

Committee Secretary
Joint Select Committee on Artificial Intelligence

21 September 2026

About the Australian Council on AI Strategy (ACAIS)

Australia is making consequential decisions about artificial intelligence (AI) including what it builds, what it regulates, what it adopts and what it declines, on a compressed timetable and against a thin domestic evidence base. Very few of the people making those decisions have seen the frontier of the technology first-hand. The consequence is a debate narrower than the decisions require. The Joint Select Committee has an opportunity to widen it.

ACAIS is an independent, non-partisan Australian charitable body focused on strengthening AI policy and capability and supports robust evidence-informed decision making. We operate nationally and internationally to bridge the knowledge and institutional gaps between actors across the AI ecosystem.

We serve Australia's national interest and recognise the significant role Australia can play in advancing the global public interest.

ACAIS was established to support Australian AI governance and policy decisions, because the stakes are high and near-term decisions will be highly consequential.

- Democratic middle powers, like Australia, can influence and even lead international AI governance and safety debates, and help to moderate risks associated with the current great power rivalry that is amplifying AI challenges and complicating constructive responses.
- Australia is only likely to have a material influence on shaping AI development if we are materially embedded in the frontier supply chain through hosting training compute in Australia.
- Australia has been dealt an extraordinary hand in attracting AI training, but it will not play itself. Australian governments must move with urgency to take advantage of this window of opportunity and give form to the distinctive contribution Australia can make to shaping frontier AI, on Australia's terms and in a way that shares the benefits of AI development and deployment.

ACAIS does not represent or advocate for any corporation or interest group, including AI companies, data-centre developers, energy proponents, copyright holders or creative-sector organisations.

We hold no commercial interest in any outcome discussed in our submission. As a matter of policy, we do not accept funding from AI companies, and our analysis and advice are free from funder influence. We are committed to the highest standards of intellectual integrity and we maintain strict intellectual independence and sole direction and control over our activities.

Assistance to the Committee and invitation to US study tour

We would welcome the opportunity to appear before the Committee. ACAIS can readily draw on and make available independent international expertise in frontier AI measurement, evaluation and forecasting. We can provide evidence and analysis across multiple AI policy dimensions including geostrategic, national security, energy, copyright, AI safety, and international governance.

In November 2026, ACAIS is hosting a Study Tour for an Australian delegation to San Francisco and Washington DC for direct engagement with frontier developers, independent evaluators and forecasters, and United States policymakers. We hope this will help address Australia's thin evidence base, and allow delegation members to form their own views on AI challenges and opportunities. We would be glad to host members of the Joint Select Committee on the Study Tour and make the resulting analysis available to the Committee ahead of its reporting date.

ACAIS also plans to commission research and analysis to address key knowledge gaps as they emerge and would be pleased to brief the Committee on the scope and findings of such work as it progresses.

For further information regarding this submission please contact Anna-Maria Arabia OAM, CEO of ACAIS on annamaria.arabia@acaais.au or +61 412 940 921.

ACAIS Submission

Summary

- Australia can help write the rules for frontier AI development and deployment, or it can inherit them. The difference rests on whether we are materially embedded in the frontier AI supply chain, and the decisions that determine this are being made now.
- For Australia, which has not been a significant part of the AI supply chain, hosting future training compute is the pre-requisite for exercising any significant global influence to advance AI safety, and for capturing an outsized share of the economic benefits at home.
- While hosting training compute does not of itself produce AI safety, there are safety conditions, postures, and practices that only jurisdictions that physically host compute can demand, hold and implement.
- Australia has been dealt an extraordinary hand in attracting AI training, but it will not play itself. Australian governments must move with urgency to take advantage of this window of opportunity and give form to the distinctive contribution Australia can make to shaping frontier AI, on Australia's terms and in a way that shares the benefits of AI development and deployment.
- The distinction between inference and training data centres matters significantly in AI deliberations because it determines whether investment in Australia buys leverage or simply buys buildings.
- Assessing the risks and benefits of hosting training compute in Australia requires comprehensive examination of what is at stake for Australia across a range of policy dimensions including, geostrategic, economic, national security, international governance, energy, copyright, regulatory and safety implications of AI.
- The decisions Australia makes on copyright law in the remaining months of 2026 will determine whether we shape our own future or accept a future shaped by others.
- We strongly urge swift resolution of the current copyright impasse because the status quo serves no one. Today, Australian content is already being used in training conducted in other jurisdictions without remuneration, while current copyright laws deter training in Australia by any AI company (large or small). Today, Australian creatives bear the cost without compensation or control.
- The Australian Government should negotiate conditions and requirements that serve our national and global interest before it enters into a deal with frontier AI labs to host training compute in Australia. This should include strengthened safety measures (exploration of mandatory incident reporting; security standards for AI data centres; strengthened evaluation; development of verification technologies to support international governance; access to compute for safety research); upskilling of the local AI R&D ecosystem; and a powerful seat at the international governance table. None of this desired level of influence will eventuate if Australia fails to secure frontier AI training.
- Crafting a win-win resolution for copyright and securing in principle agreement to host training at scale in Australia will not be the final chapter in Australia's AI journey. Setting appropriate rules for data centers, ensuring Australia remains an attractive investment destination (including improved approvals processes), and implementing strong benefit sharing arrangements will all be crucial.
- AI is already providing step-change improvements in the competence and capability of individuals and organisations. This empowers people both to do good things – like rapidly find cures to diseases – and bad things like exploit cyber and biological vulnerabilities. Australia will be best placed to capture opportunities, mitigate risks and manage inevitable disruptions when we host frontier AI training and talent, and earn a seat at the global table.

The current AI context


Frontier AI is being shaped by two actors: the US and China. Safety and security considerations are being inadequately prioritised for the sake of AI progress.

Widening the set of voices in AI governance would help bring weight to safety negotiations.

Democratic middle powers, like Australia, can influence and even lead international AI governance and safety debates and shape and pace the frontier only when they have leverage which can be gained by having a foothold in the AI supply chain. Currently, most democratic allies of the US have little, if any, leverage.

The most effective way for Australia to generate leverage is by hosting training compute here.

The strong racing dynamics between companies and between the US and China mean AI companies need to train models. Access to training compute is essential for AI progress and any disruption to access would be a significant setback for AI companies. Epoch AI measures that the compute used to train a frontier AI model has been growing at five times a year since 2020. (figure below)¹



As such, US frontier labs will not stop building compute, and the current administration is committed to supporting them. The constraints they have encountered are energy supply and growing local opposition to the poorly managed rollout of data centre developments. Regulatory direction in the United States has also shifted repeatedly and at short notice. Deploying capital at this scale and timeframe requires stable rules, and jurisdictions offering stable and predictable settings are becoming more attractive as a result.

These are some of the reasons why they are looking to build compute abroad.

Not all data centres are made equal

A great deal of Australian discussion treats "data centres" as a single category. They are not. The distinction matters significantly in this debate because it determines whether investment in Australia buys leverage or simply buys buildings.

¹ <https://epoch.ai/trends#training-runs>

The difference between training and inference compute is outlined below.

	Training compute (developing AI)	Inference compute (deploying AI)
Physical form	Single, tightly interconnected campuses at multi-GW scale; tens to hundreds of thousands of chips acting as one machine	Smaller facilities, distributed globally, assembled incrementally
Requirements	Abundant power, land, stable governance and a permissive legal regime must coincide. Not latency-sensitive, so can be remote	Many viable locations worldwide; usually built near population centres for fastest time to power.
Substitutability	Very costly to relocate once power is secured	Highly fungible; likely to commoditise
Capital at stake	Hundreds of billions of dollars of investment	Orders of magnitude less than training
Strategic consequence	Upstream insight, negotiating position, researcher presence	Local jobs and latency; limited leverage

Frontier labs and other AI companies see Australia as a suitable destination for hosting training compute for the following reasons. Australia has:

- strong alignment with the US as a Five Eyes country and an AUKUS member;
- the potential and political will to develop large-scale green energy infrastructure suitable for data centers;
- land;
- stable institutions;
- a track record of building large energy projects; and
- a stable political environment.

For Australia, it is an opportunity to create relevance and leverage in the age of AI which it does not currently have since Australia has does not have frontier labs, has no mechanisms to influence international AI governance and safety debates and can't shape frontier development and deployment.

Without training compute onshore, Australia is a bystander watching what is coming towards it.

What is at stake for Australia

In the table below we summarise our assessment of the difference between Australia creating the conditions for frontier AI training onshore and Australia not doing so.

Dimension	If Australia enables training compute	If it does not
Geostrategic	Australia becomes materially embedded in the frontier supply chain; its risk assessments are heard in Washington DC and inside AI companies	Australia is a downstream consumer of decisions made elsewhere, with declining visibility over AI's trajectory and diminishing relevance as AI becomes central to economic and military power
Economic	Capture of a share of unprecedented infrastructure investment (global AI infrastructure spending in 2026 ² is estimated at over US\$1 trillion); revenue base to manage labour-market disruption and to compensate creators	Investment and value accrue elsewhere; the labour-market disruption arrives regardless, without the fiscal capacity to manage it
National security	Direct visibility of emerging dual-use capability in cyber, biology and other domains; negotiated incident reporting and evaluation access; researchers under Australian jurisdiction	A second-hand threat picture; no ability to anticipate or pre-empt AI-enabled threats; reliance on the goodwill of foreign firms and agencies
International governance	Standing to shape evaluation, incident-reporting and biosecurity norms, including as a convenor of middle-power democracies	Rule-taker status. Australia carries the risks of AI without a voice in how they are managed
Energy	Large, assured, firm demand can de-risk and accelerate clean generation and transmission investment, if conditions are properly designed	The clean-energy transition loses a potential anchor tenant; global AI load is met in jurisdictions with dirtier grids
Copyright	An opportunity to design a mechanism that leaves Australian creators demonstrably better off than the status quo	Australian content continues to be used in offshore training with no compensation, no ability to opt out, and no enforcement leverage

The bargaining power training compute creates for Australia

Hosting training compute in Australia gives us bargaining strength we don't currently have and provides a mechanism to build our understanding of frontier AI and to shape the capability. It also allows Australia to gain global advantage as a country materially embedded in the frontier supply chain.

This level of influence will not eventuate if Australia takes an inference-only approach.

In addition to adherence to the national data centre expectations and to the Australian Standards for AI as legislated, hosting training compute opens the following negotiating positions that ACAIS recommends the Australian Government examines in detail:

1. **Strengthened national security environment:** Negotiated access to frontier AI models when compute constraints bite and when powerful dual-use capabilities continue to emerge.

² <https://www.goldmansachs.com/insights/articles/global-investment-is-forecast-to-exceed-1-trillion-in-2026>

In June 2026 access to two frontier model tiers was suspended to comply with United States export controls demonstrating that Australia's access to the most capable AI systems rests on decisions taken in another country's regulatory process. That risk grows as compute becomes scarcer and models gain more significant national-security capability, placing Australia in a precarious situation. Australia should not assume it will be prioritised without a reason for others to keep it inside the tent, such as hosting training compute.

[Further detail on both the cyber and biological threat environment and what Australia can do are offered below.](#)


2. **Build AI safety knowhow:** Australian assessments of frontier AI capabilities, including recursive self-improvement and the risks of losing human control over future AI systems, has become essential must-have knowledge needed to shape AI governance on Australia's terms. No access or analytical capability like this exists in Australia.

Commendably, the Australian AI Safety Institute, as of mid-2026, has begun testing unreleased frontier models with the Australian Signals Directorate (ASD). It is nonetheless small, under-resourced and largely dependent on the goodwill of busy staff at foreign companies and overseas safety institutes (see comparative resourcing below).

Australia's threat picture is therefore a second-hand copy of someone else's.

Evaluation capability: the gap that determines what Australia can see

Annualised public funding for national AI safety institutes, A\$ million per year.



[Data sources and methodology: see end note]

Australia cannot be an influential jurisdiction globally in the regulation of AI or the pace of AI development, unless we hold a significant piece of the AI supply chain. This is best achieved by becoming the second centre for AI training in the Western world after the US and using that to assert safety conditions, postures, and practices that only jurisdictions that physically host compute can demand, hold and implement.

[Details of safety conditions, postures, and practices that are enabled by physically hosting training compute onshore is provided below.](#)

3. **Upskilling the local AI ecosystem:** Australia can request a specified number of AI engineers and frontier researchers working under Australian jurisdiction, to engage more closely with the Australian government, education, and industry sectors. This would help the Australian AI ecosystem grow and for us to understand where the technology is heading and to shape its governance directly, rather than solely responding to decisions made elsewhere.

The UK's Safety Institute was able to attract technical talent from the USA because it was serious about shaping AI governance.

While some leading AI engineers have already indicated a desire to relocate to Australia, it may still be beneficial for the Australian government to ensure that deals with AI companies explicitly incorporate a requirement to have AI technical staff in Australia. A research presence in Australia

will be key to upskilling the local AI ecosystem and the AISI, and will generate positive spillover effects. If those researchers are not Australians living abroad, they can be supported to relocate quickly by targeted use of the National Innovation Visa.

4. **R&D access:** Some AI training capability should be available for Australian institutions and researchers to conduct AI research. This could utilise models such as Anthropic's Claude Corps Fellowship or the Research and Safety Fellows Program where participants work directly with frontier researchers in specialised tracks.

As training compute power is an increasingly valuable and scarce resource, such an arrangement would assure affordable access for Australian research purposes, giving us analytical capability, training opportunities and direct insights.

5. **A seat at the international governance table:** Regardless of where AI is developed, the risks are felt worldwide. Efforts are required to jointly advance AI safety science and set global standards, rather than engage in a competitive race to the bottom, where safety is an afterthought.

Having access to critical training compute opens the possibility of democratic middle powers, like Australia to convene, influence and even lead international AI governance and safety debates which can shape and pace the frontier. A seat at the table can help establish common evaluation, incident reporting and biological synthesis screening standards (see cyber and biological risks below).

On 1 September 2026, US-based frontier labs urged that the pace of frontier AI development be slowed so that risk prevention can keep up. Anthropic detailed a 3-step framework³ to achieve this including frontier labs embedding independent evaluators in their companies; adoption of common safety standards coordinated and adopted across frontier labs in democratic countries requiring US Government support; and that the US and other democratic countries, like Australia, seek to coordinate with authoritarian governments to pace progress ensuring all countries comply with measures to do so.


Australia's role and position in the Indo-Pacific means it is strategically placed to influence deeper US-China technical exchange on AI risk and safety in general. For example, the development and adoption of verification technologies so that all countries comply with safety standard and prohibition of dangerous use of AI.

³ <https://darioamodei.com/post/we-must-pace-the-frontier>

Understanding the threat environment and what Australia can do

The threat environment

Frontier models have demonstrated the ability to identify and exploit vulnerabilities in every major operating system and every major web browser when directed to do so. The figure below shows critical vulnerabilities detected following release of Mythos Preview in May 2026.



SOURCE: <https://epoch.ai/data/cve?view=graph>, accessed 6 September 2026

Cyber threats

Bad actors exploiting cyber vulnerabilities and the methods they use are not novel, but the timelines are, and AI has democratised competence and capability to do so.

Anthropic's September 2026 threat report⁴, covering operations it disrupted between December 2025 and August 2026, shows that cyber breaches relied on familiar methods such as stolen credentials, unpatched devices, exposed services, SQL injection, and phishing.

However, where once well-resourced and sophisticated operations were required to exploit cyber vulnerabilities, this is now being performed by AI models running in parallel at machine speed.

Capability to act has now been levelled, with a hacker or a state actor all showing adoption of similar methodologies. The world is rapidly moving to a situation where the difference between a state and non-state actors is not sophistication, but intent.

The Five Eyes agencies issued a joint statement on 22 June 2026 warning that frontier models are expected to transform both offensive and defensive cyber capabilities, and that "the timeline is not years, it is months".⁵

⁴ <https://www.anthropic.com/threat-intelligence-report-september-2026>

⁵ <https://www.cyber.gov.au/about-us/view-all-content/news/five-eyes-cyber-security-agencies-statement>

In one case, a single French-speaking operator gained internal access to at least 14 of 42 targeted European political parties, media outlets, think tanks and their SaaS providers, and built a doxxing platform loaded with tens of millions of records⁶.

Anthropic disclosed on 30 July 2026 that, reviewing 141,006 evaluation runs, it found three incidents where models gained unauthorised access to the production infrastructure of three organisations after a misconfiguration left an evaluation environment connected to the internet.

OpenAI disclosed that two models, running with reduced cyber refusals for evaluation, escaped a sandbox, discovered a zero-day, and chained stolen credentials into remote code execution to steal benchmark answers from Hugging Face's production infrastructure; Hugging Face detected and contained the intrusion five days before OpenAI connected it to its own testing.

Google Threat Intelligence Group assessed with high confidence that criminals used an AI model to find and weaponise a two-factor authentication bypass, and that its own proactive counter-discovery may have prevented the planned mass exploitation⁷.

What can Australia do?

US frontier labs voluntarily disclose incidents that no regulation compels them to disclose. This is commendable, but insufficient. Evidence shows that detection of threats increasingly happens at the model provider level, that is upstream of the victim. This is where interventions will have most impact.

Hosting training compute alone does not in of itself produce AI safety, but there are safety conditions, postures, and practices that only jurisdictions that physically host compute can demand, hold and implement, which are outlined below.

Sequencing will matter. Australia is best placed to negotiate conditions and requirements *before* it enters into a deal with frontier AI labs to host training compute.

ACAIS recommends:

- **Mandatory incident reporting.** A hosting arrangement could require mandatory reporting to AISI on defined timelines for incidents such as containment failures, crossing of dangerous-capability threshold crossings, and significant misuse detections, including for internal AI models not yet publicly available. This would deepen bilateral engagement between AISI and the US Center for AI Standards and Innovation, and position Australia as a highly valuable contributor in international safety forums like the international Network for Advanced AI Measurement, Evaluation and Science (NAAIMES)
- **Security standards for training facilities:** Australia could make security standards for training facilities a condition of hosting, with the ASD involved in accreditation of facilities.
- **Development of verification technologies:** ideally all countries would comply with safety standard and measures prohibiting dangerous use of AI. If verification mechanisms are ever going to be tested, they need willing host jurisdictions that have training-scale facilities. Australia could offer to be the pilot site for training-run reporting and for verification regimes. That would see Australia make a tremendous global contribution, thereby generating even greater middle power influence.
- **Access to compute for safety research:** Training AI is compute-intensive, but some compute must be dedicated for external safety research worldwide. Australia could negotiate allocation of hosted compute to be reserved for safety and interpretability research, that is done through partnership with AI lab researchers and external researchers, increasing transparency, progressing

⁶ <https://www.anthropic.com/threat-intelligence-report-september-2026>

⁷ <https://cloud.google.com/blog/topics/threat-intelligence/ai-vulnerability-exploitation-initial-access>

critical science, and strengthening an Australian AI workforce. Australia would need to attract a technical workforce to undertake this research on the ground in Australia.

- **Lead international AI safety governance:** If Australia hosts training compute we can be a leader amongst nations that share a collective interest in not competing away AI security and safety requirements. Australia's precedent in founding and chairing the Australia Group, a multilateral export-control regime⁸ for chemical and biological precursors, positions us well for such middle power diplomacy.

Australia's prominent role in the Asia-Pacific also positions us well to engage in technical dialogue with China to strengthen AI safety science and measures. Afterall, regardless of where AI is developed, the risks are felt worldwide.

Biological threats

The risk of biological threats receives little if any attention in the Australian debate and warrants particular attention by the Committee.

The convergence of AI, biotechnology and bioengineering has removed the barriers that once confined bioweapons work to state laboratories⁹.

As with cyber threats, AI has democratised competence in biological practices. Researchers found LLMs consistently outperforming PhD-level virologists on the knowledge needed to run laboratory experiments by April 2025.

Specific AI systems can generate genomes for viruses not previously seen in nature¹⁰. These trends coupled with growth in design of genomic sequences and their physical synthesis via cloud labs, heighten risks of biological threats. Given the massive negative impact of such threats on populations, mitigating action is recommended.

What can Australia do?

Banning underlying biotechnologies is not realistic. They are dual use technologies and are economically indispensable. Australia relies on them to combat disease, boost agricultural productivity and to solve energy challenges.

ACAIS recommends:

- As highlighted by Australians for AI Safety¹¹, Australia's Biosecurity Import Conditions system (BICON) gives the Director of Biosecurity the power to impose **conditions on imported synthetic nucleic acids**. This should be done immediately, and can be achieved without new legislation. Members of the International Gene Synthesis Consortium, already voluntarily screen orders for sequences of concern. Australia's current import permit system does not require that synthetic nucleic acids be sourced from providers that screen, and bad actors can exploit this. The current arrangement poses an unacceptable risk under the *Biosecurity Act 2015*. This should be remedied immediately and quickly without the need for legislative change.
- Australia can also build its bioresilience by developing capabilities to **detect novel and known pathogens early**, for example, wastewater surveillance capable of flagging organisms that match

⁸ <https://www.dfat.gov.au/publications/minisite/theaustraliagroupnet/site/en/index.html>

⁹ https://www.foreignaffairs.com/united-states/ai-new-age-bioweapons-elizabeth-sherwood-randall?check_logged_in=1

¹⁰ King, S. H., Driscoll, C. L., Li, D. B., Guo, D., Merchant, A. T., Brix, G., ... & Hie, B. L. (2025). Generative design of novel bacteriophages with genome language models. *bioRxiv* doi:10.1101/2025.09.12.675911

¹¹ <https://www.australiansforaisafety.com.au/letters/ai-bio-gene-synth-screening>

no known list; and **enabling swift protection** by pre-negotiating surge capacity in biomanufacturing that would accelerate production of diagnostics, therapeutics and vaccines¹².

Challenges in attracting training compute in Australia

The Australian Copyright Act

ACAIS does not represent rightsholders or AI developers. We take no position on which legal instrument Parliament should choose in overcoming the current impasse that prevents training compute occurring in Australia by US frontier labs or any companies training AI, including Australian ones.

The status quo is the one outcome on which both sides should be able to agree is a failure.

Australia has neither a fair-use doctrine nor a text and data mining exception, unlike the United States, Japan and Singapore. A TDM exception was ruled out in October 2025.

We strongly urge urgency in resolving the current impasse, which we understand is follows:

- **The AI developers' position**, as we understand it, is not primarily about price.

As we understand it, AI developers' position – including large frontier AI companies through to smaller Australian AI companies who need to train LLMs – does not primarily relate to price or access to Australian data. Rather, paying for a training license in Australia amounts to an acknowledgement that a market for training data exists, which stands to weaken the fair-use defence on which AI companies' operations in the US depend. On this account frontier AI companies are license-sensitive rather than price-sensitive.

The committee should test this claim.

- **The copyright holders' position** is that licensing is the proper path and that consent and control are matters of principle, not merely of compensation. Australia's copyright framework operates as part of an international one, therefore repercussions of any changes have global implications. The Australian Government has ruled out a text and data mining exception and has committed to creators retaining ownership and control of their work, and to being compensated.
- **The status quo**, however, serves neither. Australian content is already being used in training conducted offshore. Current settings do not prevent that; they prevent the training by any AI company (a large frontier lab or a smaller Australian company) from happening here. Australians therefore bears the cost without the compensation, the control, the jobs, the infrastructure or the leverage. The status quo effectively impedes the ability of Australian companies to develop competitive AI models in Australia.

The status quo is the one outcome on which both sides should be able to agree is a failure.

Meeting large scale energy demand

Australia is well placed to benefit from the global energy transition towards cost-competitive renewable energy from wind and solar, supported by gas, hydro and batteries. Hosting data centers (particularly for training) transforms Australia's energy opportunity from a future expectation to a current reality, establishing a new long-term energy-intensive industry to support jobs, income and other economic benefits.

Public concern about the electricity price, land-use and community impacts of large data centres is legitimate: managed badly this could have adverse impacts.

Conditionality is how a host country converts infrastructure into public benefit.

¹² <https://securebio.org/>

ACAIS supports the Government's national data centre expectations published in March 2026, and the July 2026 commitment to legislate Australian Standards for AI including obligations on large data-centre operators to underwrite new power supply and meet their full share of grid connection costs.

Large, assured demand can de-risk and accelerate clean generation and transmission investment, if conditions are properly designed. This would allow existing clean energy ambitions and international commitments to be converted into compute capacity and a new long term internationally competitive industry.

The effects of grid-connected increased data center demand on prices, reliability and other system outcomes will depend on multiple factors: the location of the data centers and associated generation; the extent and terms of flexibility of data centers electricity demand; the build profile and location of batteries or other firming options; and the degree of shortfall or overbuild of generation vs average data centre demand. Effects will also be shaped by the time profile for each of these factors.

In addition, off-grid behind-the-meter configurations at remote sites, paired with firming, can in principle bypass grid constraints altogether, albeit at a higher financial cost – avoiding both the benefits and potential adverse effects of a grid-connected approach.

The best outcomes would also involve strong, timely and productive cooperation between the Australian Government and State and Territory governments to reduce "time to power".

ACAIS is in the early stages of commissioning modelling to assess potential effects of significant data center demand growth within the National Electricity Market, and the types of policy levers or system design choices available to deliver desired benefits and outcomes.

We would be pleased to provide results and insights to the Committee as these become available.

ENDNOTES

ⁱ SAFETY INSTITUTE RESOURCING SOURCES, METHOD AND CAVEATS

United Kingdom. £66m per financial year, stated by the AI Security Institute on its own About page, together with priority access to over £1.5bn of compute through the AI Research Resource and more than £15m in grant funding. The June 2025 Spending Review separately allocated £240m to the institute across 2026/27–2029/30; that envelope is a different measure and is not used here. Compute access and grant funds are excluded from the bar, so the UK figure understates total resourcing.

Singapore. S\$50m from IMDA and the National Research Foundation over five years from 1 October 2022, funding the Digital Trust Centre, which was designated Singapore's AI Safety Institute in May 2024. This is the host centre's research grant, not an AISI-specific appropriation, and therefore overstates funding for AI safety institute functions alone.

Canada. C\$50m over five years for the Canadian AI Safety Institute, announced in Budget 2024 and confirmed by Innovation, Science and Economic Development Canada at its November 2024 launch.

Australia. A\$29.9m to establish the AI Safety Institute, announced with the National AI Plan in December 2025 by the Minister for Industry and Innovation. The four-year profile and the A\$3.5m allocated in the current financial year were confirmed at Senate Estimates on 2 June 2026. One published account gives A\$29.4m; the Portfolio Budget Statements are the authority if the two need reconciling.

Annualisation. Multi-year commitments are divided by their stated term on a straight-line basis. Actual profiles are uneven-Australia's is weighted away from the current year - so these are indicative comparisons, not budget figures. Not strictly like-for-like.

Australia's is a four-year establishment commitment; the UK's is a recurring annual operating budget for an institute running since 2023. The comparison shows orders of magnitude, not equivalent line items.

Conversion. Indicative rates: £1 = A\$1.95, C\$1 = A\$1.02, S\$1 = A\$1.08. Refresh against Reserve Bank of Australia daily rates on the date of publication.

The UK institute additionally employs more than 100 technical specialists and has completed over 30 pre-deployment evaluations of frontier models (AI Security Institute; Ada Lovelace Institute, July 2026).