Capturing AI Training and Inference Under the SOCI Act

Submission to the Proposed Amendments to Streamline and Modernise the Security of Critical Infrastructure Act 2018

Good Ancestors is an Australian charity providing evidence-based policy recommendations for Australia's biggest challenges. We work with experts around the world and help organise Australians for AI Safety.

In our December 2025 submission, we argued that the SOCI Act did not capture Australia's growing dependency on AI, or the risks associated with AI infrastructure.[1] AI capabilities have increased significantly since then.[2] We welcome the new facility-based and service-based capture mechanisms proposed under Measure 8, and the clarification under Measure 5 that automated, AI-enabled, and materially uncontrolled systems fall within the cyber security incident definition.

This submission recommends three further changes to ensure the SOCI Act captures AI inference services, entities conducting AI training, and AI service providers that critical infrastructure depends on.

AI inference should not be excluded from service-based capture

Australian hospitals, banks, law firms, government agencies, and members of the public increasingly use AI models through inference services.[3] A customer sends data to a provider, the provider runs that data through a trained AI model held on its own infrastructure, and the provider returns an output.

This is a "data storage or processing service", as defined in the SOCI Act. The provider holds the model, performs the computation on its own infrastructure, and receives the customer's data to do so. The customer does not hold a copy of the model—if the provider fails or is compromised, the service stops.

Question 42 nonetheless raises the possibility that AI inference services could be excluded. It asks which service types should be excluded from service-based capture, and gives "pure software licensing" as a starting point.[4] Inference is not mentioned in the consultation paper, and any exclusion drafted for software licensing should not capture it.

Inference is not pure software licensing. Under a software licence, the customer receives a copy of the model and runs it on their own systems. With an inference service, the provider runs the model and returns a result. Excluding inference would leave services that Australia depends on outside risk management and incident reporting obligations.

We recommend that **AI inference services are not excluded from the service-based pathway.**

---

[1] Grundy, E. & Sadler, G. (2025). *Extending SOCI Act Coverage to AI Infrastructure*. Good Ancestors.
[2] For example: Anthropic. (2026). *Project Glasswing: An initial update*; UK AI Security Institute. (2026). *Our evaluation of OpenAI's GPT-5.5 cyber capabilities*.
[3] For example: KPMG & University of Melbourne. (2025). *Trust, Attitudes and Use of AI: Global Report*.; Department of Industry, Science and Resources. (2025). *AI Adoption Tracker*. Australian Government; Department of Industry, Science and Resources. (2025). *National AI Plan* (p. 11). Australian Government.
[4] Department of Home Affairs. (2026). *Consultation on Proposed Amendments to Streamline and Modernise the Security of Critical Infrastructure Act 2018* (p. 41). Australian Government,.

AI training entities should be captured under Measure 8

Training an AI model is the compute- and data-intensive process that produces a model's weights, before that model is made available to anyone. The resulting general-purpose models are integrated into systems used across the Australian economy, including by critical infrastructure.

Anthropic is reportedly seeking 1.4 GW of Australian data centre capacity, including to train its Claude models.[5] This would make Australia its second training location outside the United States, and mean that capturing AI training entities under the SOCI Act is even more important.

No proposed pathway appears to capture the *entity* conducting an AI training run. Where training takes place in a leased facility, facility-based capture reaches the operator but not the entity doing the training. Service-based capture also appears not to reach the entity: a training run is not clearly a data processing service provided on a 'commercial basis'.[6] The commercial activity comes later, at inference.

Specific hazards arise during model training, such as theft of the model weights or 'poisoning' of the training data.[7] A data centre operator holds information about the physical characteristics of the facility, but the entity conducting the training holds information relevant to model and training risks. They are the only party able to manage those risks.

We recommend that **entities conducting AI training are captured**, in addition to the facility operator.

AI services should be captured where critical infrastructure depends on them

Question 40 asks which metrics would identify significant service-layer providers. The metrics described in Measure 8 are based on the size and scale of the provider or their service. Some services, however, may process relatively small amounts of data, but still be significant because other critical infrastructure entities depend on that processing.

AI services are being adopted into critical functions across the economy, and model capability is uneven across domains.[8] An AI model can perform at expert level in a narrow area while remaining unreliable in others. Critical functions in health, energy, water, or finance may come to depend on narrow AI services that meet none of the thresholds in Measure 8.

Section 12F currently captures providers that serve critical infrastructure entities. We support replacing that as the ordinary basis for capture, but no proposed pathway appears to retain it in any form. The reserve Ministerial designation pathway reaches providers outside the ordinary thresholds, but it is confined to those Government depends on. A provider that a hospital or an energy operator depends on—but not Government—appears to fall outside all four pathways.

We recommend that **at least one capture pathway reach AI service providers that critical infrastructure depends on**.

---

[5] Macdonald, A. & Wiggins, J. (2026). *Anthropic wants Australian data centres, copyright deal ahead of IPO*. Australian Financial Review.

[6] The reserve Ministerial designation and certification-based capture also appear not to apply to an *entity* conducting a training run.

[7] Nevo, S., et al. (2024). *Securing AI Model Weights: Preventing Theft and Misuse of Frontier Models*. RAND Corporation; Hubinger, E., et al. (2024). *Sleeper agents: Training deceptive LLMs that persist through safety training*. arXiv preprint arXiv:2401.05566.

[8] Bengio, Y., et al. (2026). *International AI Safety Report 2026* (pp. 26, 29).

Conclusion

Measure 8 responds to a gap we identified in our December 2025 submission. The SOCI Act must keep pace with AI developments by ensuring:

- Exclusions from service-based capture do not extend to AI inference services;
- The capture pathways under Measure 8 reach entities conducting AI training; and
- At least one capture pathway reaches AI service providers that critical infrastructure depends on.

Australia is attracting significant data centre investment, and developers have signalled their interest in training general-purpose AI models here. The SOCI Act should ensure those facilities, and the entities operating models within them, are subject to appropriate security and reporting obligations.

Submitted

31 July 2026

Authors

Daniel Marns, Emily Grundy

About Good Ancestors

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

Contact

If you would like to discuss this submission, please let us know at contact@goodancestors.org.au.