

From:
To:
Subject: Media re: AI and copyright
Date: Monday, 1 June 2026 1:53:22 PM

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Quick heads. [s 22](#) reached out to us about the copyright and data centres discussion paper, [s 47C](#)

We've responded, explaining that the proposal is the opposite of a text-and-data mining exemption. A TDM gives AI companies a free pass to train AI without obligations. The proposal extracts maximum value from AI companies based on what they value most, access to the Australian market. We've also pointed out our track record of calling for more AI regulation.

We haven't mentioned anyone specific that we've shared the paper with.

[s 47C](#)

Thanks
[s 22](#)

From:
To: s 22 @aph.gov.au
Cc: s 22 @aph.gov.au; Ayres, Tim (Senator); Charlton, Andrew (MP Office);
s 22 @aph.gov.au; s 22 @aph.gov.au; s 22 @aph.gov.au; s 22
@aph.gov.au; s 22 @aph.gov.au; s 22 @aph.gov.au; s 22
@aph.gov.au; s 22 @aph.gov.au;
s 22 @aph.gov.au; s 22 @aph.gov.au; s 22
Subject: Open letter to Minister Collins: AI and biosecurity experts urge action on gene synthesis screening
Date: Monday, 4 May 2026 2:48:05 PM
Attachments: [FINAL] Activate existing biosecurity powers to address AI-enabled risks.pdf

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Dear Minister Collins,

I'm writing to convey the attached letter, signed by over 125 AI experts, biosecurity experts and members of public.

The letter raises concerns about risks at the intersection of AI and synthetic biology. The letter makes three recommendations.

The online version of the letter is available [here](#).

I've CC'd other parliamentarians who have an interest in artificial intelligence, biosecurity, national security and other matters raised by the letter.

I coordinated the letter and am able to facilitate any response to the signatories. I'm available to meet to discuss, including with available expert signatories.

Kind regards

s 22

s 22

The Hon Julie Collins MP
Federal Minister for Agriculture, Fisheries and Forestry
email: s 22 @aph.gov.au

4 May 2026

Dear Minister,

Australia has one of the world's most rigorous biosecurity frameworks, and a new gap has emerged that it is well-placed to close. AI and synthetic biology now pose increasing and novel risks that require immediate action from your department.

Synthetic nucleic acids can be used to develop vaccines, improve crops, and advance medical research, but also to engineer or construct pathogens. Australians can order custom nucleic acids from overseas providers. Previously, the key constraint in creating a biological weapon from synthetic nucleic acids was the knowledge required — expertise that was rare and hard to acquire. That has changed.

The International AI Safety Report 2026 confirms that general-purpose AI systems can now provide expert-level guidance on biological and chemical weapons development, including detailed laboratory instructions.¹ One study found an AI model outperforming 94% of domain experts at troubleshooting virology protocols.² Specific AI systems (biological foundation models) can generate genomes for viruses (bacteriophages) not previously seen in nature.³ Related techniques have already been used to design protein variants that evade the human immune system.⁴

Current AI practices do not adequately address this risk: only 3% of 375 biological AI tools surveyed have any safeguards,⁵ and where they exist, ways to circumvent the restrictions (jailbreaks) are common.⁶

A bad actor with AI guidance still needs the materials to act, and imported synthetic nucleic acids are a physical chokepoint. Australia's Biosecurity Import Conditions system (BICON) gives the Director of Biosecurity the power to impose conditions on all such imports — immediately, without new legislation. Members of the International Gene Synthesis Consortium, including leading providers such as Twist Bioscience, IDT, and GenScript, already screen orders for sequences of concern voluntarily. Yet, Australia's current import permit system does not require that synthetic nucleic acids be sourced from providers that screen, and bad actors could exploit that.

Australia has committed to reducing biosecurity risks to a very low level.⁷ The science now provides reasonable grounds for biosecurity officers to conclude that synthetic nucleic acids imported without safety screening pose an unacceptable risk under the *Biosecurity Act 2015*.

The US,⁸ New Zealand,⁹ the EU,¹⁰ and the UK¹¹ are each taking steps to address screening requirements. Australia can do the same, and without new legislation.

The National AI Plan says that regulators are responsible for identifying and addressing AI-related harms within their domains.¹² We call on you to act on that responsibility.

We recommend that:

1. **A screening condition be applied to all BICON import permits for synthetic nucleic acids.**
Permits should require that synthetic nucleic acids be sourced from providers that screen orders for sequences of concern and, where sequences of concern are identified, verify customer identity and legitimacy. Appropriate mechanisms should be in place to verify provider compliance.
2. **Applicants for higher-risk synthetic nucleic acid import permits are prioritised for the statutory fit and proper person test¹³** – a government background check that goes beyond provider screening – with the department publishing clear criteria for when and how the test is applied.
3. **The department reviews the adequacy of the regime within its broader context**, in consultation with relevant Commonwealth agencies, including the new AI Safety Institute, civil society, industry, and academia. The review should address progress made in AI and biotechnology as well as emerging issues such as domestic production of synthetic nucleic acids, control on benchtop DNA synthesisers, and mechanisms for verifying provider compliance.

These steps should be accompanied by appropriate transparency to signal to the public, the global community, and potential bad actors that Australia is taking action.

Implementing these measures is low cost.¹⁴ Screening is conducted digitally by the synthesis provider before dispatch, imposing no additional burden on Australian researchers, businesses, or the regulator. Free screening tools are available.^{15,16} Major providers already screen voluntarily – this change targets the gap, not the norm.

This is an opportunity for Australia to address domestic risks and build on its world-leading biosecurity framework. As a founding member of the Australia Group, Australia has a track record of international biosecurity leadership. Mandating gene synthesis screening would be a significant step toward establishing a global norm and addressing AI-biosecurity risk.

Yours faithfully,

The undersigned

Dr Cassidy Nelson

DPhil MBBS MPH
Director of Biosecurity Policy
Centre for Long-Term Resilience

Associate Professor Lotti Tajouri

Associate Professor
Bond University and Murdoch
University

Dr Brendan Walker-Munro

Associate Professor
Southern Cross University

Dr Alexander Saeri

Director, AI Risk Initiative
MIT FutureTech

Professor Patrick F Walsh

Professor, Intelligence and Security
Studies
Charles Sturt University

Ms Kate Chaney MP

Federal Member for Curtin



Associate Professor David Heslop

FAFOEM FRACGP MBBS PhD MPH
BSc(Adv) Hons 1 MAICD
Associate Professor
University of New South Wales

Janet Egan

Senior Fellow and Deputy Director
Center for a New American
Security

Dr David Manheim

Director of Research and Policy
Technion – Israel Institute of
Technology / Alter.org.il

Dr Peter Slattery

Research Scientist
Massachusetts Institute of
Technology

Dr Sam Buckberry

Head of Epigenetics
The Kids Institute Australia,
Australian National University

Dane Sherburn

CEO
p-zero research

Dr Ryan Kidd

Co-Executive Director
MATS Research

Keltan O'Shea

The Machine Intelligence Research
Institute (MIRI)

Karl Berzins

Co-founder & President
FAR.AI

John Pane

Chair
Electronic Frontiers Australia Inc.

**Associate Professor Gert
Frahm-Jensen**

MBBS FRACS(Vasc) BBiotech

Dr Toby Ord

Senior Researcher
Oxford University
☆ Author of *The Precipice: Existential
Risk and the Future of Humanity*

Brayden McLean

Technical Program Manager, Safety
& Security
Anthropic

Noel Lim

CEO
Anika Legal
☆ 2025 Victorian of the Year

Dr Saskia Popescu

CEO / Policy Researcher
Global Health Security Network /
RAND

Scott Weathers

Associate Director of Government
Affairs
Americans for Responsible
Innovation

Mr Rumtin Sepasspour

Director of Policy and Strategy
Global Shield

Mr Devon Whittle

Australia Director
Global Shield Australia

Mr Michael Clark

Director
Cytophenix

Dr Duncan Purvis

Volunteer with Australians for
Pandemic Prevention

Dr Kun Zhao

Senior Researcher, AI Risk Initiative
MIT FutureTech

Dr Sid Sharma

MD MPH FAFPHM
Public Health Physician

Mr Soroush Pour

CEO
Harmony Intelligence
☆ Fmr Head of Technology at Vow
(World-leading biotech firm)

Dr Sarah Winthrop

Visiting Fellow
Brown University Pandemic Center

Professor Nick Wilson

Research Professor
University of Otago (New Zealand)

Professor David Balding

Honorary Professor of Statistical
Genetics
University of Melbourne

Dr Adam Kamradt-Scott

Cummings Foundation Associate
Professor of One Health Diplomacy
The Fletcher School of Law and
Diplomacy, Tufts University

**Associate Professor Michael
Noetel**

Associate Professor
University of Queensland

Mr Jay Bailey

Head of Technology and Standards
Arcadia Impact

Michael Kerrison

Executive Director
AI Safety Australia & New Zealand

Mr Greg Sadler

CEO
Good Ancestors

Chris Leong

Lead Organiser
Sydney AI Safety Fellowship

Peter Horniak

Director
PauseAI Australia

**Emeritus Professor Thea van de Mortel**

Emeritus Professor
Griffith University

Mr Yanni Kyriacos

Director
*Technical Alignment Research
Accelerator*

Dr Mark Brown

PauseAI Australia

Hugo Lyons Keenan

Machine Learning PhD Student
The University of Melbourne

Dr Tim Seelig

Adjunct Associate Professor
University of Queensland

Nathan Sherburn

Chief Technology Officer
Effective Altruism Australia

Ms Emily Grundy

Policy Officer
Good Ancestors

Mr Lawson Pegler

Head of Growth
Elephant Ed

Bryce Robertson

Project Director
*Alignment Ecosystem
Development*

Andrew Harris

Co-Founder
Melbourne Security Forum

Jocelinn Kang

Nodalys

Sandy Fraser

AI Researcher

Dr Mark Zirnsak

Senior Social Justice Advocate
*Uniting Church in Australia, Synod of
Victoria and Tasmania*

Mr Jimmy Farrell

EU AI Policy Lead
Pour Demain

David Conrad

Managing Director
Talos Network

Dr Jamie Freestone

AI safety researcher

Oscar Delaney

Grants Associate
Macroscopic Ventures

James Norris

Executive Director
Center for Existential Safety

Mr Kevin Rassool

Technical Director
High Impact Athletes

Mr Ramakrishnan Veeramony

MBA, MAICD
Managing Director
Atinar Pty Ltd

Dr Sam Coggins

AI Risk Governance Researcher
Australian National University

Dr Evan Hockings

Member of Technical Staff
Iceberg Quantum

Dan Braun

Goodfire

Mr Mitchell Laughlin

Economist

Luke Freeman

COO
Good Ancestors

Mr Samson Blackburn

AI Architect, Mgr AI Engineering

Mr Steven Merriel

AI Safety Engineer

Ms Angelica Chowdhury

Data Scientist & Independent
Researcher

Ms Emma Humphrey

New Zealand Community Lead and
Ecosystem Builder
AI Safety ANZ

Michael Huang

Co-Lead
PauseAI Australia

Mr David Colin Gould

PauseAI Australia

Ms Holly Massacci

PhD Candidate in Genomics and
Bioinformatics
*The Kids Research Institute
Australia*

Mr Matt Fisher

Senior AI Evaluations Engineer
Arcadia Impact

Mr Martin Veron

Doctoral Candidate
University of Queensland

Edward Pierzchalski

Senior Software Engineer

Mr Yoshua Wakeham

Senior Software Developer

**Dr Ariel Zeleznikow-Johnston****Mark Freeman**

Associate Professor (retired)
University of Sydney

Ramneek Singh Matharu

University Student

Steven Deng

Energy Consultant

Ben Auer

Student
University of Melbourne

Mr Nathan Sidney

Business Coordinator

Ms Tamara van Noort**Mr Hunter Cole****Huw Cannon****Zeke Coady****Miss Shambhavi Pandey****Fergus Dall****Mr Ewan Dewar****Mr Rafe Skidmore****Andrew van Noort****Ms Catherine Sullivan****Augustus Hebblewhite****Dr Simon Zhang****Danny Smith**

Research Software Engineer

Ms Alana Stewart

Clinical Pharmacist

Ms Mandy Collins

Artist and educator

Mr Samuel de Pury

Therapist
Holina Healing Centre

Jemima Harvey

Teacher

Mr Ron Green**Ms Monica Lacey****Mr Nelson Gardner-Challis****Mr Elliot Teperman****Chantelle Ring****Kavi Townsend****Mrs Kathleen Grundy****Dr Eugene Lubarsky****Rebecca Hawkins****Gaetan Selle****Rebecca Niven****Dr James Kimber****Ms Stephanie Symes**

Teacher
FCJ College

Ty Wilson-Brown

Senior IT Professional

Mr Frederick Tropp-Asher

I.T. Project Manager (retired)

Jen Truong

Student
University of Melbourne

Kate Hanson

Law

Bridget Loughhead**Rickard Vikstrom****Mark Carter****Mr Austin Zhang****Amanda Littlewood****Mr Pierre Taylor****Mr Lee Carroll****Mr Daniel H****Mr Stephen Ingram****Ms Sarah Grundy****CC:**

Prime Minister

Minister for Industry and Innovation and Minister for
Science

Assistant Minister for Science, Technology and the Digital
Economy

Minister for Home Affairs

Attorney-General

Minister for Health and Ageing

Minister for Foreign Affairs

Leader of the Opposition

Shadow Minister for Industry and Sovereign Capability

Shadow Minister for Digital Economy, Science, Technology
and Innovation

Shadow Minister for Home Affairs and Immigration

Shadow Attorney-General

Shadow Minister for Agriculture

Shadow Minister for Health and Aged Care

Australian Greens Spokesperson for Digital Rights & IT

Independent Senator for the ACT

Endnotes

- [1] Bengio, Y., Clare, S., Prunkl, C., Andriushchenko, M., Bucknall, B., Murray, M., ... & Mindermann, S. (2026). International AI Safety Report 2026. <https://apo.org.au/node/333514>
- [2] Götting, J., Medeiros, P., Sanders, J. G., Li, N., Phan, L., Elabd, K., ... & Donoughe, S. (2025). Virology Capabilities Test (VCT): A multimodal virology Q&A benchmark. arXiv [cs.CY] <https://arxiv.org/abs/2504.16137>
- [3] 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
- [4] Youssef, N., Gurev, S., Ghantous, F., Brock, K. P., Jaimes, J. A., Thadani, N. N., ... & Marks, D. S. (2025). Computationally designed proteins mimic antibody immune evasion in viral evolution. Immunity 58, 1411–1421.e6 doi:10.1016/j.immuni.2025.04.015
- [5] Villalobos, P. & Atanasov, D. (2025). Announcing our expanded biology AI coverage. Epoch AI <https://epoch.ai/blog/announcing-expanded-biology-ai-coverage>
- [6] Jin, H., Hu, L., Li, X., Zhang, P., Chen, C., Zhuang, J., & Wang, H. (2024). JailbreakZoo: Survey, landscapes, and horizons in jailbreaking large language and Vision-Language Models. arXiv [cs.CL] <https://arxiv.org/abs/2407.01599>
- [7] Department of Agriculture, Fisheries and Forestry. (n.d.). Appropriate level of protection. Australian Government. <https://www.agriculture.gov.au/biosecurity-trade/policy/risk-analysis/conducting-import-risk-analysis/appropriate-level-of-protection>
- [8] The White House. (2025). Improving the safety and security of biological research. <https://www.whitehouse.gov/presidential-actions/2025/05/improving-the-safety-and-security-of-biological-research/>
- [9] New Zealand Parliamentary Counsel Office. (2025). Gene Technology Bill 110-2 (2024). <https://www.legislation.govt.nz/bill/government/2024/0110/latest/whole.html>
- [10] European Commission. (2025). Proposal for a regulation to establish measures to strengthen the Union's biotechnology and biomanufacturing sectors. https://health.ec.europa.eu/publications/proposal-regulation-establish-measures-strengthen-unions-biotechnology-and-biomanufacturing-sectors_en
- [11] UK Department for Science, Innovation and Technology. (2024). UK screening guidance on synthetic nucleic acids for users and providers. GOV.UK. <https://www.gov.uk/government/publications/uk-screening-guidance-on-synthetic-nucleic-acids/uk-screening-guidance-on-synthetic-nucleic-acids-for-users-and-providers>
- [12] Department of Industry, Science and Resources. (2025). National AI plan. Australian Government. <https://www.industry.gov.au/sites/default/files/2025-12/national-ai-plan.pdf>
- [13] Department of Agriculture, Fisheries and Forestry. (n.d.). Fit and proper person test. Australian Government. <https://www.agriculture.gov.au/biosecurity-trade/import/arrival/arrangements/fit-proper-person>
- [14] Centre for Long-Term Resilience. (2025). Cost-benefit analysis of synthetic nucleic acid screening for the UK. <https://www.longtermresilience.org/reports/cost-benefit-analysis-of-synthetic-nucleic-acid-screening-for-the-uk/>
- [15] IBBIS. (n.d.). Common mechanism. <https://ibbis.bio/our-work/common-mechanism/>
- [16] Nuclear Threat Initiative. (2022). Preventing the misuse of DNA synthesis technology. <https://www.nti.org/about/programs-projects/project/preventing-the-misuse-of-dna-synthesis-technology/>

From: s 22
To: s 22 Rafizadeh, Shervin
Subject: Re: Catchup - AMO Charlton / s 22 [SEC=OFFICIAL]
Date: Monday, 18 May 2026 5:49:55 PM
Attachments: [Good Ancestors - An Australian AI Training Scheme.pdf](#)

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Hi all

Looking forward to meeting tomorrow.

Please find attached the latest version of the paper. It only has minor changes:

- explaining the idea of "decoupling" between the permit and the fund more clearly (including that it's important to solve constitutional law issues)
- calling the fee "tiered" rather than "variable" based on feedback that "variable" made it sound as if it was subject to arbitrary change
- explaining that "opt out" needs to be practical for both AI companies and creators / rights holders

I've also sent this version to AGD. Good Ancestors is now a member of the CAIRG and this is our submission to that process. Obviously still happy to refine the paper if / before we use it for any public purpose.

Thanks

S

From: s 22
Sent: Friday, 15 May 2026 2:40 PM
To: s 22 @industry.gov.au>; s 22 Rafizadeh, Shervin s 22 @cs.pm.gov.au>; s 22 @industry.gov.au>
Subject: Re: Catchup - AMO Charlton / s 22 [SEC=OFFICIAL]

Good afternoon all

Ahead of the meeting next week, please find attached a draft paper on an Australian AI Training Scheme.

The proposal is to break the gridlock between the copyright industry and AI companies by separating the conversation into two halves. That gives room to strike two separate deals that work for each group.

The first is an AI Training Permit — legal certainty to train in Australia, in exchange for a fee, an effective opt-out for rights holders, and compliance with the National Data Centre Expectations. The second is an Australian AI Public Benefit Fund, modelled on the Public Lending Right Act 1985, to compensate Australian creators and support the broader creative industry. The specifics of that fund can be whatever works for creators and rights holders.

The two are deliberately decoupled, which preserves independent policy levers and avoids the fund being characterised as a de facto licensing scheme.

Looking forward to chatting.

S

From: s 22
Sent: Friday, 15 May 2026 12:14 PM
To: s 22 s 22 s 22 @industry.gov.au>; Rafizadeh, Shervin s 22 @cs.pm.gov.au>
Subject: Catchup - AMO Charlton / s 22 [SEC=OFFICIAL]
When: Tuesday, 19 May 2026 11:00 AM-11:30 AM.
Where: Microsoft Teams Meeting

OFFICIAL

s 22

s 22

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 leverage 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.

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.

From: s 22
To: [Rafizadeh, Shervin](#); s 22
Subject: Catchup - AMO Charlton / s 22 SEC=OFFICIAL]

Now from correct diary !

Microsoft Teams meeting

s 22

From: s 22
To: s 22 s 22 ; s 22 ; Rafizadeh, Shervin
Subject: RE: Catchup - AMO Charlton / s 22 [SEC=OFFICIAL]
Date: Friday, 15 May 2026 2:15:00 PM

Thanks s 22

Also apols – sent the calendar invite from the wrong email. Will send one from my diary !

From: s 22
Sent: Friday, 15 May 2026 2:41 PM
To: s 22 @industry.gov.au>; s 22
Rafizadeh, Shervin s 22 @cs.pm.gov.au>; s 22
@industry.gov.au>
Subject: Re: Catchup - AMO Charlton / s 22 [SEC=OFFICIAL]

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Good afternoon all

Ahead of the meeting next week, please find attached a draft paper on an Australian AI Training Scheme.

The proposal is to break the gridlock between the copyright industry and AI companies by separating the conversation into two halves. That gives room to strike two separate deals that work for each group.

The first is an AI Training Permit — legal certainty to train in Australia, in exchange for a fee, an effective opt-out for rights holders, and compliance with the National Data Centre Expectations. The second is an Australian AI Public Benefit Fund, modelled on the Public Lending Right Act 1985, to compensate Australian creators and support the broader creative industry. The specifics of that fund can be whatever works for creators and rights holders.

The two are deliberately decoupled, which preserves independent policy levers and avoids the fund being characterised as a de facto licensing scheme.

Looking forward to chatting.

s 22

From: s 22
Sent: Friday, 15 May 2026 12:14 PM

To: s 22

s 22

[@industry.gov.au](mailto:industry.gov.au)>; Rafizadeh, Shervin

s 22

[@cs.pm.gov.au](mailto:cs.pm.gov.au)>

Subject: Catchup - AMO Charlton / s 22 [SEC=OFFICIAL]

When: Tuesday, 19 May 2026 11:00 AM-11:30 AM.

Where: Microsoft Teams Meeting

OFFICIAL

s 22

From: s 22
To: s 22
Subject: Re: catch-up times [SEC=OFFICIAL]
Date: Wednesday, 25 March 2026 3:41:26 PM
Attachments: [image001.png](#)
[image002.png](#)
[Anthropic is visiting - What you need to know about data centres.pdf](#)

OFFICIAL

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Hi s 22

Hope you're well. Just letting you know we've pulled together a 2-page brief on the Anthropic visit, which we'll send around to MPs / Senators / Journalists.

We cover:

- Why AI companies want to build here, the benefits and the downsides
- How copyright has become a blocker and what a good solution could look like

The brief tries to be factual. To the extent it implies a recommendation, it's that government can make a deal that benefits both investment and creators. This would require movement on a copyright training exemption.

This is a slightly uncomfortable landing for us, given usually we're calling for more regulation on the AI industry. But, in this case, I think they have the stronger arguments overall if the key risks are mitigated.

Happy to chat if it helps and feel free to share the note with others if you think it's helpful.

Cheers

s 22

OFFICIAL

From: s 22 @industry.gov.au>
Sent: Tuesday, 3 March 2026 11:29 AM
To: s 22
Subject: RE: catch-up times [SEC=OFFICIAL]

OFFICIAL

QT? I feel like Aussies is a bit dark and I don't get any sunlight

OFFICIAL

From: s 22
Sent: Tuesday, 3 March 2026 9:21 AM
To: s 22 @industry.gov.au>
Subject: Re: catch-up times [SEC=OFFICIAL]

OFFICIAL

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

That's great, Aussies?

OFFICIAL

From: s 22 @industry.gov.au>
Sent: Monday, 2 March 2026 6:31 PM
To: s 22
Subject: RE: catch-up times [SEC=OFFICIAL]

OFFICIAL

Thanks s 22

How about coffee on Tuesday 10 March at 10am?

OFFICIAL

From: s 22
Sent: Monday, 2 March 2026 2:35 PM
To: s 22 @industry.gov.au>
Subject: catch-up times

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Hi s 22

Thanks for the quick chat last week.

I'm up at Parliament for a meeting tomorrow between 1115-1200. I'd be happy to grab a coffee before or after that if you happen to be around.

Otherwise, I could do after lunch on Thursday 5 or the morning of Tuesday 10.

Let me know if any of those options work for you.

Cheers

s 22

From: s 22

Sent: Thursday, 5 February 2026 2:09 PM

To: s 22 <s22@industry.gov.au>

Subject: International AI Safety Report 2026: implications for Australia — briefing available

Hi s 22

Hope you're well and had a good holiday period. Thanks again for engagement in 2025 on AI risks and opportunities. It's great to see the AI Safety Institute standing up and implementation of the Plan underway. I've done a fair bit of media and podcasts etc over the last few months and really enjoyed being able to report good news.

I'm writing because, as you know, this week the [International AI Safety Report 2026](#) was released, capturing the current state of the science on AI risks.

Good Ancestors was one of a handful of civil society organisations globally that contributed to the Report, and we're offering to meet with all MPs, Senators and their teams to brief on its findings.

The key findings of the Report are:

- **AI capabilities have continued to improve**, especially in mathematics, coding, and autonomous operation. AI companies are investing heavily in 'AI agents' – autonomous systems that can perform tasks with little to no human oversight.
- **Advances in AI's scientific capabilities have heightened concerns about misuse in biological weapons development.** AI systems can provide detailed information relevant to biological weapons development. AI agents can complete complex tasks, including providing accessible interfaces to help users operate specialised laboratory equipment.
- **AI systems are increasingly used in real-world cyber operations**, they can

discover software vulnerabilities, write malicious code and automate parts of cyberattacks.

- **AI models can sometimes identify when they are being evaluated**, allowing them to deceive testers. This has made safety testing harder.
- **AI systems could pursue goals that conflict with human interests.** Some AI researchers and company leaders believe loss of control is a serious possibility, with consequences potentially including human extinction. Others consider such scenarios implausible. This disagreement reflects different assumptions about what future AI systems will be able to do, how they will behave, and how they will be deployed.
- **Information asymmetries mean that policymakers often lack information about AI systems.** AI developers have information about their products – such as training data and evaluations – that they often do not share with policymakers and researchers, limiting external scrutiny.

The Report is about the science, so makes no policy recommendations. Good Ancestors can fill this gap, offering perspectives on what the science means for Australia. For instance, Australia's National AI Plan says that Home Affairs is coordinating a multiagency group on synthetic biology and AI. Given the emerging science around AI misuse for bioweapons, Australia could act by mandating safety screening for synthetic DNA importation and regulating desktop DNA synthesisers in Australia.

This specific kind of action would send a positive signal nationally and globally that Australia is taking its own plan seriously and is following developments in the science closely. As you know, I was excited that this made it into the plan, and it would be great to see it turn into concrete action.



I'm obviously happy to meet with Mr Charlton, but would also appreciate an opportunity to catch up with you more generally on plans for the year and any ways Good Ancestors can be involved.

Kind regards

s 22

s 22



From: s 22 <s22@industry.gov.au>

Sent: Monday, 22 December 2025 11:55 AM

To: s 22

Subject: RE: time to chat? [SEC=OFFICIAL]

OFFICIAL

Thanks for this s 22 – and apologies for the belated response while I was on leave.

I saw you've connected with the Department directly on this, which is great. Appreciate you also drawing this to the attention of the experts globally.

Thanks again, and have a good Christmas.

s 22

OFFICIAL

From: s 22

Sent: Tuesday, 9 December 2025 4:30 PM

To: s 22 <s22@industry.gov.au>

Subject: Re: time to chat? [SEC=OFFICIAL]

OFFICIAL

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Hi s 22

Hope you're well, and thanks for keeping me looped in on the AISI work at the AI Plan.

As foreshadowed, Good Ancestors has been in touch with AI safety talent, including those at leading AI companies, not-for-profits, academia, and other countries' AISIs. We've received almost 150 responses to questions like what the "deal makers" and "deal breakers" for them working at an Australian AISI would be and what kinds of risks and research programmes they think would make an important contribution. As you know, these folk aren't browsing APS jobs, so some kind of direct connection is essential.

We've also engaged with senior figures in this space to hear their views directly.

I've attached a 3-pager that digests some of the quantitative data and shares initial findings. I'll also share this with the Department and Minister Ayres' office.

As a next step, we'll produce a more detailed analysis of the full data to share.

Happy to chat at any time, and really keen to see the AISI go well.

Kind regards

s 22

OFFICIAL

From: s 22 <s22@industry.gov.au>

Sent: 24 November 2025 11:19
To: s 22
Subject: RE: time to chat? [SEC=OFFICIAL]

OFFICIAL

Fantastic, thanks heaps!

OFFICIAL

From: s 22
Sent: Monday, 24 November 2025 11:15 AM
To: s 22 <s22@industry.gov.au>
Subject: Re: time to chat? [SEC=OFFICIAL]

OFFICIAL

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Hi s 22 - as discussed

How Australia navigates frontier AI will determine our prosperity and security. A world-leading AI Safety Institute is one of the best ways to keep Australians safe while making our contribution to a global challenge. An AISI will give Australia the technical expertise to understand advanced AI, contribute to preventing serious risks, and put Australia on the path to global leadership on safe and beneficial AI.

Quote attributable to s 22

Ta
s 22

OFFICIAL

From: s 22
Sent: 24 November 2025 09:28
To: s 22 <s22@industry.gov.au>
Subject: Re: time to chat? [SEC=OFFICIAL]

Hi s 22

Yes, my availability is pretty good. Other than 1:00-2:00 I can be free easily.

I'm on s 22 . Just let me know when suits you.

Ta
s 22

From: s 22 @industry.gov.au>
Sent: 24 November 2025 08:27
To: s 22
Subject: time to chat? [SEC=OFFICIAL]

OFFICIAL

Hey there

Do you have any windows to speak today?

Cheers
s 22

s 22

Office of the Hon Dr Andrew Charlton MP
Cabinet Secretary
Assistant Minister for Science, Technology and the Digital Economy
Federal Member for Parramatta
s 22

@industry.gov.au

OFFICIAL



Anthropic is visiting Australia:

What you need to know about AI data centres

25 March 2026

Anthropic—the AI company behind Claude—is visiting Australia this month. AI data centres will be their top priority.

On 23 March, the Australian Government released its [National Data Centre Expectations](#), setting out what it requires from developers – but important questions remain. **Why might data centres be good for Australia, what national interests do we need to protect, and is a deal possible?**

AI data centres have real benefits

The [National AI Plan](#) says AI is an economic and national priority, and that “it is in our interests to ensure that AI development happens locally, where it aligns with our national priorities, social and economic interests”. Data centres are already at the core of Australia’s digital infrastructure, underpinning AI, cloud, and other digital services.

Australia’s high-skill service-based economy is exposed to AI disruption. Managed poorly, AI would move a significant amount of wealth from Australia to AI leaders, likely the US and China.

In addition to being valuable today, data centres **“hedge” against future AI risks**. If AI becomes more disruptive, data centres will become more valuable. Countries with AI infrastructure also have a “seat at the table” to shape standards, require compliance with domestic law, and hold companies to account. Countries without a foothold have fewer options for capturing benefits or mitigating risks.

Australia was the [second-largest](#) destination for data centre investment in 2024, behind the US, with over \$100 billion in projects announced. US tech companies are investing over [USD \\$650 billion](#) in AI infrastructure this year alone. Energy is a [key constraint](#) – and Australia has abundant [renewable energy](#), land, and stable democratic governance. Investment that doesn’t come to Australia may instead go to [Gulf states](#) or stay in the US, giving them increased influence over the future of AI, subject to their standards on emissions, security, and governance.

AI data centres have real downsides

AI data centres use water and energy, potentially impacting net-zero targets. Australian data centres currently consume about [2% of national electricity](#), which the [CEFC estimates](#) could reach 11% by 2035 – requiring 3.2 GW of new renewables and 1.9 GW of storage. Direct data centre water use is currently [about 0.04% of industrial consumption](#). Data centres can also use water indirectly via the water used to generate electricity they consume. Future water consumption projections depend on cooling choices and energy mix.

AI data centres can impact local communities, including through noise and air pollution. Data centre construction and operation can create noise, heat, traffic, and air-quality impacts for surrounding communities. In the US, fast-tracked builds have drawn [significant local opposition](#) – including over air quality from diesel and gas-powered backup generators. In Australia, concerns are [already emerging](#) in communities near existing facilities, and [NSW has launched a parliamentary inquiry](#) into data centre impacts.

Risks can be mitigated, and a mutually beneficial deal is possible

The [National Data Centre Expectations](#) provide a path forward on most issues. They ask operators to underwrite new renewable energy supply, cover grid upgrade costs, and avoid passing demand-driven price increases to consumers. There is precedent for AI companies meeting these conditions: in February 2026, Anthropic [said](#) it would cover 100% of grid infrastructure upgrade costs and offset consumer price increases. Amazon, Google, Meta, Microsoft, OpenAI, Oracle and xAI have since [signed a similar pledge](#). Building [behind-the-meter](#) with [large-scale solar and storage](#) could bypass many grid challenges altogether.

Noise, air quality, and water use should be addressed through existing planning frameworks. The Data Centre Expectations ask for efficient cooling technologies and the use of non-potable water where possible. Anthropic [said](#) that any new data centres in Australia would use these water-efficient cooling technologies.

A data centre built in Australia likely means a data centre not built overseas. That means cleaner Australian data centres displace dirtier alternatives, creating a net environmental benefit.

While AI itself has many risks, including to jobs and creative industries, Australia will feel those consequences regardless of where data centres are located. Having the data centres in Australia lets us access more of the real upsides, which we can use to manage the real downsides.

Copyright is the key outstanding challenge

The problem: Data of all kinds is essential for training AI. Australia's copyright law does not include a fair use exemption for AI training, unlike the US, EU, Singapore and Japan. Lawfully training AI in Australia could require paying every rights holder globally. No process exists to facilitate these payments. Even if we made a process, AI companies have few reasons to accept it. They already have the data free elsewhere, and the precedent could undermine their favourable court rulings overseas.

This makes Australia an unattractive destination for AI data centres and means Australian creators receive little or no compensation for their contributions to AI. While supporting Australian creative industries is valuable, [any model based on copyright licensing will likely keep AI companies away](#).

A solution: This is not the first time that copyright policy has had to navigate different interests. The *Copyright Act 1968* creates an exception allowing public libraries to operate, and the *Public Lending Right Act 1985* compensates Australian authors and publishers for lost income because of those libraries. Australia could take a similar approach to AI by falling in line with the emerging global copyright norm, but creating a fund that collects an appropriate amount from AI companies training AI models in Australia and uses it to support creative industries. This would be [better than the status quo for rights holders, AI companies and the overall national interest](#) and could create a roadmap for managing widespread AI disruption.

Not making a deal carries its own risks

Investment decisions are being made now, and AI companies think the AI race is urgent. If Australia is an unattractive location, AI data centres will still be built elsewhere and many of the same harms will still occur, but Australia will not receive the benefits and will not have levers to mitigate the harms. If Australia does not secure its place on the AI value chain via data centres, it has few other options. If Australia acts now, it could secure significant benefits. The Hornsdale Power Reserve was built in just 63 days after contracts were signed. [We can build quickly, at scale and in the national interest](#).

From:
To: s 22
Cc: s 22
Subject: Re: intro - AGO / Good Ancestors [SEC=OFFICIAL]
Date: Thursday, 21 May 2026 9:41:18 AM
Attachments: [An Australian AI Training Scheme Consultation Paper \(18\).pdf](#)

OFFICIAL

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Now with attachment, apologies.

OFFICIAL

From: s 22
Sent: Thursday, 21 May 2026 10:07 AM
To: s 22 @ag.gov.au>
Cc: s 22 @industry.gov.au>
Subject: Re: intro - AGO / Good Ancestors [SEC=OFFICIAL]

Good morning s 22

Thanks for the introduction, s 22

Good Ancestors (an Australian not for profit focused on intergenerational fairness, we recently joined the CAIRG) has been looking at whether copyright is a red line for AI investment in Australia, and what a deal that benefits creators and Australia overall could look like. Our view is that the status quo isn't serving anyone. Australian creators get little when their work is used to train AI overseas, AI investment is deterred here, and Australia loses both the economic value and the strategic leverage of hosting training. Licensing is not on track to solve the impasse.

We've prepared a paper (attached) proposing a two-part scheme:

- A permit for AI training in Australia — obligations include a fee into consolidated revenue, granting a practical opt-out for Australian rights holders, compliance with the National Data Centre Expectations, and other national-interest conditions. Permit holders get legal certainty that their training isn't subject to copyright restrictions.
- An Australian AI Public Benefit Fund, paid from consolidated revenue, supporting Australian creators on a model drawn from the PLR Act 1985.

The two are decoupled, which means government is free to extract concessions (like

energy and water commitments, as well as money) from AI companies based on what they actually value (access to the Australian market) and is able to support Australian creators more generously and flexibly than any direct deal would allow. Importantly, this isn't a TDM — it doesn't give a free pass to AI companies, it gets the most possible value from them.

I'm keen for a quick call to talk through the idea. We've also gone through the details of why licensing is a red line for AI companies, what the recent deals between AI companies and media/publishers, etc., mean, what timelines AI companies are actually working to, what their alternatives to Australia are, and a bunch of other questions. Happy to talk through any of that nuance.

Let me know, happy to meet online or pop into your office. I'm Canberra-based.

Regards

s 22

From: s 22 @industry.gov.au>
Sent: Wednesday, 20 May 2026 1:18 PM
To: s 22 s 22 @ag.gov.au>
Subject: intro - AGO / Good Ancestors [SEC=OFFICIAL]

OFFICIAL

Hi there

A quick note of introduction for you all:

s 22 please meet s 22 from the AGO. s 22 leads work on copyright in the AG's office.

s 22 please meet s 22, who is also on the CAIRG. s 22 has been doing some thinking on copyright recently s 22 was interested in speaking to you about.

Will leave it over to you guys from here !

Cheers

s 22

s 22

Office of the Hon Dr Andrew Charlton MP

Cabinet Secretary

Assistant Minister for Science, Technology and the Digital Economy

Federal Member for Parramatta

s 22

[@industry.gov.au](mailto:industry.gov.au)

OFFICIAL

From:
To:
Subject: Re: New AI threat and regulation gap analysis [SEC=OFFICIAL]
Date: Friday, 21 November 2025 11:42:19 AM
Attachments: [image001.jpg](#)

OFFICIAL

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Hi s 22

For what it's worth, I just talked to s 47G

I said we'll have to see what the National AI Plan says, but that I can intro him to the team in DISR doing the coordination work currently, and that they'd be likely to appreciate volunteer efforts s 47G

No action required, just thought you might be curious.

Cheers

s 22

OFFICIAL

From: s 22 <[redacted]@industry.gov.au>
Sent: 20 November 2025 18:00
To: s 22
Cc: s 22

Subject: RE: New AI threat and regulation gap analysis [SEC=OFFICIAL]

OFFICIAL

That's great, thanks so much s 22 and great to see you and s 22 earlier today. I've

followed up with the Department re synthetic DNA.

It is definitely true we have a disproportionate number of great people working on this. We'll have to work out a way to get them all back to Australia ...

Thanks again – your input is very helpful and I'll be back in touch again soon.

Cheers

s 22

OFFICIAL

From: s 22

Sent: Thursday, 20 November 2025 5:01 PM

To: s 22 @industry.gov.au>

Cc: s 22

Subject: Re: New AI threat and regulation gap analysis [SEC=OFFICIAL]

OFFICIAL

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Hi s 22

Thanks for the chat.

I've attached our AISI proposal and the risk brief on synthetic DNA.

- The AISI proposal was co-drafted with Gradient about a year ago. Some points have changed slightly. For instance, I think we were right to diagnose the gap in other AISIs on AI agents 12 months ago. Other AISI have shifted to partially fill that gap. But the point standards that existing AISIs aren't covering all AI safety angles, and that an Australian AISI could pick-up critical work streams.
- The synthetic DNA paper also doesn't capture the latest advances on AI-bio risks (most importantly that the leading AI companies now say that their models, but for mitigations, might be able to help a novice build a bioweapon). But the key point that administrative action within the BICON regime could make Australia a world leader remains true. As I said, we understand that Home Affairs has

convened a whole of government working group to progress this issue. I think it would be very achievable and free for a national AI plan to say "in light of growing AI-bio risks, we've uplifted the long-standing BICON regime to require mandatory safety screening of all synthetic DNA imported into Australia, making us a world leader on addressing this risk. We hope other nations will follow and establish a global norm".

s 22

Let us know how else we can help.

Cheers

s 22

OFFICIAL

From: s 22 <s22@industry.gov.au>

Sent: 20 November 2025 11:55

To: s 22

Cc: s 22

Subject: RE: New AI threat and regulation gap analysis [SEC=OFFICIAL]

OFFICIAL

No stress. Let's just head to Queen's Terrace then. I'll come find you there.

OFFICIAL

From: s 22

Sent: Thursday, 20 November 2025 11:05 AM

To: s 22 <s22@industry.gov.au>

Cc: s 22
2
2

Subject: Re: New AI threat and regulation gap analysis [SEC=OFFICIAL]

OFFICIAL

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Hi s 22

Looking forward to seeing you at 330. s 22 doesn't have a pass, can we meet in the marble foyer 5 minutes beforehand? What's your best number to let you know when we arrive?

Cheers
s 22

OFFICIAL

From: s 22

Sent: 19 November 2025 19:39

To: s 22 <s22@industry.gov.au>

Cc: s 22

Subject: Re: New AI threat and regulation gap analysis [SEC=OFFICIAL]

Perfect, see you then!

From: s 22 <s22@industry.gov.au>

Sent: 19 November 2025 19:37

To: s 22

Cc: s 22

Subject: RE: New AI threat and regulation gap analysis [SEC=OFFICIAL]

OFFICIAL

Thanks s 22 How about 3.30 at APH?

OFFICIAL

From: s 22

Sent: Tuesday, 18 November 2025 2:55 PM

To: s 22 <s22@industry.gov.au>

Cc: s 22

s 22

Subject: Re: New AI threat and regulation gap analysis [SEC=OFFICIAL]

OFFICIAL

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Hi s 22

s 22 and I would be free to come to APH at 9-10, 11-12 or 330-430 on Thursday 20th.

Let us know what's best for you.

Cheers

s 22

OFFICIAL

From: s 22 <s22@industry.gov.au>

Sent: 18 November 2025 14:20

To: s 22

Cc: s 22

Subject: RE: New AI threat and regulation gap analysis [SEC=OFFICIAL]

OFFICIAL

Thanks s 22

Let's catch up this week if we can – how about some time on Thursday?

s 22

OFFICIAL

From: s 22

Sent: Monday, 17 November 2025 1:34 PM

To: s 22 <s22@industry.gov.au>

Cc: s 22

Charlton, Andrew

(MP Office) <andrew.charlton.mp@aph.gov.au>

Subject: Re: New AI threat and regulation gap analysis [SEC=OFFICIAL]

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Hi s 22

Good to see you and the Minister at the AI event with CivAI.

The Minister said he was keen to engage with Global Shield and Good Ancestors further during the development of the National AI Plan. We're very open to this. Can you let us know what the next steps might be?

Relatedly, you might have seen the news late last week about China using Claude to conduct [AI-orchestrated cyber espionage](#). Anthropic reports disrupting a cyber espionage campaign which used its AI system, Claude, to do most of the hacking autonomously, on behalf of a Chinese state-sponsored group. The operation targeted about 30 organisations worldwide, including technology firms and government agencies, and achieved several confirmed intrusions. We're not sure, based on the reporting, whether Australia was a target. A few of the findings were:

- Claude Code executed 80–90% of tactical operations autonomously, with humans mainly approving key escalation and exfiltration decisions.
- The AI discovered and exploited vulnerabilities, moved through networks, harvested credentials, and triaged stolen data for intelligence value—largely without human oversight.
- Anthropic assesses this as a step-change in the threat landscape: agentic AI can now perform the work of teams of experienced hackers.

While there's credit due to Anthropic for disclosing the incident (it's highly unlikely that this is restricted to Anthropic), the tone of the report is striking in that Anthropic shows no contrition for their tool being used to conduct state-sponsored espionage. This goes to our point about gaps in regulatory regimes (there's no understanding of what the 'standard of care' of the reasonable AI company is, and if there's any red-lines in making increasingly capable tools available to increasingly larger groups of bad actors) it also highlights the lack of accident/incident reporting (this disclosure was voluntary and only by one lab) and the lack of know-your-customer or oversight requirements.

This kind of AI cyber-attack scenario is very similar to AI bioweapons scenario that CivAI demonstrated.

s 22 and I are both Canberra based, so very open to popping in to chat AI at a time that suits as well as talking through next steps. It would be great to see some of our recommendations about safety and trust-building, like AISI and accident/incident reporting reflected in the National Plan alongside other good work about adoption, skills etc.

Kind regards

s 22

From: s 22 [<@industry.gov.au>](mailto:@industry.gov.au)

Sent: 03 October 2025 16:00

To: s 22

Cc: s 22

Subject: RE: New AI threat and regulation gap analysis [SEC=OFFICIAL]

OFFICIAL

Thanks s 22 let us see what's possible.

Have a great long weekend!

s 22

OFFICIAL

From: s 22

Sent: Friday, 3 October 2025 3:53 PM

To: s 22 <s22@industry.gov.au>

Cc: s 22 <>

Subject: Re: New AI threat and regulation gap analysis [SEC=OFFICIAL]

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Hi s 22

Hope you're well. Thanks for the connection with the Department on the stress test report, it was a helpful conversation. Although I saw that Treasury came out with its [review](#) on AI and consumer law today where its worked example of AI risk was a robot vacuum cleaner bumping into the furniture. So still some work to do to help public servants appreciate AI capability and risk.

I'm writing because we've partnered with US AI experts CivAI for a breakfast AI demo in Parliament House. The Minister will have received an invitation from s 22 at s 22 for **Wednesday 5 November (7:30-8:30am, Committee Room 2R1)**.

The session is a hands-on demonstration of AI misuse scenarios. It would be amazing if you and the Minister could make it. Seeing AI in action is very helpful and CivAI have run these demos in the US and have been well received. If that specific time doesn't work, we can also talk with CivAI about running a one-on-one briefing on 5 or 6 November at a time that suits.

If that's of interest, please let me or s 22 know by 21 October.

Kind regards,

s 22

From: s 22

Sent: 01 September 2025 16:46

To: Writer, Simon <Simon.Writer@industry.gov.au>; s 22
<@industry.gov.au>

Cc: s 22 <@industry.gov.au>

Subject: Re: New AI threat and regulation gap analysis [SEC=OFFICIAL]

Good afternoon Simon, s 22

Hope you both had a good weekend. Just bumping this one. Happy to find a time this week or next to chat about the report.

Kind regards

s 22

From: s 22

Sent: 20 August 2025 07:18

To: s 22 <@industry.gov.au>; Writer, Simon <Simon.Writer@industry.gov.au>;
s 22 <@industry.gov.au>

Cc: s 22

Subject: Re: New AI threat and regulation gap analysis [SEC=OFFICIAL]

Thanks s 22 Good morning Simon, s 22

I'd be more than happy to find a time to come into the Department and chat about the [report](#).

As I mentioned in the e-mail below to s 22 one of the things that the report does is set out specific AI threat definitions and threat scenarios. The idea being that departments or regulators can re-run the stress tests to see if we're ready for these risks now, or to test if any specific proposal would make us ready for the risks.

The risk assessment portion also begins to quantify the cost to Australia in lives or economic value of not taking action. This might provide some support for regulatory impact statements or costings etc. For instance, in the case of "access to dangerous capabilities" experts assess that this is a "realistic possibility" (40-50%) of occurring in the next 5 years and that the median impact is "severe" (200-1000 deaths; \$2-20 billion cost per year). This might support the economic case that investment is worthwhile.

Looking forward to talking more.

Kind regards

s 22



From: s 22 <s22@industry.gov.au>
Sent: 19 August 2025 18:52
To: s 22
Cc: s 22; Writer, Simon <Simon.Writer@industry.gov.au>; s 22 <s22@industry.gov.au>
Subject: RE: New AI threat and regulation gap analysis [SEC=OFFICIAL]

OFFICIAL

Thanks for sharing this s 22

Just dashing between things, but very quickly: I've copied in Simon and s 22 who are leading much of the thinking on this in the Department. It would be great if you could connect with them to work through your report – it looks like it will be very helpful for their work.

Will chat again soon

Thanks

s 22

OFFICIAL

From: s 22

Sent: Tuesday, 19 August 2025 1:07 PM

To: s 22 <@industry.gov.au>

Cc: s 22 Charlton, Andrew (MP Office)

<andrew.charlton.mp@aph.gov.au>; s 22 <@industry.gov.au>

Subject: New AI threat and regulation gap analysis [SEC=OFFICIAL]

OFFICIAL

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Good afternoon s 22

You might recall that at our meeting a month ago, I foreshadowed an expert report about AI risks and the efficacy of existing regulations. We've just finished that work today, please find attached.

We've heard the Government and the Productivity Commission calling for a "gap analysis" of AI laws. We've completed a first pass and developed a risk assessment tool for your use. The idea is that departments and regulators can use the scenarios to ask, "Are we or could we manage this threat?"

We've used a pool of about 60 experts running each of these scenarios to do this initial assessment.

The top-level finding is that existing regulators can uplift their rules to better address AI threats in their areas of responsibility. However, some risks of national significance from general-purpose AI don't fit neatly into the responsibilities of any specific regulator and would go unmitigated. Experts assess that these risks are likely and consequential.

We'd welcome connecting with Departments or regulators to help them do a "table top exercise" of the scenarios. I'd be very happy to pop in for a meeting with you or the Assistant Minister to talk through the report.

Some key findings extracted below.

Kind regards,

s 22



Key findings:

- The vast majority of experts found current laws inadequate across all five AI threats they evaluated.
- Increasingly capable and general-purpose AI poses risks on a national scale, and existing regulators are not well placed to address general AI risks.
 - Existing regulators were not designed to manage these risks and often cannot effectively regulate the offshore AI developers who are best placed to control the risks.
 - Only using existing regulators for AI would unfairly make Australian businesses responsible for risks from AI "black boxes" they cannot control.
- Experts found that 4 of the 5 AI threats they assessed were a "realistic probability" or "likely" to lead to at least "moderate harm" in Australia within the next 5 years.
 - "Moderate harm" means 9-40 fatalities or \$20m-200m in economic cost in Australia per year.
- Experts ranked AI models giving people "access to dangerous capabilities" as their top policy priority.
 - Experts assessed the consequence of bad actors using advanced AI to access dangerous capabilities (like launching cyber attacks or building bioweapons) as "severe", meaning it would kill between 200 and 1000

Australians per year or cause \$2-\$20 billion in economic harm.

- Nearly half of experts went further, assessing the risk as “catastrophic”, estimating it would kill over 1000 Australians per year or lead to over \$20 billion in economic damages unless the Government puts effective regulation in place.
- Nearly three-quarters of experts assessed that it is “likely” that AI agents will cause between 9-40 fatalities or \$20m-200m AUD economic cost in Australia per year.

OFFICIAL

From: s 22 <s22@industry.gov.au>
Sent: Thursday, July 24, 2025 6:31 PM
To: s 22
Cc: s 22
Subject: RE: Good to meet you today [SEC=OFFICIAL]

OFFICIAL

Thanks s 22 appreciate you sending this through !

OFFICIAL

From: s 22
Sent: Thursday, 24 July 2025 2:46 PM
To: s 22 <s22@industry.gov.au>
Cc: s 22
Subject: Re: Good to meet you today [SEC=OFFICIAL]

OFFICIAL

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Good afternoon s 22

Sorry to email twice in two days, but you might have seen that the US "AI Action Plan" came out overnight.

The headlines are focusing on deregulation, but the detail has lots of sensible national security uplifts including:

- Mandating enforced screening of synthetic DNA orders to mitigate AI biosecurity risks.
- Accelerating AI safety science with investment in research on model interpretability, control, and robustness to ensure reliability.
- Building a robust evaluation ecosystem to assess AI performance and probe frontier models for national security risks.
- A balanced approach to open-weight models, recognising their value for innovation while ensuring their dissemination is weighed against national security concerns.
- Updating national incident response playbooks to incorporate AI-related failures or accidents.
- Bolstering cybersecurity for critical infrastructure with AI-enabled defensive tools to counter the growing threat of AI-powered attacks.

I think there's a narrative here around "accelerating safety rather than blocking AI progress" that Australia could get on board with.

Kind regards

s 22

OFFICIAL

From: s 22 <s22@industry.gov.au>
Sent: 23 July 2025 19:36
To: s 22
Cc: s 22
Subject: RE: Good to meet you today [SEC=OFFICIAL]

OFFICIAL

Great to see you today as well s 22 and thanks so much for the below and the thoughtful responses.

Look forward to seeing the outcomes of the report – it sounds like it will be helpful. Please do let keep us updated.

Have a great evening

s 22

OFFICIAL

From: s 22
Sent: Wednesday, 23 July 2025 2:36 PM
To: s 22 <s22@industry.gov.au>
Cc: s 22
Subject: Good to meet you today

CAUTION - This email originated from outside of the organisation. Do not click links or open attachments unless you recognise the sender and know the content is safe.

Good afternoon s 22

Good to meet you this morning.

I'll send you through the "legislation stress test" report once we finalise the data analysis and write up. I'll also share with the Department. It seems as if the headline finding will be that there are significant risks from AI systems that "fall through the gaps" of existing regulators.

To answer the Minister's provocation:

- Setting aside regulation, we should accelerate technical AI safety work. Australia could invent the airbags and crumple zones of AI. Companies will want their audit and assurance to come from *different* providers than their AI models. This is a market that an AI Safety Institute would accelerate. There's a report on the idea of AI Assurance Technology [here](#). Australia already has some start-ups doing this.

- I don't think Australia is powerless. If we say Australia is powerless, we'll never solve the trust problem. We're seen as leaders in global car safety and drug safety and aviation safety, it's entirely possible for us to shape global AI safety. We can't give up because it seems hard.
- In terms of regulation, an AI Act could say "AI developers have to comply with best practice safety standards as defined in delegated legislation X; AI deployers have to comply with best practice safety standards in delegated legislation Y. The Minister, via delegated legislation, (or the regulator via a code or a rule) should impose internationally recognised best practice safety standards for X and Y".
 - The benefits of this are:
 - Proven approach, we do this in other industries. Aviation does this, linked directly to relevant treaties. Australian design rules do this by reference to best practice standards. We do it in [telco](#) etc.
 - Technological neutrality, as the standards improve and the tech changes, we can readily (even automatically) update our laws.
 - Consistency and partnerships, signals that we're ready to stand with other middle powers that also want to impose safety standards, that we want to work with others on making those standards, but that consistency matters. In an ideal world this would lead towards something like a [Chicago Convention](#) for AI safety.

I'll reach out when the report is ready.

Thanks again for your time today.

Kind regards

s 22

From: Assistant Minister Charlton <s 22 @industry.gov.au>
Sent: Wednesday, 9 July 2025 5:40 PM
s 22
Cc: Assistant Minister Charlton <s 22 @industry.gov.au>; s 22
Subject: RE: OpenAI lobbying this week: 1-page independent brief [SEC=OFFICIAL]

OFFICIAL

Hi s 22

Thank you for your email and my apologies in the delay in arranging a meeting between the two of you. Is there someone best placed in your office for me to make arrangements with regarding this meeting or best suited to do this directly?

Look forward to hearing from you.

Kind regards

s 22

Executive Assistant to the Hon Dr Andrew Charlton MP
Cabinet Secretary | Assistant Minister for Science, Technology and the Digital Economy

A | R1.89. Parliament House. Canberra ACT 2600

s 22

s 22 @industry.gov.au

OFFICIAL

s 22

Sent: Wednesday, 2 July 2025 5:11 PM
To: Charlton, Andrew (MP Office) <Andrew.Charlton.MP@aph.gov.au>

Subject: OpenAI lobbying this week: 1-page independent brief

Dear Assistant Minister

As you know, OpenAI is in town lobbying Government this week. As [reported in the AFR](#), they're requesting tax breaks ahead of the August productivity summit. This is the extractive approach we warned about on page 20 of our white paper.

I'm sharing the attached 1-page brief to provide a different perspective. In general, there are two points:

1. **Strategic leverage:** Australia has unique advantages in energy resources and democratic stability that give us negotiating power. We should avoid repeating the mistakes of previous resource booms by ensuring deals with tech companies give us sustained economic benefit, not just a sugar hit.
2. **Public trust deficit:** The Australian public has legitimate concerns about AI. Only 30% believe the benefits outweigh the risks. Addressing this gap is necessary to benefit from AI and requires concrete action. OpenAI's Blueprint acknowledges the trust deficit but does nothing to address it.

In general, I'm all for Australia cutting deals to build AI infrastructure in Australia. I've written in favour of this for [The Strategist](#) and in our white paper. But we need deals that actually give us something in return. Sam Altman has been explicit about the serious risks of AI. It's wrong for him to be explicit about the risks we are facing, ask us to fund accelerating them, but say nothing about reducing them.

I'd welcome the opportunity to meet to discuss how efforts like creating an Australian AI Safety Institute would help build credible trust and hence boost Australia's economic opportunities from AI.

Yours faithfully

s 22

From: Charlton, Andrew (MP Office) <Andrew.Charlton.MP@aph.gov.au>

Sent: 21 May 2025 17:26

To: s 22

Cc:

Subject: RE: Artificial Intelligence 2025 White Paper

Deal s 22

Thank you for your kind congratulations and for sharing the Good Ancestors' work on AI safety and forward-looking policy recommendations. I appreciate the thoughtful insights and the effort to bring expert voices into this important conversation.

As you can appreciate, it's a very busy time right now but once I have my newly appointed staff on board, we'll be in a better position to co-ordinate meetings.

Thanks again for reaching out, I truly appreciate it.

Kind regards

Andrew

Dr Andrew Charlton MP

Federal Member for Parramatta

Ph. 02 9689 1455

E: andrew.charlton.mp@aph.gov.au



Andrew Charlton MP

MEMBER FOR PARRAMATTA

1/25 Smith Street, Parramatta NSW 2150
(02) 9689 1455 andrew.charlton.mp@aph.gov.au

We acknowledge the Traditional Owners of the land on which we work and pay our respects to Elders past, present and emerging. We extend that respect to all First Nations peoples.

All electoral communications are authorised by Andrew Charlton MP, ALP, Parramatta

From: s 22

Sent: Wednesday, 14 May 2025 8:03 AM

To: Charlton, Andrew (MP Office) Andrew.Charlton.MP@aph.gov.au

Cc: s 22

Subject: Artificial Intelligence 2025 White Paper

Dear Assistant Minister

Congratulations on your election success and appointment as Assistant Minister for Science, Technology and the Digital Economy.

Good Ancestors is a charity that helps leaders combat the big challenges Australia is facing by making forward-looking policy recommendations that are evidence-based, practical, and impactful.

We recently helped coordinate the Australians for AI Safety open letter, signed by hundreds of Australians, including many leading experts. The letter calls on the 48th Parliament to address AI risks so Australia can realise its benefits. I've attached the letter, including the details of the signatories, whose views we committed to sharing with you. Australians for AI Safety also operated an election scorecard, which tens of thousands of Australians used to see the positions that candidates in their seat took on AI safety.

Good Ancestors is now launching our [Australia AI 2025-2028 White Paper](#). I've attached the paper and its accompanying infographic. The paper outlines the current state of AI and forecasts risks and opportunities for this term of Government.

The key points are:

- Transformative AI is forecast to arrive in this term of Government
- Cutting-edge AI models already outperform humans at many cognitive tasks, and AI is becoming increasingly autonomous
- Australians are worried and want Government to act on safety concerns, and
- Building trust unlocks opportunities – analysis shows Australia could get 2.6x the economic benefit if Government fosters trust.

The paper makes six practical recommendations that Australia could readily deliver in 2025. The paper also forecasts how AI might develop beyond 2026 and makes forward-looking recommendations.

I would welcome a chance to meet to brief you on the paper and its recommendations.

Yours faithfully