in and compensation when AI is trained in the US.	 AI training does not happen in Australia, meaning reduced economic benefit and reduced ability to secure the national interest.
4 Alternative to a negotiated solution: AI companies double-down on US investment Creators lose court cases in the US or the Trump administration intervenes to favour AI companies.	 Australian creators have no opt-out, no credit, and no compensation when AI is trained in the US.	 AI training does not happen in Australia, meaning reduced economic benefit and reduced ability to secure the national interest.
5 Alternative to a negotiated solution: AI companies train in other jurisdictions where copyright laws are favourable to them.	 Australian creators have no opt-out, no credit, and no compensation when AI is trained in other jurisdictions.	 AI training does not happen in Australia, meaning reduced economic benefit and reduced ability to secure the national interest.

 The most desirable outcome  An outcome with pros and cons  The least desirable outcome

Likelihood of possible scenarios

How stakeholders approach the negotiated options depends heavily on how likely they think each alternative is.

A directly negotiated solution is very unlikely

Scenario 1—direct negotiations between AI companies and rights holders—is **very unlikely**. Rights holders and AI companies do not agree on key issues, and many have said they don't trust the other parties and/or don't want to negotiate a direct agreement.

A Government-negotiated solution is possible

Scenario 2—Government intervention to create a path or to change incentives or frameworks—is **possible**. Government has the authority to find a negotiated solution. Only a Government-negotiated solution allows Government to protect Australia's broader national interests.

The alternatives to a negotiated solution depend on the United States

The other scenarios (3-5) are what happens if no deal is reached. Which one might eventuate depends on the US courts and US administration. For rights holders, winning in the US courts (Scenario 3), is the best alternative to a deal because it could secure opt-in, compensation, and potentially damages. For rights holders, winning in the US might be better than any negotiated solution.

If rights holders don't win in US courts, Scenarios 4 and 5 are the default.

In every alternative, training happens offshore and Australia's national interests are not secured.

What this means for policy makers

All stakeholders should work to avoid scenarios 4 and 5.

In almost all cases, negotiated solutions are better for all stakeholders than not reaching a deal.

Policy makers, regardless of their priorities, who believe that the US courts and the White House will favour AI companies should prefer a negotiated solution (1 or 2) in Australia.

Not reaching a negotiated solution is only good for policy makers who both prioritise the interests of rights holders **and** believe that rights holders are likely to win in the US courts and receive support from President Trump (Scenario 3).

Policy-makers who agree that a directly-negotiated solution (Scenario 1) is not on track should pursue a Government-negotiated solution (Scenario 2).