Extending SOCI Act Coverage to AI Infrastructure

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.

The National AI Plan highlighted that artificial intelligence (AI) is a "critical technolog[y] in the national interest" that is "already shaping our economy and society".[1] Already, half of Australians, and 40% of Australian small and medium businesses, are adopting AI.[2,3] The Tech Council of Australia estimates AI could contribute $45 billion and $115 billion annually by 2030—equivalent to 2-5 per cent of the Australian economy.[4] The Government is also positioning Australia as a "leading destination for data centre investment", with companies announcing plans to invest upwards of $100 billion in Australian data centres between 2023-2025.[5]

Industry and Government, as well as investment trends and capability evaluations, all point to AI becoming an essential service which underpins Australian society and economy. This will also expose Australia to new threats and hazards.

The SOCI Act protects essential services and systems from failure or disruption.[6] Functionally, the Act is Australia's main tool for managing AI-related critical infrastructure risk. Therefore, it should:

- address risks associated with facilities that train and operate general-purpose models that are used across the economy, and
- prepare for when AI models and systems are essential to Australian society and thus become critical infrastructure themselves.

This submission focuses on the fourth question raised by the Independent Reviewer: "*Are there new or emergent threats the SOCI Act is unable to manage in its current form?*". We outline gaps in the SOCI Act that render it unable to address risks from, and dependencies on, rapidly advancing AI.

AI is becoming critical infrastructure

AI is evolving from an optional tool to foundational infrastructure.

Currently, businesses are under pressure to integrate AI to remain competitive, and Australians increasingly rely on AI-enabled services in daily life.

In the near future, AI may be embedded in most critical infrastructure operations—including healthcare, banking, logistics, and government services—and augment consequential decision-making. Major AI system failures or outages could cause significant disruption, with recovery taking days and depending on cooperation and coordination with AI companies and data centres.

[1] Department of Industry, Science and Resources. (2025, December 2). *National AI Plan*. Australian Government.
[2] KPMG and University of Melbourne. (2025, May). *Trust, Attitudes and Use of AI: Global Report*. KPMG.
[3] Department of Industry, Science and Resources. (2024, June 4). *AI adoption in Australian businesses for 2024 Q4*. Australian Government.
[4] Microsoft and Tech Council of Australia. (2023, July). *Australia's Generative AI Opportunity.* Microsoft and Tech Council of Australia.
[5] Department of Industry, Science and Resources. (2025, December 2). *National AI Plan*. Australian Government.
[6] Department of Home Affairs. (n.d.). *Security of Critical Infrastructure Act 2018*. Australian Government.

In the medium-term future, AI could be as essential as telecommunications or electricity. Major AI system failures or outages may severely hamper Australia's economic or social stability. Recovery could take weeks to months, with no ready alternatives or substitutes.

Increasingly advanced and integrated AI creates unique threats and hazards, including to national security.[7] Large-scale AI infrastructure creates vulnerabilities as both a target and vector (i.e., it could be both used to cause, and the recipient of, harm). For example, data centres training or operating general-purpose AI models could be compromised through cyberattacks or physical incidents. They could also cause harm through poisoned training data or manipulated outputs that could propagate to downstream users, including critical infrastructure. As AI becomes increasingly advanced and integrated into Australian society, the potential consequences of inadequate management and protection grow. See the Scenario box below for examples.

ℹ️ Scenario 1: Sabotage disrupts data centre operations

A Tasmanian data centre operates GPT6, a frontier AI model. The model is integrated into systems that are then used across critical infrastructure, such as healthcare and banking. However, because the AI data centre is serving a general user base, and not "wholly or primarily" critical infrastructure entities, it falls outside SOCI Act coverage.

As a result, the data centre is not required to implement a Critical Infrastructure Risk Management Program, and conducts inadequate assessment of physical risks that could disrupt the facility.

During the peak of summer, a coordinated sabotage disables the facility's cooling and power systems. Servers overheat and within hours, the model becomes unavailable. Millions of Australians who rely on this tool—including for work—experience significant disruption. Healthcare diagnostic tools and bank fraud detection systems that operate downstream of the foundational model fail.

While Government can direct affected critical infrastructure entities on managing the outage, it cannot issue action directions to the data centre itself under s35AQ.

ℹ️ Scenario 2: Malicious customer rents compute

A foreign actor rents GPU compute at a Sydney data centre and uses it to train and operate an AI model designed for sophisticated cyber attacks. The data centre, which operates general-purpose AI infrastructure that serves multiple customers, falls outside SOCI Act coverage because it doesn't primarily serve specific critical infrastructure entities. The facility is therefore not required, under the SOCI Act pt 2, to register "interest and control information" about customers renting compute.

The trained model is deployed and begins conducting automated attacks against Australian financial institutions. The Australian Signals Directorate attributes the attacks' origin to the Sydney data centre running inference compute. However, the data centre operator has limited visibility into which customer is responsible, or even exactly who its customers are, leading to delays in identification. Section 35AQ of the SOCI Act cannot be invoked to direct the data centre to isolate or terminate the malicious workload because the facility isn't classified as critical infrastructure.

[7] Grundy, E., Sadler, G., & Freeman, L. (2025, October 28). *Artificial Intelligence and National Security*. Good Ancestors.

Gaps in the SOCI Act

Currently, the SOCI Act does not capture Australia's growing dependency on AI. There are two key gaps in the framework.

Gap 1: Limited coverage of AI data centres under SOCI Act

On its face, the SOCI Act captures data centres only when (a) used, wholly or primarily, to serve Government or other critical infrastructure, and (b) relate to business-critical data.

The "wholly or primarily" stipulation typically excludes general-purpose AI models

Section 12F excludes assets from the definition of "critical data storage or processing assets" unless they primarily provide services to:

1. Commonwealth/State/Territory entities, or
2. Entities responsible for other critical infrastructure (e.g., bank, hospital, electricity grid)

Data centres that **train**[8] general-purpose AI models that are used across the economy don't provide services *directly* to specific critical infrastructure entities. Instead, they train general-purpose models that are subsequently built into AI systems, which are then used directly and indirectly by critical infrastructure. Data centres that **operate**[9] general-purpose AI systems may also not meet the "*wholly or primarily*" definition because they are serving a diverse pool of users.

The SOCI Act currently treats the criticality of data centres based on who they serve. If a data centre is explicitly serving critical infrastructure, the law treats it as critical infrastructure. This approach made sense through the lens of cloud storage, where data centres, in a relatively structured way, were used to outsource a critical component of business that historically had been done internally. This distinction is undermined as data centres increasingly train or operate AI models. These models are a general product that provide a general service across the economy.

Viewed through the lens of AI, and its widespread adoption in society and growing dependency, these data centres become critical infrastructure in their own right—rather than only through a direct relationship to other infrastructure. They create or operate capabilities that are used across the economy, rather than only via other critical infrastructure or systems of national significance.

The below case study outlines a current example of how the Act may exclude data centres training or operating general-purpose AI models.

[8] "Training" refers to the compute- and data-intensive optimisation process that produces a set of model weights by repeatedly updating them to reduce prediction error on an objective.
[9] "Operating" refers to the comparatively less compute- and data-intensive process of running a trained model on new inputs in response to a user query or API request, using inference compute to produce an output.

ℹ️ Case study: "Nation-building digital infrastructure" not captured by the SOCI Act

OpenAI, partnering with NEXTDC, has announced plans to develop “a next-generation hyperscale AI campus and large-scale GPU supercluster” in Sydney.[10] The facility will train models that become embedded across the economy.

The 550MW facility will “support sensitive and mission-critical workloads”.[11] NEXTDC CEO Craig Scroggie calls it "nation-building digital infrastructure" that will "provide sovereign compute capability for government, finance, defence, research and enterprise."

Yet, this $7 billion data centre would likely not be captured by the SOCI Act. When training an AI model like GPT6, NEXTDC and OpenAI likely aren’t "providing a data storage or processing service" wholly or primarily to Government or critical infrastructure entities under s12F. Instead, it's building general-purpose capabilities that those entities, among many others, would likely then use or deploy. Even if government and other critical infrastructure assets subsequently use OpenAI models trained in the facility, it’s unlikely that the facility would meet the “wholly or primarily” threshold set out in s12F.

As a result, such a data centre would not be subject to requirements under the SOCI Act, including risk management programs, cyber incident reporting, and ownership transparency. OpenAI and NEXTDC may voluntarily offer to comply with the SOCI Act. While this is helpful, it illustrates that stakeholders agree that these are the kinds of assets that the SOCI Act should cover and that they’re incentivised to comply as a trust-building exercise.

“Business-critical data” definitions assume cloud storage models

Section 12F also stipulates that the critical data storage or processing assets must relate to business-critical data (BCD). The definition of BCD assumes the cloud-storage model of a data centre, i.e. a data centre that is housing, storing, or hosting information, and that the importance of the data centre depends on the importance of the information, as it relates to personal information or critical infrastructure assets.

This definition may not work as intended for data centres that train or operate AI models. These data centres may only contain model weights, which are complex algorithms, and may not contain any information or data in the sense of the BCD definition, despite forming a critical input to the asset or system.

Gap 2: General-purpose AI models will need to be captured under SOCI

Capturing data centres that train and operate general-purpose AI models under SOCI will help address some of the risks associated with advanced AI. However, the models themselves will also need to be covered as they become critical infrastructure in their own right.

Australians, businesses, and critical infrastructure operators are increasingly reliant on AI models being capable, safe, and secure. This extends beyond ensuring the safety and security of the physical facilities hosting these models.

[10] OpenAI. (2025, December 4.). *Introducing OpenAI for Australia*. OpenAI.
[11] NEXTDC. (2025, December 5). *NEXTDC to Join OpenAI in Australia as an Infrastructure Partner*. NEXTDC.

AI model developers and providers have a critical responsibility in identifying, evaluating, preventing, and mitigating AI-specific vulnerabilities and hazards. These risks are wide-ranging, from discrimination and toxicity to malicious actors and misuse, socioeconomic disruptions, environmental harms, and catastrophic risks, including loss of control (see MIT's AI Risk Repository for a synthesised taxonomy of over 1,600 AI risks). An advanced technology that is unpredictable, difficult to understand, and widely deployed could cause significant damage. Appropriate safeguards and risk management programs must ensure these systems are trustworthy, reliable, controllable, and beneficial.

Bringing AI models directly into the regime would apply the sensible risk mitigation practices Australia applies to established forms of critical infrastructure to this emerging form of critical infrastructure.

Implications of SOCI Act gaps

Exclusion from the SOCI Act creates gaps in risk prevention and management. If data centres training or operating general-purpose AI models, for example, are not classified as critical infrastructure, the Act:

- **Does not require ownership transparency**, and cannot track foreign investment in facilities training or operating models used across critical sectors. This threat vector is particularly critical for AI models and systems because of the "AI Sleeper Agents" problem, whereby models behave normally during testing but have hidden, malicious behaviours that can be triggered later.[12] Foreign ownership and investment could be a key vector through which sleeper agents are executed.
- **Does not require incident reporting,** and operators may not be obligated to report cyber breaches, model theft, or training data poisoning that could compromise downstream systems.
- **Does not require risk management programs** (Part 2A), whereby the data centre would develop and follow processes to identify, minimise, and mitigate hazards.
- **Does not enable crisis coordination.** If a model operating out of an Australian data centre is causing widespread harm (e.g., being used to conduct automated espionage at scale,[13] or other harms), Government cannot issue directions under s35AQ. The SOCI Act's crisis coordination mechanisms, including Systems of National Significance designation under s52B, only apply if data centres training or operating general-purpose AI models are first brought under coverage. Australia's AI Plan acknowledges the need for a national-level AI crisis plan, recognising that the nature of AI makes AI accidents and incidents increasingly divergent from other classes of crises. Including relevant assets in SOCI will be an essential precondition to success in AGCMF reform.

Recommendations

Recommendation 1. Bring data centres training and operating general-purpose AI models within SOCI coverage

The SOCI Act should treat infrastructure associated with AI training and operations as critical infrastructure, regardless of the customers served or whether they relate to 'business critical data'. This should apply to general-purpose AI intended for widespread adoption across the economy, where disruption or compromise would have significant social and economic consequences.

[12] Hubinger, E., Denison, C., Mu, J., Lambert, M., Tong, M., MacDiarmid, M., ... & Perez, E. (2024). *Sleeper agents: Training deceptive LLMs that persist through safety training*. arXiv preprint arXiv:2401.05566.
[13] Anthropic. (2025, November 14). *Disrupting the first reported AI-orchestrated cyber espionage campaign*. Anthropic.

Data centres serving specific critical infrastructure entities are already captured under s12F. However, facilities creating or operating AI capabilities that become foundational across society, including to critical infrastructure, currently fall outside the Act's scope because they serve general user bases rather than specific entities.

One way this coverage could be achieved is through expanding s12F. A new subsection could identify data centres used to train or operate general-purpose AI models intended for widespread use as "critical data storage or processing assets". Coverage could be determined based on model size, compute used, or the potential consequences of disruption or malfunction.

Facilities that are particularly consequential to social or economic stability, defence, or national security could be designated as Systems of National Significance under s52B. This would appropriately trigger enhanced cyber security obligations.

The below 'International precedents' box outlines how different countries and regions are treating data centres as critical infrastructure.

ℹ️ **Data centres as critical infrastructure: international precedents**

Major jurisdictions are recognising data centres as critical infrastructure, independent of who they serve. The UK, Europe, Germany, and Singapore all provide models for Australia. Two examples are outlined below.

United Kingdom: In September 2024, the UK designated data centres as Critical National Infrastructure.[14] This means data centres now have risk management and incident reporting requirements, as well as priority access to security agencies and emergency services during crises.

UK data centres are captured if they are above defined capacity thresholds. Unlike Australia's SOCI Act, inclusion criteria does not centre around serving specific customers.

Singapore: In 2024, Singapore amended its Cybersecurity Act to create a "Foundational Digital Infrastructure" category.[15] This explicitly recognises that digital infrastructure services (e.g., cloud services, data centres) that are foundational to society or the economy carry systemic risk, distinct from traditional critical infrastructure.

Foundational digital infrastructure service providers face regulatory obligations when "disruption or deterioration of the operation of a large number of businesses or organisations in Singapore which rely on or are enabled by that foundational digital infrastructure service" (ss18G(b)).[16] This definition relies on the consequences of disruption to the service—not the specific Government or critical infrastructure users they are serving.

Recommendation 2: Create a pathway to cover general-purpose AI models as critical infrastructure.

This should ensure risk management obligations fall upon the appropriate entity, like developers and operators. While data centres are important, they do not have visibility of, or the ability to mitigate, all relevant AI risks. Integrating AI models into the SOCI Act should involve consultation and co-design with industry and AI safety experts (e.g., on definitions, thresholds, and which SOCI Act obligations are appropriate for model developers and providers).

[14] UK Department for Science, Innovation and Technology. (2025, November 12). *Policy paper: Data centres*. UK Government.
[15] Cyber Security Agency of Singapore. (2025, April 2). *Cybersecurity Act*. Government of Singapore.
[16] Parliament of Singapore. (2024, April 3). *Cybersecurity (Amendment) Bill*. Bill No. 15/2024.

Conclusion

Highly capable AI systems may soon bring profound economic and social transformations. Yet, general-purpose AI is an emergent threat that the SOCI Act structurally misses. AI infrastructure, and the underlying models, should be captured to ensure they are subject to necessary obligations like risk management and cyber incident reporting.

Submitted

22 December 2025

Authors

[redacted]

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 [redacted]