

**Australian Academy of Science submission on the
Getting it right: Building AI infrastructure that works for Australia consultation**

The Australian Academy of Science welcomes the opportunity to comment on the Office of AI's consultation: *Getting it right: Building AI infrastructure that works for Australia*.

The Australian research and development (R&D) sector require a strong policy framework to ensure it is prepared for AI. This framework is needed to guide the development of the necessary infrastructure, ethical guidelines, and educational programs for Australian researchers to harness AI's potential responsibly and effectively. A coordinated national approach to AI governance will provide clarity and consistency for researchers, support responsible adoption of AI technologies, and strengthen trust in AI-enabled research.

The consultation paper identifies the right risk – that Australia could host AI infrastructure and training while the research outputs, intellectual property and commercial returns accrue offshore. This submission focuses on three critical areas needed to address that risk: high-performance computing infrastructure, mechanisms to support the research and innovation system, and opportunities to strengthen skills, training and education.

The Academy's central point is that hosting data centres onshore does not bolster AI capability on its own. Capability depends on Australia investing in the high-performance computing and data (HPCD) infrastructure that underpins AI and data-intensive R&D by both the research sector and industry.

The Academy welcomes the consideration of mechanisms to support Australia's research and innovation system. As Australia is increasingly seen as an attractive destination for foreign data centre investment, this should be leveraged to strengthen domestic research, innovation and industry capability.

The Academy recommends:

- The Australian Government should develop a 10-year national high-performance computing strategy and investment plan to secure computing capability for AI and data-intensive research. The Academy proposes an investment of \$200 million per year over 10 years to support sector planning, upgrade facilities, and coordinate investment in Tier-0 (regional) capability.
- Mechanisms to secure contributions to Australian research should include a levy on data centres to ensure this extractive industry builds local capability, and to provide merit-based access to compute capacity under favourable terms for all Australian researchers.
- Hyperscalers and frontier AI developers must contribute to the development of a skilled workforce in Australia. They should also be required to make commitments to uplift researcher AI literacy through education initiatives or funding for programs.
- Mandatory assurance and incident-reporting requirements should be implemented for high-risk AI-agent training and evaluation conducted in, or using infrastructure located in, Australia.
- Public reporting on data centre energy use, water use and emissions should be frequent and standardised, and published in open, machine-readable formats consistent with FAIR data principles.

High-performance computing is a vital part of AI infrastructure

The consultation paper refers to investment in "facilities to support AI training", but makes no mention of high-performance computing and data (HPCD) more broadly, and the need for traditional computing for many modern services including those used by the Australian Government. This is a significant omission, particularly given that the [2026 National Research Infrastructure Roadmap Exposure Draft](#), released for public consultation

on 28 September, identifies HPCD as a national priority and highlights the urgent need for substantial investment in Australia's HPCD facilities.

Access to HPCD and software is a critical enabler for developing and adopting AI tools and underpins advances in AI and science. High-performance computers are specifically built to solve large, complex problems by processing massive datasets and running simulations at speeds far beyond standard computing.

Australian research and industry urgently need more investment and capacity for HPCD, for which demand has grown faster than government programs designed to support compute. This will not be addressed by the massive increase in planned data centre construction, which is largely driven by demand for commercial AI and data infrastructure. Without sufficient onshore HPCD capability, Australia will be dependent on other nations, hindering scientific progress and our ability to tailor AI tools to national challenges.

Australia's national HPCD facilities, the National Computational Infrastructure (NCI) and Pawsey Supercomputing Research Centre, require expansion and upgrades to ensure they are equipped to meet the growing demands of data-intensive research needs. Both host and train complex scientific AI models, such as predicting medical outcomes in intensive care, processing massive astronomical imaging datasets, and modelling extreme weather events. However, both facilities are already oversubscribed, with demand expected to continue increasing exponentially. Dedicated investment is needed to scale up GPU capacity for AI-driven research, but also CPU capacity for traditional computing and expanded data storage infrastructure.

Long-term, coordinated government investment is essential to ensure infrastructure keeps pace with research and industry needs and supports national productivity and innovation. The Australian Government should develop a 10-year national high-performance computing strategy and investment plan to secure computing capability for AI and data-intensive research.

The Academy has proposed an investment of \$200 million per year over 10 years to support sector planning, upgrade facilities, and coordinate investment in Tier-0¹ (regional) capability.² This figure is based on extensive consultation to determine the needs of the scientific community, and reflects the required infrastructure, skilled workforce, and operating costs determined by computing experts to be required to sustain a sovereign computing ecosystem over the next decade.

Contributions that will add the greatest value to Australia's R&D

The Academy welcomes the consideration of mechanisms that will add value to Australia's R&D system. Australia has an opportunity to create policy settings to ensure that private investment in establishing data centres delivers lasting value and creates a thriving digital ecosystem that supports research including research-industry partnerships. These would capture longer term benefits, beyond the short-term investment injection in the construction industry to build data centres.

Enforceable mechanisms to ensure that infrastructure providers maintain a skilled onshore workforce are limited, however building an onshore skilled workforce is important for Australia's resilience and sovereignty. Therefore, finding an option to secure long term benefit through investment in our workforce from these extractive industries is of critical importance.

One of the options proposed in the consultation paper is a public-interest research payment. This could be implemented through a levy on the revenue of hyperscalers, data centres and frontier AI developers. This would work similarly to the Rural Research and Development Corporation levy model already used in Australia's agricultural, fisheries and forestry sector. Funds raised from the public-interest research payments could be used for purposes such as investment in public national computational infrastructure, or establishment of an AI research fund to build Australia's sovereign AI capability and support research in the

¹ **Tiers of HPCD facilities:** Tier-0 (international or multinational facilities); Tier-1 (national facilities); Tier-2 (institution-based facilities).

² Australian Academy of Science. (2025, February 18). *Bringing Australia's supercomputing up to speed*. <https://science.org.au/our-work/resources-reports/reports-publications/bringing-australias-supercomputing-speed>

national interest. This way, the investment would grow a workforce as well as national infrastructure and the research and industrial base.

The Academy has proposed an indicative rate of 2.5% on company revenue. This requires further analysis by the Australian Treasury to determine the most favourable arrangement. The levy could potentially be offset by quantifiable onshore R&D investments eligible for the R&D Tax Incentive (RDTI). The Academy has previously modelled a [broad R&D levy](#) that could be used as a basis.

Another contribution proposed in the consultation paper is access to compute capacity through a reservation scheme. The Academy supports the provision of compute capacity at no cost for non-commercial, publicly funded research. This should be allocated through an equitable, merit-based process.

This could be allocated through expanding the existing National Computational Merit Allocation Scheme (NCMAS). NCMAS can enable a consistent, equitable application process. This would efficiently help to coordinate access, rather than placing a new burden on researchers to apply to various individual providers. Dedicated technical expertise will be needed to support researchers to use the compute. This would require governance and operational uplift to NCMAS, that would require further investment.

Opportunities for frontier AI developers to support skills, training and education

While AI is an important opportunity, its value depends upon maintaining strong underlying scientific capability and human expertise. The influence of AI on Australia's economy and security will depend not only on infrastructure, but also on the research workforce needed to understand and use these technologies effectively and responsibly.

The Academy has published a series of discussion papers titled '*Is Australian science ready for AI?*' that explore the preparedness of the Australian science sector for AI advances.³ This series highlighted a growing need to equip researchers with AI skills, suggesting the value of integrating these capabilities into research training programs.

Frontier AI developers should measurably support skills, training and education opportunities through funding for programs such as:

- Teaching resources (e.g. the [Australian Academy of Science's education resources](#) and the [Australian Academy of Technological Sciences and Engineering STELR program](#));
- Support for workforce, diversity and equity programs (e.g. [Science & Technology Australia's Superstars of STEM program](#) – this could be targeted towards supporting researchers working in computing and/or AI);
- Scholarships and cadetships;
- Partnerships with organisations;
- Funding industry PhDs; and
- Data literacy programs for researchers.

These could be delivered in partnership with the National AI Centre or Office of AI to target the areas of greatest need.

The standards should include mandatory assurance and incident-reporting requirements

The Academy agrees with the need stated in the consultation paper for mandatory assurance and incident-reporting requirements for high-risk AI-agent training and evaluation conducted in, or using infrastructure located in, Australia.

Recent incidents suggest that Australia's AI standards should explicitly cover agent training and evaluation environments. In the recent Hugging Face incident, OpenAI reported that models operating with reduced safeguards during internal cybersecurity evaluations circumvented isolation controls, communicated through

³ Australian Academy of Science (2025). *Is Australian science ready for AI?* <https://science.org.au/our-work/resources-reports/reports-publications/australian-science-ready-ai>

unauthorised channels, gained internet access and affected third-party systems. Separately, an internal OpenAI research agent gained unauthorised access to public and non-public files in Services Australia's Medicare Statistics Reporting Service after encountering repeated access blocks, and wrote files to an internal server, according to the Australian Government.

These incidents illustrate an important gap. Model-level behavioural safeguards may sometimes be deliberately reduced during evaluation to test capabilities and worst-case scenarios. In those circumstances, mandatory compensating controls should remain in place, including:

- i. documented authorisation, scope and safety cases for high-risk evaluations;
- ii. isolated sandboxes, default-deny network access, allow listed external services, least-privilege credentials and human approval for consequential actions;
- iii. monitoring independent of the agent, tamper-evident logs, automated pause or termination mechanisms, and safe-stop behaviour when a task cannot be completed within its authorised boundaries; and
- iv. clear allocation of responsibilities among model developers, compute providers and infrastructure operators.

The standards should also define reportable AI incidents and material near misses before deployment. These should include unauthorised boundary crossing, side-channel communication between agents, unauthorised third-party access or data transfer, exposure of credentials or model weights, and failures of containment or monitoring. Requirements should cover prompt notification to affected organisations within a defined period of time and to an appropriate government body, which should be responsible for preservation of forensic evidence, confidential handling of sensitive technical details, and aggregated public reporting where appropriate.

Public reporting of data collected

Public reporting on data centre energy use, water use and emissions should be frequent and standardised, and published in open, machine-readable formats consistent with FAIR data principles. This will allow for researchers to be able to examine the environmental impact, including how data centres being built are impacting our energy mix and transition to clean energy.

To discuss or clarify any aspect of this submission, please contact Lauren Sullivan, Science Policy and Advice Manager at science.policy@science.org.au.