

# Artificial Intelligence: Managing Risks and Capturing Value

2026–27 Pre-Budget Submission

January 2026

## Submitted

29 January 2026

## Authors

Emily Grundy, Luke Freeman, Greg Sadler

## 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](mailto:contact@goodancestors.org.au).



Good  
Ancestors

## Table of Contents

<b>AI presents significant economic opportunity for Australia.....</b>	<b>3</b>
<b>Inadequate risk management stifles opportunity.....</b>	<b>4</b>
Public distrust restricts adoption.....	4
Businesses lack adequate risk management tools.....	4
<b>The risks threaten more than adoption.....</b>	<b>5</b>
<b>Government has chosen AISI as its primary response.....</b>	<b>5</b>
The AISI is not adequately resourced.....	6
<b>Capturing value through AI infrastructure.....</b>	<b>7</b>
<b>Capturing value through AI Assurance Technology.....</b>	<b>8</b>
<b>Recommendations.....</b>	<b>9</b>
1. Increase AISI funding to match its extensive mandate.....	9
2. Ensure data centre investments serve national interests.....	9
3. Invest in Australia's AI Assurance Technology industry.....	9

In Good Ancestors' 2025-26 pre-budget submission, we argued Australia needed an AI Safety Institute (AISI) to deliver on existing commitments made by the Government and prepare for economic transformation over the next decade and beyond.<sup>1</sup> We welcomed Government's November announcement of the AISI as an excellent step in addressing harms from AI.<sup>2</sup> However, the AISI's \$30 million funding over four years is insufficient for its extensive mandate: serving as the whole-of-government coordination hub, advising ministers and regulators, supporting compliance across sectors, and addressing both upstream risks and downstream harms.

This submission calls for **increased AISI funding to match its scope**. It describes the economic opportunity of AI adoption, explains why trust determines whether Australia unlocks that opportunity, and examines the resourcing necessary for the AISI to succeed as Government's key trust-building initiative. The submission also outlines opportunities for AI data centres and the AI Assurance Technology (AIAT) industry to secure Australia's place in the AI value chain.

This submission makes three recommendations:

1. **Increase AISI funding to \$50 million per year** to match its extensive mandate and be proportionate to the UK AISI.
2. **Require AI companies striking large data centre deals to commit to transparent information sharing with the AISI and make compute resources available for AI safety research**, ensuring infrastructure investments serve national interests rather than extracting wealth from Australian soil.
3. **Establish an AIAT Grant Programme with \$90 million in cornerstone funding**, jointly administered by the National AI Centre and the AISI, to build Australia's AI Assurance Technology industry and capture value in a market forecast to reach USD 276 billion by 2030.

## AI presents significant economic opportunity for Australia

Industry and Government, as well as investment trends and capability evaluations, all point to AI becoming an essential service that underpins Australian society and economy. Government recognises AI as a "critical technolog[y] in the national interest",<sup>3</sup> with the National AI Plan emphasising productivity growth across the economy. By 2030, AI could contribute between \$45 billion and \$115 billion annually to Australia's economy—equivalent to 2-5 per cent of GDP.<sup>4</sup>

This economic potential is attracting major investment. Between 2023 and 2025, companies announced plans to invest upwards of \$100 billion in Australian data centres, with Government positioning Australia as a "leading destination for data centre investment".<sup>5</sup> Over a third of Australian SMEs have adopted AI,<sup>6</sup> and Australia ranks fourth globally in the use of the popular AI tool, Claude.<sup>7</sup> 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.

---

<sup>1</sup> Good Ancestors. (2025). [AI Safety as Economic Opportunity 2025-26 Pre-Budget Submission](#).

<sup>2</sup> Department of Industry, Science and Resources. (2025). [Australia to establish new institute to strengthen AI safety](#). Australian Government.

<sup>3</sup> Department of Industry, Science and Resources. (2025). [National AI Plan](#). Australian Government.

<sup>4</sup> Microsoft and Tech Council of Australia. (2023). [Australia's Generative AI Opportunity](#). Microsoft and Tech Council of Australia.

<sup>5</sup> Department of Industry, Science and Resources. (2025). [National AI Plan](#). Australian Government.

<sup>6</sup> Department of Industry, Science and Resources. (2025). [AI adoption in Australian businesses for 2025 Q1](#).

<sup>7</sup> Anthropic. (Updated 2026, January). [Anthropic Economic Index Understanding AI's effects on the economy: Australia](#).

Beyond adoption, Australia has opportunities in the AI Assurance Technology (AIAT) industry. The global AIAT market is predicted to reach USD 276 billion by 2030,<sup>8</sup> and Australia is well-placed to secure a significant share. A properly resourced AISI could help attract talent and catalyse this sector.

## Inadequate risk management stifles opportunity

While we welcome the Government's acknowledgement of AI's economic opportunities, these opportunities depend on trust and risk management.

### Public distrust restricts adoption

Australians are among the least trusting of AI globally. 96% hold concerns about generative AI, and only 36% trust AI systems.<sup>9</sup> In data collected before the National AI Plan, only 30% of Australians believed Government measures were adequate for making AI safe, while 77% supported AI regulation.<sup>10</sup> Distrust impacts adoption rates, and adoption rates shape whether the economic value Australia captures by 2030 is more like \$45 billion or \$115 billion.<sup>11</sup> As the National AI Plan states, "We cannot seize the innovation and economic opportunities of AI if people do not trust it" (p. 28).<sup>12</sup>

### Businesses lack adequate risk management tools

Australian businesses want to adopt AI but lack adequate tools, frameworks, and knowledge to manage the risks. AI systems are already causing harm: Noosa Council lost \$1.9 million in 2025 when AI was likely used to convince staff to approve fraudulent transactions.<sup>13</sup> Commonwealth Bank cut 45 customer service roles for an AI voice bot, only to reverse the decision when call volumes spiked rather than declined.<sup>14</sup> A Bunnings AI chatbot provided detailed instructions on replacing an extension cord plug in Queensland—work that is illegal in that state if you're not a licensed electrician.<sup>15</sup>

Without adequate upstream regulation, assurance tools, and risk identification and mitigation research, Australian businesses are in an impossible situation. Government and market forces are pressuring them to adopt AI, but they're not equipped to anticipate and mitigate risks from "black box" AI systems often beyond their technical expertise or control.<sup>16</sup>

The AIAT industry could provide businesses with the tools to identify and manage these risks, but Australia's nascent AIAT sector needs support to scale. A properly resourced AISI could accelerate this by conducting technical evaluations and working with regulators, but it's underfunded. Current voluntary evaluation approaches adopted by frontier AI companies have also proved inadequate.<sup>17</sup> The result is that businesses are hesitant to deploy AI despite its potential benefits.

---

<sup>8</sup> AIAT Report. (2024). [Risk & Reward: AI Assurance Technology Market Report 2024](#).

<sup>9</sup> KPMG and University of Melbourne. (2025). [Trust, Attitudes and Use of AI: Global Report](#).

<sup>10</sup> Ibid.

<sup>11</sup> Microsoft and Tech Council of Australia. (2023). [Australia's Generative AI Opportunity](#). Microsoft and Tech Council of Australia.

<sup>12</sup> Department of Industry, Science and Resources. (2025). [National AI Plan](#). Australian Government.

<sup>13</sup> Grace, M., Howells, S., & Wong-See, T. (2025). [Noosa mayor says fraudsters used AI imitation in \\$2.3m council scam](#). ABC News.

<sup>14</sup> Chalmers, S. (2025). [Commonwealth Bank backtracks on AI job cuts, apologises for 'error' as call volumes rise](#). ABC News.

<sup>15</sup> Jolly, N. (2025). [A terrible few days for AI bots, Bunnings, and SCA](#). Mumbrella.

<sup>16</sup> Sadler, G., Grundy, E., Freeman, L., & Sherburn, N. (2025). [AI Legislation Stress Test](#). Good Ancestors.

<sup>17</sup> Future of Life Institute. (2025). [AI Safety Index – Summer 2025](#).

## The risks threaten more than adoption

Beyond slowing adoption, inadequate AI risk management brings social, economic, and national security implications. 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 1,700 AI risks).<sup>18</sup> An advanced technology that is unpredictable, difficult to understand, and widely deployed could cause immense damage.

In 2025, both OpenAI and Google warned that their leading models crossed chemical, biological, radiological, and nuclear (CBRN) risk thresholds. Google assessed that Gemini 2.5 Deep Think reached the "early warning threshold" for its CBRN risk standard—models that "can be used to significantly assist a low-resourced actor with dual-use scientific protocols, resulting in a substantial increase in ability to cause a mass casualty event."<sup>19</sup> OpenAI issued similar warnings for ChatGPT Agent<sup>20</sup> and GPT-5.<sup>21</sup>

Cyber capabilities are also increasing. Google demonstrated that an AI agent discovered a previously unknown vulnerability in widely used software.<sup>22</sup> One AI agent ranked in the top 5% of over 400 teams in cybersecurity competitions,<sup>23</sup> while actors with only basic coding skills are misusing widely available tools, like Claude Code, for large-scale extortion and AI-generated ransomware.<sup>24</sup> Government-backed threat actors are also integrating general-purpose AI into cyber operations.<sup>25</sup>

AI systems are approaching and surpassing human performance across problem-solving, scientific reasoning, and persuasion.<sup>26</sup> Leading labs and forecasters predict Artificial General Intelligence—AI with human-like cognitive capabilities—could arrive as early as this year,<sup>27</sup> with the median forecast of 1,700 forecasters pointing to November 2027.<sup>28</sup> Without adequate resources to evaluate and mitigate these emerging threats, Australia will remain vulnerable. Emphasising AI's economic benefits without assessing potential harms risks distorting Government's investment priorities. Australia needs cost-benefit analyses, not benefit-only analyses.

## Government has chosen AISI as its primary response

Australia's AISI will be established in early 2026. In our 2025 pre-budget submission, we recommended an AISI focus on evaluating advanced AI systems, driving foundational AI safety research, and coordinating internationally. Government intends for the AISI to cover these functions, but has also assigned it a broader remit. The AISI will serve as the whole-of-government coordination hub, advising ministers and regulators, supporting compliance across sectors, and addressing both upstream risks and downstream harms.<sup>29</sup>

<sup>18</sup> MIT FutureTech. (2025). [MIT AI Risk Repository](#). Massachusetts Institute of Technology.

<sup>19</sup> Google DeepMind. (2025). [Gemini 2.5 Deep Think Model Card](#).

<sup>20</sup> OpenAI. (2025). [ChatGPT Agent System Cardsafety eggs](#).

<sup>21</sup> OpenAI. (2025). [GPT-5 System Card](#).

<sup>22</sup> Big Sleep Team. (2024). [From Naptime to Big Sleep: Using large language models to catch vulnerabilities in real-world code](#). Google Project Zero.

<sup>23</sup> Bertollo. (2025). [AI vs Human: CTF results show AI agents can rival top hackers](#). Hack The Box.

<sup>24</sup> Anthropic. (2025). [Threat intelligence report](#).

<sup>25</sup> Google Threat Intelligence Group. (2025). [Adversarial Misuse of Generative AI](#).

<sup>26</sup> Bengio, Y., et al. (2025). [International AI Safety Report](#).

<sup>27</sup> Anthropic. (2025). [Anthropic's recommendations: OSTP U.S. AI action plan](#).

<sup>28</sup> Metaculus. (Accessed 28 Jan, 2026). [When Will Weakly General AI Arrive?](#)

<sup>29</sup> Department of Industry, Science and Resources. (2025). [National AI Plan](#). Australian Government.

This broader remit comes without the regulatory support we recommended. Government rejected mandatory guardrails for AI systems, choosing instead to rely on existing laws and "respond[ing] to challenges as they arise." Without mandatory obligations on developers to test for dangerous capabilities or meet safety standards, the coordination burden on the AISI increases. The AISI must advise dozens of portfolio agencies and regulators across every sector—each grappling with AI risks—rather than addressing many risks centrally through AI-specific regulation. If the AISI is the safety basket Government is putting all its eggs in, it must be adequately resourced.

## The AISI is not adequately resourced

Government allocated \$30 million over four years to the AISI, averaging \$7.5 million per year.<sup>30</sup> The AISI's funding also reduces over time, contrary to Government's view that the impacts of AI will increase over time. \$7.5 million per year falls below expert recommendations, and is insufficient to service the AISI's extensive remit.

Good Ancestors surveyed 139 professionals with expertise in AI safety, governance, and related fields, to share their views about the establishment of Australia's AI Safety Institute. Participants indicated substantially higher funding was needed for the AISI to "make a meaningful contribution to AI safety"—53.3% recommended over \$50 million per year, and 77% recommended at least \$25 million annually.<sup>31</sup>

The funding is inadequate relative to the economic stakes. For every \$1,000 of predicted AI economic benefit to Australia by 2030, Government is investing only 7 cents in the AISI.<sup>32</sup> If trust determines whether Australia captures \$45 billion or \$115 billion annually from AI, a \$7.5 million annual investment in a key institution responsible for building that trust seems disproportionate to the opportunity. By comparison, Government earmarked \$166 million for GovAI Chat, its internal AI chat tool for the Australian Public Service—more than five times the AISI budget for an internal productivity tool.<sup>33</sup>

The AISI is directed to support Government across the full range of portfolio issues, engaging with agencies and regulators. Taking into account that there are 105 regulators listed on the stocktake, this would work out to approximately \$71,400 per regulator annually to build their AI risk expertise and coordinate responses. That's also assuming the entire AISI budget is spent assisting regulators.

An AISI with the current level of resourcing cannot properly fulfil its mandate. It will lack the technical staff needed to conduct evaluations of frontier AI systems, the capacity to stay ahead of advancing capabilities, or the personnel to meaningfully support dozens of regulators managing AI risks across healthcare, finance, transport, critical infrastructure, and other sectors. To satisfy its mandate, **AISI funding should be at least proportionate to the UK's AISI—\$50 million per year** (see below comparison).

---

<sup>30</sup> The Commonwealth of Australia. (2025). *Mid-Year Economic and Fiscal Outlook 2025–26* (p. 266).

<sup>31</sup> Grundy, E., Sadler, G., & Freeman, L. (2026). *Designing Australia's AI Safety Institute: Expert Survey*. Good Ancestors.

<sup>32</sup> Hendry, J. (2026). *Red tape threat to top talent joining AI Safety Institute*. InnovationAus.

<sup>33</sup> Allen, D. (2025). *Govt sets aside \$166m for public service AI chatbot*. InnovationAus.

### International comparison: UK AI Security Institute

Australia can draw on the **United Kingdom** as the primary international reference point for AI Safety Institute funding. The UK has allocated **GBP 66 million annually** (AUD 132 million) for its AISI, which is about **\$1.89 annually per capita** or **0.0022%** of GDP. Adjusting for Australia's population (27 million) or GDP (AUD 2.73 trillion), the scale of investment below emerges:

Benchmark	Australian Equivalent Investment (AUD)	Rationale
UK Per Capita (AUD 1.89)	<b>\$51 million</b>	$\$1.89 \times 27 \text{ million}$
UK % of GDP (0.0022%)	<b>\$60 million</b>	$0.0022\% \times \$2.73 \text{ trillion}$

## Capturing value through AI infrastructure

As AI companies build out physical and digital infrastructure, Australia has an opportunity to secure a position in the AI value chain. Between 2023 and 2025, companies announced plans for investments in Australian data centres that could scale to more than \$100 billion. Recent commitments include Amazon's \$20 billion investment (June 2025),<sup>34</sup> Microsoft's \$5 billion (October 2023),<sup>35</sup> and OpenAI's \$7 billion partnership with NEXTDC (December 2025).<sup>36</sup>

Any investments or Australian Government support should ensure Australia captures enduring benefits from this infrastructure, and that it doesn't just extract wealth and resources. This means securing Australian participation in potential "superprofits" if data centre operations become vastly valuable.

The AI infrastructure build-out also requires protecting against energy costs becoming prohibitive for households and businesses. In 2025, data centres consumed 3.9 TWh of electricity—2% of grid-supplied power—across the National Electricity Market. This demand is projected to triple by 2030.<sup>37</sup> While Australia has potential to become a clean energy superpower, feeding data centres could become more profitable than supplying households and businesses while driving up other energy infrastructure costs without adequate planning. Energy planning must keep up to avoid driving up costs for other users.

Government is developing national data centre principles with states and territories to ensure investments "align with Australia's overall national interests".<sup>38</sup> **At a minimum, AI companies striking large AI data-centre deals in Australia should engage in transparent information sharing with the AISI, and make compute resources available to advancing AI safety science.**

<sup>34</sup> Albanese, A. (2025). [Amazon data centre investment in Australia](#) [Press release]. Prime Minister of Australia.

<sup>35</sup> Microsoft. (2023). [Microsoft announces A\\$5 billion investment in computing capacity and capability to help Australia seize the AI era](#).

<sup>36</sup> Chalmers, J., & Charlton, A. (2025). [\\$7 billion infrastructure deal to boost AI in Australia](#) [Press release]. Department of Industry, Science and Resources.

<sup>37</sup> Australian Energy Market Operator. (2025). [2025 Inputs, Assumptions and Scenarios Report August 2025](#).

<sup>38</sup> Department of Industry, Science and Resources. (2025). [National AI Plan](#) (p.11). Australian Government.

## Capturing value through AI Assurance Technology

Australia can build a world-leading AI Assurance Technology (AIAT) industry that addresses safety concerns and helps secure a substantial position in the AI value chain. Australia's success in safety-critical industries like mining safety, aviation standards, and food biosecurity provides a strong foundation. Australia has demonstrated capability to develop and export robust assurance frameworks. A successful AIAT industry would address safety concerns slowing Australian AI adoption while also positioning Australian companies as global exporters of safety solutions.

AIAT is the software, hardware, and services that allow organisations to understand and mitigate AI risks. The global AIAT industry was worth USD 1.6 billion in 2023 and is forecast to reach USD 276 billion by 2030.<sup>39</sup> This represents an "upstream" economic opportunity and would position Australia as a direct contributor to AI products consumed globally. In the same way that companies want their auditor to be independent from their accountant, companies may look for AIAT solutions that are independent from the large AI companies the AIAT products will oversee.

Government has demonstrated an appetite for frontier technology investments that position Australia in global supply chains. The Federal and Queensland governments, for example, jointly committed \$940 million to PsiQuantum for quantum computing capability that remains years from commercial viability.<sup>40</sup> AI Assurance Technology is a more immediate opportunity: AI systems are already deployed at scale, businesses need safety tools now, and the global AIAT market is forecast to reach USD 276 billion by 2030. Australia's strengths in safety-critical industries make this a natural fit.

Australia already has promising AIAT startups, but they need support to become global players. **Government should establish an AIAT Grant Programme, jointly administered by the National AI Centre and the AISI, to provide Australian startups and universities with funding to establish and scale AIAT in Australia.**

This model follows CSIRO's Main Sequence Ventures, which translated Australian research into companies now valued at over \$6.8 billion.<sup>41</sup> Main Sequence was founded with \$70 million in Government cornerstone investment and attracted over \$1 billion in total funds under management. The programme "reinvest[s] economic returns from commercialisation back into science".<sup>42</sup>

**Government should provide \$90 million in cornerstone funding** (matching the inflation-adjusted Main Sequence investment) to establish the AIAT Grant Programme. The UK AISI provides a precedent for joint administration and partnering with research councils to deliver grants.<sup>43</sup>

Without intervention, Australia may lose talent and fall behind while other nations capitalise on this emerging industry.

---

<sup>39</sup> AIAT Report. (2024). [Risk & Reward: AI Assurance Technology Market Report 2024](#).

<sup>40</sup> Albanese, A., Miles, S., Husic, E., & Dick, C. (2024). [Delivering a Future Made in Australia with 400 new technology jobs in Brisbane](#). Prime Minister of Australia.

<sup>41</sup> CSIRO. (2023). [Main Sequence surpasses \\$1 billion funds under management for deep tech as it launches new Fund](#).

<sup>42</sup> Main Sequence. (n.d.). [Introducing Main Sequence](#). Retrieved January 28, 2026.

<sup>43</sup> Department for Science, Innovation and Technology. (2024). [Research programme to ensure UK economy uses AI to grow safely](#). GOV.UK.

## Recommendations

### 1. Increase AISI funding to match its extensive mandate.

Government assigned the AISI responsibility for whole-of-government coordination and supporting regulators across all sectors, in addition to evaluating frontier models and engaging internationally. AISI funding should increase to match that remit.

**We recommend the scale of investment be at least proportionate to the UK's AISI—\$50 million per year.** This would bring the Australian AISI roughly in line with its UK peer and expert recommendations.

### 2. Ensure data centre investments serve national interests.

Structure arrangements that secure Australian participation in potential superprofits and protect against energy costs becoming prohibitive for households and businesses.

**At minimum, AI companies striking large AI data centre deals should commit to transparent information sharing with the AISI and make compute resources available for AI safety research.**

Infrastructure investment should avoid no-strings-attached subsidies to foreign corporations that build wealth extraction devices on Australian soil.

### 3. Invest in Australia's AI Assurance Technology industry.

The global AIAT market will reach USD 276 billion by 2030, and Australia has advantages in safety-critical industries that mean we're well-placed to capture significant value.

**Government should establish an AIAT grant programme, jointly administered by the National AI Centre and AISI, with \$90 million in cornerstone funding.** This matches the inflation-adjusted investment that launched CSIRO's Main Sequence Ventures (\$70 million in 2017), which attracted over \$1 billion in total funds and created companies now valued at \$6.8 billion.